

LEARNING CURVES FOR DEEP NEURAL NETWORKS: A FIELD THEORY PERSPECTIVE

Anonymous authors

Paper under double-blind review

ABSTRACT

A series of recent works established a rigorous correspondence between very wide deep neural networks (DNNs), trained in a particular manner, and noiseless Bayesian Inference with a certain Gaussian Process (GP) known as the Neural Tangent Kernel (NTK). Here we extend a known field-theory formalism for GP inference to get a detailed understanding of learning-curves in DNNs trained in the regime of this correspondence (NTK regime). In particular, a renormalization-group approach is used to show that noiseless GP inference using NTK, which lacks a good analytical handle, can be well approximated by noisy GP inference on a related kernel we call the renormalized NTK. Following this, a perturbation-theory analysis is carried in one over the dataset-size yielding analytical expressions for the (fixed-teacher/fixed-target) leading and sub-leading asymptotics of the learning curves. At least for uniform datasets, a coherent picture emerges wherein fully-connected DNNs have a strong implicit bias towards functions which are low order polynomials of the input.

1 INTRODUCTION

Several pleasant features underlay the success of deep learning: The scarcity of bad minima encountered in their optimization [Draxler et al. (2018); Choromanska et al. (2014)], their ability to generalize well despite being heavily over-parameterized [Neysshabur et al. (2018; 2014)] and expressive [Zhang et al. (2016)], and their ability to generate internal representations which generalize across different domains and tasks [Yosinski et al. (2014); Sermanet et al. (2013)].

Due to the complexity of DNNs our current understanding of these features is still largely empirical. Notwithstanding, progress has been made recently in the highly over-parametrized regime [Daniely et al. (2016); Jacot et al. (2018)] due to the fact that the networks' parameters, in all non-linear layers, change in a minor yet important manner during training. This facilitated the derivation of various bounds [Allen-Zhu et al. (2018); Cao & Gu (2019b;a)] on generalization and, more relevant for this work, the following correspondence with GPs: Considering finite-depth DNNs which are much wider than the dataset-size, trained with MSE loss, no weight decay, and at vanishing learning rate (the NTK-regime) one finds that the initialization-averaged predictions are the same as those of Gaussian Processes Regression (GPR) with a kernel known as the NTK. Several subsequent works corroborated these results empirically [Lee et al. (2018); Lee et al. (2019); Arora et al. (2019)] and extended them [Arora et al. (2019)]. For fully-connected DNNs, the NTK-regime (and GPs associated with DNNs in general [Lee et al. (2018); Novak et al. (2018)]) seems to faithfully capture the generalization power of DNNs trained with MSE loss [Lee et al. (2019)].

One of the most detailed objects quantifying generalization are learning-curves: graphs of how the test error diminishes with the number of datapoints (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. In contrast, for GPR many available analytical tools have yielded, in the past, high accuracy predictions for learning curves. One of the most transparent ones is the equivalence kernel (EK) Rasmussen & Williams (2005): Given a GP kernel ($K(x, x')$) along with its expansion in terms of features ($K(x, x') = \sum_n \lambda_n \phi_n(x) \phi_n(x')$), GPR on a target function ($g(x) = \sum_n g_n \phi_n(x)$) using N datapoints will yield (approximately, at large N) the function $g_{EK}^*(x) = \sum_n \frac{\lambda_n}{\lambda_n + \sigma^2/N} g_n \phi_n(x)$, where σ^2 is the variance of an observation noise on the target function.

Clearly such a detailed understanding of generalization in DNNs is desirable. However, several technical issues prohibit the application of the EK and related results [Rasmussen & Williams (2005); Malzahn & Opper (2001)] to DNNs trained in the NTK-regime. First, the NTK-regime corresponds to noiseless GPR ($\sigma^2 = 0$) where the DNN and corresponding GP both fit the training dataset exactly. In this case various approximations for generalization break down. For instance in the above EK result and other fixed-target results [Malzahn & Opper (2001)], it would appear that we can learn the function perfectly using only a single data-point at $\sigma^2 = 0$. Second, the features ($\phi_n(x)$) and eigenvalues λ_n of the NTK are needed so that the EK can be interpreted. Third, as we see EK results can be misleading, it is important to estimate the validity range of this approaches and, in a related manner, derive sub-leading corrections.

In this work we make the following contributions:

- I we extend the field-theory formalism of Malzahn & Opper (2001) for GPR and obtain closed expression for the leading and sub-leading asymptotics of learning curves for any fixed target function (fixed-teacher learning curves).
- II For uniform datasets these expression simplify considerably and, together with our results on the eigenvalues, lead to clear relations between deep fully-connected networks and polynomial regression.
- III We establish that noiseless GP inference using NTK can be well approximated by noisy GP inference on a certain renormalized NTK. In addition explicit expressions are given for the eigenvalues and features of renormalized NTKs of any depth. Also we point to a simple universal bound on the eigenvalues of all these NTKs.

Apart from facilitating further transfer of knowledge between the physics, deep learning, and GP communities, our predictions for learning curves have several merits which distinguish them from other recent works on generalization: **1.** We provide leading and sub-leading asymptotic behaviors and allow computing further sub-sub-leading corrections. **2.** Considering uniform dataset input distributions, we believe our learning curves estimates stand-out in terms of accuracy and get to within 3% accuracy in value. **3.** Our results are predictions for the curves rather than bounds and **4.** They apply for fully-connected DNNs of any depth trained in the NTK-regime.

2 PRIOR WORKS

Learning curves for GPs have been analyzed using a variety of techniques [see Rasmussen & Williams (2005) for a review] most of which focus on a GP-teacher averaged case where the target/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 similar grand-canonical/Poisson-averaged approach Malzahn & Opper (2001) as our, however, the treatment of the resulting partition function was variational whereas we take a perturbation-theory approach. In addition previous cited results for MSE-loss breakdown in the noiseless limit [Malzahn & Opper (2001)]. To the best of our knowledge, noiseless GPs learning-curves have been analyzed analytically only in the teacher-averaged case and in the following settings: For matched priors, exact results are known for one dimensional data Williams & Vivarelli (2000); Rasmussen & Williams (2005) and two dimensional data with some limitations of how one samples the inputs (in the context of optimal design) Ritter (2007; 1996). In addition Micchelli & Wahba (1979) 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 dataset averaging involved in obtained learning curves has been performed analytically and the resulting matrix traces analyzed numerically Sollich (2001). Notably none of these cited results apply in any straightforward manner in the NTK-regime.

Considering kernel eigenvalues, explicit expression for the features and eigenvalue for dot-product kernels $(x \cdot x')^n$ where given in [Azevedo & Menegatto (2015)]. The d^{-l} scaling of eigenvalues of the kernels of the type $f(x \cdot x')$ which we used in our derivation of the bound has been noticed in [Sollich (2001)]. Kernels with a trimmed spectrum where the spectrum is trimmed after the first N 's leading eigenvalues, has previously been suggested as a way of reducing the computational cost of GP inference Ferrari-Trecate et al. (1998). 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/renormalization procedure.

Several interesting recent works give bounds on generalization [Allen-Zhu et al. (2018); Cao & Gu (2019b;a)] which show $O(1/\sqrt{N})$ asymptotic decay of the learning-curve (at best). In contrast our predictions are typically well below this bound.

3 FIELD THEORY FORMULATION OF GP LEARNING-CURVES

Here we describe a field theory formalism for exploring learning curves. We begin with standard definitions of GPs and Bayesian Inference on GPs. Being Gaussian, the probability distribution on a function $f(x)$ drawn from GPs is determined by its first and second moments. The first is typically taken to be zero and second is known as the covariance function or the kernel ($K_{xx'} = E[f(x)f(x')]$, where E denotes expectation under the GP distribution). Notably, $K_{xx'}$ of both the NNGP and NTK type can be calculated analytically for many activation functions [Cho & Saul (2009); Jacot et al. (2018)]. Furthermore, Bayesian Inference on GPs drawn from DNNs is tractable [Lee et al. (2018); Cho & Saul (2009)] and explicitly given by

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

where x_* is a new datapoint, g^* is the prediction, g_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 dataset (D)), σ^2 is a regulator corresponding to a noisy measurement of the GP and I is the identity matrix. Some intuition for this formula can be gained by verifying that $x_* = x_q$ yields $g^* = g_q$ when $\sigma^2 = 0$.

While the above equation determines the predictions and therefore the learning-curves, it does not do so in any clear or computationally accessible manner. This fact is due to the (potentially very) large matrix inversion involved, and the additional averaging over D required.

To facilitate the analysis of Eq. 1 we turn to a statistical-field-theory/path-integral viewpoint [Schulman (1996)]. These are well-studied, powerful approaches for performing integrations over a space of functions (the jargon is "paths" when x is one dimensional and "fields" when x is higher dimensional). To get some familiarity with this formalism, consider first averages over the (centered) GP itself with no dataset. Using the path-integral formalism we write it as

$$P_0[f] = \frac{\exp\left(-\frac{1}{2} \int dx dx' f(x) K^{-1}(x, x') f(x')\right)}{\int D\tilde{f} \exp\left(-\frac{1}{2} \int dx dx' \tilde{f}(x) K^{-1}(x, x') \tilde{f}(x')\right)} \quad (2)$$

where $\int Df$ denotes integration over the space of functions, for concreteness we limit $\int dx'$ to some compact domain such as the hyper-sphere, $K^{-1}(x, x')$ is the inverse covariance function ($\int dx' K(x, x') K^{-1}(x', x'') = \delta(x - x'')$). To define the path-integrals one first chooses an orthonormal basis of functions $\phi_i(x)$ (with respect to $\int dx$) arranged in order of likeliness $P_0[\phi_i] \geq P_0[\phi_j]$ for $i > j$ (note that this comparison doesn't require calculating the path integral in (2)). Second, one expands $f = \sum_i f_i \phi_i(x)$, and defines the path-integral as a series of simple integrals

$$\int Df \mathcal{F}[f] = \int df_1 \int df_2 \dots \mathcal{F} \left[\sum_i f_i \phi_i \right] \quad (3)$$

where \mathcal{F} is some functional of f . Finally, one makes this last expression well-defined by taking a limit procedure where the number of integrals is gradually taken to infinity [Schulman (1996)].

Performing the above procedure we show in App. F, $K_{x_1 x_2} = \int Df P_0[f] f(x_1) f(x_2)$. Notably, all other higher correlation functions split into products of the above correlation function due to standard properties of Gaussian integrals (Wick's/Isserlis' theorem). Following a similar procedure, and denoting $\|f\|_K^2 = \int dx dx' f(x) K^{-1}(x, x') f(x')$ one can show [Rasmussen & Williams (2005)]

$$g^*(x_*) = Z^{-1} \int Df \cdot f(x_*) \cdot \exp\left(-\frac{1}{2} \|f\|_K^2 - \frac{1}{2\sigma^2} \sum_{n=1}^N (f(x_n) - g_n)^2\right) \quad (4)$$

$$Z = \int Df \exp\left(-\frac{1}{2} \|f\|_K^2 - \frac{1}{2\sigma^2} \sum_{n=1}^N (f(x_n) - g_n)^2\right)$$

where Z is known as the partition function.

