

---

# Neural Tangent Kernel Beyond the Infinite-Width Limit: Effects of Depth and Initialization

---

Mariia Seleznova<sup>1</sup> Gitta Kutyniok<sup>1</sup>

## Abstract

Neural Tangent Kernel (NTK) is widely used to analyze overparametrized neural networks due to the famous result by Jacot et al. (2018): in the infinite-width limit, the NTK is deterministic and constant during training. However, this result cannot explain the behavior of deep networks, since it generally does not hold if depth and width tend to infinity simultaneously. In this paper, we study the NTK of fully-connected ReLU networks with depth comparable to width. We prove that the NTK properties depend significantly on the depth-to-width ratio and the distribution of parameters at initialization. In fact, our results indicate the importance of the three phases in the hyperparameter space identified in Poole et al. (2016): *ordered*, *chaotic* and the *edge of chaos* (EOC). We derive exact expressions for the NTK dispersion in the infinite-depth-and-width limit in all three phases and conclude that the NTK variability grows exponentially with depth at the EOC and in the chaotic phase but not in the ordered phase. We also show that the NTK of deep networks may stay constant during training only in the ordered phase and discuss how the structure of the NTK matrix changes during training.

## 1. Introduction

Despite the widespread use of Deep Neural Networks (DNNs), the theory behind their success is still poorly understood. For instance, no present theory can explain why highly overparametrized DNNs generalize very well in practice, contrary to classical statistical learning theory predictions. Likewise, it is surprising that optimizing a highly non-convex loss function of a DNN with a variant of Gradient

---

<sup>1</sup>Department of Mathematics, Ludwig-Maximilians-Universität München, Munich, Germany. Correspondence to: Mariia Seleznova <selez@math.lmu.de>.

Descent (GD) typically yields a good local minimum.

Although training dynamics and generalization capabilities of DNNs stand among the biggest open problems of deep learning theory, it is possible to address these challenges in the special case of infinitely-wide DNNs using the so-called *Neural Tangent Kernel* (NTK). This kernel captures the first-order approximation of DNN's evolution during GD training. Consider the gradient flow dynamics of the DNN's parameters:

$$\dot{\mathbf{w}} = -\nabla_{\mathbf{w}}\mathcal{L}(\mathcal{D}) = -\sum_{(x_i, y_i) \in \mathcal{D}} \nabla_{\mathbf{w}}f(x_i) \frac{\partial \mathcal{L}(\mathcal{D})}{\partial f(x_i)}, \quad (1)$$

where  $\mathbf{w}$  is the vector of all the trainable parameters,  $f(\cdot)$  is the DNN's output function (defined in Section 2),  $\mathcal{L}(\cdot)$  is the loss function and  $\mathcal{D}$  is the dataset. Then the dynamics of the DNN's output function is given by:

$$\dot{f}(x) = \nabla_{\mathbf{w}}f(x) \cdot \dot{\mathbf{w}} = -\sum_{(x_i, y_i) \in \mathcal{D}} \Theta(x, x_i) \frac{\partial \mathcal{L}(\mathcal{D})}{\partial f(x_i)}, \quad (2)$$

where the kernel  $\Theta(x_i, x_j) := \langle \nabla_{\mathbf{w}}f(x_i), \nabla_{\mathbf{w}}f(x_j) \rangle$  is called the NTK.

A famous result by Jacot et al. (2018) states that in the infinite-width limit, the NTK is deterministic under proper random initialization and stays constant during training. Thereby, the dynamics in (2) is equivalent to kernel regression and has an analytical solution expressed in terms of the kernel. It is then possible to derive properties of trained infinitely-wide DNNs theoretically by means of their NTKs. Hence, many recent works used the NTK to explain empirically known properties of DNNs (Huang et al., 2020; Adlam & Pennington, 2020; Wang et al., 2022; Tirer et al., 2021; Geiger et al., 2019). Numerous contributions also derived the infinite-width limit of the NTK for popular DNN architectures (Yang, 2020; Du et al., 2019; Alemohammad et al., 2021). Other papers established some non-asymptotic results on the concentration of the NTK at initialization (Arora et al., 2019; Buchanan et al., 2021) and stability of the NTK during training (Huang & Yau, 2020; Lee et al., 2019).

However, the extent to which the results in the infinite-width limit extrapolate to realistic DNNs remains largely

an open question. Indeed, multiple authors have argued that the NTK regime and, in general, the infinite-width limit cannot explain the success of DNNs (Chizat et al., 2019; Hanin & Nica, 2020; Aitchison, 2020; Li et al., 2021; Seleznova & Kutyniok, 2021; Bai et al., 2020; Huang & Yau, 2020). The first argument in this direction is that no feature learning occurs if the NTK stays constant during training. Moreover, several works showed that the infinite-width limit of the NTK becomes completely data-independent as depth increases (Xiao et al., 2020; Hayou et al., 2019), which suggests poor generalization performance for deep networks in the NTK regime. Finally, numerous empirical results demonstrated that the performance of trained DNNs and the corresponding kernel methods often differs in practice (Fort et al., 2020; Lee et al., 2020; Arora et al., 2020). That is why it is essential to understand the statistical properties of the NTK and how they depend on the myriad of settings of a given DNN to assess if the infinite-width limit provides a reasonable approximation for this network. We contribute to this line of research by exploring the combined effect of two factors on the NTK: the network’s depth and initialization hyperparameters.

**Network’s depth** Most results on the NTK are derived in the setting where the network’s depth is kept constant while the width tends to infinity. This limit can only model very wide and shallow networks since the depth-to-width ratio tends to zero in it. Indeed, several recent papers demonstrated that infinite-width approximations often get worse as the depth increases (Li et al., 2021; de G. Matthews et al., 2018; Yang & Schoenholz, 2017). In particular, Hanin & Nica (2020) first showed that the NTK of fully-connected ReLU DNNs may be random and change during training if depth and width are comparable. Hu & Huang (2021) also studied the effects of depth on the NTK distribution and derived an upper bound for the NTK moments. We expand on these results by precisely characterizing the variability of the NTK at initialization and generalizing to different initialization settings described below.

**Initialization hyperparameters** There are three phases in the initialization hyperparameter space where the properties of untrained infinitely-wide DNNs differ significantly: *ordered*, *chaotic* and the *edge of chaos* (EOC) (Poole et al., 2016). In the ordered phase, the gradient norms decrease with depth, whereas in the chaotic phase the gradient norms increase, and the edge of chaos is the initialization at the border between these two phases (Schoenholz et al., 2017). The results by Hanin & Nica (2020) concerned the statistical properties of the NTK of wide and deep ReLU networks at the EOC. At the same time, several contributions demonstrated that the properties of the infinite-width NTK depend significantly on the phase of initialization (Xiao et al., 2020; Hayou et al., 2019). However, these results do not apply to networks with depth comparable to width since they assume

infinite width before considering the effects of growing depth. We fill this gap by deriving statistical properties of the NTK for wide and deep ReLU networks in all three phases of initialization.

### 1.1. Contributions

We study the **variability of the NTK at initialization** for fully-connected ReLU DNNs with depth comparable to width and varying initialization hyperparameters in Section 3. Our contributions are as follows:

- We precisely characterize the dispersion of the diagonal elements  $\Theta(x, x)$  of the NTK (for arbitrary input  $x$ ) in the **infinite-depth-and-width limit** and conclude that the variability of the NTK grows exponentially with the depth-to-width ratio at the EOC and in the chaotic phase. Conversely, the variance of  $\Theta(x, x)$  tends to zero in the same limit in the ordered phase. Our results allow to evaluate the variance of the NTK for a given DNN with any depth-to-width ratio and initialization.
- We provide non-asymptotic expressions for the first two moments of  $\Theta(x, x)$  and illustrate **finite-width effects** that follow. We show that the variance of the finite-width NTK in the ordered phase gradually increases as the initialization approaches the EOC, which describes the transition between the two kinds of behavior in the limit. We also notice that the NTK dispersion depends on the architecture, i.e. on the varying widths of the fully-connected layers. Notably, the dispersion of  $\Theta(x, x)$  decreases with depth in the ordered phase if the DNN increases the dimensionality in consecutive layers. This enables us to conclude that deeper networks are more robust to random initialization in this setting.
- We lower-bound the ratio of the expected **non-diagonal elements** of the NTK, i.e.  $\Theta(x, \tilde{x})$  with  $x \neq \tilde{x}$ , and the diagonal elements  $\Theta(x, x)$  in the infinite-depth-and-width limit. We also upper-bound the dispersion of the non-diagonal elements. In the ordered phase, our results allow to ensure that the whole NTK matrix is approximately deterministic and thus can be approximated by the infinite-width limit.
- We provide extensive **numerical experiments** to verify our theoretical results. We use JAX (Bradbury et al., 2018) and Flax (neural network library for JAX) (Heek et al., 2020) to compute the NTK of fully-connected ReLU networks effortlessly. Source code to reproduce the presented results is available at: [https://github.com/mselezniova/ntk\\_beyond\\_limit](https://github.com/mselezniova/ntk_beyond_limit).

We study the **training dynamics of the NTK** for fully-connected ReLU DNNs with depth comparable to width and varying initialization hyperparameters in Section 4. Our contributions are as follows:

- We show that the expected relative change of  $\Theta(x, x)$  in **the first GD step** tends to infinity in the chaotic phase and to zero in the ordered phase in the infinite-depth-and-width limit. Combined with the result by Hanin & Nica (2020), which states that the expected relative change of  $\Theta(x, x)$  in the first GD step is exponential in the depth-to-width ratio at the EOC, we can conclude that the NTK of deep networks can stay approximately constant during GD training only in the ordered phase.
- We discuss how the **structure of the NTK** matrix changes during training outside of the NTK regime. The NTK matrix at initialization has a diagonal structure with larger values on the main diagonal as compared to the non-diagonal ones. We speculate that the training process changes the NTK structure to block-diagonal with blocks of larger values corresponding to classes and provide experiments to support this sentiment.

## 2. Preliminaries

We consider fully-connected ReLU DNNs of depth  $L \in \mathbb{N}$  with linear output layer and widths  $(n_\ell)_{0 \leq \ell \leq L}$ , where  $n_0 \in \mathbb{N}$  is the input dimension and  $n_L = 1$  is the output dimension. Forward propagation in such a network is defined as follows:

$$\begin{aligned} \mathbf{x}^\ell(x) &:= \phi(\mathbf{h}^\ell(x)), \quad \mathbf{x}^0(x) := x \in \mathbb{R}^{n_0}, \\ \mathbf{h}^\ell(x) &:= \mathbf{W}^\ell \mathbf{x}^{\ell-1}(x) + \mathbf{b}^\ell, \quad 1 \leq \ell \leq L-1, \\ f(x) &:= \mathbf{W}^L \mathbf{x}^{L-1}(x) + \mathbf{b}^L \in \mathbb{R}, \end{aligned} \quad (3)$$

where  $\phi(x) := x \mathbb{1}_{\{x > 0\}}$  denotes the ReLU function,  $\mathbf{W}^\ell \in \mathbb{R}^{n_\ell \times n_{\ell-1}}$  and  $\mathbf{b}^\ell \in \mathbb{R}^{n_\ell}$  are the weights and the biases and  $f(x)$  is the output function of the DNN. The NTK of this network on a pair of inputs  $(x, \tilde{x})$  is given by:

$$\begin{aligned} \Theta(x, \tilde{x}) &:= \Theta_W(x, \tilde{x}) + \Theta_b(x, \tilde{x}), \\ \Theta_W(x, \tilde{x}) &:= \sum_{\ell=1}^L \sum_{j=1}^{n_\ell} \sum_{i=1}^{n_{\ell-1}} \frac{\partial f(x)}{\partial \mathbf{W}_{ij}^\ell} \frac{\partial f(\tilde{x})}{\partial \mathbf{W}_{ij}^\ell}, \\ \Theta_b(x, \tilde{x}) &:= \sum_{\ell=1}^L \sum_{j=1}^{n_\ell} \frac{\partial f(x)}{\partial \mathbf{b}_j^\ell} \frac{\partial f(\tilde{x})}{\partial \mathbf{b}_j^\ell}, \end{aligned} \quad (4)$$

where  $\Theta_W(x, \tilde{x})$  comprises the gradients w.r.t. the weights and  $\Theta_b(x, \tilde{x})$  — the gradients w.r.t. the biases.

When we consider wide networks with unequal widths in the hidden layers, we define a width scale parameter  $M$  and constants  $\lambda, (\alpha_\ell)_{0 \leq \ell \leq L-1}$  such that:

$$\frac{L}{M} = \lambda \in \mathbb{R}, \quad \frac{n_\ell}{M} = \alpha_\ell \in \mathbb{R}, \quad 0 \leq \ell \leq L-1. \quad (5)$$

Then we can describe the asymptotic behavior of the NTK in terms of  $M$  and the constants defined above.

### 2.1. Initialization and parametrization

We consider random i.i.d. initialization given by:

$$\mathbf{W}_{ij}^\ell \sim \mathcal{N}\left(0, \frac{\sigma_w^2}{n_{\ell-1}}\right), \quad \mathbf{b}_i^\ell \sim \mathcal{N}(0, \sigma_b^2), \quad (6)$$

where  $(\sigma_w, \sigma_b)$  are the initialization hyperparameters. This initialization corresponds to the so-called standard parametrization (SP), where the weights and the biases defined in (3) are the trainable parameters. We note that the NTK is often considered in the so-called NTK parametrization (NTP), where the weights in (3) are the scaled versions of trainable parameters:  $\mathbf{W}_{ij}^\ell = \sigma_w / \sqrt{n_{\ell-1}} \mathbf{w}_{ij}^\ell$  for trainable  $\mathbf{w}^\ell \in \mathbb{R}^{n_\ell \times n_{\ell-1}}$  initialized as  $\mathbf{w}_{ij}^\ell \sim \mathcal{N}(0, 1)$  i.i.d. This reparametrization, of course, does not change the distribution of the DNN's components. However, it scales the gradients by  $O(1/M)$  and gives the NTK a well-defined infinite-width limit for fixed  $L$ . At the same time, NTP is equivalent to setting an individual learning rate in each layer inverse-proportionally to width, as explained, e.g., in Yang & Hu (2021). In this paper, we focus on the NTK in SP since this parametrization is more common in practice (indeed, SP is the default setting in PyTorch). However, our results can be generalized to NTP straightforwardly.

### 2.2. Information propagation in DNNs

Results on information propagation in infinitely-wide DNNs established that the initialization hyperparameters  $(\sigma_w, \sigma_b)$  determine the evolution of the variances  $\mathbb{E}[(\mathbf{x}_i^\ell(x))^2]$  and the covariances  $\mathbb{E}[\mathbf{x}_i^\ell(x) \mathbf{x}_i^\ell(\tilde{x})]$  as they propagate through the DNN's layers. Based on this, Poole et al. (2016) identified three phases with distinct properties in the hyperparameter space: *ordered*, *chaotic* and the *edge of chaos* (EOC). Schoenholz et al. (2017) subsequently showed that the ordered phase corresponds to vanishing gradients and the chaotic phase corresponds to exploding gradients, i.e. the gradient norm decreases with depth in the ordered phase and increases in the chaotic phase. The edge of chaos is the initialization at the border between these two phases, which allows deeper signal propagation through a DNN by avoiding vanishing or exploding gradients. Consider backpropagation equations given by:

$$\begin{aligned} \frac{\partial f(x)}{\partial \mathbf{W}_{ij}^\ell} &= \delta_i^\ell \mathbf{x}_j^{\ell-1}, \quad \frac{\partial f(x)}{\partial \mathbf{b}_i^\ell} = \delta_i^\ell, \\ \delta_i^\ell &:= \frac{\partial f(x)}{\partial \mathbf{h}_i^\ell} = \phi'(\mathbf{h}_i^\ell) \sum_j \delta_j^{\ell+1} \mathbf{W}_{ji}^{\ell+1}, \end{aligned} \quad (7)$$

Schoenholz et al. (2017) studied the evolution of  $\mathbb{E}[(\delta_i^\ell)^2]$  along with  $\mathbb{E}[(\mathbf{x}_i^\ell)^2]$  to find the distribution of DNNs' gradients. Some recent publications used these results to derive the properties of the infinite-width NTK in all the three phases of initialization (Karakida et al., 2019; Xiao et al.,




Figure 1. Ratio  $\mathbb{E}[\Theta^2(x, x)]/\mathbb{E}^2[\Theta(x, x)]$  at initialization for fully-connected ReLU networks of constant width  $M \in \{100, 200, 500\}$  with  $\alpha_0 = 1$ . The dashed lines represent the experimental results and the solid lines correspond to the theoretical predictions from Theorem 3.1. For each DNN configuration, we sampled 500 random initializations and computed an unbiased estimator for the ratio (see details in Appendix C.3). The error bars (indicated by the vertical lines) show the bootstrap estimation of the standard error (only in a subset of points to keep the figure readable). We provide additional figures with continuous error bars in Appendix C.2.

2020). Hayou et al. (2019) also showed that the infinite-depth limit of the infinite-width NTK (when first the limit  $M \rightarrow \infty$  is taken with fixed  $L$  and then  $L \rightarrow \infty$ ) yields a data-independent kernel and thus cannot explain properties of finite DNNs. Although our approach is different from the mentioned results since we do not assume infinite width before increasing depth, we show that the statistical properties of  $\delta^\ell$  and  $\mathbf{x}^\ell$  can still be derived and lead to results on the NTK in our setting.

The initialization hyperparameters that comprise each phase differ depending on the chosen activation function. Since we are interested in ReLU networks, we note that the ordered phase corresponds to  $\sigma_w^2 < 2$  and the chaotic phase — to  $\sigma_w^2 > 2$  for this activation function. The EOC is the initialization with  $\sigma_w^2 = 2$ . We refer, e.g., to Schoenholz et al. (2017) for a method to compute the border between phases for a given activation function.

### 3. Variability of the NTK

In the infinite-width limit, the NTK is deterministic under random initialization, which is one of the main results of the NTK theory. We investigate when this result holds outside of the NTK limit and, consequently, when the infinite-width behavior of the NTK gives a good approximation for realistic DNNs.

#### 3.1. Infinite-depth-and-width limit

Most results on the NTK assume that the network’s depth is fixed as the width tends to infinity, i.e.,  $L/M \rightarrow 0$  in the limit. This setting, of course, does not describe deep finite-width networks since their depth-to-width ratio is bounded away from zero. Indeed, some recent works demonstrated that infinite-width approximations often get worse as the

network’s depth increases (Li et al., 2021; de G. Matthews et al., 2018; Yang & Schoenholz, 2017). In particular, Hanin & Nica (2020) considered this effect for the NTK and derived bounds for the ratio  $\mathbb{E}[\Theta^2(x, x)]/\mathbb{E}^2[\Theta(x, x)]$  in case of ReLU DNNs initialized at the EOC ( $\sigma_w = 2$ ). This ratio characterizes the dispersion of the NTK: it is close to one if the NTK is approximately deterministic and is larger than two if the NTK’s distribution is of high variance. Our first main result characterizes this ratio in the infinite-depth-and-width limit under different initializations:

**Theorem 3.1** (Dispersion of the NTK at initialization in the limit). *Consider a ReLU DNN as defined in (3) with constant width of hidden layers  $M \in \mathbb{N}$ , input dimension  $n_0 = \alpha_0 M$ ,  $\alpha_0 \in \mathbb{R}$  and output dimension  $n_L = 1$ . The initialization is given by (6) and the biases are initialized to zero, i.e.  $\sigma_b = 0$ . Then, in the infinite-depth-and-width limit  $M \rightarrow \infty$ ,  $L \rightarrow \infty$ ,  $L/M \rightarrow \lambda \in \mathbb{R}$ , the following holds for the dispersion of the NTK:*

1. In the **chaotic phase** ( $a := \sigma_w^2/2 > 1$ ), the NTK dispersion grows exponentially with depth-to-width ratio  $\lambda := L/M$  as follows:

$$\frac{\mathbb{E}[\Theta^2(x, x)]}{\mathbb{E}^2[\Theta(x, x)]} \rightarrow \frac{1}{2\lambda} e^{5\lambda} \left( 1 - \frac{1}{4\lambda} (1 - e^{-4\lambda}) \right). \quad (8)$$

2. At the **EOC** ( $a = 1$ ), the NTK dispersion grows exponentially with depth-to-width ratio  $\lambda$  as well, but with a slower rate given by:

$$\begin{aligned} \frac{\mathbb{E}[\Theta^2(x, x)]}{\mathbb{E}^2[\Theta(x, x)]} &\rightarrow \frac{1}{(1 + \alpha_0)^2} \left[ e^{5\lambda} \left( \frac{1}{2\lambda} + \frac{2\alpha_0^2 - 8\alpha_0}{25\lambda^2} \right) \right. \\ &\quad \left. + (e^\lambda - e^{5\lambda}) \frac{1 - 4\alpha_0}{8\lambda^2} + \frac{2\alpha_0}{5\lambda} \left( \frac{4 - \alpha_0}{5\lambda} - 1 - \alpha_0 \right) \right]. \end{aligned} \quad (9)$$




Figure 2. Ratio  $\mathbb{E}[\Theta^2(x, x)]/\mathbb{E}^2[\Theta(x, x)]$  at initialization for fully-connected ReLU networks of constant width  $M = 200$  with the ratio  $\alpha_0 := n_0/M \in \{2.0, 0.5, 0.1\}$ . The initialization hyperparameter  $\sigma_w^2$  is close to the EOC for all the lines. The dashed lines represent the experimental results (computed as described in Figure 1) and the solid lines show the theoretical predictions given by Theorem 3.2. The error bars are shown only for a subset of points to keep the figure readable. We provide additional figures with continuous error bars in Appendix C.2.

3. In the *ordered phase* ( $a < 1$ ), the NTK dispersion tends to one:

$$\frac{\mathbb{E}[\Theta^2(x, x)]}{\mathbb{E}^2[\Theta(x, x)]} \rightarrow 1. \quad (10)$$

Our numerical experiments in Figure 1 demonstrate that Theorem 3.1 provides accurate approximations for the behavior of sufficiently deep and wide DNNs. Indeed, the proofs listed in Appendix A.1 show that the expressions in the above theorem are true up to the approximation given by  $(1 + c/M)^L \approx e^{c\lambda}$  and  $O(1/\sqrt{M})$  in the coefficients of the exponents in case of finite width and depth.

**Remark 1.** The EOC expression in Theorem 3.1 tends to the chaotic phase expression if  $\alpha_0 := n_0/M$  tends to zero (i.e. when the input dimension is fixed). We discuss this effect in Appendix B.1.

**Remark 2.** Model scaling introduced in papers on the so-called “lazy training” phenomenon (Chizat et al., 2019) does not change the results of Theorem 3.1. We discuss lazy training and its effects on our analysis in Appendix B.3.

### 3.2. Finite depth and width effects

We notice that some features of the NTK dispersion are still not visible in the infinite-depth-and-width limit. One can see in Figure 1 that the NTK variance in the ordered phase is not exactly zero for finite-width DNNs, contrary to the prediction in the limit. This is especially noticeable for initialization close to the EOC, where the transition between the two kinds of limiting behavior occurs. Moreover, Theorem 3.1 cannot reveal the effects of the architecture since it considers only DNNs of constant width. Therefore, we provide non-asymptotic expressions for the first two moments of the NTK at initialization in the following theorem and show that these expressions accurately describe the behavior of finite-width DNNs.

**Theorem 3.2** (Moments of the NTK at initialization). *Consider a ReLU DNN defined in (3) with widths scaling defined in (5) and the output dimension  $n_L = 1$ . The initialization is given by (6) and  $\sigma_b = 0$ . Then the expectation of the NTK is determined by the following terms:*

$$\mathbb{E}[\Theta_W(x, x)] = \|\mathbf{x}^0\|^2 a^{L-1} \sum_{\ell=1}^L \frac{n_{\ell-1}}{n_0}, \quad (11)$$

$$\mathbb{E}[\Theta_b(x, x)] = \sum_{\ell=1}^L a^{L-\ell}, \quad (12)$$

where the NTK components  $\Theta_W$  and  $\Theta_b$  are defined in (4). Moreover, the second moment of the NTK is determined by:

$$\begin{aligned} \frac{\mathbb{E}[\Theta_W^2(x, x)]}{\|\mathbf{x}^0\|^4 a^{2(L-1)}} &= \mathcal{X}_{(1,L)} \left[ \sum_{\ell=1}^L \frac{n_{\ell-1}^2}{n_0^2} \right. \\ &\quad \left. + \sum_{\ell_1 < \ell_2} \frac{n_{\ell_2-1} n_{\ell_1-1}}{n_0^2} \frac{\mathcal{C}_{(\ell_1, \ell_2)}}{\mathcal{X}_{(\ell_1, \ell_2)}} \right], \end{aligned} \quad (13)$$

$$\frac{\mathbb{E}[\Theta_b(x, x)^2]}{a^{2L}} = \sum_{\ell=1}^L \frac{\mathcal{X}_{(\ell,L)}}{a^{2\ell}} + 2 \sum_{\ell_1 < \ell_2} \frac{\mathcal{X}_{(\ell_2,L)}}{a^{\ell_1 + \ell_2}}, \quad (14)$$

$$\begin{aligned} \frac{\mathbb{E}[\Theta_W(x, x)\Theta_b(x, x)]}{\|\mathbf{x}^0\|^2 a^{2L-1}} &= \sum_{\ell=1}^L \frac{n_{\ell-1}}{n_0} \frac{\mathcal{X}_{(\ell,L)}}{a^\ell} \\ &\quad + \sum_{\ell_1 < \ell_2} \frac{\mathcal{X}_{(\ell_2,L)}}{a^{\ell_1}} \frac{n_{\ell_1-1}}{n_0} \left( \frac{n_{\ell_2-1}}{n_{\ell_1-1}} \mathcal{C}_{(\ell_1, \ell_2)} + \frac{a^{\ell_1}}{a^{\ell_2}} \right), \end{aligned} \quad (15)$$

where we denoted  $\mathcal{X}_{(i,j)} := \prod_{k=i}^{j-1} \left(1 + \frac{5}{n_k} + O(M^{-3/2})\right)$ ,  $\mathcal{C}_{(i,j)} := \prod_{k=i}^{j-1} \left(1 + \frac{1}{n_k} + O(M^{-3/2})\right)$  and  $a := \sigma_w^2/2$ .

