
The Surprising Simplicity of the Early-Time Learning Dynamics of Neural Networks

Wei Hu*

Lechao Xiao[†]

Ben Adlam[‡]

Jeffrey Pennington[§]

Abstract

Modern neural networks are often regarded as complex black-box functions whose behavior is difficult to understand owing to their nonlinear dependence on the data and the nonconvexity in their loss landscapes. In this work, we show that these common perceptions can be completely false in the early phase of learning. In particular, we formally prove that, for a class of well-behaved input distributions, the early-time learning dynamics of a two-layer fully-connected neural network can be mimicked by training a simple linear model on the inputs. We additionally argue that this surprising simplicity can persist in networks with more layers and with convolutional architecture, which we verify empirically. Key to our analysis is to bound the spectral norm of the difference between the Neural Tangent Kernel (NTK) at initialization and an affine transform of the data kernel; however, unlike many previous results utilizing the NTK, we do not require the network to have disproportionately large width, and the network is allowed to escape the kernel regime later in training.

1 Introduction

Modern deep learning models are enormously complex function approximators, with many state-of-the-art architectures employing millions or even billions of trainable parameters [Radford et al., 2019, Adiwardana et al., 2020]. While the raw parameter count provides only a crude approximation of a model’s capacity, more sophisticated metrics such as those based on PAC-Bayes [McAllester, 1999, Dziugaite and Roy, 2017, Neyshabur et al., 2017b], VC dimension [Vapnik and Chervonenkis, 1971], and parameter norms [Bartlett et al., 2017, Neyshabur et al., 2017a] also suggest that modern architectures have very large capacity. Moreover, from the empirical perspective, practical models are flexible enough to perfectly fit the training data, even if the labels are pure noise [Zhang et al., 2017]. Surprisingly, these same high-capacity models generalize well when trained on real data, even without any explicit control of capacity.

These observations are in conflict with classical generalization theory, which contends that models of intermediate complexity should generalize best, striking a balance between the bias and the variance of their predictive functions. To reconcile theory with observation, it has been suggested that deep neural networks may enjoy some form of implicit regularization induced by gradient-based training algorithms that biases the trained models towards simpler functions. However, the exact notion of simplicity and the mechanism by which it might be achieved remain poorly understood except in certain simplistic settings.

*Princeton University. Work partly performed at Google. Email: huwei@cs.princeton.edu

[†]Google Research, Brain Team. Email: xlc@google.com

[‡]Google Research, Brain Team. Work done as a member of the Google AI Residency program (<http://g.co/brainresidency>). Email: adlam@google.com

[§]Google Research, Brain Team. Email: jpennin@google.com

One concrete mechanism by which such induced simplicity can emerge is the hypothesis that neural networks learn simple functions early in training, and increasingly build up their complexity in later time. In particular, recent empirical work [Nakkiran et al. \[2019\]](#) found that, intriguingly, in some natural settings the simple function being learned in the early phase may just be a linear function of the data.

In this work, we provide a novel theoretical result to support this hypothesis. Specifically, we formally prove that, for a class of well-behaved input distributions, the early-time learning dynamics of gradient descent on a two-layer fully-connected neural network with any common activation can be mimicked by training a simple model of the inputs. When training the first layer only, this simple model is a linear function of the input features; when training the second layer or both layers, it is a linear function of the features and their ℓ_2 norm. This result implies that neural networks do not fully exercise their nonlinear capacity until late in training.

Key to our technical analysis is a bound on the spectral norm of the difference between the Neural Tangent Kernel (NTK) [[Jacot et al., 2018](#)] of the neural network at initialization and that of the linear model; indeed, a weaker result, like a bound on the Frobenius norm, would be insufficient to establish our result. Although the NTK is usually associated with the study of ultra-wide networks, our result only has a mild requirement on the width and allows the network to leave the kernel regime later in training. While our formal result focuses on two-layer fully-connected networks and data with benign concentration properties (specified in Assumption 3.1), we argue with theory and provide empirical evidence that the same linear learning phenomenon persists for more complex architectures and real-world datasets.

Related work. The early phase of neural network training has been the focus of considerable recent research. [Frankle and Carbin \[2019\]](#) found that sparse, trainable subnetworks – “lottery tickets” – emerge early in training. [Achille et al. \[2017\]](#) showed the importance of early learning from the perspective of creating strong connections that are robust to corruption. [Gur-Ari et al. \[2018\]](#) observed that after a short period of training, subsequent gradient updates span a low-dimensional subspace. [Li et al. \[2019a\]](#), [Lewkowycz et al. \[2020\]](#) showed that an initial large learning rate can benefit late-time generalization performance.

Implicit regularization of (stochastic) gradient descent has also been studied in various settings, suggesting a bias towards large-margin, low-norm, or low-rank solutions [[Gunasekar et al., 2017, 2018](#), [Soudry et al., 2018](#), [Li et al., 2018](#), [Ji and Telgarsky, 2019a,b](#), [Arora et al., 2019a](#), [Lyu and Li, 2019](#), [Chizat and Bach, 2020](#), [Razin and Cohen, 2020](#)]. These results mostly aim to characterize the final solutions at convergence, while our focus is on the early-time learning dynamics. Another line of work has identified that deep linear networks gradually increase the rank during training [[Arora et al., 2019a](#), [Saxe et al., 2014](#), [Lampinen and Ganguli, 2018](#), [Gidel et al., 2019](#)].

A line of work adopted the Fourier perspective and demonstrated that low-frequency functions are often learned first [[Rahaman et al., 2018](#), [Xu, 2018](#), [Xu et al., 2019a,b](#)]. Based on the NTK theory, [Arora et al. \[2019c\]](#) showed that for very wide networks, components lying in the top eigenspace of the NTK are learned faster than others. Using this principle, [Su and Yang \[2019\]](#), [Cao et al. \[2019\]](#) analyzed the spectrum of the infinite-width NTK. However, in order to obtain precise characterization of the spectrum these papers require special data distributions such as uniform distribution on the sphere.

Most relevant to our work is the finding of [Nakkiran et al. \[2019\]](#) that a neural network learned in the early phase of training can be almost fully explained by a linear function of the data. They supported this claim empirically by examining an information theoretic measure between the predictions of the neural network and the linear model. Our result formally proves that neural network and a corresponding linear model make similar predictions in early time, thus providing a theoretical explanation of their empirical finding.

Paper organization. In Section 2, we introduce notation and briefly recap the Neural Tangent Kernel. In Section 3, we present our main theoretical results on two-layer neural networks as well as empirical verification. In Section 4, we discuss extensions to more complicated architecture from both theoretical and empirical aspects. We conclude in Section 5, and defer additional experimental results and all the proofs to the appendices.

2 Preliminaries

Notation. We use bold lowercases $\mathbf{a}, \mathbf{b}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \dots$ to represent vectors, bold uppercases $\mathbf{A}, \mathbf{B}, \dots$ to represent matrices, and unbold letters $a, b, \alpha, \beta, \dots$ to represent scalars. We use $[\mathbf{A}]_{i,j}$ or $[\mathbf{a}]_i$ to index the entries in matrices or vectors. We denote by $\|\cdot\|$ the spectral norm (largest singular value) of a matrix or the ℓ_2 norm of a vector, and denote by $\|\cdot\|_F$ the Frobenius norm of a matrix. We use $\langle \cdot, \cdot \rangle$ to represent the standard Euclidean inner product between vectors or matrices, and use \odot to denote the Hadamard (entry-wise) product between matrices. For a positive semidefinite (psd) matrix \mathbf{A} , let $\mathbf{A}^{1/2}$ be the psd matrix such that $(\mathbf{A}^{1/2})^2 = \mathbf{A}$; let $\lambda_{\max}(\mathbf{A})$ and $\lambda_{\min}(\mathbf{A})$ be the maximum and minimum eigenvalues of \mathbf{A} .

Let $[n] := \{1, 2, \dots, n\}$. For $a, b \in \mathbb{R}$ ($b > 0$), we use $a \pm b$ to represent any number in the interval $[a - b, a + b]$. Let \mathbf{I}_d be the $d \times d$ identity matrix, $\mathbf{0}_d$ be the all-zero vector in \mathbb{R}^d , and $\mathbf{1}_d$ be the all-one vector in \mathbb{R}^d ; we write $\mathbf{I}, \mathbf{0}, \mathbf{1}$ when their dimensions are clear from context. We denote by $\text{Unif}(A)$ the uniform distribution over a set A , and by $\mathcal{N}(\mu, \sigma^2)$ or $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ the univariate/multivariate Gaussian distribution. Throughout the paper we let g be a random variable with the standard normal distribution $\mathcal{N}(0, 1)$.

We use the standard $O(\cdot)$, $\Omega(\cdot)$ and $\Theta(\cdot)$ notation to only hide universal constant factors. For $a, b \geq 0$, we also use $a \lesssim b$ or $b \gtrsim a$ to mean $a = O(b)$, and use $a \ll b$ or $b \gg a$ to mean $b \geq Ca$ for a sufficiently large universal constant $C > 0$. Throughout the paper, “high probability” means a large constant probability arbitrarily close to 1 (such as 0.99).

Recap of Neural Tangent Kernel (NTK) [Jacot et al., 2018]. Consider a single-output neural network $f(\mathbf{x}; \boldsymbol{\theta})$ where \mathbf{x} is the input and $\boldsymbol{\theta}$ is the collection of parameters in the network. Around a reference network with parameters $\bar{\boldsymbol{\theta}}$, we can do a local first-order approximation:

$$f(\mathbf{x}; \boldsymbol{\theta}) \approx f(\mathbf{x}; \bar{\boldsymbol{\theta}}) + \langle \nabla_{\boldsymbol{\theta}} f(\mathbf{x}; \bar{\boldsymbol{\theta}}), \boldsymbol{\theta} - \bar{\boldsymbol{\theta}} \rangle.$$

Thus when $\boldsymbol{\theta}$ is close to $\bar{\boldsymbol{\theta}}$, for a given input \mathbf{x} the network can be viewed as linear in $\nabla_{\boldsymbol{\theta}} f(\mathbf{x}; \bar{\boldsymbol{\theta}})$. This gradient feature map $\mathbf{x} \mapsto \nabla_{\boldsymbol{\theta}} f(\mathbf{x}; \bar{\boldsymbol{\theta}})$ induces a kernel $K_{\bar{\boldsymbol{\theta}}}(\mathbf{x}, \mathbf{x}') := \langle \nabla_{\boldsymbol{\theta}} f(\mathbf{x}; \bar{\boldsymbol{\theta}}), \nabla_{\boldsymbol{\theta}} f(\mathbf{x}'; \bar{\boldsymbol{\theta}}) \rangle$ which is called the NTK at $\bar{\boldsymbol{\theta}}$. Gradient descent training of the neural network can be viewed as kernel gradient descent on the function space with respect to the NTK. We use *NTK matrix* to refer to an $n \times n$ matrix that is the NTK evaluated on n datapoints.

While in general the NTK is random at initialization and can vary significantly during training, it was shown that, for a suitable network parameterization (known as the “NTK parameterization”), when the width goes to infinity or is sufficiently large, the NTK converges to a deterministic limit at initialization and barely changes during training [Jacot et al., 2018, Lee et al., 2019, Arora et al., 2019b, Yang, 2019], so that the neural network trained by gradient descent is equivalent to a kernel method with respect to a fixed kernel. However, for networks with practical widths, the NTK does usually stray far from its initialization.

3 Two-Layer Neural Networks

We consider a two-layer fully-connected neural network with m hidden neurons defined as:

$$f(\mathbf{x}; \mathbf{W}, \mathbf{v}) := \frac{1}{\sqrt{m}} \sum_{r=1}^m v_r \phi(\mathbf{w}_r^\top \mathbf{x} / \sqrt{d}) = \frac{1}{\sqrt{m}} \mathbf{v}^\top \phi(\mathbf{W} \mathbf{x} / \sqrt{d}), \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^d$ is the input, $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_m]^\top \in \mathbb{R}^{m \times d}$ is the weight matrix in the first layer, and $\mathbf{v} = [v_1, \dots, v_m]^\top \in \mathbb{R}^m$ is the weight vector in the second layer.⁵ Here $\phi: \mathbb{R} \rightarrow \mathbb{R}$ is an activation function that acts entry-wise on vectors or matrices.

Let $\{(\mathbf{x}_i, y_i)\}_{i=1}^n \subset \mathbb{R}^d \times \mathbb{R}$ be n training samples where \mathbf{x}_i ’s are the inputs and y_i ’s are their associated labels. Denote by $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top \in \mathbb{R}^{n \times d}$ the data matrix and by $\mathbf{y} = [y_1, \dots, y_n]^\top \in \mathbb{R}^n$ the label vector. We assume $|y_i| \leq 1$ for all $i \in [n]$.

⁵The scaling factors $\frac{1}{\sqrt{d}}$ and $\frac{1}{\sqrt{m}}$ are due to the NTK parameterization such that the weights can be initialized from $\mathcal{N}(0, 1)$. The standard parameterization can also be equivalently realized with the NTK parameterization by properly setting different learning rates in different layers [Lee et al., 2019], which we do allow here.

We consider the following ℓ_2 training loss:

$$L(\mathbf{W}, \mathbf{v}) := \frac{1}{2n} \sum_{i=1}^n (f(\mathbf{x}_i; \mathbf{W}, \mathbf{v}) - y_i)^2, \quad (2)$$

and run vanilla gradient descent (GD) on the objective (2) starting from random initialization. Specifically, we use the following *symmetric initialization* for the weights (\mathbf{W}, \mathbf{v}) :

$$\begin{aligned} \mathbf{w}_1, \dots, \mathbf{w}_{m/2} &\stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d), & \mathbf{w}_{i+m/2} &= \mathbf{w}_i \ (\forall i \in [m/2]), \\ v_1, \dots, v_{m/2} &\stackrel{\text{i.i.d.}}{\sim} \text{Unif}(\{1, -1\}), & v_{i+m/2} &= -v_i \ (\forall i \in [m/2]). \end{aligned} \quad (3)$$

The above initialization scheme was used by Chizat et al. [2019], Zhang et al. [2019], Hu et al. [2020], Bai and Lee [2020], etc. It initializes the network to be the difference between two identical (random) networks, which has the benefit of ensuring zero output: $f(\mathbf{x}; \mathbf{W}, \mathbf{v}) = 0$ ($\forall \mathbf{x} \in \mathbb{R}^d$), without altering the NTK at initialization. An alternative way to achieve the same effect is to subtract the function output at initialization [Chizat et al., 2019].

Let $(\mathbf{W}(0), \mathbf{v}(0))$ be a set of initial weights drawn from the symmetric initialization (3). Then the weights are updated according to GD:

$$\mathbf{W}(t+1) = \mathbf{W}(t) - \eta_1 \nabla_{\mathbf{W}} L(\mathbf{W}(t), \mathbf{v}(t)), \quad \mathbf{v}(t+1) = \mathbf{v}(t) - \eta_2 \nabla_{\mathbf{v}} L(\mathbf{W}(t), \mathbf{v}(t)), \quad (4)$$

where η_1 and η_2 are the learning rates. Here we allow potentially different learning rates for flexibility.

Now we state the assumption on the input distribution used in our theoretical results.

