
Wide Neural Networks of Any Depth Evolve as Linear Models Under Gradient Descent

Jaehoon Lee^{*12} Lechao Xiao^{*12} Samuel S. Schoenholz¹ Yasaman Bahri¹
Jascha Sohl-Dickstein¹ Jeffrey Pennington¹

Abstract

A longstanding goal in deep learning research has been to precisely characterize training and generalization. However, the often complex loss landscapes of neural networks have made a theory of learning dynamics elusive. In this work, we show that for wide neural networks the learning dynamics simplify considerably and that, in the infinite width limit, they are governed by a linear model obtained from the first-order Taylor expansion of the network around its initial parameters. Furthermore, mirroring the correspondence between wide Bayesian neural networks and Gaussian processes, gradient-based training of wide neural networks with a squared loss produces test set predictions drawn from a Gaussian process with a particular compositional kernel. While these theoretical results are only exact in the infinite width limit, we nevertheless find excellent empirical agreement between the predictions of the original network and those of the linearized version even for finite practically-sized networks. This agreement is robust across different architectures, optimization methods, and loss functions.

1. Introduction

Machine learning models based on deep neural networks have achieved unprecedented performance across a wide range of tasks (Krizhevsky et al., 2012; He et al., 2016; Devlin et al., 2018). Typically, these models are regarded as complex systems for which many types of theoretical analyses are intractable. Moreover, characterizing the gradient-based training dynamics of these models is challenging owing to the typically high-dimensional non-convex loss surfaces governing the optimization. As is common in the physical sciences, investigating the extreme limits of such

systems can often shed light on these hard problems. For neural networks, one such limit is that of infinite width, which refers either to the number of hidden units in a fully-connected layer or to the number of channels in a convolutional layer. Under this limit, the output of the network at initialization is a draw from a Gaussian process (GP); moreover, the network output remains governed by a GP after exact Bayesian training using squared loss (Neal, 1994; Lee et al., 2018; Matthews et al., 2018; Novak et al., 2019; Garriga-Alonso et al., 2018). Aside from its theoretical simplicity, the infinite-width limit is also of practical interest as wider networks have been found to generalize better (Neyshabur et al., 2015; Novak et al., 2018; Lee et al., 2018; Novak et al., 2019; Neyshabur et al., 2019).

In this work, we explore the learning dynamics of wide neural networks under gradient descent and find that the weight-space description of the dynamics becomes surprisingly simple: as the width becomes large, the neural network can be effectively replaced by its first-order Taylor expansion with respect to its parameters at initialization. For this induced linear model, the dynamics of gradient descent become *analytically tractable*. While the linearization is only exact in the infinite width limit, we nevertheless find excellent agreement between the predictions of the original network and those of the linearized version even for finite width configurations. The agreement persists across different architectures, optimization methods, and loss functions.

For squared loss, the exact learning dynamics admit a closed-form solution that allows us to characterize the evolution of the predictive distribution in terms of a GP. This result can be thought of as extension of “sample-then-optimize” posterior sampling (Matthews et al., 2017) to the training of deep neural networks. Our empirical simulations confirm that the result accurately models the variation in predictions across an ensemble of finite-width models with different random initializations.

1.1. Our Contribution

We begin by building on a recent result by Jacot et al. (2018) that characterizes the exact dynamics of network outputs throughout gradient descent training in the infinite width

^{*}Equal contribution ¹Google Brain ²Work done as part of the Google AI Residency. Correspondence to: Jaehoon Lee <jahlee@google.com>, Lechao Xiao <xlc@google.com>.

limit. Their results establish that gradient descent in parameter space corresponds to kernel gradient descent in function space with respect to a new kernel, the Neural Tangent Kernel (NTK). One may ask what this tells us about the nature of the dynamics in parameter space, where training updates are actually made. A key contribution of our work is to show that dynamics in parameter space are equivalent to the training dynamics of a model which is affine in the collection of all network parameters, the weights and biases. This result holds regardless of the choice of loss function. In the case of squared loss, the dynamics admit a closed-form solution as a function of time.

The output of an infinitely wide neural network is Gaussian at initialization (Lee et al., 2018; Matthews et al., 2018), and as mentioned in Jacot et al. (2018), for squared loss it remains Gaussian throughout training¹. We derive explicit time-dependent expressions for the mean and covariance functions of this GP, and provide a novel interpretation of the result. In particular, it offers a quantitative understanding of the mechanism by which gradient descent differs from Bayesian posterior sampling of the parameters: while both methods generate draws from a GP, gradient descent does not generate samples from the posterior of any probabilistic model. This observation is in contrast to the “sample-then-optimize” framework of (Matthews et al., 2017) in which only the top-layer weights are trained and gradient descent *does* sample from the Bayesian posterior. These observations establish a framework with which to analyze the long-standing questions as to if, how, and in what contexts gradient descent provides concrete benefits relative to Bayesian inference.

As argued by Chizat & Bach (2018b), these theoretical results may appear too simple to be applicable to realistic neural networks. Nonetheless, we empirically investigate the applicability of the theory in the finite-width setting and find that it gives an accurate characterization of both the learning dynamics and posterior function distributions across a variety of conditions, including some practical network architectures such as the wide residual network (Zagoruyko & Komodakis, 2016).

1.2. Additional related work

Daniely et al. (2016) study the relationship between neural networks and kernels at initialization. They bound the difference between the infinite width kernel and the empirical kernel at finite width n , which diminishes as $\mathcal{O}(1/\sqrt{n})$. Daniely (2017) uses the same kernel perspective to study stochastic gradient descent (SGD) training of neural networks.

Saxe et al. (2013) study the training dynamics of deep linear

networks, in which the nonlinearities are treated as identity functions. Deep linear networks are linear in their inputs, but not their parameters. In contrast, we show that the outputs of sufficiently wide neural networks are linear in their parameters but not usually their inputs.

Du et al. (2018); Allen-Zhu et al. (2018a;b); Zou et al. (2018) study the convergence of gradient descent to global minima. They proved that for i.i.d. Gaussian initialization, the parameters of sufficiently wide networks move little from their initial values during SGD. This small motion of the parameters is crucial to the effect we present, where wide neural networks behave linearly in terms of their parameters throughout training.

Mei et al. (2018); Chizat & Bach (2018a); Rotskoff & Vanden-Eijnden (2018); Sirignano & Spiliopoulos (2018) analyze the mean field SGD dynamics of training neural networks in the large-width limit. Their mean field analysis describes distributional dynamics of network parameters via a PDE. However, their analysis is restricted to one hidden layer networks with a scaling limit $(1/n)$ different from ours $(1/\sqrt{n})$, which is commonly used in modern networks He et al. (2016); Glorot & Bengio (2010).

Finally, Zhang et al. (2019) observed that some of the layers in trained neural networks are robust to re-initialization, but not to re-randomization. Our framework provides some theoretical support for this empirical finding.

2. Theoretical Results

2.1. Notation and setup

Let $\mathcal{D} \subseteq \mathbb{R}^{n_0} \times \mathbb{R}^k$ denote the training set and $\mathcal{X} = \{x : (x, y) \in \mathcal{D}\}$ and $\mathcal{Y} = \{y : (x, y) \in \mathcal{D}\}$ denote the inputs and labels, respectively. Consider a fully-connected feed-forward network with L hidden layers with width n_l , for $l = 1, \dots, L$ and a readout layer with $n_{L+1} = k$. For each $x \in \mathbb{R}^{n_0}$, we use $h^l(x), x^l(x) \in \mathbb{R}^{n_l}$ to represent the pre- and post-activation functions at layer l with input x . The recurrence relation for a feed-forward network is defined as

$$\begin{cases} h^{l+1} &= x^l W^{l+1} + b^{l+1} \\ x^{l+1} &= \phi(h^{l+1}) \end{cases} \quad \text{and} \quad \begin{cases} W_{i,j}^l &= \frac{\sigma_\omega}{\sqrt{n_l}} \omega_{i,j}^l \\ b_j^l &= \sigma_b \beta_j^l \end{cases}, \quad (1)$$

where ϕ is a point-wise activation function, $W^{l+1} \in \mathbb{R}^{n_l \times n_{l+1}}$ and $b^{l+1} \in \mathbb{R}^{n_{l+1}}$ are the weights and biases, $\omega_{i,j}^l$ and β_j^l are the trainable variables, drawn i.i.d. from standard Gaussian $\omega_{i,j}^l, \beta_j^l \sim \mathcal{N}(0, 1)$ at initialization, and σ_ω^2 and σ_b^2 are weight and bias variances. Note that this parametrization method is non-standard (see SM E for further analysis), and we will refer to it as the NTK parameterization. It has already been adopted in several recent works (van Laarhoven, 2017; Karras et al., 2018; Jacot et al., 2018; Du et al., 2018;

¹The setting is full-batch training under gradient flow.

Park et al., 2018). Unlike the standard parameterization that only normalizes the forward dynamics of the network, the NTK-parameterization also normalizes its backward dynamics. We note that the predictions and training dynamics of NTK-parameterized networks are identical to those of standard networks, up to a width-dependent scaling factor in the learning rate for each parameter tensor. We compare the test performance of these two parameterization methods in Figure S1 of SM. Our results (linearity in weights, GP predictions) also hold for infinitely wide networks with a standard parameterization (see SM Section E Figure S2).

