

The NTK approximation is valid for longer than you think

Anonymous authors
Paper under double-blind review

Abstract

We study when the neural tangent kernel (NTK) approximation is valid for training a model with the square loss. In the lazy training setting of [Chizat et al. \(2019\)](#), we show that rescaling the model by a factor of $\alpha = O(T)$ suffices for the NTK approximation to be valid until training time T . Our bound is tight and improves on the previous bound of [Chizat et al. \(2019\)](#) which required a larger rescaling factor of $\alpha = O(T^2)$.

1 Introduction

In the modern machine learning paradigm, practitioners train the weights \mathbf{w} of a large neural network model $f_{\mathbf{w}} : \mathbb{R}^{d_{in}} \rightarrow \mathbb{R}^{d_{out}}$ via a gradient-based optimizer. Theoretical understanding lags behind, since the training dynamics are non-linear and hence difficult to analyze. To address this, [Jacot et al. \(2018\)](#) proposed an approximation to the dynamics called the *NTK approximation*, and proved it was valid for infinitely-wide networks trained by gradient descent¹. The NTK approximation has been extremely influential, leading to theoretical explanations for a range of questions, including why deep learning can memorize training data ([Du et al., 2018; 2019; Allen-Zhu et al., 2019a;b; Arora et al., 2019a; Cao & Gu, 2019; Lee et al., 2019](#)), why neural networks exhibit spectral bias ([Cao et al., 2019; Basri et al., 2020; Canatar et al., 2021](#)), and why different architectures generalize differently ([Bietti & Mairal, 2019; Mei et al., 2021; Wang et al., 2021](#)). Nevertheless, in practice the training dynamics of neural networks often diverge from the predictions of the NTK approximation (see, e.g., [Arora et al. \(2019b\)](#)) Therefore, it is of interest to understand exactly under which conditions the NTK approximation holds. In this paper, we ask the following question:

Can we give tight conditions for when the NTK approximation is valid?

1.1 The “lazy training” setting of [Chizat et al. \(2019\)](#)

The work of [Chizat et al. \(2019\)](#) showed that the NTK approximation actually holds for training any differentiable model, as long as the model’s outputs are *rescaled* so that the model’s outputs change by a large amount even when the weights change by a small amount. The correctness of the NTK approximation for infinite-width models is a consequence of this observation, because by the default the model is rescaled as the width tends to infinity; see the related work in [Section 1.3](#) for more details.

Rescaling the model Let $h : \mathbb{R}^p \rightarrow \mathcal{F}$ be a smoothly-parametrized model, where \mathcal{F} is a separable Hilbert space. Let $\alpha > 0$ be a parameter which controls the rescaling of the model and which should be thought of as large. We train the rescaled model αh with gradient flow to minimize a smooth loss function $R : \mathcal{F} \rightarrow \mathbb{R}_+$ ². Namely, the weights $\mathbf{w}(t) \in \mathbb{R}^p$ are initialized at $\mathbf{w}(0) = \mathbf{w}_0$ and evolve according to the gradient flow

$$\frac{d\mathbf{w}}{dt} = -\frac{1}{\alpha^2} \nabla_{\mathbf{w}} R(\alpha h(\mathbf{w}(t))). \quad (1)$$

¹Under a specific scaling of the initialization and learning rate as width tends to infinity.

²We use the Hilbert space notation as in [Chizat et al. \(2019\)](#). We can recover the setting of training a neural network $f_{\mathbf{w}} : \mathbb{R}^d \rightarrow \mathbb{R}$ on a finite training dataset $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\} \subseteq \mathbb{R}^d \times \mathbb{R}$ with empirical loss function $\mathcal{L}(\mathbf{w}) = \frac{1}{n} \sum_{i=1}^n \ell(f_{\mathbf{w}}(\mathbf{x}_i), y_i)$ as follows. Let $\mathcal{H} = \mathbb{R}^n$ be the Hilbert space, let $h(\mathbf{w}) = [f_{\mathbf{w}}(\mathbf{x}_1), \dots, f_{\mathbf{w}}(\mathbf{x}_n)]$, and let $R(v) = \frac{1}{n} \sum_{i=1}^n \ell(v_i, y_i)$.

NTK approximation Define the linear approximation of the model around the initial weights \mathbf{w}_0 by

$$\bar{h}(\mathbf{w}) = h(\mathbf{w}_0) + Dh(\mathbf{w}_0)(\mathbf{w} - \mathbf{w}_0), \quad (2)$$

where Dh is the first derivative of h in \mathbf{w} . Let $\bar{\mathbf{w}}(t)$ be weights initialized at $\bar{\mathbf{w}}(0) = \mathbf{w}_0$ that evolve according to the gradient flow from training the rescaled linearized model $\alpha\bar{h}$:

$$\frac{d\bar{\mathbf{w}}}{dt} = -\frac{1}{\alpha^2} \nabla_{\bar{\mathbf{w}}} R(\alpha\bar{h}(\bar{\mathbf{w}}(t))). \quad (3)$$

The NTK approximation states that

$$\alpha h(\mathbf{w}(t)) \approx \alpha \bar{h}(\bar{\mathbf{w}}(t)).$$

In other words, it states that the linearization of the model h is valid throughout training. This allows for much simpler analysis of the training dynamics since the model \bar{h} is linear in its parameters, and so the evolution of $\bar{h}(\bar{\mathbf{w}})$ can be understood via a kernel gradient flow in function space.

When is the NTK approximation valid? Chizat et al. (2019) proves that if the rescaling parameter α is large, then the NTK approximation is valid. The intuition is that the weights do not need to move far from their initialization in order to change the output of the model significantly, so the linearization (2) is valid for longer. Since the weights stay close to initialization, Chizat et al. (2019) refer to this regime of training as “lazy training.” The following bound is proved.³

Proposition 1.1 (Theorem 2.3 of Chizat et al. (2019)). *Let $R(y) = \frac{1}{2}\|y - y^*\|^2$ be the square loss, where $y^* \in \mathcal{F}$ are the target labels. Assume that h is $\text{Lip}(h)$ -Lipschitz and that Dh is $\text{Lip}(Dh)$ -Lipschitz in a ball of radius ρ around \mathbf{w}_0 . Then, for any time $0 \leq T \leq \alpha\rho/(\text{Lip}(h)\sqrt{R_0})$,*

$$\|\alpha h(\mathbf{w}(T)) - \alpha \bar{h}(\bar{\mathbf{w}}(T))\| \leq T \text{Lip}(h)^2 \kappa \sqrt{R_0}, \quad (4)$$

where $R_0 := R(\alpha h(\mathbf{w}_0))$ is the loss at initialization, and κ is

$$\kappa = \frac{T}{\alpha} \text{Lip}(Dh) \sqrt{R_0}.$$

Notice that if we take α large, then the right-hand-side of (4) is small, so the NTK approximation is valid.

1.2 Our results

Our contribution is to refine the bound of Chizat et al. (2019) for large time scales. We prove:

Theorem 1.2 (NTK approximation error bound). *Let $R(y) = \frac{1}{2}\|y - y^*\|^2$ be the square loss. Assume that Dh is $\text{Lip}(Dh)$ -Lipschitz in a ball of radius ρ around \mathbf{w}_0 . Then, at any time $0 \leq T \leq \alpha^2 \rho^2 / R_0$,*

$$\|\alpha h(\mathbf{w}(T)) - \alpha \bar{h}(\bar{\mathbf{w}}(T))\| \leq \min(6\kappa\sqrt{R_0}, \sqrt{8R_0}). \quad (5)$$

Furthermore, the converse is true. Our bound is tight up to a constant factor.

Theorem 1.3 (Converse to Theorem 1.2). *For any $\alpha, T, \text{Lip}(Dh)$, and R_0 , there is a model $h : \mathbb{R} \rightarrow \mathbb{R}$, an initialization $w_0 \in \mathbb{R}$, and a target $y^* \in \mathbb{R}$ such that, for the risk $R(y) = \frac{1}{2}(y - y^*)^2$, the initial risk is $R(\alpha h(w_0)) = R_0$, the derivative map Dh is $\text{Lip}(Dh)$ -Lipschitz, and*

$$\|\alpha h(w(T)) - \alpha \bar{h}(\bar{w}(T))\| \geq \min\left(\frac{1}{5}\kappa\sqrt{R_0}, \frac{1}{5}\sqrt{R_0}\right).$$

Comparison to Chizat et al. (2019) In contrast to our theorem, the bound (4) depends on the Lipschitz constant of h , and incurs an extra factor of $T \text{Lip}(h)^2$. So if $\text{Lip}(Dh)$, $\text{Lip}(h)$, and R_0 are bounded by constants, our result shows that the NTK approximation (up to $O(\epsilon)$ error) is valid for times $T = O(\alpha\epsilon)$, while the previously known bound is valid for $T = O(\sqrt{\alpha\epsilon})$. Since the regime of interest is training for large times $T \gg 1$, our result shows that the NTK approximation holds for much longer time horizons than previously known.

³See Section 1.3 for discussion on the other results of Chizat et al. (2019).