Assumption 3.1 (input distribution). *The datapoints $\mathbf{x}_1, \dots, \mathbf{x}_n$ are i.i.d. samples from a distribution \mathcal{D} over \mathbb{R}^d with mean $\mathbf{0}$ and covariance Σ such that $\text{Tr}[\Sigma] = d$ and $\|\Sigma\| = O(1)$. Moreover, $\mathbf{x} \sim \mathcal{D}$ can be written as $\mathbf{x} = \Sigma^{1/2} \bar{\mathbf{x}}$ where $\bar{\mathbf{x}} \in \mathbb{R}^d$ satisfies $\mathbb{E}[\bar{\mathbf{x}}] = \mathbf{0}_d$, $\mathbb{E}[\bar{\mathbf{x}} \bar{\mathbf{x}}^\top] = \mathbf{I}_d$, and $\bar{\mathbf{x}}$'s entries are independent and are all $O(1)$ -subgaussian.⁷*

Note that a special case that satisfies Assumption 3.1 is the Gaussian distribution $\mathcal{N}(\mathbf{0}, \Sigma)$, but we allow a much larger class of distributions here. The subgaussian assumption is made due to the probabilistic tail bounds used in the analysis, and it can be replaced with a weaker bounded moment condition. The independence between $\bar{\mathbf{x}}$'s entries may also be dropped if its density is strongly log-concave. We choose to use Assumption 3.1 as the most convenient way to present our results.

We allow ϕ to be any of the commonly used activation functions, including ReLU, Leaky ReLU, Erf, Tanh, Sigmoid, Softplus, etc. Formally, our requirement on ϕ is the following:

Assumption 3.2 (activation function). *The activation function $\phi(\cdot)$ satisfies either of the followings:*

(i) *smooth activation: ϕ has bounded first and second derivatives: $|\phi'(z)| = O(1)$ and $|\phi''(z)| = O(1)$ ($\forall z \in \mathbb{R}$), or*

(ii) *piece-wise linear activation: $\phi(z) = \begin{cases} z & (z \geq 0) \\ az & (z < 0) \end{cases}$ for some $a \in \mathbb{R}$, $|a| = O(1)$.⁸*

We will consider the regime where the data dimension d is sufficiently large (i.e., larger than any constant) and the number of datapoints n is at most some polynomial in d (i.e., $n \leq d^{O(1)}$). These imply $\log n = O(\log d) < d^c$ for any constant $c > 0$.

Under Assumption 3.1, the datapoints satisfy the following concentration properties:

Claim 3.1. *Suppose $n \gg d$. Then under Assumption 3.1, with high probability we have $\frac{\|\mathbf{x}_i\|^2}{d} = 1 \pm O\left(\sqrt{\frac{\log n}{d}}\right)$ ($\forall i \in [n]$), $\frac{|\langle \mathbf{x}_i, \mathbf{x}_j \rangle|}{d} = O\left(\sqrt{\frac{\log n}{d}}\right)$ ($\forall i, j \in [n], i \neq j$), and $\|\mathbf{X} \mathbf{X}^\top\| = \Theta(n)$.*

The main result in this section is to formally prove that the neural network trained by GD is approximately a linear function in the early phase of training. As we will see, there are distinct contributions coming from the two layers. Therefore, it is helpful to divide the discussion into the cases of training the first layer only, the second layer only, and both layers together. All the omitted proofs in this section are given in Appendix D.

⁶Our results also hold for $\mathcal{N}(0, 1)$ initialization in the second layer. Here we use $\text{Unif}(\{\pm 1\})$ for simplicity.

⁷Recall that a zero-mean random variable X is σ^2 -subgaussian if $\mathbb{E}[\exp(sX)] \leq \exp(\sigma^2 s^2/2)$ ($\forall s \in \mathbb{R}$).

⁸We define $\phi'(0) = 1$ in this case.

3.1 Training the First Layer

Now we consider only training the first layer weights \mathbf{W} , which corresponds to setting $\eta_2 = 0$ in (4). Denote by $f_t^1 : \mathbb{R}^d \rightarrow \mathbb{R}$ the network at iteration t in this case, namely $f_t^1(\mathbf{x}) := f(\mathbf{x}; \mathbf{W}(t), \mathbf{v}(t)) = f(\mathbf{x}; \mathbf{W}(t), \mathbf{v}(0))$ (note that $\mathbf{v}(t) = \mathbf{v}(0)$).

The linear model which will be proved to approximate the neural network f_t^1 in the early phase of training is $f_t^{\text{lin1}}(\mathbf{x}; \boldsymbol{\beta}) := \boldsymbol{\beta}^\top \boldsymbol{\psi}_1(\mathbf{x})$, where

$$\boldsymbol{\psi}_1(\mathbf{x}) := \frac{1}{\sqrt{d}} \begin{bmatrix} \zeta \mathbf{x} \\ \nu \end{bmatrix}, \quad \text{with } \zeta = \mathbb{E}[\phi'(g)] \text{ and } \nu = \mathbb{E}[g\phi'(g)] \cdot \sqrt{\text{Tr}[\boldsymbol{\Sigma}^2]/d}. \quad (5)$$

Here recall that $g \sim \mathcal{N}(0, 1)$. We also consider training this linear model via GD on the ℓ_2 loss, this time starting from zero:

$$\boldsymbol{\beta}(0) = \mathbf{0}_{d+1}, \quad \boldsymbol{\beta}(t+1) = \boldsymbol{\beta}(t) - \eta_1 \nabla_{\boldsymbol{\beta}} \frac{1}{2n} \sum_{i=1}^n (f_t^{\text{lin1}}(\mathbf{x}_i; \boldsymbol{\beta}(t)) - y_i)^2. \quad (6)$$

We let f_t^{lin1} be the model learned at iteration t , i.e., $f_t^{\text{lin1}}(\mathbf{x}) := f_t^{\text{lin1}}(\mathbf{x}; \boldsymbol{\beta}(t))$.

We emphasize that (4) and (6) have the same learning rate η_1 . Our theorem below shows that f_t^1 and f_t^{lin1} are close to each other in the early phase of training:

Theorem 3.2 (main theorem for training the first layer). *Let $\alpha \in (0, \frac{1}{4})$ be a fixed constant. Suppose the number of training samples n and the network width m satisfy $n \gtrsim d^{1+\alpha}$ and $m \gtrsim d^{1+\alpha}$. Suppose $\eta_1 \ll d$ and $\eta_2 = 0$. Then there exists a universal constant $c > 0$ such that with high probability, for all $0 \leq t \leq T = c \cdot \frac{d \log d}{\eta_1}$ simultaneously, the learned neural network f_t^1 and the linear model f_t^{lin1} at iteration t are close on average on the training data:*

$$\frac{1}{n} \sum_{i=1}^n (f_t^1(\mathbf{x}_i) - f_t^{\text{lin1}}(\mathbf{x}_i))^2 \lesssim d^{-\Omega(\alpha)}. \quad (7)$$

Moreover, f_t^1 and f_t^{lin1} are also close on the underlying data distribution \mathcal{D} . Namely, with high probability, for all $0 \leq t \leq T$ simultaneously, we have

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\min\{(f_t^1(\mathbf{x}) - f_t^{\text{lin1}}(\mathbf{x}))^2, 1\}] \lesssim d^{-\Omega(\alpha)} + \sqrt{\frac{\log T}{n}}. \quad (8)$$

Theorem 3.2 ensures that the neural network f_t^1 and the linear model f_t^{lin1} make almost the same predictions in the early time of training. This agreement is not only on the training data, but also over the underlying input distribution \mathcal{D} . Note that this does not mean that f_t^1 and f_t^{lin1} are the same on the entire space \mathbb{R}^d – they might still differ significantly at low-density regions of \mathcal{D} . We also remark that our result has no assumption on the labels $\{y_i\}$ except they are bounded.

The width requirement in Theorem 3.2 is very mild as it only requires the width m to be larger than $d^{1+\alpha}$ for some small constant α . Note that the width is allowed to be much smaller than the number of samples n , which is usually the case in practice.

The agreement guaranteed in Theorem 3.2 is up to iteration $T = c \cdot \frac{d \log d}{\eta_1}$ (for some constant c). It turns out that for well-conditioned data, after T iterations, a near optimal linear model will have been reached. This means that *the neural network in the early phase approximates a linear model all the way until the linear model converges to the optimum*. See Corollary 3.3 below.

Corollary 3.3 (well-conditioned data). *Under the same setting as Theorem 3.2, and additionally assume that the data distribution \mathcal{D} 's covariance $\boldsymbol{\Sigma}$ satisfies $\lambda_{\min}(\boldsymbol{\Sigma}) = \Omega(1)$. Let $\boldsymbol{\beta}_* \in \mathbb{R}^{d+1}$ be the optimal parameter for the linear model that GD (6) converges to, and denote $f_*^{\text{lin1}}(\mathbf{x}) := f^{\text{lin1}}(\mathbf{x}; \boldsymbol{\beta}_*)$. Then with high probability, after $T = c \cdot \frac{d \log d}{\eta_1}$ iterations (for some universal constant c), we have*

$$\frac{1}{n} \sum_{i=1}^n (f_T^1(\mathbf{x}_i) - f_*^{\text{lin1}}(\mathbf{x}_i))^2 \lesssim d^{-\Omega(\alpha)}, \quad \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\min\{(f_T^1(\mathbf{x}) - f_*^{\text{lin1}}(\mathbf{x}))^2, 1\}] \lesssim d^{-\Omega(\alpha)} + \sqrt{\frac{\log T}{n}}.$$

3.1.1 Proof Sketch of Theorem 3.2

The proof of Theorem 3.2 consists of showing that the NTK matrix for the first layer at random initialization evaluated on the training data is close to the kernel matrix corresponding to the linear model (5), and that furthermore this agreement persists in the early phase of training up to iteration T . Specifically, the NTK matrix $\Theta_1(\mathbf{W}) \in \mathbb{R}^{n \times n}$ at a given first-layer weight matrix \mathbf{W} , and the kernel matrix $\Theta^{\text{lin1}} \in \mathbb{R}^{n \times n}$ for the linear model (5) can be computed as:

$$\Theta_1(\mathbf{W}) := (\phi'(\mathbf{X}\mathbf{W}^\top/\sqrt{d})\phi'(\mathbf{X}\mathbf{W}^\top/\sqrt{d})^\top/m) \odot (\mathbf{X}\mathbf{X}^\top/d), \quad \Theta^{\text{lin1}} := (\zeta^2\mathbf{X}\mathbf{X}^\top + \nu^2\mathbf{1}\mathbf{1}^\top)/d.$$

We have the following result that bounds the difference between $\Theta_1(\mathbf{W}(0))$ and Θ^{lin1} in spectral norm:

Proposition 3.4. *With high probability over the random initialization $\mathbf{W}(0)$ and the training data \mathbf{X} , we have $\|\Theta_1(\mathbf{W}(0)) - \Theta^{\text{lin1}}\| \lesssim \frac{n}{d^{1+\alpha}}$.*

Notice that $\|\Theta^{\text{lin1}}\| = \Theta(\frac{n}{d})$ according to Claim 3.1. Thus the bound $\frac{n}{d^{1+\alpha}}$ in Proposition 3.4 is of smaller order. We emphasize that it is important to bound the spectral norm rather than the more naive Frobenius norm, since the latter would give $\|\Theta_1(\mathbf{W}(0)) - \Theta^{\text{lin1}}\|_F \gtrsim \frac{n}{d}$, which is not useful. (See Figure 5 for a numerical verification.)

To prove Proposition 3.4, we first use the matrix Bernstein inequality to bound the perturbation of $\Theta_1(\mathbf{W}(0))$ around its expectation with respect to $\mathbf{W}(0)$: $\|\Theta_1(\mathbf{W}(0)) - \mathbb{E}_{\mathbf{W}(0)}[\Theta_1(\mathbf{W}(0))]\| \lesssim \frac{n}{d^{1+\alpha}}$. Then we perform an entry-wise Taylor expansion of $\mathbb{E}_{\mathbf{W}(0)}[\Theta_1(\mathbf{W}(0))]$, and it turns out that the top-order terms exactly constitute Θ^{lin1} , and the rest can be bounded in spectral norm by $\frac{n}{d^{1+\alpha}}$.

After proving Proposition 3.4, in order to prove Theorem 3.2, we carefully track (i) the prediction difference between f_t^1 and f_t^{lin1} , (ii) how much the weight matrix \mathbf{W} move away from initialization, as well as (iii) how much the NTK changes. To prove the guarantee on the entire data distribution we further need to utilize tools from generalization theory. The full proof is given in Appendix D.

3.2 Training the Second Layer

Next we consider training the second layer weights v , which corresponds to $\eta_1 = 0$ in (4). Denote by $f_t^2: \mathbb{R}^d \rightarrow \mathbb{R}$ the network at iteration t in this case. We will show that training the second layer is also close to training a simple linear model $f^{\text{lin2}}(\mathbf{x}; \gamma) := \gamma^\top \psi_2(\mathbf{x})$ in the early phase, where:

$$\psi_2(\mathbf{x}) := \begin{bmatrix} \frac{1}{\sqrt{d}}\zeta\mathbf{x} \\ \frac{1}{\sqrt{2d}}\nu \\ \vartheta_0 + \vartheta_1(\frac{\|\mathbf{x}\|}{\sqrt{d}} - 1) + \vartheta_2(\frac{\|\mathbf{x}\|}{\sqrt{d}} - 1)^2 \end{bmatrix}, \quad \begin{cases} \zeta \text{ and } \nu \text{ are defined in (5),} \\ \vartheta_0 = \mathbb{E}[\phi(g)], \\ \vartheta_1 = \mathbb{E}[g\phi'(g)], \\ \vartheta_2 = \mathbb{E}[(\frac{1}{2}g^3 - g)\phi'(g)]. \end{cases} \quad (9)$$

As usual, this linear model is trained with GD starting from zero:

$$\gamma(0) = \mathbf{0}_{d+2}, \quad \gamma(t+1) = \gamma(t) - \eta_2 \nabla_\gamma \frac{1}{2n} \sum_{i=1}^n (f^{\text{lin2}}(\mathbf{x}_i; \gamma(t)) - y_i)^2. \quad (10)$$

We denote by f_t^{lin2} the resulting model at iteration t .

Note that strictly speaking $f^{\text{lin2}}(\mathbf{x}; \gamma)$ is not a linear model in \mathbf{x} because the feature map $\psi_2(\mathbf{x})$ contains a nonlinear feature depending on $\|\mathbf{x}\|$ in its last coordinate. Because $\frac{\|\mathbf{x}\|}{\sqrt{d}} \approx 1$ under our data assumption according to Claim 3.1, its effect might often be invisible. However, we emphasize that in general the inclusion of this norm-dependent feature is necessary, for example when the target function explicitly depends on the norm of the input. We illustrate this in Section 3.4.

Similar to Theorem 3.2, our main theorem for training the second layer is the following:

Theorem 3.5 (main theorem for training the second layer). *Let $\alpha \in (0, \frac{1}{4})$ be a fixed constant. Suppose $n \gtrsim d^{1+\alpha}$ and $\begin{cases} m \gtrsim d^{1+\alpha}, & \text{if } \mathbb{E}[\phi(g)] = 0 \\ m \gtrsim d^{2+\alpha}, & \text{otherwise} \end{cases}$. Suppose $\begin{cases} \eta_2 \ll d/\log n, & \text{if } \mathbb{E}[\phi(g)] = 0 \\ \eta_2 \ll 1, & \text{otherwise} \end{cases}$ and $\eta_1 = 0$. Then there exists a universal constant $c > 0$ such that with high probability, for all*




Figure 1: **Two-layer neural network learns a linear model early in training.** (a) Losses of a neural network and the corresponding linear model predicted by (11). Solid (dashed) lines represent the training (test) losses. We have $d = 50$, and use 20,000 training samples and 2,000 test samples. The neural network and the linear model are indistinguishable in the first 1,000 steps, after which linear learning finishes and the network continues to make progress. (b) Evolution of logits (i.e., outputs) of 5 random test examples. We see excellent agreement between the predictions of the neural network and the linear model in early time. (c) Discrepancy (in MSE) between the outputs of the network and the linear model for various values of d . As predicted, the discrepancy becomes smaller as d increases.