The averaged generalization error is defined as $\int d\mu_{x_*} \langle (g(x_*) - g^*(x_*))^2 \rangle_{x_1, \dots, x_n \sim \mu}$ where μ is the measure from which data points are drawn. Therefore, in order to calculate learning-curves, one needs to average quantities (like g^* and g^{*2}) over all datasets of size N drawn from a probability distribution $d\mu_x = P(x)dx$. We denote this averaging by $\langle \dots \rangle_{\mu, N}$. To facilitate this we next adopt the approach of [Malzahn & Oppen (2001)] and instead consider a related quantity given by the Poisson averaging of the former one

$$\langle \dots \rangle_{\mu, \eta} = e^{-\eta} \sum_{n=0}^{\infty} \frac{\eta^n}{n!} \langle \dots \rangle_{\mu, n} \quad (5)$$

where ... can be any quantity, in particular g^* and g^{*2} . Borrowing jargon from physics we refer to the original data ensemble as the canonical ensemble and to the above as the grand-canonical. Taking $\eta = N$, means we are essentially averaging over values of N in an \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 sub-leading. In App. A we compare learning curves as a function of N and η and show that they match very well. We also believe that such learning curves based on a grand-canonical/Poisson-averaged data ensembles are as interesting as the standard ones for quantifying generalization.

Using this modification, averaging over draws from the dataset can be carried using the "replica trick" (see for instance [Gardner & Derrida (1988)]), which aids in averaging over expressions like $\log(Z)$ and their derivatives via the equality $\log(Z) =_{M \rightarrow 0} \frac{Z^M - 1}{M}$. Employing this we find that for a non-negative integer M , $\langle g^* \rangle_{\mu, \eta}$ can be written as

$$\lim_{M \rightarrow 0} M^{-1} \int Df_1 \dots Df_M \exp \left(-\frac{1}{2} \sum_{m=1}^M \|f_m\|_K^2 + \eta \int d\mu_x e^{-\frac{1}{2\sigma^2} \sum_{m=1}^M (f_m(x) - g(x))^2} \right) \sum_{m=1}^M f_m(x_*) \quad (6)$$

where, as standard in the replica formalism, the computation should be carried at positive integer M and the analytical result extrapolated to zero at the end.

The main benefit of Eq. (6) over Eq. (1) is that it allows for a controlled expansion in $1/\eta$. At large η (or similarly large N) we expect that 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

$$\int d\mu e^{-\frac{\sum_{m=1}^M (f_m(x) - g(x))^2}{2\sigma^2}} = 1 - \int d\mu \frac{\sum_{m=1}^M (f_m(x) - g(x))^2}{2\sigma^2} + \frac{1}{2} \int d\mu \left[\frac{\sum_{m=1}^M (f_m(x) - g(x))^2}{2\sigma^2} \right]^2 + \dots \quad (7)$$

as shown in App. G, dealing with the first order term in this expansion in an exact (Gaussian) manner yields the aforementioned EK results ($f_{N, \sigma^2}^*(x)$) however with the difference that N is replaced by η . The second order term and further terms render the theory non-Gaussian and cannot be dealt with exactly but rather through standard perturbation-theory/Feynman-diagrams. In App. G we perform this calculation and obtain that $\langle g^*(x_*) - g(x_*) \rangle_{\mu, \eta}$ is given up to $O(1/\eta^3)$ by

$$\sum_i \frac{\frac{\sigma^2}{\eta} g_i \phi_i(x_*)}{\lambda_i + \frac{\sigma^2}{\eta}} - \frac{\eta}{\sigma^4} \sum_{i,j,k} \frac{\frac{\sigma^2}{\eta}}{\lambda_i + \frac{\sigma^2}{\eta}} \left(\frac{1}{\lambda_j} + \frac{\eta}{\sigma^2} \right)^{-1} \left(\frac{1}{\lambda_k} + \frac{\eta}{\sigma^2} \right)^{-1} g_i \phi_j(x_*) \int d\mu_x \phi_i(x) \phi_j(x) \phi_k^2(x) \quad (8)$$

As shown App. G similar expressions for $\langle g^{*2} \rangle_{\mu, \eta}$ are obtained using two replica indices. Interestingly we find that $\langle g^{*2} \rangle_{\mu, \eta} = \langle g^* \rangle_{\mu, \eta}^2 + O(1/\eta^3)$. Hence the averaged MSE error is simply Eq. 8 squared and 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 derivation and thus well-define even without averaging of datasets.

Equation 8 and its square which is the average MSE error are our first main result. They provides us with closed expressions for the dataset-averaged MSE loss as a function of η namely, the fixed-teacher

learning curve. They hold without any limitations on the dataset or the kernel and yield a variant of the EK result along with its sub-leading 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 unwanted cross-talk is generated between features due to the second sub-leading term. Notably for the renormalized NTK introduced below, the number of non-zero 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. 1 along with an N -dimensional integral involved in dataset-averaging. In addition having the sub-leading correction allows us to estimate the range of validity of our approximation by comparing the sub-leading and leading contributions, as do for the uniform case below.

4 UNIFORM DATASETS

To make the result in Eq. (8) interpretable, $\phi_i(x)$ and λ_i are required. This can be done most readily for the case of datasets normalized to the hypersphere ($\|x_n\| = 1$) with a uniform probability measure and rotation-symmetric kernel functions. By the latter we mean $K(x, x') = K(Ox, Ox')$ for any O , where O is an orthogonal matrix over the space of inputs. Although beyond the scope of the current work obvious extensions to consider are datasets which are uniform only in a sub-space of x and/or small perturbations to uniformity.

Importantly, the NTK associated with any DNN with a fully connected first layer and weights initialized from a normal distribution, has the above symmetry under rotations. This follows from the recursion relations defining the NTK [Jacot et al. (2018)] along with fact that the kernel of the first fully-connected layer is only a function of $x \cdot x'$. It follows that the NTK can be expanded as $K(x, x') = \sum_n b_n (x \cdot x')^n$. An additional corollary [Azevedo & Menegatto (2015)] 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 w.r.t a uniform measure) basis for functions on the hypersphere. 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 $\text{deg}(l)$).¹ For a kernel of the above form the eigenvalues are independent of m and given by [Azevedo & Menegatto (2015)]

$$\lambda_l = \frac{\Gamma(\frac{d}{2})}{\sqrt{\pi} \cdot 2^l} \sum_{s=0}^{\infty} b_{2s+l} \frac{(2s+l)!}{(2s)!} \frac{\Gamma(s + \frac{1}{2})}{\Gamma(s + l + \frac{d}{2})} \quad (9)$$

For ReLU and erf activations, the b_n 's, can be obtained analytically up to any desirable order. Thus one can semi-analytically obtain the NTK eigenvalues up to any desired accuracy. For the particular case of depth 2 ReLU networks, we report in the App. H closed expression where the above summation can be carried out analytically. 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, the above result can be seen as a closed analytical expression for the eigenvalues.

Interestingly, for any fully-connected network and uniform datasets of dimension d on the hypersphere, there is a universal bound given by $\lambda_l \leq K/\text{deg}(l) \approx O(d^{-l})$, where K is $K(x, x)$ which is a constant in x . Indeed note that $K(x, x)$ is finite and therefore its integral over the hypersphere is also finite and given by $\int d\mu K(x, x) = K(x, x) = \sum_{lm} \lambda_l = \sum_l \text{deg}(l) \lambda_l$. The degeneracy ($\text{deg}(l)$) is fixed from properties of hyper spherical harmonics, and equals $\text{deg}(l) = \frac{2l+d-2}{l+d-2} \binom{l+d-2}{l}$ [Frye & Efthimiou (2012)] 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 on this feature basis $g(x) = \sum_{l,m} g_{lm} Y_{lm}(x)$ Eq. 8 simplifies to

$$g^* - g = \sum_{l,m} \left[-\frac{\sigma^2/\eta}{\lambda_l + \sigma^2/\eta} - \frac{\eta^{-1} C_{K,\sigma^2/\eta}}{\lambda_l + \sigma^2/\eta} \frac{\lambda_l}{\lambda_l + \sigma^2/\eta} \right] g_{lm} Y_{lm}(x_*) \quad (10)$$

¹Note that usually the hyperspherical harmonics are normalized w.r.t Lebesgue measure on the hypersphere, but in this context the normalization is w.r.t a probability measure on the hypersphere.

where $C_{K,\sigma^2/\eta} = \sum_{lm} (\lambda_l^{-1} + \eta/\sigma^2)^{-1}$ and notably cross-talk between features has been eliminated at this order since $\sum_m \phi_{lm}(x)^2$ is constant yielding $\sum_{\bar{m}} \int d\mu_x \phi_{lm}(x) \phi_{l'm'}(x) \phi_{l\bar{m}}^2(x) \propto \delta_{l'l'} \delta_{m\bar{m}}$. By splitting the sum $C_{K,\sigma^2/\eta}$ to cases which $\lambda_l < \sigma^2/\eta$ and its complement one finds quite tight bound $C_{K,\sigma^2/\eta} < \#F \sigma^2/\eta + \sum_{lm|\lambda_l > \sigma^2/\eta} \lambda_l$, where $\#F$ is the number of non-zero kernel eigenvalues. Thus for kernels with a finite number of non-zero λ_i 's as the renormalized NTK introduced below, $C_{K,\sigma^2/\eta}$ has a η^{-1} asymptotic. This illustrates the fact the above terms are arranged by their orders in η .

Taking the leading order term one obtains the aforementioned EK result with N replaced by η . Equating the two contributions provides an estimate of when perturbation theory breaks down. Focusing on $\lambda_l > \sigma^2/\eta$, the perturbation theory appears valid when $C_{K,\sigma^2/\eta} \ll \sigma^2$. In the limit $\sigma^2 \rightarrow 0$, and for trimmed kernels, this yield $\#F \ll \eta$. Notably it means that the original non-trimmed NTK cannot be analyzed perturbatively in the noiseless limit. In the next section we tackle this issue.

5 LEARNING CURVES IN THE NOISELESS CASE ($\sigma^2 = 0$)

As argued, in the noiseless case ones expects our Eqs. (8,10) and subsequent predictions of learning curves to fail. Technically the problem lays in the perturbative expansion we performed on the exponent of $\exp(-\frac{\sum_{m=1}^M (f_m(x) - g(x))^2}{2\sigma^2})$ which is not small anymore. To overcome this we next show that the fluctuations of $f_m(x)$ associated with low λ 's ("high-energy-sector") can be traded with noise on the fluctuations of $f_m(x)$ associated with high λ 's ("low-energy-sector"). This type of reasoning where the high-energy-sector is effectively removed from the problem at the price of changing (renormalizing) some parameters in the partition function for the low energy sector, is the essence of the renormalization-group technique common in physics.

To this end, consider the expansion $K(x, x') = \sum_q b_q (x \cdot x')^q$. For two normalized datapoints x and x' , drawn from a uniform dataset on a hypersphere of radius 1, and at large d the random variable $(x \cdot x')$ is approximately Gaussian with variance $O(d^{-1})$. Since $(x \cdot x')$ is bounded to $[-1, 1]$, the random variable $(x \cdot x')^r$ must have a standard deviation which is decaying function of r . For $r \ll d$ and large d one can estimate the magnitude this standard deviation from exact known expressions and a saddle-point approximation yielding $O((d/r)^{-r/2}) \approx O(d^{-r/2})^2$. Considering next the tail of Taylor expansion $\sum_{q>r} b_q (x \cdot x')^q$ projected on the dataset $(\sum_{q>r} b_q (x_n \cdot x_m)^q)$. The resulting N by N matrix is $\sum_{q>r} b_q$ on the diagonal but $O(d^{-(r+1)/2})$ in all other entries. As we justify next, our renormalization transformation amounts to keeping only the diagonal piece of this matrix and interpreting it as noise.