We set $\theta^l \equiv [w_{i,j}^l, \beta_j^l]_{i,j}$, the collection of parameters mapping to the l -th layer, whose cardinality is $(n_{l-1} + 1) \times n_l$. Define $\theta = \text{vec}(\bigcup_{l=1}^{L+1} \theta^l)$ and similarly $\theta^{\leq l}$ or $\theta^{> l}$. Denote by θ_t the time-dependence of the parameters and by θ_0 their initial values. We use $f_t(x) \equiv h^{L+1}(x) \in \mathbb{R}^k$ to denote the output (or logits) of the neural network at time t . Let $\ell(\hat{y}, y) : \mathbb{R}^k \times \mathbb{R}^k \rightarrow \mathbb{R}$ denote the loss function where the first argument is the prediction and the second argument the true label. In supervised learning, one is interested in learning a θ that minimizes the empirical loss²,

$$\mathcal{L} = \sum_{(x,y) \in \mathcal{D}} \ell(f_t(x, \theta), y). \quad (2)$$

Let η be the learning rate³. Via continuous time gradient descent, the evolution of the parameters θ and the logits f can be written as

$$\dot{\theta}_t = -\eta \nabla_{\theta} f_t(\mathcal{X})^T \nabla_{f_t(\mathcal{X})} \mathcal{L} \quad (3)$$

$$\dot{f}_t(\mathcal{X}) = \nabla_{\theta} f_t(\mathcal{X}) \dot{\theta}_t = -\eta \hat{\Theta}_t(\mathcal{X}, \mathcal{X}) \nabla_{f_t(\mathcal{X})} \mathcal{L} \quad (4)$$

where $f_t(\mathcal{X}) = \text{vec}([f_t(x)]_{x \in \mathcal{X}})$, the $k|\mathcal{D}| \times 1$ vector of concatenated logits for all examples, and $\nabla_{f_t(\mathcal{X})} \mathcal{L}$ is the gradient of the loss with respect to the model's output, $f_t(\mathcal{X})$. $\hat{\Theta}_t \equiv \hat{\Theta}_t(\mathcal{X}, \mathcal{X})$ is the tangent kernel at time t , which is a $k|\mathcal{D}| \times k|\mathcal{D}|$ matrix

$$\hat{\Theta}_t = \nabla_{\theta} f_t(\mathcal{X}) \nabla_{\theta} f_t(\mathcal{X})^T = \sum_{l=1}^{L+1} \nabla_{\theta^l} f_t(\mathcal{X}) \nabla_{\theta^l} f_t(\mathcal{X})^T. \quad (5)$$

One can define the tangent kernel for general arguments, e.g. $\hat{\Theta}_t(x, \mathcal{X})$ where x is test input. We refer to $\hat{\Theta}$ as the *empirical* tangent kernel.

2.2. Linearized networks

In this section, we consider the training dynamics of the linearized network. Specifically, we replace the outputs of

²To simplify the notation for later equations, we use the *total* loss here instead of the *average* loss, but for all plots in Section 3, we show the *average* loss.

³Note that compared to the conventional parameterization, η is larger by factor of width. The NTK parameterization allows usage of a universal learning rate scale irrespective of network width.

the neural network by their first order Taylor expansion,

$$f_t^{\text{lin}}(x) \equiv f_0(x) + \nabla_{\theta} f_0(x) \omega_t, \quad (6)$$

where $\omega_t \equiv \theta_t - \theta_0$ is the change in the parameters from their initial values⁴. Note that f_t^{lin} is the sum of two terms: the first term is the initial output of the network, which remains unchanged during training, and the second term captures the change to the initial value during training. The dynamics of gradient flow using this linearized function are governed by,

$$\dot{\omega}_t = -\eta \nabla_{\theta} f_0(\mathcal{X})^T \nabla_{f_t^{\text{lin}}(\mathcal{X})} \mathcal{L} \quad (7)$$

$$f_t^{\text{lin}}(x) = -\eta \hat{\Theta}_0(x, \mathcal{X}) \nabla_{f_t^{\text{lin}}(\mathcal{X})} \mathcal{L}. \quad (8)$$

As $\nabla_{\theta} f_0(x)$ remains constant throughout training, these dynamics are often quite simple. In the case of an MSE loss, i.e., $\ell(\hat{y}, y) = \frac{1}{2} \|\hat{y} - y\|_2^2$, the ODEs have closed form solutions

$$\omega_t = -\nabla_{\theta} f_0(\mathcal{X})^T \hat{\Theta}_0^{-1} (I - e^{-\eta \hat{\Theta}_0 t}) (f_0(\mathcal{X}) - \mathcal{Y}), \quad (9)$$

$$f_t^{\text{lin}}(\mathcal{X}) = (I - e^{-\eta \hat{\Theta}_0 t}) \mathcal{Y} + e^{-\eta \hat{\Theta}_0 t} f_0(\mathcal{X}). \quad (10)$$

For an arbitrary point x , $f_t^{\text{lin}}(x) = \mu_t(x) + \gamma_t(x)$, where

$$\mu_t(x) = \hat{\Theta}_0(x, \mathcal{X}) \hat{\Theta}_0^{-1} (I - e^{-\eta \hat{\Theta}_0 t}) \mathcal{Y} \quad (11)$$

$$\gamma_t(x) = f_0(x) - \hat{\Theta}_0(x, \mathcal{X}) \hat{\Theta}_0^{-1} (I - e^{-\eta \hat{\Theta}_0 t}) f_0(\mathcal{X}). \quad (12)$$

Therefore, we can obtain the time evolution of the linearized neural network without training it. We only need to compute the tangent kernel $\hat{\Theta}_0$ and the outputs f_0 at initialization and use Equations 11, 12, and 9 to compute the dynamics of the outputs and the weights.

2.3. Infinite width limit yields Gaussian processes

As the width of the hidden layers approaches infinity, the Central Limit Theorem (CLT) implies that the outputs at initialization $\{f_0(x)\}_{x \in \mathcal{X}}$ converge to a multivariate Gaussian in distribution. Informally, this can be proved by induction. Conditioning on activations at layer l , each pre-activation ($h_j^{l+1}(x)$) at layer $l+1$ is a sum of n^l i.i.d. properly normalized random variables (the weights) and a Gaussian distribution (the bias term). One can apply the CLT to conclude that $\{h_j^{l+1}(x)\}_j$ are i.i.d. Gaussian; see Poole et al. (2016); Schoenholz et al. (2017); Lee et al. (2018); Xiao et al. (2018); Yang & Schoenholz (2017) for more details, and Matthews et al. (2018); Novak et al. (2019) for a formal treatment.

Therefore, randomly initialized neural networks are in correspondence with a certain class of Gaussian processes

⁴Since $f_0^{\text{lin}}(x) = f_0(x)$, we will often drop the superscript at time 0.

(hereinafter referred to as NNGPs), which facilitates a fully Bayesian treatment of neural networks (Lee et al., 2018; Matthews et al., 2018). More precisely, let f_t^i denote the i -th output dimension and \mathcal{K} denote the sample-to-sample kernel function (of the pre-activation) of the outputs in the infinite width setting,

$$\mathcal{K}^{i,j}(x, x') = \lim_{\min(n_1, \dots, n_L) \rightarrow \infty} \mathbb{E}[f_0^i(x) \cdot f_0^j(x')], \quad (13)$$

then $f_0(\mathcal{X}) \sim \mathcal{GP}(0, \mathcal{K}(\mathcal{X}, \mathcal{X}))$, where $\mathcal{K}^{i,j}(x, x')$ denotes the covariance between the i -th output of x and j -th output of x' , which can be computed recursively (see Lee et al. (2018) Section 2.3 and Supplementary Material (SM) Section A). For an unobserved test input x , the joint output distribution $f([x, \mathcal{X}])$ is also a Gaussian process. Conditioning on the training samples⁵ $f(\mathcal{X}) = \mathcal{Y}$, the posterior predictive distribution of $f(x)$ is also Gaussian $\mathcal{N}(\mu(x), \Sigma(x))$

$$\begin{aligned} \mu(x) &= \mathcal{K}(x, \mathcal{X})\mathcal{K}^{-1}\mathcal{Y} \\ \Sigma(x) &= \mathcal{K}(x, x) - \mathcal{K}(x, \mathcal{X})\mathcal{K}^{-1}\mathcal{K}(x, \mathcal{X})^T, \end{aligned} \quad (14)$$

where $\mathcal{K} = \mathcal{K}(\mathcal{X}, \mathcal{X})$. This is the posterior predictive distribution resulting from exact Bayesian inference in an infinitely wide neural network.

2.3.1. GAUSSIAN PROCESSES FROM GRADIENT DESCENT TRAINING

If we freeze the variables $\theta^{\leq L}$ after initialization and only optimize θ^{L+1} , the original network and its linearization are identical. Letting the width approach infinity, this particular tangent kernel $\hat{\Theta}_0$ will converge to \mathcal{K} in probability and Equation 11 will converge to the posterior Equation 14 as $t \rightarrow \infty$ (for further details see SM Section D). This is a realization of the ‘‘sample-then-optimize’’ approach for evaluating the posterior of a Gaussian process proposed in Matthews et al. (2017).

If none of the variables are frozen, in the infinite width setting, $\hat{\Theta}_0$ also converges in probability to a deterministic kernel Θ (Jacot et al., 2018), which we sometimes refer to as the analytic kernel, and which can also be computed recursively (see SM Section A). Letting the width go to infinity, for any t , the output $f_t^{\text{lin}}(x)$ is also Gaussian distributed because Equations 11 and 12 describe an affine transform of the Gaussian $[f_0(x), f_0(\mathcal{X})]$. The mean($\mu(x)$) and the variance($\Sigma(x)$) are given by

$$\mu(x) = \Theta(x, \mathcal{X})\Theta^{-1}(I - e^{-\eta\Theta t})\mathcal{Y} \quad (15)$$

$$\begin{aligned} \Sigma(x) &= \mathcal{K}(x, x) - 2\Theta(x, \mathcal{X})\Theta^{-1}(I - e^{-\eta\Theta t})\mathcal{K}(x, \mathcal{X})^T \\ &\quad + \Theta(x, \mathcal{X})\Theta^{-1}(I - e^{-\eta\Theta t})\mathcal{K}\Theta^{-1}(I - e^{-\eta\Theta t})\Theta(x, \mathcal{X})^T. \end{aligned} \quad (16)$$

⁵ This imposes that h^{L+1} directly corresponds to the network predictions. In the case of softmax readout, variational or sampling methods are required to marginalize over h^{L+1} .

Unlike the case when only θ^{L+1} is optimized, Equations 15 and 16 do not admit an interpretation corresponding to the posterior sampling of a probabilistic model⁶. We contrast the predictive distributions from the NNGP, NTK-GP (i.e. Equations 15 and 16) and ensembles of NNs in Figure 3.