1.3 Additional related literature

Other results of Chizat et al. (2019) In addition to the bound of Proposition 1.1 above, Chizat et al. (2019) controls the error in the NTK approximation in two other settings: (a) for general losses, but α must be taken exponential in T , and (b) for strongly convex losses and infinite training time T , but the problem must be “well-conditioned.” We work in the setting of Proposition 1.1 instead, since it is more aligned with the situation in practice, where we have long training times and the problem is ill-conditioned. Indeed, the experiments of Chizat et al. (2019) report that for convolutional neural networks on CIFAR10 trained in the lazy regime, the problem is ill-conditioned, and training takes a long time to converge.

Other works on the validity of the NTK approximation The result of Jacot et al. (2018) on the validity of NTK for infinitely-wide neural networks can be viewed under the lazy training setting, since the hyperparameters chosen by Jacot et al. (2018) give rescalings of models in hyperparameter regimes where the NTK approximation is invalid even at infinite width (Chizat & Bach, 2018; Rotskoff & Vanden-Eijnden, 2018; Sirignano & Spiliopoulos, 2022; Mei et al., 2018; 2019; Yang & Hu, 2021). This fact is used by Woodworth et al. (2020); Geiger et al. (2020; 2021) to study the effect of interpolating between lazy and non-lazy training by varying the scaling parameter α . Most work points towards provable benefits of non-lazy training (Allen-Zhu & Li, 2019; Bai & Lee, 2019; Bai et al., 2020; Chen et al., 2020; Nichani et al., 2022; Ghorbani et al., 2020; Malach et al., 2021; Abbe et al., 2022; 2023; Mousavi-Hosseini et al., 2022; Damian et al., 2022; Bietti et al., 2022; Ba et al., 2022) although interestingly there are settings where lazy training provably outperforms non-lazy training (Petrini et al., 2022).

2 Proof ideas

2.1 Proof ideas for Theorem 1.2

Proof of Chizat et al. (2019) In order to give intuition for our proof, we first explain the idea behind the proof in Chizat et al. (2019). Define residuals $r(t), \bar{r}(t) \in \mathcal{F}$ under training the original rescaled model and the linearized rescaled model as $r(t) = y^* - \alpha h(\mathbf{w}(t))$ and $\bar{r}(t) = y^* - \alpha \bar{h}(\bar{\mathbf{w}}(t))$. It is well known that these evolve according to

$$\frac{dr}{dt} = -K_t r \quad \text{and} \quad \frac{d\bar{r}}{dt} = -K_0 \bar{r},$$

for the time-dependent kernel $K_t : \mathcal{F} \rightarrow \mathcal{F}$ which is the linear operator given by $K_t := Dh(\mathbf{w}(t))Dh(\mathbf{w}(t))^\top$. To compare these trajectories, Chizat et al. (2019) observes that, since K_0 is p.s.d.,

$$\frac{1}{2} \frac{d}{dt} \|r - \bar{r}\|^2 = -\langle r - \bar{r}, K_t r - K_0 \bar{r} \rangle \leq -\langle r - \bar{r}, (K_t - K_0)r \rangle,$$

which, dividing both sides by $\|r - \bar{r}\|$ and using that $\|r\| \leq \sqrt{R_0}$ implies

$$\frac{d}{dt} \|r - \bar{r}\| \leq \|K_t - K_0\| \|r\| \leq 2\text{Lip}(h)\text{Lip}(Dh)\|\mathbf{w} - \mathbf{w}_0\| \sqrt{R_0}. \quad (6)$$

Using the Lipschitzness of the model, Chizat et al. (2019) furthermore proves that the weight change is bounded by $\|\mathbf{w}(t) - \mathbf{w}_0\| \leq t\sqrt{R_0}\text{Lip}(h)/\alpha$. Plugging this into (6) yields the bound in Proposition 1.1,

$$\|\alpha h(\mathbf{w}(T)) - \alpha \bar{h}(\bar{\mathbf{w}}(T))\| = \|r(T) - \bar{r}(T)\| \leq \int_0^T 2\text{Lip}(h)^2 \text{Lip}(Dh) t^2 R_0 \alpha^{-1} dt = T^2 \text{Lip}(h)^2 \text{Lip}(Dh) R_0 / \alpha.$$

A simple strengthening of the bound for long time horizons Our first attempt to strengthen this bound for longer time horizons is to use an improved bound on the movement of the weights. Consider the following proposition implicit in Chizat et al. (2019).

Proposition 2.1 (Bound on weight change, implicit in proof of Theorem 2.2 in Chizat et al. (2019)).

$$\|\mathbf{w}(T) - \mathbf{w}_0\| \leq \sqrt{TR_0}/\alpha \quad \text{and} \quad \|\bar{\mathbf{w}}(T) - \mathbf{w}_0\| \leq \sqrt{TR_0}/\alpha. \quad (7)$$

Proof of Proposition 2.1. By (a) Cauchy-Schwarz, and (b) the nonnegativity of the loss R ,

$$\|\mathbf{w}(T) - \mathbf{w}(0)\| \leq \int_0^T \left\| \frac{d\mathbf{w}}{dt} \right\| dt \stackrel{(a)}{\leq} \sqrt{T \int_0^T \left\| \frac{d\mathbf{w}}{dt} \right\|^2 dt} = \sqrt{-\frac{T}{\alpha^2} \int_0^T \frac{d}{dt} R(\alpha h(\mathbf{w}(t))) dt} \stackrel{(b)}{\leq} \sqrt{TR_0}/\alpha.$$

The bound for $\bar{\mathbf{w}}$ is analogous. \square

This bound (7) has the benefit of \sqrt{t} dependence (instead of linear t dependence), and also does not depend on the Lipschitzness constant. So if we plug it into (6), we obtain

$$\|\alpha h(\mathbf{w}(T)) - \alpha \bar{h}(\bar{\mathbf{w}}(T))\| \leq \int_0^T 2\text{Lip}(h)\text{Lip}(Dh)\|\mathbf{w} - \mathbf{w}_0\|\sqrt{t}R_0\alpha^{-1}dt = \frac{4}{3}T^{3/2}\text{Lip}(h)\text{Lip}(Dh)R_0/\alpha.$$

This improves over Proposition 1.1 for long time horizons since the time dependence scales as $T^{3/2}$ instead of as T^2 . However, it still depends on the Lipschitz constant $\text{Lip}(h)$ of h , and it also falls short of the linear in T dependence of Theorem 1.2.

Intuition for the tight bound of Theorem 1.2 In order to avoid dependence on $\text{Lip}(h)$ and obtain a linear dependence in T , a new approach is required. The bound (6) is not a feasible starting point, since it depends on $\text{Lip}(h)$. Furthermore, in order to achieve linear T dependence using (6), we would need that $\|\mathbf{w} - \mathbf{w}_0\| = O(1)$ for a constant that does not depend on the time horizon, which is not true unless the problem is well-conditioned.

In the full proof in Appendix A, we bound $\|r(T) - \bar{r}(T)\| = \|\alpha h(\mathbf{w}(T)) - \alpha \bar{h}(\bar{\mathbf{w}}(T))\|$, which requires working with a product integral formulation of the dynamics of r to handle the time-varying kernels K_t (Dollard & Friedman, 1984). Here, we highlight some of the ideas by providing a proof of a simplified statement. Consider $r'(t) \in \mathcal{F}$ which is initialized as $r'(0) = r(0)$ and evolves as

$$\frac{dr'}{dt} = -K_T r'.$$

The intuition is that this is a ‘‘worst-case’’ scenario, since it is at time $t = T$ that the kernel K_t can differ the most from K_0 . So intuitively if we can prove that $r'(T)$ and $\bar{r}(T)$ are close, then the same should be true for $r(T)$ and $\bar{r}(T)$. For convenience, define the operators $A = Dh(\mathbf{w}_0)^\top$ and $B = Dh(\mathbf{w}(T))^\top - Dh(\mathbf{w}_0)^\top$. Since the kernels do not vary in time, the closed-form solution is

$$r'(t) = e^{-(A+B)^\top(A+B)t} r(0) \quad \text{and} \quad \bar{r}(t) = e^{-A^\top A t} r(0)$$

We prove that the time evolution operators for r' and \bar{r} are close in operator norm.