These expressions are derived in Appendix A.1 as a part of the proof of Theorem 3.1 and they simplify to the results in the limit by noticing that  $\mathcal{X}_{(1,L)} \rightarrow e^{5\lambda}$  and  $\mathcal{C}_{(1,L)} \rightarrow e^\lambda$ .

Figure 2 examines how well the above expressions approximate the NTK of DNNs with varying ratios  $\alpha_0 := n_0/M$  between the input dimension and the width of hidden layers. One can see that the NTK variance in the ordered phase indeed grows as the initialization approaches the EOC. This effect is due to the terms proportional to  $((a-1)M)^{-1}$  in the moments of  $\Theta_b(x, x)$  and  $\Theta_W(x, x)\Theta_b(x, x)$ . When the initialization is close enough to the EOC,  $(a-1)^{-1}$  becomes comparable with finite  $M$ , and therefore the behavior diverges from the limit.

Another remarkable observation is that the NTK dispersion may decrease with depth in the ordered phase for DNNs that increase the dimensionality (i.e.  $n_0 \leq n_1 \leq \dots \leq n_{L-1}$ ), which means that deeper networks can be more robust. Indeed, in Subfigures b) and c) of Figure 2, the dispersion reaches its peak at a certain depth and then decreases. We provide additional results characterizing this effect in DNNs with non-constant width in hidden layers in Appendix B.2.

### 3.3. Non-diagonal elements of the NTK

The results stated so far only concern the diagonal elements of the NTK. To generalize to the whole kernel, we provide the following theorem proven in Appendix A.2:

**Theorem 3.3** (Non-diagonal elements of the NTK). *Consider a ReLU DNN from Theorem 3.2. The following bounds hold for the ratio of non-diagonal and diagonal elements of the NTK:*

$$1 \geq \lim_{\substack{L \rightarrow \infty, M \rightarrow \infty \\ L/M \rightarrow \lambda \in \mathbb{R}}} \frac{\mathbb{E}[\Theta(x, \tilde{x})]}{\mathbb{E}[\Theta(x, x)]} \geq \frac{1}{4}. \quad (16)$$

Moreover, the dispersion of the non-diagonal elements is bounded by the dispersion of diagonal ones:

$$\lim_{\substack{L \rightarrow \infty \\ M \rightarrow \infty \\ L/M \rightarrow \lambda}} \frac{\mathbb{E}[\Theta^2(x, \tilde{x})]}{\mathbb{E}^2[\Theta(x, \tilde{x})]} \leq 16 \lim_{\substack{L \rightarrow \infty \\ M \rightarrow \infty \\ L/M \rightarrow \lambda}} \frac{\mathbb{E}[\Theta^2(x, x)]}{\mathbb{E}^2[\Theta(x, x)]}. \quad (17)$$

Of course, the bound in (17) is too loose for practical applications if the goal is to prove that the NTK is approximately deterministic. However, we note that the ratio of non-diagonal and diagonal elements can be close to the lower bound only in the chaotic phase. In the ordered phase, our proof suggests the following bound for sufficiently wide and deep networks:

$$\frac{\mathbb{E}[\Theta(x, \tilde{x})]}{\mathbb{E}[\Theta(x, x)]} \gtrsim \frac{\sum_{\ell=1}^L a^{L-\ell} \prod_{k=\ell}^{L-1} g(\rho_{k-1})}{\sum_{\ell=1}^L a^{L-\ell}}, \quad (18)$$

where  $g(t) := \frac{1}{\pi}(\pi/2 + \arcsin t)$  and  $\rho_k$  is the infinite-width approximation of the cosine distance between  $\mathbf{x}^k$  and  $\tilde{\mathbf{x}}^k$ ,

which only increases with depth and is given by applying the function  $r(t) := \frac{1}{\pi}(\sqrt{1-t^2} + t\pi/2 + t \arcsin t)$  to  $\langle \mathbf{x}^0, \tilde{\mathbf{x}}^0 \rangle$  consecutively  $k$  times. The function  $r(\cdot)$  arises from the expectation of a product of two correlated Gaussian variables under ReLU function.

We provide empirical results on the ratio of non-diagonal and diagonal elements of the NTK in Figure 3. We also plot the estimate for the ratio given by (18) in the same figure. One can see that the ratio quickly increases with depth in the ordered phase. Moreover, the lower bound in (18) gives a good approximation for the experimental results. Then for a given network in the ordered phase one can replace the coefficient 16 in the bound (17) with  $1/c^2$ , where  $c$  is a better estimate for the lower bound of  $\mathbb{E}[\Theta(x, \tilde{x})]/\mathbb{E}[\Theta(x, x)]$  and can be close to one in the ordered phase.

We also provide experiments on the dispersion of the non-diagonal elements in Appendix C.1. Our results indicate that, in practice, the dispersion here is only slightly higher than the prediction for the diagonal elements. The general picture stays the same as in Figure 1: the dispersion is low and does not grow with depth in the ordered phase but increases exponentially with the depth-to-width ratio at the EOC and in the chaotic phase. The finite-width effects represented in Figure 2 also remain the same for the non-diagonal elements.

### 3.4. Proof ideas

All our proofs are based on the following decomposition of the NTK:

$$\Theta(x, x) = \sum_{\ell=1}^L \|\delta^\ell(x)\|^2 (\|\mathbf{x}^{\ell-1}(x)\|^2 + 1), \quad (19)$$

which directly follows from (4) and the representation of the gradients in backpropagation (7). Using forward-propagation equations (3) and backpropagation equations (7), we derive the first two moments for the ratios  $\mathcal{N}_x^\ell := \|\mathbf{x}^\ell\|^2/\|\mathbf{x}^{\ell-1}\|^2$  and  $\mathcal{N}_\delta^\ell := \|\delta^\ell\|^2/\|\delta^{\ell+1}\|^2$  in Lemmas A.1 and A.2. We then notice that  $\mathcal{N}_x^\ell$  are uncorrelated in different layers of the networks, as well as  $\mathcal{N}_\delta^\ell$ , while  $\mathcal{N}_x^\ell$  and  $\mathcal{N}_\delta^\ell$  in the same layer can be weakly correlated and we quantify the effects of this dependence in Lemma A.4. Given the moments of  $\mathcal{N}_x^\ell$  and  $\mathcal{N}_\delta^\ell$  and the results on their correlations, we can represent summands of the NTK as the following telescopic products:

$$\|\delta^\ell\|^2 \|\mathbf{x}^{\ell-1}\|^2 = \|\mathbf{x}^0\|^2 \|\delta^L\|^2 \prod_{k=1}^{\ell-1} \mathcal{N}_x^k \prod_{p=\ell}^{L-1} \mathcal{N}_\delta^p \quad (20)$$

and use this decomposition to compute the expectation and the second moment of the NTK. We derive the first two moments for  $\Theta_W(x, x)$  and  $\Theta_b(x, x)$  separately in Lemmas A.5 and A.6. These two components have very different




Figure 3. Ratio  $\mathbb{E}[\Theta(x, \tilde{x})]/\mathbb{E}[\Theta(x, x)]$  at initialization for fully-connected ReLU networks of constant width  $M \in \{100, 200, 500\}$  with  $\alpha_0 = 1$ . Colors and markers indicate different values of  $\sigma_w^2$ . There are 5 dashed lines for each  $\sigma_w^2$  value, which correspond to 5 values of the initial angle between input samples  $\langle \mathbf{x}^0, \tilde{\mathbf{x}}^0 \rangle \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$ . Darker lines (which also display larger values of the ratio of interest) correspond to larger product  $\langle \mathbf{x}^0, \tilde{\mathbf{x}}^0 \rangle$ . Expectations are computed by sampling 500 random initializations of each DNN configuration. The solid lines show the estimate for the ratio of interest given by (18) in the ordered phase.

properties in the infinite-depth-and-width limit and, as we show in the proof of Theorem 3.1, the behavior of the NTK is determined by  $\Theta_W(x, x)$  in the chaotic phase and by  $\Theta_b(x, x)$  in the ordered phase. We also derive the expectation of  $\Theta_W(x, x)\Theta_b(x, x)$  in Lemma A.7 to complete the calculations of the second moment of the NTK.

We note that many papers on the NTK use the so-called gradient independence assumption (GIA), which leads to the independence of  $\mathcal{N}_x^\ell$  and  $\mathcal{N}_\delta^\ell$ . This assumption often leads to correct results in the infinite-width limit, as discussed in Yang (2019). However, in our case of infinite depth and width, it may have a non-negligible effect even for simple fully-connected networks with all the weights initialized independently. Thus, we have to calculate this effect explicitly in our proofs. We also note that Li et al. (2021) used a similar technique involving telescoping products of weakly-correlated variables to derive the distribution of the activation norms of ResNets.

## 4. Training dynamics of the NTK

In the infinite-width limit, the NTK stays constant during training, which allows to study the gradient flow dynamics of infinitely-wide DNNs analytically. In this section, we discuss when this result holds outside of the infinite-width limit and how the empirical NTK changes during training.

### 4.1. The first GD step

Hanin & Nica (2020) proved that the NTK of over-parametrized fully-connected ReLU networks initialized at the EOC can evolve non-trivially during GD training if depth and width of the network are comparable. In particu-

lar, their result bounds the relative change of the diagonal elements of the NTK  $\Theta(x, x)$  in the first GD step carried out on a single sample  $x$  above and below by an exponential function of the depth-to-width ratio  $\lambda$ . We generalize this result to different initializations with the following theorem proven in Appendix A.3:

**Theorem 4.1** (GD step of the NTK). *Consider a ReLU DNN from Theorem 3.2. A single GD update on a sample  $(x, y) \in \mathcal{D}$  results in the following changes of the NTK:*

1. In the **chaotic phase** ( $a := \sigma_w^2/2 > 1$ ), the changes to the NTK value are infinite in the limit for a constant learning rate  $\eta \in \mathbb{R}$ :

$$\frac{\mathbb{E}[\Delta\Theta(x, x)]}{\mathbb{E}[\Theta(x, x)]} \xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} \infty. \quad (21)$$

2. In the **ordered phase** ( $a < 1$ ), the NTK stays constant in the limit:

$$\frac{\mathbb{E}[\Delta\Theta(x, x)]}{\mathbb{E}[\Theta(x, x)]} \xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} 0. \quad (22)$$

This result shows that deep networks can potentially behave according to the NTK theory during GD training only in case of initialization in the ordered phase. We refer to experiments in Seleznova & Kutyniok (2021), which confirm that the relative change of the NTK during training on MNIST is significant and grows with depth in the chaotic phase and at the EOC but not in the ordered phase. However, it is unclear how to generalize this result to realistic scenarios of DNN training, which include randomly selected batches of

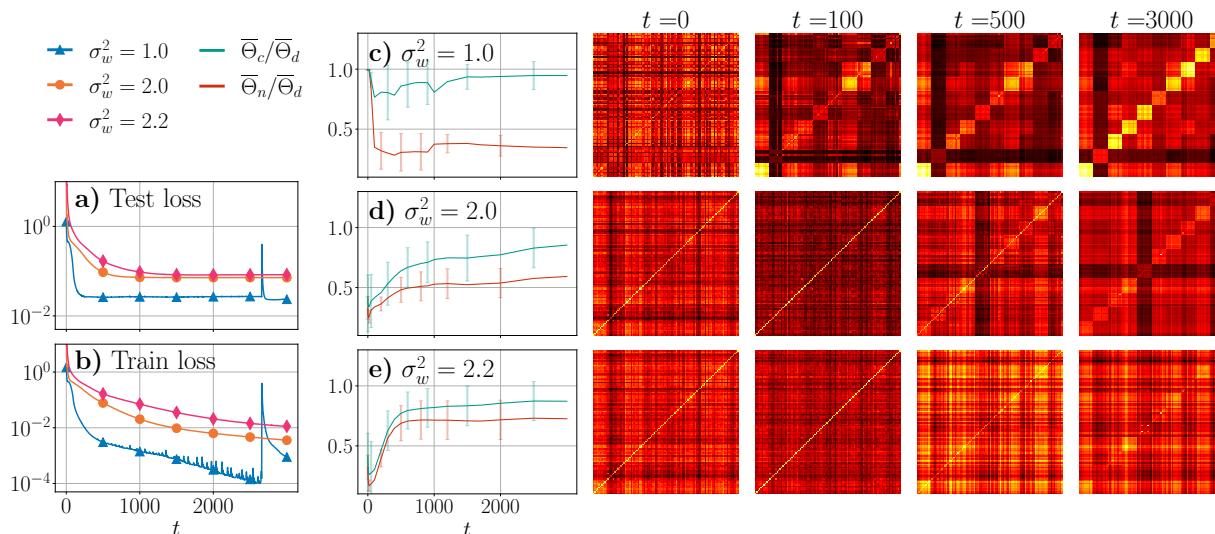


Figure 4. Structure of the NTK matrix in different stages of training for fully-connected ReLU DNNs with  $L = 20$  and  $M = 300$ . The DNNs are initialized with  $\sigma_w^2 \in \{1.0, 2.0, 2.2\}$  and trained on MNIST using Adam algorithm with learning rate  $10^{-5}$ . Subplots **a)** and **b)** show the test and the train loss achieved by each DNN. Subplots **c)**, **d)** and **e)** characterize label-awareness of the NTK. Variables  $\bar{\Theta}_d$ ,  $\bar{\Theta}_c$  and  $\bar{\Theta}_n$  are defined in (23). The heatmaps show the NTK matrix on MNIST subsample of size 100 at epoch  $t \in \{0, 100, 500, 3000\}$ . The subsample contains 10 elements of each class and is arranged so that consecutive diagonal blocks of size 10 contain pairwise NTK values on each class. The color range in the heatmaps is adjusted to include the interval between the maximal and the minimal values of the NTK in a given epoch, i.e. the colors correspond to different values for different epochs. Brighter colors indicate larger values.

arbitrary size and optimization algorithms beyond vanilla GD. Our experiments in the next subsection show that the NTK evolution is in general non-trivial even in the ordered phase.

**Remark 3.** Deep networks rescaled as in [Chizat et al. \(2019\)](#) can exhibit lazy training (with random NTK at initialization) in the chaotic phase only if the scaling parameter grows exponentially with depth  $L$ . We discuss the lazy training phenomenon and its effects on our results in Appendix B.3.

## 4.2. Changes of the NTK structure

The NTK at initialization is label-agnostic, i.e. its value on a pair  $(x, \tilde{x})$  is independent of whether the labels of  $x$  and  $\tilde{x}$  are the same or not. Clearly, label-agnostic features cannot provide an optimal representation system for an arbitrary task and many authors studied the benefits of adding label information to kernels ([Cristianini et al., 2001](#); [Gönen & Alpaydin, 2011](#); [Tishby & Zaslavsky, 2015](#)). In particular, [Chen et al. \(2020\)](#) argued that label-agnosticism can explain the performance gap between trained DNNs and the NTK and demonstrated that adding label-awareness improves the performance of the infinite-width NTK. Thus, it is important to characterize label-awareness of the empirical NTK and how the training process leads to it to understand the properties of DNNs.

We saw in Section 3.3 that the NTK at initialization has an approximately diagonal structure with the diagonal values

larger than the non-diagonal ones. On the contrary, the ”optimal kernel” for a classification task would be block-diagonal with blocks of larger values corresponding to samples of the same class. Thus, we expect the NTK to naturally change towards the block-diagonal structure during the training process. Our experiments in Figure 4 confirm this intuition in a simple setting of fully-connected ReLU networks trained on MNIST. Let us define the following variables that characterize label-awareness of the NTK matrix:

$$\begin{aligned} \bar{\Theta}_d &:= \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} \Theta(x, x), \\ \bar{\Theta}_c &:= \frac{1}{K} \sum_{k=1}^K \frac{1}{|\mathcal{X}_k|(|\mathcal{X}_k| - 1)} \sum_{\substack{x_i \neq x_j, \\ x_i, x_j \in \mathcal{X}_k}} \Theta(x_i, x_j), \\ \bar{\Theta}_n &:= \frac{1}{K} \sum_{k=1}^K \frac{1}{|\mathcal{X}_k|(|\mathcal{X}| - |\mathcal{X}_k|)} \sum_{\substack{x_i \in \mathcal{X}_k, \\ x_j \notin \mathcal{X}_k}} \Theta(x_i, x_j), \end{aligned} \quad (23)$$

where  $\mathcal{X} = \cup_{k=1}^K \mathcal{X}_k$  is the decomposition of the dataset  $\mathcal{X}$  into  $K$  classes. Then  $\bar{\Theta}_d$  is the mean diagonal value,  $\bar{\Theta}_c$  is the mean value of the NTK on samples from the same class and  $\bar{\Theta}_n$  is the mean value on samples from different classes. Figure 4 suggests that a larger gap between  $\bar{\Theta}_n/\bar{\Theta}_d$  and  $\bar{\Theta}_c/\bar{\Theta}_d$  may be related to better performance of DNNs. Moreover, the gap between 1 and the ratio  $\bar{\Theta}_c/\bar{\Theta}_d$  may characterize overfitting. Therefore, we believe that the structure



of the NTK can be a proxy for generalization of DNNs even outside of the NTK regime. One can also see that the structure of the NTK changes more rapidly in the early stages of training, which is coherent with the conclusion in Fort et al. (2020) that useful features are mostly learned in the first epochs of training. Thus, dynamics of the NTK may provide information about the state of the training process.

## 5. Conclusions and future work

This paper adds to the line of research on the statistical properties of the NTK and the correspondence between finite-width DNNs and their infinite-width approximations. Our results in Section 3 precisely quantify variability of the NTK at initialization for a given fully-connected ReLU DNN and assess how well the kernel is approximated by its infinite-width limit. Combining our findings from Section 3 with the results on the GD update of the NTK in Section 4.1, we conclude that the NTK regime can approximate trained networks with non-trivial depth-to-width ratio only in the ordered phase. At the same time, the behavior of overparametrized DNNs outside of the NTK regime is very poorly understood so far. It is unclear how to characterize DNNs' training dynamics in the general case and what role the properties of the (random and dynamic) NTK play here. We make a step into this direction in Section 4.2 by demonstrating how the NTK acquires a block-diagonal structure during training. We believe that precisely characterizing the effects of this NTK structure on the generalization of DNNs is a promising direction for future work. In general, we hope to establish new connections between the NTK and other aspects of DNN training outside of the NTK regime.

## Acknowledgements

G.K. acknowledges partial support by the NSF–Simons Research Collaboration on the Mathematical and Scientific Foundations of Deep Learning (MoDL) (NSF DMS 2031985) and DFG SPP 1798, KU 1446/27-2 and KU 1446/21-2.

## References

- Adlam, B. and Pennington, J. The neural tangent kernel in high dimensions: Triple descent and a multi-scale theory of generalization. In *International Conference on Machine Learning*, pp. 74–84. PMLR, 2020.
- Aitchison, L. Why bigger is not always better: on finite and infinite neural networks. In *International Conference on Machine Learning, ICML*, volume 119, pp. 156–164. PMLR, 2020.
- Alemohammad, S., Wang, Z., Balestrieri, R., and Baraniuk, R. G. The recurrent neural tangent kernel. In *International Conference on Learning Representations, ICLR*, 2021.
- Arora, S., Du, S. S., Hu, W., Li, Z., Salakhutdinov, R., and Wang, R. On exact computation with an infinitely wide neural net. In *Advances in Neural Information Processing Systems*, pp. 8139–8148, 2019.
- Arora, S., Du, S. S., Li, Z., Salakhutdinov, R., Wang, R., and Yu, D. Harnessing the power of infinitely wide deep nets on small-data tasks. In *International Conference on Learning Representations, ICLR*, 2020.
- Bai, Y., Krause, B., Wang, H., Xiong, C., and Socher, R. Taylorized training: Towards better approximation of neural network training at finite width. *CoRR*, abs/2002.04010, 2020. URL <https://arxiv.org/abs/2002.04010>.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Buchanan, S., Gilboa, D., and Wright, J. Deep networks and the multiple manifold problem. In *International Conference on Learning Representations, ICLR*, 2021.
- Chen, S., He, H., and Su, W. J. Label-aware neural tangent kernel: Toward better generalization and local elasticity. In *Advances in Neural Information Processing Systems*, 2020.
- Chizat, L., Oyallon, E., and Bach, F. R. On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems*, pp. 2933–2943, 2019.
- Cristianini, N., Shawe-Taylor, J., Elisseeff, A., and Kandola, J. S. On kernel-target alignment. In *Advances in Neural Information Processing Systems*, pp. 367–373. MIT Press, 2001.
- de G. Matthews, A. 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, ICLR*, 2018.
- Du, S. S., Hou, K., Salakhutdinov, R., Póczos, B., Wang, R., and Xu, K. Graph neural tangent kernel: Fusing graph neural networks with graph kernels. In *Advances in Neural Information Processing Systems*, pp. 5724–5734, 2019.
- Fort, S., Dziugaite, G. K., Paul, M., Kharaghani, S., Roy, D. M., and Ganguli, S. Deep learning versus kernel learning: an empirical study of loss landscape geometry and the time evolution of the neural tangent kernel. In *Advances in Neural Information Processing Systems*, 2020.

- Geiger, M., Jacot, A., Spigler, S., Gabriel, F., Sagun, L., d'Ascoli, S., Biroli, G., Hongler, C., and Wyart, M. Scaling description of generalization with number of parameters in deep learning. *CoRR*, abs/1901.01608, 2019. URL <http://arxiv.org/abs/1901.01608>.
- Gönen, M. and Alpaydin, E. Multiple kernel learning algorithms. *J. Mach. Learn. Res.*, 12:2211–2268, 2011.
- Hanin, B. and Nica, M. Finite depth and width corrections to the neural tangent kernel. In *International Conference on Learning Representations, ICLR*, 2020.
- Hayou, S., Doucet, A., and Rousseau, J. Mean-field behaviour of neural tangent kernel for deep neural networks. *CoRR*, abs/1905.13654, 2019. URL <https://arxiv.org/abs/1905.13654>.
- Heek, J., Levskaya, A., Oliver, A., Ritter, M., Rondepierre, B., Steiner, A., and van Zee, M. Flax: A neural network library and ecosystem for JAX, 2020. URL <http://github.com/google/flax>.
- Hu, Z. and Huang, H. On the random conjugate kernel and neural tangent kernel. In *International Conference on Machine Learning*, pp. 4359–4368. PMLR, 2021.
- Huang, J. and Yau, H. Dynamics of deep neural networks and neural tangent hierarchy. In *International Conference on Machine Learning, ICML*, volume 119, pp. 4542–4551. PMLR, 2020.
- Huang, K., Wang, Y., Tao, M., and Zhao, T. Why do deep residual networks generalize better than deep feedforward networks? — A neural tangent kernel perspective. In *Advances in Neural Information Processing System*, 2020.
- Jacot, A., Hongler, C., and Gabriel, F. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*, pp. 8580–8589, 2018.
- Karakida, R., Akaho, S., and Amari, S. Universal statistics of Fisher information in deep neural networks: Mean field approach. In *International Conference on Artificial Intelligence and Statistics, AISTATS*, volume 89, pp. 1032–1041. PMLR, 2019.
- Korotkov, N. E. and Korotkov, A. N. *Integrals Related to the Error Function*. CRC Press, 2020.
- Lee, J., Xiao, L., Schoenholz, S. S., Bahri, Y., Novak, R., Sohl-Dickstein, J., and Pennington, J. Wide neural networks of any depth evolve as linear models under gradient descent. In *Advances in Neural Information Processing Systems*, pp. 8570–8581, 2019.
- Lee, J., Schoenholz, S. S., Pennington, J., Adlam, B., Xiao, L., Novak, R., and Sohl-Dickstein, J. Finite versus infinite neural networks: an empirical study. In *Advances in Neural Information Processing Systems*, 2020.
- Li, M. B., Nica, M., and Roy, D. M. The future is log-gaussian: ResNets and their infinite-depth-and-width limit at initialization. *CoRR*, abs/2106.04013, 2021. URL <https://arxiv.org/abs/2106.04013>.
- Poole, B., Lahiri, S., Raghu, M., Sohl-Dickstein, J., and Ganguli, S. Exponential expressivity in deep neural networks through transientchaos. In *Advances in Neural Information Processing Systems*, pp. 3360–3368, 2016.
- Schoenholz, S. S., Gilmer, J., Ganguli, S., and Sohl-Dickstein, J. Deep information propagation. In *International Conference on Learning Representations, ICLR*, 2017.
- Seleznova, M. and Kutyniok, G. Analyzing finite neural networks: Can we trust neural tangent kernel theory? In *Conference on Mathematical and Scientific Machine Learning*, volume 145. PMLR, 2021.
- Tirer, T., Bruna, J., and Giryes, R. Kernel-based smoothness analysis of residual networks. In *Conference on Mathematical and Scientific Machine Learning*, volume 145. PMLR, 2021.
- Tishby, N. and Zaslavsky, N. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop, ITW*, pp. 1–5. IEEE, 2015. doi: 10.1109/ITW.2015.7133169.
- Wang, S., Yu, X., and Perdikaris, P. When and why PINNs fail to train: A neural tangent kernel perspective. *J. Comput. Phys.*, 449:110768, 2022. doi: 10.1016/j.jcp.2021.110768.
- Xiao, L., Pennington, J., and Schoenholz, S. Disentangling trainability and generalization in deep neural networks. In *International Conference on Machine Learning*, pp. 10462–10472. PMLR, 2020.
- Yang, G. Scaling limits of wide neural networks with weight sharing: Gaussian process behavior, gradient independence, and neural tangent kernel derivation. *CoRR*, abs/1902.04760, 2019. URL <http://arxiv.org/abs/1902.04760>.
- Yang, G. Tensor programs II: Neural tangent kernel for any architecture. *CoRR*, abs/2006.14548, 2020. URL <https://arxiv.org/abs/2006.14548>.
- Yang, G. and Hu, E. J. Tensor programs IV: Feature learning in infinite-width neural networks. In *International Conference on Machine Learning*, volume 139, pp. 11727–11737. PMLR, 2021.