2.4. Infinite width networks are linearized networks

The ODEs (Equation 3, 4) of the original network are unsolvable in general, since $\hat{\Theta}_t$ evolves with time. Remarkably, Jacot et al. (2018) showed that

$$\sup_{t \in [0, T]} \|\hat{\Theta}_t - \hat{\Theta}_0\|_F = O(\min\{n_1, \dots, n_L\}^{-1/2}) \quad (17)$$

under the technical assumption that the integrated functional gradient norm remains stochastically bounded as $n_1, \dots, n_L \rightarrow \infty$ sequentially, i.e.

$$\int_0^T \|\nabla_{f_t(x)} \mathcal{L}\|_2 dt < \infty, \quad (18)$$

where T is the training budget (independent from the widths), i.e. the amount of allowable training time. This assumption was verified in (Jacot et al., 2018) in some specific cases. Recent work on the convergence theory of over-parameterized neural networks provides a rigorous proof of Equation 18 in the *simultaneous* limit (i.e. $\min\{n_1, \dots, n_L\} \rightarrow \infty$) under discrete-time stochastic gradient descent with MSE loss on various architectures (fully-connected networks, residual CNNs, RNNs) (Du et al., 2018; Allen-Zhu et al., 2018b;a; Zou et al., 2018). Note that the bound in Equation 17 is just a theoretical upper bound and we observe empirically that the exponent $-\frac{1}{2}$ may be improved to -1 ; see Figure 2.

Coupling Equation 18 with Grönwall’s type arguments, in the MSE setting, we can upper bound the discrepancy between the outputs of the original network and those of its linearization,

$$\sup_{t \in [0, T]} \|f_t(\mathcal{X}) - f_t^{\text{lin}}(\mathcal{X})\|_2 = O\left(\sup_{t \in [0, T]} \|\hat{\Theta}_t - \hat{\Theta}_0\|_F\right), \quad (19)$$

which approaches 0 as the width goes to infinity. Intuitively, the ODE of the original network (Equation 4) can be considered as a $\|\hat{\Theta}_t - \hat{\Theta}_0\|_F$ -fluctuation from the linearized ODE (Equation 8). One expects the difference between the solutions of these two ODEs to be upper bounded by some functional of $\|\hat{\Theta}_t - \hat{\Theta}_0\|_F$ (refer to SM Section F for the proof). Therefore, for a large width network, the training dynamics can be well approximated by linearized dynamics.

Note that the updates for individual weights in Equation 7 vanish in the infinite width limit, which for instance can be

⁶One possible exception is when the NNGP kernel and NTK are the same up to a scalar multiplication.

seen from the explicit width dependence of the gradients in the NTK parameterization. Individual weights move by a vanishingly small amount for wide networks in this regime of dynamics, as do hidden layer activations, but they collectively conspire to provide a finite change in the final output of the network, as is necessary for training.

An additional insight gained from linearization of the network is that the dynamics derived in (Jacot et al., 2018) are equivalent to a random features method, where the features are the gradients of the model with respect to its weights.

2.5. Extensions

The linearization of the wide neural networks and its training dynamics can be generalized in multiple directions.

2.5.1. MOMENTUM

One direction is to go beyond vanilla gradient descent dynamics. We consider momentum updates⁷

$$\theta_{k+1} = \theta_k + \beta(\theta_k - \theta_{k-1}) - \eta \nabla_{\theta} \mathcal{L}|_{\theta=\theta_k}. \quad (20)$$

The discrete update to the function output becomes

$$\begin{aligned} f_{k+1}^{\text{lin}}(x) &= f_k^{\text{lin}}(x) - \eta \hat{\Theta}_0(x, \mathcal{X}) \nabla_{f_k^{\text{lin}}(\mathcal{X})} \mathcal{L} \\ &\quad + \beta(f_k^{\text{lin}}(x) - f_{k-1}^{\text{lin}}(x)) \end{aligned} \quad (21)$$

where $f_t^{\text{lin}}(x)$ is the output of the linearized network after t steps. One can take the continuous time limit as in Qian (1999); Su et al. (2014) and obtain

$$\dot{\omega}_t = \tilde{\beta} \dot{\omega}_t - \nabla_{\theta} f_0^{\text{lin}}(\mathcal{X})^T \nabla_{f_t^{\text{lin}}(\mathcal{X})} \mathcal{L} \quad (22)$$

$$\ddot{f}_t^{\text{lin}}(x) = \tilde{\beta} \dot{f}_t^{\text{lin}}(x) - \hat{\Theta}_0(x, \mathcal{X}) \nabla_{f_t^{\text{lin}}(\mathcal{X})} \mathcal{L} \quad (23)$$

where continuous time relates to steps $t = k\sqrt{\eta}$ and $\tilde{\beta} = (\beta - 1)/\sqrt{\eta}$. These equations are also amenable to analytic treatment for MSE loss. See Figure 5, 6 and 7 for experimental agreement.

2.5.2. CROSS-ENTROPY LOSS

One can extend the loss function to general functions with multiple output dimensions. Unlike for squared error, we do not have a closed form solution to the dynamics equation. However, the equations for the dynamics can be solved using an ODE solver as an initial value problem. For cross-entropy loss with softmax output (see SM Section C),

$$\ell(f, y) = - \sum_i y_i \log \sigma(f^i), \quad \sigma(f^i) \equiv \frac{\exp(f^i)}{\sum_j \exp(f^j)}. \quad (24)$$

⁷Combining the usual two stage update into a single equation.

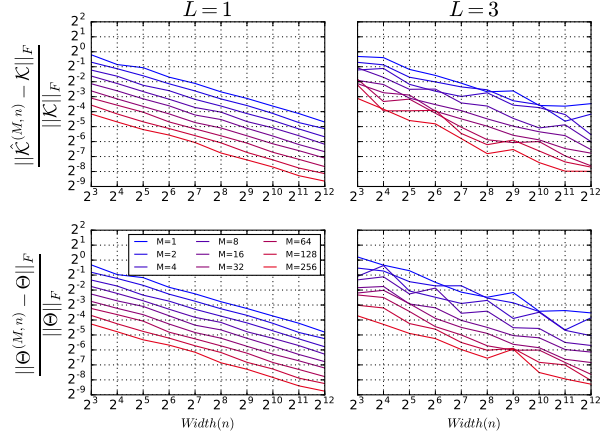


Figure 1. **Kernel convergence.** Kernels computed from randomly initialized ReLU networks with one and three hidden layers converge to the corresponding analytic kernel as width n and number of Monte Carlo samples M increases. Colors indicate averages over different numbers of Monte Carlo samples.

the dynamics equation becomes

$$\dot{f}_t^{\text{lin}}(\mathcal{X}) = -\eta \hat{\Theta}_0(\mathcal{X}, \mathcal{X}) (\sigma(f_t^{\text{lin}}(\mathcal{X})) - \mathcal{Y}) \quad (25)$$

$$\dot{f}_t^{\text{lin}}(x) = -\eta \hat{\Theta}_0(x, \mathcal{X}) (\sigma(f_t^{\text{lin}}(\mathcal{X})) - \mathcal{Y}). \quad (26)$$

2.5.3. BEYOND FULLY-CONNECTED NETWORKS

Although our theoretical analysis has focused on fully-connected architectures, there is good reason to suspect the results to extend to much broader class of models. In particular, a wealth of recent literature suggests that the mean field theory governing the wide network limit of fully-connected models (Poole et al., 2016; Schoenholz et al., 2017) extends naturally to residual networks (Yang & Schoenholz, 2017), CNNs (Xiao et al., 2018), RNNs (Chen et al., 2018), batch normalization (Yang et al., 2019), and general architectures (Yang, 2019). We postpone the development of these theoretical extensions in favor of a purely empirical investigation of linearization for a variety of architectures (see Section 3.3).

3. Experiments

In this section, we provide empirical support showing that the training dynamics of wide neural networks are well captured by linearized models. We consider fully-connected, convolutional, and wide ResNet architectures trained with full and mini- batch gradient descent using learning rates sufficiently small so that the continuous time approximation holds well. We consider two-class classification on CIFAR-10 (horses and planes) as well as ten-class classification on MNIST and CIFAR-10. When using MSE loss, we treat

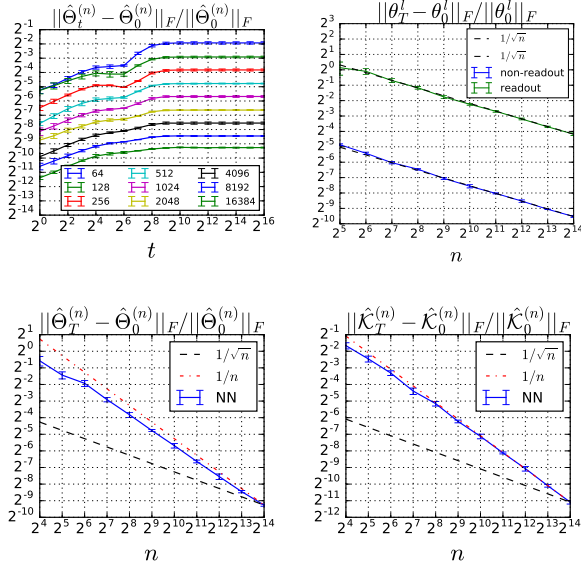


Figure 2. Relative Frobenius norm change during training. One hidden layer, ReLU networks trained with $\eta = 1.0$, $|\mathcal{D}| = 128$. We measure changes of (read-out/non read-out) weights, empirical $\hat{\Theta}$ and empirical $\hat{\mathcal{K}}$ after $T = 65, 536$ steps of gradient descent updates for varying width. We see that the change in weights scales as $\frac{1}{\sqrt{n}}$ whereas the change in $\hat{\Theta}$ and $\hat{\mathcal{K}}$ is bounded by $O(\frac{1}{\sqrt{n}})$ but close to $O(\frac{1}{n})$.

the binary classification task as regression with one class regressing to $+1$ and the other to -1 .

3.1. Convergence of empirical kernel

As in Novak et al. (2019), we can use Monte Carlo estimates of the tangent kernel (Equation 5) to probe convergence to the infinite width kernel (analytically computed using Equations S3, S6). For simplicity, we consider random inputs drawn from $\mathcal{N}(0, 1)$ with $n_0 = 1024$. In Figure 1, we observe convergence as both width n increases and the number of Monte Carlo samples M increases. For both NNGP and tangent kernels we observe $\|\hat{\Theta}^{(n)} - \Theta\|_F = O(\frac{1}{\sqrt{n}})$ and $\|\hat{\mathcal{K}}^{(n)} - \mathcal{K}\|_F = O(\frac{1}{\sqrt{n}})$ as predicted by a CLT (Daniely et al., 2016).

Moreover, as the neural network trains the change during training in the NNGP and tangent kernels, and in individual weights, becomes small as width increases, as shown in Figure 2.

3.2. Predictive output distribution

In the case of an MSE loss, the output distribution remains Gaussian throughout training. In Figure 3, the predictive output distribution for input points interpolated between two




Figure 3. Dynamics of mean and variance of trained neural network outputs follow analytic dynamics from linearization.