$0 \leq t \leq T = c \cdot \frac{d \log d}{\eta_2}$ simultaneously, we have

$$\frac{1}{n} \sum_{i=1}^n (f_t^2(\mathbf{x}_i) - f_t^{\text{lin}2}(\mathbf{x}_i))^2 \lesssim d^{-\Omega(\alpha)}, \quad \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\min\{(f_t^2(\mathbf{x}) - f_t^{\text{lin}2}(\mathbf{x}))^2, 1\}] \lesssim d^{-\Omega(\alpha)}.$$

Similar to Theorem 3.2, an important step in proving Theorem 3.5 is to prove that the NTK matrix for the second layer is close to the kernel for the linear model (9). Note that the theorem treats the case $\vartheta_0 = \mathbb{E}[\phi(g)] = 0$ differently. This is because when $\vartheta_0 \neq 0$, the second layer NTK has a large eigenvalue of size $\Theta(n)$, while when $\vartheta_0 = 0$, its largest eigenvalue is only $O(\frac{n \log n}{d})$.

We remark that if the data distribution is well-conditioned, we can also have a guarantee similar to Corollary 3.3.

3.3 Training Both Layers

Finally we consider the case where both layers are trained, in which $\eta_1 = \eta_2 = \eta > 0$ in (4). Since the NTK for training both layers is simply the sum of the first-layer NTK and the second-layer NTK, the corresponding linear model should have its kernel being the sum of the kernels for linear models (5) and (9), which can be derived easily:

$$f^{\text{lin}}(\mathbf{x}; \boldsymbol{\delta}) := \boldsymbol{\delta}^\top \boldsymbol{\psi}(\mathbf{x}), \quad \boldsymbol{\psi}(\mathbf{x}) := \begin{bmatrix} \sqrt{\frac{2}{d}} \zeta \mathbf{x} \\ \sqrt{\frac{3}{2d}} \nu \\ \vartheta_0 + \vartheta_1 \left(\frac{\|\mathbf{x}\|}{\sqrt{d}} - 1 \right) + \vartheta_2 \left(\frac{\|\mathbf{x}\|}{\sqrt{d}} - 1 \right)^2 \end{bmatrix}, \quad (11)$$

where the constants are from (9). Note that $\langle \boldsymbol{\psi}(\mathbf{x}), \boldsymbol{\psi}(\mathbf{x}') \rangle = \langle \boldsymbol{\psi}_1(\mathbf{x}), \boldsymbol{\psi}_1(\mathbf{x}') \rangle + \langle \boldsymbol{\psi}_2(\mathbf{x}), \boldsymbol{\psi}_2(\mathbf{x}') \rangle$.

Again, we can show that the neural network is close to the linear model (11) in early time. The guarantee is very similar to Theorems 3.2 and 3.5, so we defer the formal theorem to Appendix D; see Theorem D.1. Note that our result can be directly generalized to the case where $\eta_1 \neq \eta_2$, for which we just need to redefine the linear model using a weighted combination of the kernels for (5) and (9).

3.4 Empirical Verification

Verifying the early-time agreement between neural network and linear model. We verify our theory by training a two-layer neural network with erf activation and width 256 on synthetic data generated by $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $y = \text{sign}(f^*(\mathbf{x}))$, where f^* is a ground-truth two-layer erf network with width 5. In Figure 1a, we plot the training and test losses of the neural network (colored in




Figure 2: **The norm-dependent feature is necessary.** For the task of learning a norm-dependent function, test losses are shown for a **neural network** with ReLU activation, its corresponding **linear model** predicted by (11), and a **naive linear model** by resetting $\vartheta_1 = \vartheta_2 = 0$ in (11). Our predicted linear model is a much better approximation to the neural network than the naive linear model.

blue) and its corresponding linear model f^{lin} (in red).⁹ In the early training phase (up to 1,000 steps), the training/test losses of the network and the linear model are indistinguishable. After that, the optimal linear model is reached, and the network continues to make progress. In Figure 1b, we plot the evolution of the outputs (logits) of the network and the linear model on 5 random test examples, and we see excellent early-time agreement even on each individual sample. Finally, in Figure 1c, we vary the input dimension d , and for each case plot the mean squared error (MSE) of the discrepancies between the outputs of the network and the linear model. We see that the discrepancy indeed becomes smaller as d increases, matching our theoretical prediction.

The necessity of the norm-dependent feature. We now illustrate the necessity of including the norm-dependent feature in (11) and (9) through an example of learning a norm-dependent function. We generate data from $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $y = \frac{\|\mathbf{x}\|}{\sqrt{d}} + \text{ReLU}(\mathbf{a}^\top \mathbf{x})$ ($\|\mathbf{a}\| = O(1)$), and train a two-layer network with ReLU activation. We also train the corresponding linear model f^{lin} (11) as well as a “naive linear model” which is identical to f^{lin} except ϑ_1 and ϑ_2 are replaced with 0. Figure 2 shows that f^{lin} is indeed a much better approximation to the neural network than the naive linear model.

4 Extensions to Multi-Layer and Convolutional Neural Networks

In this section, we provide theoretical and empirical evidence supporting that the agreement between neural networks and linear models in the early phase of training may continue to hold for more complicated network architectures and datasets than what we analyzed in Section 3.

4.1 Theoretical Observations

Multi-layer fully-connected (FC) neural networks. For multi-layer FC networks, it was known that their infinite-width NTKs have the form $K(\mathbf{x}, \mathbf{x}') = h\left(\frac{\|\mathbf{x}\|^2}{d}, \frac{\|\mathbf{x}'\|^2}{d}, \frac{\langle \mathbf{x}, \mathbf{x}' \rangle}{d}\right)$ ($\mathbf{x}, \mathbf{x}' \in \mathbb{R}^d$) for some function $h: \mathbb{R}^3 \rightarrow \mathbb{R}$ [Yang and Salman, 2019]. Let Θ be the NTK matrix on the n training data: $[\Theta]_{i,j} = K(\mathbf{x}_i, \mathbf{x}_j)$. Under Assumption 3.1, we know from Claim 3.1 that $\frac{\|\mathbf{x}_i\|^2}{d} \approx 1$ and $\frac{\langle \mathbf{x}_i, \mathbf{x}_j \rangle}{d} \approx 0$ ($i \neq j$). Hence we can Taylor expand h around $(1, 1, 0)$ for the off-diagonal entries of Θ and around $(1, 1, 1)$ for the diagonal entries. Similar to our analysis of two-layer networks, we should be able to bound the higher-order components in the expansion, and only keep the simple ones like $\mathbf{X}\mathbf{X}^\top$, $\mathbf{1}\mathbf{1}^\top$, etc. This suggests that the early-time linear learning behavior which we showed for two-layer FC networks may persist in multi-layer FC networks.

Convolutional neural networks (CNNs). We consider a simple 1-dimensional CNN with one convolutional layer and without pooling (generalization to the commonly used 2-dimensional CNNs is straightforward):

$$f_{\text{CNN}}(\mathbf{x}; \mathbf{W}, \mathbf{V}) := \frac{1}{\sqrt{md}} \sum_{r=1}^m \mathbf{v}_r^\top \phi(\mathbf{w}_r * \mathbf{x} / \sqrt{q}). \quad (12)$$

Here $\mathbf{x} \in \mathbb{R}^d$ is the input, $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_m]^\top \in \mathbb{R}^{m \times q}$ and $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_m]^\top \in \mathbb{R}^{m \times d}$ contain the weights, where m is the number of channels (or width), and $q \leq d$ is the filter size. All the weights are initialized i.i.d from $\mathcal{N}(0, 1)$. The convolution operator $*$ is defined as: for input $\mathbf{x} \in \mathbb{R}^d$

⁹For $\phi = \text{erf}$, we have $\vartheta_0 = \vartheta_1 = \vartheta_2 = 0$, so f^{lin} in (11) is a linear model in \mathbf{x} without the nonlinear feature.




Figure 3: **Good agreement between 4-hidden-layer CNN/FC network and linear model on CIFAR-10 early in training.** (a) Decomposition of the test losses onto V_{lin} (solid lines) and V_{lin}^{\perp} (dashed lines) for **CNN**, **FC** and the corresponding **linear model**. (b) Three randomly selected test outputs for different models. (c) The relative MSE between the networks and the linear model. Note that we adjust the learning rates of **CNN** and **FC** so that their corresponding linear models are identical.

and filter $\mathbf{w} \in \mathbb{R}^q$, we have $\mathbf{w} * \mathbf{x} \in \mathbb{R}^d$ with $[\mathbf{w} * \mathbf{x}]_i := \sum_{j=1}^q [\mathbf{w}]_j [\mathbf{x}]_{i+j-1}$. We consider circular padding (as in Xiao et al. [2018], Li et al. [2019b]), so the indices in input should be understood as $[\mathbf{x}]_i = [\mathbf{x}]_{i+d}$.

We have the following result concerning the NTK of this CNN:

Proposition 4.1. *Let $\phi = \text{erf}$. Suppose $n \gtrsim d^{1+\alpha}$ and $q \gtrsim d^{\frac{1}{2}+2\alpha}$ for some constant $\alpha \in (0, \frac{1}{4})$. Consider n datapoints $\mathbf{x}_1, \dots, \mathbf{x}_n \stackrel{i.i.d.}{\sim} \text{Unif}(\{\pm 1\}^d)$. Then the corresponding NTK matrix $\Theta_{\text{CNN}} \in \mathbb{R}^{n \times n}$ of the CNN (12) in the infinite-width limit ($m \rightarrow \infty$) satisfies $\|\Theta_{\text{CNN}} - 2\zeta^2 \mathbf{X} \mathbf{X}^{\top} / d\| \lesssim \frac{n}{d^{1+\alpha}}$ with high probability, where $\zeta = \mathbb{E}[\phi'(g)]$.*

The proof is given in Appendix E. The above result shows that the NTK of a CNN can also be close to the (scaled) data kernel, which implies the linear learning behavior in the early time of training the CNN. Our empirical results will show that this behavior can even persist to multi-layer CNNs and real data beyond our analysis.

4.2 Empirical Results

We perform experiments on a binary classification task from CIFAR-10 (“cats” vs “horses”) using a multi-layer FC network and a CNN. The numbers of training and test data are 10,000 and 2,000. The original size of the images is $32 \times 32 \times 3$, and we down-sample the images into size $8 \times 8 \times 3$ using a 4×4 average pooling. Then we train a 4-hidden-layer FC net and a 4-hidden-layer CNN with erf activation. To have finer-grained examination of the evolution of the losses, we decompose the residual of the predictions on test data (namely, $f_t(\mathbf{x}) - y$ for all test data collected as a vector in \mathbb{R}^{2000}) onto V_{lin} , the space spanned by the inputs (of dimension $d = 192$), and its complement V_{lin}^{\perp} (of dimension $2000 - d$). For both networks, we observe in Figure 3a that the test losses of the networks and the linear model are almost identical up to 1,000 steps, and the networks start to make progress in V_{lin}^{\perp} after that. In Figure 3b we plot the logit evolution of 3 random test datapoints and again observe good agreement in early time. In Figure 3c, we plot the relative MSE between the network and the linear model (i.e., $\mathbb{E}_{\mathbf{x}} \|f_t(\mathbf{x}) - f_t^{\text{lin}}(\mathbf{x})\|^2 / \mathbb{E}_{\mathbf{x}} \|f_t^{\text{lin}}(\mathbf{x})\|^2$ evaluated on test data). We observe that this quantity for either network is small in the first 1,000 steps and grows afterwards. The detailed setup and additional results for full-size CIFAR-10 and MNIST are deferred to Appendix A.

5 Conclusion

This work gave a novel theoretical result rigorously showing that gradient descent on a neural network learns a simple linear function in the early phase. While we mainly focused on two-layer fully-connected neural networks, we further provided theoretical and empirical evidence suggesting that this phenomenon continues to exist in more complicated models. Formally extending our result to those settings is a direction of future work. Another interesting direction is to study the dynamics of neural networks after the initial linear learning phase.

Broader Impact

This work is theoretical and does not present any foreseeable societal consequence.

Acknowledgments and Disclosure of Funding

WH was supported by NSF, ONR, Simons Foundation, Schmidt Foundation, Amazon Research, DARPA and SRC.

References

- Alessandro Achille, Matteo Rovere, and Stefano Soatto. Critical learning periods in deep neural networks. *arXiv preprint arXiv:1711.08856*, 2017.
- Daniel Adiwardana, Minh-Thang Luong, David R So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu, et al. Towards a human-like open-domain chatbot. *arXiv preprint arXiv:2001.09977*, 2020.
- Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. In *Advances in Neural Information Processing Systems*, pages 7411–7422, 2019a.
- Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Ruslan Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. *arXiv preprint arXiv:1904.11955*, 2019b.
- Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. *arXiv preprint arXiv:1901.08584*, 2019c.
- Yu Bai and Jason D. Lee. Beyond linearization: On quadratic and higher-order approximation of wide neural networks. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rk1lGyBFPH>.
- Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- Peter L Bartlett, Dylan J Foster, and Matus J Telgarsky. Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems*, pages 6241–6250, 2017.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, and Skye Wanderman-Milne. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Yuan Cao, Zhiying Fang, Yue Wu, Ding-Xuan Zhou, and Quanquan Gu. Towards understanding the spectral bias of deep learning. *arXiv preprint arXiv:1912.01198*, 2019.
- Lenaic Chizat and Francis Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. *arXiv preprint arXiv:2002.04486*, 2020.
- Lenaic Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems*, pages 2933–2943, 2019.
- Simon S Du, Wei Hu, Sham M Kakade, Jason D Lee, and Qi Lei. Few-shot learning via learning the representation, provably. *arXiv preprint arXiv:2002.09434*, 2020.
- Gintare Karolina Dziugaite and Daniel M Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. *arXiv preprint arXiv:1703.11008*, 2017.
- Noureddine El Karoui. The spectrum of kernel random matrices. *The Annals of Statistics*, 38(1): 1–50, 2010.

- Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rJl-b3RcF7>.
- Gauthier Gidel, Francis Bach, and Simon Lacoste-Julien. Implicit regularization of discrete gradient dynamics in linear neural networks. In *Advances in Neural Information Processing Systems*, pages 3196–3206, 2019.
- Suriya Gunasekar, Blake E Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Implicit regularization in matrix factorization. In *Advances in Neural Information Processing Systems*, pages 6151–6159, 2017.
- Suriya Gunasekar, Jason Lee, Daniel Soudry, and Nathan Srebro. Implicit bias of gradient descent on linear convolutional networks. *arXiv preprint arXiv:1806.00468*, 2018.
- Guy Gur-Ari, Daniel A Roberts, and Ethan Dyer. Gradient descent happens in a tiny subspace. *arXiv preprint arXiv:1812.04754*, 2018.
- Wei Hu, Zhiyuan Li, and Dingli Yu. Simple and effective regularization methods for training on noisily labeled data with generalization guarantee. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=Hke3gyHYwH>.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *arXiv preprint arXiv:1806.07572*, 2018.
- Ziwei Ji and Matus Telgarsky. The implicit bias of gradient descent on nonseparable data. In *Conference on Learning Theory*, pages 1772–1798, 2019a.
- Ziwei Ji and Matus Jan Telgarsky. Gradient descent aligns the layers of deep linear networks. In *7th International Conference on Learning Representations, ICLR 2019*, 2019b.
- Yegor Klochkov and Nikita Zhivotovskiy. Uniform hanson-wright type concentration inequalities for unbounded entries via the entropy method. *Electronic Journal of Probability*, 25, 2020.
- Andrew K Lampinen and Surya Ganguli. An analytic theory of generalization dynamics and transfer learning in deep linear networks. *arXiv preprint arXiv:1809.10374*, 2018.
- Jaehoon Lee, Lechao Xiao, Samuel S Schoenholz, Yasaman Bahri, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. *arXiv preprint arXiv:1902.06720*, 2019.
- Aitor Lewkowycz, Yasaman Bahri, Ethan Dyer, Jascha Sohl-Dickstein, and Guy Gur-Ari. The large learning rate phase of deep learning: the catapult mechanism. *arXiv preprint arXiv:2003.02218*, 2020.
- Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In *Conference On Learning Theory*, pages 2–47, 2018.
- Yuanzhi Li, Colin Wei, and Tengyu Ma. Towards explaining the regularization effect of initial large learning rate in training neural networks. In *Advances in Neural Information Processing Systems*, pages 11669–11680, 2019a.
- Zhiyuan Li, Ruosong Wang, Dingli Yu, Simon S Du, Wei Hu, Ruslan Salakhutdinov, and Sanjeev Arora. Enhanced convolutional neural tangent kernels. *arXiv preprint arXiv:1911.00809*, 2019b.
- Kaifeng Lyu and Jian Li. Gradient descent maximizes the margin of homogeneous neural networks. *arXiv preprint arXiv:1906.05890*, 2019.
- David A McAllester. Pac-bayesian model averaging. In *Proceedings of the twelfth annual conference on Computational learning theory*, pages 164–170, 1999.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. Foundations of machine learning. *MIT Press*, 2012.

- Preetum Nakkiran, Gal Kaplun, Dimitris Kalimeris, Tristan Yang, Benjamin L Edelman, Fred Zhang, and Boaz Barak. Sgd on neural networks learns functions of increasing complexity. *arXiv preprint arXiv:1905.11604*, 2019.
- Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nathan Srebro. A PAC-Bayesian approach to spectrally-normalized margin bounds for neural networks. *arXiv preprint arXiv:1707.09564*, 2017a.
- Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems*, pages 5947–5956, 2017b.
- Roman Novak, Lechao Xiao, Jiri Hron, Jaehoon Lee, Alexander A Alemi, Jascha Sohl-Dickstein, and Samuel S Schoenholz. Neural tangents: Fast and easy infinite neural networks in python. *arXiv preprint arXiv:1912.02803*, 2019.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9, 2019.
- Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred A Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. *arXiv preprint arXiv:1806.08734*, 2018.
- Noam Razin and Nadav Cohen. Implicit regularization in deep learning may not be explainable by norms. *arXiv preprint arXiv:2005.06398*, 2020.
- Mark Rudelson and Roman Vershynin. Hanson-wright inequality and sub-gaussian concentration. *Electronic Communications in Probability*, 18, 2013.
- AM Saxe, JL McClelland, and S Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. International Conference on Learning Representations, 2014.
- Jssai Schur. Bemerkungen zur theorie der beschränkten bilinearformen mit unendlich vielen veränderlichen. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 1911(140):1–28, 1911.
- Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *Journal of Machine Learning Research*, 19(70), 2018.
- Lili Su and Pengkun Yang. On learning over-parameterized neural networks: A functional approximation perspective. In *Advances in Neural Information Processing Systems*, pages 2637–2646, 2019.
- Joel A Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015.
- VN Vapnik and A Ya Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280, 1971.
- Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- Lechao Xiao, Yasaman Bahri, Jascha Sohl-Dickstein, Samuel Schoenholz, and Jeffrey Pennington. Dynamical isometry and a mean field theory of cnns: How to train 10,000-layer vanilla convolutional neural networks. In *International Conference on Machine Learning*, pages 5393–5402, 2018.
- Zhi-Qin John Xu, Yaoyu Zhang, Tao Luo, Yanyang Xiao, and Zheng Ma. Frequency principle: Fourier analysis sheds light on deep neural networks. *arXiv preprint arXiv:1901.06523*, 2019a.
- Zhi-Qin John Xu, Yaoyu Zhang, and Yanyang Xiao. Training behavior of deep neural network in frequency domain. In *International Conference on Neural Information Processing*, pages 264–274. Springer, 2019b.

Zhiqin John Xu. Understanding training and generalization in deep learning by fourier analysis. *arXiv preprint arXiv:1808.04295*, 2018.

Greg Yang. Scaling limits of wide neural networks with weight sharing: Gaussian process behavior, gradient independence, and neural tangent kernel derivation. *arXiv preprint arXiv:1902.04760*, 2019.

Greg Yang and Hadi Salman. A fine-grained spectral perspective on neural networks. *arXiv preprint arXiv:1907.10599*, 2019.

Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.

Yaoyu Zhang, Zhi-Qin John Xu, Tao Luo, and Zheng Ma. A type of generalization error induced by initialization in deep neural networks. *arXiv preprint arXiv:1905.07777*, 2019.

A Experiment Setup and Additional Plots

We provide additional plots and describe additional experiment details in this section.

In Figure 4, we repeat the same experiments in Figure 3 on the full-size ($32 \times 32 \times 3$) CIFAR-10 as well as MNIST datasets, using the same 4-hidden-layer FC and CNN architectures. For both datasets we take two classes and perform binary classification. We see very good early-time agreement except for CNN on CIFAR-10, where the agreement only lasts for a shorter time.

For the experiments in Figures 3 and 4, the FC network has width 512 in each of the 4 hidden layers, and the CNN uses circular padding and has 256 channels in each of the 4 hidden layers. For CIFAR-10 and MNIST images, we use standard data pre-processing, i.e., normalizing each image to have zero mean and unit variance. To ensure the initial outputs are always 0, we subtract the function output at initialization for each datapoint (as discussed in Section 3). We train and test using the ℓ_2 loss with ± 1 labels. We use vanilla stochastic gradient descent with batch size 500, and choose a small learning rate (roughly $\frac{0.01}{\|\mathbf{N}^{\text{TK}}\|}$) so that we can better observe early time of training (similar to Nakkiran et al. [2019]).

We use the Neural Tangents Library [Novak et al., 2019] and JAX [Bradbury et al., 2018] for our experiments.

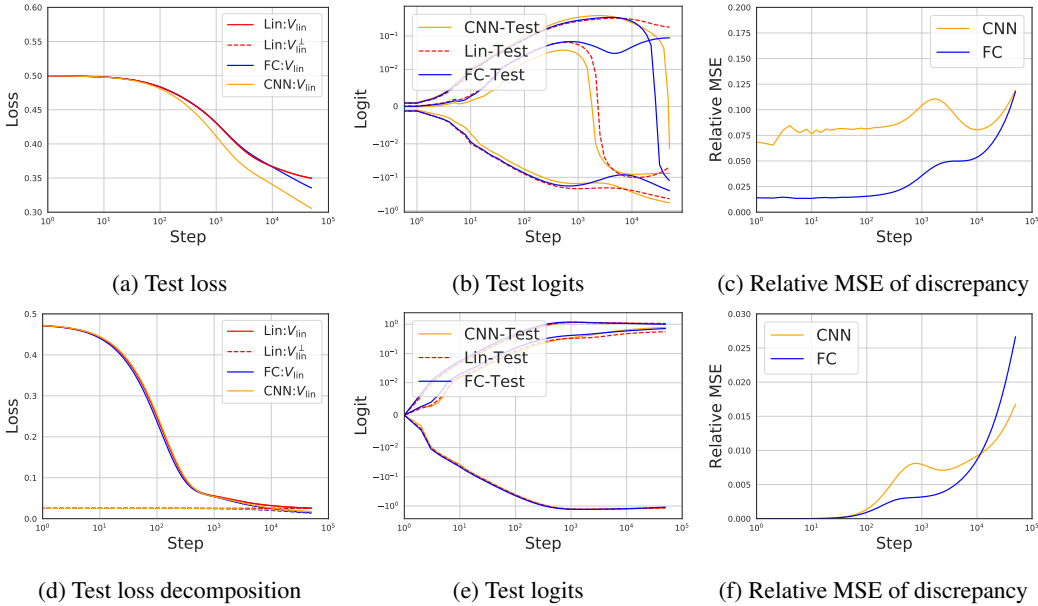


Figure 4: Replication of Figure 3 on full-size CIFAR-10 (top row) and MNIST (bottom row). In Figure 4a, there is no projection onto V_{lin}^\perp because the data dimension $32 \times 32 \times 3$ is larger than the number of test data 2,000.

B Additional Notation and Lemmas

We introduce some additional notation and lemmas that will be used in the proofs.

We use $\tilde{O}(\cdot)$ to hide poly-logarithmic factors in n (the number of training datapoints). Denote by $\mathbb{1}_{\{E\}}$ the indicator function for an event E . For a vector \mathbf{a} , we let $\text{diag}(\mathbf{a})$ be a diagonal matrix whose diagonal entries constitute \mathbf{a} . For a matrix \mathbf{A} , we use $\text{vec}(\mathbf{A})$ to denote the vectorization of \mathbf{A} in row-first order.

For a square matrix \mathbf{A} , we denote its diagonal and off-diagonal parts as \mathbf{A}_{diag} and \mathbf{A}_{off} , respectively. Namely, we have $\mathbf{A} = \mathbf{A}_{\text{diag}} + \mathbf{A}_{\text{off}}$, where $[\mathbf{A}_{\text{diag}}]_{i,j} = [\mathbf{A}]_{i,j} \mathbb{1}_{\{i=j\}}$ and $[\mathbf{A}_{\text{off}}]_{i,j} = [\mathbf{A}]_{i,j} \mathbb{1}_{\{i \neq j\}}$. Equivalently, $\mathbf{A}_{\text{diag}} = \mathbf{A} \odot \mathbf{I}$ and $\mathbf{A}_{\text{off}} = \mathbf{A} \odot (\mathbf{1}\mathbf{1}^\top - \mathbf{I})$.

Lemma B.1. For any matrix \mathbf{A} and a submatrix \mathbf{A}_1 of \mathbf{A} , we have $\|\mathbf{A}_1\| \leq \|\mathbf{A}\|$.

Proof. For simplicity we assume that \mathbf{A}_1 is in the top-left corner of \mathbf{A} , i.e. $\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 \\ \mathbf{A}_3 & \mathbf{A}_4 \end{bmatrix}$. The same proof works when \mathbf{A}_1 is any other submatrix of \mathbf{A} .

By the definition of spectral norm, we have

$$\begin{aligned} \|\mathbf{A}\| &= \max_{\|\mathbf{x}\|=\|\mathbf{y}\|=1} \mathbf{x}^\top \mathbf{A} \mathbf{y} \\ &= \max_{\|\mathbf{x}\|=\|\mathbf{y}\|=1} \mathbf{x}^\top \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 \\ \mathbf{A}_3 & \mathbf{A}_4 \end{bmatrix} \mathbf{y} \\ &\geq \max_{\|\mathbf{x}_1\|=\|\mathbf{y}_1\|=1} [\mathbf{x}_1^\top, \mathbf{0}^\top] \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 \\ \mathbf{A}_3 & \mathbf{A}_4 \end{bmatrix} \begin{bmatrix} \mathbf{y}_1 \\ \mathbf{0} \end{bmatrix} \\ &= \max_{\|\mathbf{x}_1\|=\|\mathbf{y}_1\|=1} \mathbf{x}_1^\top \mathbf{A}_1 \mathbf{y}_1 \\ &= \|\mathbf{A}_1\|. \end{aligned} \quad \square$$

Lemma B.2. For any square matrix \mathbf{A} , we have $\|\mathbf{A}_{\text{diag}}\| \leq \|\mathbf{A}\|$ and $\|\mathbf{A}_{\text{off}}\| \leq 2\|\mathbf{A}\|$.

Proof. From Lemma B.1 we know that $|[\mathbf{A}]_{i,i}| \leq \|\mathbf{A}\|$ for all i since $[\mathbf{A}]_{i,i}$ can be viewed as a submatrix of \mathbf{A} . Thus we have

$$\|\mathbf{A}_{\text{diag}}\| = \max_i |[\mathbf{A}]_{i,i}| \leq \|\mathbf{A}\|.$$

It follows that

$$\|\mathbf{A}_{\text{off}}\| = \|\mathbf{A} - \mathbf{A}_{\text{diag}}\| \leq \|\mathbf{A}\| + \|\mathbf{A}_{\text{diag}}\| \leq 2\|\mathbf{A}\|. \quad \square$$

Lemma B.3 (Schur [1911]). For any two positive semidefinite matrices \mathbf{A}, \mathbf{B} , we have

$$\|\mathbf{A} \odot \mathbf{B}\| \leq \|\mathbf{A}\| \cdot \max_i [\mathbf{B}]_{i,i}.$$

C General Result on the Closeness between Two Dynamics

We present a general result that shows how the GD trajectory for a non-linear least squares problem can be simulated by a linear one. Later we will specialize this result to the settings considered in the paper.

We consider an objective function of the form:

$$F(\boldsymbol{\theta}) = \frac{1}{2n} \|\mathbf{f}(\boldsymbol{\theta}) - \mathbf{y}\|^2,$$

where $\mathbf{f} : \mathbb{R}^N \mapsto \mathbb{R}^n$ is a general differentiable function, and $\mathbf{y} \in \mathbb{R}^n$ satisfies $\|\mathbf{y}\| \leq \sqrt{n}$. We denote by $\mathbf{J} : \mathbb{R}^N \mapsto \mathbb{R}^{n \times N}$ the Jacobian map of \mathbf{f} . Then starting from some $\boldsymbol{\theta}(0) \in \mathbb{R}^N$, the GD updates for minimizing F can be written as:

$$\boldsymbol{\theta}(t+1) = \boldsymbol{\theta}(t) - \eta \nabla F(\boldsymbol{\theta}(t)) = \boldsymbol{\theta}(t) - \frac{1}{n} \eta \mathbf{J}(\boldsymbol{\theta}(t))^\top (\mathbf{f}(\boldsymbol{\theta}(t)) - \mathbf{y}).$$

Consider another linear least squares problem:

$$G(\boldsymbol{\omega}) = \frac{1}{2n} \|\Phi \boldsymbol{\omega} - \mathbf{y}\|^2,$$

where $\Phi \in \mathbb{R}^{n \times M}$ is a fixed matrix. Its GD dynamics started from $\boldsymbol{\omega}(0) \in \mathbb{R}^M$ can be written as:

$$\boldsymbol{\omega}(t+1) = \boldsymbol{\omega}(t) - \eta \nabla G(\boldsymbol{\omega}(t)) = \boldsymbol{\omega}(t) - \frac{1}{n} \eta \Phi^\top (\Phi \boldsymbol{\omega}(t) - \mathbf{y}).$$

Let $\mathbf{K} := \Phi \Phi^\top$, and let

$$\mathbf{u}(t) := \mathbf{f}(\boldsymbol{\theta}(t)),$$

$$\mathbf{u}^{\text{lin}}(t) := \Phi\omega(t),$$

which stand for the predictions of these two models at iteration t .

The linear dynamics admit a very simple analytical form, summarized below.