Yang, G. and Schoenholz, S. S. Mean field residual networks: On the edge of chaos. In *Advances in Neural Information Processing Systems*, pp. 7103–7114, 2017.

## A. Proofs

### A.1. Variability of the NTK at initialization

**Lemma A.1** (Forward-propagation of variance). *Consider a fully-connected DNN defined in (3) initialized as in (6). The activation function in the hidden layers is ReLU, i.e.  $\phi(x) = x\mathbb{1}\{x > 0\}$ . Assume further that the biases are initialized to zero, i.e.  $\sigma_b = 0$ . Then the following holds for the ratios of the activation norms in consecutive layers of the network, denoted  $\mathcal{N}_x^\ell := \|\mathbf{x}^\ell\|^2 / \|\mathbf{x}^{\ell-1}\|^2$ ,  $\ell = 1, \dots, L-1$ :*

$$\mathbb{E}[\mathcal{N}_x^\ell] = \frac{\sigma_w^2}{2} \frac{n_\ell}{n_{\ell-1}}, \quad \mathbb{E}[(\mathcal{N}_x^\ell)^2] = \left(\frac{\sigma_w^2}{2}\right)^2 \left(\frac{n_\ell}{n_{\ell-1}}\right)^2 \left(1 + \frac{5}{n_\ell}\right). \quad (24)$$

$$\frac{\mathcal{N}_x^\ell - \mathbb{E}[\mathcal{N}_x^\ell]}{\sqrt{\mathbb{V}[\mathcal{N}_x^\ell]}} \xrightarrow[n_\ell \rightarrow \infty]{d} \mathcal{N}(0, 1), \quad (25)$$

where  $\mathcal{N}(0, 1)$  is the standard normal distribution. Moreover, random variables  $\{\mathcal{N}_x^\ell\}_{\ell=0, \dots, L-1}$  are mutually independent.

*Proof.* The squared norm of the activation vector in layer  $\ell$  is given by

$$\|\mathbf{x}^\ell\|^2 = \sum_{i=1}^{n_\ell} \phi^2(\mathbf{W}_{i \cdot}^\ell \mathbf{x}^{\ell-1} + b_i^\ell)$$

Here  $\mathbf{x}^{\ell-1}$  depends only on  $\{(\mathbf{W}^j, \mathbf{b}^j)\}_{j=1, \dots, \ell-1}$ , therefore  $\mathbf{x}^{\ell-1}$  is independent of  $(\mathbf{W}^\ell, \mathbf{b}^\ell)$ . Since elements of  $\mathbf{W}^\ell$  are i.i.d Gaussian, the distribution of  $\mathbf{W}^\ell \mathbf{x}^{\ell-1}$  depends only on the norm of  $\mathbf{x}^{\ell-1}$  and not on the direction. Then we can write the following equalities in distribution:

$$\begin{aligned} \mathbf{W}_{i \cdot}^\ell \mathbf{x}^{\ell-1} &= \sqrt{\frac{\sigma_w^2}{n_{\ell-1}}} \|\mathbf{x}^{\ell-1}\| \mathcal{U}_i^\ell, \\ \phi^2(\mathbf{W}_{i \cdot}^\ell \mathbf{x}^{\ell-1} + b_i^\ell) &= \left(\sqrt{\frac{\sigma_w^2}{n_{\ell-1}}} \|\mathbf{x}^{\ell-1}\| + \sigma_b\right)^2 \phi^2(\mathcal{U}_i^\ell), \end{aligned}$$

where we introduced i.i.d random variables  $\mathcal{U}_i^\ell = \left\langle \sqrt{\frac{n_{\ell-1}}{\sigma_w^2}} (\mathbf{W}_{i \cdot}^\ell)^T, \frac{\mathbf{x}^{\ell-1}}{\|\mathbf{x}^{\ell-1}\|} \right\rangle \sim \mathcal{N}(0, 1)$ ,  $i = 1, \dots, n_\ell$ , which are independent of  $\mathbf{x}^{\ell-1}$ , and used the fact that  $\phi(\alpha x) = \alpha \phi(x)$  for  $\alpha \in \mathbb{R}^+$ . Therefore for the norm of the activation vector we have the following:

$$\|\mathbf{x}^\ell\|^2 = \sum_{i=1}^{n_\ell} \phi^2(\mathbf{W}_{i \cdot}^\ell \mathbf{x}^{\ell-1} + b_i^\ell) = \left(\sqrt{\frac{\sigma_w^2}{n_{\ell-1}}} \|\mathbf{x}^{\ell-1}\| + \sigma_b\right)^2 \sum_{i=1}^{n_\ell} \phi^2(\mathcal{U}_i^\ell),$$

where only the first bracket depends on  $\mathbf{x}^{\ell-1}$ . Then in case of zero biases, i.e.  $\sigma_b = 0$ , for the ratio between the norms of consecutive activation vectors we have

$$\mathcal{N}_x^\ell = \frac{\sigma_w^2}{n_{\ell-1}} \sum_{i=1}^{n_\ell} \phi^2(\mathcal{U}_i^\ell),$$

where the variables  $\mathcal{U}_i^\ell, i = 1, \dots, n_\ell$  depend only on the weights in the given layer  $\mathbf{W}^\ell$ . Then the ratios  $\mathcal{N}_x^\ell$  in different layers are independent and we can obtain the desired moments of  $\mathcal{N}_x^\ell$  as follows:

$$\begin{aligned} \mathbb{E}[\mathcal{N}_x^\ell] &= \frac{\sigma_w^2}{n_{\ell-1}} \sum_{i=1}^{n_\ell} \mathbb{E}[\phi^2(\mathcal{U}_i^\ell)] = \frac{\sigma_w^2}{2} \frac{n_\ell}{n_{\ell-1}}, \\ \mathbb{E}[(\mathcal{N}_x^\ell)^2] &= \left(\frac{\sigma_w^2}{n_{\ell-1}}\right)^2 \sum_{i=1}^{n_\ell} \mathbb{V}[\phi^2(\mathcal{U}_i^\ell)] + \mathbb{E}^2[\phi^2(\mathcal{U}_i^\ell)] = \left(\frac{\sigma_w^2}{2}\right)^2 \left(\frac{n_\ell}{n_{\ell-1}}\right)^2 \left(1 + \frac{5}{n_\ell}\right), \end{aligned}$$

where we used the moments of variables  $\phi(\mathcal{U}_i)$ , which can be calculated by integration:

$$\mathbb{E}[\phi^2(\mathcal{U}_i^\ell)] = \frac{1}{2}, \quad \mathbb{V}[\phi^2(\mathcal{U}_i^\ell)] = \frac{5}{4}, \quad i = 1, \dots, n_\ell.$$

Moreover, by the central limit theorem we have

$$\frac{\mathcal{N}_x^\ell - \mathbb{E}[\mathcal{N}_x^\ell]}{\sqrt{\mathbb{V}[\mathcal{N}_x^\ell]}} = \frac{2\left(\frac{1}{n_\ell} \sum_{i=1}^{n_\ell} \phi^2(\mathcal{U}_i^\ell) - \frac{1}{2}\right)}{\sqrt{5/n_\ell}} \xrightarrow[n_\ell \rightarrow \infty]{d} \mathcal{N}(0, 1).$$

□

**Lemma A.2** (Backpropagation of variance). *Consider the same setting as in Lemma A.1. Then the following holds for the ratios of norms of backpropagated errors (defined in (7)) in consecutive layers, denoted  $\mathcal{N}_\delta^\ell := \|\delta^\ell\|^2 / \|\delta^{\ell+1}\|^2$ ,  $\ell = 1, \dots, L-1$ :*

$$\mathbb{E}[\mathcal{N}_\delta^\ell] = \frac{\sigma_w^2}{2}, \quad \mathbb{E}[(\mathcal{N}_\delta^\ell)^2] = \left(\frac{\sigma_w^2}{2}\right)^2 \left(1 + \frac{5}{n_\ell}\right). \quad (26)$$

$$\frac{\mathcal{N}_\delta^\ell - \mathbb{E}[\mathcal{N}_\delta^\ell]}{\sqrt{\mathbb{V}[\mathcal{N}_\delta^\ell]}} \xrightarrow[n_\ell \rightarrow \infty]{d} \mathcal{N}(0, 1), \quad (27)$$

where  $\mathcal{N}(0, 1)$  is the standard normal distribution.

*Proof.* The recursive formula for the backpropagated errors is given by

$$\delta_i^\ell = \phi'(\mathbf{h}_i^\ell) \sum_{j=1}^{n_{\ell+1}} \mathbf{W}_{ji}^{\ell+1} \delta_j^{\ell+1} = \phi'(\mathbf{h}_i^\ell) (\mathbf{W}^{\ell+1})_i^T \delta^{\ell+1}.$$

Then, in the same way as in Lemma A.1, we have the following for the squared norm of  $\delta^\ell$ :

$$\|\delta^\ell\|^2 = \sum_{i=1}^{n_\ell} (\phi'(\mathbf{h}_i^\ell))^2 \left( (\mathbf{W}^{\ell+1})_i^T \delta^{\ell+1} \right)^2 = \frac{\sigma_w^2}{n_\ell} \|\delta^{\ell+1}\|^2 \sum_{i=1}^{n_\ell} (\phi'(\mathbf{h}_i^\ell))^2 (\mathcal{V}_i^{\ell+1})^2,$$

where we introduced i.i.d. random variables  $\mathcal{V}_i^{\ell+1} = \left\langle \sqrt{\frac{n_\ell}{\sigma_w^2}} \mathbf{W}_{i \cdot}^{\ell+1}, \frac{\delta^{\ell+1}}{\|\delta^{\ell+1}\|} \right\rangle \sim \mathcal{N}(0, 1)$ ,  $i = 1, \dots, n_\ell$ , which are independent of  $\delta^{\ell+1}$ . One can also see that  $\phi'(\mathbf{h}^\ell)$  can only depend on  $\{(\mathbf{W}^j, \mathbf{b}^j)\}_{j=1, \dots, \ell}$ , therefore it is independent of  $\|\delta\|^{\ell+1}$  and of  $\mathcal{V}_i^{\ell+1}$ ,  $i = 1, \dots, n_\ell$ . Moreover,  $\phi'(\mathbf{h}_i^\ell) = \phi'(\mathbf{W}_i^\ell \mathbf{x}^{\ell-1}) = \phi'(\mathcal{U}_i^\ell)$  for all  $i = 1, \dots, n_\ell$ , therefore  $\phi'(\mathbf{h}^\ell)$  depends only on  $\mathbf{W}^\ell$ . Then we can write the following for the ratio of interest and its moments:

$$\mathcal{N}_\delta^\ell = \frac{\sigma_w^2}{n_\ell} \sum_{i=1}^{n_\ell} \phi'(\mathcal{U}_i^\ell) (\mathcal{V}_i^{\ell+1})^2,$$

$$\mathbb{E}[\mathcal{N}_\delta^\ell] = \frac{\sigma_w^2}{2}, \quad \mathbb{E}[(\mathcal{N}_\delta^\ell)^2] = \left(\frac{\sigma_w^2}{2}\right)^2 \left(1 + \frac{5}{n_\ell}\right),$$

where we calculated the moments of the summands as

$$\mathbb{E}[(\phi'(\mathbf{h}_i^\ell))^2 \mathcal{V}_i^2] = \mathbb{E}[(\phi'(\mathbf{h}_i^\ell))^2] \mathbb{E}[\mathcal{V}_i^2] = \frac{1}{2}, \quad \mathbb{V}[(\phi'(\mathbf{h}_i^\ell))^2 \mathcal{V}_i^2] = \mathbb{E}[(\phi'(\mathbf{h}_i^\ell))^4] \mathbb{E}[\mathcal{V}_i^4] - \frac{1}{4} = \frac{5}{4}.$$

Here we used that, in case of ReLU activation,  $\phi'(\mathbf{h}_i^\ell)$ ,  $i = 1, \dots, n_\ell$  are Bernoulli variables with probability of 1 and 0 equal to 1/2, since  $\mathbf{h}^\ell$  is symmetric around zero. Therefore,  $\mathbb{E}[(\phi'(\mathbf{h}_i^\ell))^2] = \mathbb{E}[(\phi'(\mathbf{h}_i^\ell))^4] = 1/2$ .

Same as in Lemma A.1, the limiting distribution of  $\mathcal{N}_\delta^\ell$  is given by the central limit theorem.

□

As we note in Section 3.4, many papers that study the NTK adopt the following assumption:

**Assumption A.3** (Gradient independence assumption (GIA)). Matrix  $(\mathbf{W}^\ell)^T$  in backpropagation equations (7) and matrix  $\mathbf{W}^\ell$  in forward-propagation equations (3) are independent for all  $\ell \in \{1, \dots, L\}$ .

This assumption is of course not true; however, the products  $(\mathbf{W}^\ell)_i^T \mathbf{x} = \sum_{k=1}^{n_\ell} \mathbf{W}_{ki}^\ell \mathbf{x}_k$  and  $\mathbf{W}_j^\ell \mathbf{x} = \sum_{k=1}^{n_{\ell-1}} \mathbf{W}_{jk}^\ell \mathbf{x}_k$  are only dependent through the single summand containing  $\mathbf{W}_{ij}^\ell$ . Thus, the correlations caused by this dependence are of order  $O(1/M)$  and can be disregarded in the infinite-width limit. However, in our case of the infinite-depth-and-width limit terms of order  $O(1/M)$  can have a non-trivial impact on the computations. Therefore, we calculate the effects of the dependence between the forward-propagated chain and the backpropagated chain in the following lemma.

**Lemma A.4** (Gradient independence assumption (GIA)). *Consider the same setting as in Lemma A.1. Then the following statements hold:*

1. GIA does not change the expectation of  $\|\delta^\ell\|^2 / \|\delta^{\ell+k+1}\|^2$ :

$$\mathbb{E} \left[ \prod_{p=0}^k \mathcal{N}_\delta^{\ell+p} \right] = \prod_{p=0}^k \mathbb{E}[\mathcal{N}_\delta^{\ell+p}] \quad (28)$$

2. GIA changes the expectation of  $\|\delta^\ell\|^2 / \|\delta^{\ell+k+1}\|^2 \cdot \|\mathbf{x}^{\ell+k}\|^2 / \|\mathbf{x}^{\ell-1}\|^2$  by a term that has a non-trivial depth-and-width limit where  $M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}$ . In particular, we have:

$$\mathbb{E} \left[ \prod_{p=0}^k \mathcal{N}_\delta^{\ell+p} \mathcal{N}_x^{\ell+p} \right] = \prod_{p=0}^k \mathbb{E}[\mathcal{N}_\delta^{\ell+p}] \mathbb{E}[\mathcal{N}_x^{\ell+p}] \left( 1 + \frac{1}{n_{\ell+p}} + O\left(\frac{1}{M^{3/2}}\right) \right), \quad (29)$$

where  $n_\ell = \alpha_\ell M, \alpha_\ell \in \mathbb{R}, \ell = 1, \dots, L-1$

3. GIA does not change the expectation of  $(\|\delta^\ell\|^2 / \|\delta^{\ell+k+1}\|^2)^2$  in the infinite-depth-and-width limit where  $M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}$ . In particular, we have:

$$\mathbb{E} \left[ \prod_{p=0}^k (\mathcal{N}_\delta^{\ell+p})^2 \right] = \prod_{p=0}^k \mathbb{E}[(\mathcal{N}_\delta^{\ell+p})^2] \left( 1 + O\left(\frac{1}{M^{3/2}}\right) \right), \quad (30)$$

where  $n_\ell = \alpha_\ell M, \alpha_\ell \in \mathbb{R}, \ell = 1, \dots, L-1$ :

*Proof.* In Lemmas A.1 and A.2 we derived the following equations for  $\mathcal{N}_\delta^\ell$  and  $\mathcal{N}_x^\ell$ :

$$\begin{aligned} \mathcal{N}_\delta^\ell &= \frac{\sigma_w^2}{n_\ell} \sum_{i=1}^{n_\ell} (\phi'(\mathbf{h}_i^\ell))^2 (\mathcal{V}_i^{\ell+1})^2 = \frac{\sigma_w^2}{n_\ell} \sum_{i=1}^{n_\ell} (\phi'(\mathcal{U}_i^\ell))^2 (\mathcal{V}_i^{\ell+1})^2, \\ \mathcal{N}_x^\ell &= \frac{\sigma_w^2}{n_{\ell-1}} \sum_{i=1}^{n_\ell} \phi^2(\mathcal{U}_i^\ell), \end{aligned}$$

where  $\mathcal{U}_i^\ell$  depends only on the  $i$ -th row of the weights matrix  $\mathbf{W}_i^\ell$ , and  $\mathcal{V}_j^\ell$  depends only on  $j$ -th column of the same matrix  $\mathbf{W}_{\cdot j}^\ell$  for  $i = 1, \dots, n_\ell, j = 1, \dots, n_{\ell-1}, \ell = 1, \dots, L-1$ . Therefore, variables  $\mathcal{U}_i^\ell$  and  $\mathcal{V}_j^\ell$  are only dependent through the single weight  $\mathbf{W}_{ij}^\ell$ , which nevertheless makes  $\mathcal{N}_\delta^\ell$  and  $\mathcal{N}_x^{\ell+1}$  dependent for any  $\ell = 1, \dots, L-2$ . One can also see that  $\mathcal{N}_\delta^\ell$  and  $\mathcal{N}_x^\ell$  are dependent through  $\{\mathcal{U}_i^\ell\}_{i=1, \dots, n_\ell}$ . The objective of this lemma is to determine the effects of these weak dependencies on the expectation of products that appear in the NTK.

**Part 1.** We first consider the product of ratios of the backpropagated errors:

$$\begin{aligned} \mathbb{E} \left[ \prod_{p=0}^k \mathcal{N}_\delta^{\ell+p} \right] &= \prod_{p=0}^k \frac{\sigma_w^2}{n_{\ell+p}} \sum_{i_0=1}^{n_\ell} \dots \sum_{i_k=1}^{n_{\ell+k}} \mathbb{E}[\phi'(\mathcal{U}_{i_0}^\ell) (\mathcal{V}_{i_0}^{\ell+1})^2 \phi'(\mathcal{U}_{i_1}^{\ell+1}) (\mathcal{V}_{i_1}^{\ell+2})^2 \dots \phi'(\mathcal{U}_{i_k}^{\ell+k}) (\mathcal{V}_{i_k}^{\ell+k+1})^2] \\ &= \prod_{p=0}^k \frac{\sigma_w^2}{n_{\ell+p}} \sum_{i_0=1}^{n_\ell} \dots \sum_{i_k=1}^{n_{\ell+k}} \mathbb{E}[\phi'(\mathcal{U}_{i_0}^\ell)] \mathbb{E}[(\mathcal{V}_{i_k}^{\ell+k+1})^2] \prod_{p=1}^k \mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \phi'(\mathcal{U}_{i_p}^{\ell+p})] \\ &= \frac{1}{2} \prod_{p=0}^k \frac{\sigma_w^2}{n_{\ell+p}} \sum_{i_0=1}^{n_\ell} \dots \sum_{i_k=1}^{n_{\ell+k}} \prod_{p=1}^k \mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \phi'(\mathcal{U}_{i_p}^{\ell+p})]. \end{aligned}$$

As  $\mathcal{U}_{i_p}^{\ell+p}$  that  $\mathcal{V}_{i_{p-1}}^{\ell+p}$  depend only through  $\mathbf{W}_{i_p i_{p-1}}^{\ell+p}$ , we can condition the expectation of their product as follows:

$$\mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \phi'(\mathcal{U}_{i_p}^{\ell+p})] = \mathbb{E}[\mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \mid \mathbf{W}_{i_p i_{p-1}}^{\ell+p}] \cdot \mathbb{E}[\phi'(\mathcal{U}_{i_p}^{\ell+p}) \mid \mathbf{W}_{i_p i_{p-1}}^{\ell+p}]].$$

To simplify the notation, let us denote  $w_{i_p i_{p-1}} := \sqrt{\frac{n_{\ell+p-1}}{\sigma_w^2}} \mathbf{W}_{i_p i_{p-1}}^{\ell+p} \sim \mathcal{N}(0, 1)$ ,  $a_j := \mathbf{x}_j^{\ell+p} / \|\mathbf{x}^{\ell+p}\|$  and  $b_k := \boldsymbol{\delta}_k^{\ell+p} / \|\boldsymbol{\delta}^{\ell+p}\|$ . Then we have  $\mathcal{V}_{i_{p-1}}^{\ell+p} = \sum_{k=1}^{n_{\ell+p}} w_{k i_{p-1}} b_k = w_{i_p i_{p-1}} b_{i_p} + \sum_{k \neq i_p} w_{k i_{p-1}} b_k$  and  $\mathcal{U}_{i_p}^{\ell+p} = \sum_{j=1}^{n_{\ell+p-1}} w_{j i_{p-1}} a_j = w_{i_p i_{p-1}} a_{i_{p-1}} + \sum_{j \neq i_{p-1}} w_{i_p j} a_j$ . We can then open the conditional expectations:

$$\begin{aligned} \mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \mid \mathbf{W}_{i_p i_{p-1}}^{\ell+p}] &= w_{i_p i_{p-1}}^2 b_{i_p}^2 + \mathbb{E}[(\sum_{k \neq i_p} w_{k i_{p-1}} b_k)^2] = 1 - b_{i_p}^2 (1 - w_{i_p i_{p-1}}^2), \\ \mathbb{E}[\phi'(\mathcal{U}_{i_p}^{\ell+p}) \mid \mathbf{W}_{i_p i_{p-1}}^{\ell+p}] &= \mathbb{P}[\sum_{j \neq i_{p-1}} w_{i_p j} b_k > -w_{i_p i_{p-1}} a_{i_{p-1}}] = \Phi\left(\frac{w_{i_p i_{p-1}} a_{i_{p-1}}}{\sqrt{1 - a_{i_{p-1}}^2}}\right), \end{aligned}$$

where  $\Phi(\cdot)$  is the CDF of the standard normal distribution. Here we used that  $\sum_{k \neq i_p} w_{k i_{p-1}} b_k \sim \mathcal{N}(0, 1 - b_{i_p}^2)$  and  $\sum_{j \neq i_{p-1}} w_{i_p j} a_j \sim \mathcal{N}(0, 1 - a_{i_{p-1}}^2)$ . Then we have:

$$\mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \phi'(\mathcal{U}_{i_p}^{\ell+p})] = (1 - b_{i_p}^2) \mathbb{E}[\Phi(A \cdot w_{i_p i_{p-1}})] + b_{i_p}^2 \mathbb{E}[w_{i_p i_{p-1}}^2 \Phi(A \cdot w_{i_p i_{p-1}})] = \frac{1}{2},$$

where we used the following integrals:

$$\begin{aligned} \mathbb{E}[\Phi(A \cdot w_{i_p i_{p-1}})] &= \frac{1}{2} + \frac{1}{2} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} \operatorname{erf}\left(\frac{A}{\sqrt{2}} w_{i_p i_{p-1}}\right) \exp\left(-\frac{w_{i_p i_{p-1}}^2}{2}\right) dw_{i_p i_{p-1}} = \frac{1}{2}, \\ \mathbb{E}[w_{i_p i_{p-1}}^2 \Phi(A \cdot w_{i_p i_{p-1}})] &= \frac{1}{2} \mathbb{E}[w_{i_p i_{p-1}}^2] \\ &\quad + \frac{1}{2} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} w_{i_p i_{p-1}}^2 \operatorname{erf}\left(\frac{A}{\sqrt{2}} w_{i_p i_{p-1}}\right) \exp\left(-\frac{w_{i_p i_{p-1}}^2}{2}\right) dw_{i_p i_{p-1}} = \frac{1}{2}. \end{aligned}$$

Thus, the expectation of the product of the ratios of backpropagated errors is exactly equal to the product of their expectations:

$$\mathbb{E}\left[\prod_{p=0}^k \mathcal{N}_{\delta}^{\ell+p}\right] = \frac{1}{2} \prod_{p=0}^k \frac{\sigma_w^2}{n_{\ell+p}} \sum_{i_0=1}^{n_{\ell}} \cdots \sum_{i_k=1}^{n_{\ell+k}} \frac{1}{2^k} = \left(\frac{\sigma_w^2}{2}\right)^{k+1} = \prod_{p=0}^k \mathbb{E}[\mathcal{N}_{\delta}^{\ell+p}],$$

which completes the proof of the first statement.

