

Learning curves for overparametrized deep neural networks: A field theory perspective

Omry Cohen^{*,†}, Or Malka^{*,‡} and Zohar Ringel[§]

The Racah Institute of Physics, The Hebrew University of Jerusalem, Jerusalem, Israel

 (Received 12 March 2020; revised 20 February 2021; accepted 15 January 2021; published 9 April 2021)

In the past decade, deep neural networks (DNNs) came to the fore as the leading machine-learning algorithms for a variety of tasks. Their rise was founded on market needs and engineering craftsmanship, the latter based more on trial and error than on theory. While still far behind the application forefront, the theoretical study of DNNs has recently made important advancements in analyzing the highly overparametrized regime where some exact results have been obtained. Leveraging these ideas and adopting a more physicslike approach, here we construct a versatile field theory formalism for supervised deep learning, involving renormalization group, Feynman diagrams, and replicas. In particular, we show that our approach leads to highly accurate predictions of learning curves of truly deep DNNs trained on polynomial regression problems. It also explains in a concrete manner why DNNs generalize well despite being highly overparametrized, this due to an entropic bias to simple functions which, for the case of fully connected DNNs with data sampled on the hypersphere, are low-order polynomials in the input vector. Being a complex interacting system of artificial neurons, we believe that such tools and methodologies borrowed from condensed matter physics would prove essential for obtaining an accurate quantitative understanding of deep learning.

DOI: [10.1103/PhysRevResearch.3.023034](https://doi.org/10.1103/PhysRevResearch.3.023034)

I. INTRODUCTION

Deep artificial neural networks (DNNs) have been rapidly advancing the state of the art in machine learning, showing human and sometimes superhuman performance in image recognition [1], speech recognition [2], reinforcement learning [3], and natural language processing tasks [4]. Their rise to prominence was largely results driven, with little theoretical support or guarantee [5]. Such mode of invention is very different from how, say, the transistor was discovered, and more akin to how new materials, such as lithium-ion batteries, are discovered. Indeed, being huge interacting systems of artificial neurons, DNNs are more analogous to a complex metamaterial than to an electronic component [6]. Due to this complexity, a general theory of deep learning with predictive power is still lacking.

Notwithstanding, recently several results were obtained in the highly overparametrized regime [7,8] where the role played by any specific DNN weight is small. This facilitated the proofs of various bounds [9–11] on generalization for shallow networks and, more relevant for this work, two correspondences between fully trained DNNs and a different type of inference model called Gaussian processes (GPs) [12]. As

shown below, these can be thought of as noninteracting scalar field theories with disorder and a nonlocal action.

The first such correspondence [8] between GPs and trained DNNs is known as the neural tangent kernel result, which we would refer to here as the NTK correspondence. It holds when highly overparametrized DNNs are initialized according to standard practice and trained with mean-square-error (MSE) loss at vanishing learning rate and without weight decay.

The second correspondence [13] (the NNSP, neural networks stochastic process correspondence) applies when DNNs are trained using a similar protocol which involves random noise, roughly mimicking the stochastic gradient descent (SGD) optimization, as well as weight decay. It relates the outputs of the trained DNN to a stochastic process (SP) which, in the highly overparametrized limit, tends to a GP. It thus yields an additional training protocol, complementary in some ways to the previous one, which is analytically tractable.

How much of deep learning can be explained through such correspondences remains to be seen. On the one hand, some aspects such as learning sharp filters (features) in the first DNN layers seem out of reach as specific DNN weights change only infinitesimally in the NTK case and remain largely random, apart from a small bias, in the NNSP case. In addition, learning in the NTK regime, sometimes dubbed “lazy learning,” often lags behind state-of-the-art training protocols (see [14] and references therein) where finite learning rates, widths, and minibatches are used. On the other hand, lazy learning or more generally GP methods are being extended and improved [15–18] by importing technologies such as pooling and data augmentation. Currently, GP models corresponding to DNNs are competitive with deep learning on the UCI data sets [15] as well as Fashion-MNIST [16], whereas on the CIFAR-10 data set the performance of the best GPs currently lags 5% behind celebrated DNNs such

* O.C. and O.M. contributed equally to this work.

† omry.cohen@mail.huji.ac.il

‡ or.malka@mail.huji.ac.il

§ zohar.ringel@mail.huji.ac.il

as AlexNet [16,18], while surpassing pre-AlexNet nondeep methods by 8% [19]. In addition, there is the prospect of extending these correspondences by including nonlinearities coming from finite-width [13] and finite-learning rates [20,21] settings. These results and prospects invite further study of how such DNNs trained in the NTK and NNSP regimes make predictions.

In this work we introduce a versatile field theory formalism for analyzing deep neural networks, which involves replicas, Feynman diagrams, and renormalization group (RG) techniques. In its most basic version, studied in depth below, it applies to DNNs trained using the protocols for which the NTK and NNSP correspondences hold exactly and lead to GP models. For these cases we provide expressions for the generalization power of fully connected DNNs in the form of learning curves. These learning curves depend on the data-set distribution and the target which we learn. For uniform data sets on the hypersphere and any target function, our learning curves become fully explicit and provide a clear picture of how such DNNs generalize. This includes the more challenging case of the NTK correspondence where certain infinities in the action are removed by our renormalization group transformation. To the best of our knowledge, the accuracy at which our learning curves capture the empirical ones far exceeds the current theoretical state of the art.

In addition, our formalism can also accommodate various extensions of these correspondences. For the case of the NNSP correspondence, we can work with loss functions different than MSE as well as corrections to the infinite overparametrization limit. Furthermore, recent results on extensions of the NTK correspondence [20] suggest that high learning rate leads to a renormalized NTK correspondence whose performance can again be analyzed using our approach. Such extensions may prove useful in addressing the gap [22] between GPs and their DNN counterparts.

We hope that the results and formalism introduced here would aid in developing a more physicslike paradigm for studying DNNs, complementary to the proof-based approach common in theoretical computer science (see also [23–27]). Such a paradigm should fill in the gap, typically large in complex systems, between what can be predicted following some reasonable assumptions and what can be proven rigorously.

This paper is structured as follows. In Sec. III we provide the necessary background on deep neural networks, Gaussian processes, and the correspondences between the two. Section IV describes our field theory approach and analytical results. Section V considers the case of uniformly distributed data on the hypersphere, where further analytical simplifications can be carried. Section VI introduces the RG approach used to tackle the noiseless NTK case. Section VII applies our results to concrete examples and compares them with empirical results. Section VIII shows how our results can be used to perform efficient hyperparameter optimization on actual DNNs, and Sec. IX summarizes the results and discusses possible directions for future work.

II. PRIOR WORKS

Learning curves for GPs have been analyzed using a variety of techniques (see [12] for a review), most of which focus on a GP-teacher averaged case where the target or teacher

is drawn from the same GP used for inference (matched priors) and is furthermore averaged over. Fixed-teacher or fixed-target learning curves have been analyzed using a grand-canonical and Poisson-averaged approach [28] similar to the one we used. However, their treatment of the resulting partition function was variational whereas we take a different, perturbation-theory-based, approach. In addition, previous cited results for MSE loss break in the noiseless limit [28]. To the best of our knowledge, noiseless GP learning curves have been analyzed analytically only in the teacher-averaged case and limited to the following settings: For matched priors, exact results are known for one-dimensional data [12,29] and two-dimensional data with some limitations of how one samples the inputs (in the context of optimal design) [30,31]. In addition [32] derived a lower bound on generalization. For noiseless inference with partially mismatched priors (matching features, mismatching eigenvalues) and at large input dimension the teacher and data-set averaging involved in obtaining learning curves, has been performed analytically and the resulting matrix traces analyzed numerically [33]. Notably, none of these cited results apply in any straightforward manner in the NTK regime.

Considering kernel eigenvalues, explicit expression for the features and eigenvalues of dot-product kernels [$K = K(x \cdot x')$] were given in [34]. The fact, that the l th eigenvalue of such kernels scales as d^{-l} (d being the input dimension), which we used in our derivation of the bound, has been noticed in [33]. Kernels with a trimmed spectrum where the spectrum is trimmed after the first r 's leading eigenvalues has previously been suggested as a way of reducing the computational cost of GP inference [35]. In contrast, we trim the Taylor expansion of the kernel function rather than the spectrum (which has a very different effect) and show that an effective observation noise compensates for our trimming and renormalization procedure.

Several interesting recent works give bounds on generalization [9–11] which show $O(1/\sqrt{N})$ asymptotic decay of the learning curve (at best). In contrast, our predictions are typically well below this bound.

III. THEORETICAL BACKGROUND

A. DNNs, expected error, and learning curves

We begin with the standard definitions of DNNs as they apply to this work. While the majority of this work is applicable to many network architectures, we will focus on a simple feed-forward network for the sake of simplicity. A fully connected feed-forward DNN with L hidden layers of width n_l for $l = 1, \dots, L$ and readout layer $n_{L+1} = k$ is a function f defined recursively by

$$\begin{aligned} h^{l+1} &= x^l W^{l+1} + b^{l+1}, \\ x^{l+1} &= \phi(h^{l+1}), \\ f(x; W^1, \dots, W^L, b^1, \dots, b^L) &= x^{L+1}, \end{aligned} \quad (1)$$

where ϕ is a pointwise activation function, $x^0 \in \mathbb{R}^d$ is the input of the network, and $W^{l+1} \in \mathbb{R}^{n_l \times n_{l+1}}$, $b^{l+1} \in \mathbb{R}^{n_{l+1}}$ are trainable weights and biases, which will be collectively referred to as weights from here on. Each component of the

weights is usually initialized randomly from a normal distribution $\mathcal{N}(0, \sigma_w^2)$ for the weights and $\mathcal{N}(0, \sigma_b^2)$ for the biases.

In the usual setting one starts with a training set, a set of input points $D = \{x_n\}_{n=1}^N$ where $x_n \in \mathbb{R}^d$ along with their labels $\{l_n\}_{n=1}^N$ where $l_n \in \mathbb{R}^k$. One then picks weights for the network by minimizing a loss function $\mathcal{L}(f(D), \{l_n\})$ which compares the values of network function over D to the labels $\{l_n\}_{n=1}^N$, assigning a smaller value to points where the network function and labels are similar. One then finds weights which minimize the loss by some variation of gradient descent, usually stochastic gradient descent (SGD), wherein one approximates the gradient at each iteration using a random batch of the training set (see [36] for details). The performance of the network is then evaluated by computing the loss function over a set of labeled points, different from the training set, known as the test set. This is known as the test error of the network, and is used as a proxy for the *expected error*, the loss averaged over draws from the data-set distribution.

One of the most detailed objects quantifying the performance of a machine-learning algorithm, and the main focus of this work, is its learning curve, a graph of how the expected error diminishes with the number of data points (N). There are currently no analytical predictions or bounds we are aware of for DNN learning curves which are tight even just in terms of their scaling with N , let alone tight in an absolute sense (see Sec. II)

B. Gaussian processes regression

In this work we will investigate the properties of DNNs by their correspondence with GPs. We supply here some standard definitions of GPs and their usage in regression tasks. Regression here simply means approximating a function $[g(x)]$ based on discrete samples $[\{g(x_n)\}_{n=1}^N]$. A GP is commonly defined as a stochastic process of which any finite subset of random variables follow a multivariate normal distribution [12]. In a similar fashion to multivariate normal variables, GPs are also determined by their first and second moments. The first is typically taken to be zero, and second is known as the covariance function or the kernel $K_{x,x'} = \mathbb{E}[f(x)f(x')]$, where $\mathbb{E}[\cdot]$ here denotes expectation with respect to the GP distribution. The main appeal of GPs is that Bayesian inference with GP priors is tractable [12]. In GP inference we use the mean of the GP distribution conditioned on the data (posterior) as the predictor g^* , and it is given by

$$g^*(x_*) = \sum_{n,m=1}^N K_{x_*,x_n} [K(D) + \sigma^2 I]_{nm}^{-1} l_m, \quad (2)$$

where x_* is a new data point, l_m are the training targets, x_n are the training data points, $[K(D)]_{nm} = K_{x_n,x_m}$ is the covariance matrix (the covariance function projected on the training data set D), and σ^2 is the variance of the assumed Gaussian noise of the labels, which also acts as a regulator of the prediction. Some intuition for this formula can be gained by verifying that in the noiseless case ($\sigma^2 = 0$) the prediction at some training point $x_* = x_q$ coincides with that point's label $g^* = l_q$.

The quantity of interest in this paper, which we define now, is the expected error averaged over all the possible data sets. Throughout this paper we will assume that both train and test

points are drawn from a probability measure $d\mu_x = P(x)dx$. With this in mind, we define the expected error of a prediction g^* as

$$\|g - g^*\|^2 = \int d\mu_x [g(x) - g^*(x)]^2. \quad (3)$$

Note that g^* is itself a function of N draws from μ which make up the training set D_N . Our quantity of interest, the data-set-averaged expected error (DAEE), is Eq. (3) averaged over the ensemble of all possible N -sized training sets. We denote this average as $\langle \cdot \rangle_{D_N}$, so the DAEE is given by $\langle \|g^* - g\|^2 \rangle_{D_N}$. The learning curve is the dependence of the DAEE on N . We see that in order to calculate learning curves, one needs to calculate quantities like $\langle g^* \rangle_{D_N}$ and $\langle g^{*2} \rangle_{D_N}$.

Equation (2) determines the predictions, and therefore the learning curves, but it is not very convenient for analytic exploration of the expected predictions. This fact is due to the (potentially very) large matrix inversion involved, and the additional averaging over D_N required. Nonetheless, there are some approximations for the expected prediction $\langle g^* \rangle_{D_N}$, the most famous of which is the equivalence kernel (EK) result [12]

$$\langle g^*(x) \rangle_{D_N} \approx g_{EK,N}^*(x) = \sum_n \frac{\lambda_n}{\lambda_n + \frac{\sigma^2}{N}} g_n \phi_n(x), \quad (4)$$

where λ_n and $\phi_n(x)$ here are the eigenvalues and eigenfunctions of the kernel with respect to the input probability measure μ , and $g(x) = \sum_n g_n \phi_n(x)$ is the target function. One notices immediately that this approximation breaks down completely in the noiseless case where Eq. (4) implies perfect estimation of the target with just one data point. To gain some intuition as to why having $\sigma^2 = 0$ hinders predictions of $\langle g^* \rangle_{D_N}$, one can view it as a hard constraint $[f(x_n) = g(x_n)]$, and hard constraints are typically less tractable than soft ones. In a related view, finite σ^2 can be seen as a form of averaging which smooths and regulates analytical expressions making them more tractable. Another limitation of the EK result is that (to the best of our knowledge) there is no systematic way to extend it in orders of $1/N$ and get a more detailed picture of generalization in GP regression (GPR).

C. From DNNs to GPs through Langevin dynamics

Here we review, for completeness, several recent correspondences between DNNs and GPs. It has long been known [37] that randomly initialized, infinitely wide DNNs with independent and identically distributed weights are equivalent to samples from a GP known as the neural network GP (NNGP). More recently it was shown that training only the last layer of a network with gradient descent is equivalent to posterior sampling of the NNGP [38], and consequently averaging the prediction of many networks trained on the same data set is equivalent to GPR. Turning to more standard training of the entire DNN, it has been recently established [8] that fully training a network with vanishing learning rate for infinitely long time and MSE loss yields the same predictions as a noiseless GPR with a different kernel, the neural tangent kernel (NTK), along with an additional initialization-dependent term. Averaging over many initialization seeds gives an exact correspondence with a GP whose kernel is the NTK.