Consider then Eq. 1 for g^* in two scenarios: (I) g_∞^* with the full NTK ($K(x, x')$) and no noise and (II) g_r^* with the NTK trim 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 $O(\sqrt{N} d^{-(r+1)/2})$ and bounded by $O(N d^{-(r+1)/2})$. The typical eigenvalues of $[K_r + \sigma_m^2 \mathbf{I}]^{-1}$ 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 $O(\sqrt{N} d^{-(r+1)/2} / K)$ and bounded from above by $O(N d^{-(r+1)/2} / \sigma_r^2)$. The NTK has the desirable property that σ_r^2 decays very slowly. Thus certainly in the typical case but even in the worse case scenario we expect good agreement at large r . In Fig. 1, right panel, we provide supporting numerical evidence.

²A more accurate estimate is $\left(\frac{r}{r+d-3}\right)^{r/2} \left(\frac{d-3}{r+d-3}\right)^{d/4}$

We refer to $K_r(x, x')$ as the renormalized NTKs at the scale r . As follows from Eq. (9), 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 involve 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 $\sqrt{N}d^{-(r+1)/2} \ll 1$).

Lastly we comment that our renormalization NTK approach is not limited to uniform datasets. The entire logic relies on having a rapidly decaying ratio of off-diagonal 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 multi-dimension Gaussian data distribution the input dimension (d) traded by an effective dimension (d_{eff}) defined by the variance of $(x_m \cdot x_n)$. In App. B we show an excellent agreement between the g_∞^* and g_r^* on the CIFAR10 dataset. We also provide evidence that as far as GP inference goes, CIFAR10's input distribution is well approximated by a multi-dimension Gaussian.

A numerical study of the average $(g_r^*(x_*) - g_\infty^*)^2$ averaged over x_* , for both a uniform dataset at $d = 50$ and CIFAR10 (where g^* becomes a vector of length 10 due to the one-hot encoding of the labels) are reported in App. B The DNN was a fully connected with depth 4, $\sigma_w^2 = \sigma_b^2 = 1$, and ReLU activations.

6 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 datasets with a uniform distribution normalized to the hypersphere.

To make more specific statements we now focus on the NTK kernel implied by a fully connected network of depth 4 with $\sigma_w^2 = \sigma_b^2 = 1$ and ReLU activations. We take $\eta = 3500$, $d = 50$, a target function with equal spectral weights at $l = 1, 2$. Accordingly we choose the scale $r = 3$. Experimental learning curves along with our leading and sub-leading estimates are shown in Fig. 1. left panel. See App. D for technical details on how integration of x_* , averaging over datasets, and Poisson averaging was carried.

Our analytical expressions following Eq. 9 combined with known results Jacot et al. (2018); Cho & Saul (2009) about the Taylor coefficients (b_n) yield $\lambda_0, \dots, \lambda_3 = \{3.19, 7.27e-3, 5.98e-6, 1.62e-7\}$ and $\sigma_r^2 = 0.018$. Since $\lambda_0, \lambda_1 \gg \sigma^2/\eta \gg \lambda_2, \lambda_3$ for $50 < \eta < 3500$, $C_{K_r, \sigma_r^2/\eta} \sigma^{-2} < [deg(0) + deg(1)]\sigma^2/\eta + O(deg(2)10^{-6})$, thus $C_{K_r, \sigma_r^2/\eta} \sigma^{-2} \approx 51/\eta$. Thus we expect perturbation theory to be valid for $\eta \gg 50$. At $\eta = 100$ the $l = 1$ features are learned well since $\sigma^2/\eta = 1.8e-4 \gg \lambda_1$ and the $l = 2$ features neglected, at $\eta = 1000$ they are learned but suppressed by a factor of a factor of about 3. Had the target contained $l = 3$ features, they would have been entirely neglected at these η scale.

Notably no actual DNNs were optimized in the reported learning-curve as we saw no value in re-establishing that the NTK correspondence works in the NTK-regime Jacot et al. (2018); Lee et al. (2019); Arora et al. (2019). Furthermore since our aim was to predict what the DNNs would predict rather reach SOTA predictions, we focus on reasonable hyper-parameter but did not perform any hyper-parameter optimization.

Lastly we argue that the asymptotic behavior of learning-curve we predict is more accurate than the recent PAC based bounds [Allen-Zhu et al. (2018); Cao & Gu (2019b;a)]. In App. C we show a log-log plot of the learning-curves contrasted 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.

7 DISCUSSION AND OUTLOOK

In this work we laid out a formalism based on field theory tools for predicting learning-curves in the NTK regime. Well within the validly regime of our perturbative analysis we find excellent 3% accuracy between our best estimate and the experimental curves. Central to our analysis was the

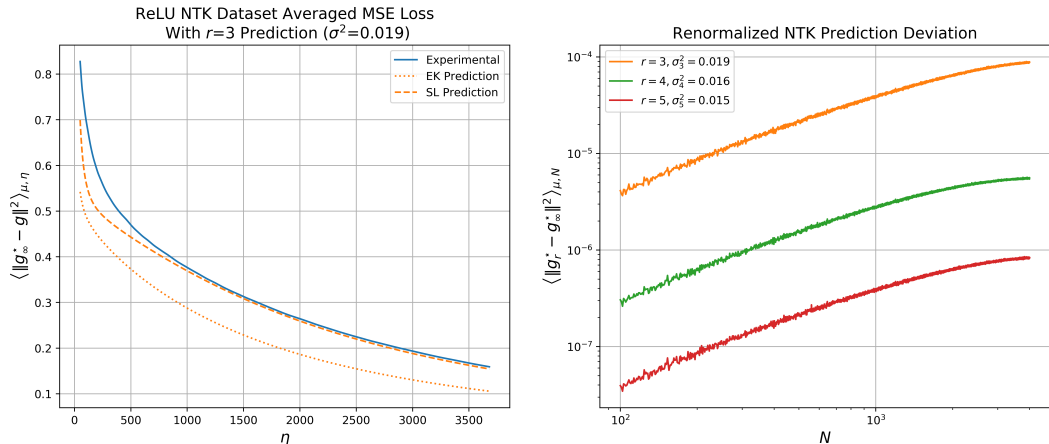


Figure 1: **Left panel:** The experimental learning curve (solid line) for a depth 4 ReLU network trained in the NTK regime on quadratic target function on a $d = 50$ hypersphere is shown along with our analytical predictions for the leading (dotted line) and leading plus sub-leading behavior (dashed line). **Right panel:** For the same dataset, we plot the dataset-averaged difference between predictions based on NTK (g_∞^*) and the renormalized NTK at scale r (g_r^*) showing an excellent agreement as r increases.

renormalization-group transformation on the NTK leading to effective observation noise on the target. Our analysis could be readily extend in several ways: Going beyond the uniform dataset case should be possible for multi-variate Gaussian input distribution with a set of similar finite variances and a set of nearly zero variances. Adding weak randomness to $K(x, x')$ to study the difference between empirical and averaged NTKs. It would also be interesting to extend our analysis to simple CNNs. The renormalized kernel can also be used for spectral analysis of the NTK and other kernels associated with DNNs.

REFERENCES

- Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and Generalization in Overparameterized Neural Networks, Going Beyond Two Layers. *arXiv e-prints*, art. arXiv:1811.04918, Nov 2018.
- Sanjeev Arora, Simon S. Du, Wei Hu, Zhiyuan Li, Ruslan Salakhutdinov, and Ruosong Wang. On Exact Computation with an Infinitely Wide Neural Net. *arXiv e-prints*, art. arXiv:1904.11955, Apr 2019.
- Douglas Azevedo and Valdir A. Menegatto. Eigenvalues of dot-product kernels on the sphere. *ArXiv e-prints*, 2015.
- Yuan Cao and Quanquan Gu. Generalization Bounds of Stochastic Gradient Descent for Wide and Deep Neural Networks. *arXiv e-prints*, art. arXiv:1905.13210, May 2019a.
- Yuan Cao and Quanquan Gu. Generalization Error Bounds of Gradient Descent for Learning Over-parameterized Deep ReLU Networks. *arXiv e-prints*, art. arXiv:1902.01384, Feb 2019b.
- Youngmin Cho and Lawrence K. Saul. Kernel methods for deep learning. In *Proceedings of the 22Nd International Conference on Neural Information Processing Systems, NIPS'09*, pp. 342–350, USA, 2009. Curran Associates Inc. ISBN 978-1-61567-911-9. URL <http://dl.acm.org/citation.cfm?id=2984093.2984132>.
- Anna Choromanska, Mikael Henaff, Michael Mathieu, Gérard Ben Arous, and Yann LeCun. The Loss Surfaces of Multilayer Networks. *arXiv e-prints*, art. arXiv:1412.0233, Nov 2014.
- A. Daniely, R. Frostig, and Y. Singer. Toward Deeper Understanding of Neural Networks: The Power of Initialization and a Dual View on Expressivity. *ArXiv e-prints*, February 2016.