**Part 2.** We now consider the expectation of products of the activations' ratios and the backpropagated errors' ratios for the same layers. The product in a single layer is given by:

$$\mathcal{N}_{\delta}^{\ell} \mathcal{N}_x^{\ell} = \frac{\sigma_w^2}{n_{\ell-1}} \frac{\sigma_w^2}{n_{\ell}} \sum_{i=1}^{n_{\ell}} \sum_{j=1}^{n_{\ell}} \phi^2(\mathcal{U}_j^{\ell}) (\phi'(\mathcal{U}_i^{\ell}))^2 (\mathcal{V}_i^{\ell+1})^2 = \frac{\sigma_w^2}{n_{\ell-1}} \frac{\sigma_w^2}{n_{\ell}} \sum_{i=1}^{n_{\ell}} \sum_{j=1}^{n_{\ell}} \phi'(\mathcal{U}_i^{\ell}) \phi'(\mathcal{U}_j^{\ell}) (\mathcal{U}_j^{\ell})^2 (\mathcal{V}_i^{\ell+1})^2,$$

where we noticed that  $\phi(\mathcal{U}_i^{\ell}) \phi'(\mathcal{U}_i^{\ell}) = \phi(\mathcal{U}_i^{\ell}) = \mathcal{U}_i^{\ell} \phi'(\mathcal{U}_i^{\ell})$ . Then for the product involving multiple layers we have:

$$\begin{aligned} \prod_{p=0}^k \mathcal{N}_{\delta}^{\ell+p} \mathcal{N}_x^{\ell+p} &= \frac{\sigma_w^2}{n_{\ell-1}} \frac{\sigma_w^2}{n_{\ell+k}} \prod_{j=0}^{k-1} \left(\frac{\sigma_w^2}{n_{\ell+j}}\right)^2 \sum_{i_0=1}^{n_{\ell}} \sum_{j_0=1}^{n_{\ell}} \cdots \\ &\quad \cdots \sum_{i_k=1}^{n_{\ell+k}} \sum_{j_k=1}^{n_{\ell+k}} \phi'(\mathcal{U}_{i_0}^{\ell}) \phi'(\mathcal{U}_{j_0}^{\ell}) (\mathcal{U}_{j_0}^{\ell})^2 (\mathcal{V}_{i_0}^{\ell+1})^2 \cdots \phi'(\mathcal{U}_{i_k}^{\ell+k}) \phi'(\mathcal{U}_{j_k}^{\ell+k}) (\mathcal{U}_{j_k}^{\ell+k})^2 (\mathcal{V}_{i_k}^{\ell+k+1})^2, \end{aligned}$$

And the expectation can be decomposed into products as follows:

$$\begin{aligned}
 \mathbb{E}\left[\prod_{p=0}^k \mathcal{N}_{\delta}^{\ell+p} \mathcal{N}_x^{\ell+p}\right] &= \frac{\sigma_w^2}{n_{\ell-1}} \frac{\sigma_w^2}{n_{\ell+k}} \prod_{j=0}^{k-1} \left(\frac{\sigma_w^2}{n_{\ell+j}}\right)^2 \sum_{i_0=1}^{n_{\ell}} \sum_{j_0=1}^{n_{\ell}} \dots \\
 &\dots \sum_{i_k=1}^{n_{\ell+k}} \sum_{j_k=1}^{n_{\ell+k}} \mathbb{E}[\phi'(\mathcal{U}_{i_0}^{\ell}) \phi'(\mathcal{U}_{j_0}^{\ell}) (\mathcal{U}_{j_0}^{\ell})^2] \mathbb{E}[(\mathcal{V}_{i_k}^{\ell+k+1})^2] \cdot \prod_{p=1}^k \mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \phi'(\mathcal{U}_{i_p}^{\ell+p}) \phi'(\mathcal{U}_{j_p}^{\ell+p}) (\mathcal{U}_{j_p}^{\ell+p})^2] \\
 &= \frac{n_{\ell}}{4} \left(1 + \frac{1}{n_{\ell}}\right) \frac{\sigma_w^2}{n_{\ell-1}} \frac{\sigma_w^2}{n_{\ell+k}} \prod_{j=0}^{k-1} \left(\frac{\sigma_w^2}{n_{\ell+j}}\right)^2 \sum_{i_0=1}^{n_{\ell}} \dots \sum_{i_k=1}^{n_{\ell+k}} \sum_{j_k=1}^{n_{\ell+k}} \prod_{p=1}^k \mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \phi'(\mathcal{U}_{i_p}^{\ell+p}) \phi'(\mathcal{U}_{j_p}^{\ell+p}) (\mathcal{U}_{j_p}^{\ell+p})^2],
 \end{aligned}$$

where we used that  $\mathbb{E}[(\mathcal{V}_{i_k}^{\ell+k+1})^2] = 1$  and  $\sum_{j_0=1}^{n_{\ell}} \mathbb{E}[\phi'(\mathcal{U}_{i_0}^{\ell}) \phi'(\mathcal{U}_{j_0}^{\ell}) (\mathcal{U}_{j_0}^{\ell})^2] = \frac{1}{4}(n_{\ell} - 1) + \frac{1}{2} = \frac{n_{\ell}}{4} \left(1 + \frac{1}{n_{\ell}}\right)$ . If  $j_p \neq i_p$ , we also know that the terms in  $\mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \phi'(\mathcal{U}_{i_p}^{\ell+p}) \phi'(\mathcal{U}_{j_p}^{\ell+p}) (\mathcal{U}_{j_p}^{\ell+p})^2]$  are only dependent through  $\{\mathbf{W}_{i_p, i_{p-1}}^{\ell+p}, \mathbf{W}_{j_p, i_{p-1}}^{\ell+p}\}$ . Same as in Part 1, we can condition the product on these weights:

$$\begin{aligned}
 \mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \phi'(\mathcal{U}_{i_p}^{\ell+p}) \phi'(\mathcal{U}_{j_p}^{\ell+p}) (\mathcal{U}_{j_p}^{\ell+p})^2] &= \mathbb{E}[\mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \mid w_{i_p, i_{p-1}}, w_{j_p, i_{p-1}}] \cdot \mathbb{E}[\phi'(\mathcal{U}_{i_p}^{\ell+p}) \mid w_{i_p, i_{p-1}}] \\
 &\quad \cdot \mathbb{E}[\phi'(\mathcal{U}_{j_p}^{\ell+p}) \mid w_{j_p, i_{p-1}}] \cdot \mathbb{E}[(\mathcal{U}_{j_p}^{\ell+p})^2 \mid w_{j_p, i_{p-1}}]].
 \end{aligned}$$

And we can again write the conditional expectations in case  $j_p \neq i_p$  as follows:

$$\begin{aligned}
 \mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \mid w_{i_p, i_{p-1}}, w_{j_p, i_{p-1}}] &= 1 - b_{i_p}^2 (1 - w_{i_p, i_{p-1}}^2) - b_{j_p}^2 (1 - w_{j_p, i_{p-1}}^2) + 2b_{i_p} b_{j_p} w_{i_p, i_{p-1}} w_{j_p, i_{p-1}}, \\
 \mathbb{E}[\phi'(\mathcal{U}_{i_p}^{\ell+p}) \mid w_{i_p, i_{p-1}}] &= \Phi\left(\frac{w_{i_p, i_{p-1}} a_{i_{p-1}}}{\sqrt{1 - a_{i_{p-1}}^2}}\right), \\
 \mathbb{E}[\phi'(\mathcal{U}_{j_p}^{\ell+p}) \mid w_{j_p, i_{p-1}}] &= \Phi\left(\frac{w_{j_p, i_{p-1}} a_{i_{p-1}}}{\sqrt{1 - a_{i_{p-1}}^2}}\right), \\
 \mathbb{E}[(\mathcal{U}_{j_p}^{\ell+p})^2 \mid w_{j_p, i_{p-1}}] &= 1 - a_{i_{p-1}}^2 (1 - w_{j_p, i_{p-1}}^2).
 \end{aligned}$$

To calculate the expectation of the product here we will need to use that  $\mathbb{E}[\Phi(A \cdot w)] = \mathbb{E}[w^2 \Phi(A \cdot w)] = \frac{1}{2}$ , which we already computed in Part 1. One can also easily see that  $\mathbb{E}[w^4 \Phi(A \cdot w)] = \frac{3}{2}$ . The other expectations involved in the product can be calculated as follows:

$$\begin{aligned}
 \mathbb{E}[w \Phi(A \cdot w)] &= \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\infty} w \operatorname{erf}\left(\frac{A}{\sqrt{2}} w\right) \exp\left(-\frac{w^2}{2}\right) dw = \frac{1}{\sqrt{2\pi}} \frac{A}{\sqrt{A^2 + 1}}, \\
 \mathbb{E}[w^3 \Phi(A \cdot w)] &= \frac{1}{2\sqrt{2\pi}} \int_{-\infty}^{\infty} w^3 \operatorname{erf}\left(\frac{A}{\sqrt{2}} w\right) \exp\left(-\frac{w^2}{2}\right) dw = \frac{1}{\sqrt{2\pi}} \frac{A}{\sqrt{A^2 + 1}} \left(2 + \frac{1}{A^2 + 1}\right).
 \end{aligned}$$

Expressions for the integrals above can be found e.g. in (Korotkov & Korotkov, 2020). Using all the above expressions, we can obtain the following expression for the considered expectation in case  $j_p \neq i_p$ :

$$\begin{aligned}
 \mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \phi'(\mathcal{U}_{i_p}^{\ell+p}) \phi'(\mathcal{U}_{j_p}^{\ell+p}) (\mathcal{U}_{j_p}^{\ell+p})^2] &= \frac{1}{4}(1 - b_{i_p}^2 - b_{j_p}^2)(1 - a_{i_{p-1}}^2) + \frac{1}{4}(b_{i_p}^2 + b_{j_p}^2)(1 - a_{i_{p-1}}^2) \\
 &\quad + \frac{1}{4}(1 - b_{i_p}^2 - b_{j_p}^2)a_{i_{p-1}}^2 + \frac{1}{4}b_{i_p}^2 a_{i_{p-1}}^2 + \frac{3}{4}b_{j_p}^2 a_{i_{p-1}}^2 \\
 &\quad + 2b_{i_p} b_{j_p} (1 - a_{i_{p-1}}^2) \frac{4A^2}{A^2 + 1} + 2b_{i_p} b_{j_p} a_{i_{p-1}}^2 \frac{4A^2}{A^2 + 1} \left(2 + \frac{1}{A^2 + 1}\right) \\
 &= \frac{1}{4} + \frac{1}{2}b_{j_p}^2 a_{i_{p-1}}^2 + 8b_{i_p} b_{j_p} a_{i_{p-1}}^2 (1 + 2a_{i_{p-1}}^2 - a_{i_{p-1}}^4)
 \end{aligned}$$

On the other hand, if  $j_p = i_p$  we have  $\phi'(\mathcal{U}_{i_p}^{\ell+p}) \phi'(\mathcal{U}_{j_p}^{\ell+p}) (\mathcal{U}_{j_p}^{\ell+p})^2 = \phi'(\mathcal{U}_{i_p}^{\ell+p}) (\mathcal{U}_{i_p}^{\ell+p})^2$  and therefore the expectation is given by:

$$\mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \phi'(\mathcal{U}_{i_p}^{\ell+p}) (\mathcal{U}_{i_p}^{\ell+p})^2] = \frac{1}{2} + a_{i_{p-1}}^2 b_{i_p}^2.$$



We now notice that index  $j_p$  appears only in one expectation term in the product for each  $p$ . Therefore, we can sum over  $j_p$  independently for all  $p$ :

$$\begin{aligned} \sum_{j_p=1}^{n_{\ell+p}} \mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \phi'(U_{i_p}^{\ell+p}) \phi'(U_{j_p}^{\ell+p}) (U_{j_p}^{\ell+p})^2] &= \frac{n_{\ell+p}}{4} \left(1 + \frac{1}{n_{\ell+p}}\right) + \frac{1}{2} a_{i_{p-1}}^2 + \frac{1}{2} a_{i_{p-1}}^2 b_{i_p}^2 \\ &+ \left(\sum_{j_p=1}^{n_{\ell+p}} b_{j_p}\right) 8b_{i_p} a_{i_{p-1}}^2 (1 + 2a_{i_{p-1}}^2 - a_{i_{p-1}}^4). \end{aligned}$$

On the other hand, we need to sum over  $i_{p-1}$  values sequentially over different values of  $p$ . First, we can calculate the following sum:

$$\begin{aligned} \sum_{i_{p-1}=1}^{n_{\ell+p-1}} \sum_{j_p=1}^{n_{\ell+p}} \mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \phi'(U_{i_p}^{\ell+p}) \phi'(U_{j_p}^{\ell+p}) (U_{j_p}^{\ell+p})^2] &= \frac{n_{\ell+p-1} n_{\ell+p}}{4} \left(1 + \frac{1}{n_{\ell+p}}\right) + \frac{1}{2} + \frac{1}{2} b_{i_p}^2 \\ &+ \left(\sum_{j_p=1}^{n_{\ell+p}} b_{j_p}\right) 8b_{i_p} \left(1 + 2 \sum_{i_{p-1}=1}^{n_{\ell+p-1}} a_{i_{p-1}}^4 - \sum_{i_{p-1}=1}^{n_{\ell+p-1}} a_{i_{p-1}}^6\right). \end{aligned}$$

Then we can obtain the following bounds for the sum of  $b_{j_p}$  and  $a_{i_{p-1}}$  given by Hölder's inequality:

$$\begin{aligned} \left| \sum_{j_p=1}^{n_{\ell+p}} b_{j_p} \right| &\leq \sum_{j_p=1}^{n_{\ell+p}} |b_{j_p}| = \|\mathbf{b}\|_1 \leq \sqrt{n_{\ell+p}} \|\mathbf{b}^2\| = \sqrt{n_{\ell+p}}, \\ 0 &\leq 1 + 2 \sum_{i_{p-1}=1}^{n_{\ell+p-1}} a_{i_{p-1}}^4 - \sum_{i_{p-1}=1}^{n_{\ell+p-1}} a_{i_{p-1}}^6 \leq 3, \\ |b_{i_p}| &\leq 1, \quad b_{i_p}^2 \leq 1. \end{aligned}$$

Therefore, we can rewrite the previous sum as

$$\sum_{i_{p-1}=1}^{n_{\ell+p-1}} \sum_{j_p=1}^{n_{\ell+p}} \mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \phi'(U_{i_p}^{\ell+p}) \phi'(U_{j_p}^{\ell+p}) (U_{j_p}^{\ell+p})^2] = \frac{n_{\ell+p-1} n_{\ell+p}}{4} \left(1 + \frac{1}{n_{\ell+p}} + O\left(\frac{1}{M^{3/2}}\right)\right)$$

Finally, for the expectation of the whole product we have:

$$\begin{aligned} \mathbb{E}\left[\prod_{p=0}^k \mathcal{N}_{\delta}^{\ell+p} \mathcal{N}_x^{\ell+p}\right] &= \frac{\sigma_w^2}{n_{\ell-1}} \frac{\sigma_w^2}{n_{\ell+k}} \prod_{j=0}^{k-1} \left(\frac{\sigma_w^2}{n_{\ell+j}}\right)^2 \frac{n_{\ell} n_{\ell+k}}{4} \left(1 + \frac{1}{n_{\ell}}\right) \prod_{p=1}^k \frac{n_{\ell+p-1} n_{\ell+p}}{4} \left(1 + \frac{1}{n_{\ell+p}} + O\left(\frac{1}{M^{3/2}}\right)\right) \\ &= \left(\frac{\sigma_w^2}{2}\right)^{2(k+1)} \frac{n_{\ell+k}}{n_{\ell-1}} \prod_{p=0}^k \left(1 + \frac{1}{n_{\ell+p}} + O\left(\frac{1}{M^{3/2}}\right)\right) \\ &= \prod_{p=0}^k \mathbb{E}[\mathcal{N}_{\delta}^{\ell+p}] \mathbb{E}[\mathcal{N}_x^{\ell+p}] \left(1 + \frac{1}{n_{\ell+p}} + O\left(\frac{1}{M^{3/2}}\right)\right) \end{aligned}$$

**Part 3.** Finally, we consider the expectation of a product of squared ratios of the backpropagated errors. In a single layer we have:

$$(\mathcal{N}_{\delta}^{\ell})^2 = \left(\frac{\sigma_w^2}{n_{\ell}}\right)^2 \sum_{i=1}^{n_{\ell}} \sum_{j=1}^{n_{\ell}} \phi'(U_i^{\ell}) \phi'(U_j^{\ell}) (\mathcal{V}_i^{\ell+1})^2 (\mathcal{V}_j^{\ell+1})^2.$$

And for the product in multiple layers we have:

$$\begin{aligned} \prod_{p=0}^k (\mathcal{N}_\delta^{\ell+p})^2 &= \prod_{p=0}^k \left( \frac{\sigma_w^2}{n_{\ell+j}} \right)^2 \sum_{i_0=1}^{n_\ell} \sum_{j_0=1}^{n_\ell} \dots \\ &\quad \dots \sum_{i_k=1}^{n_{\ell+k}} \sum_{j_k=1}^{n_{\ell+k}} \phi'(\mathcal{U}_{i_0}^\ell) \phi'(\mathcal{U}_{j_0}^\ell) (\mathcal{V}_{i_0}^{\ell+1})^2 (\mathcal{V}_{j_0}^{\ell+1})^2 \dots \phi'(\mathcal{U}_{i_k}^{\ell+k}) \phi'(\mathcal{U}_{j_k}^{\ell+k}) (\mathcal{V}_{i_k}^{\ell+k+1})^2 (\mathcal{V}_{j_k}^{\ell+k+1})^2, \\ \mathbb{E} \left[ \prod_{p=0}^k (\mathcal{N}_\delta^{\ell+p})^2 \right] &= \prod_{p=0}^k \left( \frac{\sigma_w^2}{n_{\ell+p}} \right)^2 \sum_{i_0=1}^{n_\ell} \sum_{j_0=1}^{n_\ell} \dots \\ &\quad \dots \sum_{i_k=1}^{n_{\ell+k}} \sum_{j_k=1}^{n_{\ell+k}} \mathbb{E}[\phi'(\mathcal{U}_{i_0}^\ell) \phi'(\mathcal{U}_{j_0}^\ell)] \mathbb{E}[(\mathcal{V}_{i_k}^{\ell+k+1})^2 (\mathcal{V}_{j_k}^{\ell+k+1})^2] \cdot \prod_{p=1}^k \mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 (\mathcal{V}_{j_{p-1}}^{\ell+p})^2 \phi'(\mathcal{U}_{i_p}^{\ell+p}) \phi'(\mathcal{U}_{j_p}^{\ell+p})] \end{aligned}$$

Here the expectations under product are more complicated since the variables in  $\mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 (\mathcal{V}_{j_{p-1}}^{\ell+p})^2 \phi'(\mathcal{U}_{i_p}^{\ell+p}) \phi'(\mathcal{U}_{j_p}^{\ell+p})]$  are dependent through  $\{\mathbf{W}_{i_p i_{p-1}}^{\ell+p}, \mathbf{W}_{j_p i_{p-1}}^{\ell+p}, \mathbf{W}_{i_p j_{p-1}}^{\ell+p}, \mathbf{W}_{j_p j_{p-1}}^{\ell+p}\}$ . Nevertheless, we can still decompose the expectation as before into the following terms:

$$\begin{aligned} \mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 (\mathcal{V}_{j_{p-1}}^{\ell+p})^2 \phi'(\mathcal{U}_{i_p}^{\ell+p}) \phi'(\mathcal{U}_{j_p}^{\ell+p})] &= \mathbb{E}[\mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \mid w_{i_p i_{p-1}}, w_{j_p i_{p-1}}] \cdot \mathbb{E}[(\mathcal{V}_{j_{p-1}}^{\ell+p})^2 \mid w_{i_p j_{p-1}}, w_{j_p j_{p-1}}] \\ &\quad \mathbb{E}[\phi'(\mathcal{U}_{i_p}^{\ell+p}) \mid w_{i_p i_{p-1}}, w_{i_p j_{p-1}}] \cdot \mathbb{E}[\phi'(\mathcal{U}_{j_p}^{\ell+p}) \mid w_{j_p i_{p-1}}, w_{j_p j_{p-1}}]]. \end{aligned}$$

And each conditional expectations can again be calculated explicitly:

$$\begin{aligned} \mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 \mid w_{i_p i_{p-1}}, w_{j_p i_{p-1}}] &= 1 - b_{i_p}^2 (1 - w_{i_p i_{p-1}}^2) - b_{j_p}^2 (1 - w_{j_p i_{p-1}}^2) + 2b_{i_p} b_{j_p} w_{i_p i_{p-1}} w_{j_p i_{p-1}}, \\ \mathbb{E}[(\mathcal{V}_{j_{p-1}}^{\ell+p})^2 \mid w_{i_p j_{p-1}}, w_{j_p j_{p-1}}] &= 1 - b_{i_p}^2 (1 - w_{i_p j_{p-1}}^2) - b_{j_p}^2 (1 - w_{j_p j_{p-1}}^2) + 2b_{i_p} b_{j_p} w_{i_p j_{p-1}} w_{j_p j_{p-1}}, \\ \mathbb{E}[\phi'(\mathcal{U}_{i_p}^{\ell+p}) \mid w_{i_p i_{p-1}}, w_{i_p j_{p-1}}] &= \Phi \left( \frac{w_{i_p i_{p-1}} a_{i_{p-1}} + w_{i_p j_{p-1}} a_{j_{p-1}}}{\sqrt{1 - a_{i_{p-1}}^2 - a_{j_{p-1}}^2}} \right), \\ \mathbb{E}[\phi'(\mathcal{U}_{j_p}^{\ell+p}) \mid w_{j_p i_{p-1}}, w_{j_p j_{p-1}}] &= \Phi \left( \frac{w_{j_p i_{p-1}} a_{i_{p-1}} + w_{j_p j_{p-1}} a_{j_{p-1}}}{\sqrt{1 - a_{i_{p-1}}^2 - a_{j_{p-1}}^2}} \right). \end{aligned}$$

We open the expectation using the following expressions, which, as before, are integrals involving the error function computed e.g. in (Korotkov & Korotkov, 2020):

$$\begin{aligned} \mathbb{E}[\Phi(A_i w_i + A_j w_j)] &= \mathbb{E}[w_i^2 \Phi(A_i w_i + A_j w_j)] = \frac{1}{2}, \\ \mathbb{E} \left[ w_i \Phi \left( \frac{w_i a_{i_{p-1}} + w_j a_{j_{p-1}}}{\sqrt{1 - a_{i_{p-1}}^2 - a_{j_{p-1}}^2}} \right) \right] &= \sqrt{\frac{1}{2\pi}} a_{i_{p-1}}, \\ \mathbb{E} \left[ w_i w_j \Phi \left( \frac{w_i a_{i_{p-1}} + w_j a_{j_{p-1}}}{\sqrt{1 - a_{i_{p-1}}^2 - a_{j_{p-1}}^2}} \right) \right] &= 0, \\ \mathbb{E} \left[ w_i^2 w_j^2 \Phi \left( \frac{w_i a_{i_{p-1}} + w_j a_{j_{p-1}}}{\sqrt{1 - a_{i_{p-1}}^2 - a_{j_{p-1}}^2}} \right) \right] &= \frac{1}{2}, \\ \mathbb{E} \left[ w_i w_j^2 \Phi \left( \frac{w_i a_{i_{p-1}} + w_j a_{j_{p-1}}}{\sqrt{1 - a_{i_{p-1}}^2 - a_{j_{p-1}}^2}} \right) \right] &= \frac{1}{\sqrt{2\pi}} a_{i_{p-1}} (1 - a_{j_{p-1}}^2). \end{aligned}$$

Using all of the above, we get the following expression for the expectation in case  $i_{p-1} \neq j_{p-1}$ :

$$\mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 (\mathcal{V}_{j_{p-1}}^{\ell+p})^2 \phi'(\mathcal{U}_{i_p}^{\ell+p}) \phi'(\mathcal{U}_{j_p}^{\ell+p})] = \begin{cases} \frac{1}{2} & i_p = j_p, \\ \frac{1}{4} + \frac{1}{\pi} b_{i_p} b_{j_p} (a_{i_{p-1}}^2 + a_{j_{p-1}}^2 - 2a_{i_{p-1}}^2 a_{j_{p-1}}^2 (b_{i_p}^2 + b_{j_p}^2)) & i_p \neq j_p, \end{cases}$$

In case  $i_{p-1} = j_{p-1}$ , we have

$$\begin{aligned}\mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^4 | w_{i_p i_{p-1}}, w_{j_p i_{p-1}}] &= 1 - b_{i_p}^2 (1 - w_{i_p i_{p-1}}^2) - b_{j_p}^2 (1 - w_{j_p i_{p-1}}^2) + 2b_{i_p} b_{j_p} w_{i_p i_{p-1}} w_{j_p i_{p-1}}, \\ \mathbb{E}[\phi'(\mathcal{U}_{i_p}^{\ell+p}) | w_{i_p i_{p-1}}] &= \Phi\left(\frac{w_{i_p i_{p-1}} a_{i_{p-1}}}{\sqrt{1 - a_{i_{p-1}}^2}}\right), \\ \mathbb{E}[\phi'(\mathcal{U}_{j_p}^{\ell+p}) | w_{j_p i_{p-1}}] &= \Phi\left(\frac{w_{j_p i_{p-1}} a_{i_{p-1}}}{\sqrt{1 - a_{i_{p-1}}^2}}\right).\end{aligned}$$

Therefore, the expectation in this case is given by:

$$\mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^4 \phi'(\mathcal{U}_{i_p}^{\ell+p}) \phi'(\mathcal{U}_{j_p}^{\ell+p})] = \begin{cases} \frac{3}{2} & i_p = j_p, \\ \frac{3}{4} + \frac{2}{\pi} b_{i_p} b_{j_p} a_{i_{p-1}}^2 (3 - a_{i_{p-1}}^2 (b_{i_p}^2 + b_{j_p}^2)) & i_p \neq j_p. \end{cases}$$

To compute the sum, we now notice that equality of indices in one layer ( $i_p = j_p$ ) amounts to multiplying the product  $\prod_{p=1}^k \mathbb{E}[(\mathcal{V}_{i_{p-1}}^{\ell+p})^2 (\mathcal{V}_{j_{p-1}}^{\ell+p})^2 \phi'(\mathcal{U}_{i_p}^{\ell+p}) \phi'(\mathcal{U}_{j_p}^{\ell+p})]$  by  $6 + O(1/M^{3/2})$  and for every pair of indices  $(i_p, j_p)$  there are only  $n_{\ell+p}$  summands with this multiplier and  $n_{\ell+p}(n_{\ell+p} - 1)$  summands without it. We can also see that if we computed the sum with all the pairs of indices not equal, we would get  $\prod_{p=0}^k (n_{\ell+p}/2)^2 (1 - 1/n_{\ell+p} + O(1/M^{3/2}))$ . Therefore, we get the desired expression for the expectation of the product:

$$\begin{aligned}\mathbb{E}\left[\prod_{p=0}^k (\mathcal{N}_\delta^{\ell+p})^2\right] &= \prod_{p=0}^k \left(\frac{\sigma_w^2}{n_{\ell+p}}\right)^2 \prod_{p=0}^k \left(\frac{n_{\ell+p}}{2}\right)^2 \left(1 - \frac{1}{n_{\ell+p}} + \frac{6}{n_{\ell+p}} + O\left(\frac{1}{M^{3/2}}\right)\right) \\ &= \prod_{p=0}^k \left(\frac{\sigma_w^2}{2}\right)^2 \left(1 + \frac{5}{n_{\ell+p}} + O\left(\frac{1}{M^{3/2}}\right)\right) \\ &= \prod_{p=0}^k \mathbb{E}[(\mathcal{N}_\delta^{\ell+p})^2] \left(1 + O\left(\frac{1}{M^{3/2}}\right)\right).\end{aligned}$$

□

**Lemma A.5** (Dispersion of  $\Theta_W(x, x)$ ). *Consider a fully-connected DNN of depth  $L$  defined in (3) initialized as in (6). The input dimension is given by  $n_0 = \alpha_0 M$ , the output dimension is 1, and the hidden layers have constant width  $M$ . The activation function in the hidden layers is ReLU, i.e.  $\phi(x) = x \mathbb{1}\{x > 0\}$ , and the output layer is linear. Assume also that the biases are initialized to zero, i.e.  $\sigma_b = 0$ , and the input data is normalized. Then the component of the NTK corresponding to the weights  $\Theta_W(x, x) := \sum_{\ell=1}^L \sum_{ij} \left(\frac{\partial f(x)}{\partial \mathbf{W}_{ij}^\ell}\right)^2$  has the following properties at initialization:*

$$\mathbb{E}[\Theta_W(x, x)] = \left(\frac{\sigma_w^2}{2}\right)^{L-1} \left(1 + \frac{M}{n_0} (L-1)\right), \quad (31)$$

$$\frac{\mathbb{E}[\Theta_W^2(x, x)]}{\mathbb{E}^2[\Theta_W(x, x)]} \xrightarrow[\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}]{\quad} \frac{1}{2\lambda} e^{5\lambda} \left(1 - \frac{1}{4\lambda} (1 - e^{-4\lambda})\right). \quad (32)$$

*Proof.* Using backpropagation formulas for the gradients, we can rewrite the NTK as follows:

$$\begin{aligned}
 \Theta_W(x, x) &= \sum_{\ell=1}^L \sum_{i=1}^{n_\ell} \sum_{j=1}^{n_{\ell-1}} (\boldsymbol{\delta}_i^\ell)^2 (\mathbf{x}_j^{\ell-1})^2 \\
 &= \sum_{\ell=1}^L \|\boldsymbol{\delta}^\ell \times \mathbf{x}^{\ell-1}\|^2 = \sum_{\ell=1}^L \|\boldsymbol{\delta}^\ell\|^2 \|\mathbf{x}^{\ell-1}\|^2 \\
 &= \sum_{\ell=1}^L \|\boldsymbol{\delta}^L\|^2 \|\mathbf{x}^0\|^2 \prod_{j=\ell}^{L-1} \frac{\|\boldsymbol{\delta}^j\|^2}{\|\boldsymbol{\delta}^{j+1}\|^2} \prod_{k=1}^{\ell-1} \frac{\|\mathbf{x}^k\|^2}{\|\mathbf{x}^{k-1}\|^2} \\
 &= \sum_{\ell=1}^L \|\boldsymbol{\delta}^L\|^2 \|\mathbf{x}^0\|^2 \prod_{j=\ell}^{L-1} \mathcal{N}_\delta^j \prod_{k=1}^{\ell-1} \mathcal{N}_x^k.
 \end{aligned}$$

Here for the simplicity of notation we omit the dependence on  $\boldsymbol{\delta}^\ell$  and  $\mathbf{x}^{\ell-1}$  on the input  $x$ .