Recently, another correspondence between DNNs and GPs has been introduced [13]. Due to its simplicity we shall rederive it here. Consider the training of a DNN using gradient descent (full-batch SGD) with weight decay and added white noise, in the limit of vanishing learning rate. For sufficiently small learning rate, and making the reasonable assumption that the gradients of the loss are globally Lifshitz, the SGD equations are ergodic and converge to the same invariant measure (equilibrium distribution) as the following Langevin equation [39,40]:

$$\frac{dw_i}{dt} = -\partial_{w_i} \left(\mathcal{L}[z_W] + \sum_j \frac{T w_j^2}{2\sigma_w^2} \right) + \sqrt{2T} \xi_i(t), \quad (5)$$

where $\xi_i(t)$ being a set of Gaussian white noise [$\langle \xi_i(t) \xi_j(t') \rangle = \delta_{ij} \delta(t - t')$], T accounting for the strength of the noise, w_i being the set of network parameters (W), z_W is the network output for a given configuration of W , and \mathcal{L} is the loss function. The equilibrium distribution or invariant measure describing the steady state of the above equation is the Boltzmann distribution [40] $P(W) \propto e^{-\frac{1}{2\sigma^2} \mathcal{L}[z_W] - \frac{1}{2\sigma_w^2} \sum_i w_i^2}$ with $T = 2\sigma^2$. Notably, various works argue that at low learning rates, discrete SGD dynamics running for long enough times reaches the above equilibrium, approximately [41,42].

Next, we adopt the approach of [8] and describe the dynamics in function space (f) instead of weight space (W). Using the Boltzmann distribution described above, the post-training probability density function for some function f is given by

$$\begin{aligned} P[f] &= \int dW P(W) \delta[f - z_W] \\ &\propto e^{-\frac{1}{2\sigma^2} \mathcal{L}[f]} \int dW e^{-\frac{1}{2\sigma_w^2} \sum_i w_i^2} \delta[f - z_W] \\ &\propto P_{nd}[f] e^{-\frac{1}{2\sigma^2} \mathcal{L}[f]}, \end{aligned} \quad (6)$$

where we identify $P_{nd}[f] \propto \int dW e^{-\frac{1}{2\sigma_w^2} \sum_i w_i^2} \delta[f - z_W]$ as the distribution of the output of the network after being trained with no data (or, equivalently, a vanishing loss function). $\delta[\dots]$ is a *functional* delta function, which can be thought of as the limit of a large product of regular delta functions on each Fourier component of the argument. As we will discuss, for an infinitely overparametrized network P_{nd} coincides with the prior of a NNGP with the weights and biases variance determined by training parameters rather than by initialization. However, for finite overparametrization it becomes a more generic stochastic process determined by the neural network (an NNSP). In [13], the leading finite-width corrections were calculated and shown to result in f^4 corrections to the prior.

Clearly, the practical use of the above result hinges on how quickly the dynamics mixes or reaches ergodicity. While ergodicity in its full sense (for any weight-space observable) seems unrealistic in this nonconvex scenario, reaching ergodicity in the mean of the outputs of the DNNs [for low-order polynomials in $f(x)$] may be quicker. This milder form of ergodicity was shown numerically for fully connected DNNs trained on regression problems similar to those studied here as well as CNNs trained on CIFAR-10 [13].

From now on we shall focus on the infinite overparametrized limit.

IV. FIELD THEORY FORMULATION OF GP LEARNING CURVES

A. Rephrasing GPs as a field theory

We begin by phrasing inference with GPs in the language of field theory. To this end, we first write a Gaussian distribution over the space of functions that leads to a two-point correlation function equal to $K_{xx'}$. This is given by

$$\begin{aligned} P_0[f] &\propto e^{-\frac{1}{2} \|f\|_K^2}, \\ \|f\|_K^2 &= \int d\mu_x d\mu_{x'} f(x) K^{-1}(x, x') f(x'), \end{aligned} \quad (7)$$

where $K^{-1}(x, x')$ is the inverse kernel function, meaning that $\int d\mu_{x''} K(x, x'') K^{-1}(x'', x') = \delta(x - x')/P(x)$ where $d\mu = P(x)dx$. This formalism is sometimes referred to as information field theory (IFT) [43].

A different viewpoint on $P_0[f]$ comes from viewing $f(x)$ as the outputs of a wide DNN with weights drawn from an independent and identically distributed Gaussian distribution $P_0(W)$. It is well known [44] that correlations between the outputs of random DNNs are Gaussian and governed by some kernel $K_{xx'}$. This kernel is determined, in a tractable manner, by the DNN architecture. From a field theory viewpoint this can be stated as

$$P_0[f] = \int dW P_0(W) \delta[f - z_W] \quad (8)$$

at infinite width, where $z_W(x)$ is the output of a DNN with weights W on an input point x . The keen reader may be alarmed by the fact that this definition of $P_0[f]$ does not involve the measure $\mu(x)$. However, as shown in [12], the norm $\|f\|_K^2$ (called the RKHS norm), and therefore Eq. (7) are in fact the same for any two probability measures with identical support.

Performing Bayesian inference in the context of GPs means conditioning Eq. (7) using Bayes' theorem and assuming Gaussian noise with amplitude σ^2 on our target function $g(x)$. This yields the additional factor

$$P[f] \propto e^{-\frac{1}{2} \|f\|_K^2 - \frac{1}{2\sigma^2} \sum_{i=1}^N [f(x_i) - g(x_i)]^2}. \quad (9)$$

It can be checked that the expectation value of $f(x_*)$ under the above probability yields Eq. (2).

Notably, by taking into account Eq. (8), the above expression coincides with that obtained via the NNSP correspondence (6) in the infinite overparametrization limit where $P_{nd}[f] = P_0[f]$ for MSE loss and a suitably chosen $K_{xx'}$. In the NNSP context, the data term came out quadratic when training using MSE loss and more generally it could be replaced with a general loss function $\mathcal{L}[f]$, so the DNNs predictive distribution becomes

$$P[f] \propto e^{-\frac{1}{2} \|f\|_K^2 - \frac{1}{2\sigma^2} \mathcal{L}[f]}, \quad (10)$$

where K in this context is the kernel of the DNN trained with no data. Although not necessarily Gaussian, this expression can still be treated using mean-field or perturbative




FIG. 1. A physical picture of supervised deep learning. The output of the DNN, as a function of input data, can be seen as an elastic membrane (surface) which relaxes to its equilibrium distribution during training. In this steady state it fluctuates (green surface) so to maximize its entropy while minimizing its energy. Its energy consists of a data term pinning it to its target values (yellow surface) on the training points (red points). In addition, an elastic energy term determined by the DNN architecture affects its behavior between the training points. For infinitely overparametrized DNNs, this elastic energy is quadratic and the average surface (blue surface) can be calculated analytically, up to a large matrix inversion, using Gaussian processes regression.

approaches. A more detailed treatment of different loss functions, most notably as cross-entropy loss, is left for future work.

Denoting $S[f] = \frac{1}{2} \|f\|_K^2 + \frac{1}{2\sigma^2} \mathcal{L}[f]$ (the “information Hamiltonian,” in IFT terminology), Eq. (9) gives rise to the partition function

$$Z[\alpha] = \int \mathcal{D}f e^{-S[f] + \int dx \alpha(x) f(x)}, \quad (11)$$

where $\int dx \alpha(x) f(x)$ is a source term used to calculate cumulants of $P[f]$, and specifically the average prediction of the network:

$$g^*(x_*) = \left. \frac{\delta \log(Z[\alpha])}{\delta \alpha(x_*)} \right|_{\alpha=0} = \frac{1}{Z[0]} \int \mathcal{D}f f(x_*) e^{-S[f]}, \quad (12)$$

where $\delta/\delta\alpha$ stands for functional derivative. As shown visually in Fig. 1, this expression leads to a tangible physical picture of how DNNs learn. Their output as a function of the input can be seen as a fluctuating elastic membrane over input space which, in the highly overparametrized limit, is in its linear elastic regime. The training data appear as isolated points at which this membrane is pinned down to certain value by (loss-dependent) springs whose constant is proportional to $1/\sigma^2$, the inverse of the noise on the gradients during training. The membrane then interpolates and extrapolates between these pinning points in a way which, on average, minimizes its elastic energy. This elastic energy differs considerably from that of physical membrane and in particular has a nonlocal dependence on the shape of the membrane. Different DNNs correspond to different elastic energies. Finite networks entail nonlinear corrections to the elastic energy which may be beneficial for learning in the case of CNNs [13].

B. Predictions in the grand-canonical ensemble

As mentioned, in order to calculate the learning curve one needs to calculate quantities like $\langle g^* \rangle_{D_N}$ and $\langle g^{*2} \rangle_{D_N}$. These averages involve multidimensional integrations over all possible data sets. To facilitate their computation, we adopt the approach of [28] and instead consider a related quantity given by the Poisson averaging of the former

$$\langle \dots \rangle_\eta = e^{-\eta} \sum_{n=0}^{\infty} \frac{\eta^n}{n!} \langle \dots \rangle_{D_n}, \quad (13)$$

where the ellipsis can be any quantity, in particular g^* and g^{*2} . This average can be thought of as a grand-canonical ensemble, though a nonstandard one since we average the observables and not the partition function. Taking $\eta = N$ means we are essentially averaging over values of N in a \sqrt{N} vicinity of N . This means that as far as the leading asymptotic behavior is concerned, one can safely exchange N and η as the differences would be subleading. We therefore focus on calculating the grand-canonical DAEE $\langle \|g^* - g\|^2 \rangle_\eta$. In Appendix A we compare learning curves as a function of N and η and show that they match very well.

By using the grand-canonical ensemble, averaging over draws from the data set can be carried out as follows. First, using the replica trick

$$\langle g^*(x_*) \rangle_\eta = \lim_{M \rightarrow 0} \left. \frac{1}{M} \frac{\delta \langle Z^M \rangle_\eta}{\delta \alpha(x_*)} \right|_{\alpha=0}, \quad (14)$$

where for integer M and assuming that the loss function acts pointwise on the training set $\mathcal{L}[f] = \sum_{i=1}^n \mathcal{L}_f(x_i)$, we have

$$\begin{aligned} \langle Z^M \rangle_\eta &= e^{-\eta} \int \prod_{m=1}^M \mathcal{D}f_m \\ &\times e^{-\sum_{m=1}^M (\frac{1}{2} \|f_m\|_K^2 - \int dx \alpha f_m) + \eta \int dx e^{-\frac{\sum_{m=1}^M \mathcal{L}_{f_m}(x)}{2\sigma^2}}}. \end{aligned} \quad (15)$$

As shown in Appendix G2, a Taylor expansion in η of the above right-hand side yields the $\langle \dots \rangle_\eta$ averaging appearing on the left-hand side.

Second, we notice that the main benefit of Eqs. (14) and (15) over Eq. (2) is that it allows for a controlled expansion in $1/\eta$. At large η (or similarly large N) we expect the fluctuations in $f_m(x)$ to be small and centered around $g(x)$. Indeed, such a behavior is encouraged by the term multiplied by η in the exponent. We can therefore systematically Taylor expand the inner exponent

$$\begin{aligned} e^{-\frac{\sum_{m=1}^M \mathcal{L}_{f_m}(x)}{2\sigma^2}} &= 1 - \frac{\sum_{m=1}^M \mathcal{L}_{f_m}(x)}{2\sigma^2} \\ &+ \frac{1}{2} \left[\frac{\sum_{m=1}^M \mathcal{L}_{f_m}(x)}{2\sigma^2} \right]^2 + \dots \end{aligned} \quad (16)$$

and each term will yield a higher order of $\langle g^*(x_*) \rangle_\eta$ in $1/\eta$.

C. EK as a free theory

Notably, so far the choice of a loss function was largely arbitrary. The advantage of choosing MSE loss, $\mathcal{L}_f(x) = [f(x) - g(x)]^2$, is that $P[f]$ also becomes a GP or,

equivalently, has a quadratic action. From now on we shall focus on MSE loss.

Aiming for standard perturbative calculations, we wish to perform diagrammatic calculations with respect to a free quadratic theory. Expanding Eq. (16) to first order and substituting in Eq. (15) we obtain

$$\begin{aligned} \langle Z^M \rangle_\eta &= Z_{\text{EK}}^M + O(1/\eta^2), \\ Z_{\text{EK}}[\alpha] &= \int \mathcal{D}f e^{-S_{\text{EK}}[f] + \int dx \alpha(x) f(x)}, \end{aligned} \quad (17)$$

where $S_{\text{EK}}[f] = \frac{1}{2} \|f\|_K^2 + \frac{\eta}{2\sigma^2} \int d\mu_x [f(x) - g(x)]^2$, which is quadratic in f and therefore induces a Gaussian field.

Substituting Eq. (17) in (14) we get

$$\begin{aligned} \langle g^*(x_*) \rangle_\eta &= \left. \frac{\delta \log(Z_{\text{EK}})}{\delta \alpha(x_*)} \right|_{\alpha=0} + O(1/\eta^2) \\ &= \arg \min [S_{\text{EK}}[f]] + O(1/\eta^2) \\ &= g_{\text{EK},\eta}^*(x_*) + O(1/\eta^2), \end{aligned} \quad (18)$$

where the second equality is due to the fact that for Gaussian distributions the expectation value coincides with the most probable value, and the third equality is due to [12], with the subtle change that N is being replaced by η .

Let us denote by $\langle \dots \rangle_0$ the free-theory average, that is an average with respect to Z_{EK} . We therefore get $\langle f \rangle_0 = g_{\text{EK},\eta}^* \neq 0$, meaning that our free theory, though Gaussian, is not centered. The correlations of the free theory are

$$\text{Cov}_0[f(x), f(y)] = \sum_i \left(\frac{1}{\lambda_i} + \frac{\eta}{\sigma^2} \right)^{-1} \phi_i(x) \phi_i(y), \quad (19)$$

where again, λ_i and ϕ_i are the eigenvalues and eigenfunctions of the kernel.

D. Next-order corrections

We now wish to perform perturbative calculations with respect to the free (Gaussian) EK theory, and obtain a subleading (SL) correction for the EK result in the inverse data-set size:

$$\langle g^*(x_*) \rangle_\eta = g_{\text{EK},\eta}^*(x_*) + g_{\text{SL},\eta}^*(x_*) + O(1/\eta^3). \quad (20)$$

Expanding Eq. (16) to second order, substituting in Eq. (15) and keeping only $O(1/\eta^2)$ terms, the calculation can be carried using Feynman diagrams with respect to the free EK Gaussian theory. Leaving the details to Appendix I, the subleading correction is

$$\begin{aligned} g_{\text{SL},\eta}^*(x_*) &= \frac{\eta}{\sigma^4} \int d\mu_x [g_{\text{EK},\eta}^*(x) - g(x)] \\ &\quad \times \text{Cov}_0[f(x), f(x)] \text{Cov}_0[f(x), f(x_*)] \end{aligned} \quad (21)$$

or explicitly

$$\begin{aligned} g_{\text{SL},\eta}^*(x_*) &= -\frac{\eta}{\sigma^4} \sum_{i,j,k} \Lambda_{i,j,k} g_i \phi_j(x_*) \int d\mu_x \phi_i(x) \phi_j(x) \phi_k^2(x), \\ \Lambda_{i,j,k} &= \frac{\sigma^2}{\lambda_i + \frac{\sigma^2}{\eta}} \left(\frac{1}{\lambda_j} + \frac{\eta}{\sigma^2} \right)^{-1} \left(\frac{1}{\lambda_k} + \frac{\eta}{\sigma^2} \right)^{-1}. \end{aligned} \quad (22)$$

As shown in Appendix G2, similar expressions for $\langle g^{*2} \rangle_\eta$ are obtained using two replica indices. Interestingly, we find that $\langle g^{*2} \rangle_\eta = \langle g^* \rangle_\eta^2 + O(1/\eta^3)$. Hence, up to $O(1/\eta^3)$ corrections, the averaged MSE error is $[\langle g^*(x_*) \rangle_\eta - g(x_*)]^2$ integrated over x_* . Since the variance of g^* came out to be $O(1/\eta^3)$ one finds that $g^* - g$, which is $O(1/\eta)$, is asymptotically much larger than its standard deviation. This implies self-averaging at large η or, equivalently, that our data-set-averaged results capture the behavior of a single fixed data set.