- Felix Draxler, Kambis Veschgini, Manfred Salmhofer, and Fred A. Hamprecht. Essentially No Barriers in Neural Network Energy Landscape. *arXiv e-prints*, art. arXiv:1803.00885, March 2018.
- Giancarlo Ferrari-Trecate, Christopher K. I. Williams, and Manfred Opper. Finite-dimensional approximation of gaussian processes. In *NIPS*, 1998.
- Christopher Frye and Costas J. Efthimiou. Spherical Harmonics in p Dimensions. *ArXiv e-prints*, May 2012.
- E. Gardner and B. Derrida. Optimal storage properties of neural network models. *Journal of Physics A Mathematical General*, 21:271–284, January 1988. doi: 10.1088/0305-4470/21/1/031.
- A. Jacot, F. Gabriel, and C. Hongler. Neural Tangent Kernel: Convergence and Generalization in Neural Networks. *ArXiv e-prints*, June 2018.
- Jaehoon Lee, Jascha Sohl-dickstein, Jeffrey Pennington, Roman Novak, Sam Schoenholz, and Yasaman Bahri. Deep neural networks as gaussian processes. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=B1EA-M-0Z>.
- Jaehoon Lee, Lechao Xiao, Samuel S. Schoenholz, Yasaman Bahri, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide Neural Networks of Any Depth Evolve as Linear Models Under Gradient Descent. *arXiv e-prints*, art. arXiv:1902.06720, Feb 2019.
- Dörthe Malzahn and Manfred Opper. A variational approach to learning curves. In *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic*, NIPS’01, pp. 463–469, Cambridge, MA, USA, 2001. MIT Press. URL <http://dl.acm.org/citation.cfm?id=2980539.2980600>.
- Charles A. Micchelli and Grace Wahba. Design problems for optimal surface interpolation. 1979.
- Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In Search of the Real Inductive Bias: On the Role of Implicit Regularization in Deep Learning. *arXiv e-prints*, art. arXiv:1412.6614, December 2014.
- Behnam Neyshabur, Zhiyuan Li, Srinadh Bhojanapalli, Yann LeCun, and Nathan Srebro. Towards Understanding the Role of Over-Parametrization in Generalization of Neural Networks. *arXiv e-prints*, art. arXiv:1805.12076, May 2018.
- Roman Novak, Lechao Xiao, Jaehoon Lee, Yasaman Bahri, Greg Yang, Daniel A. Abolafia, Jeffrey Pennington, and Jascha Sohl-Dickstein. Bayesian Deep Convolutional Networks with Many Channels are Gaussian Processes. *arXiv e-prints*, art. arXiv:1810.05148, October 2018.
- Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005. ISBN 026218253X.
- K. Ritter. *Average-Case Analysis of Numerical Problems*. Lecture Notes in Mathematics. Springer Berlin Heidelberg, 2007. ISBN 9783540455929. URL https://books.google.co.il/books?id=X_l6CwAAQBAJ.
- Klaus Ritter. Asymptotic optimality of regular sequence designs. *Ann. Statist.*, 24(5):2081–2096, 10 1996. doi: 10.1214/aos/1069362311. URL <https://doi.org/10.1214/aos/1069362311>.
- L.S. Schulman. Techniques and applications of path integration. 1996. URL <https://books.google.co.il/books?id=Cuc9AQAAIAAJ>.
- Pierre Sermanet, David Eigen, Xiang Zhang, Michael Mathieu, Rob Fergus, and Yann LeCun. OverFeat: Integrated Recognition, Localization and Detection using Convolutional Networks. *arXiv e-prints*, art. arXiv:1312.6229, December 2013.
- Peter Sollich. Gaussian Process Regression with Mismatched Models. *arXiv e-prints*, art. cond-mat/0106475, Jun 2001.


Christopher K. I. Williams and Francesco Vivarelli. Upper and lower bounds on the learning curve for gaussian processes. *Mach. Learn.*, 40(1):77–102, July 2000. ISSN 0885-6125. doi: 10.1023/A:1007601601278. URL <https://doi.org/10.1023/A:1007601601278>.

Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS'14*, pp. 3320–3328, Cambridge, MA, USA, 2014. MIT Press. URL <http://dl.acm.org/citation.cfm?id=2969033.2969197>.

Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv e-prints*, art. arXiv:1611.03530, November 2016.


A POISSON AVERAGING DEMONSTRATION

Here we demonstrate that Poisson averaging has no substantial effect on the learning curve. To this end we show the experimental learning curve from the main text pre- and post-averaging. 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.




B COMPARISON OF NTK AND RENORMALIZED NTK PREDICTIONS ON NON-UNIFORM DATASET

While our lack of knowledge of the NTK eigenvalues and eigenfunctions with respect to a non-uniform measure prevents us from predicting learning curves, we would like to show that the renormalized NTK is still a valid approximation in this setting. To this end we compare the prediction of the NTK and renormalized NTK on the one-hot encoding of the cifar-10 dataset.



C 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. Here we show that given a functional behavior of the experimental learning curves such a bound cannot be nearly as tight as our predictions.



D FURTHER DETAILS ON THE EXPERIMENTS

In Fig. 1. of the main text we generated datasets D_N of N uniformly distributed points on the unit sphere S^{d-1} . For each such dataset we obtained the expected MSE loss ($\|g_\infty^* - g\|^2$) of the NTK GPR (g_∞^*) by numerical integration over x_* . Repeating this process many times we obtained the dataset averaged loss ($\langle \|g_\infty^* - g\|^2 \rangle_{\mu, N}$) for $N = 1, 2, \dots, N_{\max}$, within an estimated 5% accuracy (this typically required averaging over 10 datasets). For direct comparison with our prediction for the learning curve, we computed the Poisson averaged learning curve in accordance with Eq. 5 as $\langle \|g_\infty^* - g\|^2 \rangle_{\mu, \eta} = e^{-\eta} \sum_{n=0}^{\infty} \eta^n \langle \|g_\infty^* - g\|^2 \rangle_{\mu, n} / n!$, neglecting the terms $n > N_{\max}$. We restricted ourselves to $\eta_{\max} \leq N_{\max} - 5\sqrt{N_{\max}}$ to make tail effects negligible.

E NOTATIONS FOR THE FIELD THEORY DERIVATION.

For completeness, here we re-state the notations used in the main-text.

Ω - Space space of inputs.

x, x', x^* - Inputs (in Ω).

μ_x - Measure on Ω .

$K(x, x')$ - Kernel function (covariance) of a Gaussian process. Assumed to be symmetric and positive-semi-definite.

$\phi_i(x)$ - i 'th eigenfunction of $K(x, x')$. By the spectral theorem, the set $\{\phi_i\}_{i=1}^{\infty}$ can be assumed to be orthonormal:

$$\int_{x \in \Omega} d\mu_x \phi_i(x) \phi_j(x) = \delta_{ij}$$

λ_i - i 'th eigenvalue of $K(x, x')$.

$$\int_{x' \in \Omega} d\mu_{x'} K(x, x') \phi_i(x') = \lambda_i \phi_i(x)$$

$\|\cdot\|_{\mathcal{H}_K}$ - RKHS norm. If $f(x) = \sum_i f_i \phi_i(x)$ then $\|f\|_{\mathcal{H}_K} = \sum_i \frac{f_i^2}{\lambda_i}$ (where ϕ_i is an orthonormal set). Note that this norm is independent of μ_x .

$g(x)$ - The target function.

σ^2 - Noise variance. The noise is assumed to be Gaussian.

N - Number of inputs in the data-set.

D_N - Data-set of size N , $D_N = \{x_1, \dots, x_N\}$.

$f_{D_N, \sigma^2}^*(x)$ - The prediction function.

F EXPLICIT PATH INTEGRAL COMPUTATIONS

Here we wish to prove the probability function defined in Eq. 2. of the main text yields the GP defined by a given kernel using explicit computation of the path integral. Denoting

$\int d\mu_x d\mu_{x'} f(x) K^{-1}(x, x') f(x')$ as $\|f\|_K^2$ and noting that $\|f\|_K^2 = \sum_i \frac{f_i^2}{\lambda_i}$:

$$\begin{aligned} & \frac{\int Df \cdot f(x) \cdot f(y) \cdot \exp\left(-\frac{1}{2} \|f\|_K^2\right)}{\int Df \exp\left(-\frac{1}{2} \|f\|_K^2\right)} = \frac{\int \prod_i df_i \cdot \sum_i f_i \phi_i(x) \cdot \sum_j f_j \phi_j(y) \cdot \exp\left(-\frac{1}{2} \sum_l \frac{f_l^2}{\lambda_l}\right)}{\int \prod_i df_i \exp\left(-\frac{1}{2} \sum_l \frac{f_l^2}{\lambda_l}\right)} = \\ & = \sum_i \underbrace{\frac{\int df \cdot f^2 \cdot \exp\left(-\frac{f^2}{2\lambda_i}\right)}{\int df \exp\left(-\frac{f^2}{2\lambda_i}\right)}}_{\lambda_i} \phi_i(x) \phi_i(y) + \sum_{i \neq j} \underbrace{\frac{\int df \cdot f \cdot \exp\left(-\frac{f^2}{2\lambda_i}\right)}{\int df \exp\left(-\frac{f^2}{2\lambda_i}\right)}}_0 \cdot \underbrace{\frac{\int df \cdot f \cdot \exp\left(-\frac{f^2}{2\lambda_j}\right)}{\int df \exp\left(-\frac{f^2}{2\lambda_j}\right)}}_0 = \\ & = \sum_i \lambda_i \phi_i(x) \phi_i(y) = K(x, y) \end{aligned}$$

G GAUSSIAN PROCESS PREDICTION AS A FIELD THEORY

Let us assume a Gaussian process (GP) with mean 0 and co-variance function $K(x, x')$. For a data-set D_N of size N and noisy targets $\{g(x_i)\}_{i=1}^N$, it is known that the posterior mean obtained by

Bayesian inference is

$$f_{D_N, \sigma^2}^* = \arg \min_f \left[\frac{1}{2} \|f\|_{\mathcal{H}_K}^2 + \sum_{x_i \in D_N} \frac{(f(x_i) - g(x_i))^2}{2\sigma^2} \right]$$

For a data-set D_N of size N and noisy targets $\{g(x_i)\}_{i=1}^N$, we present the GP canonical partition function:

$$Z_{D_N, \sigma^2} [\alpha(x)] \stackrel{def}{=} \int Df \exp \left(-\frac{1}{2} \|f\|_{\mathcal{H}_K}^2 + \int \alpha(x) f(x) dx - \sum_{x_i \in D_N} \frac{(f(x_i) - g(x_i))^2}{2\sigma^2} \right)$$

Where the Df notation stands for path integral. Notice that the functional derivative of $\log(Z_{D_N, \sigma^2}[\alpha(x)])$ w.r.t $\alpha(x^*)$ at $\alpha(x) = 0$ yields:

$$\begin{aligned} \frac{\partial}{\partial \alpha(x^*)} \Big|_{\alpha(x)=0} \log(Z_{D_N, \sigma^2}[\alpha(x)]) &= \frac{1}{Z_{D_N, \sigma^2}[\alpha(x)=0]} \cdot \frac{\partial}{\partial \alpha(x^*)} \Big|_{\alpha(x)=0} (Z_{D_N, \sigma^2}[\alpha(x)]) = \\ &= \frac{\int Df \cdot f(x^*) \exp \left(-\frac{1}{2} \|f\|_{\mathcal{H}_K}^2 - \sum_{x_i \in D_N} \frac{(f(x_i) - g(x_i))^2}{2\sigma^2} \right)}{\int Df \exp \left(-\frac{1}{2} \|f\|_{\mathcal{H}_K}^2 - \sum_{x_i \in D_N} \frac{(f(x_i) - g(x_i))^2}{2\sigma^2} \right)} = \arg \min_{f|x^*} \left[\frac{1}{2} \|f\|_{\mathcal{H}_K}^2 + \sum_{i=1}^N \frac{(f(x_i) - g(x_i))^2}{2\sigma^2} \right] \end{aligned}$$

where the last equality is due to the fact that for Gaussian distributions, the expected value coincides with the most probable value. Therefore, the exact Bayesian inference mean:

$$f_{D_N, \sigma^2}^*(x^*) = \frac{\partial}{\partial \alpha(x^*)} \Big|_{\alpha(x)=0} \log(Z_{D_N, \sigma^2}[\alpha(x)])$$

G.1 CANONICAL ENSEMBLE FORMALISM

for evaluating the quality of a certain GP, we're interested in the average prediction for all the data-sets of size N , meaning:

$$f_{N, \sigma^2}^C(x^*) \stackrel{def}{=} \langle f_{D_N, \sigma^2}^*(x^*) \rangle_{D_N \sim \mu_x^N} = \int d\mu_{x_1} \int d\mu_{x_2} \dots \int d\mu_{x_N} f_{D_N = \{x_1, \dots, x_N\}, \sigma^2}^*(x^*)$$