Lemma 2.2. *For any $t \geq 0$, we have $\|e^{-(A+B)^\top(A+B)t} - e^{-A^\top A t}\| \leq 2\|B\|\sqrt{t}$.*

Proof of Lemma 2.2. Define $Z(\zeta) = -(A + \zeta B)^\top(A + \zeta B)t$. By the fundamental theorem of calculus

$$\|e^{-(A+B)^\top(A+B)t} - e^{-A^\top A t}\| = \|e^{Z(1)} - e^{Z(0)}\| = \left\| \int_0^1 \frac{d}{d\zeta} e^{Z(\zeta)} d\zeta \right\| \leq \sup_{\zeta \in [0,1]} \left\| \frac{d}{d\zeta} e^{Z(\zeta)} \right\|.$$

Using the integral representation of the exponential map (Dollard & Friedman, 1984),

$$\begin{aligned} \left\| \frac{de^{Z(\zeta)}}{d\zeta} \right\| &= \left\| \int_0^1 e^{(1-\tau)Z(\zeta)} \left(\frac{d}{d\zeta} Z(\zeta) \right) e^{\tau Z(\zeta)} d\tau \right\| \\ &= \left\| t \int_0^1 e^{(1-\tau)Z(\zeta)} (A^\top B + B^\top A + 2\zeta B^\top B) e^{\tau Z(\zeta)} d\tau \right\| \\ &\leq \underbrace{\left\| t \int_0^1 e^{(1-\tau)Z(\zeta)} (A + \zeta B)^\top B e^{\tau Z(\zeta)} d\tau \right\|}_{\text{(Term 1)}} + \underbrace{\left\| t \int_0^1 e^{(1-\tau)Z(\zeta)} B^\top (A + \zeta B) e^{\tau Z(\zeta)} d\tau \right\|}_{\text{(Term 2)}} \end{aligned}$$

By symmetry under transposing and reversing time, (Term 1) = (Term 2), so it suffices to bound the first term. Since $\|e^{\tau Z(\zeta)}\| \leq 1$,

$$\begin{aligned}
(\text{Term 1}) &\leq t \int_0^1 \|e^{(1-\tau)Z(\zeta)}(A + \zeta B)^\top\| \|B\| \|e^{\tau Z(\zeta)}\| d\tau \\
&\leq t \|B\| \int_0^1 \|e^{(1-\tau)Z(\zeta)}(A + \zeta B)^\top\| d\tau \\
&= t \|B\| \int_0^1 \sqrt{\|e^{(1-\tau)Z(\zeta)}(A + \zeta B)^\top(A + \zeta B)e^{(1-\tau)Z(\zeta)}\|} d\tau \\
&= \sqrt{t} \|B\| \int_0^1 \sqrt{\|e^{(1-\tau)Z(\zeta)} Z(\zeta) e^{(1-\tau)Z(\zeta)}\|} d\tau \\
&\leq \sqrt{t} \|B\| \int_0^1 \sup_{\lambda \geq 0} \sqrt{\lambda e^{-2(1-\tau)\lambda}} d\tau \\
&= \sqrt{t} \|B\| \int_0^1 \sqrt{1/(2e(1-\tau))} d\tau \\
&= \sqrt{2t/e} \|B\|.
\end{aligned}$$

where in the third-to-last line we use the Courant-Fischer-Weyl theorem and the fact that $Z(\zeta)$ is negative semidefinite. Combining these bounds $\|e^{-(A+B)^\top(A+B)t} - e^{-A^\top A t}\| \leq 2\sqrt{2t/e} \|B\| \leq 2\|B\|\sqrt{t}$. \square

Using Lemma 2.2, we can show that the dynamics of r' are close to those of \bar{r} :

$$\|r'(T) - \bar{r}(T)\| \leq \min(3\kappa\sqrt{R_0}, \sqrt{8R_0}). \quad (8)$$

To see why, notice that the weight-change bound in Proposition 2.1 implies

$$\|B\| \leq \text{Lip}(Dh) \|\mathbf{w}(T) - \mathbf{w}_0\| \leq \text{Lip}(Dh) \sqrt{TR_0}/\alpha.$$

So Lemma 2.2 implies

$$\|r'(T) - \bar{r}(T)\| \leq 2\text{Lip}(Dh)T\sqrt{R_0}\alpha^{-1}\|r(0)\| = 2\kappa\|r(0)\|.$$

Combining this with $\|r'(T) - \bar{r}(T)\| \leq \|r'(T)\| + \|\bar{r}(T)\| \leq 2\|r(0)\| = 2\sqrt{2R_0}$ implies (8). Thus, we have shown the result of Theorem 1.2 if we replace r by r' . The actual proof of the theorem handles the time-varying kernel K_t , and is in Appendix A.

2.2 Proof ideas for Theorem 1.3

The converse in Theorem 1.3 is achieved in the simple case where $h(w) = aw + \frac{1}{2}bw^2$ for $a = \frac{1}{\sqrt{T}}$ and $b = \text{Lip}(Dh)$, and $w_0 = 0$ and $R(y) = \frac{1}{2}(y - \sqrt{2R_0})^2$, as we show in Appendix B by direct calculation.

3 Discussion

A limitation of our result is that it applies only to the gradient flow, which corresponds to SGD with infinitesimally small step size. However, larger step sizes are beneficial for generalization in practice (see, e.g., Li et al. (2021); Andriushchenko et al. (2022)), so it would be interesting to understand the validity of the NTK approximation in that setting. Another limitation is that our result applies only to the square loss, and not to other popular losses such as the cross-entropy loss. Indeed, the known bounds in the setting of general losses require either a ‘‘well-conditioning’’ assumption, or taking α exponential in the training time T (Chizat et al., 2019). Can one prove bounds analogous to Theorem 1.2 for more general losses, with α depending polynomially on T , and without conditioning assumptions?

A natural question raised by our bounds in Theorems 1.2 and 1.3 is: how do the dynamics behave just outside the regime where the NTK approximation is valid? For models h where $\text{Lip}(h)$ and $\text{Lip}(Dh)$ are bounded by a constant, can we understand the dynamics in the regime where $T \geq C\alpha$ for some large constant C and $\alpha \gg C$, at the edge of the lazy training regime?

References

- Emmanuel Abbe, Enric Boix-Adsera, and Theodor Misiakiewicz. The merged-staircase property: a necessary and nearly sufficient condition for sgd learning of sparse functions on two-layer neural networks. In *Conference on Learning Theory*, pp. 4782–4887. PMLR, 2022.
- Emmanuel Abbe, Enric Boix-Adsera, and Theodor Misiakiewicz. Sgd learning on neural networks: leap complexity and saddle-to-saddle dynamics. *arXiv preprint arXiv:2302.11055*, 2023.
- Zeyuan Allen-Zhu and Yuanzhi Li. What can resnet learn efficiently, going beyond kernels? *Advances in Neural Information Processing Systems*, 32, 2019.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. *Advances in neural information processing systems*, 32, 2019a.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via overparameterization. In *International Conference on Machine Learning*, pp. 242–252. PMLR, 2019b.
- Maksym Andriushchenko, Aditya Varre, Loucas Pillaud-Vivien, and Nicolas Flammarion. Sgd with large step sizes learns sparse features. *arXiv preprint arXiv:2210.05337*, 2022.
- Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *Proceedings of the 36th International Conference on Machine Learning*, pp. 322–332, 2019a.
- Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Russ R Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. In *Advances in Neural Information Processing Systems*, 2019b.
- Jimmy Ba, Murat A Erdogdu, Taiji Suzuki, Zhichao Wang, Denny Wu, and Greg Yang. High-dimensional asymptotics of feature learning: How one gradient step improves the representation. *arXiv preprint arXiv:2205.01445*, 2022.
- Yu Bai and Jason D Lee. Beyond linearization: On quadratic and higher-order approximation of wide neural networks. *arXiv preprint arXiv:1910.01619*, 2019.
- Yu Bai, Ben Krause, Huan Wang, Caiming Xiong, and Richard Socher. Taylorized training: Towards better approximation of neural network training at finite width. *arXiv preprint arXiv:2002.04010*, 2020.
- Ronen Basri, Meirav Galun, Amnon Geifman, David Jacobs, Yoni Kasten, and Shira Kritchman. Frequency bias in neural networks for input of non-uniform density. In *International Conference on Machine Learning*, pp. 685–694. PMLR, 2020.
- Alberto Bietti and Julien Mairal. On the inductive bias of neural tangent kernels. *Advances in Neural Information Processing Systems*, 32, 2019.
- Alberto Bietti, Joan Bruna, Clayton Sanford, and Min Jae Song. Learning single-index models with shallow neural networks. *arXiv preprint arXiv:2210.15651*, 2022.
- Abdulkadir Canatar, Blake Bordelon, and Cengiz Pehlevan. Spectral bias and task-model alignment explain generalization in kernel regression and infinitely wide neural networks. *Nature communications*, 12(1):2914, 2021.
- Yuan Cao and Quanquan Gu. Generalization bounds of stochastic gradient descent for wide and deep neural networks. In *Advances in Neural Information Processing Systems*, 2019.
- Yuan Cao, Zhiying Fang, Yue Wu, Ding-Xuan Zhou, and Quanquan Gu. Towards understanding the spectral bias of deep learning. *arXiv preprint arXiv:1912.01198*, 2019.