Claim C.1. *For all $t \geq 0$ we have $\mathbf{u}^{\text{lin}}(t) - \mathbf{y} = (\mathbf{I} - \frac{1}{n}\eta\mathbf{K})^t (\mathbf{u}^{\text{lin}}(0) - \mathbf{y})$. As a consequence, if $\eta \leq \frac{2n}{\|\mathbf{K}\|}$, then we have $\|\mathbf{u}^{\text{lin}}(t) - \mathbf{y}\| \leq \|\mathbf{u}^{\text{lin}}(0) - \mathbf{y}\|$ for all $t \geq 0$.*

Proof. By definition we have $\mathbf{u}^{\text{lin}}(t+1) = \mathbf{u}^{\text{lin}}(t) - \frac{1}{n}\eta\mathbf{K}(\mathbf{u}^{\text{lin}}(t) - \mathbf{y})$, which implies $\mathbf{u}^{\text{lin}}(t+1) - \mathbf{y} = (\mathbf{I} - \frac{1}{n}\eta\mathbf{K})(\mathbf{u}^{\text{lin}}(t) - \mathbf{y})$. Thus the first statement follows directly. Then the second statement can be proved by noting that $\|\mathbf{I} - \frac{1}{n}\eta\mathbf{K}\| \leq 1$ when $\eta \leq \frac{2n}{\|\mathbf{K}\|}$. \square

We make the following assumption that connects these two problems:

Assumption C.1. *There exist $0 < \epsilon < \|\mathbf{K}\|$, $R > 0$ such that for any $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathbb{R}^N$, as long as $\|\boldsymbol{\theta} - \boldsymbol{\theta}(0)\| \leq R$ and $\|\boldsymbol{\theta}' - \boldsymbol{\theta}(0)\| \leq R$, we have*

$$\|\mathbf{J}(\boldsymbol{\theta})\mathbf{J}(\boldsymbol{\theta}')^\top - \mathbf{K}\| \leq \epsilon.$$

Based on the above assumption, we have the following theorem showing the agreement between $\mathbf{u}(t)$ and $\mathbf{u}^{\text{lin}}(t)$ as well as the parameter boundedness in early time.

Theorem C.2. *Suppose that the initializations are chosen so that $\mathbf{u}(0) = \mathbf{u}^{\text{lin}}(0) = \mathbf{0}$, and that the learning rate satisfies $\eta \leq \frac{n}{\|\mathbf{K}\|}$. Suppose that Assumption C.1 is satisfied with $R^2\epsilon < n$. Then there exists a universal constant $c > 0$ such that for all $0 \leq t \leq c\frac{R^2}{\eta}$:*

- (closeness of predictions) $\|\mathbf{u}(t) - \mathbf{u}^{\text{lin}}(t)\| \lesssim \frac{\eta t \epsilon}{\sqrt{n}}$;
- (boundedness of parameter movement) $\|\boldsymbol{\theta}(t) - \boldsymbol{\theta}(0)\| \leq R, \|\omega(t) - \omega(0)\| \leq R$.

Proof. We first prove the first two properties, and will prove the last property $\|\omega(t) - \omega(0)\| \leq R$ at the end.

We use induction to prove $\|\mathbf{u}(t) - \mathbf{u}^{\text{lin}}(t)\| \lesssim \frac{\eta t \epsilon}{\sqrt{n}}$ and $\|\boldsymbol{\theta}(t) - \boldsymbol{\theta}(0)\| \leq R$. For $t = 0$, these statements are trivially true. Now suppose for some $1 \leq t \leq c\frac{R^2}{\eta}$ we have $\|\mathbf{u}(\tau) - \mathbf{u}^{\text{lin}}(\tau)\| \lesssim \frac{\eta \tau \epsilon}{\sqrt{n}}$ and $\|\boldsymbol{\theta}(\tau) - \boldsymbol{\theta}(0)\| \leq R$ for $\tau = 0, 1, \dots, t-1$. We will now prove $\|\mathbf{u}(t) - \mathbf{u}^{\text{lin}}(t)\| \lesssim \frac{\eta t \epsilon}{\sqrt{n}}$ and $\|\boldsymbol{\theta}(t) - \boldsymbol{\theta}(0)\| \leq R$ under these induction hypotheses.

Notice that from $\|\mathbf{u}(\tau) - \mathbf{u}^{\text{lin}}(\tau)\| \lesssim \frac{\eta \tau \epsilon}{\sqrt{n}} \leq \frac{cR^2\epsilon}{\sqrt{n}} \lesssim \sqrt{n}$ and Claim C.1 we know $\|\mathbf{u}(\tau) - \mathbf{y}\| \lesssim \sqrt{n}$ for all $\tau < t$.

Step 1: proving $\|\boldsymbol{\theta}(t) - \boldsymbol{\theta}(0)\| \leq R$. We define

$$\mathbf{J}(\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}') := \int_0^1 \mathbf{J}(\boldsymbol{\theta} + x(\boldsymbol{\theta}' - \boldsymbol{\theta}))dx.$$

We first prove $\|\boldsymbol{\theta}(t-1) - \boldsymbol{\theta}(0)\| \leq \frac{R}{2}$. If $t = 1$, this is trivially true. Now we assume $t \geq 2$. For each $0 \leq \tau < t-1$, by the fundamental theorem for line integrals we have

$$\begin{aligned} \mathbf{u}(\tau+1) - \mathbf{u}(\tau) &= \mathbf{J}(\boldsymbol{\theta}(\tau) \rightarrow \boldsymbol{\theta}(\tau+1)) \cdot (\boldsymbol{\theta}(\tau+1) - \boldsymbol{\theta}(\tau)) \\ &= -\frac{\eta}{n} \mathbf{J}(\boldsymbol{\theta}(\tau) \rightarrow \boldsymbol{\theta}(\tau+1)) \mathbf{J}(\boldsymbol{\theta}(\tau))^\top (\mathbf{u}(\tau) - \mathbf{y}). \end{aligned}$$

Let $\mathbf{E}(\tau) := \mathbf{J}(\boldsymbol{\theta}(\tau) \rightarrow \boldsymbol{\theta}(\tau+1)) \mathbf{J}(\boldsymbol{\theta}(\tau))^\top - \mathbf{K}$. Since $\|\boldsymbol{\theta}(\tau) - \boldsymbol{\theta}(0)\| \leq R$ and $\|\boldsymbol{\theta}(\tau+1) - \boldsymbol{\theta}(0)\| \leq R$, from Assumption C.1 we know that $\|\mathbf{E}(\tau)\| \leq \epsilon$. We can write

$$\begin{aligned} \mathbf{u}(\tau+1) - \mathbf{y} &= \left(\mathbf{I} - \frac{\eta}{n} \mathbf{J}(\boldsymbol{\theta}(\tau) \rightarrow \boldsymbol{\theta}(\tau+1)) \mathbf{J}(\boldsymbol{\theta}(\tau))^\top \right) (\mathbf{u}(\tau) - \mathbf{y}) \\ &= \left(\mathbf{I} - \frac{\eta}{n} \mathbf{K} \right) (\mathbf{u}(\tau) - \mathbf{y}) - \frac{\eta}{n} \mathbf{E}(\tau) (\mathbf{u}(\tau) - \mathbf{y}). \end{aligned} \tag{13}$$

It follows that

$$\begin{aligned}
& \|\mathbf{u}(\tau+1) - \mathbf{y}\|^2 \\
& \leq \left\| \left(\mathbf{I} - \frac{\eta}{n} \mathbf{K} \right) (\mathbf{u}(\tau) - \mathbf{y}) \right\|^2 + 2 \left\| \left(\mathbf{I} - \frac{\eta}{n} \mathbf{K} \right) (\mathbf{u}(\tau) - \mathbf{y}) \right\| \cdot \left\| \frac{\eta}{n} \mathbf{E}(\tau) (\mathbf{u}(\tau) - \mathbf{y}) \right\| \\
& \quad + \left\| \frac{\eta}{n} \mathbf{E}(\tau) (\mathbf{u}(\tau) - \mathbf{y}) \right\|^2 \\
& \leq \left\| \left(\mathbf{I} - \frac{\eta}{n} \mathbf{K} \right) (\mathbf{u}(\tau) - \mathbf{y}) \right\|^2 + O \left(\sqrt{n} \cdot \frac{\eta}{n} \epsilon \sqrt{n} + \left(\frac{\eta}{n} \epsilon \sqrt{n} \right)^2 \right) \\
& = \left\| \left(\mathbf{I} - \frac{\eta}{n} \mathbf{K} \right) (\mathbf{u}(\tau) - \mathbf{y}) \right\|^2 + O(\eta \epsilon) \tag{$\eta \epsilon \lesssim n$} \\
& = \|\mathbf{u}(\tau) - \mathbf{y}\|^2 - \frac{2\eta}{n} (\mathbf{u}(\tau) - \mathbf{y})^\top \mathbf{K} (\mathbf{u}(\tau) - \mathbf{y}) + \frac{\eta^2}{n^2} \|\mathbf{K} (\mathbf{u}(\tau) - \mathbf{y})\|^2 + O(\eta \epsilon) \\
& \leq \|\mathbf{u}(\tau) - \mathbf{y}\|^2 - \frac{2\eta}{n} (\mathbf{u}(\tau) - \mathbf{y})^\top \mathbf{K} (\mathbf{u}(\tau) - \mathbf{y}) + \frac{\eta^2}{n^2} \|\mathbf{K}\| \cdot \left\| \mathbf{K}^{1/2} (\mathbf{u}(\tau) - \mathbf{y}) \right\|^2 + O(\eta \epsilon) \\
& \leq \|\mathbf{u}(\tau) - \mathbf{y}\|^2 - \frac{\eta}{n} (\mathbf{u}(\tau) - \mathbf{y})^\top \mathbf{K} (\mathbf{u}(\tau) - \mathbf{y}) + O(\eta \epsilon). \tag{$\frac{\eta^2 \|\mathbf{K}\|}{n^2} \leq \frac{\eta}{n}$}
\end{aligned}$$

On the other hand, we have

$$\begin{aligned}
& \|\boldsymbol{\theta}(\tau+1) - \boldsymbol{\theta}(\tau)\|^2 \\
& = \frac{\eta^2}{n^2} \|\mathbf{J}(\boldsymbol{\theta}(\tau))^\top (\mathbf{u}(\tau) - \mathbf{y})\|^2 \\
& = \frac{\eta^2}{n^2} (\mathbf{u}(\tau) - \mathbf{y})^\top \mathbf{J}(\boldsymbol{\theta}(\tau)) \mathbf{J}(\boldsymbol{\theta}(\tau))^\top (\mathbf{u}(\tau) - \mathbf{y}) \\
& \leq \frac{\eta^2}{n^2} \left((\mathbf{u}(\tau) - \mathbf{y})^\top \mathbf{K} (\mathbf{u}(\tau) - \mathbf{y}) + \|\mathbf{u}(\tau) - \mathbf{y}\|^2 \|\mathbf{J}(\boldsymbol{\theta}(\tau)) \mathbf{J}(\boldsymbol{\theta}(\tau))^\top - \mathbf{K}\| \right) \\
& \leq \frac{\eta^2}{n^2} \left((\mathbf{u}(\tau) - \mathbf{y})^\top \mathbf{K} (\mathbf{u}(\tau) - \mathbf{y}) + O(n\epsilon) \right). \tag{14}
\end{aligned}$$

Combining the above two inequalities, we obtain

$$\begin{aligned}
& \|\mathbf{u}(\tau+1) - \mathbf{y}\|^2 - \|\mathbf{u}(\tau) - \mathbf{y}\|^2 \\
& \leq -\frac{\eta}{n} \left(\frac{n^2}{\eta^2} \|\boldsymbol{\theta}(\tau+1) - \boldsymbol{\theta}(\tau)\|^2 - O(n\epsilon) \right) + O(\eta \epsilon) \\
& = -\frac{n}{\eta} \|\boldsymbol{\theta}(\tau+1) - \boldsymbol{\theta}(\tau)\|^2 + O(\eta \epsilon).
\end{aligned}$$

Taking sum over $\tau = 0, \dots, t-2$, we get

$$\|\mathbf{u}(t-1) - \mathbf{y}\|^2 - \|\mathbf{u}(0) - \mathbf{y}\|^2 \leq -\frac{n}{\eta} \sum_{\tau=0}^{t-2} \|\boldsymbol{\theta}(\tau+1) - \boldsymbol{\theta}(\tau)\|^2 + O(\eta t \epsilon),$$

which implies

$$\frac{n}{\eta} \sum_{\tau=0}^{t-2} \|\boldsymbol{\theta}(\tau+1) - \boldsymbol{\theta}(\tau)\|^2 \leq \|\mathbf{y}\|^2 + O(\eta t \epsilon) \leq \|\mathbf{y}\|^2 + O(R^2 \epsilon) = O(n).$$

Then by the Cauchy-Schwartz inequality we have

$$\begin{aligned}
\|\boldsymbol{\theta}(t-1) - \boldsymbol{\theta}(0)\| & \leq \sum_{\tau=0}^{t-2} \|\boldsymbol{\theta}(\tau+1) - \boldsymbol{\theta}(\tau)\| \leq \sqrt{(t-1) \sum_{\tau=0}^{t-2} \|\boldsymbol{\theta}(\tau+1) - \boldsymbol{\theta}(\tau)\|^2} \\
& \leq \sqrt{t \cdot O(\eta)} \leq \sqrt{c \frac{R^2}{\eta} \cdot O(\eta)}.
\end{aligned}$$

Choosing c sufficiently small, we can ensure $\|\boldsymbol{\theta}(t-1) - \boldsymbol{\theta}(0)\| \leq \frac{R}{2}$.

Now that we have proved $\|\boldsymbol{\theta}(t-1) - \boldsymbol{\theta}(0)\| \leq \frac{R}{2}$, to prove $\|\boldsymbol{\theta}(t) - \boldsymbol{\theta}(0)\| \leq R$ it suffices to bound the one-step deviation $\|\boldsymbol{\theta}(t) - \boldsymbol{\theta}(t-1)\|$ by $\frac{R}{2}$. Using the exact same method in (14), we have

$$\|\boldsymbol{\theta}(t) - \boldsymbol{\theta}(t-1)\| \leq \frac{\eta}{n} \sqrt{n\|\mathbf{K}\| + O(n\epsilon)} \lesssim \eta \sqrt{\|\mathbf{K}\|/n} = \sqrt{\eta\|\mathbf{K}\|/n} \sqrt{\eta} \leq \sqrt{c}R,$$

where we have used $\eta \leq \frac{n}{\|\mathbf{K}\|}$ and $\eta \leq \eta t \leq cR^2$. Choosing c sufficiently small, we can ensure $\|\boldsymbol{\theta}(t) - \boldsymbol{\theta}(t-1)\| \leq \frac{R}{2}$. Therefore we conclude that $\|\boldsymbol{\theta}(t) - \boldsymbol{\theta}(0)\| \leq R$.