Black lines indicate the time evolution of the predictive output distribution from an ensemble of 100 trained neural networks (NNs). The blue region indicates the analytic prediction of the output distribution throughout training (Equations 15, 16). Finally, the red region indicates the prediction that would result from training only the top layer, corresponding to an NNGP (Equations S30, S31). The trained network has 3 hidden layers of width 8192, tanh activation functions, $\sigma_w^2 = 1.5$, no bias, and $\eta = 0.5$. The output is computed for inputs interpolated between two training points (denoted with black dots) $x(\alpha) = \alpha x^{(1)} + (1 - \alpha)x^{(2)}$. The shaded region and dotted lines denote 2 standard deviations ($\sim 95\%$ quantile) from the mean denoted in solid lines. Training was performed with full-batch gradient descent with dataset size $|\mathcal{D}| = 128$. For dynamics for individual draw of functions, see SM Figure S4

training points is shown for an ensemble of neural networks and their corresponding GPs. The interpolation is given by $x(\alpha) = \alpha x^{(1)} + (1 - \alpha)x^{(2)}$ where $x^{(1,2)}$ are two training inputs with different classes. We observe that the mean and variance dynamics of neural network outputs during gradient descent training follow the analytic dynamics from linearization well (Equations 15, 16). Moreover the NNGP posterior which corresponds to exact Bayesian inference, while similar, is noticeably *different* from the predictive distribution at the end of gradient descent training. For dynamics for individual function draws, see SM Figure S4.

3.3. Comparison of training dynamics of linearized network to original network

For a particular realization of a finite width network, one can analytically predict the dynamics of the weights and outputs over the course of training using the empirical tangent kernel at initialization. In Figures 5,6,7, we compare these linearized dynamics (Equations 9, 10) with the result of training the actual network. In all cases we see

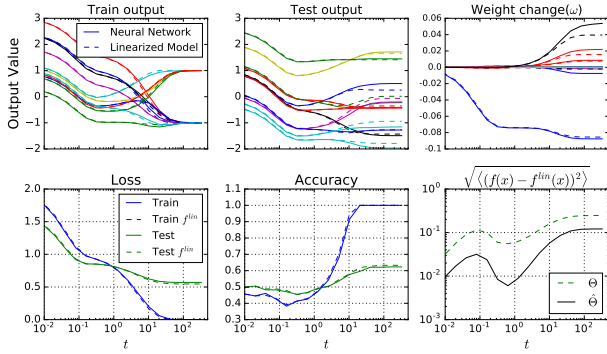


Figure 4. Full batch gradient descent on a model behaves similarly to analytic dynamics on its linearization, both for network outputs, and also for individual weights. A binary CIFAR classification task with MSE loss and a ReLU fully connected network with 5 hidden layers of width $n = 2048$, $\eta = 0.01$, $|\mathcal{D}| = 256$, $k = 1$, $\sigma_w^2 = 2.0$, and $\sigma_b^2 = 0.1$. All three panes in the first row show dynamics for a randomly selected subset of datapoints or parameters. The first two panes in the second row show that the dynamics of loss and accuracy for training and test points agree well between the original and linearized model. The bottom right pane shows the dynamics of RMSE between the two models on test points. We observe that the empirical kernel $\hat{\Theta}$ gives more accurate dynamics for finite width networks.

remarkably good agreement. We also observe that for finite networks, dynamics predicted using the empirical kernel $\hat{\Theta}$ better match the data than those obtained using the infinite-width, analytic, kernel Θ . To understand this we note that $\|\hat{\Theta}_T^{(n)} - \hat{\Theta}_0^{(n)}\|_F = O(\frac{1}{n}) < O(\frac{1}{\sqrt{n}}) = \|\hat{\Theta}_0^{(n)} - \Theta\|_F$, as plotted in Figure 2.

For general loss, e.g. cross-entropy with softmax output, we need to rely on solving the ODE Equations 25 and 26. We use the `dopri5` method for ODE integration, which is the default integrator in TensorFlow (`tf.contrib.integrate.odeint`). In Figure 4, we see that the learning dynamics for the CIFAR-10 all class classification task with cross-entropy loss are well described by the linearized model. In Figure 6, we tested full MNIST digit classification with cross-entropy loss, and trained with a momentum optimizer. For cross-entropy loss with softmax output, some logits at late times grow indefinitely, in contrast to MSE loss where logits converge to target value. The error between original and linearized model for cross entropy loss becomes much worse at late times if the two models deviate significantly before the logits enter their late-time steady-growth regime (See Figure S5.).

One can directly optimize parameters of f^{lin} instead of solving the ODE induced by the tangent kernel $\hat{\Theta}$. Standard neural network optimization techniques such as mini-batching,

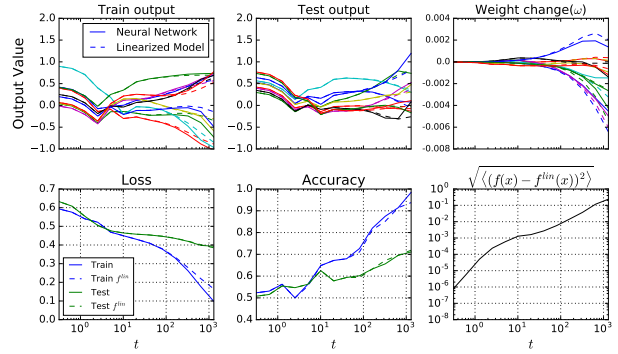


Figure 5. A convolutional network and its linearization behave similarly when trained using full batch gradient descent with a momentum optimizer Binary CIFAR classification task with MSE loss, tanh convolutional network with 3 hidden layers of channel size $n = 512$, 3×3 size filters, average pooling after last convolutional layer, $\eta = 0.1$, $\beta = 0.9$, $|\mathcal{D}| = 128$, $\sigma_w^2 = 2.0$ and $\sigma_b^2 = 0.1$. The linearized model is trained directly by full batch gradient descent with momentum, rather than by integrating its continuous time analytic dynamics. Panes are the same as in Figure 4.

weight decay, and data augmentation can be directly applied. In Figure 5 and 7, we compared the training dynamics of the linearized and original network while directly training both networks.

As discussed in Section 2.5.3, the linearized dynamics successfully describes the training of networks beyond vanilla fully connected models. To demonstrate the generality of this procedure we show we can predict the learning dynamics of Wide Residual Networks (WRN) (Zagoruyko & Komodakis, 2016). WRNs are a class of model that are popular in computer vision and leverage convolutions, batch normalization, skip connections, and average pooling. In Figure 7, we show a comparison between the linearized dynamics and the true dynamics for a wide residual network trained with MSE loss and SGD with momentum. We slightly modified the block structure described in Figure 7 so that each layer has a constant number of channels (1024 in this case) and otherwise followed the original implementation. As elsewhere, we see strong agreement between the predicted dynamics and the result of training.

3.4. Effects of depth and dataset size

The training dynamics of a neural network match those of its linearization when the width is infinite and the dataset is finite. In previous experiments, we chose sufficiently wide networks to achieve small error between neural networks and their linearization for smaller datasets. Here we investigate how the agreement between the linearized dynamics




Figure 6. A neural network and its linearization behave similarly when both are trained via SGD with momentum on cross entropy loss on MNIST. Experiment is for 10 class MNIST classification using a ReLU fully connected network with 2 hidden layers of width $n = 2048$, $\eta = 1.0$, $\beta = 0.9$, $|\mathcal{D}| = 50,000$, $k = 10$, $\sigma_w^2 = 2.0$, and $\sigma_b^2 = 0.1$. Both models are trained using stochastic minibatching with batch size 64. Panes are the same as in Figure 4, except that the top row shows all ten logits for a single randomly selected datapoint.

and the true dynamics behaves as a function of width and dataset size across a wide range of models. We consider Root Mean Squared Error (RMSE) between the predicted outputs and the true outputs of the network over the test set. Generally, the RMSE will increase with time until it plateaus at the end of training at some final value. In Figure 8 we plot this plateau RMSE for a range of models as a function of both width and dataset size. Overall, we observe that as the width grows the error decreases. This decrease goes approximately as $1/N$ for fully-connected networks, with more ambiguous scaling for convolutional and WRN architectures. Additionally, we see that the error grows approximately linearly in the size of the dataset. Thus, although error grows with dataset this can be counterbalanced by a corresponding increase in the model size.

4. Discussion

We showed theoretically that the learning dynamics in parameter space of deep nonlinear neural networks are exactly described by a linearized model in the infinite width limit. Empirical investigation revealed that this agrees very well with actual training dynamics and predictive distributions across fully-connected, convolutional, and even wide residual network architectures, as well as with different optimizers (gradient descent, momentum, mini-batching) and loss functions (MSE, cross-entropy). Our results suggest that a surprising number of realistic neural networks may be operating in the regime we studied.

In the regime which we study, since the learning dynam-

GROUP NAME	OUTPUT SIZE	BLOCK TYPE
CONV1	32×32	$[3 \times 3, \text{CHANNEL SIZE}]$
CONV2	32×32	$\begin{bmatrix} 3 \times 3, \text{CHANNEL SIZE} \\ 3 \times 3, \text{CHANNEL SIZE} \end{bmatrix} \times N$
CONV3	16×16	$\begin{bmatrix} 3 \times 3, \text{CHANNEL SIZE} \\ 3 \times 3, \text{CHANNEL SIZE} \\ 3 \times 3, \text{CHANNEL SIZE} \end{bmatrix} \times N$
CONV4	8×8	$\begin{bmatrix} 3 \times 3, \text{CHANNEL SIZE} \\ 3 \times 3, \text{CHANNEL SIZE} \\ 3 \times 3, \text{CHANNEL SIZE} \end{bmatrix} \times N$
AVG-POOL	1×1	$[8 \times 8]$




Figure 7. A wide residual network and its linearization behave similarly when both are trained by SGD with momentum on MSE loss on CIFAR-10. (top) We adopt the network architecture from Zagaruyko & Komodakis (2016). In the residual block, we follow Batch Normalization-ReLU-Conv ordering. We use $N = 1$, channel size 1024, $\eta = 1.0$, $\beta = 0.9$, $k = 10$, $\sigma_w^2 = 1.0$, and $\sigma_b^2 = 0.0$. Both linearized and original model are trained directly on full CIFAR-10 ($|\mathcal{D}| = 50,000$), using stochastic minibatching with batch size 8. (bottom) Output dynamics for a randomly selected subset of train and test points are shown in the first row. The second row shows training and accuracy curves for original and linearized networks.

ics is fully captured by the kernel $\hat{\Theta}$ and the target signal, studying the properties of $\hat{\Theta}$ to determine trainability and generalization are interesting future directions. Furthermore, the infinite width limit gives us a simple characterization of both gradient descent and Bayesian inference. Some preliminary observations in Lee et al. (2018) showed that wide neural networks trained with SGD perform similarly to the corresponding GPs as width increase, while Novak et al. (2019) found the opposite in the case of convopostlutional networks without pooling. By studying properties of the NNGP kernel \mathcal{K} and the tangent kernel Θ , we may shed light on the inductive bias of gradient descent.