- Minshuo Chen, Yu Bai, Jason D Lee, Tuo Zhao, Huan Wang, Caiming Xiong, and Richard Socher. Towards understanding hierarchical learning: Benefits of neural representations. *Advances in Neural Information Processing Systems*, 33:22134–22145, 2020.
- Lenaïc Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. *Advances in neural information processing systems*, 31, 2018.
- Lenaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. *Advances in Neural Information Processing Systems*, 32, 2019.
- Alexandru Damian, Jason Lee, and Mahdi Soltanolkotabi. Neural networks can learn representations with gradient descent. In *Conference on Learning Theory*, pp. 5413–5452. PMLR, 2022.
- John Day Dollard and Charles N. Friedman. *Product Integration with Application to Differential Equations*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1984. doi: 10.1017/CBO9781107340701.
- Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *International conference on machine learning*, pp. 1675–1685. PMLR, 2019.
- Simon S Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. *arXiv preprint arXiv:1810.02054*, 2018.
- Mario Geiger, Stefano Spigler, Arthur Jacot, and Matthieu Wyart. Disentangling feature and lazy training in deep neural networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2020(11):113301, 2020.
- Mario Geiger, Leonardo Petrini, and Matthieu Wyart. Landscape and training regimes in deep learning. *Physics Reports*, 924:1–18, 2021.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. When do neural networks outperform kernel methods? *Advances in Neural Information Processing Systems*, 33:14820–14830, 2020.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. *Advances in neural information processing systems*, 32, 2019.
- Zhiyuan Li, Sadhika Malladi, and Sanjeev Arora. On the validity of modeling sgd with stochastic differential equations (sdes). *Advances in Neural Information Processing Systems*, 34:12712–12725, 2021.
- Eran Malach, Pritish Kamath, Emmanuel Abbe, and Nathan Srebro. Quantifying the benefit of using differentiable learning over tangent kernels. In *International Conference on Machine Learning*, pp. 7379–7389. PMLR, 2021.
- Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.
- Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. In *Conference on Learning Theory*, pp. 2388–2464. PMLR, 2019.
- Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Learning with invariances in random features and kernel models. In *Conference on Learning Theory*, pp. 3351–3418. PMLR, 2021.
- Alireza Mousavi-Hosseini, Sejun Park, Manuela Girotti, Ioannis Mitliagkas, and Murat A Erdogdu. Neural networks efficiently learn low-dimensional representations with sgd. *arXiv preprint arXiv:2209.14863*, 2022.
- Eshaan Nichani, Yu Bai, and Jason D Lee. Identifying good directions to escape the ntk regime and efficiently learn low-degree plus sparse polynomials. *arXiv preprint arXiv:2206.03688*, 2022.

Leonardo Petrini, Francesco Cagnetta, Eric Vanden-Eijnden, and Matthieu Wyart. Learning sparse features can lead to overfitting in neural networks. *arXiv preprint arXiv:2206.12314*, 2022.

Grant M Rotskoff and Eric Vanden-Eijnden. Neural networks as interacting particle systems: Asymptotic convexity of the loss landscape and universal scaling of the approximation error. *stat*, 1050:22, 2018.

Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of deep neural networks. *Mathematics of Operations Research*, 47(1):120–152, 2022.

Terence Tao. Topics in random matrix theory. *Graduate Studies in Mathematics*, 132, 2011.

Sifan Wang, Hanwen Wang, and Paris Perdikaris. On the eigenvector bias of fourier feature networks: From regression to solving multi-scale pdes with physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 384:113938, 2021.

Blake Woodworth, Suriya Gunasekar, Jason D Lee, Edward Moroshko, Pedro Savarese, Itay Golan, Daniel Soudry, and Nathan Srebro. Kernel and rich regimes in overparametrized models. In *Conference on Learning Theory*, pp. 3635–3673. PMLR, 2020.

Greg Yang and Edward J Hu. Tensor programs iv: Feature learning in infinite-width neural networks. In *International Conference on Machine Learning*, pp. 11727–11737. PMLR, 2021.

A Proof of Theorem 1.2

A.1 Notations

We let $R_0 := R(\alpha h(\mathbf{w}_0))$ denote the loss at initialization. We define the residuals $r(t), \bar{r}(t) \in \mathcal{F}$ under training the original model and the linearized model as

$$r(t) = y^* - \alpha h(\mathbf{w}(t)), \text{ and } \bar{r}(t) = y^* - \alpha \bar{h}(\bar{\mathbf{w}}(t)).$$

Since the evolution of \mathbf{w} and $\bar{\mathbf{w}}$ is given by the gradient flow in (1) and (3), the residuals evolve as follows. We write $Dh^\top = \left(\frac{dh}{d\mathbf{w}}\right)^\top$ to denote the adjoint of $Dh = \frac{dh}{d\mathbf{w}}$.

$$\frac{dr}{dt} = \frac{dr}{d\mathbf{w}} \frac{d\mathbf{w}}{dt} = -\alpha Dh(\mathbf{w})(-\nabla F_\alpha(\mathbf{w})) = \alpha D(\mathbf{w})\left(\frac{1}{\alpha} Dh(\mathbf{w})^\top \nabla R(\alpha h(\mathbf{w}))\right) = -Dh(\mathbf{w})Dh(\mathbf{w})^\top r,$$

since $\nabla R(\alpha h(\mathbf{w})) = -(y^* - \alpha h(\mathbf{w})) = -r$. An analogous result can be derived for the residual \bar{r} under the linearized dynamics:

$$\frac{d\bar{r}}{dt} = -D\bar{h}(\bar{\mathbf{w}})D\bar{h}(\bar{\mathbf{w}})^\top \bar{r}.$$

For any time $t \geq 0$, define the kernel $K_t : \mathcal{F} \rightarrow \mathcal{F}$ as

$$K_t := Dh(\mathbf{w}(t))Dh(\mathbf{w}(t))^\top.$$

Since $K_0 = Dh(\mathbf{w}(0))Dh(\mathbf{w}(0))^\top = D\bar{h}(\bar{\mathbf{w}}(0))D\bar{h}(\bar{\mathbf{w}}(0))^\top$, we can write the dynamics in compact form:

$$\frac{dr}{dt} = -K_t r \quad \text{and} \quad \frac{d\bar{r}}{dt} = -K_0 \bar{r}. \quad (9)$$

A.2 Basic facts about boundedness and continuity of the kernel

We begin with a series of simple propositions. First, the weights $\mathbf{w}(t)$ stay within the ball of radius ρ around \mathbf{w}_0 , allowing us to use the bounds $\text{Lip}(h)$ and $\text{Lip}(Dh)$ on the Lipschitz constants of h and Dh .

Lemma A.1. *For any time $0 \leq t \leq T$, we have $\|\mathbf{w}(t) - \mathbf{w}_0\| \leq \rho$.*

Proof. Immediate from Proposition 2.1 and the fact that $T \leq \rho^2 \alpha^2 / R_0$. \square

Now we show that the kernels K_t and K_0 stay close during training, in the sense that the difference between Dh and $D\bar{h}$ during training is bounded.

Lemma A.2. *For any time $0 \leq t \leq T$, we have*

$$\|Dh(\mathbf{w}(t)) - D\bar{h}(\bar{\mathbf{w}}(t))\| \leq \text{Lip}(Dh) \sqrt{tR_0}/\alpha$$

Proof. Since (a) \bar{h} is the linearization of h at \mathbf{w}_0 , and (b) $\|\mathbf{w}(t) - \mathbf{w}_0\| \leq \min(\rho, \sqrt{tR_0}/\alpha)$ by Lemma A.1 and Proposition 2.1,

$$\|Dh(\mathbf{w}(t)) - D\bar{h}(\bar{\mathbf{w}}(t))\| \stackrel{(a)}{=} \|Dh(\mathbf{w}(t)) - Dh(\mathbf{w}_0)\| \stackrel{(b)}{\leq} \text{Lip}(Dh) \sqrt{tR_0}/\alpha.$$

\square

This allows us to show that the kernel K_t is bounded at all times $t \in [0, T]$ in operator norm.

Lemma A.3. *For any $t \in [0, T]$, we have $\|K_t\| \leq 3\|Dh(\mathbf{w}(0))\|^2 + 2\text{Lip}(Dh)tR_0/\alpha^2$.*

Proof. By triangle inequality and Lemma A.2,

$$\begin{aligned} \|K_t - K_0\| &= \|Dh(\mathbf{w}(t))Dh(\mathbf{w}(t))^\top - Dh(\mathbf{w}(0))Dh(\mathbf{w}(0))^\top\| \\ &\leq \|Dh(\mathbf{w}(t))\|^2 + \|Dh(\mathbf{w}(0))\|^2 \\ &\leq (\|Dh(\mathbf{w}(t)) - Dh(\mathbf{w}(0))\| + \|Dh(\mathbf{w}(0))\|)^2 + \|Dh(\mathbf{w}(0))\|^2 \\ &\leq 3\|Dh(\mathbf{w}(0))\|^2 + 2\text{Lip}(Dh)tR_0/\alpha^2. \end{aligned}$$

\square

And finally we note that the kernel evolves continuously in time.