Step 2: proving $\|\mathbf{u}(t) - \mathbf{u}^{\text{lin}}(t)\| \lesssim \frac{\eta t \epsilon}{\sqrt{n}}$. Same as (13) we have

$$\mathbf{u}(t) - \mathbf{y} = \left(\mathbf{I} - \frac{\eta}{n}\mathbf{K}\right) (\mathbf{u}(t-1) - \mathbf{y}) - \frac{\eta}{n}\mathbf{E}(t-1)(\mathbf{u}(t-1) - \mathbf{y}),$$

where $\mathbf{E}(t-1) = \mathbf{J}(\boldsymbol{\theta}(t-1), \boldsymbol{\theta}(t))\mathbf{J}(\boldsymbol{\theta}(t-1))^\top - \mathbf{K}$. Since $\|\boldsymbol{\theta}(t-1) - \boldsymbol{\theta}(0)\| \leq R$ and $\|\boldsymbol{\theta}(t) - \boldsymbol{\theta}(0)\| \leq R$, we know from Assumption C.1 that $\|\mathbf{E}(t-1)\| \leq \epsilon$. Moreover, from Claim C.1 we know

$$\mathbf{u}^{\text{lin}}(t) - \mathbf{y} = \left(\mathbf{I} - \frac{\eta}{n}\mathbf{K}\right) (\mathbf{u}^{\text{lin}}(t-1) - \mathbf{y}).$$

It follows that

$$\mathbf{u}(t) - \mathbf{u}^{\text{lin}}(t) = \left(\mathbf{I} - \frac{\eta}{n}\mathbf{K}\right) (\mathbf{u}(t-1) - \mathbf{u}^{\text{lin}}(t-1)) - \frac{\eta}{n}\mathbf{E}(t-1)(\mathbf{u}(t-1) - \mathbf{y}),$$

which implies

$$\begin{aligned} \|\mathbf{u}(t) - \mathbf{u}^{\text{lin}}(t)\| &\leq \left\| \left(\mathbf{I} - \frac{\eta}{n}\mathbf{K}\right) (\mathbf{u}(t-1) - \mathbf{u}^{\text{lin}}(t-1)) \right\| + \left\| \frac{\eta}{n}\mathbf{E}(t-1)(\mathbf{u}(t-1) - \mathbf{y}) \right\| \\ &\leq \|\mathbf{u}(t-1) - \mathbf{u}^{\text{lin}}(t-1)\| + O\left(\frac{\eta}{n}\epsilon\sqrt{n}\right) \\ &= \|\mathbf{u}(t-1) - \mathbf{u}^{\text{lin}}(t-1)\| + O\left(\frac{\eta\epsilon}{\sqrt{n}}\right). \end{aligned}$$

Therefore from $\|\mathbf{u}(t-1) - \mathbf{u}^{\text{lin}}(t-1)\| \lesssim \frac{\eta(t-1)\epsilon}{\sqrt{n}}$ we know $\|\mathbf{u}(t) - \mathbf{u}^{\text{lin}}(t)\| \lesssim \frac{\eta t \epsilon}{\sqrt{n}}$, completing the proof.

Finally, we prove the last statement in the theorem, i.e., $\|\boldsymbol{\omega}(t) - \boldsymbol{\omega}(0)\| \leq R$. In fact we have already proved this – notice that we have proved $\|\boldsymbol{\theta}(t) - \boldsymbol{\theta}(0)\| \leq R$ and that a special instance of this problem is when $\boldsymbol{\theta}(t) = \boldsymbol{\omega}(t)$, i.e., the two dynamics are the same. Applying our result on that problem instance, we obtain $\|\boldsymbol{\omega}(t) - \boldsymbol{\omega}(0)\| \leq R$. \square

D Omitted Details in Section 3

In Appendix D.1, we present the formal theoretical guarantee (Theorem D.1) for the case of training both layers.

In Appendix D.2, we calculate the formulae of various Jacobians and NTKs that will be used in the analysis.

In Appendix D.3, we prove Theorem 3.2 (training the first layer).

In Appendix D.4, we prove Corollary 3.3 (training the first layer with well-conditioned data).

In Appendix D.5, we prove Theorem 3.5 (training the second layer).

In Appendix D.6, we prove Theorem D.1 (training both layers).

In Appendix D.7, we prove Claim 3.1 (data concentration properties).

D.1 Guarantee for Training Both Layers

Now we state our guarantee for the case of training both layers, continuing from Section 3.3. Recall that the neural network weights $(\mathbf{W}(t), \mathbf{v}(t))$ are updated according to GD (4) with learning rate $\eta_1 = \eta_2 = \eta$. The linear model $f^{\text{lin}}(\mathbf{x}; \boldsymbol{\delta})$ in (11) is also trained with GD:

$$\boldsymbol{\delta}(0) = \mathbf{0}_{d+2}, \quad \boldsymbol{\delta}(t+1) = \boldsymbol{\delta}(t) - \eta \nabla_{\boldsymbol{\delta}} \frac{1}{2n} \sum_{i=1}^n (f^{\text{lin}}(\mathbf{x}_i; \boldsymbol{\delta}(t)) - y_i)^2.$$

We let f_t and f_t^{lin} be the neural network and the linear model at iteration t , i.e., $f_t(\mathbf{x}) := f(\mathbf{x}; \mathbf{W}(t), \mathbf{v}(t))$ and $f_t^{\text{lin}}(\mathbf{x}) := f^{\text{lin}}(\mathbf{x}; \boldsymbol{\delta}(t))$.

Theorem D.1 (main theorem for training both layers). *Let $\alpha \in (0, \frac{1}{4})$ be a fixed constant. Suppose $n \gtrsim d^{1+\alpha}$ and $m \gtrsim d^{2+\alpha}$. Suppose $\begin{cases} \eta \ll d/\log n, & \text{if } \mathbb{E}[\phi(g)] = 0 \\ \eta \ll 1, & \text{otherwise} \end{cases}$. Then there exists a universal constant $c > 0$ such that with high probability, for all $0 \leq t \leq T = c \cdot \frac{d \log d}{\eta}$ simultaneously, we have*

$$\frac{1}{n} \sum_{i=1}^n (f_t(\mathbf{x}_i) - f_t^{\text{lin}}(\mathbf{x}_i))^2 \lesssim d^{-\Omega(\alpha)}, \quad \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\min\{(f_t(\mathbf{x}) - f_t^{\text{lin}}(\mathbf{x}))^2, 1\}] \lesssim d^{-\Omega(\alpha)} + \sqrt{\frac{\log T}{n}}.$$

We remark that if the data distribution is well-conditioned, we can also have a guarantee similar to Corollary 3.3.

D.2 Formulae of Jacobians and NTKs

We first calculate the Jacobian of the network outputs at the training data \mathbf{X} with respect to the weights in the network. The Jacobian for the first layer is:

$$\mathbf{J}_1(\mathbf{W}, \mathbf{v}) := [\mathbf{J}_1(\mathbf{w}_1, v_1), \mathbf{J}_1(\mathbf{w}_2, v_2), \dots, \mathbf{J}_1(\mathbf{w}_m, v_m)] \in \mathbb{R}^{n \times md}, \quad (15)$$

where

$$\mathbf{J}_1(\mathbf{w}_r, v_r) := \frac{1}{\sqrt{md}} v_r \text{diag} \left(\phi'(\mathbf{X} \mathbf{w}_r / \sqrt{d}) \right) \mathbf{X} \in \mathbb{R}^{n \times d}, \quad r \in [m].$$

The Jacobian for the second layer is:

$$\mathbf{J}_2(\mathbf{W}) := \frac{1}{\sqrt{m}} \phi(\mathbf{X} \mathbf{W}^\top / \sqrt{d}) \in \mathbb{R}^{n \times m}. \quad (16)$$

Here we omit \mathbf{v} in the notation since it does not affect the Jacobian. The Jacobian for both layers is simply $\mathbf{J}(\mathbf{W}, \mathbf{v}) := [\mathbf{J}_1(\mathbf{W}, \mathbf{v}), \mathbf{J}_2(\mathbf{W})] \in \mathbb{R}^{n \times (md+m)}$.

After calculating the Jacobians, we can calculate the NTK matrices for the first layer, the second layer, and both layers as follows:

$$\begin{aligned} \Theta_1(\mathbf{W}, \mathbf{v}) &:= \mathbf{J}_1(\mathbf{W}, \mathbf{v}) \mathbf{J}_1(\mathbf{W}, \mathbf{v})^\top = \frac{1}{m} \sum_{r=1}^m v_r^2 \left(\phi'(\mathbf{X} \mathbf{w}_r / \sqrt{d}) \phi'(\mathbf{X} \mathbf{w}_r / \sqrt{d})^\top \right) \odot \frac{\mathbf{X} \mathbf{X}^\top}{d}, \\ \Theta_2(\mathbf{W}) &:= \mathbf{J}_2(\mathbf{W}) \mathbf{J}_2(\mathbf{W})^\top = \frac{1}{m} \phi(\mathbf{X} \mathbf{W}^\top / \sqrt{d}) \phi(\mathbf{X} \mathbf{W}^\top / \sqrt{d})^\top, \\ \Theta(\mathbf{W}, \mathbf{v}) &:= \mathbf{J}(\mathbf{W}, \mathbf{v}) \mathbf{J}(\mathbf{W}, \mathbf{v})^\top = \Theta_1(\mathbf{W}, \mathbf{v}) + \Theta_2(\mathbf{W}). \end{aligned} \quad (17)$$

We also denote the expected NTK matrices at random initialization as:

$$\begin{aligned} \Theta_1^* &:= \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), v \sim \text{Unif}\{\pm 1\}} \left[v^2 \left(\phi'(\mathbf{X} \mathbf{w} / \sqrt{d}) \phi'(\mathbf{X} \mathbf{w} / \sqrt{d})^\top \right) \right] \odot \frac{\mathbf{X} \mathbf{X}^\top}{d} \\ &= \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left(\phi'(\mathbf{X} \mathbf{w} / \sqrt{d}) \phi'(\mathbf{X} \mathbf{w} / \sqrt{d})^\top \right) \right] \odot \frac{\mathbf{X} \mathbf{X}^\top}{d}, \\ \Theta_2^* &:= \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\phi(\mathbf{X} \mathbf{w} / \sqrt{d}) \phi(\mathbf{X} \mathbf{w} / \sqrt{d})^\top \right], \\ \Theta^* &:= \Theta_1^* + \Theta_2^*. \end{aligned} \quad (18)$$

These are also the NTK matrices at infinite width ($m \rightarrow \infty$).

Next, for the three linear models (5), (9) and (11) defined in Section 3, denote their feature/Jacobian matrices by:

$$\begin{aligned} \Psi_1 &:= [\psi_1(\mathbf{x}_1), \dots, \psi_1(\mathbf{x}_n)]^\top, \\ \Psi_2 &:= [\psi_2(\mathbf{x}_1), \dots, \psi_2(\mathbf{x}_n)]^\top, \\ \Psi &:= [\psi(\mathbf{x}_1), \dots, \psi(\mathbf{x}_n)]^\top. \end{aligned} \quad (19)$$

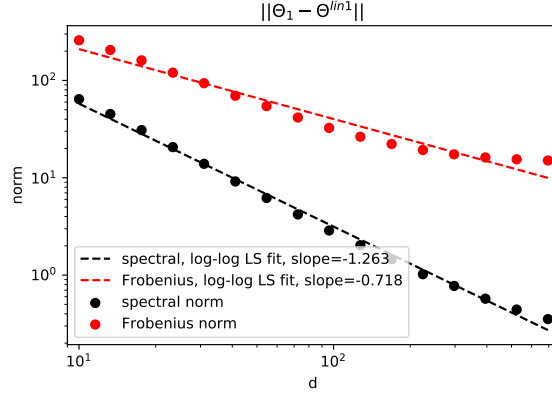


Figure 5: **Verification of Proposition 3.4/D.2.** We simulate the dependence of the spectral and Frobenius norms of $\Theta_1(\mathbf{W}(0)) - \Theta^{\text{lin1}}$ on d . We set $\phi = \text{erf}$, $n = 10^4$ and $m = 2 \times 10^4$, and generate data from $\mathcal{N}(\mathbf{0}, \mathbf{I})$ for various d . We perform a linear least-squares fit on the log mean norms against $\log(d)$. Numerically we find $\|\Theta_1(\mathbf{W}(0)) - \Theta^{\text{lin1}}\| \propto d^{-1.263}$ and $\|\Theta_1(\mathbf{W}(0)) - \Theta^{\text{lin1}}\|_F \propto d^{-0.718}$.

Consequently, their corresponding kernel matrices are:

$$\begin{aligned}
 \Theta^{\text{lin1}} &:= \Psi_1 \Psi_1^\top = \frac{1}{d} (\zeta^2 \mathbf{X} \mathbf{X}^\top + \nu^2 \mathbf{1} \mathbf{1}^\top), \\
 \Theta^{\text{lin2}} &:= \Psi_2 \Psi_2^\top = \frac{1}{d} \left(\zeta^2 \mathbf{X} \mathbf{X}^\top + \frac{1}{2} \nu^2 \mathbf{1} \mathbf{1}^\top \right) + \mathbf{q} \mathbf{q}^\top, \\
 \Theta^{\text{lin}} &:= \Psi \Psi^\top = \frac{1}{d} \left(2\zeta^2 \mathbf{X} \mathbf{X}^\top + \frac{3}{2} \nu^2 \mathbf{1} \mathbf{1}^\top \right) + \mathbf{q} \mathbf{q}^\top.
 \end{aligned} \tag{20}$$

Here the constants are defined in (9), and $\mathbf{q} \in \mathbb{R}^n$ is defined as $[\mathbf{q}]_i := \vartheta_0 + \vartheta_1 \left(\frac{\|\mathbf{x}_i\|}{\sqrt{d}} - 1 \right) + \vartheta_2 \left(\frac{\|\mathbf{x}_i\|}{\sqrt{d}} - 1 \right)^2$ for each $i \in [n]$.

D.3 Proof of Theorem 3.2 (Training the First Layer)

For convenience we let $\mathbf{v} = \mathbf{v}(0)$ which is the fixed second layer. Since we have $v_r \in \{\pm 1\}$ ($\forall r \in [m]$), we can write the first-layer NTK matrix as

$$\Theta_1(\mathbf{W}, \mathbf{v}) = \frac{1}{m} \sum_{r=1}^m \left(\phi'(\mathbf{X} \mathbf{w}_r / \sqrt{d}) \phi'(\mathbf{X} \mathbf{w}_r / \sqrt{d})^\top \right) \odot \frac{\mathbf{X} \mathbf{X}^\top}{d}.$$

Because it does not depend on \mathbf{v} , we denote $\Theta_1(\mathbf{W}) := \Theta_1(\mathbf{W}, \mathbf{v})$ for convenience.

D.3.1 The NTK at Initialization

Now we prove Proposition 3.4, restated below:

Proposition D.2 (restatement of Proposition 3.4). *With high probability over the random initialization $\mathbf{W}(0)$ and the training data \mathbf{X} , we have*

$$\|\Theta_1(\mathbf{W}(0)) - \Theta^{\text{lin1}}\| \lesssim \frac{n}{d^{1+\alpha}}.$$

We perform a simulation to empirically verify Proposition D.2 in Figure 5. Here we fix n and m to be large and look at the dependence of $\|\Theta_1(\mathbf{W}(0)) - \Theta^{\text{lin1}}\|$ on d . We find that $\|\Theta_1(\mathbf{W}(0)) - \Theta^{\text{lin1}}\|$ indeed decays faster than $\frac{1}{d}$. In contrast, $\|\Theta_1(\mathbf{W}(0)) - \Theta^{\text{lin1}}\|_F$ decays slower than $\frac{1}{d}$, indicating that bounding the Frobenius norm is insufficient.

To prove Proposition D.2, we will prove $\Theta_1(\mathbf{W}(0))$ is close to its expectation Θ_1^* (defined in (18)), and then prove Θ_1^* is close to Θ^{lin1} . We do these steps in the next two propositions.

Proposition D.3. *With high probability over the random initialization $\mathbf{W}(0)$ and the training data \mathbf{X} , we have*

$$\|\Theta_1(\mathbf{W}(0)) - \Theta_1^*\| \leq \frac{n}{d^{1+\alpha}}.$$

Proof. For convenience we denote $\mathbf{W} = \mathbf{W}(0)$ and $\Theta_1 = \Theta_1(\mathbf{W}) = \Theta_1(\mathbf{W}(0))$ in this proof.

From Claim 3.1 we know $\|\mathbf{X}\mathbf{X}^\top\| = O(n)$ with high probability. For the rest of the proof we will be conditioned on \mathbf{X} and on Claim 3.1, and only consider the randomness in \mathbf{W} .