Equations (20) and (22) and their application to the calculation of the grand-canonical DAEE are one of our key results. They provide us with closed expressions for the DAEE as a function of η , namely, the fixed-teacher learning curve. They hold without any limitations on the data set or the kernel and yield a variant of the EK result along with its subleading correction. From an analytic perspective, once λ_i and $\phi_i(x)$ are known, the above expressions provide clear insights to how well the GP learns each feature and what crosstalk is generated between features due to the second subleading term. Notably for the renormalized NTK introduced below, the number of nonzero λ_i 's is finite, and so the above infinite summations reduce to finite ones. This makes these expressions computationally superior to directly performing the matrix inversion in Eq. (2) along with an N -dimensional integral involved in data-set averaging. In addition, having the subleading correction allows us to estimate the range of validity of our approximation by comparing the subleading and leading contributions, as we shall do for the uniform case below.

V. UNIFORM DATA SETS

To make Eq. (22) interpretable, $\phi_i(x)$ and λ_i are required. This can be done most readily for the case of data sets normalized to the hypersphere ($\|x_\eta\| = 1$) with a uniform probability measure and rotation-symmetric kernel functions. By the latter we mean $K_{x,x'} = K_{Ox,Ox'}$ for any orthogonal matrix O with the same dimension as the inputs. Although beyond the scope of this work, obvious extensions to consider are data sets which are uniform only in a subspace of x and/or small perturbations to uniformity.

Importantly, both NNGP and NTK associated with any DNN with a fully connected first layer and weights initialized from a normal distribution have the above symmetry under rotations (see Appendix E). It follows that such a kernel can be expanded as $K_{x,x'} = \sum_n b_n (x \cdot x')^n$. An additional corollary [34] is that its features are hyperspherical harmonics $[Y_{lm}(x)]$ as these are the features of all dot product kernels. Hyperspherical harmonics are a complete and orthonormal basis with respect to a uniform probability measure on the hypersphere. Note that this implies a nonstandard normalization for the Y_{lm} 's in this context, as they are usually normalized with respect to Lebesgue measure. For each l these can be written as a sum of polynomials in the input coordinates of degree l . The extra index m enumerates an orthogonal set of such polynomials [of size $\text{deg}(l)$]. For a kernel of the above form the eigenvalues are independent of m and given by [34]

$$\lambda_l = \frac{\Gamma(\frac{d}{2})}{\sqrt{\pi} 2^l} \sum_{s=0}^{\infty} b_{2s+l} \frac{(2s+l)!}{(2s)!} \frac{\Gamma(s+\frac{1}{2})}{\Gamma(s+l+\frac{d}{2})}. \quad (23)$$

For ReLU and erf activations, the b_n 's can be obtained analytically up to any desirable order [44]. Thus, one can semianalytically obtain the eigenvalues up to any desired accuracy. For the particular case of depth 2 ReLU networks with no biases, we report in Appendix J closed expression where the above summation can be carried out analytically for the NNGP and NTK kernels. However, as we shall argue soon, it is in fact desirable to trim the NTK in the sense of cutting off its Taylor expansion at some order m , resulting in what we call the renormalized NTK. For such kernels, which would be our main focus next, Eq. (23) can be seen as a closed analytical expression for the eigenvalues.

Interestingly, for any dot-product kernel and uniform data of dimension d on the hypersphere, there is a universal bound given by $\lambda_l \leq K_{x,x} / \deg(l) \approx O(d^{-l})$, where $K_{x,x}$ is a constant in x . Indeed, $K_{x,x} = \sum_{lm} \lambda_l = \sum_l \deg(l) \lambda_l$. The degeneracy $[\deg(l)]$ is fixed from properties of hyperspherical harmonics, and equals $\deg(l) = \frac{2l+d-2}{l+d-2} \binom{l+d-2}{l}$ [45] which goes as $O(d^l)$ for $l \ll d$. This combined with the positivity of the λ_l 's implies the above bound.

Expressing our target in this feature basis $g(x) = \sum_{l,m} g_{lm} Y_{lm}(x)$, Eq. (22) simplifies to

$$g_{\text{SL},\eta}^*(x_*) = - \sum_{l,m} \frac{\eta^{-1} \lambda_l C_{K,\sigma^2/\eta}}{(\lambda_l + \sigma^2/\eta)^2} g_{lm} Y_{lm}(x_*), \quad (24)$$

where $C_{K,\sigma^2/\eta} = \sum_l \deg(l) (\lambda_l^{-1} + \eta/\sigma^2)^{-1}$ and notably crosstalk between features has been eliminated at this order since $\sum_m Y_{lm}^2(x) = \deg(l)$ is independent of x , yielding $\sum_{\tilde{m}} \int d\mu_x Y_{lm}(x) Y_{l'm'}(x) Y_{\tilde{l}\tilde{m}}^2(x) = \deg(\tilde{l}) \delta_{ll'} \delta_{mm'}$.

By splitting the sum in $C_{K,\sigma^2/\eta}$, to cases in which $\lambda_l < \sigma^2/\eta$ and their complement, one has the bound $C_{K,\sigma^2/\eta} < \#F \sigma^2/\eta + \sum_{l|m|\lambda_l < \sigma^2/\eta} \lambda_l$, where $\#F$ is the number of eigenvalues such that $\lambda_l > \sigma^2/\eta$. Thus, for kernels with a finite number of nonzero λ_l 's (as the renormalized NTK introduced below), and for large enough η , $\#F$ becomes the number of nonzero eigenvalues and $C_{K,\sigma^2/\eta} = \#F \sigma^2/\eta$ has a η^{-1} asymptotic. This illustrates the fact that the above terms are arranged by their orders in η .

We can use Eq. (24) to understand the validity of the EK result. We therefore look for sufficient conditions for $g_{\text{EK},\eta}^* \gg g_{\text{SL},\eta}^*$ to hold. By a termwise comparison, for some l we obtain $C_{K,\sigma^2/\eta} \ll \eta(\lambda_l + \sigma^2/\eta)$ which holds for $C_{K,\sigma^2/\eta} \ll \sigma^2$. For trimmed kernels, this yields $\#F \ll \eta$. Notably, it means that the original nontrimmed NTK cannot be analyzed perturbatively since with $\sigma^2 = 0$, $\#F$ becomes infinite. In the next section we tackle this issue.

VI. GENERALIZATION IN THE NOISELESS CASE AND THE RENORMALIZED NTK

The correspondence between DNNs trained in the NTK regime and GPR using NTK implies noiseless GPR ($\sigma^2 = 0$) for which the perturbative analysis carried in previous sections fails. Here we show that the fluctuations of f associated with small λ_l 's can be traded for noise on the fluctuations of f associated with large λ_l 's, thereby making our perturbative analysis applicable. As shown in the previous section, for uniform data sets, the smaller λ_l 's correspond to higher

spherical harmonics (higher l) and hence have higher oscillatory components. We argue that these higher oscillatory modes can be marginalized over in a controlled manner to generate both noise and corrections to the large λ_l 's. This is very much in the spirit of the renormalization group technique, wherein high oscillatory modes are integrated over to generate changes (renormalization) of some parameters in the probability distribution of the low oscillatory modes.

We begin by defining a set of renormalized NTKs. As argued, an NTK of any fully connected DNN can be expanded as $K_{x,x'} = \sum_{q=0}^{\infty} b_q (x \cdot x')^q$. The renormalized NTK at scale r is then simply $K_{x,x'}^{(r)} = \sum_{q=0}^r b_q (x \cdot x')^q$. Harmoniously with this notation we denote the prediction of GPR with the original kernel as g_{∞}^* . Our claim is that GPR with K and a noise of σ^2 can be well approximated by GPR with $K^{(r)}$ and noise $\sigma^2 + \sigma_r^2$ (where $\sigma_r^2 = \sum_{q=r+1}^{\infty} b_q$), for sufficiently large r . Specifically, our claim is that the discrepancy between the original vs truncated GPR predictions scales as $O(\sqrt{N} d^{-(r+1)/2} / K_{x,x})$, where d is the effective data-input dimension. Importantly, as can be inferred from Eq. (23), the renormalized NTK at scale r has zero eigenvalues for all spherical harmonics with $l > r$, as well as modified eigenvalues for spherical harmonics with $l \leq r$ (compared to the nontruncated NTK). Thus, as advertised, these high Fourier modes have been removed from the problem in exchange for a renormalized theory with a modified low-energy spectrum, and augmented noise. In a related manner, trimming the Taylor expansion after $(x \cdot x')^r$ effectively reduces our angular resolution and coarse grains the fine angular features captured by these spherical harmonics with $l > r$.

To justify this approximation, we consider the difference matrix $A_{nm} = K_{x_n,x_m} - K_{x_n,x_m}^{(r)}$, given a data set $\{x_n\}_{n=1}^N$ drawn from a uniform distribution on a hypersphere of dimension d . The terms $b_q (x_n \cdot x_m)^q$ scale roughly as $d^{-q/2}$ (see Appendix K for a more accurate expression) due to the tendency of random vectors in high dimensions to be orthogonal. Consequently, the above difference diminishes very quickly with r . Notably, this also applies for the entries of $K_{x_*,x_n} - K_{x_*,x_n}^{(r)}$, provided x_* is a test point and not a train point. In contrast, the diagonal part of A is $A_{nn} = \sigma_r^2$ and may diminish more slowly depending on the coefficients $b_{q>r}$. Upon neglecting $K_{x_*,x_n} - K_{x_*,x_n}^{(r)}$ and the off-diagonal elements of A , one finds that Eq. (2) with these two GPRs yields identical predictions. As shown in Appendix K, these neglected off-diagonal elements yield a discrepancy which scales as $\sqrt{N} d^{-(r+1)/2}$. Consequently, the MSE error between the two GPRs should scale as N times an exponentially small factor (d^{-r-1}). This scaling with N should saturate when the accuracy is nearly perfect since then the predictions remain largely constant as N is increased.

Focusing back on the question of how to tackle noiseless GPR, we thus find that as long as the b_q 's decay slowly enough with q , then at any finite N we can choose a large enough r such that two desirable properties are maintained: (a) The discrepancy between the GPRs is small and (b) σ_r^2 is large enough to ensure convergence to our perturbative analysis. The required slow decay of b_q is harmonious with the intuition that DNNs should be initialized at the edge of

chaos [46] where the output of the network has a fine and multiscale sensitivity to small changes in the input. As $K_{x,x'}$ is the correlation of two outputs with inputs x and x' , having a power law decaying b_q implies such fine and multiscale sensitivity. Establishing relations between good initialization and effectiveness of our renormalized NTK is left for future work.

We have tested the accuracy of approximating noiseless NTK GPR with renormalized NTK GPR with the appropriate σ_r^2 , both on artificial data sets (see next section) and on real-world data set such as CIFAR-10 (see Appendix B). In both cases we found an excellent agreement between the two GPRs for r 's as small as 3 and 4.

VII. GENERALIZATION IN THE NTK REGIME

Collecting the results of all the preceding sections, we can obtain a detailed and clear picture of generalization in fully connected DNNs trained in the NTK regime on data sets with a uniform distribution normalized to some hypersphere in input space. We begin with a qualitative discussion and consider some renormalized NTK at scale r . From Sec. V, we have that the features of this kernel are hyperspherical harmonics and that λ_l scales as d^{-l} . We also recall that Y_{lm} is a polynomial of degree l and that all the hyperspherical harmonics up to degree l span all polynomials on the hypersphere with degree up to l . Examining Eq. (24) we find that features with $\lambda_l \gg \sigma^2/\eta$ are learnable and via the above scaling we find that we learn polynomials of degree $O[\log(\eta/\sigma^2)/\log(d)]$ or less. In particular, a function like parity, which is a polynomial of degree d is very hard to learn whereas a linear function is the easiest to learn. Thus, despite having infinitely more parameters than data points (due to infinite width) and despite being able to span almost any function (due to the richness of the kernel's features), the DNN here avoids overfitting by having a strong bias towards low-degree polynomials.

To make more quantitative statements, we now focus on a specific setting. We consider input data in dimension $d = 50$ and a scalar target function $g(x) = \sum_{l=1,2,m} g_{lm} Y_{lm}(x)$ such that the vectors $(g_{l,1}, g_{l,2}, \dots, g_{l,\deg(l)})^T$ for $l = 1, 2$ are drawn from a uniform measure on the $\deg(l)$ sphere of radius $1/\sqrt{2}$. We generate several toy data sets D_N consisting of N data points (x_n) uniformly distributed on the hypersphere S^{d-1} and their corresponding targets $[g(x_n)]$. We consider the GP equivalent to training a fully connected DNN consisting of four layers with ReLU activations and width W which we initialize with variance $(\sigma_w^2 = \sigma_b^2 = 1/d)$ for the input layer and $(\sigma_w^2 = \sigma_b^2 = 1/W)$ for the hidden layers (see for instance [38] for how to compute the kernel, and notice there is a factor of $1/W$ between our convention for σ_w^2 and [38]). To be in the NTK correspondence regime we consider training such a network at vanishing learning rate, MSE loss, and with $W \gg N$. One then has that the predictions of the DNN are given by GPR with $\sigma^2 = 0$ and the K given by the NTK kernel [8]. (To be more precise, [8] predicts correspondence with GPR up to a random initialization factor, so to get an exact match with GPR one would also need to average over initialization seeds. Recent research [38] suggests this caveat can be avoided under some conditions.)

For each such DNN we obtained the expected MSE loss $\|g_\infty^* - g\|^2$ of GPR with the NTK kernel by numerical integration over x_* . Repeating this process multiple times we obtained the DAEE for $N = 1, 2, \dots, N_{\max}$ with a relative standard error of less than 5% (this typically required averaging over 10 data sets). For direct comparison with our prediction of the learning curve, we computed the Poisson-averaged learning curve $(\|g_\infty^* - g\|^2)_\eta$ in accordance with Eq. (13), neglecting the terms $n > N_{\max}$. We restricted ourselves to $\eta_{\max} \leq N_{\max} - 5\sqrt{N_{\max}}$ to make tail effects negligible. Notably, the Poisson averaging makes the final statistical error negligible relative to the discrepancies coming from our large- η approximations (see Appendix A). To make it easier to appreciate the power of GPs over simpler regression models, we also provide the Poisson-averaged DAEE for ordinary least-squares method (OLS) as a yardstick.

To pick the renormalization scale r we must consider two factors. On the one hand, we want the discrepancy between the renormalized and regular NTK to be small, this scales as $O(\sqrt{N}d^{-(r+1)/2}/K_{x,x})$. On the other hand, we want the effective noise σ_r^2 to be as large as possible to ensure the accuracy of the prediction. We found that $r = 3$ strikes a good balance for the range of N values used in the experiment, but $r = 4, 5, 6$ also produced adequate predictions since σ_r^2 shrinks slowly with r for the architecture used.