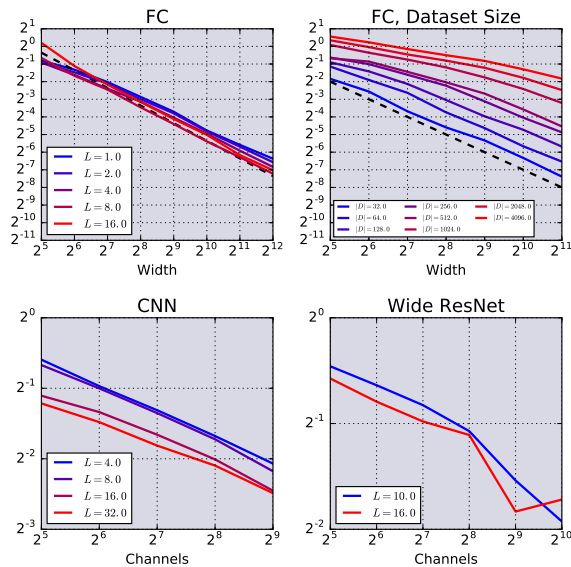


Figure 8. Error dependence on depth and dataset size. Final value of the RMSE for fully-connected, convolutional, wide residual network as networks becomes wider for varying depth and dataset size. Top row: error in fully connected networks as the depth is varied from 1 to 16 (left) and the dataset size is varied from 32 to 4096 (right). Bottom row: error in convolutional networks (left) as the depth is varied between 1 and 32 and WRN (right) for depths 10 and 16 corresponding to $N=1,2$ described in Figure 7. Network is critically initialized $\sigma_w^2 = 2.0$, $\sigma_b^2 = 0.1$, trained with gradient descent on MSE loss.

Acknowledgements

We thank Roman Novak and Greg Yang for useful discussions and feedback.

References

- Allen-Zhu, Z., Li, Y., and Song, Z. A convergence theory for deep learning via over-parameterization. *arXiv preprint arXiv:1811.03962*, 2018a.
- Allen-Zhu, Z., Li, Y., and Song, Z. On the convergence rate of training recurrent neural networks. *arXiv preprint arXiv:1810.12065*, 2018b.
- Chen, M., Pennington, J., and Schoenholz, S. Dynamical isometry and a mean field theory of RNNs: Gating enables signal propagation in recurrent neural networks. In *International Conference on Machine Learning*, 2018.
- Chizat, L. and Bach, F. On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Advances in neural information processing systems*, pp. 3040–3050, 2018a.
- Chizat, L. and Bach, F. A note on lazy training in supervised differentiable programming. *arXiv preprint arXiv:1812.07956*, 2018b.
- Cho, Y. and Saul, L. K. Kernel methods for deep learning. In *Advances in neural information processing systems*, 2009.
- Daniely, A. SGD learns the conjugate kernel class of the network. In *Advances in Neural Information Processing Systems 30*. 2017.
- Daniely, A., Frostig, R., and Singer, Y. Toward deeper understanding of neural networks: The power of initialization and a dual view on expressivity. In *Advances In Neural Information Processing Systems*, 2016.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Dragomir, S. S. *Some Gronwall type inequalities and applications*. Nova Science Publishers New York, 2003.
- Du, S. S., Lee, J. D., Li, H., Wang, L., and Zhai, X. Gradient descent finds global minima of deep neural networks. *arXiv preprint arXiv:1811.03804*, 2018.
- Garriga-Alonso, A., Aitchison, L., and Rasmussen, C. E. Deep convolutional networks as shallow gaussian processes. *arXiv preprint arXiv:1808.05587*, 2018.
- Glorot, X. and Bengio, Y. Understanding the difficulty of training deep feedforward neural networks. In *International Conference on Artificial Intelligence and Statistics*, pp. 249–256, 2010.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems 31*. 2018.
- Karras, T., Aila, T., Laine, S., and Lehtinen, J. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*. 2012.

- Lee, J., Bahri, Y., Novak, R., Schoenholz, S., Pennington, J., and Sohl-dickstein, J. Deep neural networks as gaussian processes. In *International Conference on Learning Representations*, 2018.
- Matthews, A. G. d. G., Hron, J., Turner, R. E., and Ghahramani, Z. Sample-then-optimize posterior sampling for bayesian linear models. In *NeurIPS Workshop on Advances in Approximate Bayesian Inference*, 2017. URL <http://approximateinference.org/2017/accepted/MatthewsEtAl2017.pdf>.
- Matthews, A. G. d. G., Hron, J., Rowland, M., Turner, R. E., and Ghahramani, Z. Gaussian process behaviour in wide deep neural networks. In *International Conference on Learning Representations*, 2018.
- Mei, S., Montanari, A., and Nguyen, P.-M. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33): E7665–E7671, 2018.
- Neal, R. M. Priors for infinite networks (tech. rep. no. crg-tr-94-1). *University of Toronto*, 1994.
- Neyshabur, B., Tomioka, R., and Srebro, N. In search of the real inductive bias: On the role of implicit regularization in deep learning. In *International Conference on Learning Representations workshop track*, 2015.
- Neyshabur, B., Li, Z., Bhojanapalli, S., LeCun, Y., and Srebro, N. The role of over-parametrization in generalization of neural networks. In *International Conference on Learning Representations*, 2019.
- Novak, R., Bahri, Y., Abolafia, D. A., Pennington, J., and Sohl-Dickstein, J. Sensitivity and generalization in neural networks: an empirical study. In *International Conference on Learning Representations*, 2018.
- Novak, R., Xiao, L., Lee, J., Bahri, Y., Yang, G., Hron, J., Abolafia, D. A., Pennington, J., and Sohl-Dickstein, J. Bayesian deep convolutional networks with many channels are gaussian processes. In *International Conference on Learning Representations*, 2019.
- Park, D. S., Smith, S. L., Sohl-dickstein, J., and Le, Q. V. Optimal SGD hyperparameters for fully connected networks. In *NeurIPS Bayesian Deep Learning Workshop*, 2018. URL <http://bayesiandeeplearning.org/2018/papers/86.pdf>.
- Poole, B., Lahiri, S., Raghu, M., Sohl-Dickstein, J., and Ganguli, S. Exponential expressivity in deep neural networks through transient chaos. In *Advances In Neural Information Processing Systems*, pp. 3360–3368, 2016.
- Qian, N. On the momentum term in gradient descent learning algorithms. *Neural networks*, 12(1):145–151, 1999.
- Rotskoff, G. M. and Vanden-Eijnden, E. Neural networks as interacting particle systems: Asymptotic convexity of the loss landscape and universal scaling of the approximation error. *arXiv preprint arXiv:1805.00915*, 2018.
- Saxe, A. M., McClelland, J. L., and Ganguli, S. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*, 2013.
- Schoenholz, S. S., Gilmer, J., Ganguli, S., and Sohl-Dickstein, J. Deep information propagation. *International Conference on Learning Representations*, 2017.
- Sirignano, J. and Spiliopoulos, K. Mean field analysis of neural networks. *arXiv preprint arXiv:1805.01053*, 2018.
- Su, W., Boyd, S., and Candes, E. A differential equation for modeling nesterovs accelerated gradient method: Theory and insights. In *Advances in Neural Information Processing Systems*, pp. 2510–2518, 2014.
- van Laarhoven, T. L2 regularization versus batch and weight normalization. *arXiv preprint arXiv:1706.05350*, 2017.
- Williams, C. K. Computing with infinite networks. In *Advances in neural information processing systems*, pp. 295–301, 1997.
- Xiao, L., Bahri, Y., Sohl-Dickstein, J., Schoenholz, S., and Pennington, J. Dynamical isometry and a mean field theory of CNNs: How to train 10,000-layer vanilla convolutional neural networks. In *International Conference on Machine Learning*, 2018.
- Yang, G. Scaling limits of wide neural networks with weight sharing: Gaussian process behavior, gradient independence, and neural tangent kernel derivation. *arXiv preprint arXiv:1902.04760*, 2019.
- Yang, G. and Schoenholz, S. Mean field residual networks: On the edge of chaos. In *Advances in Neural Information Processing Systems*. 2017.
- Yang, G., Pennington, J., Rao, V., Sohl-Dickstein, J., and Schoenholz, S. S. A mean field theory of batch normalization. In *International Conference on Learning Representations*, 2019.
- Zagoruyko, S. and Komodakis, N. Wide residual networks. In *British Machine Vision Conference*, 2016.
- Zhang, C., Bengio, S., and Singer, Y. Are all layers created equal? *arXiv preprint arXiv:1902.01996*, 2019.
- Zou, D., Cao, Y., Zhou, D., and Gu, Q. Stochastic gradient descent optimizes over-parameterized deep relu networks. *arXiv preprint arXiv:1811.08888*, 2018.