If the last layer has a linear activation and the input data is normalized, we also have that  $\|\boldsymbol{\delta}^L\|^2 \|\mathbf{x}^0\|^2 = 1$ . Then, using the results about expectations of  $\mathcal{N}_\delta^\ell$  and  $\mathcal{N}_x^\ell$  from Lemma A.1 and Lemma A.2, as well as the results about correlations from Lemma A.4, we can write the following for the expectation of  $\Theta_W(x, x)$ :

$$\mathbb{E}[\Theta_W(x, x)] = \sum_{\ell=1}^L \prod_{j=\ell}^{L-1} \frac{\sigma_w^2}{2} \prod_{k=1}^{\ell-1} \frac{\sigma_w^2}{2} \frac{n_k}{n_{k-1}} = \left(\frac{\sigma_w^2}{2}\right)^{L-1} \sum_{\ell=1}^L \frac{n_{\ell-1}}{n_0}$$

And for constant width of hidden layers, i.e.  $n_\ell = M, \ell = 1, \dots, L-1$ , this simplifies to

$$\mathbb{E}[\Theta_W(x, x)] = \left(\frac{\sigma_w^2}{2}\right)^{L-1} \left(1 + \frac{M}{n_0}(L-1)\right) \propto \left(\frac{\sigma_w^2}{2}\right)^L \frac{ML}{n_0}$$

Now we consider the second moment of the NTK, which is given by:

$$\begin{aligned}
 \mathbb{E}[\Theta_W^2(x, x)] &= \sum_{\ell=1}^L \mathbb{E}[\theta_\ell^2] + 2 \sum_{1 \leq \ell_1 < \ell_2 \leq L} \mathbb{E}[\theta_{\ell_1} \theta_{\ell_2}], \\
 \theta_\ell &= \prod_{j=\ell}^{L-1} \mathcal{N}_\delta^j \prod_{k=1}^{\ell-1} \mathcal{N}_x^k, \quad \ell = 1, \dots, L.
 \end{aligned}$$

We can open the expectation of the squared terms defined above as follows:

$$\begin{aligned}
 \mathbb{E}[\theta_\ell^2] &= \mathbb{E}\left[\prod_{j=\ell}^{L-1} (\mathcal{N}_\delta^j)^2\right] \mathbb{E}\left[\prod_{k=1}^{\ell-1} (\mathcal{N}_x^k)^2\right] \\
 &= \left(\frac{\sigma_w^2}{2}\right)^{2(L-1)} \left(\frac{n_{\ell-1}}{n_0}\right)^2 \prod_{j=\ell}^{L-1} \left(1 + \frac{5}{n_j} + O\left(\frac{1}{M^{3/2}}\right)\right) \prod_{k=1}^{\ell-1} \left(1 + \frac{5}{n_k}\right),
 \end{aligned}$$

which simplifies to the following expressions in case of constant width  $M$ :

$$\begin{aligned}
 \mathbb{E}[\theta_\ell^2] &= \left(\frac{\sigma_w^2}{2}\right)^{2(L-1)} \left(\frac{M}{n_0}\right)^2 \left(1 + \frac{5}{M} + O\left(\frac{1}{M^{3/2}}\right)\right)^{L-1}, \quad \ell > 1, \\
 \mathbb{E}[\theta_1^2] &= \left(\frac{\sigma_w^2}{2}\right)^{2(L-1)} \left(1 + \frac{5}{M} + O\left(\frac{1}{M^{3/2}}\right)\right)^{L-1}.
 \end{aligned}$$

And the mixed terms with  $1 \leq \ell_1 < \ell_2 \leq L$  can be calculated as follows:

$$\begin{aligned} \mathbb{E}[\theta_{\ell_1} \theta_{\ell_2}] &= \mathbb{E} \left[ \prod_{j=\ell_2}^{L-1} (\mathcal{N}_\delta^j)^2 \right] \mathbb{E} \left[ \prod_{p=\ell_1}^{\ell_2-1} \mathcal{N}_\delta^p \mathcal{N}_x^p \right] \prod_{k=1}^{\ell_1-1} \mathbb{E}[(\mathcal{N}_x^k)^2] \\ &= \left( \frac{\sigma_w^2}{2} \right)^{2(L-1)} \left( \frac{n_{\ell_2-1} n_{\ell_1-1}}{n_0^2} \right) \prod_{j=\ell_2}^{L-1} \left( 1 + \frac{5}{n_j} + O\left( \frac{1}{M^{3/2}} \right) \right) \prod_{p=\ell_1}^{\ell_2-1} \left( 1 + \frac{1}{n_p} + O\left( \frac{1}{M^{3/2}} \right) \right) \prod_{k=1}^{\ell_1-1} \left( 1 + \frac{5}{n_k} \right), \end{aligned}$$

which for constant width  $M$  simplifies to

$$\begin{aligned} \mathbb{E}[\theta_{\ell_1} \theta_{\ell_2}] &= \left( \frac{\sigma_w^2}{2} \right)^{2(L-1)} \left( \frac{M}{n_0} \right)^2 \left( 1 + \frac{5}{M} + O\left( \frac{1}{M^{3/2}} \right) \right)^{L-1-\Delta_\ell} \left( 1 + \frac{1}{M} + O\left( \frac{1}{M^{3/2}} \right) \right)^{\Delta_\ell}, \quad \ell_1 > 1, \\ \mathbb{E}[\theta_1 \theta_{\ell_2}] &= \left( \frac{\sigma_w^2}{2} \right)^{2(L-1)} \left( \frac{M}{n_0} \right) \left( 1 + \frac{5}{M} + O\left( \frac{1}{M^{3/2}} \right) \right)^{L-\ell_2} \left( 1 + \frac{1}{M} + O\left( \frac{1}{M^{3/2}} \right) \right)^{\ell_2-1}. \end{aligned}$$

To make the notation lighter, we will denote  $x := 1 + 5/M + O(M^{-3/2})$ ,  $y := 1 + 1/M + O(M^{-3/2})$ ,  $a := \sigma_w^2/2$  and  $\lambda := L/M$  here and in the following proofs. Then we can rewrite the two sums that comprise the second moment of the NTK as follows:

$$\begin{aligned} \sum_{\ell=1}^L \mathbb{E}[\theta_\ell^2] &= a^{2(L-1)} x^{L-1} \left( \frac{M^2}{n_0^2} (L-1) + 1 \right) \\ &= a^{2(L-1)} \frac{M^2 L^2}{n_0^2} \left[ x^{L-1} \frac{1}{\lambda M} + O\left( \frac{1}{M^{3/2}} \right) \right] = a^{2(L-1)} \frac{M^2 L^2}{n_0^2} \left[ x^{L-1} \frac{1}{\lambda M} + O\left( \frac{1}{M^{3/2}} \right) \right], \\ \sum_{1 \leq \ell_1 < \ell_2 \leq L} \mathbb{E}[\theta_{\ell_1} \theta_{\ell_2}] &= a^{2(L-1)} \frac{M^2}{n_0^2} \sum_{\Delta_\ell=1}^{L-2} (L-1-\Delta_\ell) x^{L-1-\Delta_\ell} y^{\Delta_\ell} + a^{2(L-1)} \frac{M}{n_0} \sum_{\ell_2=2}^L x^{L-\ell_2} y^{\ell_2-1} \\ &= a^{2(L-1)} \frac{M^4}{16n_0^2} ((L-2)yx^L - (L-1)y^2x^{L-1} + xy^L) + a^{2(L-1)} \frac{M^2}{4n_0} (yx^{L-1} - y^L) \\ &= a^{2(L-1)} \frac{M^2 L^2}{n_0^2} \left[ x^L \left( \frac{1}{4\lambda} \left( 1 - \frac{5}{M} \right) - \frac{1}{16\lambda^2} \left( 1 + \frac{5-4\alpha_0}{M} \right) + O\left( \frac{1}{M^{3/2}} \right) \right) + \right. \\ &\quad \left. + y^L \left( \frac{1}{16\lambda^2} \left( 1 + \frac{5-4\alpha_0}{M} \right) + O\left( \frac{1}{M^{3/2}} \right) \right) \right] \end{aligned}$$

Therefore, the complete expression for the second moment of  $\Theta_W(x, x)$  is given by:

$$\begin{aligned} \sum_{\ell=1}^L \mathbb{E}[\theta_\ell^2] + 2 \sum_{1 \leq \ell_1 < \ell_2 \leq L} \mathbb{E}[\theta_{\ell_1} \theta_{\ell_2}] &= a^{2(L-1)} \frac{M^2 L^2}{n_0^2} \left[ x^L \left( \frac{1}{2\lambda} \left( 1 - \frac{3}{M} \right) - \frac{1}{8\lambda^2} \left( 1 + \frac{5-4\alpha_0}{M} \right) + O\left( \frac{1}{M^{3/2}} \right) \right) + \right. \\ &\quad \left. + y^L \left( \frac{1}{8\lambda^2} \left( 1 + \frac{5-4\alpha_0}{M} \right) + O\left( \frac{1}{M^{3/2}} \right) \right) \right] \\ &\quad a^{2(L-1)} \frac{M^2 L^2}{n_0^2} \left[ x^L \left( \frac{1}{2\lambda} - \frac{1}{8\lambda^2} + O\left( \frac{1}{M} \right) \right) + y^L \frac{1}{8\lambda^2} + O\left( \frac{1}{M} \right) \right] \end{aligned}$$

One can see that in the limit  $L \rightarrow \infty$ ,  $M \rightarrow \infty$ ,  $L/M \rightarrow \lambda \in \mathbb{R}$ , we have  $x^L \rightarrow e^{5\lambda}$  and  $y^L \rightarrow e^\lambda$ . Therefore, we can find the limit of the desired ratio:

$$\frac{\mathbb{E}[\Theta_W^2(x, x)]}{\left( \frac{\sigma_w^2}{2} \right)^{2(L-1)} \frac{L^2 M^2}{n_0^2}} \xrightarrow[L/M \rightarrow \lambda \in \mathbb{R}]{M \rightarrow \infty, L \rightarrow \infty} \frac{1}{2\lambda} e^{5\lambda} \left( 1 - \frac{1}{4\lambda} (1 - e^{-4\lambda}) \right).$$

□

**Lemma A.6** (Dispersion of  $\Theta_b(x, x)$ ). *Consider a fully-connected DNN of depth  $L$  defined in (3) initialized as in (6). The input dimension is given by  $n_0 = \alpha_0 M$ , the output dimension is 1, and the hidden layers have constant width  $M$ . The activation function in the hidden layers is ReLU, i.e.  $\phi(x) = x \mathbb{1}\{x > 0\}$ , and the output layer is linear. Assume also that the biases are initialized to zero, i.e.  $\sigma_b = 0$ , and the input data is normalized. Then the component of the NTK corresponding to the biases  $\Theta_b(x, x) := \sum_{\ell=1}^L \sum_i \left( \frac{\partial f(x)}{\partial \mathbf{b}_i^\ell} \right)^2$  has the following properties at initialization:*

$$\mathbb{E}[\Theta_b(x, x)] = \begin{cases} \frac{\left(\frac{\sigma_w^2}{2}\right)^L - 1}{\frac{\sigma_w^2}{2} - 1} & \text{if } \frac{\sigma_w^2}{2} \neq 1 \\ L & \text{if } \frac{\sigma_w^2}{2} = 1 \end{cases} \quad (33)$$

$$\frac{\mathbb{E}[\Theta_b^2(x, x)]}{\mathbb{E}^2[\Theta_b(x, x)]} \xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} \begin{cases} 1 & \text{if } \frac{\sigma_w^2}{2} < 1 \\ \frac{2}{25\lambda^2}(e^{5\lambda} - 1) - \frac{2}{5\lambda} & \text{if } \frac{\sigma_w^2}{2} = 1 \\ e^{5\lambda} & \text{if } \frac{\sigma_w^2}{2} > 1 \end{cases} \quad (34)$$

*Proof.* Using backpropagation equations (7), we can obtain the following expression for  $\Theta_b(x, x)$ :

$$\Theta_b(x, x) = \sum_{\ell=1}^L \|\delta^\ell\|^2 = \|\delta^L\|^2 \sum_{\ell=1}^L \prod_{j=\ell}^{L-1} \frac{\|\delta^j\|^2}{\|\delta^{j+1}\|^2} = \sum_{\ell=1}^L \prod_{j=\ell}^{L-1} \mathcal{N}_\delta^j.$$

In this lemma, we will again denote  $a := \sigma_w^2/2$  and  $x := 1 + 5/M + O(1/M^{3/2})$ . And in the following computations, we will need to consider cases with  $a \neq 1$  and  $a = 1$  separately.

**Case 1:**  $a \neq 1$ . In this case the expectation is given by a sum of a geometric progression:

$$\mathbb{E}[\Theta_b(x, x)] = \sum_{\ell=1}^L \prod_{j=\ell}^{L-1} \mathbb{E}[\mathcal{N}_\delta^j] = \sum_{\ell=1}^L a^{L-\ell} = \frac{a^L - 1}{a - 1}$$

And for the second moment we can write:

$$\begin{aligned} \mathbb{E}[\Theta_b(x, x)^2] &= \mathbb{E}\left[\left(\sum_{\ell=1}^L \prod_{j=\ell}^{L-1} \mathcal{N}_\delta^j\right)^2\right] = \sum_{\ell=1}^L \mathbb{E}\left[\prod_{j=\ell}^{L-1} (\mathcal{N}_\delta^j)^2\right] + 2 \sum_{1 \leq \ell_1 < \ell_2 \leq L} \mathbb{E}\left[\prod_{j=\ell_1}^{\ell_2-1} \mathcal{N}_\delta^j \prod_{k=\ell_2}^{L-1} (\mathcal{N}_\delta^k)^2\right] \\ &= \sum_{\ell=1}^L a^{2(L-\ell)} \prod_{j=\ell}^{L-1} \left(1 + \frac{5}{n_j} + O\left(\frac{1}{M^{3/2}}\right)\right) + 2 \sum_{1 \leq \ell_1 < \ell_2 \leq L} a^{2L-\ell_1-\ell_2} \prod_{k=\ell_2}^{L-1} \left(1 + \frac{5}{n_k} + O\left(\frac{1}{M^{3/2}}\right)\right) \end{aligned}$$

For constant width  $M$  the above expression simplifies to the following sum:

$$\begin{aligned} \mathbb{E}[\Theta_b(x, x)^2] &= \sum_{\ell=1}^L a^{2(L-\ell)} x^{L-\ell} + 2 \sum_{1 \leq \ell_1 < \ell_2 \leq L} a^{2L-\ell_1-\ell_2} x^{L-\ell_2} \\ &= \sum_{\ell=1}^L a^{2(L-\ell)} x^{L-\ell} + 2 \sum_{\ell_1=1}^{L-1} a^{2(L-\ell_1)} x^{L-\ell_1} \sum_{\Delta_\ell=1}^{L-\ell_1} a^{-\Delta_\ell} x^{-\Delta_\ell} \end{aligned}$$

And the involved terms can be further calculated explicitly as follows:

$$\begin{aligned}
 \mathbb{E}[\Theta_b(x, x)^2] &= \frac{a^{2L}x^L - 1}{a^2x - 1} + \frac{2}{ax - 1} \sum_{\ell=1}^{L-1} a^{2(L-\ell)} x^{L-\ell} (1 - a^{L-\ell} x^{L-\ell}) \\
 &= \frac{a^{2L}x^L - 1}{a^2x - 1} + \frac{2}{ax - 1} \left( \frac{a^{2L}x^L - 1}{a^2x - 1} - 1 - \frac{a^L - a}{a - 1} \right) \\
 &= \frac{a^{2L}x^L - 1}{a^2x - 1} \frac{ax + 1}{ax - 1} - \frac{2}{ax - 1} \frac{a^L - 1}{a - 1} \\
 &= \frac{1}{(a - 1)^2} \left[ a^{2L}x^L \left( 1 + O\left(\frac{1}{M}\right) \right) - 2a^L \left( 1 + O\left(\frac{1}{M}\right) \right) + 1 + O\left(\frac{1}{M}\right) \right]
 \end{aligned}$$

If  $a < 1$ , the expectation and the second moment have finite limits:

$$\begin{aligned}
 \mathbb{E}[\Theta_b(x, x)] &\xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} \frac{1}{1 - a} \\
 \mathbb{E}[\Theta_b(x, x)^2] &\xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} \frac{-1}{a^2 - 1} \frac{a + 1}{a - 1} - \frac{2}{a - 1} \frac{-1}{a - 1} = \frac{1}{(a - 1)^2}
 \end{aligned}$$

Therefore for  $a < 1$  we have

$$\frac{\mathbb{E}[\Theta_b^2(x, x)]}{\mathbb{E}^2[\Theta_b(x, x)]} \xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} 1$$

On the other hand, if  $a > 1$  then the limits are infinite but there is a finite limit of the ratio:

$$\frac{\mathbb{E}[\Theta_b(x, x)^2]}{a^{2L}/(a - 1)^2} \xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} e^{5\lambda}$$

**Case 2:**  $a = 1$ . In this case, the expectation is just a sum of ones, so we have

$$\mathbb{E}[\Theta_b(x, x)] = \sum_{\ell=1}^L 1 = L$$

And the second moment can be calculated as follows:

$$\begin{aligned}
 \mathbb{E}[\Theta_b(x, x)^2] &= \sum_{\ell=1}^L x^{L-\ell} + 2 \sum_{\ell=1}^{L-1} x^{L-\ell} \sum_{\Delta_\ell=1}^{L-\ell} x^{-\Delta_\ell} = \frac{x^L - 1}{x - 1} + \frac{2}{x - 1} \left( \sum_{\ell=1}^{L-1} x^{L-\ell} - \sum_{\ell=1}^{L-1} 1 \right) \\
 &= \frac{x^L - 1}{x - 1} + \frac{2}{x - 1} \left( \frac{x^L - x}{x - 1} - L + 1 \right) = M^2 \left[ x^L \left( \frac{1}{5M} + \frac{2}{25} \right) - \frac{1}{5M} - \frac{10\lambda + 2}{25} \right] \\
 &= \frac{x^L - 1}{x - 1} + \frac{2}{x - 1} \left( \frac{x^L - x}{x - 1} - L + 1 \right) = L^2 \left[ x^L \left( \frac{2}{25\lambda^2} + O\left(\frac{1}{M}\right) \right) - \frac{2}{5\lambda} - \frac{2}{25\lambda^2} + O\left(\frac{1}{M}\right) \right]
 \end{aligned}$$

Then for the desired ratio we have the following result in the limit:

$$\frac{\mathbb{E}[\Theta_b(x, x)^2]}{L^2} \xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} \frac{2}{25\lambda^2} (e^{5\lambda} - 1) - \frac{2}{5\lambda},$$

which completes the proof for all the cases.  $\square$

**Lemma A.7** (Dispersion of  $\Theta_W(x, x)\Theta_b(x, x)$ ). *Consider a fully-connected DNN of depth  $L$  defined in (3) initialized as in (6). The input dimension is given by  $n_0 = \alpha_0 M$ , the output dimension is 1, and the hidden layers have constant width  $M$ . The activation function in the hidden layers is ReLU, i.e.  $\phi(x) = x\mathbb{1}\{x > 0\}$ , and the output layer is linear. Assume also that the biases are initialized to zero, i.e.  $\sigma_b = 0$ , and the input data is normalized. Then the following statements hold:*

1. In the chaotic phase, i.e. if  $\sigma_w^2/2 > 1$ :

$$\frac{\mathbb{E}[\Theta_W(x, x)\Theta_b(x, x)]}{\mathbb{E}^2[\Theta_W(x, x)]} \xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} 0 \quad (35)$$

2. In the ordered phase, i.e. if  $\sigma_w^2/2 < 1$ :

$$\frac{\mathbb{E}[\Theta_W(x, x)\Theta_b(x, x)]}{\mathbb{E}^2[\Theta_b(x, x)]} \xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} 0 \quad (36)$$

3. At the EOC, i.e. if  $\sigma_w^2/2 = 1$ :

$$\frac{\mathbb{E}[\Theta_W(x, x)\Theta_b(x, x)]}{L^2} \xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} \frac{1}{4\alpha_0} \left( e^{5\lambda} \frac{9}{25\lambda^2} - e^\lambda \frac{1}{\lambda^2} - \frac{4}{5\lambda} + \frac{16}{25\lambda^2} \right) \quad (37)$$

*Proof.* We can decompose  $\Theta_W(x, x)\Theta_b(x, x)$  into telescopic products as follows:

$$\begin{aligned} \Theta_W(x, x)\Theta_b(x, x) &= \sum_{\ell=1}^L \|\delta^\ell\|^2 \|\mathbf{x}^{\ell-1}\|^2 \sum_{\ell'=1}^L \|\delta^{\ell'}\|^2 \\ &= \|\delta^L\|^4 \|\mathbf{x}^0\|^2 \sum_{\ell=1}^L \prod_{j=\ell}^{L-1} \frac{\|\delta^j\|^4}{\|\delta^{j+1}\|^4} \prod_{k=1}^{\ell-1} \frac{\|\mathbf{x}^k\|^2}{\|\mathbf{x}^{k-1}\|^2} \\ &\quad + \|\delta^L\|^4 \|\mathbf{x}^0\|^2 \sum_{1 \leq \ell_1 < \ell_2 \leq L} \prod_{p=\ell_1}^{\ell_2-1} \frac{\|\delta^p\|^2}{\|\delta^{p+1}\|^2} \prod_{j=\ell_2}^{L-1} \frac{\|\delta^j\|^4}{\|\delta^{j+1}\|^4} \prod_{k=1}^{\ell_2-1} \frac{\|\mathbf{x}^k\|^2}{\|\mathbf{x}^{k-1}\|^2} \\ &\quad + \|\delta^L\|^4 \|\mathbf{x}^0\|^2 \sum_{1 \leq \ell_1 < \ell_2 \leq L} \prod_{p=\ell_1}^{\ell_2-1} \frac{\|\delta^p\|^2}{\|\delta^{p+1}\|^2} \prod_{j=\ell_2}^{L-1} \frac{\|\delta^j\|^4}{\|\delta^{j+1}\|^4} \prod_{k=1}^{\ell_1-1} \frac{\|\mathbf{x}^k\|^2}{\|\mathbf{x}^{k-1}\|^2} \end{aligned}$$

Then, as in the previous lemmas, we can calculate the expectation using the results of Lemmas A.1, A.2 and A.4:

$$\begin{aligned} \mathbb{E}[\Theta_W(x, x)\Theta_b(x, x)] &= \sum_{\ell=1}^L \mathbb{E} \left[ \prod_{j=\ell}^{L-1} (\mathcal{N}_\delta^j)^2 \right] \prod_{k=1}^{\ell-1} \mathbb{E}[\mathcal{N}_x^k] + \sum_{1 \leq \ell_1 < \ell_2 \leq L} \prod_{p=\ell_1}^{\ell_2-1} \mathbb{E}[\mathcal{N}_\delta^p \mathcal{N}_x^p] \mathbb{E} \left[ \prod_{j=\ell_2}^{L-1} (\mathcal{N}_\delta^j)^2 \right] \prod_{k=1}^{\ell_1-1} \mathbb{E}[\mathcal{N}_x^k] \\ &\quad + \sum_{1 \leq \ell_1 < \ell_2 \leq L} \prod_{p=\ell_1}^{\ell_2-1} \mathbb{E}[\mathcal{N}_\delta^p] \mathbb{E} \left[ \prod_{j=\ell_2}^{L-1} (\mathcal{N}_\delta^j)^2 \right] \prod_{k=1}^{\ell_1-1} \mathbb{E}[\mathcal{N}_x^k] \\ &= \sum_{\ell=1}^L a^{2L-\ell-1} \frac{n_{\ell-1}}{n_0} \prod_{j=\ell}^{L-1} \left( 1 + \frac{5}{n_j} + \left( \frac{1}{M^{3/2}} \right) \right) \\ &\quad + \sum_{1 \leq \ell_1 < \ell_2 \leq L} a^{2L-\ell_1-1} \frac{n_{\ell_2-1}}{n_0} \prod_{p=\ell_1}^{\ell_2-1} \left( 1 + \frac{1}{n_p} + \left( \frac{1}{M^{3/2}} \right) \right) \prod_{j=\ell_2}^{L-1} \left( 1 + \frac{5}{n_j} + \left( \frac{1}{M^{3/2}} \right) \right) \\ &\quad + \sum_{1 \leq \ell_1 < \ell_2 \leq L} a^{2L-\ell_2-1} \frac{n_{\ell_1-1}}{n_0} \prod_{j=\ell_2}^{L-1} \left( 1 + \frac{5}{n_j} + \left( \frac{1}{M^{3/2}} \right) \right), \end{aligned}$$

where we denoted  $a := \sigma_w^2/2$ . As in Lemma A.6, we will need to consider the cases with  $a \neq 1$  and  $a = 1$  separately here.