Lemma A.4. *The map $t \mapsto K_t$ is continuous (in the operator norm topology) in the interval $[0, T]$.*

Proof. First, $t \mapsto \mathbf{w}(t)$ is continuous in time, since it solves the gradient flow. Second, we know that $\mathbf{w}(t)$ is in the ball of radius ρ around \mathbf{w}_0 by Lemma A.1, and in this ball the map $\mathbf{w} \mapsto Dh(\mathbf{w})$ is continuous because $\text{Lip}(Dh) < \infty$. Finally, $Dh \mapsto DhDh^\top$ is continuous. \square

A.3 Product integral formulation of dynamics

Now we can present an equivalent formulation of the training dynamics (9) in terms of product integration. For any $0 \leq x \leq y \leq T$, let $P(y, x) : \mathbb{R}^p \rightarrow \mathcal{F}$ solve the operator integral equation

$$P(y, x) = I - \int_x^y K_t P(t, x) dt. \quad (10)$$

A solution $P(y, x)$ are guaranteed to exist and to be unique:

Lemma A.5. *The unique solution to the integral equation (10) is given as follows. For any $0 \leq x \leq y \leq T$ define $s_{m,j} = (y-x)(j/m) + x$ and $\delta = (y-x)/m$, and let $P(y, x)$ be the product integral*

$$P(y, x) := \prod_x^y e^{-K_s ds} := \lim_{m \rightarrow \infty} \prod_{j=1}^m e^{-\delta K_{s_j}} = \lim_{m \rightarrow \infty} e^{-\delta K_{s_m}} e^{-\delta K_{s_{m-1}}} \dots e^{-\delta K_{s_2}} e^{-\delta K_{s_1}}.$$

Proof. Existence, uniqueness, and the expression as an infinite product are guaranteed by Theorems 3.4.1, 3.4.2, and 3.5.1 of Dollard & Friedman (1984), since $t \mapsto K_t$ lies in $L_s^1(0, T)$, which is the space of “strongly integrable” functions on $[0, T]$ defined in Definition 3.3.1 of Dollard & Friedman (1984). This fact is guaranteed by the separability of \mathcal{F} and the continuity and boundedness of $t \mapsto K_t$ (Lemmas A.3 and A.4). \square

The operators $P(y, x)$ are the time-evolution operators corresponding to the differential equation (9) for the residual error r_t . Namely, for any time $0 \leq t \leq T$,

$$r_t = P(t, 0)r_0.$$

On the other hand, the solution to the linearized dynamics (9) is given by

$$\bar{r}_t = e^{-K_0 t} r_0,$$

since $e^{-K_0 t}$ is the time-evolution operator when the kernel does not evolve with time.

A.4 Error bound between product integrals

To prove Theorem 1.2, it suffices to bound $\|P(t, 0) - e^{-K_0 t}\|$, the difference of the time-evolution operators under the full dynamics versus the linearized dynamics. We will do this via a general theorem. To state it, we must define the total variation norm of a time-indexed sequence of operators:

Definition A.6. Let $\{C_t\}_{t \in [x, y]}$ be a sequence of time-bounded operators $C_t : \mathbb{R}^p \rightarrow \mathcal{F}$ so that $t \mapsto C_t$ is continuous in the interval $[x, y]$. Then the total variation norm of $\{C_t\}_{t \in [x, y]}$ is

$$\mathcal{V}(\{C_t\}_{t \in [x, y]}) = \sup_{P \in \mathcal{P}} \sum_{i=1}^{n_P-1} \|C_{t_i} - C_{t_{i-1}}\|,$$

where the supremum is taken over partitions $\mathcal{P} = \{P = \{x = t_1 \leq t_2 \leq \dots \leq t_{n_P-1} \leq t_{n_P} = y\}\}$ of the interval $[x, y]$.

We may now state the general result.

Theorem A.7. Let \mathcal{F} be a separable Hilbert space, and let $\{A_t\}_{t \in [0, T]}$, $\{B_t\}_{t \in [0, T]}$ be time-indexed sequences of bounded operators $A_t : \mathbb{R}^p \rightarrow \mathcal{F}$ and $B_t : \mathbb{R}^p \rightarrow \mathcal{F}$ such that $t \mapsto A_t$ and $t \mapsto B_t$ are continuous in $[0, T]$. Then,

$$\left\| \prod_0^T e^{-A_s A_s^\top} ds - \prod_0^T e^{-B_s B_s^\top} ds \right\| \leq \left(\sup_{t \in [0, T]} \|A_t - B_t\| \right) (2\sqrt{T} + 3T \cdot \mathcal{V}(\{A_t - B_t\}_{t \in [0, T]})).$$

If we can establish Theorem A.7, then we may prove Theorem 1.2 as follows.

Proof of Theorem 1.2. for each $t \geq 0$ we choose the linear operators $A_t, B_t : \mathbb{R}^p \rightarrow \mathcal{F}$ by $A_t = Dh(\mathbf{w}(t))$ and $B_t = Dh(\mathbf{w}(0))$ so that $A_t A_t^\top = K_t$ and $B_t B_t^\top = K_0$. We know that A_t, B_t are bounded by Lemma A.3, and that $t \mapsto A_t$ is continuous by Lemma A.4. (Also $t \mapsto B_t$ is trivially continuous). So we may apply Theorem A.7 to bound the difference in the residuals.

We first bound the total variation norm of $\{A_t - B_t\}_{t \in [0, T]}$. By (a) the fact from Lemma A.1 that $\mathbf{w}(t)$ is in a ball of radius at most ρ around \mathbf{w}_0 where $\mathbf{w} \mapsto Dh(\mathbf{w})$ is Lip(Dh)-Lipschitz; (b) the fact that $t \mapsto \mathbf{w}(t)$ is differentiable, since it solves a gradient flow; and (c) Proposition 2.1,

$$\begin{aligned} \mathcal{V}(\{A_t - B_t\}_{t \in [0, T]}) &= \sup_{P \in \mathcal{P}} \sum_{i=1}^{n_P-1} \|Dh(\mathbf{w}(t_{i+1})) - Dh(\mathbf{w}(0)) - Dh(\mathbf{w}(t_i)) + Dh(\mathbf{w}(0))\| \\ &= \sup_{P \in \mathcal{P}} \sum_{i=1}^{n_P-1} \|Dh(\mathbf{w}(t_{i+1})) - Dh(\mathbf{w}(t_i))\| \\ &\stackrel{(a)}{\leq} \text{Lip}(Dh) \sup_{P \in \mathcal{P}} \sum_{i=1}^{n_P-1} \|\mathbf{w}(t_{i+1}) - \mathbf{w}(t_i)\| \\ &\stackrel{(b)}{=} \text{Lip}(Dh) \int_0^T \left\| \frac{d\mathbf{w}}{dt} \right\| dt \\ &\stackrel{(c)}{\leq} \text{Lip}(Dh) \sqrt{TR_0} / \alpha \end{aligned}$$

When we plug the above expression into Theorem A.7, along with the bound $\|A_t - B_t\| \leq \text{Lip}(Dh)\sqrt{tR_0}/\alpha$ from Lemma A.2, we obtain:

$$\begin{aligned} \|r_T - \bar{r}_T\| &= \left\| \left(\prod_{s=0}^T e^{-K_s ds} - \prod_{s=0}^T e^{-K_0 ds} \right) r_0 \right\| \leq (\text{Lip}(Dh)\sqrt{TR_0}/\alpha)(2\sqrt{T} + 3\text{Lip}(Dh)T^{3/2}\sqrt{R_0}/\alpha)\sqrt{2R_0} \\ &= (2\kappa + 3\kappa^2)\sqrt{2R_0}, \end{aligned}$$

where $\kappa = \frac{T}{\alpha}\text{Lip}(Dh)\sqrt{R_0}$.

Also note that

$$\|r_T - \bar{r}_T\| \leq \|r_T\| + \|\bar{r}_T\| = \sqrt{2R(\alpha h(\mathbf{w}(T)))} + \sqrt{2R(\alpha \bar{h}(\bar{\mathbf{w}}(T)))} \leq 2\sqrt{2R_0},$$

since the gradient flow does not increase the risk. So

$$\|r_T - \bar{r}_T\| \leq \min(2\kappa + 3\kappa^2, 2)\sqrt{2R_0} \leq 6\kappa\sqrt{R_0}.$$

□

It remains only to prove Theorem A.7.

A.5 Reduction to the finite-dimensional case

We first show that in order to prove Theorem A.7, it suffices to consider the case where \mathcal{F} is a finite-dimensional Hilbert space. The argument is standard, and uses the fact that \mathcal{F} has a countable orthonormal basis.