Our analytical expressions following Eq. (23) combined with known results [8,44] about the Taylor coefficients (b_n) yield $\lambda_0, \dots, \lambda_3 = \{3.19, 7.27 \times 10^{-3}, 5.98 \times 10^{-6}, 1.62 \times 10^{-7}\}$ and $\sigma_r^2 = 0.018$. Since $\lambda_0, \lambda_1 \gg \sigma^2/\eta \gg \lambda_2, \lambda_3$ for $50 < \eta < 3500$, $C_{K_r, \sigma^2/\eta} \sigma^{-2} < [\deg(0) + \deg(1)]\sigma^2/\eta + O[\deg(2)10^{-6}]$, thus, $C_{K_r, \sigma^2/\eta} \sigma^{-2} \approx 51/\eta$. Thus, we expect perturbation theory to be valid for $\eta \gg 50$. At $\eta = 1000$ the $l = 1$ features are learned well since $\sigma^2/\eta = 1.8 \times 10^{-4} \gg \lambda_1$ and the $l = 2$ features neglected, at $\eta = 1000$ they are learned but suppressed by a factor of about 3. Had the target contained $l = 3$ features, they would have been entirely neglected at these η scales. Experimental learning curves along with our leading and subleading estimates are shown in Fig. 2 (left panel) showing an excellent agreement between theory and experiment.

While no actual DNNs were trained in the above experiments, the NTK correspondence means that this would be the exact behavior of a DNN trained in the NTK regime [8,22,38]. Furthermore, since our aim was to estimate what the DNNs would predict rather than reach state of the art predictions, we focus on reasonable hyperparameters but did not perform any hyperparameter optimization. The complementary case of noisy GPR, which one encounters in the NNSP correspondence, is studied in Appendix C.

To demonstrate that our results work with more complex functions, we repeated the experiment with a different target function which cannot be expressed as a finite-order polynomial. We drew a uniformly distributed vector w on the $(d - 1)$ sphere of radius \sqrt{d} and set the target as $g(x) = |w \cdot x|$. Figure 2 (right) shows good agreement between theory and experiment here as well.

Lastly, we argue that the asymptotic behavior of learning curve we predict is more accurate than the recent probably approximately correct (PAC) based bounds [9–11]. In Appendix D we show a log-log plot of the learning curves contrasted




FIG. 2. The experimental learning curves (solid lines) for a depth of four ReLU networks trained in the NTK regime on different target functions on a $d = 50$ hypersphere are shown along with our analytical predictions for the leading (dotted line) and leading plus subleading behavior (dashed line). Left panel shows the results for a second-order polynomial in the input, and the right panel show results for the function $|w \cdot x|$ (where w is a random vector of norm \sqrt{d}) which cannot be expressed as a finite linear combination of eigenfunctions. The learning curves of ordinary least squares (OLS) on the same regression tasks are provided to help compare the performance of GPs with simpler regression methods.

with a $1/\sqrt{\eta}$ which is the most rapidly decaying bound appearing in those works. It can be seen that such an asymptotic cannot be made to fit the experimental learning curve with any precision close to ours.

VIII. APPLICATION OF RESULTS TO HYPERPARAMETER OPTIMIZATION

As with most machine-learning algorithms, when training a neural network for a particular task one needs to choose a number of hyperparameters such as the network’s width at each level W_l , depth L , variance of weights at initialization σ_b^2, σ_w^2 , activation function, and optimizer-related parameters such as batch size, learning rate, etc. There are many considerations for hyperparameter selection such as training time, memory footprint, and optimizer convergence, but here we will focus on the expected loss of the network. While there are some accepted heuristics, there is no *a priori* way to predict the best performing architecture other than an expensive process of trial and error. In this section we introduce a scheme for picking theoretically advantageous parameters, given minimal information on the target function and data-set distribution. As recent research suggests [47] W should be increased as much as possible to put the network in the interpolation regime, we will assume that choice was made. We also note that σ_b^2, σ_w^2 are typically thought to be related more to convergence issues, for example, via the exploding and vanishing gradient problem, than to the performance of the network. However, as this work as well as [48] suggest, these parameters have an important effect on the network performance by changing the NTK spectrum.

We suggest the following scenario: we have $N = 1000$ data points uniformly distributed on S^9 . We are also given the spectral weight of the target in each eigenspace, that is $w_\ell^2 = \|\Pi_\ell(g)\|^2$ where Π_ℓ is the projection operator on the ℓ subspace. For the case of uniform measure on the hypersphere,

the projection operator is simply $\Pi_\ell = P_\ell(\langle xy \rangle) \text{deg}(\ell)$ where P_ℓ is the Legendre polynomial of degree ℓ and $\text{deg}(\ell)$ is the dimension of the eigenspace, so finding w_ℓ^2 is a much simpler task than finding the $\text{deg}(\ell)$ coefficients of the target (which scale as d^ℓ) and can be accomplished with a few numeric integrals. In this case, we focus on a target with $w_\ell^2 = \frac{1}{3}(\delta_{1,\ell} + \delta_{2,\ell} + \delta_{3,\ell})$. Given this setting, we would like to find a network architecture with minimal expected error. For computational efficiency reasons, we decide to focus on ReLU networks with one hidden layer, so we need to choose four hyperparameters $\sigma = (\sigma_{w_1}, \sigma_{w_2}, \sigma_{b_1}, \sigma_{b_2})$.

We present two typical ways used to choose σ , then propose a better way to do so based our theory. The naive and most prevalent way to choose hyperparameters is to simply take $\sigma_{\text{typical}} = (\sqrt{2}, \sqrt{2}, 0.05, 0.05)$ which roughly correspond to He initialization [49], a common heuristic for avoiding gradient propagation issues. A more diligent approach would be to draw some random values in the vicinity of σ_{typical} , train the network, evaluate the test loss, and pick the best performing hyperparameters σ_{best} .

Next, we suggest a different approach which utilizes our analytical results. We construct a symbolic expression for the expected loss using the formalism outlined in the paper. By taking $\eta = N$ and applying the renormalization scheme with appropriate r we get an estimator for the expected loss $\hat{L}(\sigma) = \int dx \langle [f(x) - g(x)]^2 \rangle$ which we can use to predict the performance of different hyperparameters without training the network. Moreover, we can use standard numerical optimization algorithms to minimize the predicted loss and obtain $\sigma_{\text{optimized}} = \text{argmin}_\sigma \hat{L}(\sigma)$.

To test the scheme experimentally, we drew 21 random hyperparameter proposals $\{\sigma_i\}$ uniformly distributed in the hyper-rectangle defined by $\frac{1}{2}\sigma_{\text{typical}}, \frac{3}{2}\sigma_{\text{typical}}$. We constrained the optimization algorithm to this hyper-rectangle as well to avoid convergence issues in the training procedure, as unconstrained optimization leads to very small values of $\sigma_w \ll \sqrt{2}$ which in turn lead to vanishing gradients and

TABLE I. Comparison of the performance of networks trained in the NTK regime using different hyperparameters (σ). The “Test” column shows an estimate of the DNN test loss, the “Prediction” column the loss as predicted by our learning-curve, and the “GPR” column an estimate of the data-set-averaged expected loss of the corresponding GP. Worst, median, and best refer to one of 21 networks with random hyperparameters ranked by test loss. Typical and optimized refer to networks with σ_{typical} (defined in the text) and the optimal hyperparameters following our optimization scheme. For more experiment results see Appendix L.

	Test	Prediction	GPR
Worst	0.413	0.406	0.382
Median	0.313	0.316	0.319
Best	0.175	0.198	0.214
Typical	0.307	0.307	0.317
Optimized	0.078	0.110	0.141

long training times. We defined networks corresponding to $\sigma_{\text{optimized}}$, σ_{typical} , $\{\sigma_i\}$ and trained each network on the same data set with full-batch gradient descent and learning rate 1.0 until the train loss was smaller than $0.1\sigma_r^2$. A summary of the experiment is outlined in Table I.

The results clearly demonstrate the effectiveness of our scheme which reduces the test loss by a factor of 4 relative to the typical hyperparameter choice and a factor of 2 over the best performing random hyperparameters. In terms of computational complexity, it took approximately 2.5 h to train each network using Google’s neural tangent package [50] on a 20-core CPU¹ with $W = 2^{14}$. In comparison, the time it takes to build and optimize \hat{L} is completely negligible at about 30 s. The best random hyperparameters were found on the 15th attempt, so had we stopped then we would have wasted 35 computer hours relative to our scheme and gotten inferior test loss. Note also that we focused on $L = 2$ in order to speed up training, which scales exponentially with depth, but the optimization procedure is not nearly as sensitive to depth and could have been done for any reasonable L . Moreover, increasing the depth would have also enlarged the hyperparameter space, making random search even less effective. For each network, we also experimentally obtained the data-set-averaged expected loss using GPR with the associated NTK. The fair agreement between the test loss and data-set-averaged expected loss (GPR in Table I) further solidifies previous results and demonstrates our claim of self-averaging.

As expected, the above results required some knowledge of the target function, in particular its spectral weight within each angular momentum space. Alternatively, one can capitalize on the fact that our learning-curves prediction are quadratic in the target, average them over a target function ensemble, and optimize with respect to this average case. Another option is to consider a min-max optimization scheme in which hyperparameters are optimized for the worst-case target within some

domain. The scheme can also be extended to nonuniform data sets and different activation functions as long as some way of computing the eigenvalues is provided.

IX. DISCUSSION AND OUTLOOK

In this work we laid out a formalism based on field theory tools for predicting learning curves in the NTK and NNSP correspondence regimes. Despite DNNs’ black-box reputation, well within the validly range of our perturbative analysis, we obtained very low relative mismatch between our best estimate and the experimental curves, with good agreement extending well into regions with low amounts of data compared to that needed to learn the target. A potential use of such learning curves in hyperparameter optimization was also demonstrated.

Central to our analysis was a renormalization group transformation leading to effective observation noise on the target and to a simpler renormalized quadratic action or kernel. Notably, this RG transformation implied that wide fully connected networks, even ones working on real-world data sets such as CIFAR-10, could be effectively described by very few parameters being the noise level and the $O(1)$ -first Taylor expansion parameters of the kernel.

Our analysis provides a laboratory setting in which deep learning can be understood. In its training phase, DNNs avoid local-minima issues and glassy behavior due to their high overparametrization which makes the optimization problem highly underconstrained [13,51,52]. As a result, many different solutions or weights which fit perfectly the training data are possible. While each such solution will behave differently on a test point, this arbitrariness does not entail an erratic behavior. The reason is the implicit bias DNNs have towards simple functions. In the case of the NNSP correspondence, a simple function is, by definition, a function that can be generated, up to some small noise, by a large phase space of weights.

Simplicity is therefore strongly architecture and data set dependent. For fully connected DNNs trained in the regime of the NTK or the NNSP correspondences, as well as data uniformly sampled from the hypersphere, simplicity amounts to low-order polynomials over that hypersphere. These are the hyperspherical harmonics with low l , which are the leading eigenfunctions with respect to such a uniform measure of a generic kernel associated with a fully connected DNN. As long as the DNN has at least one nonlinear layer and biases, depth has only a quantitative effect as it modifies the eigenvalues (λ_l) but does not change their scale. Generally, the eigenvalues and eigenfunctions vary with architecture and data distribution. Convolutional neural networks (CNNs) require further study, however, one can argue on a qualitative level that simple functions would be polynomials with certain spatial hierarchy. Moreover, one expects that qualitative details of this hierarchy would depend on depth as it controls the input-fan-in of the hidden activation in the last CNN layer.

It seems unrealistic that a purely analytical approach such as ours would describe well the predictions of state-of-the-art DNNs such as VGG-19 trained on a real-world data set such as ImageNet. Similarly unrealistic is to expect an analytical

¹While GPUs are generally faster, fully connected DNNs do not gain the full benefit of GPU parallelism and we expect the computation time would only improve by a factor of $O(1)$.

computation based on thermodynamics to capture the efficiency of a modern car engine or one based on Navier-Stoke’s equations to output a better shape of a wing. Still, scientific experience shows that understanding toy models, especially rich enough ones, has value. Indeed, toy models provide an analytical laboratory where theories could be refined or refuted, algorithms could be benchmarked and improved, and wider ranging conjectures and intuitions could be formed. Such models are useful whenever domain knowledge possesses some degree of universality or independence from detail. In converse, when all details matter, knowledge is nothing more than a log of all experiences. The fact that DNNs work well in variety of different architecture and data-set settings suggests that some degree of universality worth exploring is present. Further research would thus tell if the tools and methodologies that have enabled us to comprehend our physical world can help us comprehend the artificial world of deep learning.

Many extensions of this work, aimed at approaching real-world settings, can be considered. First and foremost, much of the recent excitement about DNNs comes from either CNNs or long short-term memory networks (LSTMs). Considering CNNs, while much of our formalism applies, the spectrum of CNN kernels is more challenging to obtain as their kernels are less symmetric compared to fully connected DNNs. For similar reasons the RG approach presented here requires a more elaborate trimming of the CNN kernel since the latter would not consist of only powers of dot products. Furthermore, CNNs trained with SGD show rather large gaps in performance compared to their NNGP or NTK. The culprit here might very well be the finite-width or finite-number-of-channels corrections to the NNGP or NTK priors. Leading finite-width corrections, considered in Ref. [13], amount to adding quartic terms to $P_0[f]$. Those could be dealt with straightforwardly using our perturbation theory formalism. Interestingly, at least for CNNs without pooling, these corrections introduce a qualitative change to the prior, making it reflect the weight-sharing property of CNNs, which is lost at the level of the NNGP or NTK [13,53]. Other viable directions are handling richer data-set distributions, extending EK results to the more common cross-entropy loss, applying RG reasoning on finite-width DNNs, and using the above insights for developing DNN architecture design principles.

ACKNOWLEDGMENT

Z.R. and O.M. acknowledge support from ISF Grant No. 2250/19.

APPENDIX A: POISSON-AVERAGING DEMONSTRATION

Here we demonstrate that Poisson averaging has no substantial effect on the learning curve. To this end, Fig. 3 shows the experimental learning curve from the main text preaveraging and postaveraging. It is evident that other than the unintended consequence of eliminating the experimental noise, the averaged learning curve is equivalent to the original for all practical intents.




FIG. 3. Data-set-averaged generalization before (gray) and after (blue) Poisson averaging.

APPENDIX B: COMPARISON OF NTK AND RENORMALIZED NTK PREDICTIONS ON SYNTHETIC AND REAL-WORLD DATA SETS

In Sec. VII we used the renormalized NTK as a proxy for the regular NTK; the purpose of this Appendix is to affirm the validity of this approximation. Moreover, while our lack of knowledge of the NTK eigenvalues and eigenfunctions with respect to a nonuniform measure prevents us from predicting learning curves, we would like to show that the renormalized NTK is a valid approximation in this setting as well.

To this end, we used the following procedure. We took the NTK kernel defined in the paper and its associated renormalized kernels at different scales and trained them over the same training set D_N . In Fig. 4 (top) the training set and target function were the ones defined in the main text. In Fig. 4 (bottom) D_N consisted of uniform draws without replacements from the CIFAR-10 training set, standardized and normalized to unit vectors, and the target function was the one hot encoding of the labels standardized to have zero mean and $K_{y,x}$ variance.

For each training set we logged the average squared deviation of each renormalized kernel estimation g_r^* from the estimation of the nonrenormalized kernel g_∞^* . This is the quantity $\|g_r^* - g_\infty^*\|^2$ (where in the CIFAR-10 case $\|\cdot\|$ implies both the Euclidean norm in \mathbb{R}^{10} and integration over the input measure, which we approximated by averaging over the CIFAR-10 test set). We averaged this quantity over different draws of training sets to obtain $\langle \|g_r^* - g_\infty^*\|^2 \rangle_{D_N}$. The results show good agreement between g_∞^* and g_r^* as r is increased.