Using the replica trick we obtain:

$$f_{N, \sigma^2}^C(x^*) = \lim_{M \rightarrow 0} \frac{\partial}{\partial \alpha(x^*)} \Big|_{\alpha(x)=0} \frac{\langle Z_{D_N, \sigma^2}^M[\alpha(x)] \rangle_{D_N \sim \mu_x^N} - 1}{M}$$

Now, let us calculate $\langle Z_{D_N, \sigma^2}^M[\alpha(x)] \rangle_{D_N \sim \mu_x^N}$ for an integer M :

$$\begin{aligned}
Z_{D_N, \sigma^2}^M [\alpha(x)] &= \underbrace{\int \int \dots \int}_M \prod_{j=1}^M Df_j \\
&\exp \left(-\frac{1}{2} \sum_{j=1}^M \|f_j\|_{\mathcal{H}_K}^2 + \sum_{j=1}^M \int \alpha(x) f_j(x) dx - \sum_{j=1}^M \sum_{x_i \in D_N} \frac{(f_j(x_i) - g(x_i))^2}{2\sigma^2} \right) \\
Z_{N, M, \sigma^2} &= \langle Z_{D_N, \sigma^2}^M [\alpha(x)] \rangle_{D_N \sim \mu_x^N} = \int d\mu_{x_1} \dots \int d\mu_{x_N} Z_{D_N, \sigma^2}^M [\alpha(x)] = \\
&= \underbrace{\int \dots \int}_M \prod_{j=1}^M Df_j \exp \left(-\frac{1}{2} \sum_{j=1}^M \|f_j\|_{\mathcal{H}_K}^2 + \sum_{j=1}^M \int \alpha(x) f_j(x) dx \right) \left\langle \exp \left(-\sum_{j=1}^M \frac{(f_j(x) - g(x))^2}{2\sigma^2} \right) \right\rangle_{x \sim \mu_x}^N
\end{aligned}$$

so

$$f_{N, \sigma^2}^C(x^*) = \lim_{M \rightarrow 0} \frac{\frac{\partial}{\partial \alpha(x^*)} \Big|_{\alpha(x)=0} \langle Z_{D_N, \sigma^2}^M [\alpha(x)] \rangle_{D_N \sim \mu_x^N}}{M}$$

G.2 GRAND CANONICAL ENSEMBLE FORMALISM

We now wish to allow fluctuations in the value of N , meaning averaging over $f_{N, \sigma^2}^C(x^*)$ for different values of N . The motivation is to simplify the calculations, while averaging around a relatively confined region of N s. Let us average the canonical prediction while weighting N according to Poisson distribution with expected value η :

$$f_{\eta, \sigma^2}^{GC}(x^*) \stackrel{def}{=} \sum_{N=0}^{\infty} \frac{e^{-\eta} \eta^N}{N!} f_{N, \sigma^2}^C(x^*).$$

and defining:

$$\mathcal{Z}_{\eta, M, \sigma^2} [\alpha(x)] = \sum_{N=0}^{\infty} \frac{e^{-\eta} \eta^N}{N!} \langle Z_{D_N, \sigma^2}^M [\alpha(x)] \rangle_{D_N \sim \mu_x^N}$$

we get:

$$f_{\eta, \sigma^2}^{GC}(x^*) = \frac{\partial}{\partial \alpha(x^*)} \Big|_{\alpha(x)=0} \lim_{M \rightarrow 0} \frac{\mathcal{Z}_{\eta, M, \sigma^2} [\alpha(x)]}{M}$$

That is, the functional derivative w.r.t $\alpha(x^*)$ at $\alpha(x) = 0$ yields the average prediction, averaged over different data-set sizes (the canonical averaging) and different data-sets for each size (the grand canonical averaging).

For a given η , the standard deviation of N is η , so the relative error is $\frac{1}{\sqrt{\eta}}$, decreases with η .

Substituting $\langle Z_{D_N, \sigma^2}^M [\alpha(x)] \rangle_{D_N \sim \mu_x^N}$ in the expression for $\mathcal{Z}_{\eta, M, \sigma^2} [\alpha(x)]$ we obtain:

$$\begin{aligned}
& \langle Z_{D_N, \sigma^2}^M [\alpha(x)] \rangle_{D_N \sim \mu_x^N} = \\
& = \sum_{N=0}^{\infty} \frac{e^{-\eta} \eta^N}{N!} \underbrace{\int \dots \int}_{M} \prod_{j=1}^M Df_j \exp \left(-\frac{1}{2} \sum_{j=1}^M \|f_j\|_{\mathcal{H}_K}^2 + \sum_{j=1}^M \int \alpha(x) f_j(x) dx \right) \\
& \quad \left\langle \exp \left(-\sum_{j=1}^M \frac{(f_j(x) - g(x))^2}{2\sigma^2} \right) \right\rangle_{x \sim \mu_x}^N = \\
& = \underbrace{\int \dots \int}_{M} \prod_{j=1}^M Df_j \exp \left(-\frac{1}{2} \sum_{j=1}^M \|f_j\|_{\mathcal{H}_K}^2 + \sum_{j=1}^M \int \alpha(x) f_j(x) dx \right) \\
& \quad \sum_{N=0}^{\infty} \frac{e^{-\eta} \eta^N}{N!} \left\langle \exp \left(-\sum_{j=1}^M \frac{(f_j(x) - g(x))^2}{2\sigma^2} \right) \right\rangle_{x \sim \mu_x}^N = \\
& \quad = e^{-\eta} \underbrace{\int \dots \int}_{M} Df_1 \dots Df_M \\
& \exp \left(-\frac{1}{2} \sum_{j=1}^M \|f_j\|_{\mathcal{H}_K}^2 + \sum_{j=1}^M \int \alpha(x) f_j(x) dx + \eta \left\langle \exp \left(-\sum_{j=1}^M \frac{(f_j(x) - g(x))^2}{2\sigma^2} \right) \right\rangle_{x \sim \mu_x} \right)
\end{aligned}$$

G.3 DERIVING THE EQUIVALENCE KERNEL USING THE GRAND CANONICAL FORMALISM

We wish to get rid of the exponent inside the exponent. Expanding it using (first order) Taylor series:

$$\begin{aligned}
& \mathcal{Z}_{\eta, M, \sigma^2} [\alpha(x)] = \\
& = e^{-\eta} \underbrace{\int \dots \int}_{M} \prod_{j=1}^M Df_j \\
& \exp \left(-\frac{1}{2} \sum_{j=1}^M \|f_j\|_{\mathcal{H}_K}^2 + \sum_{j=1}^M \int \alpha(x) f_j(x) dx + \eta \left\langle \exp \left(-\sum_{j=1}^M \frac{(f_j(x) - g(x))^2}{2\sigma^2} \right) \right\rangle_{x \sim \mu_x} \right) \approx \\
& \approx e^{-\eta} \underbrace{\int \dots \int}_{M} \prod_{j=1}^M Df_j \\
& \exp \left(-\frac{1}{2} \sum_{j=1}^M \|f_j\|_{\mathcal{H}_K}^2 + \sum_{j=1}^M \int \alpha(x) f_j(x) dx + \eta \left\langle 1 - \sum_{j=1}^M \frac{(f_j(x) - g(x))^2}{2\sigma^2} \right\rangle_{x \sim \mu_x} \right) = \\
& = \left[\int Df \exp \left(-\frac{1}{2} \|f\|_{\mathcal{H}_K}^2 + \int \alpha(x) f(x) dx - \eta \left\langle \frac{(f(x) - g(x))^2}{2\sigma^2} \right\rangle_{x \sim \mu_x} \right) \right]^M \stackrel{def}{=} (\mathcal{Z}_{\eta, \sigma^2}^{EK} [\alpha(x)])^M \\
& \mathcal{Z}_{\eta, \sigma^2}^{EK} [\alpha(x)] = \int Df \exp \left(-\frac{1}{2} \|f\|_{\mathcal{H}_K}^2 + \int \alpha(x) f(x) d\mu_x - \frac{\eta}{2\sigma^2} \int d\mu_x (f(x) - g(x))^2 \right)
\end{aligned}$$

and without any additional approximation:

$$\begin{aligned}
f_{\eta, \sigma^2}^{GC}(x^*) &= \left. \frac{\partial}{\partial \alpha(x^*)} \right|_{\alpha(x)=0} \lim_{M \rightarrow 0} \frac{\mathcal{Z}_{\eta, M, \sigma^2}[\alpha(x)] - 1}{M} \approx \left. \frac{\partial}{\partial \alpha(x^*)} \right|_{\alpha(x)=0} \lim_{M \rightarrow 0} \frac{\left(\mathcal{Z}_{\eta, \sigma^2}^{EK}[\alpha(x)]\right)^M - 1}{M} = \\
&= \left. \frac{\partial}{\partial \alpha(x^*)} \right|_{\alpha(x)=0} \log \left(\mathcal{Z}_{\eta, \sigma^2}^{EK}[\alpha(x)]\right) = \frac{1}{\mathcal{Z}_{\eta, \sigma^2}^{EK}[\alpha(x)=0]} \cdot \left. \frac{\partial}{\partial \alpha(x^*)} \right|_{\alpha(x)=0} \left(\mathcal{Z}_{\eta, \sigma^2}^{EK}[\alpha(x)]\right) = \\
&= \frac{1}{\mathcal{Z}_{\eta, \sigma^2}^{EK}[\alpha(x)=0]} \cdot \int Df \cdot f(x^*) \\
&\exp \left(-\frac{1}{2} \|f\|_{\mathcal{H}_K}^2 + \int \alpha(x) f(x) dx - \frac{\eta}{2\sigma^2} \int d\mu_x (f(x) - g(x))^2 \right)_{\alpha(x)=0} = \\
&= \frac{\int Df \cdot f(x^*) \exp \left(-\frac{1}{2} \|f\|_{\mathcal{H}_K}^2 - \frac{\eta}{2\sigma^2} \int d\mu_x (f(x) - g(x))^2 \right)}{\int Df \exp \left(-\frac{1}{2} \|f\|_{\mathcal{H}_K}^2 - \frac{\eta}{2\sigma^2} \int d\mu_x (f(x) - g(x))^2 \right)} = \\
&= \arg \min_{f|_{x^*}} \left[\frac{1}{2} \|f\|_{\mathcal{H}_K}^2 + \frac{\eta}{2\sigma^2} \int d\mu_x (f(x) - g(x))^2 \right] \stackrel{def}{=} f_{\eta, \sigma^2}^{EK}(x^*)
\end{aligned}$$

and that is exactly the result for the equivalence kernel, where η is the data-set size (we regarded it as the mean of the data-set size).