Lemma A.8. *Theorem A.7 is true for general separable Hilbert spaces \mathcal{F} if it is true whenever \mathcal{F} is finite-dimensional.*

Proof. Suppose that \mathcal{F} is a separable Hilbert space and $A_t, B_t : \mathbb{R}^p \rightarrow \mathcal{F}$ satisfy the hypotheses of Theorem A.7. Let $\{f_i\}_{i \in \mathbb{N}}$ be a countable orthonormal basis for \mathcal{F} , which is guaranteed by separability of \mathcal{F} . For any n , let $\mathcal{P}_n : \mathcal{F} \rightarrow \mathcal{F}$ be the linear projection operator defined by

$$\mathcal{P}_n(f_i) = \begin{cases} f_i, & 1 \leq i \leq n \\ 0, & \text{otherwise} \end{cases}.$$

By (a) Duhamel's formula in Theorem 3.5.8 of [Dollard & Friedman \(1984\)](#), (b) the fact that $\|e^{-A_s A_s^\top}\| \leq 1$ and $\|e^{-\mathcal{P}_n A_s A_s^\top \mathcal{P}_n^\top}\| \leq 1$ because $A_s A_s^\top$ and $\mathcal{P}_n A_s A_s^\top \mathcal{P}_n^\top$ are positive semidefinite.

$$\begin{aligned} & \left\| \prod_0^T e^{-A_s A_s^\top ds} - \prod_0^T e^{-\mathcal{P}_n A_s A_s^\top \mathcal{P}_n^\top ds} \right\| \\ & \stackrel{(a)}{=} \left\| \int_0^T \left(\prod_\tau^T e^{-A_s A_s^\top ds} \right) (A_\tau A_\tau^\top - \mathcal{P}_n A_\tau A_\tau^\top \mathcal{P}_n^\top) \left(\prod_0^\tau e^{-\mathcal{P}_n A_s A_s^\top \mathcal{P}_n^\top ds} \right) d\tau \right\| \\ & \leq \int_0^T \left\| \prod_\tau^T e^{-A_s A_s^\top ds} \right\| \|A_\tau A_\tau^\top - \mathcal{P}_n A_\tau A_\tau^\top \mathcal{P}_n^\top\| \left\| \prod_0^\tau e^{-\mathcal{P}_n A_s A_s^\top \mathcal{P}_n^\top ds} \right\| d\tau \\ & \stackrel{(b)}{\leq} \int_0^T \|A_\tau A_\tau^\top - \mathcal{P}_n A_\tau A_\tau^\top \mathcal{P}_n^\top\| d\tau \end{aligned} \tag{11}$$

We have chosen \mathcal{P}_n so that for any bounded linear operator $M : \mathcal{F} \rightarrow \mathcal{F}$, we have $\lim_{n \rightarrow \infty} \mathcal{P}_n M \mathcal{P}_n = M$. Since $\tau \mapsto A_\tau A_\tau^\top$ is continuous in τ , and A_τ is bounded for each τ , the expression in (11) converges to 0 as $n \rightarrow \infty$. By triangle inequality, we conclude that

$$\left\| \prod_0^T e^{-A_s A_s^\top ds} - \prod_0^T e^{-B_s B_s^\top ds} \right\| \leq \limsup_{n \rightarrow \infty} \left\| \prod_0^T e^{-\mathcal{P}_n A_s A_s^\top \mathcal{P}_n^\top ds} - \prod_0^T e^{-\mathcal{P}_n B_s B_s^\top \mathcal{P}_n^\top ds} \right\|.$$

Notice that $\mathcal{P}_n A_t$ and $\mathcal{P}_n B_t$ are bounded maps from \mathbb{R}^p to $\text{span}\{f_1, \dots, f_n\}$, and $t \mapsto \mathcal{P}_n A_t$ and $t \mapsto \mathcal{P}_n B_t$ are continuous and bounded. So using the theorem in the case where \mathcal{F} is finite-dimensional, the right-hand side can be bounded by

$$\begin{aligned} & \limsup_{n \rightarrow \infty} \left\| \prod_0^T e^{-\mathcal{P}_n A_s A_s^\top \mathcal{P}_n^\top ds} - \prod_0^T e^{-\mathcal{P}_n B_s B_s^\top \mathcal{P}_n^\top ds} \right\| \\ & \leq \limsup_{n \rightarrow \infty} \left(\sup_{t \in [0, T]} \|\mathcal{P}_n A_t - \mathcal{P}_n B_t\| \right) (2\sqrt{T} + 3T \cdot \mathcal{V}(\{\mathcal{P}_n A_t - \mathcal{P}_n B_t\}_{t \in [0, T]})) dt \\ & = \left(\sup_{t \in [0, T]} \|A_t - B_t\| \right) (2\sqrt{T} + 3T \cdot \mathcal{V}(\{A_t - B_t\}_{t \in [0, T]})) dt. \end{aligned}$$

□

A.6 Bound for the finite-dimensional case

We conclude by proving Theorem A.7 in the finite-dimensional case, where we write $A_t, B_t \in \mathbb{R}^{n \times d}$ as time-indexed matrices. In order to prove a bound, we will interpolate between the dynamics we wish to compare. Let

$$C_t = B_t - A_t \in \mathbb{R}^{n \times d}$$

For any $\zeta \in [0, 1]$ and $t \in [0, T]$, we define the kernel

$$K_{t, \zeta} = (A_t + \zeta C_t)(A_t + \zeta C_t)^\top \in \mathbb{R}^{n \times n}.$$

This interpolates between the kernels of interest, since on the one hand, $K_{t, 0} = A_t A_t^\top$ and on the other $K_{t, 1} = B_t B_t^\top$. For any $x, y \in [0, T]$, let

$$P(y, x; \zeta) = \prod_x^y e^{-K_{t, \zeta} dt} \in \mathbb{R}^{n \times n}.$$

The derivative $\frac{\partial}{\partial \zeta} P(y, x; \zeta)$ exists by Theorem 1.5.3 of [Dollard & Friedman \(1984\)](#), since (i) $K_{t, \zeta}$ is continuous in t for each fixed ζ , (ii) $K_{t, \zeta}$ is differentiable in ζ in the L^1 sense⁴, and (iii) has a partial derivative $\frac{\partial K_{t, \zeta}}{\partial \zeta}$ that is integrable in t . The formula for $\frac{\partial}{\partial \zeta} P(y, x; \zeta)$ is given by the following formula, which generalizes the integral representation of the exponential map:⁵

$$\frac{\partial}{\partial \zeta} P(y, x; \zeta) = - \int_x^y P(y, t; \zeta) \frac{\partial K_{t, \zeta}}{\partial \zeta} P(t, x; \zeta) dt \quad (12)$$

Our main lemma is:

Lemma A.9. *For all $\zeta \in [0, 1]$, we have the bound*

$$\left\| \frac{\partial P(T, 0; \zeta)}{\partial \zeta} \right\| \leq \left(\sup_{t \in [0, T]} \|C_t\| \right) (2\sqrt{T} + 3T \cdot \mathcal{V}(\{C_t\}_{t \in [0, T]})).$$

This lemma suffices to prove Theorem A.7.

Proof of Theorem A.7. Using the fundamental theorem of calculus,

$$\left\| \prod_0^T e^{-A_t A_t^\top dt} - \prod_0^T e^{-B_t B_t^\top dt} \right\| = \|P(T, 0; 1) - P(T, 0; 0)\| \leq \int_0^1 \left\| \frac{\partial P(T, 0; \zeta)}{\partial \zeta} \right\| d\zeta,$$

which combined with Lemma A.9 proves Theorem A.7. □

⁴For any $\zeta \in [0, 1]$, $\lim_{\zeta' \rightarrow \zeta} \int_0^T \left\| \frac{K_{t, \zeta'} - K_{t, \zeta}}{\zeta' - \zeta} - \frac{\partial K_{t, \zeta}}{\partial \zeta} \right\| dt = 0$, since the matrices A_t, B_t are uniformly bounded.

⁵The proof is a consequence of Duhamel's formula. This a tool used in a variety of contexts, including perturbative analysis of path integrals in quantum mechanics [Dollard & Friedman \(1984\)](#).

A.7 Proof of Lemma A.9

By a direct calculation,

$$\frac{\partial K_{t,\zeta}}{\partial \zeta} = (A_t + \zeta C_t)C_t^\top + C_t(A_t + \zeta C_t)^\top,$$

so, by (12),

$$\begin{aligned} \frac{\partial}{\partial \zeta} P(T, 0; \zeta) &= - \int_0^T P(T, t; \zeta) ((A_t + \zeta C_t)C_t^\top + C_t(A_t + \zeta C_t)^\top) P(t, 0; \zeta) dt \\ &= - \underbrace{\int_0^T P(T, t; \zeta) (A_t + \zeta C_t)C_t^\top P(t, 0; \zeta) dt}_{M_1} - \underbrace{\int_0^T P(T, t; \zeta) C_t(A_t + \zeta C_t)^\top P(t, 0; \zeta) dt}_{M_2}. \end{aligned}$$

The arguments are similar for bounding M_1 and M_2 , so we only bound M_1 . We will need two technical bounds, whose proofs are deferred to Section A.8.

Claim 1. For any $0 \leq t \leq T$, we have $\|P(t, 0; \zeta)\| \leq 1$.

Claim 2. For any $0 \leq t \leq T$, we have $\|P(T, t; \zeta)(A_t + \zeta C_t)\| \leq \frac{1}{2\sqrt{T-t}} + 3\mathcal{V}(\{C_s\}_{s \in [t, T]})$.

Using (a) Claim 1, and (b) Claim 2,

$$\|M_1\| \stackrel{(a)}{\leq} \left(\sup_{t \in [0, T]} \|C_t\| \right) \int_0^T \|P(T, t; \zeta)(A_t + \zeta C_t)\| dt \quad (13)$$

$$\stackrel{(b)}{\leq} \left(\sup_{t \in [0, T]} \|C_t\| \right) \left(\int_0^T \frac{1}{2\sqrt{T-t}} + 3\mathcal{V}(\{C_s\}_{s \in [t, T]}) dt \right) \quad (14)$$

$$= \left(\sup_{t \in [0, T]} \|C_t\| \right) (\sqrt{T} + 3 \int_0^T \mathcal{V}(\{C_s\}_{s \in [t, T]}) dt). \quad (15)$$

Lemma A.9 is proved by noting that, symmetrically,

$$\|M_2\| \leq \left(\sup_{t \in [0, T]} \|C_t\| \right) (\sqrt{T} + 3 \int_0^T \mathcal{V}(\{C_s\}_{s \in [0, t]}) dt),$$

and, for any $t \in [0, T]$,

$$\mathcal{V}(\{C_s\}_{s \in [0, t]}) + \mathcal{V}(\{C_s\}_{s \in [t, T]}) = \mathcal{V}(\{C_s\}_{s \in [0, T]}).$$

A.8 Deferred proofs of Claims 1 and 2

Claim 3 (Deferred Claim 1). For any $0 \leq t \leq T$, we have $\|P(t, 0; \zeta)\| \leq 1$.