Supplemental Material

A. Computing NTK and NNGP Kernel

For completeness, we reproduce, informally, the recursive formula of the NNGP kernel and the tangent kernel from (Lee et al., 2018) and (Jacot et al., 2018), respectively. Let the activation function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ be absolutely continuous. Let \mathcal{T} and $\dot{\mathcal{T}}$ be functions from 2×2 positive semi-definite matrices Σ to \mathbb{R} given by

$$\begin{cases} \mathcal{T}(\Sigma) = \mathbb{E}[\phi(u)\phi(v)] \\ \dot{\mathcal{T}}(\Sigma) = \mathbb{E}[\phi'(u)\phi'(v)] \end{cases} \quad (u, v) \sim \mathcal{N}(0, \Sigma). \quad (\text{S1})$$

In the infinite width limit, the NNGP and tangent kernel can be computed recursively. Let x, x' be two inputs in \mathbb{R}^{n_0} . Then $h^l(x)$ and $h^l(x')$ converge in distribution to a joint Gaussian as $\min\{n_1, \dots, n_{l-1}\}$. The mean is zero and the variance $\mathcal{K}^l(x, x')$ is

$$\mathcal{K}^l(x, x') = \tilde{\mathcal{K}}^l(x, x') \otimes \mathbf{Id}_{n_l} \quad (\text{S2})$$

$$\tilde{\mathcal{K}}^l(x, x') = \sigma_\omega^2 \mathcal{T} \left(\begin{bmatrix} \tilde{\mathcal{K}}^{l-1}(x, x) & \tilde{\mathcal{K}}^{l-1}(x, x') \\ \tilde{\mathcal{K}}^{l-1}(x, x') & \tilde{\mathcal{K}}^{l-1}(x', x') \end{bmatrix} \right) + \sigma_b^2 \quad (\text{S3})$$

with base case

$$\mathcal{K}^1(x, x') = \sigma_\omega^2 \cdot \frac{1}{n_0} x^T x' + \sigma_b^2. \quad (\text{S4})$$

Using this one can also derive the tangent kernel for gradient descent training. We will use induction to show that

$$\Theta^l(x, x') = \tilde{\Theta}^l(x, x') \otimes \mathbf{Id}_{n_l} \quad (\text{S5})$$

where

$$\tilde{\Theta}^l(x, x') = \tilde{\mathcal{K}}^l(x, x') + \sigma_\omega^2 \tilde{\Theta}^{l-1}(x, x') \dot{\mathcal{T}} \left(\begin{bmatrix} \tilde{\mathcal{K}}^{l-1}(x, x) & \tilde{\mathcal{K}}^{l-1}(x, x') \\ \tilde{\mathcal{K}}^{l-1}(x, x') & \tilde{\mathcal{K}}^{l-1}(x', x') \end{bmatrix} \right) \quad (\text{S6})$$

with $\tilde{\Theta}^1 = \tilde{\mathcal{K}}^1$. Let

$$J^l(x) = \nabla_{\theta \leq l} h_0^l(x) = [\nabla_{\theta^l} h_0^l(x), \nabla_{\theta < l} h_0^l(x)]. \quad (\text{S7})$$

Then

$$J^l(x) J^l(x')^T = \nabla_{\theta^l} h_0^l(x) \nabla_{\theta^l} h_0^l(x')^T + \nabla_{\theta < l} h_0^l(x) \nabla_{\theta < l} h_0^l(x')^T \quad (\text{S8})$$

Letting $n_1, \dots, n_{l-1} \rightarrow \infty$ sequentially, the first term converges to the NNGP kernel $\mathcal{K}^l(x, x')$. By applying the chain rule and the induction step (letting $n_1, \dots, n_{l-2} \rightarrow \infty$ sequentially), the second term is

$$\nabla_{\theta < l} h_0^l(x) \nabla_{\theta < l} h_0^l(x')^T = \frac{\partial h_0^l(x)}{\partial h_0^{l-1}(x)} \nabla_{\theta \leq l-1} h_0^{l-1}(x) \nabla_{\theta \leq l-1} h_0^{l-1}(x')^T \frac{\partial h_0^l(x')}{\partial h_0^{l-1}(x')} \quad (\text{S9})$$

$$\rightarrow \frac{\partial h_0^l(x)}{\partial h_0^{l-1}(x)} \tilde{\Theta}^{l-1}(x, x') \otimes \mathbf{Id}_{n_{l-1}} \frac{\partial h_0^l(x')}{\partial h_0^{l-1}(x')} \quad (n_1, \dots, n_{l-2} \rightarrow \infty) \quad (\text{S10})$$

$$\rightarrow \sigma_\omega^2 \left(\mathbb{E} \phi'(h_{0,i}^{l-1}(x)) \phi'(h_{0,i}^{l-1}(x')) \right) \tilde{\Theta}^{l-1}(x, x') \otimes \mathbf{Id}_{n_l} \quad (n_{l-1} \rightarrow \infty) \quad (\text{S11})$$

$$= \left(\sigma_\omega^2 \tilde{\Theta}^{l-1}(x, x') \dot{\mathcal{T}} \left(\begin{bmatrix} \tilde{\mathcal{K}}^{l-1}(x, x) & \tilde{\mathcal{K}}^{l-1}(x, x') \\ \tilde{\mathcal{K}}^{l-1}(x, x') & \tilde{\mathcal{K}}^{l-1}(x', x') \end{bmatrix} \right) \right) \otimes \mathbf{Id}_{n_l} \quad (\text{S12})$$

B. Tangent kernel for ReLU and erf

For ReLU and erf activation functions, the tangent kernel can be computed analytically. We begin with the case $\phi = \text{ReLU}$; using the formula from [Cho & Saul \(2009\)](#), we can compute \mathcal{T} and $\dot{\mathcal{T}}$ in closed form. Let Σ be a 2×2 PSD matrix. We will use

$$k_n(x, y) = \int \phi^n(x \cdot w) \phi^n(y \cdot w) e^{-\|w\|^2/2} dw \cdot (2\pi)^{-d/2} = \frac{1}{2\pi} \|x\|^n \|y\|^n J_n(\theta) \quad (\text{S13})$$

where

$$\begin{aligned} \phi(x) &= \max(x, 0), \quad \theta(x, y) = \arccos\left(\frac{x \cdot y}{\|x\| \|y\|}\right), \\ J_0(\theta) &= \pi - \theta, \quad J_1(\theta) = \sin \theta + (\pi - \theta) \cos \theta = \sqrt{1 - \left(\frac{x \cdot y}{\|x\| \|y\|}\right)^2} + (\pi - \theta) \left(\frac{x \cdot y}{\|x\| \|y\|}\right). \end{aligned} \quad (\text{S14})$$

Let $d = 2$ and $u = (x \cdot w, y \cdot w)^T$. Then u is a mean zero Gaussian with $\Sigma = [[x \cdot x, x \cdot y]; [x \cdot y, y \cdot y]]$. Then

$$\mathcal{T}(\Sigma) = k_1(x, y) = \frac{1}{2\pi} \|x\| \|y\| J_1(\theta) \quad (\text{S15})$$

$$\dot{\mathcal{T}}(\Sigma) = k_0(x, y) = \frac{1}{2\pi} J_0(\theta) \quad (\text{S16})$$

For $\phi = \text{erf}$, let Σ be the same as above. Following [Williams \(1997\)](#), we get

$$\mathcal{T}(\Sigma) = \frac{2}{\pi} \sin^{-1}\left(\frac{2x \cdot y}{\sqrt{(1 + 2x \cdot x)(1 + 2y \cdot y)}}\right) \quad (\text{S17})$$

$$\dot{\mathcal{T}}(\Sigma) = \frac{4}{\pi} \det(I + 2\Sigma)^{-1/2} \quad (\text{S18})$$

C. Multi-dimensional output and cross-entropy loss

For completeness, we include the derivation for cross entropy loss for softmax output,

$$\ell(f, y) = - \sum_i y^i \log \sigma(f^i), \quad \sigma(f^i) \equiv \frac{\exp(f^i)}{\sum_j \exp(f^j)}. \quad (\text{S19})$$

Recall that $\frac{\partial \ell}{\partial \hat{y}^i} = \sigma(\hat{y}^i) - y^i$. For general input point x and for an arbitrary parameterized function $f^i(x)$ parameterized by θ , gradient flow dynamics is given by

$$\dot{f}_t^i(x) = \nabla_{\theta} f_t^i(x) \frac{d\theta}{dt} = -\eta \nabla_{\theta} f_t^i(x) \sum_j \sum_{(z, y) \in \mathcal{D}} \left[\nabla_{\theta} f_t^j(z)^T \frac{\partial \ell(f_t, y)}{\partial \hat{y}^j} \right] \quad (\text{S20})$$

$$= -\eta \sum_{(z, y) \in \mathcal{D}} \sum_j \nabla_{\theta} f_t^i(x) \nabla_{\theta} f_t^j(z)^T \left(\sigma(f_t^j(z)) - y^j \right) \quad (\text{S21})$$

Let $\Theta^{ij}(x, \mathcal{X}) = \nabla_{\theta} f^i(x) \nabla_{\theta} f^j(\mathcal{X})^T$. The above is

$$\dot{f}_t(\mathcal{X}) = -\eta \Theta_t(\mathcal{X}, \mathcal{X}) (\sigma(f_t(\mathcal{X})) - \mathcal{Y}) \quad (\text{S22})$$

$$\dot{f}_t(x) = -\eta \Theta_t(x, \mathcal{X}) (\sigma(f_t(\mathcal{X})) - \mathcal{Y}). \quad (\text{S23})$$

The linearization is

$$\dot{f}_t^{\text{lin}}(\mathcal{X}) = -\eta \Theta_0(\mathcal{X}, \mathcal{X}) (\sigma(f_t^{\text{lin}}(\mathcal{X})) - \mathcal{Y}) \quad (\text{S24})$$

$$\dot{f}_t^{\text{lin}}(x) = -\eta \Theta_0(x, \mathcal{X}) (\sigma(f_t^{\text{lin}}(\mathcal{X})) - \mathcal{Y}). \quad (\text{S25})$$

D. Gradient flow dynamics for training only the readout-layer

The connection between Gaussian processes and Bayesian wide neural networks can be extended to the setting when only the readout layer parameters are being optimized. More precisely, we show that when training only the readout layer, the outputs of the network form a Gaussian process (over an ensemble of draws from the parameter prior) throughout training, where that output is an interpolation between the GP prior and GP posterior.