APPENDIX C: LEARNING CURVES IN THE NNSP PROTOCOL

We report here the results of a similar experiment to the one presented in the main text, but with the NTK kernel replaced with the NNGP kernel as appropriate for the NNSP correspondence. In this case we used a kernel simulating a network with a single hidden layer and $\sigma_w^2 = 1/W, \sigma_b^2 = 0$, and a target function equivalent to the one in the main text. In the NNSP protocol the renormalization group approach is not necessary to introduce noise to the observations, as it comes

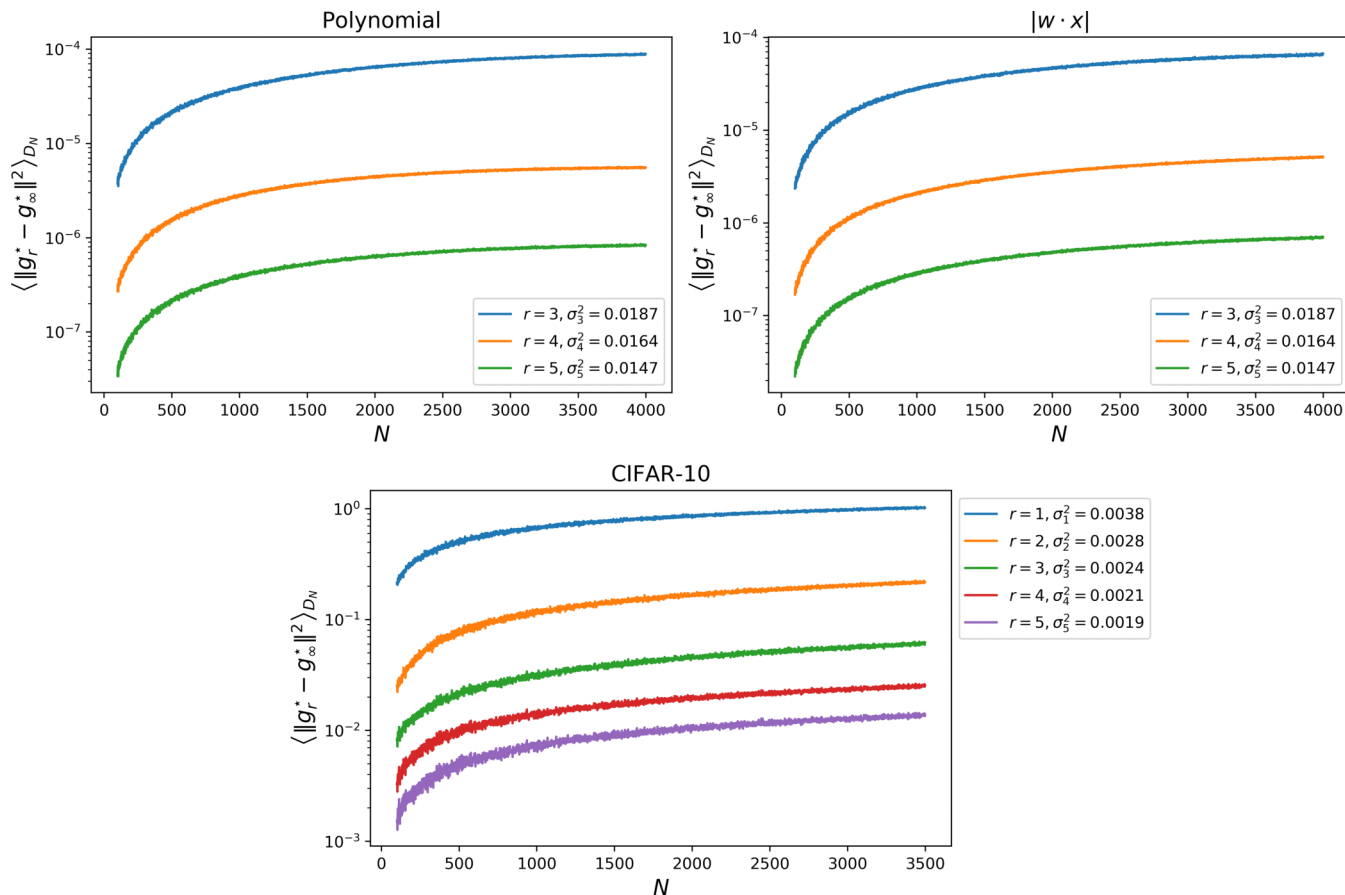


FIG. 4. Expected deviation between renormalized and nonrenormalized kernel predictions in terms of the squared norm induced by the input measure. The deviation is averaged over different training set draws.

into play naturally via the temperature-dependent fluctuations, so we can choose arbitrary σ^2 . Notwithstanding, the renormalization group approach can aid in analyzing low-temperature behavior.

Notice, in Fig. 5, that the subleading prediction significantly improves upon the EK prediction. As the inset plot demonstrates, when the data-set size is small the expected error actually increases. Surprisingly, the subleading correction manages to capture this behavior even though the data-set size is small, demonstrating its superiority.




FIG. 5. Experimental data-set-averaged learning curves for regression with NNGP kernel, along with leading and subleading predictions.

APPENDIX D: COMPARISON WITH RECENT BOUNDS

As mentioned in the main text, various recent bounds, relevant to the NTK regime, have been derived recently. Notwithstanding importance and rigor of these works, their bounds have at best a $1/\sqrt{N}$ asymptotic scaling. Figure 6 shows that given a functional behavior of the experimental




FIG. 6. Asymptotic comparison of our prediction for the learning curve with the recent lower bound of $1/\sqrt{N}$.




FIG. 7. Correlation plots between GPR loss, DNN test loss, and our prediction. The correlation between GPR and DNN test loss (middle) affirms the relation between GPR and neural networks in the NTK regime, while the plots affirm the relation between our predictions and experimentally obtained quantities.

learning curves such a bound cannot be nearly as tight as our predictions.

APPENDIX E: NNGP AND NTK ARE ROTATIONALLY INVARIANT

Let us prove that the NNGP and NTK kernels associated with any network whose first layer is fully connected are rotationally invariant. Indeed, let $h_w(x)$ be the output vector of the first layer $[h_w(x)]_i = \phi(\sum_j w_{ij}x_j + b)$ where x_j is the j th component of the input vector x . Let $z_{w'}(h)$ be the output of the rest of the network given h . The covariance function of NNGPs are defined by [44]

$$K(x, y) = \int dw dw' P_0(w, w') z_{w'}(h_w(x)) z_{w'}(h_w(y)), \tag{E1}$$

where $P_0(w, w')$ is a prior over the weights, typically taken to be independent and identically distributed Gaussian for each layer $[P_0(w, w') = P_0(w)P_0(w')$ and $P_0(w) \propto e^{-\sum_{ij} w_{ij}^2/(2\sigma^2)}$]. Following this one can show

$$\begin{aligned} K(Ox, Oy) &= \int dw dw' P_0(w, w') z_{w'}(h_w(Ox)) z_{w'}(h_w(Oy)) \\ &= \int dw dw' P_0(w, w') z_{w'}(h_{O^T w}(x)) z_{w'}(h_{O^T w}(y)) \\ &= \int dw dw' P_0(Ow, w') z_{w'}(h_w(x)) z_{w'}(h_w(y)) \\ &= \int dw dw' P_0(w, w') z_{w'}(h_w(x)) z_{w'}(h_w(y)) \\ &= K(x, y), \end{aligned} \tag{E2}$$

where the second equality uses the definition of $h_w(x)$, the third results from an orthogonal change of integration variable $w \rightarrow O^T w$, and the fourth is a property of our prior over

w . Since the NTK relates to the NNGP kernel in a recursive manner [8], it inherits that symmetry as well.

APPENDIX F: NOTATIONS FOR THE FIELD THEORY DERIVATION

For completeness, here we restate the notations used in the main text.

- (i) x, x', x^* : inputs.
- (ii) μ_x : measure on input space.
- (iii) $K(x, x')$: kernel function (covariance) of a Gaussian process; assumed to be symmetric and positive semidefinite.
- (iv) $\phi_i(x)$: i th eigenfunction of $K(x, x')$. By the spectral theorem, the set $\{\phi_i\}_{i=1}^\infty$ can be assumed to be orthonormal,

$$\int d\mu_x \phi_i(x) \phi_j(x) = \delta_{ij}.$$

- (v) λ_i : i th eigenvalue of $K(x, x')$,

$$\int d\mu_{x'} K(x, x') \phi_i(x') = \lambda_i \phi_i(x).$$

- (vi) $\|\cdot\|_K$: RKHS norm

$$\|\cdot\|_K = \int d\mu_x d\mu_{x'} f(x) K^{-1}(x, x') f(x').$$

- (vii) If $f(x) = \sum_i f_i \phi_i(x)$, then $\|f\|_K = \sum_i \frac{f_i^2}{\lambda_i}$ (where ϕ_i is an orthonormal set). Note that this norm is independent of μ_x [12].
- (viii) $g(x)$: The target function.
- (ix) σ^2 : Noise variance.
- (x) N : Number of inputs in the data set.
- (xi) D_N : Data set of size N , $D_N = \{x_1, \dots, x_N\}$.
- (xii) g^* : The prediction function.

APPENDIX G: PHRASING THE PROBLEM AS A FIELD THEORY PROBLEM

1. Without data

We start by establishing the exact equivalence between a prior of a centered GP and the corresponding partition function. For a kernel function K , let us define the partition function

$$Z[\alpha] = \int \mathcal{D}f \exp\left(-\frac{1}{2}\|f\|_K^2 + \int dx \alpha(x)f(x)\right). \tag{G1}$$

Since the RKHS norm is quadratic in f , the distribution over the space of functions induced by Z is Gaussian (a GP). Since a GP is determined by its mean and kernel, it is sufficient to show those equalities. For the mean we get

$$\langle f(x_*) \rangle = \frac{\delta \log(Z[\alpha])}{\delta \alpha(x_*)} \Big|_{\alpha=0} = \frac{\int \mathcal{D}f f(x_*) \exp\left(-\frac{1}{2}\|f\|_K^2\right)}{\int \mathcal{D}f \exp\left(-\frac{1}{2}\|f\|_K^2\right)} = \arg \min \left[\frac{1}{2}\|f\|_K^2 \right] \Big|_{x_*} = 0 \tag{G2}$$

since for Gaussian distributions it holds that the average case is also the most probable case. For the covariance we get

$$\begin{aligned} \langle f(x)f(y) \rangle &= \frac{\delta^2 \log(Z_0[\alpha])}{\delta \alpha(x)\delta \alpha(y)} \Big|_{\alpha=0} = \frac{\int \mathcal{D}f f(x)f(y) \exp\left(-\frac{1}{2}\|f\|_K^2\right)}{\int \mathcal{D}f \exp\left(-\frac{1}{2}\|f\|_K^2\right)} = \frac{\int \prod_i df_i \sum_i f_i \phi_i(x) \sum_j f_j \phi_j(y) \exp\left(-\frac{1}{2} \sum_l \frac{f_l^2}{\lambda_l}\right)}{\int \prod_i df_i \exp\left(-\frac{1}{2} \sum_l \frac{f_l^2}{\lambda_l}\right)} \\ &= \sum_i \underbrace{\frac{\int df f^2 \exp\left(-\frac{f^2}{2\lambda_i}\right)}{\int df \exp\left(-\frac{f^2}{2\lambda_i}\right)}}_{\lambda_i} \phi_i(x)\phi_i(y) + \sum_{i \neq j} \underbrace{\frac{\int df f \exp\left(-\frac{f^2}{2\lambda_i}\right)}{\int df \exp\left(-\frac{f^2}{2\lambda_i}\right)}}_0 \underbrace{\frac{\int df f \exp\left(-\frac{f^2}{2\lambda_j}\right)}{\int df \exp\left(-\frac{f^2}{2\lambda_j}\right)}}_0 \\ &= \sum_i \lambda_i \phi_i(x)\phi_i(y) = K(x, y). \end{aligned} \tag{G3}$$

Indeed, Z is the partition function corresponding to a centered GP with kernel K .

2. With data

We continue by establishing the exact equivalence between Bayesian inference on a GP and the corresponding partition function. From (G1) we get that

$$P[f] \propto \exp\left(-\frac{1}{2}\|f\|_K^2\right). \tag{G4}$$

For given target function g and a sampled data point $(x_1, g(x_1))$, assuming that f is our prediction it holds that $g(x_i) \sim \mathcal{N}(f(x_i), \sigma^2)$ since g distributes normally around f with variance σ^2 . Therefore, $p(g(x_i)|f) \propto \exp(-[g(x_i) - f(x_i)]^2/2\sigma^2)$, so

$$P[D|f] = \prod_{i=1}^N p(g(x_i)|f, M) \propto \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^N [g(x_i) - f(x_i)]^2\right) \tag{G5}$$

and using Bayes' theorem we get

$$P[f|D] \propto \exp\left(-\frac{1}{2}\|f\|_K^2 - \frac{1}{2\sigma^2} \sum_{i=1}^N [g(x_i) - f(x_i)]^2\right) \tag{G6}$$

which gives rise to the posterior partition function

$$Z[\alpha] = \int \mathcal{D}f \exp\left(-\frac{1}{2}\|f\|_K^2 - \frac{1}{2\sigma^2} \sum_{i=1}^N [g(x_i) - f(x_i)]^2 + \int dx \alpha(x)f(x)\right) \tag{G7}$$

and again, the exponent is quadratic in f leading to a Gaussian distribution over the space of functions. Indeed, for the mean we get

$$g^*(x_*) = \langle f(x_*) \rangle = \frac{\delta \log(Z[\alpha])}{\delta \alpha(x_*)} \Big|_{\alpha=0} = \arg \min \left[\frac{1}{2}\|f\|_K^2 + \frac{1}{2\sigma^2} \sum_{i=1}^N [f(x_i) - g(x_i)]^2 \right] \Big|_{x_*} \tag{G8}$$

in agreement with [12].

3. Calculating observables

a. Averaging g^*

Applying the replica trick to Eq. (G8) and averaging over all the data sets of size N we obtain

$$\langle g^*(x_*) \rangle_{D_N} = \lim_{M \rightarrow 0} \frac{1}{M} \left. \frac{\delta \langle Z^M[\alpha] \rangle_{D_N}}{\delta \alpha(x_*)} \right|_{\alpha=0}, \tag{G9}$$

for integer M we get

$$Z^M[\alpha] = \underbrace{\int \dots \int}_M \prod_{j=1}^M \mathcal{D}f_j \exp \left(-\frac{1}{2} \sum_{j=1}^M \|f_j\|_K^2 - \sum_{j=1}^M \sum_{i=1}^N \frac{[f_j(x_i) - g(x_i)]^2}{2\sigma^2} + \sum_{j=1}^M \int \alpha(x) f_j(x) dx \right) \tag{G10}$$

and after averaging

$$\langle Z^M[\alpha] \rangle_{D_N} = \underbrace{\int \dots \int}_M \prod_{j=1}^M \mathcal{D}f_j \exp \left(-\frac{1}{2} \sum_{j=1}^M \|f_j\|_K^2 + \sum_{j=1}^M \int \alpha(x) f_j(x) dx \right) \left\langle \exp \left(-\sum_{j=1}^M \frac{[f_j(x) - g(x)]^2}{2\sigma^2} \right) \right\rangle_{x \sim \mu}, \tag{G11}$$

where $\langle \dots \rangle_{x \sim \mu} = \int \dots d\mu_x$.