Let us derive it explicitly. For $f(x) = \sum_i f_i \phi_i(x)$ and $g(x) = \sum_i g_i \phi_i(x)$:

$$\begin{aligned}
f_{\eta, \sigma^2}^{EK}(x^*) &= \frac{\int Df \cdot f(x^*) \exp \left(-\frac{1}{2} \|f\|_{\mathcal{H}_K}^2 - \frac{\eta}{2\sigma^2} \int d\mu_x (f(x) - g(x))^2 \right)}{\int Df \exp \left(-\frac{1}{2} \|f\|_{\mathcal{H}_K}^2 - \frac{\eta}{2\sigma^2} \int d\mu_x (f(x) - g(x))^2 \right)} = \\
&= \frac{\int \prod_i df_i \cdot \sum_i f_i \phi_i(x^*) \cdot \exp \left(-\frac{1}{2} \sum_i \left(\frac{f_i^2}{\lambda_i} + \frac{\eta}{\sigma^2} (f_i - g_i)^2 \right) \right)}{\int \prod_i df_i \exp \left(-\frac{1}{2} \sum_i \left(\frac{f_i^2}{\lambda_i} + \frac{\eta}{\sigma^2} (f_i - g_i)^2 \right) \right)} = \\
&= \sum_i \phi_i(x^*) \frac{\int df_i \cdot f_i \cdot \exp \left(-\frac{f_i^2}{2\lambda_i} - \frac{\eta}{2\sigma^2} (f_i - g_i)^2 \right)}{\int df_i \exp \left(-\frac{f_i^2}{2\lambda_i} - \frac{\eta}{2\sigma^2} (f_i - g_i)^2 \right)} = \sum_i \frac{\lambda_i}{\lambda_i + \frac{\sigma^2}{\eta}} g_i \phi_i(x^*)
\end{aligned}$$

G.4 EQUIVALENCE KERNEL AS FREE FIELD THEORY

Regarding the Equivalence Kernel as the free (quadratic) theory, we can denote $f_{\eta, \sigma^2}^{EK}(x^*) = \langle f(x^*) \rangle_{f \sim EK} = \langle f(x^*) \rangle_0 = \sum_i \frac{\lambda_i}{\lambda_i + \frac{\sigma^2}{\eta}} g_i \phi_i(x^*)$.

Let us calculate the correlations in the free theory:

$$\begin{aligned}
\langle f(x) f(y) \rangle_0 &= \frac{\int Df \cdot f(x) f(y) \exp\left(-\frac{1}{2} \|f\|_{\mathcal{H}_K}^2 - \frac{\eta}{2\sigma^2} \int d\mu_x (f(x) - g(x))^2\right)}{\int Df \exp\left(-\frac{1}{2} \|f\|_{\mathcal{H}_K}^2 - \frac{\eta}{2\sigma^2} \int d\mu_x (f(x) - g(x))^2\right)} = \\
&= \frac{\int \prod_i df_i \cdot \sum_{i,j} f_i f_j \phi_i(x) \phi_j(y) \cdot \exp\left(-\frac{1}{2} \sum_i \left(\frac{f_i^2}{\lambda_i} + \frac{\eta}{\sigma^2} (f_i - g_i)^2\right)\right)}{\int \prod_i df_i \exp\left(-\frac{1}{2} \sum_i \left(\frac{f_i^2}{\lambda_i} + \frac{\eta}{\sigma^2} (f_i - g_i)^2\right)\right)} = \\
&= \sum_i \underbrace{\frac{\int df_i \cdot f_i^2 \cdot \exp\left(-\frac{f_i^2}{2\lambda_i} - \frac{\eta}{2\sigma^2} (f_i - g_i)^2\right)}{\int df_i \exp\left(-\frac{f_i^2}{2\lambda_i} - \frac{\eta}{2\sigma^2} (f_i - g_i)^2\right)}}_{\frac{\lambda_i^2 g_i^2}{(\lambda_i + \frac{\sigma^2}{\eta})^2} + \left(\frac{1}{\lambda_i} + \frac{\eta}{\sigma^2}\right)^{-1}} \phi_i(x) \phi_i(y) + \sum_{i \neq j} \frac{\lambda_i g_i \lambda_j g_j}{\left(\lambda_i + \frac{\sigma^2}{\eta}\right) \left(\lambda_j + \frac{\sigma^2}{\eta}\right)} \phi_i(x) \phi_j(y) = \\
&= \sum_i \left(\frac{1}{\lambda_i} + \frac{\eta}{\sigma^2}\right)^{-1} \phi_i(x) \phi_i(y) + \underbrace{\sum_{i,j} \frac{\lambda_i g_i \lambda_j g_j}{\left(\lambda_i + \frac{\sigma^2}{\eta}\right) \left(\lambda_j + \frac{\sigma^2}{\eta}\right)} \phi_i(x) \phi_j(y)}_{\langle f(x) \rangle_0 \langle f(y) \rangle_0}
\end{aligned}$$

Therefore:

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

and we see that the correlations are $O\left(\frac{1}{\eta}\right)$.

For rotationally invariant kernel, we get that

$$\text{Var}[f(x)] = \sum_{l=0}^{\infty} \sum_{m=0}^{\text{deg}(l)} \left(\frac{1}{\lambda_l} + \frac{\eta}{\sigma^2}\right)^{-1} Y_{l,m}^2(x) \stackrel{\text{def}}{=} C_{K,\eta,\sigma^2}$$

is independent of x since

$$\sum_{m=0}^{\text{deg}(l)} Y_{l,m}^2(x) = \text{deg}(l)$$

so

$$C_{K,\eta,\sigma^2} = \sum_{l=0}^{\infty} \sum_{m=0}^{\text{deg}(l)} \frac{1}{\lambda_l^{-1} + \eta/\sigma^2}$$

G.5 PERTUBATIVE CORRECTION FOR THE EQUIVALENCE KERNEL

G.5.1 AVERAGING f

Going to the next order in the expansion:

$$\begin{aligned}
& \mathcal{Z}_{\eta, M, \sigma^2} [\alpha(x)] = \\
& = e^{-\eta} \underbrace{\int \dots \int}_M \prod_{j=1}^M Df_j \\
& \exp \left(-\frac{1}{2} \sum_{j=1}^M \|f_j\|_{\mathcal{H}_K}^2 + \sum_{j=1}^M \int \alpha(x) f_j(x) dx + \eta \left\langle \exp \left(-\sum_{j=1}^M \frac{(f_j(x) - g(x))^2}{2\sigma^2} \right) \right\rangle_{x \sim \mu_x} \right) \approx \\
& \approx e^{-\eta} \underbrace{\int \dots \int}_M \prod_{j=1}^M Df_j \\
& \exp \left(-\frac{1}{2} \sum_{j=1}^M \|f_j\|_{\mathcal{H}_K}^2 + \sum_{j=1}^M \int \alpha(x) f_j(x) dx \right) \\
& \exp \left(\eta \left\langle 1 - \sum_{j=1}^M \frac{(f_j(x) - g(x))^2}{2\sigma^2} + \frac{1}{2} \left(\sum_{j=1}^M \frac{(f_j(x) - g(x))^2}{2\sigma^2} \right)^2 \right\rangle_{x \sim \mu_x} \right) = \\
& = \underbrace{\int \dots \int}_M \prod_{j=1}^M Df_j \\
& \exp \left(\sum_{j=1}^M \left(-\frac{1}{2} \|f_j\|_{\mathcal{H}_K}^2 + \int \alpha(x) f_j(x) dx - \frac{\eta}{2\sigma^2} \int d\mu_x (f_j(x) - g(x))^2 \right) \right) \\
& \exp \left(\frac{\eta}{8\sigma^4} \sum_{j=1}^M \sum_{l=1}^M \int d\mu_x (f_j(x) - g(x))^2 \cdot (f_l(x) - g(x))^2 \right)
\end{aligned}$$

Note that:

$$\begin{aligned}
f_{\eta, \sigma^2}^{GC}(x^*) &= \frac{\partial}{\partial \alpha(x^*)} \Big|_{\alpha(x)=0} \lim_{M \rightarrow 0} \frac{\mathcal{Z}_{\eta, M, \sigma^2} [\alpha(x)] - 1}{M} = \\
&= \lim_{M \rightarrow 0} \left[\left(\mathcal{Z}_{\eta, \sigma^2}^{EK} [\alpha(x) = 0] \right)^M \cdot \frac{\frac{\partial}{\partial \alpha(x^*)} \Big|_{\alpha(x)=0} \mathcal{Z}_{\eta, M, \sigma^2} [\alpha(x)]}{\left(\mathcal{Z}_{\eta, \sigma^2}^{EK} [\alpha(x) = 0] \right)^M} \cdot \frac{1}{M} \right] = \\
&= \underbrace{\lim_{M \rightarrow 0} \left(\mathcal{Z}_{\eta, \sigma^2}^{EK} [\alpha(x) = 0] \right)^M}_1 \cdot \lim_{M \rightarrow 0} \left[\frac{\frac{\partial}{\partial \alpha(x^*)} \Big|_{\alpha(x)=0} \mathcal{Z}_{\eta, M, \sigma^2} [\alpha(x)]}{\left(\mathcal{Z}_{\eta, \sigma^2}^{EK} [\alpha(x) = 0] \right)^M} \cdot \frac{1}{M} \right]
\end{aligned}$$

Calculating the first order perturbative corrections:

$$\begin{aligned}
f_{\eta, \sigma^2}^{GC}(x^*) &= \lim_{M \rightarrow 0} \frac{\frac{\partial}{\partial \alpha(x^*)} \Big|_{\alpha(x)=0} \mathcal{Z}_{\eta, M, \sigma^2}[\alpha(x)]}{M \cdot \left(\mathcal{Z}_{\eta, \sigma^2}^{EK}[\alpha(x)=0] \right)^M} = \\
&= \lim_{M \rightarrow 0} \frac{1}{M \mathcal{Z}_{\eta, \sigma^2}^{EK}[0]^M} \underbrace{\int \dots \int}_{M} \prod_{j=1}^M Df_j \\
&\quad \exp \left(-\frac{1}{2} \sum_{j=1}^M \|f_j\|_{\mathcal{H}_K}^2 - \frac{\eta}{2\sigma^2} \sum_{j=1}^M \int d\mu_x (f_j(x) - g(x))^2 \right) \\
&\quad \exp \left(\frac{\eta}{8\sigma^4} \sum_{j=1}^M \sum_{l=1}^M \int d\mu_x (f_j(x) - g(x))^2 \cdot (f_l(x) - g(x))^2 \right) \cdot \sum_{i=1}^M f_i(x^*) = \\
&= \lim_{M \rightarrow 0} \frac{1}{M \mathcal{Z}_{\eta, \sigma^2}^{EK}[0]^M} \underbrace{\int \dots \int}_{M} \prod_{j=1}^M Df_j \exp \left(-\frac{1}{2} \sum_{j=1}^M \|f_j\|_{\mathcal{H}_K}^2 - \frac{\eta}{2\sigma^2} \sum_{j=1}^M \int d\mu_x (f_j(x) - g(x))^2 \right) \\
&\quad \left(1 + \frac{\eta}{8\sigma^4} \sum_{j=1}^M \sum_{l=1}^M \int d\mu_x (f_j(x) - g(x))^2 \cdot (f_l(x) - g(x))^2 \right) \cdot \sum_{i=1}^M f_i(x^*) + O\left(\frac{1}{\eta^3}\right) = \\
&= \lim_{M \rightarrow 0} \frac{1}{M} \left\langle \left(1 + \frac{\eta}{8\sigma^4} \sum_{j=1}^M \sum_{l=1}^M \int d\mu_x (f_j(x) - g(x))^2 \cdot (f_l(x) - g(x))^2 \right) \cdot \sum_{i=1}^M f_i(x^*) \right\rangle_{f_1, \dots, f_M \sim EK} \\
&\quad + O\left(\frac{1}{\eta^3}\right) = \\
&= \langle f(x^*) \rangle_0 + \lim_{M \rightarrow 0} \frac{1}{M} \left\langle \left(\frac{\eta}{8\sigma^4} \sum_{j=1}^M \sum_{l=1}^M \int d\mu_x (f_j(x) - g(x))^2 \cdot (f_l(x) - g(x))^2 \right) \cdot \sum_{i=1}^M f_i(x^*) \right\rangle_{f_1, \dots, f_M \sim EK} \\
&\quad + O\left(\frac{1}{\eta^3}\right) = \\
&\quad = \langle f(x^*) \rangle_0 \\
&+ \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 \cdot (f_l(x) - g(x))^2 f_i(x^*) \right\rangle_{f_1, \dots, f_M \sim EK} + O\left(\frac{1}{\eta^3}\right).
\end{aligned}$$