Note that for any $x, x' \in \mathbb{R}^{n_0}$, in the infinite width limit $\bar{x}(x) \cdot \bar{x}(x') = \hat{\mathcal{K}}(x, x') \rightarrow \mathcal{K}(x, x')$ in probability, where for notational simplicity we assign $\bar{x}(x) = [\frac{\sigma_w x^L(x)}{\sqrt{n_L}}, \sigma_b]$. The regression problem is specified with mean-squared loss

$$\mathcal{L} = \frac{1}{2} \|f(\mathcal{X}) - \mathcal{Y}\|_2^2 = \frac{1}{2} \|\bar{x}(\mathcal{X})\theta^{L+1} - \mathcal{Y}\|_2^2, \quad (\text{S26})$$

and applying gradient flow to optimize the readout layer (and freezing all other parameters),

$$\dot{\theta}^{L+1} = -\eta \bar{x}(\mathcal{X})^T (\bar{x}(\mathcal{X})\theta^{L+1} - \mathcal{Y}), \quad (\text{S27})$$

where η is the learning rate. The solution to this ODE gives the evolution of the output of an arbitrary x^* . So long as the empirical kernel $\bar{x}(\mathcal{X})\bar{x}(\mathcal{X})^T$ is invertible, it is

$$f_t(x^*) = f_0(x^*) + \hat{\mathcal{K}}(x, \mathcal{X})\hat{\mathcal{K}}(\mathcal{X}, \mathcal{X})^{-1} \left(\exp(-\eta t \hat{\mathcal{K}}(\mathcal{X}, \mathcal{X})) - I \right) (f_0(\mathcal{X}) - \mathcal{Y}) \quad (\text{S28})$$

For any $x, x' \in \mathbb{R}^{n_0}$, letting $n_l \rightarrow \infty$ for $l = 1, \dots, L$, one has the convergence in distribution in probability and distribution respectively

$$\bar{x}(x)\bar{x}(x') \rightarrow \mathcal{K}(x, x') \quad \text{and} \quad \bar{x}(\mathcal{X})\theta_0^{L+1} \rightarrow \mathcal{N}(0, \mathcal{K}(\mathcal{X}, \mathcal{X})). \quad (\text{S29})$$

Moreover $\bar{x}(\mathcal{X})\theta_0^{L+1}$ and the term containing $f_0(\mathcal{X})$ are the only stochastic term over the ensemble of network initializations, therefore for any t the output $f(x^*)$ throughout training converges to a Gaussian distribution in the infinite width limit, with

$$\mathbb{E}[f_t(x^*)] = \mathcal{K}(x^*, \mathcal{X})\mathcal{K}^{-1}(I - e^{-\eta \mathcal{K}t})\mathcal{Y}, \quad (\text{S30})$$

$$\text{Var}[f_t(x^*)] = \mathcal{K}(x^*, x^*) - \mathcal{K}(x^*, \mathcal{X})\mathcal{K}^{-1}(I - e^{-\eta \mathcal{K}t})\mathcal{K}(x^*, \mathcal{X})^T. \quad (\text{S31})$$

Thus the output of the neural network is also a GP and the asymptotic solution (i.e. $t \rightarrow \infty$) is identical to the posterior of the NNGP (Equation 14). Therefore, in the infinite width case, the optimized neural network is performing posterior sampling if only the readout layer is being trained. This result is a realization of sample-then-optimize equivalence identified in [Matthews et al. \(2017\)](#).

E. Results for NTK parameterization transfer to standard parameterization

Here we present a sketch for why the linearization results, derived for NTK parameterized networks, also apply to networks with a standard parameterization

$$\begin{cases} h^{l+1} &= x^l W^{l+1} + b^{l+1} \\ x^{l+1} &= \phi(h^{l+1}) \end{cases} \quad \text{and} \quad \begin{cases} W_{i,j}^l &= \omega_{ij}^l \sim \mathcal{N}(0, \frac{\sigma_\omega^2}{n_l}) \\ b_j^l &= \beta_j^l \sim \mathcal{N}(0, \sigma_b^2) \end{cases}. \quad (\text{S32})$$

The NTK parameterization in Equation 1 is not commonly used for training neural networks. While the function that the network represents is the same for both NTK and standard parameterization, training dynamics under gradient descent are generally different for the two parameterizations. However, for a particular choice of layer-dependent learning rate training dynamics also become identical. Let $\eta_{\text{NTK},w}^l$ and $\eta_{\text{NTK},b}^l$ be layer-dependent learning rate for W^l and b^l in the NTK parameterization, and $\eta_{\text{std}} = \frac{1}{n_{\text{max}}}\eta_0$ be the learning rate for all parameters in the standard parameterization, where $n_{\text{max}} = \max_l n_l$. Recall that gradient descent training in standard neural networks requires a learning rate that scales with width like $\frac{1}{n_{\text{max}}}$, so η_0 defines a width-invariant learning rate ([Park et al., 2018](#)). If we choose

$$\eta_{\text{NTK},w}^l = \frac{n_l}{n_{\text{max}}\sigma_\omega^2}\eta_0, \quad \text{and} \quad \eta_{\text{NTK},b}^l = \frac{1}{n_{\text{max}}\sigma_b^2}\eta_0, \quad (\text{S33})$$




Figure S1. NTK vs Standard parameterization. Across different choices of dataset, activation function and loss function, models obtained from (S)GD training for both parameterization (circle and triangle denotes NTK and standard parameterization respectively) get similar performance.

then learning dynamics are identical for networks with NTK and standard parameterizations. With only extremely minor modifications, consisting of incorporating the multiplicative factors in Equation S33 into the per-layer contributions to the Jacobian, the arguments in Section 2.4 go through for an NTK network with learning rates defined in Equation S33. Since an NTK network with these learning rates exhibits identical training dynamics to a standard network with learning rate η_{std} , the result in Section 2.4 that sufficiently wide NTK networks are linear in their parameters throughout training also applies to standard networks.

We can verify this property of networks with the standard parameterization experimentally. In Figure S1, we see that for different choices of dataset, activation function and loss function, final performance of two different parameterization leads to similar quality model for similar value of normalized learning rate $\eta_{\text{std}} = \eta_{\text{NTK}}/n$. Also, in Figure S2, we observe that our results is not due to the parameterization choice and holds for wide networks using the standard parameterization.

F. Bounding the discrepancy between the original and the linearized network: MSE loss

To simplify the notation, let $g^{\text{lin}}(t) \equiv f_t^{\text{lin}}(\mathcal{X}) - \mathcal{Y}$ and $g(t) \equiv f_t(\mathcal{X}) - \mathcal{Y}$.

$$\frac{d}{dt} (\exp(\eta\Theta_0 t)(g^{\text{lin}}(t) - g(t))) = \eta (\Theta_0 \exp(\eta\Theta_0 t)(g^{\text{lin}}(t) - g(t)) + \exp(\eta\Theta_0 t)(-\Theta_0 g^{\text{lin}}(t) + \Theta_t g(t))) \quad (\text{S34})$$

$$= \eta (\exp(\eta\Theta_0 t)(\Theta_t - \Theta_0)g(t)) \quad (\text{S35})$$

Integrating both sides and using the fact $g^{\text{lin}}(0) = g(0)$,

$$(g^{\text{lin}}(t) - g(t)) = - \int_0^t \eta (\exp(\eta\Theta_0(s-t))(\Theta_s - \Theta_0)(g^{\text{lin}}(s) - g(s))) ds \quad (\text{S36})$$

$$+ \int_0^t \eta (\exp(\eta\Theta_0(s-t))(\Theta_s - \Theta_0)g^{\text{lin}}(s)) ds \quad (\text{S37})$$

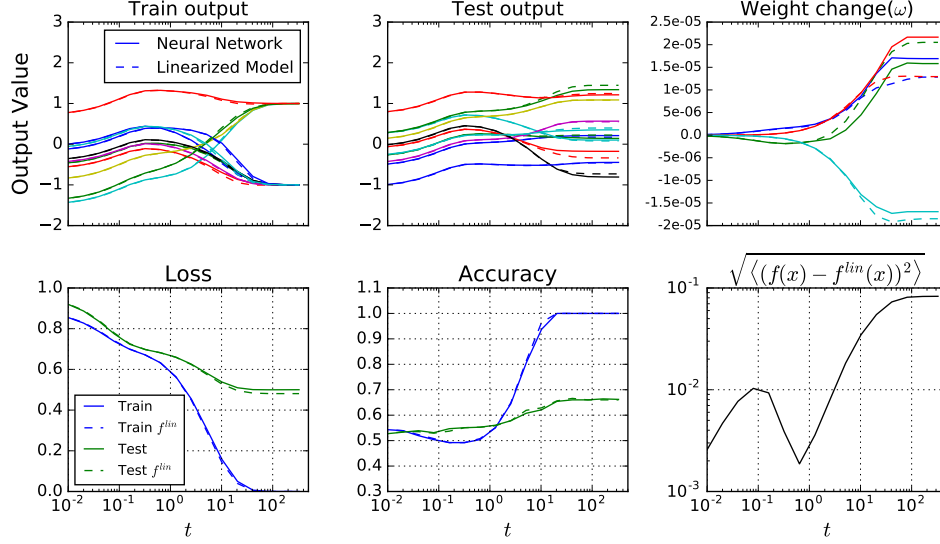


Figure S2. **Exact and experimental dynamics are nearly identical for network outputs, and are similar for individual weights (Standard parameterization).** Experiment is for a MSE loss, ReLU network with 5 hidden layers of width $n = 2048$, $\eta = 0.005/2048$, $|\mathcal{D}| = 256$, $k = 1$, $\sigma_w^2 = 2.0$, and $\sigma_b^2 = 0.1$. All three panes in the first row show dynamics for a randomly selected subset of datapoints or parameters. First two panes in the second row show dynamics of loss and accuracy for training and test points agree well between original and linearized model. Bottom right pane shows the dynamics of RMSE between the two models on test points using empirical kernel.