Performing the Poissonic averaging we get

$$\begin{aligned} \langle Z^M[\alpha] \rangle_\eta &= e^{-\eta} \sum_{N=0}^\infty \frac{\eta^N}{N!} \langle Z^M[\alpha] \rangle_{D_N} \\ &= \underbrace{\int \dots \int}_{M \text{ times}} \mathcal{D}f_1 \dots \mathcal{D}f_M \exp \left[-\frac{1}{2} \sum_{j=1}^M \|f_j\|_K^2 + \sum_{j=1}^M \int \alpha(x) f_j(x) dx + \eta \left\langle \exp \left(-\sum_{j=1}^M \frac{[f_j(x) - g(x)]^2}{2\sigma^2} \right) - 1 \right\rangle_{x \sim \mu} \right], \end{aligned} \tag{G12}$$

so overall

$$\langle g^*(x_*) \rangle_\eta = \lim_{M \rightarrow 0} \frac{1}{M} \left. \frac{\delta \langle Z^M[\alpha] \rangle_\eta}{\delta \alpha(x_*)} \right|_{\alpha=0}. \tag{G13}$$

b. Averaging g^{*2}

From (G9) we get that

$$\langle g^{*2}(x_*) \rangle_{D_N} = \lim_{M \rightarrow 0} \lim_{W \rightarrow 0} \frac{1}{MW} \left. \frac{\delta^2 \langle Z^M[\alpha] Z^W[\beta] \rangle_{D_N}}{\delta \alpha(x_*) \delta \beta(x_*)} \right|_{\alpha, \beta=0}. \tag{G14}$$

Therefore,

$$\langle g^{*2}(x_*) \rangle_\eta = \lim_{M \rightarrow 0} \lim_{W \rightarrow 0} \frac{1}{MW} \left. \frac{\delta^2 \langle Z^M[\alpha] Z^W[\beta] \rangle_\eta}{\delta \alpha(x_*) \delta \beta(x_*)} \right|_{\alpha, \beta=0}. \tag{G15}$$

APPENDIX H: EQUIVALENCE KERNEL AS FREE THEORY

Expanding the nested exponent in Eq. (G12) using (first-order) Taylor series we get

$$\begin{aligned} \langle Z^M[\alpha] \rangle_\eta &= e^{-\eta} \underbrace{\int \dots \int}_M \prod_{j=1}^M \mathcal{D}f_j \exp \left[-\frac{1}{2} \sum_{j=1}^M \|f_j\|_K^2 + \sum_{j=1}^M \int \alpha(x) f_j(x) dx + \eta \left\langle \exp \left(-\sum_{j=1}^M \frac{[f_j(x) - g(x)]^2}{2\sigma^2} \right) \right\rangle_{x \sim \mu} \right] \\ &= \underbrace{\int \dots \int}_M \prod_{j=1}^M \mathcal{D}f_j \exp \left(-\frac{1}{2} \sum_{j=1}^M \|f_j\|_K^2 + \sum_{j=1}^M \int \alpha(x) f_j(x) dx - \eta \left\langle \sum_{j=1}^M \frac{[f_j(x) - g(x)]^2}{2\sigma^2} \right\rangle_{x \sim \mu} \right) + O(1/\eta^2) \\ &= \left[\int \mathcal{D}f \exp \left(-\frac{1}{2} \|f\|_K^2 + \int \alpha(x) f(x) d\mu_x - \frac{\eta}{2\sigma^2} \int d\mu_x [f(x) - g(x)]^2 \right) \right]^M + O(1/\eta^2) \\ &= (Z_{EK}[\alpha])^M + O(1/\eta^2), \end{aligned} \tag{H1}$$

where we defined

$$Z_{EK}[\alpha] \stackrel{\text{def}}{=} \int \mathcal{D}f \exp \left(-\frac{1}{2} \|f\|_K^2 + \int \alpha(x) f(x) d\mu_x - \frac{\eta}{2\sigma^2} \int d\mu_x [f(x) - g(x)]^2 \right) \tag{H2}$$

under this approximation we get

$$\lim_{M \rightarrow 0} \frac{\langle Z^M[\alpha] \rangle_\eta - 1}{M} = \lim_{M \rightarrow 0} \frac{(Z_{EK}[\alpha])^M - 1}{M} + O(1/\eta^2) = \log(Z_{EK}[\alpha]) + O(1/\eta^2). \tag{H3}$$

Denoting the average with respect to Z_{EK} as $\langle \dots \rangle_0$. The mean of the distribution induced by Z_{EK} is

$$\begin{aligned} \langle f(x_*) \rangle_0 &= \frac{\delta \log(Z_{EK}[\alpha])}{\delta \alpha(x_*)} \Big|_{\alpha=0} = \frac{\int \mathcal{D}f f(x_*) \exp\left(-\frac{1}{2}\|f\|_K^2 - \frac{\eta}{2\sigma^2} \int d\mu_x [f(x) - g(x)]^2\right)}{\int \mathcal{D}f \exp\left(-\frac{1}{2}\|f\|_K^2 - \frac{\eta}{2\sigma^2} \int d\mu_x [f(x) - g(x)]^2\right)} \\ &= \arg \min \left[\frac{1}{2}\|f\|_K^2 + \frac{\eta}{2\sigma^2} \int d\mu_x [f(x) - g(x)]^2 \right] \Big|_{x_*} = \delta_{EK,\eta}^*(x_*), \end{aligned} \tag{H4}$$

where the last equality is due to [12].

The covariance induced by Z_{EK} is

$$\begin{aligned} \text{Cov}_0[f(x), f(y)] &= \langle f(x)f(y) \rangle_0 - \langle f(x) \rangle_0 \langle f(y) \rangle_0 = \frac{\delta^2 \log(Z_{EK}[\alpha])}{\delta \alpha(x) \delta \alpha(y)} \Big|_{\alpha=0} \\ &\stackrel{*}{=} \frac{\int \mathcal{D}f f(x)f(y) \exp\left(-\frac{1}{2}\|f\|_K^2 - \frac{\eta}{2\sigma^2} \int d\mu_x f^2(x)\right)}{\int \mathcal{D}f \exp\left(-\frac{1}{2}\|f\|_K^2 - \frac{\eta}{2\sigma^2} \int d\mu_x f^2(x)\right)} \\ &\stackrel{**}{=} \frac{\int \prod_i df_i \sum_{i,j} f_i f_j \phi_i(x) \phi_j(y) \exp\left[-\frac{1}{2} \sum_i \left(\frac{1}{\lambda_i} + \frac{\eta}{\sigma^2}\right) f_i^2\right]}{\int \prod_i df_i \exp\left[-\frac{1}{2} \sum_i \left(\frac{1}{\lambda_i} + \frac{\eta}{\sigma^2}\right) f_i^2\right]} \\ &\stackrel{***}{=} \sum_i \frac{\int df f^2 \exp\left[-\frac{1}{2} \left(\frac{1}{\lambda_i} + \frac{\eta}{\sigma^2}\right) f^2\right]}{\int df \exp\left[-\frac{1}{2} \left(\frac{1}{\lambda_i} + \frac{\eta}{\sigma^2}\right) f^2\right]} \phi_i(x) \phi_i(y) \\ &= \sum_i \left(\frac{1}{\lambda_i} + \frac{\eta}{\sigma^2}\right)^{-1} \phi_i(x) \phi_i(y), \end{aligned} \tag{H5}$$

where in (*) the noncentered part of the distribution was deleted, in (**) the eigenfunctions of K were chosen as a base for the path integration, and in (***) we used the fact that $\int df f \exp\left[-\frac{1}{2} \left(\frac{1}{\lambda_i} + \frac{\eta}{\sigma^2}\right) f^2\right] = 0$ since it is the mean of a centered (unnormalized) Gaussian distribution.

For a rotationally invariant kernel, the eigenfunctions are Y_{lm} and the eigenvalues are λ_l (independent of m) so Eq. (H5) becomes

$$\text{Cov}_0[f(x), f(y)] = \sum_l \left(\frac{1}{\lambda_l} + \frac{\eta}{\sigma^2}\right)^{-1} \sum_m Y_{lm}(x) Y_{lm}(y) \tag{H6}$$

and the variance

$$\text{Var}_0[f(x)] = \sum_l \left(\frac{1}{\lambda_l} + \frac{\eta}{\sigma^2}\right)^{-1} \underbrace{\sum_m Y_{lm}^2(x)}_{\text{deg}(l)} \stackrel{\text{def}}{=} C_{K,\sigma^2/\eta}, \tag{H7}$$

which is a constant (independent of x).

APPENDIX I: NEXT-ORDER CORRECTION

We now wish to perform the first-order correction to the free theory. Expanding Eq. (G12) to the next order [keeping terms up to $O(1/\eta^2)$] we get

$$\begin{aligned} \langle Z^M[\alpha] \rangle_\eta &= \underbrace{\int \dots \int}_{M \text{ times}} \mathcal{D}f_1 \dots \mathcal{D}f_M \exp\left(-\frac{1}{2} \sum_{j=1}^M \|f_j\|_K^2 + \sum_{j=1}^M \int \alpha(x) f_j(x) dx + \eta \left\langle -\sum_{j=1}^M \frac{[f_j(x) - g(x)]^2}{2\sigma^2} \right\rangle_{x \sim \mu_x}\right) \\ &\quad \times \exp\left[\frac{\eta}{2} \left\langle \left(\sum_{j=1}^M \frac{[f_j(x) - g(x)]^2}{2\sigma^2}\right)^2 \right\rangle_{x \sim \mu_x}\right] + O(1/\eta^3) \\ &= \underbrace{\int \dots \int}_{M \text{ times}} \mathcal{D}f_1 \dots \mathcal{D}f_M \exp\left[\sum_{i=1}^M \left(-\frac{1}{2}\|f_i\|_K^2 + \int \alpha(x) f_i(x) dx - \eta \left\langle \frac{[f_i(x) - g(x)]^2}{2\sigma^2} \right\rangle_{x \sim \mu}\right)\right] \end{aligned}$$

$$\begin{aligned}
 & \times \exp \left(\frac{\eta}{8\sigma^4} \sum_{i=1}^M \sum_{j=1}^M [f_j(x) - g(x)]^2 [f_i(x) - g(x)]^2 \right) + O(1/\eta^3) \\
 &= \underbrace{\int \cdots \int}_{M \text{ times}} \mathcal{D}f_1 \dots \mathcal{D}f_M \exp \left[\sum_{i=1}^M \left(-\frac{1}{2} \|f_i\|_K^2 + \int \alpha(x) f_i(x) dx - \eta \left\langle \frac{[f_i(x) - g(x)]^2}{2\sigma^2} \right\rangle_{x \sim \mu} \right) \right] \\
 & \times \left(1 + \frac{\eta}{8\sigma^4} \sum_{i=1}^M \sum_{j=1}^M [f_j(x) - g(x)]^2 [f_i(x) - g(x)]^2 \right) + O(1/\eta^3). \tag{11}
 \end{aligned}$$

1. Calculating $\langle g^* \rangle_\eta$

We now wish to calculate the correction to $\langle g^* \rangle_\eta$ given by Eq. (11). From Eq. (G13) we get

$$\langle g^* \rangle_\eta = g_{EK,\eta}^*(x_*) + \lim_{M \rightarrow 0} \frac{1}{M} \frac{\eta}{8\sigma^4} \int d\mu_x \left\langle \sum_{j=1}^M \sum_{l=1}^M \sum_{i=1}^M [f_j(x) - g(x)]^2 [f_l(x) - g(x)]^2 f_i(x_*) \right\rangle_{f_1, \dots, f_M \sim EK} + O(1/\eta^3). \tag{12}$$

Simplifying the average of the multiple sums we get

$$\begin{aligned}
 & \left\langle \sum_{j=1}^M \sum_{l=1}^M \sum_{i=1}^M [f_j(x) - g(x)]^2 [f_l(x) - g(x)]^2 f_i(x_*) \right\rangle_{f_1, \dots, f_M \sim EK} \\
 &= M \langle [f(x) - g(x)]^4 f(x_*) \rangle_0 \\
 & \quad + M(M-1) \langle [f(x) - g(x)]^2 \rangle_0 \langle [f(x) - g(x)]^2 f(x_*) \rangle_0 + \langle [f(x) - g(x)]^4 \rangle_0 \langle f(x_*) \rangle_0 \\
 & \quad + M(M-1)(M-2) \langle [f(x) - g(x)]^2 \rangle_0^2 \langle f(x_*) \rangle_0. \tag{13}
 \end{aligned}$$

Since f has a Gaussian distribution ($f \sim EK$), such averages can be calculated using Feynman diagrams.

Let us denote $f(x) - g(x)$ by filled squares and $f(x_*)$ by filled circles. Since our free theory is not centered ($\langle f \rangle_0 = g_{EK,\eta}^* \neq 0$), we allow edges in the diagrams to be connected at only one side, representing the average of the vertex with respect to the EK distribution. An edge connected to vertices on both sides represents the covariance. Note that since we divide by M and take the limit $M \rightarrow 0$, we do not care about diagrams which are not connected to $f(x_*)$ since they scale as M^2 .

Calculating the averages we get

$$\begin{aligned}
 \langle (f(x) - g(x))^4 f(x_*) \rangle_0 &= \text{Diagram 1} + \text{Diagram 2} + \text{disconnected diagrams} \\
 &= 12 (g_{EK,\eta}^*(x) - g(x)) \text{Var}_0 [f(x)] \text{Cov}_0 [f(x), f(x_*)] + 4 (g_{EK,\eta}^*(x) - g(x))^3 \text{Cov}_0 [f(x), f(x_*)] \\
 & \quad + \text{disconnected diagrams} \tag{14}
 \end{aligned}$$

$$\begin{aligned}
 \langle (f(x) - g(x))^2 \rangle_0 \langle (f(x) - g(x))^2 f(x_*) \rangle_0 &= \left(\text{Diagram 3} + \text{Diagram 4} \right) \left(\text{Diagram 5} + \text{disconnected diagrams} \right) \\
 &= 2 \text{Cov}_0 [f(x), f(x_*)] (g_{EK,\eta}^*(x) - g(x))^3 + 2 \text{Var}_0 [f(x)] \text{Cov}_0 [f(x), f(x_*)] (g_{EK,\eta}^*(x) - g(x)) \\
 & \quad + \text{disconnected diagrams} \tag{15}
 \end{aligned}$$

$$\langle [f(x) - g(x)]^4 \rangle_0 \langle f(x_*) \rangle_0 = \text{disconnected diagrams}, \tag{16}$$

$$\langle [f(x) - g(x)]^2 \rangle_0^2 \langle f(x_*) \rangle_0 = \text{disconnected diagrams}. \tag{17}$$

Taking the limit $M \rightarrow 0$ and summing everything together we get

$$\lim_{M \rightarrow 0} \frac{1}{M} \left\langle \sum_{j=1}^M \sum_{l=1}^M \sum_{i=1}^M [f_j(x) - g(x)]^2 [f_l(x) - g(x)]^2 f_i(x_*) \right\rangle_{f_1, \dots, f_M \sim \text{EK}} = 8[g_{\text{EK}, \eta}^*(x) - g(x)] \text{Var}_0[f(x)] \text{Cov}_0[f(x), f(x_*)], \quad (18)$$

so finally

$$\langle g^* \rangle_\eta = g_{\text{EK}, \eta}^*(x_*) + \frac{\eta}{\sigma^4} \int d\mu_x [g_{\text{EK}, \eta}^*(x) - g(x)] \text{Var}_0[f(x)] \text{Cov}_0[f(x), f(x_*)] + O(1/\eta^3). \quad (19)$$

Substituting the expressions for the free variance and the covariance [Eq. (H5)] we get

$$\langle g^* \rangle_\eta = g_{\text{EK}, \eta}^*(x_*) - \frac{\eta}{\sigma^4} \sum_{i,j,k} \frac{\sigma^2}{\lambda_i + \sigma^2/\eta} \left(\frac{1}{\lambda_j} + \frac{\eta}{\sigma^2} \right)^{-1} \left(\frac{1}{\lambda_k} + \frac{\eta}{\sigma^2} \right)^{-1} g_i \phi_j(x_*) \int d\mu_x \phi_i(x) \phi_j(x) \phi_k^2(x) + O(1/\eta^3). \quad (110)$$