Calculating the correction:

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

Note that we eventually divide by M and take the limit $M \rightarrow 0$, so we only care about $O(M)$ terms:

$$\begin{aligned}
& f_{\eta, \sigma^2}^{GC}(x^*) = f_{\eta, \sigma^2}^{EK}(x^*) \\
& + \frac{\eta}{8\sigma^4} \int d\mu_x \left[\left\langle (f(x) - g(x))^4 f(x^*) \right\rangle_0 - 2 \left\langle (f(x) - g(x))^2 \right\rangle_0 \left\langle (f(x) - g(x))^2 f(x^*) \right\rangle_0 \right. \\
& \quad \left. - \left\langle (f(x) - g(x))^4 \right\rangle_0 \langle f(x^*) \rangle_0 + 2 \left\langle (f(x) - g(x))^2 \right\rangle_0^2 \langle f(x^*) \rangle_0 \right] + O\left(\frac{1}{\eta^3}\right)
\end{aligned}$$

These correlations can be calculated using Feynman diagrams, since the free theory (EK) is quadratic (Gaussian):

$$\begin{aligned}
& \left\langle (f(x) - g(x))^4 f(x^*) \right\rangle_0 = \\
& = 3f_{\eta, \sigma^2}^{EK}(x^*) \text{Var}[f(x)]^2 + 6f_{\eta, \sigma^2}^{EK}(x^*) (f_{\eta, \sigma^2}^{EK}(x) - g(x))^2 \text{Var}[f(x)] \\
& + f_{\eta, \sigma^2}^{EK}(x^*) (f_{\eta, \sigma^2}^{EK}(x) - g(x))^4 + 4(f_{\eta, \sigma^2}^{EK}(x) - g(x))^3 \text{Cov}[f(x), f(x^*)] \\
& + 12(f_{\eta, \sigma^2}^{EK}(x) - g(x)) \text{Var}[f(x)] \text{Cov}[f(x), f(x^*)]
\end{aligned}$$

$$\begin{aligned}
& \left\langle (f(x) - g(x))^4 \right\rangle_0 \langle f(x^*) \rangle_0 = \\
& = 3f_{\eta, \sigma^2}^{EK}(x^*) \text{Var}[f(x)]^2 \\
& + 6f_{\eta, \sigma^2}^{EK}(x^*) (f_{\eta, \sigma^2}^{EK}(x) - g(x))^2 \text{Var}[f(x)] + f_{\eta, \sigma^2}^{EK}(x^*) (f_{\eta, \sigma^2}^{EK}(x) - g(x))^4
\end{aligned}$$

$$\begin{aligned}
& \left\langle (f(x) - g(x))^2 \right\rangle_0 \left\langle (f(x) - g(x))^2 f(x^*) \right\rangle_0 = \\
& = 2\text{Var}[f(x)] \text{Cov}[f(x), f(x^*)] (f_{\eta, \sigma^2}^{EK}(x) - g(x)) + f_{\eta, \sigma^2}^{EK}(x^*) \text{Var}[f(x)]^2 \\
& + 2\text{Cov}[f(x), f(x^*)] (f_{\eta, \sigma^2}^{EK}(x) - g(x))^3 \\
& + 2f_{\eta, \sigma^2}^{EK}(x^*) \text{Var}[f(x)] (f_{\eta, \sigma^2}^{EK}(x) - g(x))^2 + f_{\eta, \sigma^2}^{EK}(x^*) (f_{\eta, \sigma^2}^{EK}(x) - g(x))^4
\end{aligned}$$

$$\begin{aligned}
& \left\langle (f(x) - g(x))^2 \right\rangle_0^2 \langle f(x^*) \rangle_0 = \\
& = f_{\eta, \sigma^2}^{EK}(x^*) \text{Var}[f(x)]^2 \\
& + 2f_{\eta, \sigma^2}^{EK}(x^*) \text{Var}[f(x)] \cdot (f_{\eta, \sigma^2}^{EK}(x) - g(x))^2 + f_{\eta, \sigma^2}^{EK}(x^*) (f_{\eta, \sigma^2}^{EK}(x) - g(x))^4
\end{aligned}$$

Summing everything up:

$$\begin{aligned}
& \left\langle (f(x) - g(x))^4 f(x^*) \right\rangle_0 - \left\langle (f(x) - g(x))^4 \right\rangle_0 \langle f(x^*) \rangle_0 \\
& - 2 \left\langle (f(x) - g(x))^2 \right\rangle_0 \left\langle (f(x) - g(x))^2 f(x^*) \right\rangle_0 + 2 \left\langle (f(x) - g(x))^2 \right\rangle_0^2 \langle f(x^*) \rangle_0 = \\
& = 8(f_{\eta, \sigma^2}^{EK}(x) - g(x)) \text{Var}[f(x)] \text{Cov}[f(x), f(x^*)]
\end{aligned}$$

and all the bubble diagrams cancel as expected.

So we get:

$$f_{\eta, \sigma^2}^{GC}(x^*) = f_{\eta, \sigma^2}^{EK}(x^*) + \frac{\eta}{\sigma^4} \int d\mu_x (f_{\eta, \sigma^2}^{EK}(x) - g(x)) \text{Var}[f(x)] \text{Cov}[f(x), f(x^*)] + O\left(\frac{1}{\eta^3}\right)$$

Substituting the expressions for the variance and covariance:

$$\begin{aligned} & f_{\eta, \sigma^2}^{GC}(x^*) = \\ & = f_{\eta, \sigma^2}^{EK}(x^*) - \frac{\eta}{\sigma^4} \sum_{i,j,k} \frac{\frac{\sigma^2}{\eta}}{\lambda_i + \frac{\sigma^2}{\eta}} \left(\frac{1}{\lambda_j} + \frac{\eta}{\sigma^2}\right)^{-1} \left(\frac{1}{\lambda_k} + \frac{\eta}{\sigma^2}\right)^{-1} g_i \phi_j(x^*) \int d\mu_x \phi_i(x) \phi_j(x) \phi_k^2(x) + O\left(\frac{1}{\eta^3}\right) \end{aligned}$$

G.5.2 AVERAGING f^2

This time we must use two different replica indices:

$$\begin{aligned} & \left\langle [f_{D_N, \sigma^2}^*(x^*)]^2 \right\rangle_{D_N \sim \mu_x^N} = \left\langle \left[\frac{\partial \log(Z_{D_N, \sigma^2}[\alpha(x)])}{\partial \alpha(x^*)} \Big|_{\alpha(x)=0} \right]^2 \right\rangle_{D_N \sim \mu_x^N} = \\ & = \left\langle \left(\lim_{M \rightarrow 0} \frac{1}{M} \cdot \frac{\partial Z_{D_N, \sigma^2}^M[\alpha(x)]}{\partial \alpha(x^*)} \Big|_{\alpha(x)=0} \right)^2 \right\rangle_{D_N \sim \mu_x^N} = \\ & = \lim_{M \rightarrow 0} \lim_{\tilde{M} \rightarrow 0} \frac{1}{M \tilde{M}} \cdot \int Df_1 \dots \int Df_M \int D\tilde{f}_1 \dots \int D\tilde{f}_{\tilde{M}} \exp\left(-\frac{1}{2} \sum_{m=1}^M \|f_m\|_{\mathcal{H}_K}^2 - \frac{1}{2} \sum_{\tilde{m}=1}^{\tilde{M}} \|\tilde{f}_{\tilde{m}}\|_{\mathcal{H}_K}^2\right) \\ & \sum_{m=1}^M f_m(x^*) \sum_{\tilde{m}=1}^{\tilde{M}} \tilde{f}_{\tilde{m}}(x^*) \left\langle \exp\left(-\sum_{m=1}^M \frac{(f_m(x) - g(x))^2}{2\sigma^2} - \sum_{\tilde{m}=1}^{\tilde{M}} \frac{(\tilde{f}_{\tilde{m}}(x) - g(x))^2}{2\sigma^2}\right) \right\rangle_{x \sim \mu_x} \end{aligned}$$

Averaging w.r.t poisson distribution:

$$\begin{aligned}
& \left\langle \left\langle [f_{D_N, \sigma^2}^*(x^*)]^2 \right\rangle_{D_N \sim \mu_x^N} \right\rangle_{N \sim \text{Poi}(\eta)} = \sum_{N=0}^{\infty} \frac{e^{-\eta} \eta^N}{N!} \left\langle [f_{D_N, \sigma^2}^*(x^*)]^2 \right\rangle_{D_N \sim \mu_x^N} = \\
& = \lim_{M \rightarrow 0} \lim_{\tilde{M} \rightarrow 0} \frac{1}{M\tilde{M}} \cdot \int Df_1 \dots \int Df_M \int D\tilde{f}_1 \dots \int D\tilde{f}_{\tilde{M}} \\
& \quad \exp \left(-\eta - \frac{1}{2} \sum_{m=1}^M \|f_m\|_{\mathcal{H}_K}^2 - \frac{1}{2} \sum_{\tilde{m}=1}^{\tilde{M}} \|\tilde{f}_{\tilde{m}}\|_{\mathcal{H}_K}^2 \right) \\
& \exp \left(\eta \left\langle \exp \left(-\sum_{m=1}^M \frac{(f_m(x) - g(x))^2}{2\sigma^2} - \sum_{\tilde{m}=1}^{\tilde{M}} \frac{(\tilde{f}_{\tilde{m}}(x) - g(x))^2}{2\sigma^2} \right) \right\rangle_{x \sim \mu_x} \right) \sum_{m=1}^M f_m(x^*) \sum_{\tilde{m}=1}^{\tilde{M}} \tilde{f}_{\tilde{m}}(x^*) \approx \\
& \approx \lim_{M \rightarrow 0} \lim_{\tilde{M} \rightarrow 0} \frac{1}{M\tilde{M}} \cdot \int Df_1 \dots \int Df_M \int D\tilde{f}_1 \dots \int D\tilde{f}_{\tilde{M}} \\
& \exp \left(-\frac{1}{2} \sum_{m=1}^M \|f_m\|_{\mathcal{H}_K}^2 - \frac{1}{2} \sum_{\tilde{m}=1}^{\tilde{M}} \|\tilde{f}_{\tilde{m}}\|_{\mathcal{H}_K}^2 + \eta \left\langle \left(-\sum_{m=1}^M \frac{(f_m(x) - g(x))^2}{2\sigma^2} - \sum_{\tilde{m}=1}^{\tilde{M}} \frac{(\tilde{f}_{\tilde{m}}(x) - g(x))^2}{2\sigma^2} \right) \right\rangle_{x \sim \mu_x} \right) \\
& + \frac{\eta}{2} \left\langle \left(-\sum_{m=1}^M \frac{(f_m(x) - g(x))^2}{2\sigma^2} - \sum_{\tilde{m}=1}^{\tilde{M}} \frac{(\tilde{f}_{\tilde{m}}(x) - g(x))^2}{2\sigma^2} \right)^2 \right\rangle_{x \sim \mu_x} \sum_{m=1}^M f_m(x^*) \sum_{\tilde{m}=1}^{\tilde{M}} \tilde{f}_{\tilde{m}}(x^*) = \\
& \approx \lim_{M \rightarrow 0} \lim_{\tilde{M} \rightarrow 0} \frac{1}{M\tilde{M}} \cdot \int Df_1 \dots \int Df_M \int D\tilde{f}_1 \dots \int D\tilde{f}_{\tilde{M}} \\
& \exp \left(-\frac{1}{2} \sum_{m=1}^M \|f_m\|_{\mathcal{H}_K}^2 - \frac{1}{2} \sum_{\tilde{m}=1}^{\tilde{M}} \|\tilde{f}_{\tilde{m}}\|_{\mathcal{H}_K}^2 + \eta \left\langle \left(-\sum_{m=1}^M \frac{(f_m(x) - g(x))^2}{2\sigma^2} - \sum_{\tilde{m}=1}^{\tilde{M}} \frac{(\tilde{f}_{\tilde{m}}(x) - g(x))^2}{2\sigma^2} \right) \right\rangle_{x \sim \mu_x} \right) \\
& \left(1 + \frac{\eta}{2} \left\langle \left(\sum_{m=1}^M \frac{(f_m(x) - g(x))^2}{2\sigma^2} + \sum_{\tilde{m}=1}^{\tilde{M}} \frac{(\tilde{f}_{\tilde{m}}(x) - g(x))^2}{2\sigma^2} \right)^2 \right\rangle_{x \sim \mu_x} \right) \sum_{m=1}^M f_m(x^*) \sum_{\tilde{m}=1}^{\tilde{M}} \tilde{f}_{\tilde{m}}(x^*) = \\
& = (f_{\eta, \sigma^2}^{EK}(x^*))^2 \\
& + \lim_{M \rightarrow 0} \lim_{\tilde{M} \rightarrow 0} \frac{1}{M\tilde{M}} \cdot \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}^{\tilde{M}} (\tilde{f}_b(x) - g(x))^2 \right)^2 \sum_{c=1}^M f_c(x^*) \sum_{d=1}^{\tilde{M}} \tilde{f}_d(x^*) \right\rangle_0 = \\
& = (f_{\eta, \sigma^2}^{EK}(x^*))^2 \\
& + \lim_{M \rightarrow 0} \lim_{\tilde{M} \rightarrow 0} \frac{1}{M\tilde{M}} \cdot \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}^{\tilde{M}} (\tilde{f}_b(x) - g(x))^2 \sum_{c=1}^M f_c(x^*) \sum_{d=1}^{\tilde{M}} \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}^{\tilde{M}} (\tilde{f}_b(x) - g(x))^2 \sum_{c=1}^M f_c(x^*) \sum_{d=1}^{\tilde{M}} \tilde{f}_d(x^*) \right\rangle_0 \right] = \\
& = (f_{\eta, \sigma^2}^{EK}(x^*))^2 + \frac{\eta}{4\sigma^4} \int d\mu_x \underbrace{\lim_{M \rightarrow 0} \frac{1}{M} \left\langle \sum_{a=1}^M (f_a(x) - g(x))^2 \sum_{b=1}^{\tilde{M}} (\tilde{f}_b(x) - g(x))^2 \sum_{c=1}^M f_c(x^*) \right\rangle_0}_{8(f_{\eta, \sigma^2}^{EK}(x) - g(x)) \text{Var}[f(x)] \text{Cov}[f(x), f(x^*)] \text{ as we saw in } \langle f \rangle} f_{\eta, \sigma^2}^{EK}(x^*) \\
& + \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}^{\tilde{M}} \tilde{f}_b(x^*) \right\rangle_0 \right)^2
\end{aligned}$$

and we're left with:

$$\begin{aligned} & \lim_{M \rightarrow 0} \frac{1}{M} \left\langle \sum_{a=1}^M (f_a(x) - g(x))^2 \sum_{b=1}^M f_b(x^*) \right\rangle_0 = \\ & \lim_{M \rightarrow 0} \frac{1}{M} \left[M \left\langle (f(x) - g(x))^2 f(x^*) \right\rangle_0 + M(M-1) \langle f(x^*) \rangle \left\langle (f(x) - g(x))^2 \right\rangle_0 \right] = \\ & = \left\langle (f(x) - g(x))^2 f(x^*) \right\rangle_0 - \langle f(x^*) \rangle_0 \left\langle (f(x) - g(x))^2 \right\rangle_0 = 2 (f_{\eta, \sigma^2}^{EK}(x) - g(x)) \text{Cov}[f(x), f(x^*)] \end{aligned}$$

but this correction gives $O\left(\frac{1}{\eta^3}\right)$ so:

$$\begin{aligned} & \left\langle \left\langle [f_{D_N, \sigma^2}^*(x^*)]^2 \right\rangle_{D_N \sim \mu_x^N} \right\rangle_{N \sim \text{Poi}(\eta)} = \\ & (f_{\eta, \sigma^2}^{EK}(x^*))^2 + \frac{2\eta}{\sigma^4} f_{\eta, \sigma^2}^{EK}(x^*) \int d\mu_x (f_{\eta, \sigma^2}^{EK}(x) - g(x)) \text{Var}[f(x)] \text{Cov}[f(x), f(x^*)] + O\left(\frac{1}{\eta^3}\right) \end{aligned}$$

and notably

$$\langle f^2 \rangle = \langle f \rangle^2 + O\left(\frac{1}{\eta^3}\right)$$

G.6 PERUBATIVE CORRECTION FOR ROTATIONALLY INVARIANT KERNEL

We now wish to evaluate this expression for a rotationally invariant kernel and a uniform measure on the hypersphere. This simplifies the expression for $\langle f \rangle$ to:

$$\begin{aligned} & f_{\eta, \sigma^2}^{GC}(x^*) = \\ & = f_{\eta, \sigma^2}^{EK}(x^*) + \frac{\eta}{\sigma^4} \int d\mu_x (f_{\eta, \sigma^2}^{EK}(x) - g(x)) \text{Var}[f(x)] \text{Cov}[f(x), f(x^*)] + O\left(\frac{1}{\eta^3}\right) = \\ & = f_{\eta, \sigma^2}^{EK}(x^*) + \frac{\eta}{\sigma^4} C_{K, \eta, \sigma^2} \int d\mu_x (f_{\eta, \sigma^2}^{EK}(x) - g(x)) \text{Cov}[f(x), f(x^*)] = \\ & = f_{\eta, \sigma^2}^{EK}(x^*) - \frac{\eta}{\sigma^4} C_{K, \eta, \sigma^2} \sum_{i,j} \frac{\frac{\sigma^2}{\eta}}{\lambda_i + \frac{\sigma^2}{\eta}} \left(\frac{1}{\lambda_j} + \frac{\eta}{\sigma^2}\right)^{-1} g_i \phi_j(x^*) \underbrace{\int d\mu_x \phi_i(x) \phi_j(x)}_{\delta_{ij}} + O\left(\frac{1}{\eta^3}\right) = \\ & = f_{\eta, \sigma^2}^{EK}(x^*) - \frac{\eta}{\sigma^4} C_{K, \eta, \sigma^2} \sum_i \frac{\frac{\sigma^2}{\eta}}{\lambda_i + \frac{\sigma^2}{\eta}} \left(\frac{1}{\lambda_i} + \frac{\eta}{\sigma^2}\right)^{-1} g_i \phi_i(x^*) + O\left(\frac{1}{\eta^3}\right) = \\ & = f_{\eta, \sigma^2}^{EK}(x^*) - C_{K, \eta, \sigma^2} \sum_{l,m} \frac{g_{l,m}}{\sigma^2 \left(2 + \frac{\lambda_l \eta}{\sigma^2} + \frac{\sigma^2}{\lambda_l \eta}\right)} Y_{l,m}(x^*) + O\left(\frac{1}{\eta^3}\right). \end{aligned}$$

The expression for $\langle f^2 \rangle$ is:

$$\begin{aligned} & \left\langle \left\langle [f_{D_N, \sigma^2}^*(x^*)]^2 \right\rangle_{D_N \sim \mu_x^N} \right\rangle_{N \sim \text{Poi}(\eta)} = \\ & = (f_{\eta, \sigma^2}^{EK}(x^*))^2 - 2f_{\eta, \sigma^2}^{EK}(x^*) C_{K, \eta, \sigma^2} \sum_{l,m} \frac{g_{l,m}}{\sigma^2 \left(2 + \frac{\lambda_l \eta}{\sigma^2} + \frac{\sigma^2}{\lambda_l \eta}\right)} Y_{l,m}(x^*) + O\left(\frac{1}{\eta^3}\right) \end{aligned}$$

H VARIOUS INSIGHTS

H.1 CORRECTION MEANS WORSE GENERALIZATION

The correction always means worse generalization than what the EK suggests. Indeed

$$\begin{aligned}
f_{\eta, \sigma^2}^{GC}(x^*) &= f_{\eta, \sigma^2}^{EK}(x^*) - C_{K, \eta, \sigma^2} \sum_{l, m} \frac{g_{l, m}}{\sigma^2 \left(2 + \frac{\lambda_l \eta}{\sigma^2} + \frac{\sigma^2}{\lambda_l \eta}\right)} Y_{l, m}(x^*) + O\left(\frac{1}{\eta^3}\right) = \\
&= \sum_{l, m} \frac{\lambda_l}{\lambda_l + \frac{\sigma^2}{\eta}} g_{l, m} Y_{l, m}(x^*) - C_{K, \eta, \sigma^2} \sum_{l, m} \frac{g_{l, m}}{\sigma^2 \left(2 + \frac{\lambda_l \eta}{\sigma^2} + \frac{\sigma^2}{\lambda_l \eta}\right)} Y_{l, m}(x^*) + O\left(\frac{1}{\eta^3}\right) = \\
&= \sum_{l, m} \underbrace{\left(\frac{\lambda_l}{\lambda_l + \frac{\sigma^2}{\eta}} - \frac{C_{K, \eta, \sigma^2}}{\sigma^2 \left(2 + \frac{\lambda_l \eta}{\sigma^2} + \frac{\sigma^2}{\lambda_l \eta}\right)} \right)}_{\substack{\text{positive} \\ < \frac{\lambda_l}{\lambda_l + \frac{\sigma^2}{\eta}} < 1}} g_{l, m} Y_{l, m}(x^*)
\end{aligned}$$

H.2 EXACT EIGENVALUES FOR 2-LAYER RELU NTK WITH $\sigma_b^2 = 0$

For the NTK associated with a 2-layer ReLU NTK without bias we were able to find an exact expression for the eigenvalues for all l :

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

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.