Let $\lambda_0 > 0$ be the smallest eigenvalue of Θ_0 . Taking the norm gives

$$\|g^{\text{lin}}(t) - g(t)\|_2 \leq \eta \left(\int_0^t \|\exp(\Theta_0 \eta(s-t))\|_{op} \|(\Theta_s - \Theta_0)\|_{op} \|g^{\text{lin}}(s) - g(s)\|_2 ds \right) \quad (\text{S38})$$

$$+ \int_0^t \|\exp(\Theta_0 \eta(s-t))\|_{op} \|(\Theta_s - \Theta_0)\|_{op} \|g^{\text{lin}}(s)\|_2 ds \quad (\text{S39})$$

$$\leq \eta \left(\int_0^t e^{\eta \lambda_0 (s-t)} \|(\Theta_s - \Theta_0)\|_{op} \|g^{\text{lin}}(s) - g(s)\|_2 ds \right) \quad (\text{S40})$$

$$+ \int_0^t e^{\eta \lambda_0 (s-t)} \|(\Theta_s - \Theta_0)\|_{op} \|g^{\text{lin}}(s)\|_2 ds \quad (\text{S41})$$

Let

$$u(t) \equiv e^{\lambda_0 \eta t} \|g^{\text{lin}}(t) - g(t)\|_2 \quad (\text{S42})$$

$$\alpha(t) \equiv \eta \int_0^t e^{\lambda_0 \eta s} \|(\Theta_s - \Theta_0)\|_{op} \|g^{\text{lin}}(s)\|_2 ds \quad (\text{S43})$$

$$\beta(t) \equiv \eta \|(\Theta_t - \Theta_0)\|_{op} \quad (\text{S44})$$

The above can be written as

$$u(t) \leq \alpha(t) + \int_0^t \beta(s) u(s) ds \quad (\text{S45})$$

Note that $\alpha(t)$ is non-decreasing. Applying an integral form of the Grönwall's inequality (see Theorem 1 in (Dragomir, 2003)) gives

$$u(t) \leq \alpha(t) \exp\left(\int_0^t \beta(s) ds\right) \quad (\text{S46})$$

Note that

$$\|g^{\text{lin}}(t)\|_2 = \|\exp(-\eta\Theta_0 t) g^{\text{lin}}(0)\|_2 \leq \|\exp(-\eta\Theta_0 t)\|_{op} \|g^{\text{lin}}(0)\|_2 = e^{-\lambda_0 \eta t} \|g^{\text{lin}}(0)\|_2. \quad (\text{S47})$$

Then

$$\|g^{\text{lin}}(t) - g(t)\|_2 \leq \eta e^{-\lambda_0 \eta t} \int_0^t e^{\lambda_0 \eta s} \|\Theta_s - \Theta_0\|_{op} \|g^{\text{lin}}(s)\|_2 ds \exp\left(\int_0^t \eta \|\Theta_s - \Theta_0\|_{op} ds\right) \quad (\text{S48})$$

$$\leq \eta e^{-\lambda_0 \eta t} \|g^{\text{lin}}(0)\|_2 \int_0^t \|\Theta_s - \Theta_0\|_{op} ds \exp\left(\int_0^t \eta \|\Theta_s - \Theta_0\|_{op} ds\right) \quad (\text{S49})$$

Let $\sigma_t = \sup_{0 \leq s \leq t} \|\Theta_s - \Theta_0\|_{op}$. Then

$$\|g^{\text{lin}}(t) - g(t)\|_2 \lesssim e^{-\lambda_0 \eta t} (\eta t \sigma_t e^{\eta t \sigma_t}) \quad (\text{S50})$$

For a fixed amount of time T , if $\sigma_T \rightarrow 0$ as $\min\{n_1, \dots, n_L\} \rightarrow \infty$, then

$$\sup_{0 \leq t \leq T} e^{\eta t \sigma_t}, \quad \sup_{0 \leq t \leq T} e^{-\lambda_0 \eta t} \eta t = O(1) \quad (\text{S51})$$

Therefore

$$\sup_{0 \leq t \leq T} \|g^{\text{lin}}(t) - g(t)\|_2 = O(\sigma_T) \rightarrow 0 \quad (\text{S52})$$

as $\min\{n_1, \dots, n_L\} \rightarrow \infty$.

Now we control the discrepancy on a test point x . Let y be its true label. Similarly,

$$\frac{d}{dt} (g^{\text{lin}}(t, x) - g(t, x)) = -\eta (\Theta_0(x, \mathcal{X}) - \Theta_t(x, \mathcal{X})) g^{\text{lin}}(t) + \eta \Theta_t(x, \mathcal{X}) (g(t) - g^{\text{lin}}(t)). \quad (\text{S53})$$

Integrating over $[0, t]$ and taking the norm imply

$$\|g^{\text{lin}}(t, x) - g(t, x)\|_2 \leq \eta \int_0^t \|\Theta_0(x, \mathcal{X}) - \Theta_s(x, \mathcal{X})\|_2 \|g^{\text{lin}}(s)\|_2 ds + \eta \int_0^t \|\Theta_s(x, \mathcal{X})\|_2 \|g(s) - g^{\text{lin}}(s)\|_2 ds \quad (\text{S54})$$

$$\leq \eta \|g^{\text{lin}}(0)\|_2 \int_0^t \|\Theta_0(x, \mathcal{X}) - \Theta_s(x, \mathcal{X})\|_2 e^{-\eta \lambda_0 s} ds \quad (\text{S55})$$

$$+ \eta \int_0^t (\|\Theta_0(x, \mathcal{X})\|_2 + \|\Theta_s(x, \mathcal{X}) - \Theta_0(x, \mathcal{X})\|_2) \|g(s) - g^{\text{lin}}(s)\|_2 ds \quad (\text{S56})$$

For fixed T , assuming $\sup_{0 \leq t \leq T} \|\Theta_0(x, \mathcal{X}) - \Theta_t(x, \mathcal{X})\|_2 \rightarrow 0$ and $\|\Theta_0(x, \mathcal{X})\|_2$ is bounded as $\min\{n_1, \dots, n_L\} \rightarrow \infty$ and using Equation S52,

$$\|g^{\text{lin}}(t, x) - g(t, x)\|_2 \rightarrow 0. \quad (\text{S57})$$

G. Additional plots

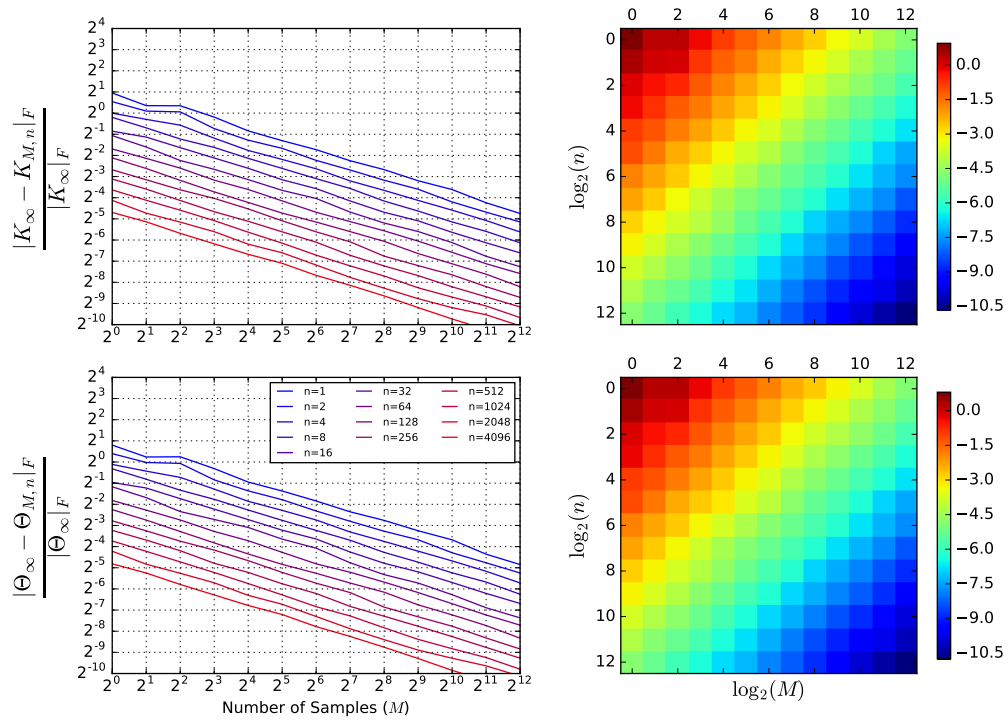


Figure S3. **Kernel convergence.** Kernels from single hidden layer randomly initialized ReLU network convergence to analytic kernel using Monte Carlo sampling (M samples).




Figure S4. **Sample of neural network outputs.** The lines correspond to the functions learned for 100 different initializations. The configuration is that same as in Figure 3.

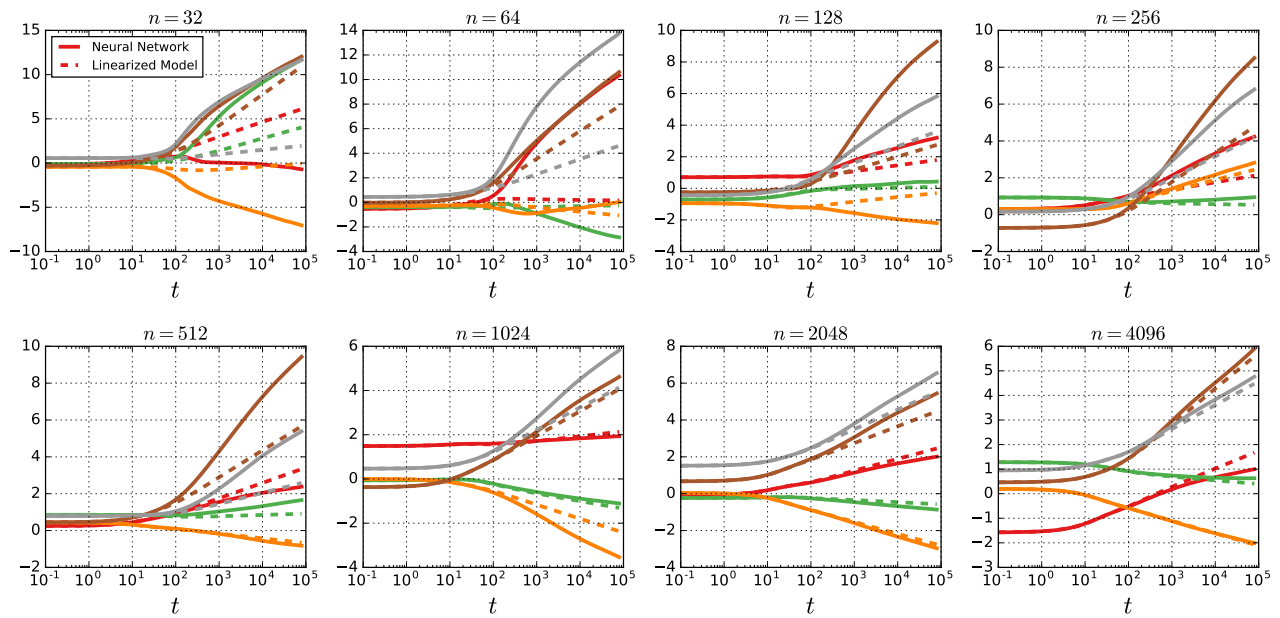


Figure S5. **Logit deviation for cross entropy loss.** Logits for models trained with cross entropy loss diverge at late times. If the deviation between the logits of the linearized model and original model are large early in training, as shown for the narrower networks (first row), logit deviation at late times can be significantly large. As network becomes wider (second row), the logit deviates at later point in training. Fully connected tanh network $L = 4$ trained on binary CIFAR classification problem.