For a rotationally invariant kernel, (110) becomes

$$\langle g^* \rangle_\eta = g_{\text{EK}, \eta}^*(x_*) - \sum_{l,m} \frac{\eta^{-1} \lambda_l C_{K, \sigma^2/\eta}}{(\lambda_l + \sigma^2/\eta)^2} g_{lm} Y_{lm}(x_*) + O(1/\eta^3). \quad (111)$$

2. Calculating $\langle g^{*2} \rangle_\eta$

Substituting (G10) in (G15) we get

$$\begin{aligned} \langle g^{*2} \rangle_\eta &= \lim_{M \rightarrow 0} \lim_{W \rightarrow 0} \frac{1}{MW (Z_{\text{EK}}[\alpha = 0])^{M+W}} \int \mathcal{D}f_1 \dots \int \mathcal{D}f_M \int \mathcal{D}\tilde{f}_1 \dots \int \mathcal{D}\tilde{f}_W \\ &\times \exp \left(-\eta - \frac{1}{2} \sum_{m=1}^M \|f_m\|_K^2 - \frac{1}{2} \sum_{w=1}^W \|\tilde{f}_w\|_K^2 \right) \\ &\times \exp \left[\eta \left(\exp \left(-\sum_{m=1}^M \frac{[f_m(x) - g(x)]^2}{2\sigma^2} - \sum_{w=1}^W \frac{[\tilde{f}_w(x) - g(x)]^2}{2\sigma^2} \right) \right) \right]_{x \sim \mu_x} \sum_{m=1}^M f_m(x_*) \sum_{w=1}^W \tilde{f}_w(x_*). \end{aligned} \quad (112)$$

By expanding to the same order we get [all equalities are up to $O(1/\eta^3)$]

$$\begin{aligned} \langle g^{*2}(x_*) \rangle_\eta &= \lim_{M \rightarrow 0} \lim_{W \rightarrow 0} \frac{1}{MW (Z_{\text{EK}}[\alpha = 0])^{M+W}} \int \mathcal{D}f_1 \dots \int \mathcal{D}f_M \int \mathcal{D}\tilde{f}_1 \dots \int \mathcal{D}\tilde{f}_W \\ &\times \exp \left[-\frac{1}{2} \sum_{m=1}^M \|f_m\|_K^2 - \frac{1}{2} \sum_{w=1}^W \|\tilde{f}_w\|_K^2 + \eta \left\langle \left(-\sum_{m=1}^M \frac{[f_m(x) - g(x)]^2}{2\sigma^2} - \sum_{w=1}^W \frac{[\tilde{f}_w(x) - g(x)]^2}{2\sigma^2} \right) \right\rangle_{x \sim \mu_x} \right] \\ &\times \left[1 + \frac{\eta}{2} \left\langle \left(\sum_{m=1}^M \frac{[f_m(x) - g(x)]^2}{2\sigma^2} + \sum_{w=1}^W \frac{[\tilde{f}_w(x) - g(x)]^2}{2\sigma^2} \right)^2 \right\rangle_{x \sim \mu_x} \right] \sum_{m=1}^M f_m(x_*) \sum_{w=1}^W \tilde{f}_w(x_*) \\ &= g_{\text{EK}, \eta}^{*2}(x_*) + \lim_{M \rightarrow 0} \lim_{W \rightarrow 0} \frac{1}{MW} \frac{\eta}{8\sigma^4} \int d\mu_x \left\langle \left(\sum_{a=1}^M [f_a(x) - g(x)]^2 + \sum_{b=1}^W [\tilde{f}_b(x) - g(x)]^2 \right)^2 \sum_{c=1}^M f_c(x_*) \sum_{d=1}^W \tilde{f}_d(x_*) \right\rangle_0 \\ &= g_{\text{EK}, \eta}^{*2}(x_*) + \lim_{M \rightarrow 0} \lim_{W \rightarrow 0} \frac{1}{MW} \frac{\eta}{4\sigma^4} \int d\mu_x \left[\left\langle \sum_{a=1}^M [f_a(x) - g(x)]^2 \sum_{b=1}^M [f_b(x) - g(x)]^2 \sum_{c=1}^M f_c(x_*) \sum_{d=1}^W \tilde{f}_d(x_*) \right\rangle_0 \right. \\ &\quad \left. + \left\langle \sum_{a=1}^M [f_a(x) - g(x)]^2 \sum_{b=1}^W [\tilde{f}_b(x) - g(x)]^2 \sum_{c=1}^M f_c(x_*) \sum_{d=1}^W \tilde{f}_d(x_*) \right\rangle_0 \right] \\ &= g_{\text{EK}, \eta}^{*2}(x_*) + \frac{\eta}{4\sigma^4} \int d\mu_x \lim_{M \rightarrow 0} \frac{1}{M} \left\langle \sum_{a=1}^M [f_a(x) - g(x)]^2 \sum_{b=1}^M [f_b(x) - g(x)]^2 \sum_{c=1}^M f_c(x_*) \right\rangle_0 g_{\text{EK}, \eta}^*(x_*) \\ &\quad + \frac{\eta}{4\sigma^4} \int d\mu_x \left(\lim_{M \rightarrow 0} \frac{1}{M} \left\langle \sum_{a=1}^M [f_a(x) - g(x)]^2 \sum_{b=1}^M f_b(x_*) \right\rangle_0 \right)^2. \end{aligned} \quad (113)$$

The first integrand was already calculated and is given in (18). For the second integrand we get

$$\begin{aligned}
 & \lim_{M \rightarrow 0} \frac{1}{M} \left\langle \sum_{a=1}^M (f_a(x) - g(x))^2 \sum_{b=1}^M f_b(x_*) \right\rangle \\
 &= \lim_{M \rightarrow 0} \frac{1}{M} [M \langle (f(x) - g(x))^2 f(x_*) \rangle + M(M-1) \langle f(x_*) \rangle \langle (f(x) - g(x))^2 \rangle] \\
 &= \langle (f(x) - g(x))^2 f(x_*) \rangle - \langle f(x_*) \rangle \langle (f(x) - g(x))^2 \rangle \\
 &= \text{---} \blacksquare \text{---} \bullet \text{---} + \text{disconnected diagrams} \\
 &= 2 (g^*_{EK,\eta}(x) - g(x)) \text{Cov}_0 [f(x), f(x_*)] = O(1/\eta^3)
 \end{aligned} \tag{I14}$$

so the correction for $\langle g^{*2}(x_*) \rangle_\eta$ is

$$\langle g^{*2}(x_*) \rangle_\eta = g^{*2}_{EK,\eta}(x_*) - 2 \underbrace{\frac{\eta}{\sigma^4} g^*_{EK,\eta}(x_*) \sum_{i,j,k} \frac{\sigma^2}{\lambda_i + \frac{\sigma^2}{\eta}} \left(\frac{1}{\lambda_j} + \frac{\eta}{\sigma^2} \right)^{-1} \left(\frac{1}{\lambda_k} + \frac{\eta}{\sigma^2} \right)^{-1} g_i \phi_j(x_*) \int d\mu_x \phi_i(x) \phi_j(x) \phi_k^2(x)}_{O(1/\eta^2)} + O(1/\eta^3) \tag{I15}$$

and for a rotationally invariant kernel we get

$$\langle g^{*2}(x_*) \rangle_\eta = g^{*2}_{EK,\eta}(x_*) - 2 g^*_{EK,\eta}(x_*) \sum_{l,m} \frac{\eta^{-1} \lambda_l C_{K,\sigma^2/\eta}}{(\lambda_l + \sigma^2/\eta)^2} g_{lm} Y_{lm}(x_*) + O(1/\eta^3). \tag{I16}$$

APPENDIX J: VARIOUS INSIGHTS

1. Correction means worse generalization

The correction always means worse generalization than what the EK suggests. Indeed, expanding Eq. (I11) we get

$$\begin{aligned}
 \langle g^* \rangle_\eta &= g^*_{EK,\eta}(x_*) - \sum_{l,m} \frac{\eta^{-1} \lambda_l C_{K,\sigma^2/\eta}}{(\lambda_l + \sigma^2/\eta)^2} g_{lm} Y_{lm}(x_*) + O(1/\eta^3) \\
 &= \sum_{l,m} \frac{\lambda_l}{\lambda_l + \frac{\sigma^2}{\eta}} g_{l,m} Y_{l,m}(x_*) \\
 &\quad - \sum_{l,m} \frac{\eta^{-1} \lambda_l C_{K,\sigma^2/\eta}}{(\lambda_l + \sigma^2/\eta)^2} g_{lm} Y_{lm}(x_*) + O(1/\eta^3) \\
 &= \sum_{l,m} \underbrace{\left(\frac{\lambda_l}{\lambda_l + \frac{\sigma^2}{\eta}} - \frac{\eta^{-1} \lambda_l C_{K,\sigma^2/\eta}}{(\lambda_l + \sigma^2/\eta)^2} \right)}_{\substack{\text{positive} \\ < \frac{\lambda_l}{\lambda_l + \frac{\sigma^2}{\eta}} < 1}} g_{l,m} Y_{l,m}(x_*) .
 \end{aligned}$$

2. Exact eigenvalues for two-layer ReLU NNGP and NTK with $\sigma_b^2 = 0$

For the NNGP associated with a two-layer ReLU NTK without bias we were able to find an exact expression for the eigenvalues for all l :

$$\begin{aligned}
 \lambda_{l=2k} &= \sigma_{w_0}^2 \sigma_{w_1}^2 \frac{d}{16\pi^2} \left(\frac{\Gamma(\frac{l-1}{2}) \Gamma(\frac{d}{2})}{\Gamma(\frac{l+d+1}{2})} \right)^2, \\
 \lambda_{l=2k+1} &= \sigma_{w_0}^2 \sigma_{w_1}^2 \frac{1}{4d} \delta_{k,0}
 \end{aligned}$$

and for NTK

$$\begin{aligned}
 \lambda_{2k} &= \frac{\sigma_{w_1}^2 \sigma_{w_2}^2}{2\pi} \frac{d(1+2k) + (1-2k)^2}{8\pi} \left(\frac{\Gamma(k - \frac{1}{2}) \Gamma(\frac{d}{2})}{\Gamma(k + \frac{d+1}{2})} \right)^2, \\
 \lambda_{2k+1} &= \frac{\sigma_{w_1}^2 \sigma_{w_2}^2}{2\pi} \frac{\pi}{d} \delta_{k,0}.
 \end{aligned}$$

It is interesting to note that for all odd $l > 1$ $\lambda_l = 0$ so the expressive power of the kernel (and hence the neural network) is greatly reduced.

APPENDIX K: ACCURACY OF THE RENORMALIZED NTK

Let us consider the random variable $t = x \cdot x'$, for two normalized data points x and x' drawn uniformly from the unit hypersphere S^{d-1} . Without loss of generality, x can be assumed to be a unit vector in the direction of the last axis, and therefore t is the last component of x' . The density at $t \in [-1, 1]$ is therefore proportional to the surface area lying at a height between t and $t + dt$ on the unit sphere. That proportion occurs within a belt of height dt and radius $\sqrt{1-t^2}$, which is a conical frustum constructed out of a $(d-2)$ -dimensional hypersphere of radius $\sqrt{1-t^2}$, of height dt , and slope $1/\sqrt{1-t^2}$. Hence, the probability is proportional to $p(t)dt \propto (1-t^2)^{(d-3)/2} dt$.

Defining $u = (t+1)/2$, it holds that $p(u)du \propto u^{(d-3)/2} (1-u)^{(d-3)/2}$, meaning that $u \sim \text{Beta}(\frac{d-1}{2}, \frac{d-1}{2})$, and for large d it holds that $\text{Var}[t] = O(d^{-1})$. Since t is bounded to $[-1, 1]$, the random variable t' must have a standard deviation which is a decaying function of r . Indeed, for $n \ll d$ and large d , approximating the integral $\int t^n (1-t^2)^{(d-3)/2} dt$ using saddle-point approximation we

get that $f(t) = n \ln(t) + \frac{d-3}{2} \ln(1-t^2)$ is maximal for $t_0 = [n/(n+d-3)]^{1/2}$, and $f''(t_0) = 2[n^2 - (d-3)^2]/(d-3)$, so overall

$$\langle t^{2r} \rangle = \frac{\int t^{2r}(1-t^2)^{(d-3)/2}}{\int t^0(1-t^2)^{(d-3)/2}} \approx \left[1 - \left(\frac{2r}{d-3} \right)^2 \right]^{-1/2} \times \left(\frac{2r}{2r+d-3} \right)^r \left(\frac{d-3}{2r+d-3} \right)^{(d-3)/2}. \quad (\text{K1})$$

This implies that for $r \ll d$, the standard deviation of t^r is $O(d^{-r/2})$. Considering next the tail of the Taylor expansion $\sum_{q>r} b_q(x \cdot x')^q$, projected on the data set $[\sum_{q>r} b_q(x_n \cdot x_m)^q]$. The resulting $N \times N$ matrix is $\sum_{q>r} b_q$ on the diagonal, but $O(d^{-(r+1)/2})$ in all other entries. As we justified in the main text, our renormalization transformation amounts to keeping only the diagonal piece of this matrix and interpreting it as noise.

Consider then (2) for g^* in two scenarios: (I) g_∞^* with the full NTK $[K(x, x')]$ and no noise and (II) g_r^* with the NTK trimmed after the r th power $[K_r(x, x')]$ but with $\sigma_r^2 = \sum_{q>r} b_q$. The first $K(x_*, x_n)$ piece, for x_* drawn from the dataset distribution, obeys $K(x_*, x_n) - K_r(x_*, x_n) = O(d^{-(r+1)/2})$. Next, we compare $K_r(x_n, x_m) + I_{nm}\sigma_r^2$ and $K(x_n, x_m)$. On their diagonal they agree exactly but their off-diagonal terms agree only up to a $O(d^{-(r+1)/2})$ discrepancy. Denoting by δK the difference between these two matrices, we

may expand $K^{-1} = [K_r + \sigma_m^2 \mathbf{I} + \delta K]^{-1} = [K_r + \sigma_r^2 \mathbf{I}]^{-1} [1 - \delta K [K_r + \sigma_r^2 \mathbf{I}]^{-1} + \delta K [K_r + \sigma_r^2 \mathbf{I}]^{-1} \delta K [K_r + \sigma_r^2 \mathbf{I}]^{-1} + \dots]$.

We next argue that $\delta K [K_r + \sigma_r^2 \mathbf{I}]^{-1}$ multiplied by target vector $[g(x_n)]$ is negligible compared to the identity for large enough r , thereby establishing the equivalence of the two scenarios. Indeed, consider the eigenvalues of $\delta K [K_r + \sigma_r^2 \mathbf{I}]^{-1}$. As δK_{nm} is $O(d^{-(r+1)/2})$ its typical eigenvalues are $O(\sqrt{Nd}^{-(r+1)/2})$ and bounded by $O(Nd^{-(r+1)/2})$. The typical eigenvalues of $[K_r + \sigma_r^2 \mathbf{I}]^{-1}$ are of the same order as $K(x_n, x_n) = K$ and bounded from below by σ_r^2 . Thus, typical eigenvalues of $\delta K [K_r + \sigma_r^2 \mathbf{I}]^{-1}$ are $O(\sqrt{Nd}^{-(r+1)/2}/K)$ and bounded from above by $O(Nd^{-(r+1)/2}/\sigma_r^2)$. The NTK has the desirable property that σ_r^2 decays very slowly. Thus, certainly in the typical case but even in the worst-case scenario we expect good agreement at large r . In Fig. 1, right panel, we provide supporting numerical evidence.