**Case 1:**  $a \neq 1$ . For constant width  $M$  the above expression for the expectation simplifies to:

$$\begin{aligned}
 \mathbb{E}[\Theta_W(x, x)\Theta_b(x, x)] &= \sum_{\ell=1}^L \frac{n_{\ell-1}}{n_0} a^{2L-\ell-1} x^{L-\ell} + \sum_{\ell_1=1}^{L-1} a^{2L-\ell_1-1} \sum_{\ell_2=\ell_1+1}^L \frac{n_{\ell_2-1}}{n_0} y^{\ell_2-\ell_1} x^{L-\ell_2} \\
 &+ \sum_{\ell_1=1}^{L-1} \frac{n_{\ell_1-1}}{n_0} \sum_{\ell_2=\ell_1+1}^L a^{2L-\ell_2-1} x^{L-\ell_2} \\
 &= a^{2(L-1)} x^{L-1} \left(1 + \frac{M}{n_0} \frac{1}{ax-1}\right) - \frac{M}{n_0} \frac{a^{L-1}}{ax-1} \\
 &+ a^{2(L-1)} x^{L-1} \frac{M}{n_0} \frac{1}{x-y} \frac{ay}{ax-1} - a^{2(L-1)} y^L \frac{M}{n_0} \frac{1}{x-y} \frac{ay}{ay-1} \\
 &+ a^L \frac{M}{n_0} \frac{1}{x-y} \left(\frac{-xy}{ax-1} + \frac{y^2}{ay-1}\right) \\
 &+ a^{2(L-1)} x^{L-1} \frac{1}{ax-1} \left(1 + \frac{M}{n_0} \frac{1}{ax-1}\right) - \frac{a^{L-1}}{ax-1} - \frac{a^{L-1}}{ax-1} \frac{M}{n_0} (L-2) - \frac{M}{n_0} \frac{a^L x}{(ax-1)^2} \\
 &= a^{2(L-1)} x^{L-1} \left[ \left(1 + \frac{M}{n_0} \frac{1}{ax-1}\right) \left(1 + \frac{1}{ax-1}\right) + \frac{M}{n_0} \frac{1}{x-y} \frac{ay}{ax-1} \right] \\
 &- a^{2(L-1)} y^L \frac{M}{n_0} \frac{1}{x-y} \frac{ay}{ay-1} \\
 &+ a^{L-1} \left[ \frac{M}{n_0} \frac{a}{x-y} \left(\frac{-xy}{ax-1} + \frac{y^2}{ay-1}\right) - \frac{1}{ax-1} \left(1 + \frac{M(L-1)}{n_0}\right) - \frac{M}{n_0} \frac{ax}{(ax-1)^2} \right] \\
 &= a^{2(L-1)} x^{L-1} \frac{M}{4\alpha_0} \frac{a}{ax-1} \left[ y + \frac{4\alpha_0 x}{M} + \frac{4x}{(ax-1)M} + O\left(\frac{1}{M}\right) \right] \\
 &- a^{2(L-1)} y^L \frac{M}{4\alpha_0} \frac{a}{ay-1} \left[ y + O\left(\frac{1}{M}\right) \right] \\
 &+ a^{L-1} \frac{M}{4\alpha_0} \frac{a}{ax-1} \left[ \frac{16}{M^2(ay-1)(ax-1)} - \frac{4\alpha_0}{Ma} \left(1 + \frac{L-1}{\alpha_0}\right) \right] \\
 &= \frac{M}{4\alpha_0} \frac{a}{a-1} \left[ a^{2(L-1)} x^{L-1} - a^{2(L-1)} y^L - 4a^{L-1} \lambda + O\left(\frac{1}{M}\right) \right],
 \end{aligned}$$

where we also denoted  $x := 1 + 5/M + O(1/M^{3/2})$  and  $y := 1 + 1/M + O(1/M^{3/2})$ . From the last expression, we see that  $\mathbb{E}[\Theta_W(x, x)\Theta_b(x, x)]$  tends to zero if  $a < 1$ , therefore in this case we get

$$\frac{\mathbb{E}[\Theta_W(x, x)\Theta_b(x, x)]}{\mathbb{E}^2[\Theta_b(x, x)]} \xrightarrow[M/L \rightarrow \lambda \in \mathbb{R}]{M \rightarrow \infty, L \rightarrow \infty} 0, \quad (38)$$

using the result of Lemma A.6 that  $\mathbb{E}[\Theta_b^2(x, x)]$  has a finite limit when  $a < 1$ .

On the other hand, if  $a > 1$ , we can see that  $\mathbb{E}[\Theta_W(x, x)\Theta_b(x, x)]$  contains polynomials of  $M$  and  $L$  of degree not larger than 1. Therefore, we have

$$\frac{\mathbb{E}[\Theta_W(x, x)\Theta_b(x, x)]}{a^{2L} M^2 L^2 / n_0^2} \xrightarrow[L/M \rightarrow \lambda \in \mathbb{R}]{M \rightarrow \infty, L \rightarrow \infty} 0, \quad (39)$$

which completes the proof for the case when  $a \neq 1$ .

**Case 2:**  $a = 1$ . In this case, the expression for the expectation with constant width  $M$  is given by:

$$\begin{aligned} \mathbb{E}[\Theta_W(x, x)\Theta_b(x, x)] &= \sum_{\ell=1}^L \frac{n_{\ell-1}}{n_0} x^{L-\ell} + \sum_{\ell_1=1}^{L-1} \sum_{\ell_2=\ell_1+1}^L \frac{n_{\ell_2-1}}{n_0} y^{\ell_2-\ell_1} x^{L-\ell_2} \\ &\quad + \sum_{\ell_1=1}^{L-1} \frac{n_{\ell_1-1}}{n_0} \sum_{\ell_2=\ell_1+1}^L x^{L-\ell_2} \\ &= \frac{M^2}{4\alpha_0} \left[ x^{L-1} \left( \frac{9}{25} + O\left(\frac{1}{M}\right) \right) - y^L \left( 1 + O\left(\frac{1}{M}\right) \right) - \frac{4\lambda}{5} + \frac{16}{25} \right] \end{aligned}$$

□

**Theorem A.8** (Dispersion of the NTK at initialization). *Consider a fully-connected DNN of depth  $L$  defined in (3) initialized as in (6). The input dimension is given by  $n_0 = \alpha_0 M$ , the output dimension is 1, and the hidden layers have constant width  $M$ . The activation function in the hidden layers is ReLU, i.e.  $\phi(x) = x\mathbb{1}\{x > 0\}$ , and the output layer is linear. Assume also that the biases are initialized to zero, i.e.  $\sigma_b = 0$ , and the input data is normalized. Then the dispersion of the NTK at initialization is given by the following expressions:*

1. In the **chaotic phase** ( $a := \sigma_w^2/2 > 1$ ), the NTK dispersion grows exponentially with the depth-to-width ratio  $\lambda := L/M$  as

$$\frac{\mathbb{E}[\Theta^2(x, x)]}{\mathbb{E}^2[\Theta(x, x)]} \xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} \frac{1}{2\lambda} e^{5\lambda} \left( 1 - \frac{1}{4\lambda} (1 - e^{-4\lambda}) \right) \quad (40)$$

2. At the **EOC** ( $a = 1$ ), the NTK dispersion grows exponentially with the depth-to-width ratio  $\lambda$  as well, but with a slower rate given by

$$\begin{aligned} \frac{\mathbb{E}[\Theta^2(x, x)]}{\mathbb{E}^2[\Theta(x, x)]} &\xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} \frac{1}{(1 + \alpha_0)^2 \lambda} \left[ e^{5\lambda} \left( \frac{1}{2} + \frac{16\alpha_0^2 + 36\alpha_0 - 25}{200\lambda} \right) \right. \\ &\quad \left. + e^\lambda \frac{1 - 4\alpha_0}{8\lambda} + \frac{2\alpha_0(4 - \alpha_0)}{25\lambda} - \frac{2\alpha_0(1 + \alpha_0)}{5} \right] \end{aligned} \quad (41)$$

3. In the **ordered phase** ( $a < 1$ ), the NTK variance does not grow with  $\lambda$  and we have

$$\frac{\mathbb{E}[\Theta^2(x, x)]}{\mathbb{E}^2[\Theta(x, x)]} \xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} 1 \quad (42)$$

*Proof.* We will consider the cases of the ordered phase ( $a := \sigma_w^2/2 < 1$ ), the chaotic phase ( $a > 1$ ) and the EOC ( $a = 1$ ) separately.

**Case 1: Chaotic phase.** Using the results of Lemmas A.5, A.6, and A.7 and taking into account that  $a > 1$ , we obtain the following limit:

$$\frac{\mathbb{E}[\Theta_b(x, x)]}{\mathbb{E}[\Theta_W(x, x)]} = \frac{\frac{a^L - 1}{a - 1}}{a^{L-1} \left( 1 + \frac{M}{n_0} (L - 1) \right)} = \frac{\frac{a - a^{-L+1}}{a - 1}}{1 + \frac{M}{n_0} (L - 1)} \xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} 0$$

Therefore, recalling that  $\Theta(x, x) = \Theta_W(x, x) + \Theta_b(x, x)$ , we get the ratio between the complete NTK and its component corresponding to weights:

$$\frac{\mathbb{E}^2[\Theta(x, x)]}{\mathbb{E}^2[\Theta_W(x, x)]} = \left( 1 + \frac{\mathbb{E}[\Theta_b(x, x)]}{\mathbb{E}[\Theta_W(x, x)]} \right)^2 \xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} 1$$

Similarly, from Lemmas A.5, A.6, A.7, we can also obtain the following limit:

$$\frac{\mathbb{E}[\Theta^2(x, x)]}{\mathbb{E}^2[\Theta_W(x, x)]} = \frac{\mathbb{E}[\Theta_W^2(x, x)]}{\mathbb{E}^2[\Theta_W(x, x)]} + \frac{\mathbb{E}[\Theta_b^2(x, x)]}{\mathbb{E}^2[\Theta_W(x, x)]} + \frac{\mathbb{E}[\Theta_W(x, x)\Theta_b(x, x)]}{\mathbb{E}^2[\Theta_W(x, x)]} \xrightarrow[\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}]{\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}} \frac{\mathbb{E}[\Theta_W^2(x, x)]}{\mathbb{E}^2[\Theta_W(x, x)]}$$

Therefore, the dispersion of the NTK is determined by  $\Theta_W(x, x)$  in the infinite-depth-and-width limit in case of the initialization in the chaotic phase. Then we have the following expression for the dispersion in the limit:

$$\frac{\mathbb{E}[\Theta^2(x, x)]}{\mathbb{E}^2[\Theta(x, x)]} \xrightarrow[\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}]{\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}} \frac{1}{2\lambda} e^{5\lambda} \left( 1 - \frac{1}{4\lambda} (1 - e^{-4\lambda}) \right),$$

which completes the first part of the proof.

**Case 2: Ordered phase.** In the ordered phase, i.e. if  $a < 1$ , we have that  $a^L \rightarrow 0$  as  $L \rightarrow \infty$ , so Lemmas A.5, A.6 and A.7 suggest different relations between the terms of the NTK:

$$\begin{aligned} \frac{\mathbb{E}[\Theta_W(x, x)]}{\mathbb{E}[\Theta_b(x, x)]} &= \frac{a^{L-1} \left( 1 + \frac{M}{n_0} (L-1) \right)}{\frac{a^L - 1}{a - 1}} \xrightarrow[\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}]{\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}} 0 \\ \mathbb{E}[\Theta_W^2(x, x)] &\xrightarrow[\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}]{\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}} 0, \\ \mathbb{E}[\Theta_W(x, x)\Theta_b(x, x)] &\xrightarrow[\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}]{\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}} 0 \\ \mathbb{E}[\Theta_b^2(x, x)] &\xrightarrow[\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}]{\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}} \frac{1}{(a-1)^2} \end{aligned}$$

Therefore, the dispersion of the NTK is determined by the component corresponding to biases  $\Theta_b$  in the limit in case of initialization in the ordered phase:

$$\frac{\mathbb{E}[\Theta^2(x, x)]}{\mathbb{E}^2[\Theta(x, x)]} \xrightarrow[\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}]{\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}} 1,$$

which completes this part of the proof.

**Case 3: EOC.** Here we have  $a^\ell = 1$  for any  $\ell \in \mathbb{N}$ . Therefore, we can simplify the expressions for expectations from Lemmas A.5, A.6 and A.7 as follows:

$$\begin{aligned} \mathbb{E}[\Theta_W(x, x)] &= 1 + \frac{1}{\alpha_0} (L-1), \\ \mathbb{E}[\Theta_b(x, x)] &= L. \end{aligned}$$

Then the expectation of the complete NTK is given by:

$$\mathbb{E}[\Theta(x, x)] = \mathbb{E}[\Theta_W(x, x)] + \mathbb{E}[\Theta_b(x, x)] = \frac{L}{\alpha_0} \left( 1 + \alpha_0 + \frac{\alpha_0}{L} - \frac{1}{L} \right) \propto \frac{L}{\alpha_0} (1 + \alpha_0).$$

The squared NTK is given by  $\Theta^2(x, x) = \Theta_W^2(x, x) + 2\Theta_W(x, x)\Theta_b(x, x) + \Theta_b^2(x, x)$ . Then we need to consider the

expectations of all the components of this sum:

$$\begin{aligned} \frac{\mathbb{E}[\Theta_W^2(x, x)]}{L^2/\alpha_0^2} &\xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} e^{5\lambda} \left( \frac{1}{2\lambda} - \frac{1}{8\lambda^2} \right) + e^\lambda \frac{1}{8\lambda^2}, \\ \frac{\mathbb{E}[\Theta_b^2(x, x)]}{L^2/\alpha_0^2} &\xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} \alpha_0^2 \left( \frac{2}{25\lambda^2} e^{5\lambda} - \frac{2}{25\lambda^2} - \frac{2}{5\lambda} \right), \\ \frac{\mathbb{E}[\Theta_W(x, x)\Theta_b(x, x)]}{L^2/\alpha_0^2} &\xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} \frac{\alpha_0}{4} \left( \frac{9}{25\lambda^2} e^{5\lambda} - \frac{1}{\lambda^2} e^\lambda - \frac{4}{5\lambda} + \frac{16}{25\lambda^2} \right). \end{aligned}$$

Putting the above expressions together, we get the following limit for the desired ratio:

$$\begin{aligned} \frac{\mathbb{E}[\Theta^2(x, x)]}{\mathbb{E}^2[\Theta(x, x)]} &\xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} \frac{1}{(1 + \alpha_0)^2 \lambda} \left[ e^{5\lambda} \left( \frac{1}{2} + \frac{16\alpha_0^2 + 36\alpha_0 - 25}{200\lambda} \right) \right. \\ &\quad \left. + e^\lambda \frac{1 - 4\alpha_0}{8\lambda} + \frac{2\alpha_0(4 - \alpha_0)}{25\lambda} - \frac{2\alpha_0(1 + \alpha_0)}{5} \right], \end{aligned}$$

which completes the proof.  $\square$

## A.2. Non-diagonal elements of the NTK

**Theorem A.9** (Non-diagonal elements of the NTK). *Consider a fully-connected DNN of depth  $L$  defined in (3) initialized as in (6). The input dimension is given by  $n_0 = \alpha_0 M$ , the output dimension is 1, and the hidden layers have constant width  $M$ . The activation function in the hidden layers is ReLU, i.e.  $\phi(x) = x \mathbb{1}\{x > 0\}$ , and the output layer is linear. Assume also that the biases are initialized to zero, i.e.  $\sigma_b = 0$ , and the input data is normalized. Then for the ratio of non-diagonal and diagonal elements of the NTK we have:*

$$1 \geq \lim_{\substack{L \rightarrow \infty, M \rightarrow \infty \\ L/M \rightarrow \lambda \in \mathbb{R}}} \frac{\mathbb{E}[\Theta(x, \tilde{x})]}{\mathbb{E}[\Theta(x, x)]} \geq \frac{1}{4}$$

Moreover, for the dispersion of the non-diagonal elements we have:

$$\lim_{\substack{L \rightarrow \infty, M \rightarrow \infty \\ L/M \rightarrow \lambda \in \mathbb{R}}} \frac{\mathbb{E}[\Theta^2(x, \tilde{x})]}{\mathbb{E}^2[\Theta(x, \tilde{x})]} \leq 16 \lim_{\substack{L \rightarrow \infty, M \rightarrow \infty \\ L/M \rightarrow \lambda \in \mathbb{R}}} \frac{\mathbb{E}[\Theta^2(x, x)]}{\mathbb{E}^2[\Theta(x, x)]}$$

*Proof.* The non-diagonal element of the NTK on point  $x$  and  $\tilde{x}$  is given by

$$\Theta(x, \tilde{x}) = \sum_{\ell=1}^L \langle \delta^\ell, \tilde{\delta}^\ell \rangle \langle \mathbf{x}^{\ell-1}, \tilde{\mathbf{x}}^{\ell-1} \rangle + \sum_{\ell=1}^L \langle \delta^\ell, \tilde{\delta}^\ell \rangle,$$

where the activations and the backpropagated errors with tilde correspond to  $\tilde{x}$ . Same as in Lemma A.1, we can write the following for the involved dot products:

$$\langle \mathbf{x}^\ell, \tilde{\mathbf{x}}^\ell \rangle = \frac{\sigma_w^2}{n_{\ell-1}} \|\mathbf{x}^{\ell-1}\| \|\tilde{\mathbf{x}}^{\ell-1}\| \sum_{i=1}^{n_\ell} \phi(\mathcal{U}_i^\ell) \phi(\tilde{\mathcal{U}}_i^\ell)$$

We notice that in this case  $\mathcal{U}_i^\ell \sim \mathcal{N}(0, 1)$  and  $\tilde{\mathcal{U}}_i^\ell \sim \mathcal{N}(0, 1)$  are correlated variables and the covariance is given by  $\rho_x^{\ell-1} := \frac{\langle \mathbf{x}^{\ell-1}, \tilde{\mathbf{x}}^{\ell-1} \rangle}{\|\mathbf{x}^{\ell-1}\| \|\tilde{\mathbf{x}}^{\ell-1}\|}$ . The distribution of  $\mathcal{U}_i^\ell$  and  $\tilde{\mathcal{U}}_i^\ell$  depends only on the angle between the activations and not on the norms.

Assuming  $\rho_x^{\ell-1}$  is given, we can calculate the expectation of  $\phi(\mathcal{U}_i^\ell) \phi(\tilde{\mathcal{U}}_i^\ell)$ :

$$\mathbb{E}[\phi(\mathcal{U}_i^\ell) \phi(\tilde{\mathcal{U}}_i^\ell) \mid \rho_x^{\ell-1}] = \frac{1}{2\pi} \left( \sqrt{1 - (\rho_x^{\ell-1})^2} + \rho_x^{\ell-1} \pi/2 + \rho_x^{\ell-1} \arcsin \rho_x^{\ell-1} \right)$$

Then, denoting  $g(x) := \frac{1}{\pi} (\sqrt{1-x^2} + x\pi/2 + x \arcsin x)$ , we have

$$\mathbb{E} \left[ \frac{\langle \mathbf{x}^\ell, \tilde{\mathbf{x}}^\ell \rangle}{\langle \mathbf{x}^{\ell-1}, \tilde{\mathbf{x}}^{\ell-1} \rangle} \right] = \frac{\sigma_w^2}{2} \frac{n_\ell}{n_{\ell-1}} \mathbb{E} \left[ \frac{g(\rho_x^{\ell-1})}{\rho_x^{\ell-1}} \right]$$

We can reason in the same way to find expected dot products of the backpropagated errors:

$$\langle \boldsymbol{\delta}^\ell, \tilde{\boldsymbol{\delta}}^\ell \rangle = \frac{\sigma_w^2}{n_\ell} \|\boldsymbol{\delta}^{\ell+1}\| \|\tilde{\boldsymbol{\delta}}^{\ell+1}\| \sum_{i=1}^{n_\ell} \phi'(\mathcal{U}_i^\ell) \phi'(\tilde{\mathcal{U}}_i^\ell) \mathcal{V}_i^{\ell+1} \tilde{\mathcal{V}}_i^{\ell+1}$$

We can also calculate the involved expectations:

$$\begin{aligned} \mathbb{E}[\phi'(\mathcal{U}_i^\ell) \phi'(\tilde{\mathcal{U}}_i^\ell) \mid \rho_x^{\ell-1}] &= \frac{1}{2\pi} \left( \frac{\pi}{2} + \arcsin \rho_x^{\ell-1} \right), \\ \mathbb{E}[\mathcal{V}_i^\ell \tilde{\mathcal{V}}_i^\ell \mid \rho_\delta^\ell] &= \rho_\delta^\ell := \frac{\langle \boldsymbol{\delta}^\ell, \tilde{\boldsymbol{\delta}}^\ell \rangle}{\|\boldsymbol{\delta}^\ell\| \|\tilde{\boldsymbol{\delta}}^\ell\|} \end{aligned}$$

And, using the above expressions, we get

$$\mathbb{E} \left[ \frac{\langle \boldsymbol{\delta}^\ell, \tilde{\boldsymbol{\delta}}^\ell \rangle}{\langle \boldsymbol{\delta}^{\ell+1}, \tilde{\boldsymbol{\delta}}^{\ell+1} \rangle} \right] = \frac{\sigma_w^2}{2} \mathbb{E} \left[ \frac{1}{\pi} \left( \frac{\pi}{2} + \arcsin \rho_x^{\ell-1} \right) \right]$$

We also need to consider the expectation of  $\rho_x^\ell$ :

$$\mathbb{E}[\rho_x^\ell \mid \rho_x^{\ell-1}] = \mathbb{E} \left[ \frac{\sum_i \phi(\mathcal{U}_i^\ell) \phi(\tilde{\mathcal{U}}_i^\ell)}{\sqrt{\sum_i \phi^2(\mathcal{U}_i^\ell)} \sqrt{\sum_i \phi^2(\tilde{\mathcal{U}}_i^\ell)}} \mid \rho_x^{\ell-1} \right] \xrightarrow{n_{\ell-1} \rightarrow \infty} g(\rho_x^{\ell-1}),$$

where the correction to the above expectation for finite width is of order  $O(1/M)$  since the components approach normality with this rate. Moreover, the estimator of correlation coefficient has a negative bias, therefore  $\mathbb{E}[\rho_x^\ell \mid \rho_x^{\ell-1}]$  approaches  $g(\rho_x^{\ell-1})$  from below with  $n_{\ell-1} \rightarrow \infty$ . Then we have

$$\mathbb{E} \left[ \langle \mathbf{x}^\ell, \tilde{\mathbf{x}}^\ell \rangle \right] = \langle \mathbf{x}^0, \tilde{\mathbf{x}}^0 \rangle \mathbb{E} \left[ \prod_{k=1}^{\ell} \frac{\langle \mathbf{x}^k, \tilde{\mathbf{x}}^k \rangle}{\langle \mathbf{x}^{k-1}, \tilde{\mathbf{x}}^{k-1} \rangle} \right] = \langle \mathbf{x}^0, \tilde{\mathbf{x}}^0 \rangle a^\ell \frac{n_\ell}{n_0} \mathbb{E} \left[ \frac{g(\rho_x^{\ell-1})}{\rho_x^0} \prod_{k=0}^{\ell-2} \frac{g(\rho_x^k)}{\rho_x^{k+1}} \right] \geq \|\mathbf{x}^0\| \|\tilde{\mathbf{x}}^0\| a^\ell \frac{n_\ell}{n_0} \mathbb{E}[\rho_x^\ell]$$

Similarly, denoting  $f(x) := \frac{1}{\pi} (\pi/2 + \arcsin x)$ , we get the following for the products of backpropagated errors:

$$\mathbb{E}[\langle \boldsymbol{\delta}^\ell, \tilde{\boldsymbol{\delta}}^\ell \rangle] \geq a^{L-\ell} \prod_{k=\ell}^{L-1} f(\mathbb{E}[\rho_x^{k-1}])$$

Now we notice that  $\mathbb{E}[\rho_x^\ell] \rightarrow g^{\circ \ell}(\rho_x^0)$  not only if  $M \rightarrow \infty$  but also if  $\ell \rightarrow \infty$  with finite  $M$ , where  $g^{\circ k}$  denotes composition of the function  $k$  times. Indeed  $g(x)$  is a monotonically increasing function with  $g(x) \geq x$  and a single fixed point at  $x = 1$ , so we have  $\mathbb{E}[\rho_x^\ell] \rightarrow 1$  and  $g^{\circ \ell}(\rho_x^0) \rightarrow 1$  if  $\ell \rightarrow \infty$ . In other words, if  $\mathbb{E}[\rho_x^\ell]/g^{\circ \ell}(\rho_x^0) = 1 + c_\ell/M$  for some coefficients  $c_\ell$ , then  $c_\ell \rightarrow 0$  as  $\ell \rightarrow \infty$ . Therefore, we can replace  $\mathbb{E}[\rho_x^\ell]$  with  $g^{\circ \ell}(\rho_x^0)$  in the above bounds to obtain the infinite-depth-and-width limit.