We define $\Theta_1^{(r)} := \left(\phi'(\mathbf{X}\mathbf{w}_r/\sqrt{d})\phi'(\mathbf{X}\mathbf{w}_r/\sqrt{d})^\top\right) \odot \frac{\mathbf{X}\mathbf{X}^\top}{d}$ for each $r \in [m]$. Then we have $\Theta_1 = \frac{1}{m} \sum_{r=1}^m \Theta_1^{(r)}$. According to the initialization scheme (3), we know that $\Theta_1^{(1)}, \Theta_1^{(2)}, \dots, \Theta_1^{(m/2)}$ are independent, $\Theta_1^{(m/2+1)}, \Theta_1^{(m/2+2)}, \dots, \Theta_1^{(m)}$ are independent, and $\mathbb{E}[\Theta_1^{(r)}] = \Theta_1^*$ for all $r \in [m]$.

Next we will apply the matrix Bernstein inequality (Theorem 1.6.2 in Tropp [2015]) to bound $\|\Theta_1 - \Theta_1^*\|$. We will first consider the first half of independent neurons, i.e. $r \in [m/2]$. For each r we have

$$\begin{aligned} \|\Theta_1^{(r)}\| &= \left\| \text{diag}\left(\phi'(\mathbf{X}\mathbf{w}_r/\sqrt{d})\right) \cdot \frac{\mathbf{X}\mathbf{X}^\top}{d} \cdot \text{diag}\left(\phi'(\mathbf{X}\mathbf{w}_r/\sqrt{d})\right) \right\| \\ &\leq \left\| \text{diag}\left(\phi'(\mathbf{X}\mathbf{w}_r/\sqrt{d})\right) \right\| \cdot \left\| \frac{\mathbf{X}\mathbf{X}^\top}{d} \right\| \cdot \left\| \text{diag}\left(\phi'(\mathbf{X}\mathbf{w}_r/\sqrt{d})\right) \right\| \\ &\leq O(1) \cdot O(n/d) \cdot O(1) \\ &= O(n/d). \end{aligned}$$

Here we have used the boundedness of $\phi'(\cdot)$ (Assumption 3.2). Since $\Theta_1^* = \mathbb{E}[\Theta_1^{(r)}]$, it follows that

$$\begin{aligned} \|\Theta_1^*\| &\leq O(n/d), \\ \|\Theta_1^{(r)} - \Theta_1^*\| &\leq O(n/d), \quad \forall r \in [m/2] \\ \left\| \sum_{r=1}^{m/2} \mathbb{E}[(\Theta_1^{(r)} - \Theta_1^*)^2] \right\| &\leq \sum_{r=1}^{m/2} \left\| \mathbb{E}[(\Theta_1^{(r)} - \Theta_1^*)^2] \right\| \leq O(mn^2/d^2). \end{aligned}$$

Therefore, from the the matrix Bernstein inequality, for any $s \geq 0$ we have:

$$\Pr \left[\left\| \sum_{r=1}^{m/2} (\Theta_1^{(r)} - \Theta_1^*) \right\| \geq s \right] \leq 2n \cdot \exp \left(\frac{-s^2/2}{O(mn^2/d^2) + sn/d} \right).$$

Letting $s = \frac{m}{2} \cdot \frac{n}{d^{1+\alpha}}$, we obtain

$$\begin{aligned} \Pr \left[\left\| \sum_{r=1}^{m/2} (\Theta_1^{(r)} - \Theta_1^*) \right\| \geq \frac{m}{2} \cdot \frac{n}{d^{1+\alpha}} \right] &\leq 2n \cdot \exp \left(-\Omega \left(\frac{m^2 n^2 / d^{2+2\alpha}}{mn^2/d^2 + mn^2/d^{2+\alpha}} \right) \right) \\ &= 2n \cdot \exp \left(-\Omega \left(\frac{m}{d^{2\alpha}} \right) \right) \\ &= d^{O(1)} \cdot e^{-\Omega(d^{1-\alpha})} \\ &\ll 1, \end{aligned}$$

where we have used $m = \Omega(d^{1+\alpha})$ and $n = d^{O(1)}$. Therefore with high probability we have

$$\left\| \sum_{r=1}^{m/2} (\Theta_1^{(r)} - \Theta_1^*) \right\| \leq \frac{m}{2} \cdot \frac{n}{d^{1+\alpha}}.$$

Similarly, for the second half of the neurons we also have with high probability

$$\left\| \sum_{r=m/2+1}^m (\Theta_1^{(r)} - \Theta_1^*) \right\| \leq \frac{m}{2} \cdot \frac{n}{d^{1+\alpha}}.$$

Finally, by the triangle inequality we have

$$\|\Theta_1 - \Theta_1^*\| = \frac{1}{m} \left\| \sum_{r=1}^m (\Theta_1^{(r)} - \Theta_1^*) \right\| \leq \frac{1}{m} \left(\frac{m}{2} \cdot \frac{n}{d^{1+\alpha}} + \frac{m}{2} \cdot \frac{n}{d^{1+\alpha}} \right) = \frac{n}{d^{1+\alpha}}$$

with high probability, completing the proof. \square

Proposition D.4. *With high probability over the training data \mathbf{X} , we have*

$$\|\Theta_1^* - \Theta^{\text{lin1}}\| \lesssim \frac{n}{d^{1+\alpha}}.$$

Proof. We will be conditioned on the high probability events stated in Claim 3.1.

By the definition of Θ_1^* , we know

$$[\Theta_1^*]_{i,j} = \frac{1}{d} \mathbf{x}_i^\top \mathbf{x}_j \cdot \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\phi'(\mathbf{w}^\top \mathbf{x}_i / \sqrt{d}) \phi'(\mathbf{w}^\top \mathbf{x}_j / \sqrt{d}) \right], \quad i, j \in [n].$$

We define

$$\Phi(a, b, c) := \mathbb{E}_{(z_1, z_2) \sim \mathcal{N}(\mathbf{0}, \mathbf{\Lambda})} [\phi'(z_1) \phi'(z_2)], \quad \text{where } \mathbf{\Lambda} = \begin{pmatrix} a & c \\ c & b \end{pmatrix}, \quad a \geq 0, b \geq 0, |c| \leq \sqrt{ab}.$$

Then we can write

$$[\Theta_1^*]_{i,j} = \frac{1}{d} \mathbf{x}_i^\top \mathbf{x}_j \cdot \Phi \left(\frac{\|\mathbf{x}_i\|^2}{d}, \frac{\|\mathbf{x}_j\|^2}{d}, \frac{\mathbf{x}_i^\top \mathbf{x}_j}{d} \right).$$

We consider the diagonal and off-diagonal entries of Θ_1^* separately.

For $i \neq j$, from Claim 3.1 we know $\frac{\|\mathbf{x}_i\|^2}{d} = 1 \pm \tilde{O}(\frac{1}{\sqrt{d}})$, $\frac{\|\mathbf{x}_j\|^2}{d} = 1 \pm \tilde{O}(\frac{1}{\sqrt{d}})$ and $\frac{\mathbf{x}_i^\top \mathbf{x}_j}{d} = \pm \tilde{O}(\frac{1}{\sqrt{d}})$. Hence we apply Taylor expansion of Φ around $(1, 1, 0)$:

$$\begin{aligned} & \Phi \left(\frac{\|\mathbf{x}_i\|^2}{d}, \frac{\|\mathbf{x}_j\|^2}{d}, \frac{\mathbf{x}_i^\top \mathbf{x}_j}{d} \right) \\ &= \Phi(1, 1, 0) + c_1 \left(\frac{\|\mathbf{x}_i\|^2}{d} - 1 \right) + c_2 \left(\frac{\|\mathbf{x}_j\|^2}{d} - 1 \right) + c_3 \frac{(\mathbf{x}_i^\top \mathbf{x}_j)^2}{d} \\ & \quad \pm O \left(\left(\frac{\|\mathbf{x}_i\|^2}{d} - 1 \right)^2 + \left(\frac{\|\mathbf{x}_j\|^2}{d} - 1 \right)^2 + \left(\frac{(\mathbf{x}_i^\top \mathbf{x}_j)^2}{d} \right)^2 \right) \\ &= \Phi(1, 1, 0) + c_1 \left(\frac{\|\mathbf{x}_i\|^2}{d} - 1 \right) + c_2 \left(\frac{\|\mathbf{x}_j\|^2}{d} - 1 \right) + c_3 \frac{(\mathbf{x}_i^\top \mathbf{x}_j)^2}{d} \pm \tilde{O} \left(\frac{1}{d} \right). \end{aligned}$$

Here $(c_1, c_2, c_3) := \nabla \Phi(1, 1, 0)$. Note that $\Phi(1, 1, 0)$ and all first and second order derivatives of Φ at $(1, 1, 0)$ exist and are bounded for activation ϕ that satisfies Assumption 3.2. In particular, we have $\Phi(1, 1, 0) = (\mathbb{E}[\phi'(g)])^2 = \zeta^2$ and $c_3 = (\mathbb{E}[g\phi'(g)])^2$. Using the above expansion, we can write

$$\begin{aligned} (\Theta_1^*)_{\text{off}} &= \zeta^2 \left(\frac{\mathbf{X}\mathbf{X}^\top}{d} \right)_{\text{off}} + c_1 \left(\text{diag}(\boldsymbol{\epsilon}) \cdot \frac{\mathbf{X}\mathbf{X}^\top}{d} \right)_{\text{off}} + c_2 \left(\frac{\mathbf{X}\mathbf{X}^\top}{d} \cdot \text{diag}(\boldsymbol{\epsilon}) \right)_{\text{off}} \\ & \quad + c_3 \left(\frac{\mathbf{X}\mathbf{X}^\top}{d} \odot \frac{\mathbf{X}\mathbf{X}^\top}{d} \right)_{\text{off}} + \mathbf{E}, \end{aligned} \quad (21)$$

where $\boldsymbol{\epsilon} \in \mathbb{R}^n$ is defined as $[\boldsymbol{\epsilon}]_i = \frac{\|\mathbf{x}_i\|^2}{d} - 1$, and $[\mathbf{E}]_{i,j} = \pm \tilde{O}(\frac{1}{d}) \cdot \frac{\mathbf{x}_i^\top \mathbf{x}_j}{d} \mathbb{1}_{\{i \neq j\}} = \pm \tilde{O}(\frac{1}{d^{1.5}})$.

Now we treat the terms in (21) separately. First, we have

$$\left\| \text{diag}(\boldsymbol{\epsilon}) \cdot \frac{\mathbf{X}\mathbf{X}^\top}{d} \right\| \leq \|\text{diag}(\boldsymbol{\epsilon})\| \cdot \left\| \frac{\mathbf{X}\mathbf{X}^\top}{d} \right\| = \max_{i \in [n]} |[\boldsymbol{\epsilon}]_i| \cdot \left\| \frac{\mathbf{X}\mathbf{X}^\top}{d} \right\| \leq \tilde{O} \left(\frac{1}{\sqrt{d}} \right) \cdot O \left(\frac{n}{d} \right)$$

$$= \tilde{O}\left(\frac{n}{d^{1.5}}\right).$$

Similarly, we have $\left\|\frac{\mathbf{X}\mathbf{X}^\top}{d} \cdot \text{diag}(\epsilon)\right\| \leq \tilde{O}\left(\frac{n}{d^{1.5}}\right)$.

Next, for $\left(\frac{\mathbf{X}\mathbf{X}^\top}{d} \odot \frac{\mathbf{X}\mathbf{X}^\top}{d}\right)_{\text{off}}$, we can use the 4th moment method in [El Karoui \[2010\]](#) to show that it is close to its mean. Specifically, the mean at each entry is $\mathbb{E}\left[\left(\frac{\mathbf{x}_i^\top \mathbf{x}_j}{d}\right)^2\right] = \frac{\text{Tr}[\boldsymbol{\Sigma}^2]}{d^2}$ ($i \neq j$), and the moment calculation in [El Karoui \[2010\]](#) shows the following bound on the error matrix $\mathbf{F} = \left(\frac{\mathbf{X}\mathbf{X}^\top}{d} \odot \frac{\mathbf{X}\mathbf{X}^\top}{d} - \frac{\text{Tr}[\boldsymbol{\Sigma}^2]}{d^2} \mathbf{1}\mathbf{1}^\top\right)_{\text{off}}$:

$$\mathbb{E}\left[\|\mathbf{F}\|^4\right] \leq \mathbb{E}\left[\text{Tr}[\mathbf{F}^4]\right] \leq \tilde{O}\left(\frac{n^4}{d^6} + \frac{n^3}{d^4}\right) \leq \tilde{O}\left(\frac{n^4}{d^5}\right),$$

where we have used $n \gtrsim d$. Therefore by Markov inequality we know that with high probability, $\|\mathbf{F}\| \leq \tilde{O}\left(\frac{n}{d^{1.25}}\right)$.

For the final term \mathbf{E} in (21), we have

$$\|\mathbf{E}\| \leq \|\mathbf{E}\|_F \leq \sqrt{n^2 \cdot \tilde{O}\left(\frac{1}{d^3}\right)} = \tilde{O}\left(\frac{n}{d^{1.5}}\right).$$

Put together, we can obtain the following bound regarding $(\boldsymbol{\Theta}_1^*)_{\text{off}}$:

$$\begin{aligned} & \left\| \left(\boldsymbol{\Theta}_1^* - \zeta^2 \frac{\mathbf{X}\mathbf{X}^\top}{d} - c_3 \frac{\text{Tr}[\boldsymbol{\Sigma}^2]}{d^2} \mathbf{1}\mathbf{1}^\top \right)_{\text{off}} \right\| \\ & \leq c_1 \cdot \tilde{O}\left(\frac{n}{d^{1.5}}\right) + c_2 \cdot \tilde{O}\left(\frac{n}{d^{1.5}}\right) + c_3 \cdot \tilde{O}\left(\frac{n}{d^{1.25}}\right) + \tilde{O}\left(\frac{n}{d^{1.5}}\right) \\ & = \tilde{O}\left(\frac{n}{d^{1.25}}\right). \end{aligned} \quad (22)$$

Here we have used [Lemma B.2](#) to bound the spectral norm of the off-diagonal part of a matrix by the spectral norm of the matrix itself. Notice $c_3 \frac{\text{Tr}[\boldsymbol{\Sigma}^2]}{d} = (\mathbb{E}[g\phi'(g)])^2 \cdot \frac{\text{Tr}[\boldsymbol{\Sigma}^2]}{d} = \nu^2$ (c.f. (5)). Hence (22) becomes

$$\|(\boldsymbol{\Theta}_1^* - \boldsymbol{\Theta}^{\text{lin1}})_{\text{off}}\| = \tilde{O}\left(\frac{n}{d^{1.25}}\right). \quad (23)$$

For the diagonal entries of $\boldsymbol{\Theta}_1^*$, we have $[\boldsymbol{\Theta}_1^*]_{i,i} = \frac{\|\mathbf{x}_i\|^2}{d} \cdot \Phi\left(\frac{\|\mathbf{x}_i\|^2}{d}, \frac{\|\mathbf{x}_i\|^2}{d}, \frac{\|\mathbf{x}_i\|^2}{d}\right)$. We denote $\bar{\Phi}(a) := \Phi(a, a, a)$ ($a \geq 0$). When ϕ is a smooth activation as in [Assumption 3.2](#), we know that $\bar{\Phi}$ has bounded derivative, and thus we get

$$[\boldsymbol{\Theta}_1^*]_{i,i} = \frac{\|\mathbf{x}_i\|^2}{d} \cdot \bar{\Phi}\left(\frac{\|\mathbf{x}_i\|^2}{d}\right) = \left(1 \pm \tilde{O}\left(\frac{1}{\sqrt{d}}\right)\right) \cdot \left(\bar{\Phi}(1) \pm \tilde{O}\left(\frac{1}{\sqrt{d}}\right)\right) = \bar{\Phi}(1) \pm \tilde{O}\left(\frac{1}{\sqrt{d}}\right). \quad (24)$$

When ϕ is a piece-wise linear activation as in [Assumption 3.2](#), $\bar{\Phi}(a)$ is a constant, so we have $\bar{\Phi}\left(\frac{\|\mathbf{x}_i\|^2}{d}\right) = \bar{\Phi}(1)$. Therefore (24) also holds. Notice that $\bar{\Phi}(1) = \mathbb{E}[(\phi'(g))^2] =: \gamma$. It follows from (24) that

$$\|(\boldsymbol{\Theta}_1^*)_{\text{diag}} - \gamma \mathbf{I}\| = \tilde{O}\left(\frac{1}{\sqrt{d}}\right).$$

Also note that

$$\|\boldsymbol{\Theta}_{\text{diag}}^{\text{lin1}} - \zeta^2 \mathbf{I}\| = \tilde{O}\left(\frac{1}{\sqrt{d}}\right).$$

Therefore we obtain

$$\|(\boldsymbol{\Theta}_1^* - \boldsymbol{\Theta}^{\text{lin1}})_{\text{diag}} - (\gamma - \zeta^2) \mathbf{I}\| = \tilde{O}\left(\frac{1}{\sqrt{d}}\right). \quad (25)$$

Combining the off-diagonal and diagonal approximations (23) and (25), we obtain

$$\|\Theta_1^* - \Theta^{\text{lin1}} - (\gamma - \zeta^2)\mathbf{I}\| = \tilde{O}\left(\frac{n}{d^{1.25}}\right).$$

Finally, when $n \gtrsim d^{1+\alpha}$ ($0 < \alpha < \frac{1}{4}$), we have $\|\mathbf{I}\| = 1 \lesssim \frac{n}{d^{1+\alpha}}$. Hence we can discard the identity component above and get

$$\|\Theta_1^* - \Theta^{\text{lin1}}\| = O\left(\frac{n}{d^{1+\alpha}}\right).$$

This completes the proof. \square

Combining Propositions D.3 and D.4 directly gives Proposition D.2.