We refer to $K_r(x, x')$ as the renormalized NTKs at the scale r . As follows from (23), λ_l 's with $l \geq r$ are zero. Therefore, as advertised, the high-energy sector has been removed and compensated by noise on the target and a change of the remaining $l < r$ (low-energy) eigenvalues. A proper choice of r involves two considerations. Requiring perturbation theory to hold well ($C_{K_r, \sigma_r^2/\eta} < \sigma_r^2$) which puts an η -dependent upper bound on r and requiring small discrepancy in predictions puts another η -dependent lower bound on r (typically, $\sqrt{Nd}^{-(r+1)/2} \ll 1$).

Lastly, we comment that our renormalization NTK approach is not limited to uniform data sets. The entire logic relies on having a rapidly decaying ratio of off-diagonal

TABLE II. Hyperparameter performance comparison.

Test	Prediction	GPR	σ_r^2	Train	σ_{w_1}	σ_{b_1}	σ_{w_2}	σ_{b_2}	
Random 1	0.389	0.400	0.364	0.068	1.29e-04	1.555	0.032	2.028	0.026
Random 2	0.316	0.287	0.331	0.004	4.17e-04	0.914	0.029	0.880	0.062
Random 3	0.191	0.219	0.250	0.003	3.36e-04	0.922	0.049	0.760	0.046
Random 4	0.300	0.324	0.306	0.045	1.39e-04	1.552	0.058	1.644	0.054
Random 5	0.268	0.338	0.273	0.070	1.24e-04	1.585	0.072	2.020	0.029
Random 6	0.413	0.406	0.382	0.065	1.28e-04	2.037	0.032	1.512	0.053
Random 7	0.332	0.297	0.335	0.010	8.76e-04	0.994	0.030	1.190	0.071
Random 8	0.228	0.245	0.262	0.015	6.41e-04	1.165	0.059	1.271	0.068
Random 9	0.337	0.308	0.332	0.018	1.37e-03	0.909	0.027	1.758	0.030
Random 10	0.371	0.392	0.355	0.069	1.19e-04	1.658	0.041	1.908	0.057
Random 11	0.313	0.316	0.319	0.032	1.31e-04	1.440	0.050	1.502	0.045
Random 12	0.335	0.340	0.329	0.042	1.33e-04	2.106	0.065	1.175	0.040
Random 13	0.397	0.336	0.373	0.018	1.71e-03	1.546	0.029	1.037	0.044
Random 14	0.236	0.253	0.271	0.017	4.18e-04	1.388	0.067	1.133	0.050
Random 15	0.293	0.288	0.306	0.018	6.37e-04	1.534	0.058	1.068	0.065
Random 16	0.175	0.198	0.214	0.011	1.04e-03	1.132	0.074	1.117	0.053
Random 17	0.206	0.226	0.246	0.013	7.66e-04	1.095	0.061	1.275	0.072
Random 18	0.264	0.264	0.292	0.011	1.00e-03	1.035	0.043	1.230	0.066
Random 19	0.239	0.248	0.282	0.006	5.34e-04	0.835	0.037	1.107	0.050
Random 20	0.385	0.378	0.364	0.054	1.43e-04	1.810	0.039	1.540	0.046
Random 21	0.318	0.300	0.323	0.018	9.58e-04	1.282	0.042	1.266	0.055
Worst	0.413	0.406	0.382	0.065	1.28e-04	2.037	0.032	1.512	0.053
Median	0.313	0.316	0.319	0.032	1.31e-04	1.440	0.050	1.502	0.045
Best	0.175	0.198	0.214	0.011	1.04e-03	1.132	0.074	1.117	0.053
Typical	0.307	0.307	0.317	0.028	1.41e-04	1.414	0.050	1.414	0.050
Optimized	0.078	0.110	0.141	0.002	1.66e-04	0.707	0.075	0.707	0.027

moments $[(x_n \cdot x_m)^{2r}]$ and diagonal moments $(x_n \cdot x_n)^{2r}$ as one increases r . We expect this to hold in real-world distributions. For instance, for a multidimensional Gaussian data distribution the input dimension (d) traded by an effective dimension (d_{eff}) defined by the variance of $(x_m \cdot x_n)$.

APPENDIX L: HYPERPARAMETER OPTIMIZATION EXPERIMENT RESULTS

Table II shows a comparison of hyperparameter optimization performance. Fig. 7 shows correlations between GPR loss, DNN test loss, and our prediction.

- [1] Q. Xie, M. T. Luong, E. Hovy, and Q. V. Le, Self-training with noisy student improves imagenet classification, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020).
- [2] G. Saon, G. Kurata, T. Sercu, K. Audhkhasi, S. Thomas, D. Dimitriadis, X. Cui, B. Ramabhadran, M. Picheny, L. L. Lim, B. Roomi, and P. Hall, English conversational telephone speech recognition by humans and machines, in *INTERSPEECH* (2017).
- [3] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, T. P. Lillicrap, K. Simonyan, and D. Hassabis, Mastering chess and shogi by self-play with a general reinforcement learning algorithm, [arXiv:1712.01815](https://arxiv.org/abs/1712.01815).
- [4] N. Kalchbrenner, L. Espeholt, K. Simonyan, A. van den Oord, A. Graves, and K. Kavukcuoglu, Neural machine translation in linear time, [arXiv:1610.10099](https://arxiv.org/abs/1610.10099).
- [5] R. Shwartz-Ziv and N. Tishby, Opening the black box of deep neural networks via information, [arXiv:1703.00810](https://arxiv.org/abs/1703.00810).
- [6] D. Hexner, A. J. Liu, and S. R. Nagel, Periodic training of creeping solids, *Proc. Nat. Acad. Sci.* **117**, 31690 (2020).
- [7] A. Daniely, R. Frostig, and Y. Singer, Toward deeper understanding of neural networks: The power of initialization and a dual view on expressivity, *Advances in Neural Information Processing Systems*, edited by D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (Curran Associates, Inc., 2016).
- [8] A. Jacot, F. Gabriel, and C. Hongler, Neural tangent kernel: Convergence and generalization in neural networks, *Advances in Neural Information Processing Systems*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran Associates, Inc., 2018), Vol. 31.
- [9] Z. Allen-Zhu, Y. Li, and Y. Liang, Learning and generalization in overparameterized neural networks, going beyond two layers, *Advances in Neural Information Processing Systems*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019), Vol. 32.
- [10] Y. Cao and Q. Gu, Generalization error bounds of gradient descent for learning over-parameterized deep ReLU networks, *Proc. AAAI Conf. Artificial Intelligence* **34**, 3349 (2020).
- [11] Y. Cao and Q. Gu, Generalization bounds of stochastic gradient descent for wide and deep neural networks, *Advances in Neural Information Processing Systems*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019), Vol. 32.
- [12] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)* (MIT Press, Cambridge, MA, 2005).
- [13] G. Naveh, O. Ben-David, H. Sompolinsky, and Z. Ringel, Predicting the outputs of finite networks trained with noisy gradients, [arXiv:2004.01190](https://arxiv.org/abs/2004.01190).
- [14] B. Ghorbani, S. Mei, T. Misiakiewicz, and A. Montanari, When do neural networks outperform kernel methods? *Advances in Neural Information Processing Systems*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Curran Associates, Inc., 2020), Vol. 33, pp. 14820–14830.
- [15] S. Arora, S. S. Du, Z. Li, R. Salakhutdinov, R. Wang, and D. Yu, Harnessing the power of infinitely wide deep nets on small-data tasks, [arXiv:1910.01663](https://arxiv.org/abs/1910.01663).
- [16] Z. Li, W. Ruosong, D. Yu, S. Du, W. Hu, R. Salakhutdinov, and S. Arora, Enhanced convolutional neural tangent kernels, [arXiv:1911.00809](https://arxiv.org/abs/1911.00809).
- [17] J. Lee, J. Sohl-Dickstein, J. Pennington, R. Novak, S. Schoenholz, and Y. Bahri, Deep neural networks as Gaussian processes, in *Proceedings of the International Conference on Learning Representations*, 2018 (unpublished).
- [18] J. Lee, S. S. Schoenholz, J. Pennington, B. Adlam, L. Xiao, R. Novak, and J. Sohl-Dickstein, Finite versus infinite neural networks: An empirical study, *Advances in Neural Information Processing Systems*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Curran Associates, Inc., 2020), Vol. 33, pp. 15156–15172.
- [19] L. Bo, X. Ren, and D. Fox, in *Advances in Neural Information Processing Systems 23*, edited by J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta (Curran Associates, Red Hook, NY, 2010), pp. 244–252.
- [20] A. Lewkowycz, Y. Bahri, E. Dyer, J. Sohl-Dickstein, and G. Gur-Ari, The large learning rate phase of deep learning: The catapult mechanism, [arXiv:2003.02218](https://arxiv.org/abs/2003.02218).
- [21] R. F. Warming and B. J. Hyett, The modified equation approach to the stability and accuracy analysis of finite-difference methods, *J. Comput. Phys.* **14**, 159 (1974).
- [22] S. Arora, S. S. Du, W. Hu, Z. Li, R. Salakhutdinov, and R. Wang, On exact computation with an infinitely wide neural net, [arXiv:1904.11955](https://arxiv.org/abs/1904.11955).
- [23] L. Zdeborová, Understanding deep learning is also a job for physicists, *Nat. Phys.* **16**, 602 (2020).
- [24] Y. Levine, O. Sharir, N. Cohen, and A. Shashua, Quantum Entanglement in Deep Learning Architectures, *Phys. Rev. Lett.* **122**, 065301 (2019).
- [25] B. Li and D. Saad, Exploring the Function Space of Deep-Learning Machines, *Phys. Rev. Lett.* **120**, 248301 (2018).
- [26] S. Becker, Y. Zhang, and A. A. Lee, Geometry of Energy Landscapes and the Optimizability of Deep Neural Networks, *Phys. Rev. Lett.* **124**, 108301 (2020).
- [27] E. W. Tramel, M. Gabrié, A. Manoel, F. Caltagirone, and F. Krzakala, Deterministic and Generalized Framework for Unsupervised Learning with Restricted Boltzmann Machines, *Phys. Rev. X* **8**, 041006 (2018).
- [28] Dörthe Malzahn and M. Opper, A variational approach to learning curves, in *Proceedings of the 14th International Con-*

- ference on Neural Information Processing Systems: Natural and Synthetic*, NIPS'01 (MIT Press, Cambridge, MA, 2001), pp. 463–469.
- [29] Christopher K. I. Williams and F. Vivarelli, Upper and lower bounds on the learning curve for Gaussian processes, *Mach. Learn.* **40**, 77 (2000).
- [30] K. Ritter, *Average-Case Analysis of Numerical Problems*, Lecture Notes in Mathematics (Springer, Berlin, 2007).
- [31] K. Ritter, Asymptotic optimality of regular sequence designs, *Ann. Stat.* **24**, 2081 (1996).
- [32] C. A. Micchelli and G. Wahba, Design problems for optimal surface interpolation (1979).
- [33] P. Sollich, Gaussian process regression with mismatched models, *Advances in Neural Information Processing Systems*, edited by T. Dietterich, S. Becker, and Z. Ghahramani (MIT Press, 2002).
- [34] D. Azevedo and V. A. Menegatto, Eigenvalues of dot-product kernels on the sphere, *Proceeding Series of the Brazilian Society of Computational and Applied Mathematics* (2015).
- [35] G. Ferrari-Trecate, Christopher K. I. Williams, and M. Opper, Finite-dimensional approximation of Gaussian processes, in NIPS (1998).
- [36] M. A. Nielsen, *Neural Networks and Deep Learning* (Determination Press, 2015).
- [37] R. M. Neal, *Bayesian Learning for Neural Networks* (Springer, New York, 2012), Vol. 118.
- [38] J. Lee, L. Xiao, S. S. Schoenholz, Y. Bahri, J. Sohl-Dickstein, and J. Pennington, Wide neural networks of any depth evolve as linear models under gradient descent, *J. Stat. Mech.* (2020) 124002.
- [39] J. C. Mattingly, A. M. Stuart, and D. J. Higham, Ergodicity for sdes and approximations: Locally lipschitz vector fields and degenerate noise, *Stochastic Process. Appl.* **101**, 185 (2002).
- [40] H. Risken and T. Frank, *The Fokker-Planck Equation: Methods of Solution and Applications*, Springer Series in Synergetics (Springer, Berlin, 1996).
- [41] S. Mandt, M. D. Hoffman, and D. M. Blei, Stochastic gradient descent as approximate bayesian inference, *J. Mach. Learn. Res.* **18**, 4873 (2017).
- [42] M. Welling and Y. W. Teh, Bayesian learning via stochastic gradient langevin dynamics, in *Proceedings of the 28th International Conference on International Conference on Machine Learning*, ICML'11 (Omnipress, Norristown, PA, 2011), pp. 681–688.
- [43] T. A. Ensslin, M. Frommert, and F. S. Kitaura, Information field theory for cosmological perturbation reconstruction and non-linear signal analysis, *Phys. Rev. D* **80**, 105005 (2009).
- [44] Y. Cho and L. K. Saul, Kernel methods for deep learning, in *Proceedings of the 22Nd International Conference on Neural Information Processing Systems*, NIPS'09 (Curran Associates, Red Hook, NY, 2009), pp. 342–350.
- [45] C. Efthimiou and C. Frye, Spherical harmonics in p dimensions, (World Scientific, Singapore, 2014).
- [46] S. S. Schoenholz, J. Gilmer, S. Ganguli, and J. Sohl-Dickstein, Deep information propagation, [arXiv:1611.01232](https://arxiv.org/abs/1611.01232).
- [47] M. Belkin, D. Hsu, S. Ma, and S. Mandal, Reconciling modern machine learning and the bias-variance trade-off, [arXiv:1812.11118](https://arxiv.org/abs/1812.11118).
- [48] N. Rahaman, D. Arpit, A. Baratin, F. Dräxler, M. Lin, F. A. Hamprecht, Y. Bengio, and A. C. Courville, On the spectral bias of deep neural networks, in *Proceedings of the 36th International Conference on Machine Learning*, (PMLR, 2019), Vol. 97, pp. 5301–5310.
- [49] K. He, X. Zhang, S. Ren, and J. Sun, Delving deep into rectifiers: Surpassing human-level performance on image net classification, in *2015 IEEE International Conference on Computer Vision (ICCV)* (IEEE, Piscataway, NJ, 2015), pp. 1026–1034.
- [50] R. Novak, L. Xiao, J. Hron, J. Lee, A. A. Alemi, J. Sohl-Dickstein, and S. S. Schoenholz, Neural tangents: Fast and easy infinite neural networks in python, in *Proceedings of the International Conference on Learning Representations*, 2020 (unpublished).
- [51] Y. N. Dauphin, R. Pascanu, C. Gulcehre, K. Cho, S. Ganguli, and Y. Bengio, Identifying and attacking the saddle point problem in high-dimensional non-convex optimization, *Advances in Neural Information Processing Systems*, edited by Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger (Curran Associates, Inc., 2014), Vol. 27.
- [52] F. Draxler, K. Veschgini, M. Salmhofer, and F. A. Hamprecht, Essentially no barriers in neural network energy landscape, in *Proceedings of the 35th International Conference on Machine Learning*, edited by J. Dy and A. Krause (PMLR, 2018), Vol. 80, pp. 1309–1318.
- [53] R. Novak, L. Xiao, J. Lee, Y. Bahri, G. Yang, D. A. Abolafia, J. Pennington, and J. Sohl-Dickstein, Bayesian deep convolutional networks with many channels are gaussian processes, [arXiv:1810.05148](https://arxiv.org/abs/1810.05148).