Proof. This follows from the definition of the product integral as an infinite product, and the fact that each term $e^{-\delta K_{t_i, \zeta}}$ in the product has norm at most 1 because $K_{t_i, \zeta}$ is positive semidefinite. \square

In order to prove Claim 2, we need two more claims:

Claim 4. For any $X \in \mathbb{R}^{n \times d}$ and $t \geq 0$,

$$\|e^{-X X^\top t} X\| \leq \frac{1}{2\sqrt{t}}$$

Proof. Since XX^\top is positive semidefinite,

$$\|e^{-XX^\top}X\| = \sqrt{\|e^{-XX^\top t}XX^\top e^{-XX^\top t}\|} \leq \sup_{\lambda \geq 0} \sqrt{e^{-\lambda t} \lambda e^{-\lambda t}} = \sup_{\lambda \geq 0} e^{-\lambda} \sqrt{\lambda/t} \leq \frac{1}{2\sqrt{t}}.$$

□

Claim 5. For any time-indexed sequence of matrices $(X_t)_{t \in [0, T]}$ in $\mathbb{R}^{n \times d}$ such that $t \mapsto X_t$ is continuous in $[0, T]$, and any $0 \leq a \leq b \leq T$,

$$\|e^{-X_b X_b^\top (b-a)} X_b - \left(\prod_a^b e^{-X_s X_s^\top ds} \right) X_a\| \leq 3\mathcal{V}(\{X_s\}_{s \in [a, b]})$$

This latter claim shows that we can approximate the product integral with an exponential. The proof is involved, and is provided in Section A.9. Assuming the previous two claims, we may prove Claim 2.

Claim 6 (Deferred Claim 2). For any $0 \leq t \leq T$, we have

$$\|P(T, t; \zeta)(A_t + \zeta C_t)\| \leq \frac{1}{2\sqrt{T-t}} + 3\mathcal{V}(\{C_s\}_{s \in [t, T]}).$$

Proof. By Claim 4, since $K_{T, \zeta} = (A_T + \zeta C_T)(A_T + \zeta C_T)^\top$,

$$\|e^{-K_{T, \zeta}(T-t)}(A_T + \zeta C_T)\| \leq \frac{1}{2\sqrt{T-t}}.$$

By the triangle inequality, it remains to prove

$$\|e^{-K_{T, \zeta}(T-t)}(A_T + \zeta C_T) - P(T, t; \zeta)(A_t + \zeta C_t)^\top\| \leq 3\mathcal{V}(\{C_s\}_{s \in [t, T]}),$$

and this is implied by Claim 5, defining $X_t = A_t + \zeta C_t$ and $a = t$, $b = T$. □

A.9 Deferred proof of Claim 5

The proof will be by interpolation, using the integral representation of the exponential map provided in (12), similarly to the main body of the proof, but of course interpolating with respect to a different parameter. We begin the following claim, which we will subsequently strengthen.

Claim 7. For any symmetric $X, Y \in \mathbb{R}^{n \times n}$ and $t \geq 0$,

$$\|e^{-(X+Y)^2 t}(X+Y) - e^{-X^2 t}X\| \leq 3\|Y\|$$

Proof. For any $\tau \in [0, 1]$, define $X(\tau) = X + \tau Y$, which interpolates between X and Y . Then by (a) the derivative of the exponential map in (12), and (b) the fact that $X'(\tau) = Y$ and $\|e^{-X^2(\tau)t/2}\| \leq 1$,

$$\begin{aligned}
& \|e^{-(X+Y)^2t}(X+Y) - e^{-X^2t}X\| \\
&= \left\| \int_0^1 \frac{d}{d\tau} \left(e^{-X^2(\tau)t} X(\tau) \right) d\tau \right\| \\
&= \left\| \int_0^1 \frac{d}{d\tau} \left(e^{-X^2(\tau)t/2} X(\tau) e^{-X^2(\tau)t/2} \right) d\tau \right\| \\
&\leq \sup_{\tau \in [0,1]} \left\| \frac{d}{d\tau} e^{-X^2(\tau)t/2} X(\tau) e^{-X^2(\tau)t/2} \right\| \\
&\stackrel{(a)}{=} \sup_{\tau \in [0,1]} \|e^{-X^2(\tau)t/2} X'(\tau) e^{-X^2(\tau)t/2} \\
&\quad - (t/2) \int_0^1 e^{-(1-s)X^2(\tau)t/2} (X'(\tau)X(\tau) + X(\tau)X'(\tau)) e^{-sX^2(\tau)t/2} X(\tau) e^{-X^2(\tau)t/2} ds \\
&\quad - (t/2) \int_0^1 e^{-X^2(\tau)t/2} X(\tau) e^{-(1-s)X^2(\tau)t/2} (X'(\tau)X(\tau) + X(\tau)X'(\tau)) e^{-sX^2(\tau)t/2} ds\| \\
&\leq \sup_{\tau \in [0,1]} \|Y\| + \underbrace{\left\| (t/2) \int_0^1 e^{-(1-s)X^2(\tau)t/2} (YX(\tau) + X(\tau)Y) e^{-sX^2(\tau)t/2} X(\tau) e^{-X^2(\tau)t/2} ds \right\|}_{T_1} \\
&\quad + \underbrace{\left\| (t/2) \int_0^1 e^{-X^2(\tau)t/2} X(\tau) e^{-(1-s)X^2(\tau)t/2} (YX(\tau) + X(\tau)Y) e^{-sX^2(\tau)t/2} ds \right\|}_{T_2},
\end{aligned}$$

and by (a) $\sup_{\lambda \geq 0} \lambda e^{-\lambda^2 t/2} = 1/\sqrt{et}$, and (b) $\|e^{-M}\| \leq 1$ if M is p.s.d.,

$$\begin{aligned}
T_1 &\leq (t/2) \int_0^1 \|e^{-(1-s)X^2(\tau)t/2} (YX(\tau) + X(\tau)Y) e^{-sX^2(\tau)t/2}\| \|X(\tau) e^{-X^2(\tau)t/2}\| ds \\
&\stackrel{(a)}{\leq} (t/2) \int_0^1 \|e^{-(1-s)X^2(\tau)t/2} (YX(\tau) + X(\tau)Y) e^{-sX^2(\tau)t/2}\| \frac{1}{\sqrt{et}} ds \\
&\leq \frac{\sqrt{t}}{2\sqrt{e}} \int_0^1 \|e^{-(1-s)X^2(\tau)t/2}\| \|Y\| \|X(\tau) e^{-sX^2(\tau)t/2}\| + \|e^{-(1-s)X^2(\tau)t/2} X(\tau)\| \|Y\| \|e^{-sX^2(\tau)t/2}\| ds \\
&\stackrel{(a)}{\leq} \frac{\sqrt{t}}{2e} \int_0^1 \|e^{-(1-s)X^2(\tau)t/2}\| \|Y\| \frac{1}{\sqrt{st}} + \frac{1}{\sqrt{(1-s)t}} \|Y\| \|e^{-sX^2(\tau)t/2}\| ds \\
&\stackrel{(b)}{\leq} \frac{\|Y\|}{2e} \int_0^1 \frac{1}{\sqrt{s}} + \frac{1}{\sqrt{(1-s)}} ds \\
&= \frac{\|Y\|}{2e} (2\sqrt{s} - 2\sqrt{1-s})_0^1 \\
&\leq \|Y\|.
\end{aligned}$$

Similarly, $T_2 \leq \|Y\|$. □

Claim 8. For any $X, Y \in \mathbb{R}^{n \times d}$ (not necessarily symmetric) and $t \geq 0$,

$$\|e^{-(X+Y)(X+Y)^\top t}(X+Y) - e^{-XX^\top t}X\| \leq 3\|Y\|.$$

Proof. We will use Claim 7, combined with a general method for lifting facts from symmetric matrices to asymmetric matrices (see, e.g., Tao (2011) for other similar arguments). Assume without loss of generality that $n = d$, since otherwise we can pad with zeros. Define $\bar{n} = 2n$ and

$$\bar{X} = \begin{bmatrix} 0 & X^\top \\ X & 0 \end{bmatrix} \in \mathbb{R}^{\bar{n} \times \bar{n}} \text{ and } \bar{Y} = \begin{bmatrix} 0 & Y^\top \\ Y & 0 \end{bmatrix} \in \mathbb{R}^{\bar{n} \times \bar{n}}.$$