Putting everything together, we can write the following bound for the expectation of a non-diagonal element of the NTK

$$\begin{aligned} \lim_{\substack{L \rightarrow \infty, M \rightarrow \infty \\ L/M \rightarrow \lambda \in \mathbb{R}}} [\Theta(x, \tilde{x})] &\geq \lim_{\substack{L \rightarrow \infty, M \rightarrow \infty \\ L/M \rightarrow \lambda \in \mathbb{R}}} \left[ \|\mathbf{x}^0\| \|\tilde{\mathbf{x}}^0\| a^{L-1} \sum_{\ell=1}^L \frac{n_{\ell-1}}{n_0} g^{\circ \ell-1}(\rho_x^0) \prod_{k=\ell}^{L-1} f(g^{\circ(k-1)}(\rho_x^0)) + \right. \\ &\quad \left. \sum_{\ell=1}^L a^{L-\ell} \prod_{k=\ell}^{L-1} f(g^{\circ(k-1)}(\rho_x^0)) \right], \end{aligned}$$

Now studying the expressions above we can find the following bounds:

$$1 \geq \lim_{\substack{L \rightarrow \infty, M \rightarrow \infty \\ L/M \rightarrow \lambda \in \mathbb{R}}} \frac{\mathbb{E}[\Theta(x, \tilde{x})]}{\mathbb{E}[\Theta(x, x)]} \geq \frac{1}{4},$$

The upper bound is trivial. We obtain the lower bound in case of initialization in the chaotic phase by noticing that  $\sum_{\ell=1}^L g^{\circ \ell-1}(\rho_x^0) \prod_{k=\ell}^{L-1} f(g^{\circ(k-1)}(\rho_x^0)) \geq L/4$  for  $L \geq 2$ , which, by Chebyshev's sum inequality, gives the maximal ratio between diagonal and non-diagonal elements of the NTK, since  $\mathbb{E}[\Theta_W(x, x)] = a^{L-1} \sum_{\ell=1}^L \frac{n_{\ell-1}}{n_0}$ . In the ordered phase, we have  $\sum_{\ell=1}^L \prod_{k=\ell}^{L-1} f(g^{\circ(k-1)}(\rho_x^0)) \geq L/4$  and  $\mathbb{E}[\Theta_b(x, x)] = \sum_{\ell=1}^L a^{L-\ell}$ , which gives the same bound.

Moreover, it is easy to see that  $\mathbb{E}[\Theta^2(x, \tilde{x})] \leq \mathbb{E}[\Theta^2(x, x)]$ , therefore we can write

$$\lim_{\substack{L \rightarrow \infty, M \rightarrow \infty \\ L/M \rightarrow \lambda \in \mathbb{R}}} \frac{\mathbb{E}[\Theta^2(x, \tilde{x})]}{\mathbb{E}^2[\Theta(x, \tilde{x})]} \leq 16 \lim_{\substack{L \rightarrow \infty, M \rightarrow \infty \\ L/M \rightarrow \lambda \in \mathbb{R}}} \frac{\mathbb{E}[\Theta^2(x, x)]}{\mathbb{E}^2[\Theta(x, x)]}$$

□

### A.3. Training dynamics of the NTK

**Theorem A.10** (GD step of the NTK). *Consider a fully-connected DNN of depth  $L$  defined in (3) initialized as in (6). The input dimension is given by  $n_0 = \alpha_0 M$ , the output dimension is 1, and the hidden layers have constant width  $M$ . The activation function in the hidden layers is ReLU, i.e.  $\phi(x) = x \mathbb{1}\{x > 0\}$ , and the output layer is linear. Assume also that the biases are initialized to zero, i.e.  $\sigma_b = 0$ , and the input data is normalized. Then, if we perform a GD step on a point  $(x, y) \in \mathcal{D}$  with learning rate  $\eta$ , the following holds for the changes of the corresponding element of the NTK:*

1. In the **chaotic phase** ( $a := \sigma_w^2/2 > 1$ ), the changes to the NTK value are infinite in the limit for a constant learning rate:

$$\frac{\mathbb{E}[\Delta\Theta(x, x)]}{\mathbb{E}[\Theta(x, x)]} \xrightarrow[\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}]{\eta} \infty \quad (43)$$

The scaling of the learning rate needed to avoid the infinite limit is given by  $\eta = O(a^{-L})$ , which tends to zero with depth.

2. In the **ordered phase** ( $a < 1$ ), the NTK stays constant in the limit:

$$\frac{\mathbb{E}[\Delta\Theta(x, x)]}{\mathbb{E}[\Theta(x, x)]} \xrightarrow[\substack{M \rightarrow \infty, L \rightarrow \infty, \\ L/M \rightarrow \lambda \in \mathbb{R}}]{\eta} 0 \quad (44)$$

*Proof.* A derivative of the NTK in gradient flow can be expanded as follows:

$$\dot{\Theta}(x, x) = \sum_{\ell=1}^L \left( \sum_{i,j} \frac{\partial \Theta(x, x)}{\partial \mathbf{W}_{ij}^\ell} \dot{\mathbf{W}}_{ij}^\ell + \sum_i \frac{\partial \Theta(x, x)}{\partial \mathbf{b}_i^\ell} \dot{\mathbf{b}}_i^\ell \right),$$

where the parameters change in the direction of the negative gradient:

$$\dot{\mathbf{W}}_{ij}^\ell = -\frac{\partial \mathcal{L}(\mathcal{D})}{\partial \mathbf{W}_{ij}^\ell}, \quad \dot{\mathbf{b}}_i^\ell = -\frac{\partial \mathcal{L}(\mathcal{D})}{\partial \mathbf{b}_i^\ell}, \quad i = 1, \dots, n_\ell, \quad j = 1, \dots, n_{\ell-1}, \quad \ell = 1, \dots, L$$

If we now assume that the gradient descent step is performed on a single point of the dataset  $x$ , which is the same point for which the NTK is calculated, we have:

$$\begin{aligned} \dot{\mathbf{W}}_{ij}^\ell &= -\frac{\partial \mathcal{L}(x)}{\partial \mathbf{W}_{ij}^\ell} = -\frac{\partial \mathcal{L}(x)}{\partial f(x)} \frac{\partial f(x)}{\partial \mathbf{W}_{ij}^\ell} = -\frac{\partial \mathcal{L}(x)}{\partial f(x)} \delta_i^\ell \mathbf{x}_j^{\ell-1}, \\ \dot{\mathbf{b}}_i^\ell &= -\frac{\partial \mathcal{L}(x)}{\partial \mathbf{b}_i^\ell} = -\frac{\partial \mathcal{L}(x)}{\partial f(x)} \frac{\partial f(x)}{\partial \mathbf{b}_i^\ell} = -\frac{\partial \mathcal{L}(x)}{\partial f(x)} \delta_i^\ell \end{aligned}$$

It remains to calculate the derivatives of the NTK with respect to the parameters. The involved terms are:

$$\begin{aligned}\frac{\partial \Theta_W(x, x)}{\partial \mathbf{W}_{ij}^\ell} &= \sum_{\ell'} \sum_{i', j'} \frac{\partial}{\partial \mathbf{W}_{ij}^\ell} \left( \frac{\partial f(x)}{\partial \mathbf{W}_{i'j'}^{\ell'}} \right)^2 = 2 \sum_{\ell'} \sum_{i', j'} \frac{\partial f(x)}{\partial \mathbf{W}_{i'j'}^{\ell'}} \frac{\partial^2 f(x)}{\partial \mathbf{W}_{ij}^\ell \partial \mathbf{W}_{i'j'}^{\ell'}}, \\ \frac{\partial \Theta_W(x, x)}{\partial \mathbf{b}_k^\ell} &= \sum_{\ell'} \sum_{i', j'} \frac{\partial}{\partial \mathbf{b}_k^\ell} \left( \frac{\partial f(x)}{\partial \mathbf{W}_{i'j'}^{\ell'}} \right)^2 = 2 \sum_{\ell'} \sum_{i', j'} \frac{\partial f(x)}{\partial \mathbf{W}_{i'j'}^{\ell'}} \frac{\partial^2 f(x)}{\partial \mathbf{b}_k^\ell \partial \mathbf{W}_{i'j'}^{\ell'}}, \\ \frac{\partial \Theta_b(x, x)}{\partial \mathbf{b}_k^\ell} &= \sum_{\ell'} \sum_{i'} \frac{\partial}{\partial \mathbf{b}_k^\ell} \left( \frac{\partial f(x)}{\partial \mathbf{b}_{i'}^{\ell'}} \right)^2 = 2 \sum_{\ell'} \sum_{i'} \frac{\partial f(x)}{\partial \mathbf{b}_{i'}^{\ell'}} \frac{\partial^2 f(x)}{\partial \mathbf{b}_k^\ell \partial \mathbf{b}_{i'}^{\ell'}}, \\ \frac{\partial \Theta_b(x, x)}{\partial \mathbf{W}_{ij}^\ell} &= \sum_{\ell'} \sum_{i', j'} \frac{\partial}{\partial \mathbf{W}_{ij}^\ell} \left( \frac{\partial f(x)}{\partial \mathbf{b}_{i'}^{\ell'}} \right)^2 = 2 \sum_{\ell'} \sum_{i'} \frac{\partial f(x)}{\partial \mathbf{b}_{i'}^{\ell'}} \frac{\partial^2 f(x)}{\partial \mathbf{W}_{ij}^\ell \partial \mathbf{b}_{i'}^{\ell'}}\end{aligned}$$

To calculate these terms, we need to find the second derivatives of the DNN's output function.

$$\frac{\partial^2 f(x)}{\partial \mathbf{W}_{ij}^\ell \partial \mathbf{W}_{i'j'}^{\ell'}} = \delta_{i'}^{\ell'} \frac{\partial \mathbf{x}_{j'}^{\ell'-1}}{\partial \mathbf{W}_{ij}^\ell} + \mathbf{x}_{j'}^{\ell'-1} \frac{\partial \delta_{i'}^{\ell'}}{\partial \mathbf{W}_{ij}^\ell} = \mathbb{1}_{\ell < \ell'} \delta_{i'}^{\ell'} \mathbf{x}_j^{\ell-1} \frac{\partial \mathbf{x}_{j'}^{\ell'-1}}{\partial \mathbf{h}_i^\ell} + \mathbb{1}_{\ell > \ell'} \delta_i^\ell \mathbf{x}_{j'}^{\ell'-1} \phi'(\mathbf{h}_j^{\ell-1}) \frac{\partial \delta_{i'}^{\ell'}}{\partial \delta_j^{\ell-1}},$$

In the above equation the first term is non-zero only in case  $\ell' > \ell$  and the second term is non-zero only if  $\ell' < \ell$ . Then we can write the following:

$$\begin{aligned}\sum_{\ell} \sum_{i, j} \frac{\partial \Theta_W(x, x)}{\partial \mathbf{W}_{ij}^\ell} \dot{\mathbf{W}}_{ij}^\ell &= - \frac{\partial \mathcal{L}(x)}{\partial f(x)} \sum_{\ell' > \ell} \|\delta^{\ell'}\|^2 \|\mathbf{x}^{\ell-1}\|^2 \sum_i \delta_i^{\ell'} \frac{\partial \|\mathbf{x}^{\ell'-1}\|^2}{\partial \mathbf{h}_i^\ell} \\ &\quad - \frac{\partial \mathcal{L}(x)}{\partial f(x)} \sum_{\ell > \ell'} \|\delta^\ell\|^2 \|\mathbf{x}^{\ell'-1}\|^2 \sum_j \mathbf{x}_j^{\ell-1} \frac{\partial \|\delta^{\ell'}\|^2}{\partial \delta_j^{\ell-1}}\end{aligned}$$

Opening the remaining parts of the derivative in the same way, we obtain the following expression:

$$\begin{aligned}\dot{\Theta}(x, x) &= - \frac{\partial \mathcal{L}(x)}{\partial f(x)} \left( \sum_{\ell' > \ell} (\|\delta^{\ell'}\|^2 \|\mathbf{x}^{\ell-1}\|^2 + \|\delta^{\ell'}\|^2) \sum_i \delta_i^{\ell'} \frac{\partial \|\mathbf{x}^{\ell'-1}\|^2}{\partial \mathbf{h}_i^\ell} \right. \\ &\quad \left. + \sum_{\ell > \ell'} (\|\delta^\ell\|^2 \|\mathbf{x}^{\ell'-1}\|^2 + \|\delta^\ell\|^2) \sum_j \mathbf{x}_j^{\ell-1} \frac{\partial \|\delta^{\ell'}\|^2}{\partial \delta_j^{\ell-1}} \right)\end{aligned}$$

**Case 1. Chaotic phase.** Let us bound the change of the NTK by computing only the terms with  $\ell' = \ell + 1$ . In this case,  $\frac{\partial \|\mathbf{x}^{\ell'-1}\|^2}{\partial \mathbf{h}_i^\ell} = 2\mathbf{x}_i^\ell$ . We then notice that  $\sum_i \mathbf{x}_i^\ell \delta_i^\ell = \sum_k \delta_k^{\ell+1} \sum_i \mathbf{W}_{ki}^{\ell+1} \mathbf{x}_i^\ell = \sum_k \delta_k^{\ell+1} \mathbf{h}_k^{\ell+1} = \sum_k \delta_k^{\ell+1} \mathbf{x}_k^{\ell+1}$ , which by induction gives  $\sum_i \mathbf{x}_i^\ell \delta_i^\ell = f(x)$ . Therefore, taking into account that for quadratic loss we have  $\partial \mathcal{L}(x)/\partial f(x) = f(x) - y$ , we can write the following bound:

$$\mathbb{E}[|\Delta \Theta(x, x)|] \geq 2\eta \mathbb{E} \left[ f(x)^2 \sum_{\ell=1}^L (\|\delta^{\ell+1}\|^2 \|\mathbf{x}^{\ell-1}\|^2 + \|\delta^{\ell+1}\|^2) \right]$$

Then, using the results of Lemmas A.1, A.2 and A.4 again, we obtain the expectation of the first part:

$$\mathbb{E}[|\Delta \Theta(x, x)|] \geq 4\eta \frac{\|\mathbf{x}^0\|^4}{n_0} a^{2L-1} \sum_{\ell=1}^L \prod_{j=\ell+1}^{L-1} \left( 1 + \frac{1}{n_j} + O\left(\frac{1}{M^{3/2}}\right) \right) \prod_{i=1}^{\ell-1} \left( 1 + \frac{5}{n_j} \right) \propto a^{2L-1}$$

And since in case of the chaotic phase  $\mathbb{E}[\Theta(x, x)] \propto \mathbb{E}[\Theta_W(x, x)] \propto a^L LM/n_0$ , we have the desired limit:

$$\frac{\mathbb{E}[|\Delta\Theta(x, x)|]}{\mathbb{E}[\Theta(x, x)]} \xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} \infty$$

**Case 2. Ordered phase.** The bound that we used for the chaotic phase above gives zero in the limit in case of the ordered phase. We will now show that the upper bound of the relative change of the NTK is also zero in the limit in this case. We notice that  $\sum_i f(x) \delta_i^\ell \frac{\partial \|\mathbf{x}^{\ell'-1}\|^2}{\partial \mathbf{h}_i^\ell} = \sum_i \delta_i^\ell \sum_k \delta_k^\ell \frac{\partial \|\mathbf{x}^{\ell'-1}\|^2}{\partial \mathbf{h}_i^\ell} \mathbf{h}_k^\ell = \sum_i (\delta_i^\ell)^2 \frac{\partial \|\mathbf{x}^{\ell'-1}\|^2}{\partial \mathbf{h}_i^\ell} \mathbf{h}_i^\ell + \sum_{i \neq k} \delta_i^\ell \delta_k^\ell \frac{\partial \|\mathbf{x}^{\ell'-1}\|^2}{\partial \mathbf{h}_i^\ell} \mathbf{h}_k^\ell$ . Then we have  $\sum_i f(x) \delta_i^\ell \frac{\partial \|\mathbf{x}^{\ell'-1}\|^2}{\partial \mathbf{h}_i^\ell} \leq \|\delta^\ell\|^2 \|\mathbf{x}^{\ell'-1}\|^2 + A$ , where the expectation of  $A$  is zero. Similarly, we have  $\sum_j f(x) \mathbf{x}_j^{\ell-1} \frac{\partial \|\delta^{\ell'}\|^2}{\partial \delta_j^{\ell-1}} \leq \|\delta^{\ell'}\|^2 \|\mathbf{x}^{\ell-1}\|^2 + B$  with a term  $B$  of zero expectation. Then we have the following bound for the change of the NTK:

$$\mathbb{E}[|\Delta\Theta(x, x)|] \leq 2\eta \mathbb{E} \left[ \sum_{\ell_1} \|\delta^{\ell_1}\|^2 \|\mathbf{x}^{\ell_1-1}\|^2 \sum_{\ell_2 < \ell_1} \|\delta^{\ell_2}\|^2 (\|\mathbf{x}^{\ell_2}\|^2 + 1) \right]$$

The expectation of  $\sum_{\ell_2 < \ell_1} \theta_W^{\ell_1} \theta_W^{\ell_2}$ , where  $\theta_W^\ell := \|\delta^\ell\|^2 \|\mathbf{x}^{\ell-1}\|^2$ , was calculated in Lemma A.5 and the expectation of  $\sum_{\ell_2 < \ell_1} \theta_W^{\ell_1} \theta_b^{\ell_2}$ , where  $\theta_b^\ell := \|\delta^\ell\|^2$ , was calculated in Lemma A.7. In particular, we have the following results for the two sums:

$$\begin{aligned} \mathbb{E} \left[ \sum_{\ell_2 < \ell_1} \theta_W^{\ell_1} \theta_W^{\ell_2} \right] &\propto a^{2L} \frac{L^2}{\alpha_0^2} \frac{1}{4\lambda} e^{5\lambda} \left( 1 - \frac{1}{4\lambda} (1 - e^{-4\lambda}) \right), \\ \mathbb{E} \left[ \sum_{\ell_2 < \ell_1} \theta_W^{\ell_1} \theta_b^{\ell_2} \right] &\propto \frac{a^{2L}}{a-1} \frac{L}{\alpha_0} \frac{1}{4\lambda} e^{5\lambda} (1 - e^{-4\lambda}). \end{aligned}$$

Then we see that the upper bound on the changes of the NTK is proportional to  $a^{2L} L^2$ , which tends to zero with depth in the ordered phase. Given that the expectation of the NTK in the ordered phase has a non-zero limit given by  $1/(1-a)$ , we can then conclude that

$$\frac{\mathbb{E}[|\Delta\Theta(x, x)|]}{\mathbb{E}[\Theta(x, x)]} \xrightarrow[M \rightarrow \infty, L \rightarrow \infty, L/M \rightarrow \lambda \in \mathbb{R}]{} 0$$

in case of initialization in the ordered phase. □

## B. Additional observations

### B.1. Effects of $\alpha_0 := n_0/M$ at the EOC

The theoretical expression for the NTK dispersion in the infinite-width limit, which we derived in Theorem 3.1, depends on the ratio  $\alpha_0 := n_0/M$  at the EOC:

$$V_{\text{EOC}} := \frac{\mathbb{E}[\Theta^2(x, x)]}{\mathbb{E}^2[\Theta(x, x)]} \rightarrow \frac{1}{(1 + \alpha_0)^2} \left[ e^{5\lambda} \left( \frac{1}{2\lambda} + \frac{2\alpha_0^2 - 8\alpha_0}{25\lambda^2} \right) + (e^\lambda - e^{5\lambda}) \frac{1 - 4\alpha_0}{8\lambda^2} + \frac{2\alpha_0}{5\lambda} \left( \frac{4 - \alpha_0}{5\lambda} - 1 - \alpha_0 \right) \right].$$

Examining this expression, one can see that it tends to the limiting expression for the NTK dispersion in the chaotic phase as the ratio  $\alpha_0$  decreases:

$$V_{\text{EOC}} \xrightarrow{\alpha_0 \rightarrow 0} \frac{1}{2\lambda} e^{5\lambda} \left( 1 - \frac{1}{4\lambda} (1 - e^{-4\lambda}) \right).$$

We illustrate this effect in Figure 5. One can see that gradually decreasing the value of  $\alpha_0$  moves the NTK dispersion at the EOC closer to the NTK dispersion in the chaotic phase.

### B.2. Effects of the architecture

In Section 3.2, we showed that constant-width DNNs that increase the input dimensionality, i.e.  $n_0 < n_1 = \dots = n_{L-1}$ , get more robust with depth in a sense that the dispersion of their NTK decreases. Here we show how the theoretical expressions






Figure 5. Effects of  $\alpha_0 := n_0/M$  on the NTK dispersion at the EOC in the infinite-depth-and-width limit. All the lines show the theoretical expressions from Theorem 3.1. The black line (uppermost) corresponds to the NTK dispersion in the chaotic phase, while all the other lines show the NTK dispersion at the EOC with varying  $\alpha_0$  values. The colors spanning from yellow to violet (from lighter to darker tones) indicate the value of  $\alpha_0$  spanning from 1 (yellow) to 0.1 (violet).

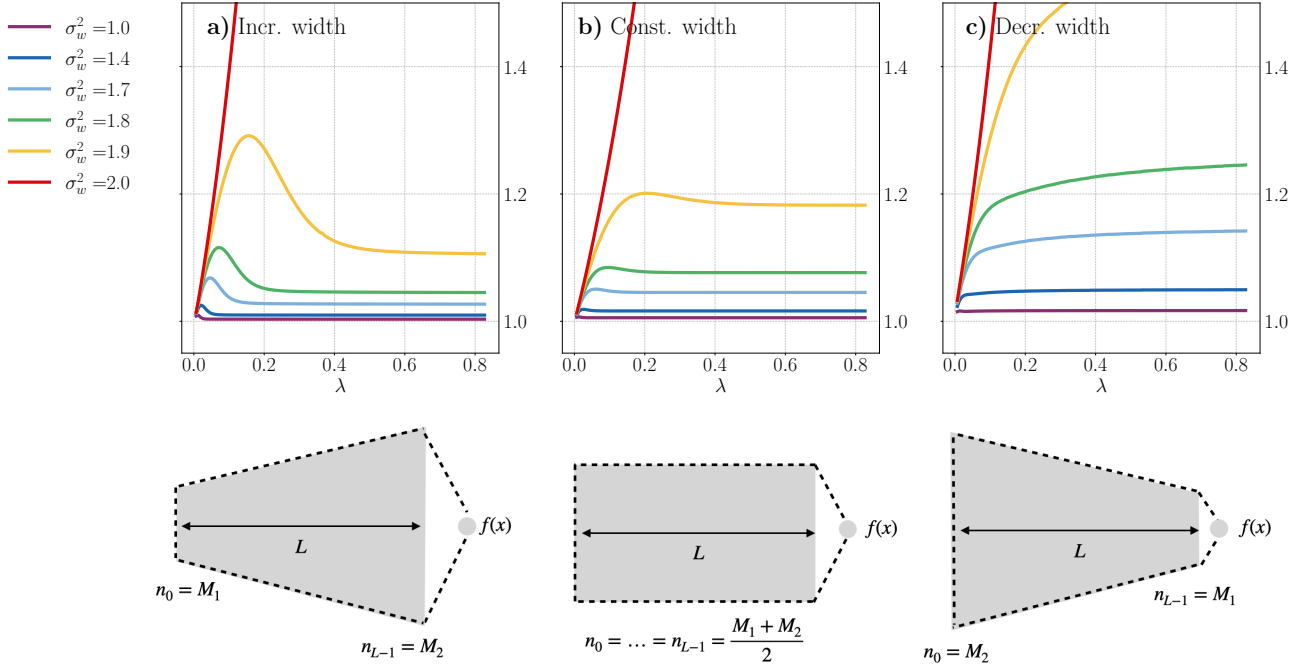


Figure 6. Effects of the architecture on the NTK dispersion ratio  $\mathbb{E}[\Theta^2(x, x)]/\mathbb{E}^2[\Theta(x, x)]$  as predicted by Theorem 3.2. The subplots show the dispersion for varying values of  $\sigma_w^2$  for three different architectures. The lower row of the figure illustrates the considered architectures. Formally, the widths for each architecture are given by: **a)**  $n_\ell = M_1 + \lceil \ell(M_2 - M_1)/L \rceil$ , **b)**  $n_\ell = \lceil (M_1 + M_2)/2 \rceil$ , **c)**  $n_\ell = M_2 + \lceil \ell(M_1 - M_2)/L \rceil$  for  $0 \leq \ell \leq L$ . The width parameters are given by  $M_1 = 100$ ,  $M_2 = 500$  and do not change as the depth grows. The depth-to-width ratio is computed for the average width, i.e.  $\lambda = 2L/(M_1 + M_2)$ .

for the NTK moments in Theorem 3.2 reveal more effects of the DNN's architecture. In Figure 6, we compute the theoretical prediction of the NTK dispersion for three architectures: the first one gradually increases width over  $L$  layers from  $n_0 = M_1$  to  $n_{L-1} = M_2$ , the second one keeps constant width, i.e.  $n_0 = \dots = n_{L-1} = (M_1 + M_2)/2$ , and the third one gradually decreases width over  $L$  layers from  $n_0 = M_2$  to  $n_{L-1} = M_1$ . We note that all the architectures have the same average width in this setting. We also note that we keep  $M_1$  and  $M_2$  fixed while varying the depth  $L$ . Therefore, we compare networks that increase or decrease the dimensionality equally but over a different number of layers. Figure 6 demonstrates that the NTK dispersion is lower for the DNNs that increase the dimensionality. Moreover, for such DNNs the peak of dispersion falls on the relatively shallow networks. Therefore, it may be beneficial to increase dimensionality over more layers if the goal is to decrease the variance of the DNN. On the contrary, the dispersion only increases with depth for DNNs that decrease the dimensionality. Thus, it may be beneficial to keep such networks shallow if one wants to keep the variance minimal.