These are symmetric matrices by construction. Furthermore,

$$\bar{X}^2 = \begin{bmatrix} X^\top X & 0 \\ 0 & XX^\top \end{bmatrix} \text{ and } (\bar{X} + \bar{Y})^2 = \begin{bmatrix} (X+Y)^\top(X+Y) & 0 \\ 0 & (X+Y)(X+Y)^\top \end{bmatrix}.$$

Because of the block-diagonal structure of these matrices,

$$e^{-\bar{X}^2 t} = \begin{bmatrix} e^{-X^\top X t} & 0 \\ 0 & e^{-XX^\top t} \end{bmatrix} \text{ and } e^{-(\bar{X} + \bar{Y})^2 t} = \begin{bmatrix} e^{-(X+Y)^\top(X+Y)t} & 0 \\ 0 & e^{-(X+Y)(X+Y)^\top t} \end{bmatrix}.$$

So

$$e^{-\bar{X}^2 t} \bar{X} = \begin{bmatrix} 0 & e^{-X^\top X t} X^\top \\ e^{-XX^\top t} X & 0 \end{bmatrix}.$$

Similarly,

$$e^{-(\bar{X} + \bar{Y})^2 t} (\bar{X} + \bar{Y}) = \begin{bmatrix} 0 & e^{-(X+Y)^\top(X+Y)t} (X+Y)^\top \\ e^{-(X+Y)(X+Y)^\top t} (X+Y) & 0 \end{bmatrix}.$$

For any matrix $M \in \mathbb{R}^{n \times n}$, we have $\|M\| = \sup_{\mathbf{v} \in \mathbb{S}^{n-1}} \|M\mathbf{v}\|$. So for any matrices $M_1, M_2 \in \mathbb{R}^{n \times n}$, we have

$$\left\| \begin{bmatrix} 0 & M_1 \\ M_2 & 0 \end{bmatrix} \right\| \geq \sup_{\mathbf{v} \in \mathbb{S}^{n-1}} \left\| \begin{bmatrix} 0 & M_1 \\ M_2 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ 0 \end{bmatrix} \right\| = \sup_{\mathbf{v} \in \mathbb{S}^{n-1}} \|M_2 \mathbf{v}\| = \|M_2\|.$$

This means that, using Claim 7,

$$\|e^{-(X+Y)(X+Y)^\top t} (X+Y) - e^{-XX^\top t} X\| \leq \|e^{-(\bar{X} + \bar{Y})^2 t} (\bar{X} + \bar{Y}) - e^{\bar{X}^2 t} \bar{X}\| \leq 3\|\bar{Y}\|.$$

Finally, using the symmetry of \bar{Y} and the block-diagonal structure of \bar{Y}^2 ,

$$\|\bar{Y}\| = \|\bar{Y}^2\|^{1/2} = \left\| \begin{bmatrix} Y^\top Y & 0 \\ 0 & YY^\top \end{bmatrix} \right\|^{1/2} = \max(\|Y^\top Y\|^{1/2}, \|YY^\top\|^{1/2}) = \|Y\|.$$

□

We conclude by using these results to prove Claim 5 with a telescoping argument.

Claim 9 (Restatement of Claim 5). *For any time-indexed sequence of matrices $(X_t)_{t \in [0, T]}$ in $\mathbb{R}^{n \times d}$ such that $t \mapsto X_t$ is continuous in $[0, T]$, and any $0 \leq a \leq b \leq T$,*

$$\|e^{-X_b X_b^\top (b-a)} X_b - \left(\prod_a^b e^{-X_s X_s^\top ds} \right) X_a\| \leq 3\mathcal{V}(\{X_s\}_{s \in [a, b]})$$

Proof. We will do this by approximating the product integral by a finite product of m matrices, and taking the limit $m \rightarrow \infty$. For any $m \geq 0$ and $j \in \{0, \dots, m\}$, and $t \in [0, T]$, let $t_{m,j} = (b-a)(j/m) + a$, and define the finite-product approximation to the product integral for any $k \in \{0, \dots, m\}$

$$P_{m,k} = \left(\prod_{j=k+1}^m Q_{m,j} \right) Q_{m,k}^k,$$

where

$$Q_{m,j} = e^{-X_{t_{m,j}} X_{t_{m,j}}^\top (b-a)/m}.$$

Notice that for $k = 0$ we have

$$P_{m,0} = \left(\prod_{j=1}^m e^{-X_{t_{m,j}} X_{t_{m,j}}^\top (b-a)/m} \right),$$

and so by continuity of $s \mapsto X_s$ for $s \in [0, T]$ and boundedness of X_s , we know from Theorem 1.1 of [Dollard & Friedman \(1984\)](#) that

$$\lim_{m \rightarrow \infty} P_{m,0} = \prod_a^b e^{-X_s X_s^\top ds}.$$

This means that

$$\begin{aligned} \|e^{-X_b X_b^\top (b-a)} X_b - \left(\prod_a^b e^{-X_s X_s^\top ds} \right) X_a\| &= \|P_{m,m} X_{t_{m,m}} - \left(\prod_a^b e^{-X_s X_s^\top ds} \right) X_a\| \\ &= \lim_{m \rightarrow \infty} \|P_{m,m} X_{t_{m,m}} - P_{m,0} X_{t_{m,0}}\|. \end{aligned}$$

Finally, telescoping and (a) using $\|Q_{m,j}\| \leq 1$ for all j , and (b) Claim 8,

$$\begin{aligned} \|P_{m,m} X_{t_{m,m}} - P_{m,1} X_{t_{m,1}}\| &\leq \sum_{k=1}^m \|P_{m,k} X_{t_{m,k}} - P_{m,k-1} X_{t_{m,k-1}}\| \\ &= \sum_{k=1}^m \left\| \left(\prod_{j=k+1}^m Q_{m,j} \right) (Q_{m,k}^k X_{t_{m,k}} - Q_{m,k} (Q_{m,k-1})^{k-1} X_{t_{m,k-1}}) \right\| \\ &\stackrel{(a)}{\leq} \sum_{k=1}^m \|(Q_{m,k})^{k-1} X_{t_{m,k}} - (Q_{m,k-1})^{k-1} X_{t_{m,k-1}}\| \\ &\stackrel{(b)}{\leq} \sum_{k=1}^m 3 \|X_{t_{m,k}} - X_{t_{m,k-1}}\| \\ &\leq 3\mathcal{V}(\{X_s\}_{s \in [a,b]}). \end{aligned}$$

The lemma follows by taking $m \rightarrow \infty$, since the bound is independent of m . \square

B Proof of Theorem 1.3

The converse bound in Theorem 1.2 is achieved in the simple case where $h : \mathbb{R} \rightarrow \mathbb{R}$ is given by $h(w) = aw + \frac{1}{2}bw^2$ for $a = \frac{1}{\sqrt{T}}$ and $b = \text{Lip}(Dh)$. We also let $w_0 = 0$ and $y^* = \sqrt{2R_0}$ so that all conditions are satisfied. The evolution of the residuals $r(t) = y^* - \alpha h(w(t))$ and $\bar{r}(t) = y^* - \alpha \bar{h}(\bar{w}(t))$ is given by

$$\frac{dr}{dt} = -K_t r \quad \text{and} \quad \frac{d\bar{r}}{dt} = -K_0 \bar{r},$$

where $K_t = Dh(w(t))Dh(w(t))^\top = (a + bw(t))^2$ and $K_0 = a^2$. Since $r = y^* - \alpha(aw + \frac{1}{2}w^2)$, we can express the evolution of the residuals r and \bar{r} as:

$$\frac{dr}{dt} = -(a^2 + 2b(y^* - r)/\alpha)r \quad \text{and} \quad \frac{d\bar{r}}{dt} = -a^2 \bar{r}. \quad (16)$$

Since $b(y^* - r)/\alpha \geq 0$ at all times, we must have $a^2 + 2b(y^* - r)/\alpha \geq a^2$. This means that at all times

$$r(t) \leq \bar{r}(t) = e^{-a^2 t} y^*.$$

So, at any time $t \geq T/2$,

$$r(t) \leq r(T/2) \leq \bar{r}(T/2) = e^{-(1/\sqrt{T})^2(T/2)} y^* = e^{-1/2} y^*.$$

By plugging this into (16), for times $t \geq T/2$,

$$\frac{dr}{dt} \leq -(a^2 + 2\text{Lip}(Dh)(1 - e^{-1/2})\sqrt{2R_0}/\alpha)r \leq -(a^2 + 1.1\kappa/T)r.$$

So, at time T , assuming that $\kappa \leq 1$ without loss of generality,

$$r(T) \leq r(T/2)e^{-(1/T+1.1\kappa/T)(T/2)} = r(T/2)e^{-1/2-0.55\kappa} \leq y^*e^{-1}e^{-0.55\kappa} \leq y^*e^{-1}(1 - 0.4\kappa).$$

So

$$|\alpha h(w(T)) - \alpha \bar{h}(\bar{w}(T))| = |r(T) - \bar{r}(T)| \geq |e^{-1}y^* - (1 - 0.4\kappa)e^{-1}y^*| \geq 0.4\kappa e^{-1}\sqrt{2R_0} \geq \sqrt{R_0}/5.$$

□