### B.3. Lazy training

The NTK regime of neural networks is often discussed in connection with the so-called lazy training phenomenon (Chizat et al., 2019). In lazy training, a model behaves as its linearization around the initial parameters due to rescaling given by:

$$\tilde{f}_\alpha(x) = \alpha f(x), \quad \tilde{\mathcal{L}}_\alpha(\mathcal{D}) = \frac{1}{\alpha^2} \mathcal{L}(\mathcal{D}), \quad \alpha \in \mathbb{R},$$

where  $f(\cdot)$  is the original model's output function and  $\mathcal{L}(\mathcal{D})$  is the training loss. Chizat et al. (2019) showed that the dynamics of a rescaled model defined by  $\tilde{f}_\alpha(\cdot)$  and  $\tilde{\mathcal{L}}_\alpha(\mathcal{D})$  is close to its linearization if the scaling factor  $\alpha$  is large. Thus, in this section we discuss the effects of the lazy training rescaling on the results presented in our paper.

One can see that the NTK changes trivially if we rescale the output function:

$$\nabla_{\mathbf{w}} \tilde{f}_\alpha(x) = \alpha \nabla_{\mathbf{w}} f(x) \Rightarrow \tilde{\Theta}_\alpha(x_1, x_2) = \alpha^2 \Theta(x_1, x_2),$$

where we denoted the NTK of the rescaled model as  $\tilde{\Theta}_\alpha$ . Therefore, all our results concerning the NTK dispersion (Theorem 3.1) and the ratios of expectations at initialization (Theorem 3.3) do not change if we rescale the model, since the constants added to the nominator and the denominator cancel each other:

$$\frac{\mathbb{E}[\tilde{\Theta}_\alpha^2(x_1, x_2)]}{\mathbb{E}^2[\tilde{\Theta}_\alpha(x_1, x_2)]} = \frac{\mathbb{E}[\Theta^2(x_1, x_2)]}{\mathbb{E}^2[\Theta(x_1, x_2)]}.$$

On the other hand, the relative change of the NTK in a gradient descent step (Theorem 4.1) is affected by the rescaling. Recall that the NTK derivative is given by:

$$\dot{\Theta}(x, x) = \sum_{\ell=1}^L \left( \sum_{i,j} \frac{\partial \Theta(x, x)}{\partial \mathbf{W}_{ij}^\ell} \dot{\mathbf{W}}_{ij}^\ell + \sum_i \frac{\partial \Theta(x, x)}{\partial \mathbf{b}_i^\ell} \dot{\mathbf{b}}_i^\ell \right).$$

Terms of the above expression change as follows due to the rescaling:

$$\begin{aligned} \frac{\partial \tilde{\Theta}_\alpha(x, x)}{\partial \mathbf{W}_{ij}^\ell} &= \alpha^2 \frac{\partial \Theta(x, x)}{\partial \mathbf{W}_{ij}^\ell}, & \frac{\partial \tilde{\Theta}_\alpha(x, x)}{\partial \mathbf{b}_i^\ell} &= \alpha^2 \frac{\partial \Theta(x, x)}{\partial \mathbf{b}_i^\ell}, \\ (\tilde{\dot{\mathbf{W}}}_{ij}^\ell)_\alpha &= -\frac{\partial \tilde{\mathcal{L}}_\alpha(\mathcal{D})}{\partial \mathbf{W}_{ij}^\ell} = -\frac{1}{\alpha^2} \frac{\partial \mathcal{L}(\mathcal{D})}{\partial \mathbf{W}_{ij}^\ell} = -\frac{1}{\alpha^2} \dot{\mathbf{W}}_{ij}^\ell, & (\tilde{\dot{\mathbf{b}}}_i^\ell)_\alpha &= \frac{1}{\alpha^2} \dot{\mathbf{b}}_i^\ell. \end{aligned}$$

Therefore, the NTK derivative is not changed by the rescaling and for the relative change of the NTK we have:

$$\frac{\mathbb{E}[\Delta \tilde{\Theta}_\alpha(x, x)]}{\mathbb{E}[\tilde{\Theta}_\alpha(x, x)]} = \frac{1}{\alpha^2} \frac{\mathbb{E}[\Delta \Theta(x, x)]}{\mathbb{E}[\Theta(x, x)]}.$$

For any constant choice of  $\alpha \in \mathbb{R}$ , this scaling does not change our limiting results in Theorem 4.1. However, if the rescaling parameter is scaled exponentially with  $L$  as  $\alpha \propto (\sigma_w/\sqrt{2})^L$  then the relative change of the NTK tends to zero in the chaotic phase, as well as in the ordered phase and at the EOC. Thus, it is possible to enforce lazy training in deep networks outside of the ordered phase but the required rescaling parameter grows exponentially with depth.

## C. Additional experiments

### C.1. Non-diagonal elements of the NTK

We provide empirical results on the dispersion of the non-diagonal elements of the NTK in this subsection. Figure 7 is analogous to Figure 1: it compares the dispersion of the NTK in different phases of initialization. One can see that the non-diagonal elements of the NTK behave similarly to the diagonal ones. In particular, the dispersion of the non-diagonal elements grows exponentially with the depth-to-width ratio and reaches high values for deep networks in the chaotic phase and at the EOC, whereas in the ordered phase the dispersion is low and does not grow with depth. Figure 8 is analogous to Figure 2: it characterizes the behavior of the NTK dispersion close to the EOC, where the finite-width effects are significant. Here the picture is again similar to the one described in Section 3.2 for the diagonal elements of the NTK. In particular, the dispersion gradually increases as  $\sigma_w^2$  grows and approaches the EOC. One can also see that the dispersion of the non-diagonal elements decreases with depth in the ordered phase for networks with  $\alpha_0 < 1$ , same as in the case of the diagonal elements. In all the figures, we provide experiments for varying initial angles between the two input samples of the NTK and conclude that the dispersion does not depend significantly on this angle.

### C.2. Additional error bars for Figures 1, 2 and 3

To keep all the figures readable, we include error bars only in a subset of points in Figures 1 and 2 in the main text. We also omit error bars in Figure 3. To give the reader a better idea about the variance observed in our experiments, we include additional figures with continuous error bars in this section. Figure 9 is analogous to Figure 1: it shows the estimated dispersion of the NTK along with the theoretical expressions in the infinite-depth-and-width limit from Theorem 3.1. We include fewer lines (values of  $\sigma_w^2$ ) in this figure to keep the continuous error bars distinguishable. Similarly, Figure 10 is analogous to Figure 2: it shows the results concerning the NTK dispersion around the EOC. Finally, Figure 11 shows a subset of lines from Figure 3 with their continuous error bars and concerns the ratio between non-diagonal and diagonal elements of the NTK.

### C.3. Estimating the NTK dispersion from a sample

In our experiments, we estimate the ratio  $r := \mathbb{E}[\Theta^2(x, x)]/\mathbb{E}^2[\Theta(x, x)]$  at initialization from a sample. To do so, we sample an element of the NTK  $N$  times with independently chosen initialization parameters and get a sample  $\{\theta_i\}_{i=1}^N$ .

In this setting, the standard estimators for the first and the second moments given by  $\hat{\mu}_1 := \sum_{i=1}^N \theta_i/N$  and  $\hat{\mu}_2 := \sum_{i=1}^N \theta_i^2/N$  are unbiased:

$$\begin{aligned}\mathbb{E}[\hat{\mu}_1] &= \mathbb{E}\left[\frac{1}{N} \sum_{i=1}^N \theta_i\right] = \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\theta_i] = \mu_1, \\ \mathbb{E}[\hat{\mu}_2] &= \mathbb{E}\left[\frac{1}{N} \sum_{i=1}^N \theta_i^2\right] = \frac{1}{N} \sum_{i=1}^N \mathbb{E}[\theta_i^2] = \mu_2,\end{aligned}$$

where we denoted the actual moments as  $\mu_1 := \mathbb{E}[\theta_i]$  and  $\mu_2 := \mathbb{E}[\theta_i^2]$ .

However,  $\hat{\mu}_1$  and  $\hat{\mu}_2$  computed on the same sample are dependent, so the estimator for the desired ratio given by  $\hat{\mu}_2/(\hat{\mu}_1)^2$  is biased and we need to correct it. First, we notice that the estimator for  $\mu_1^2$  given by the square of  $\hat{\mu}_1$  is biased as follows:

$$\mathbb{E}[\hat{\mu}_1^2] = \mathbb{E}\left[\frac{1}{N^2} \left(\sum_{i=1}^N \theta_i\right)^2\right] = \frac{\mu_2}{N} + \frac{N-1}{N} \mu_1^2$$

Therefore, an unbiased estimator for  $\mu_1^2$  can be computed as

$$\widehat{(\mu_1^2)} = \frac{N}{N-1} (\hat{\mu}_1^2 - \frac{1}{N} \hat{\mu}_2)$$

Second, we want to remove the dependence between the numerator and the denominator, which can be done simply by

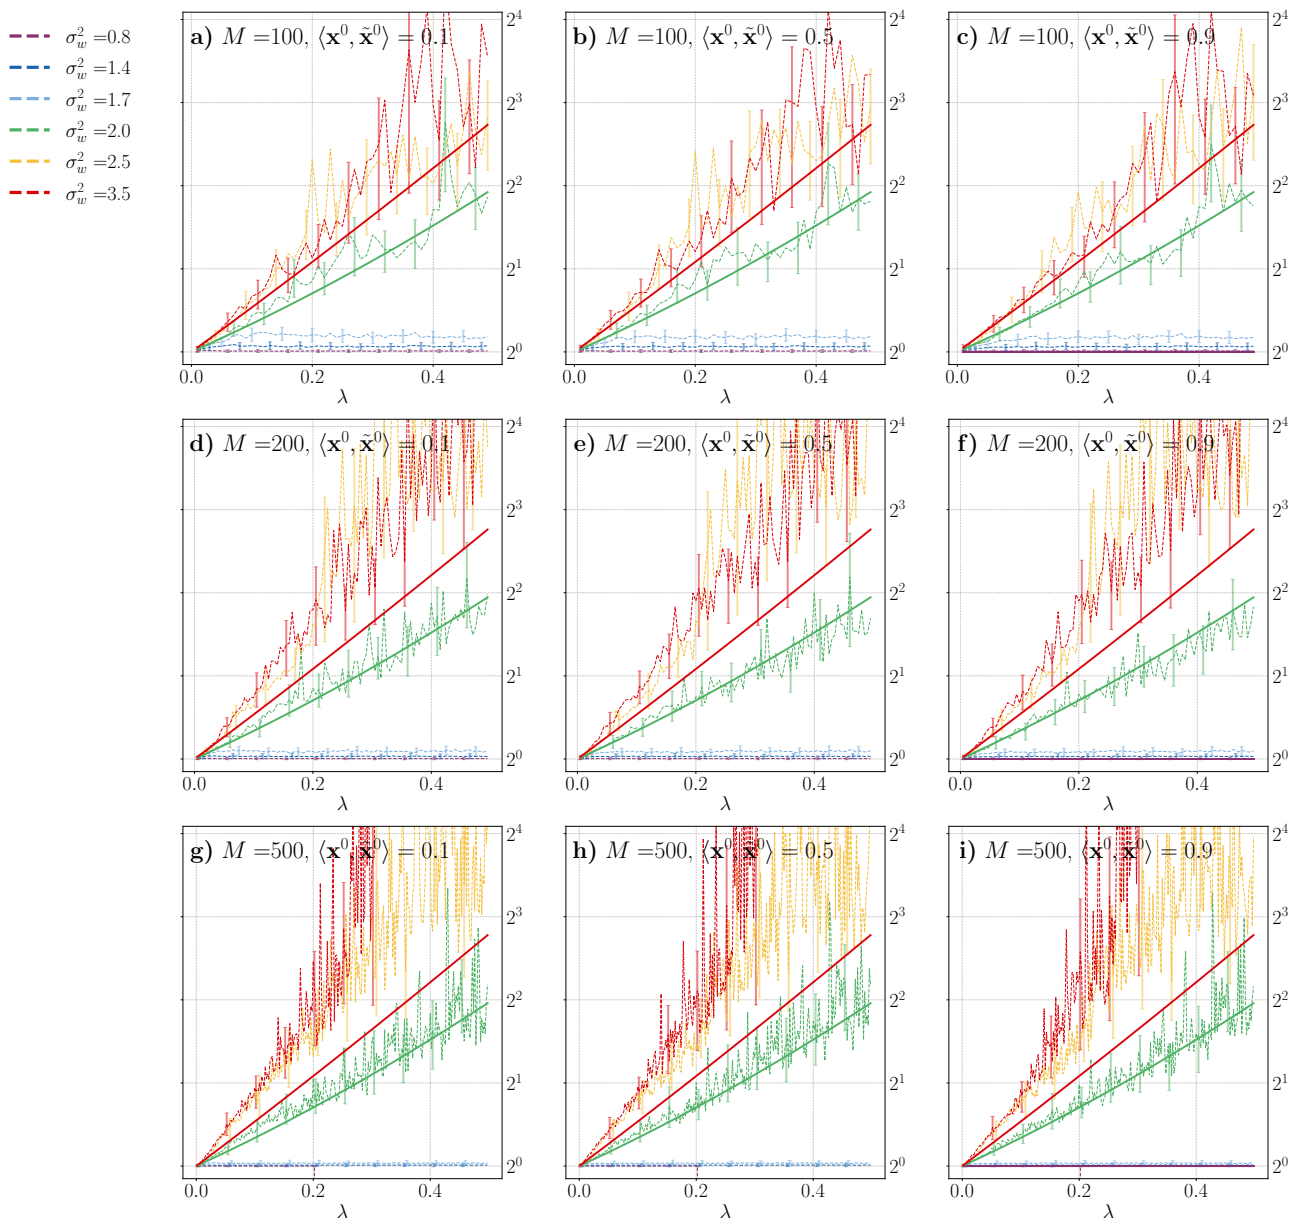


Figure 7. Ratio  $\mathbb{E}[\Theta^2(x, \tilde{x})]/\mathbb{E}^2[\Theta(x, \tilde{x})]$  at initialization (on a pair of different input samples, i.e.  $x \neq \tilde{x}$ ) for fully-connected ReLU networks of constant width  $M = 200$  with  $\alpha_0 \in \{2.0, 0.5, 0.1\}$  and varying initial angle between the input samples  $\langle \mathbf{x}^0, \tilde{\mathbf{x}}^0 \rangle \in \{0.1, 0.5, 0.9\}$ . The dashed lines show the experimental results and the solid lines show the corresponding theoretical predictions for diagonal elements of the NTK from Theorem 3.1. For each network configuration, we sampled 500 random initializations and computed an unbiased estimator for the ratio (see details in Appendix C.3). The error bars show the bootstrap estimation of the standard error.

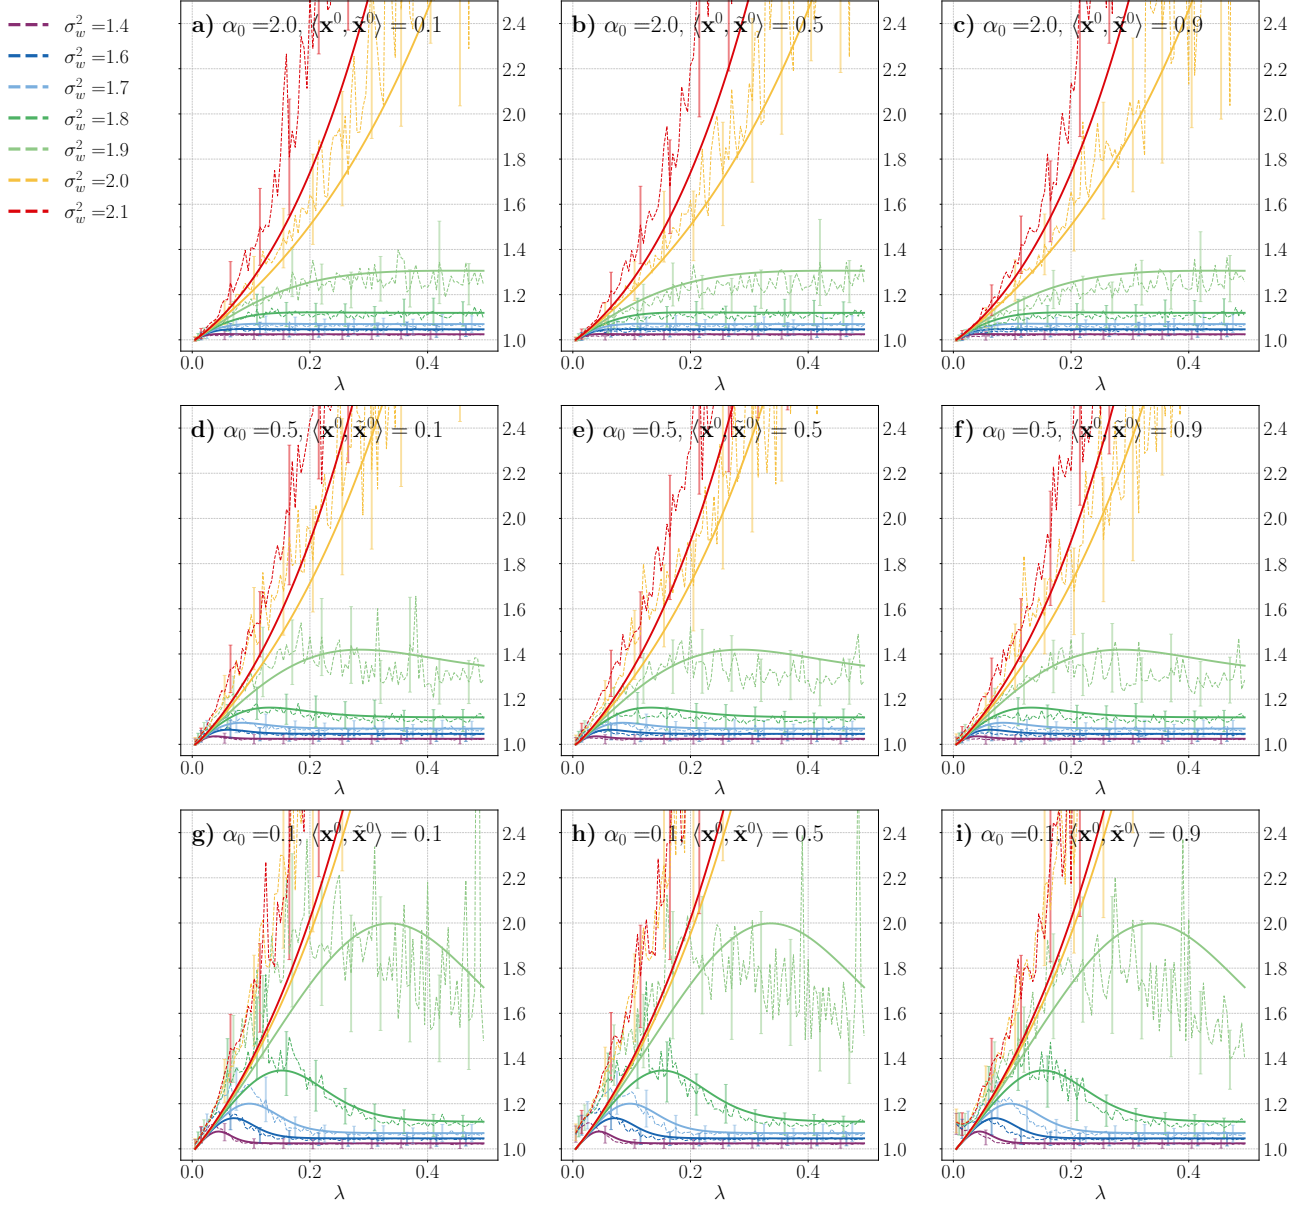


Figure 8. Ratio  $\mathbb{E}[\Theta^2(x, \tilde{x})]/\mathbb{E}^2[\Theta(x, \tilde{x})]$  at initialization (on a pair of different input samples, i.e.  $x \neq \tilde{x}$ ) for fully-connected ReLU networks of constant width  $M = 200$  with  $\alpha_0 \in \{2.0, 0.5, 0.1\}$  and varying initial angle between the input samples  $\langle \mathbf{x}^0, \tilde{\mathbf{x}}^0 \rangle \in \{0.1, 0.5, 0.9\}$ . The dashed lines show the experimental results and the solid lines show the theoretical predictions given by Theorem 3.2 for the diagonal elements of the NTK. For each network configuration, we sampled 500 random initializations and computed an unbiased estimator for the ratio (see details in Appendix C.3). The error bars show the bootstrap estimation of the standard error.




Figure 9. Ratio  $\mathbb{E}[\Theta^2(x, x)]/\mathbb{E}^2[\Theta(x, x)]$  at initialization for fully-connected ReLU DNNs of constant width  $M \in \{100, 200, 500\}$  with  $\alpha_0 = 1$ . The experiment setup is the same as in Figure 1. Continuous error bars show the bootstrap estimation of the standard error.

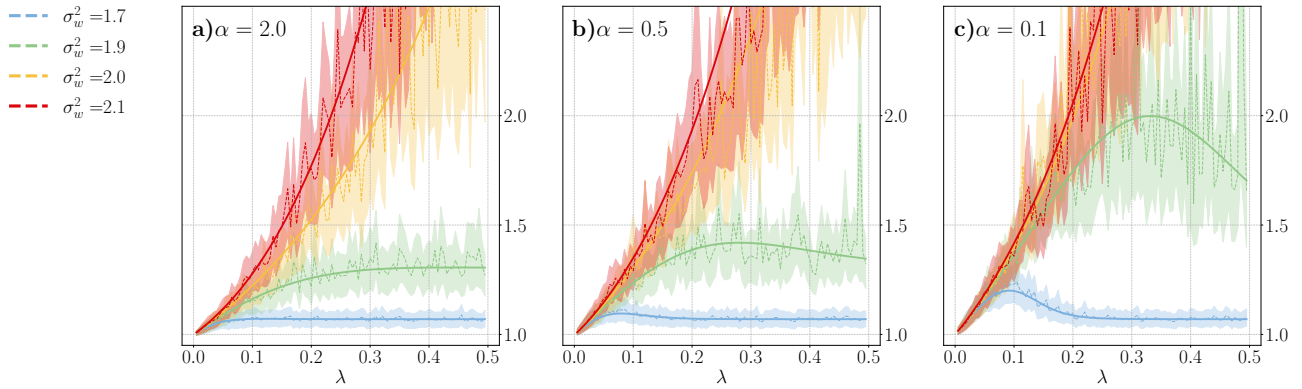


Figure 10. Ratio  $\mathbb{E}[\Theta^2(x, x)]/\mathbb{E}^2[\Theta(x, x)]$  at initialization for fully-connected ReLU networks of constant width  $M = 200$  with the ratio  $\alpha_0 := n_0/M \in \{2.0, 0.5, 0.1\}$ . The initialization hyperparameter  $\sigma_w^2$  is close to the EOC for all the lines. The experiment setup is the same as in Figure 2. Continuous error bars show the bootstrap estimation of the standard error.




Figure 11. Ratio  $\mathbb{E}[\Theta^2(x, x)]/\mathbb{E}^2[\Theta(x, x)]$  at initialization for fully-connected ReLU networks of constant width  $M = 200$  with the ratio  $\alpha_0 := n_0/M \in \{2.0, 0.5, 0.1\}$ . The initialization hyperparameter  $\sigma_w^2$  is close to the EOC for all the lines. The experiment setup is the same as in Figure 2. Continuous error bars show the bootstrap estimation of the standard error.

using disjoint parts of the sample to compute the two:

$$\hat{r} := \frac{1}{N-2} \sum_{i=1}^N \frac{\theta_i^2}{\frac{1}{(N-1)(N-2)} \left[ \left( \sum_{j \neq i} \theta_j \right)^2 - \sum_{j \neq i} \theta_j^2 \right]},$$

where we used the unbiased version of the estimator for  $\mu_1^2$  computed on a sample without  $\theta_i$  in the denominator. Then we can compute the expectation of our new estimator for the ratio as follows:

$$\mathbb{E}[\hat{r}] = (N-1) \sum_{i=1}^N \frac{\mathbb{E}[\theta_i^2]}{\mathbb{E} \left[ \left( \sum_{j \neq i} \theta_j \right)^2 - \sum_{j \neq i} \theta_j^2 \right]} = \frac{\mu_2}{\mu_1^2}.$$

Therefore,  $\hat{r}$  is an unbiased estimator of the ratio.