D.3.2 Agreement on Training Data

Now we prove the first part of Theorem 3.2, i.e., (7), which says that the neural network f_t^1 and the linear model f_t^{lin1} are close on the training data. We will use Theorem C.2, and the most important step is to verify Assumption C.1. To this end we prove the following Jacobian perturbation lemma.

Lemma D.5 (Jacobian perturbation for the first layer). *If ϕ is a smooth activation as in Assumption 3.2, then with high probability over the training data \mathbf{X} , we have*

$$\left\| \mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\widetilde{\mathbf{W}}, \mathbf{v}) \right\| \lesssim \sqrt{\frac{n}{md}} \left\| \mathbf{W} - \widetilde{\mathbf{W}} \right\|_F, \quad \forall \mathbf{W}, \widetilde{\mathbf{W}} \in \mathbb{R}^{m \times d}. \quad (26)$$

If ϕ is a piece-wise linear activation as in Assumption 3.2, then with high probability over the random initialization $\mathbf{W}(0)$ and the training data \mathbf{X} , we have

$$\left\| \mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}) \right\| \lesssim \sqrt{\frac{n}{d}} \left(\frac{\|\mathbf{W} - \mathbf{W}(0)\|^{1/3}}{m^{1/6}} + \left(\frac{\log n}{m}\right)^{1/4} \right), \quad \forall \mathbf{W} \in \mathbb{R}^{m \times d}. \quad (27)$$

Proof. Throughout the proof we will be conditioned on \mathbf{X} and on the high-probability events in Claim 3.1.

By the definition of $\mathbf{J}_1(\mathbf{W}, \mathbf{v})$ in (15), we have

$$\begin{aligned} & (\mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\widetilde{\mathbf{W}}, \mathbf{v}))(\mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\widetilde{\mathbf{W}}, \mathbf{v}))^\top \\ &= \frac{1}{md} \left(\phi' \left(\mathbf{X} \mathbf{W}^\top / \sqrt{d} \right) - \phi' \left(\mathbf{X} \widetilde{\mathbf{W}}^\top / \sqrt{d} \right) \right) \left(\phi' \left(\mathbf{X} \mathbf{W}^\top / \sqrt{d} \right) - \phi' \left(\mathbf{X} \widetilde{\mathbf{W}}^\top / \sqrt{d} \right) \right)^\top \odot (\mathbf{X} \mathbf{X}^\top). \end{aligned} \quad (28)$$

Then if ϕ is a smooth activation, we have with high probability,

$$\begin{aligned} & \left\| \mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\widetilde{\mathbf{W}}, \mathbf{v}) \right\|^2 \\ & \leq \frac{1}{md} \left\| \phi' \left(\mathbf{X} \mathbf{W}^\top / \sqrt{d} \right) - \phi' \left(\mathbf{X} \widetilde{\mathbf{W}}^\top / \sqrt{d} \right) \right\|^2 \cdot \max_{i \in [n]} \|\mathbf{x}_i\|^2 \quad ((28) \text{ and Lemma B.3}) \\ & \lesssim \frac{1}{md} \left\| \phi' \left(\mathbf{X} \mathbf{W}^\top / \sqrt{d} \right) - \phi' \left(\mathbf{X} \widetilde{\mathbf{W}}^\top / \sqrt{d} \right) \right\|_F^2 \cdot d \quad (\text{Claim 3.1}) \\ & \lesssim \frac{1}{md} \left\| \mathbf{X} \mathbf{W}^\top / \sqrt{d} - \mathbf{X} \widetilde{\mathbf{W}}^\top / \sqrt{d} \right\|_F^2 \cdot d \quad (\phi'' \text{ is bounded}) \\ & = \frac{1}{md} \left\| \mathbf{X}(\mathbf{W} - \widetilde{\mathbf{W}})^\top \right\|_F^2 \\ & \leq \frac{1}{md} \|\mathbf{X}\|^2 \left\| \mathbf{W} - \widetilde{\mathbf{W}} \right\|_F^2 \\ & \lesssim \frac{n}{md} \left\| \mathbf{W} - \widetilde{\mathbf{W}} \right\|_F^2. \quad (\text{Claim 3.1}) \end{aligned}$$

This proves (26).

Next we consider the case where ϕ is a piece-wise linear activation. From (28) and Lemma B.3 we have

$$\begin{aligned} \|\mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v})\|^2 &\leq \frac{1}{md} \|\mathbf{X}\mathbf{X}^\top\| \cdot \max_{i \in [n]} \left\| \phi'(\mathbf{W}\mathbf{x}_i/\sqrt{d}) - \phi'(\mathbf{W}(0)\mathbf{x}_i/\sqrt{d}) \right\|^2 \\ &\lesssim \frac{n}{md} \cdot \max_{i \in [n]} \left\| \phi'(\mathbf{W}\mathbf{x}_i/\sqrt{d}) - \phi'(\mathbf{W}(0)\mathbf{x}_i/\sqrt{d}) \right\|^2. \end{aligned} \quad (29)$$

For each $i \in [n]$, let

$$M_i = \{r \in [m] : \text{sign}(\mathbf{w}_r^\top \mathbf{x}_i) \neq \text{sign}(\mathbf{w}_r(0)^\top \mathbf{x}_i)\}$$

Since ϕ' is a step function that only depends on the sign of the input, we have

$$\left\| \phi'(\mathbf{W}\mathbf{x}_i/\sqrt{d}) - \phi'(\mathbf{W}(0)\mathbf{x}_i/\sqrt{d}) \right\|^2 \lesssim |M_i|, \quad \forall i \in [n]. \quad (30)$$

Therefore we need to bound $|M_i|$, i.e. how many coordinates in $\mathbf{W}\mathbf{x}_i$ and $\mathbf{W}(0)\mathbf{x}_i$ differ in sign for each $i \in [n]$.

Let $\lambda > 0$ be a parameter whose value will be determined later. For each $i \in [n]$, define

$$N_i := \{r \in [m] : |\mathbf{w}_r(0)^\top \mathbf{x}_i| \leq \lambda \|\mathbf{x}_i\|\}.$$

We have

$$|N_i| = \sum_{r=1}^m \mathbb{1}_{\{|\mathbf{w}_r(0)^\top \mathbf{x}_i| \leq \lambda \|\mathbf{x}_i\|\}} = 2 \sum_{r=1}^{m/2} \mathbb{1}_{\{|\mathbf{w}_r(0)^\top \mathbf{x}_i| \leq \lambda \|\mathbf{x}_i\|\}},$$

where the second equality is due to the symmetric initialization (3). Since $\frac{\mathbf{w}_r(0)^\top \mathbf{x}_i}{\|\mathbf{x}_i\|} \sim \mathcal{N}(0, 1)$, we have $\mathbb{E}[\mathbb{1}_{\{|\mathbf{w}_r(0)^\top \mathbf{x}_i| \leq \lambda \|\mathbf{x}_i\|\}}] = \Pr[|g| \leq \lambda] \leq \frac{2\lambda}{\sqrt{2\pi}}$. Also note that $\mathbf{w}_1(0), \dots, \mathbf{w}_{m/2}(0)$ are independent. Then by Hoeffding's inequality we know that with probability at least $1 - \delta$,

$$|N_i| \leq \sqrt{\frac{2}{\pi}} \lambda m + O\left(\sqrt{m \log \frac{1}{\delta}}\right).$$

Taking a union bound over all $i \in [n]$, we know that with high probability,

$$|N_i| \lesssim \lambda m + \sqrt{m \log n}, \quad \forall i \in [n]. \quad (31)$$

By definition, if $r \in M_i$ but $r \notin N_i$, we must have $|\mathbf{w}_r^\top \mathbf{x}_i - \mathbf{w}_r(0)^\top \mathbf{x}_i| \geq |\mathbf{w}_r(0)^\top \mathbf{x}_i| > \lambda \|\mathbf{x}_i\|$. This leads to

$$\begin{aligned} \|(\mathbf{W} - \mathbf{W}(0))\mathbf{x}_i\|^2 &= \sum_{r=1}^m |(\mathbf{w}_r - \mathbf{w}_r(0))^\top \mathbf{x}_i|^2 \geq \sum_{r \in M_i \setminus N_i} |(\mathbf{w}_r - \mathbf{w}_r(0))^\top \mathbf{x}_i|^2 \\ &\geq \sum_{r \in M_i \setminus N_i} \lambda^2 \|\mathbf{x}_i\|^2 \gtrsim \sum_{r \in M_i \setminus N_i} \lambda^2 d = \lambda^2 d |M_i \setminus N_i| \end{aligned}$$

Thus we have

$$|M_i \setminus N_i| \lesssim \frac{\|(\mathbf{W} - \mathbf{W}(0))\mathbf{x}_i\|^2}{\lambda^2 d} \leq \frac{\|\mathbf{W} - \mathbf{W}(0)\|^2 \|\mathbf{x}_i\|^2}{\lambda^2 d} \lesssim \frac{\|\mathbf{W} - \mathbf{W}(0)\|^2}{\lambda^2}, \quad \forall i \in [n]. \quad (32)$$

Combining (31) and (32) we obtain

$$|M_i| \lesssim \lambda m + \sqrt{m \log n} + \frac{\|\mathbf{W} - \mathbf{W}(0)\|^2}{\lambda^2}, \quad \forall i \in [n].$$

Letting $\lambda = \left(\frac{\|\mathbf{W} - \mathbf{W}(0)\|^2}{m}\right)^{1/3}$, we get

$$|M_i| \lesssim m^{2/3} \|\mathbf{W} - \mathbf{W}(0)\|^{2/3} + \sqrt{m \log n}, \quad \forall i \in [n]. \quad (33)$$

Finally, we combine (29), (30) and (33) to obtain

$$\begin{aligned}\|\mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v})\|^2 &\lesssim \frac{n}{md} \left(m^{2/3} \|\mathbf{W} - \mathbf{W}(0)\|^{2/3} + \sqrt{m \log n} \right) \\ &= \frac{n}{d} \left(\frac{\|\mathbf{W} - \mathbf{W}(0)\|^{2/3}}{m^{1/3}} + \sqrt{\frac{\log n}{m}} \right).\end{aligned}$$

This proves (27). \square

The next lemma verifies Assumption C.1 for the case of training the first layer.

Lemma D.6. *Let $R = \sqrt{d \log d}$. With high probability over the random initialization $\mathbf{W}(0)$ and the training data \mathbf{X} , for all $\mathbf{W}, \widetilde{\mathbf{W}} \in \mathbb{R}^{m \times d}$ such that $\|\mathbf{W} - \mathbf{W}(0)\|_F \leq R$ and $\|\widetilde{\mathbf{W}} - \mathbf{W}(0)\|_F \leq R$, we have*

$$\left\| \mathbf{J}_1(\mathbf{W}, \mathbf{v}) \mathbf{J}_1(\widetilde{\mathbf{W}}, \mathbf{v})^\top - \Theta^{\text{lin1}} \right\| \lesssim \frac{n}{d^{1+\frac{\alpha}{7}}}.$$

Proof. This proof is conditioned on all the high-probability events we have shown.

Now consider $\mathbf{W}, \widetilde{\mathbf{W}} \in \mathbb{R}^{m \times d}$ such that $\|\mathbf{W} - \mathbf{W}(0)\|_F \leq R$ and $\|\widetilde{\mathbf{W}} - \mathbf{W}(0)\|_F \leq R$. If ϕ is a smooth activation, from Lemma D.5 we have

$$\|\mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v})\| \lesssim \sqrt{\frac{n}{md}} \|\mathbf{W} - \mathbf{W}(0)\|_F \leq \sqrt{\frac{n}{md}} \cdot \sqrt{d \log d} \lesssim \sqrt{\frac{n \log d}{d^{1+\alpha}}} \ll \sqrt{\frac{n}{d^{1+\frac{\alpha}{2}}}},$$

where we have used $m \gtrsim d^{1+\alpha}$. If ϕ is a piece-wise linear activation, from Lemma D.5 we have

$$\begin{aligned}\|\mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v})\| &\lesssim \sqrt{\frac{n}{d}} \left(\frac{\|\mathbf{W} - \mathbf{W}(0)\|^{1/3}}{m^{1/6}} + \left(\frac{\log n}{m} \right)^{1/4} \right) \\ &\leq \sqrt{\frac{n}{d}} \left(\frac{(d \log d)^{1/6}}{m^{1/6}} + \left(\frac{\log n}{m} \right)^{1/4} \right) \\ &\lesssim \sqrt{\frac{n}{d}} \cdot \frac{(d \log d)^{1/6}}{d^{1/6+\alpha/6}} \\ &\ll \frac{\sqrt{n}}{d^{\frac{1}{2}+\frac{\alpha}{7}}}.\end{aligned}$$

Hence we always have $\|\mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v})\| \leq \frac{\sqrt{n}}{d^{\frac{1}{2}+\frac{\alpha}{7}}}$. Similarly, we have

$$\left\| \mathbf{J}_1(\widetilde{\mathbf{W}}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}) \right\| \leq \frac{\sqrt{n}}{d^{\frac{1}{2}+\frac{\alpha}{7}}}.$$

Note that from Proposition D.2 and Claim 3.1 we know

$$\left\| \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}) \mathbf{J}_1(\mathbf{W}(0), \mathbf{v})^\top \right\| \lesssim \|\Theta^{\text{lin1}}\| + \frac{n}{d^{1+\alpha}} \lesssim \frac{n}{d} + \frac{n}{d^{1+\alpha}} \lesssim \frac{n}{d},$$

which implies $\|\mathbf{J}_1(\mathbf{W}(0), \mathbf{v})\| \lesssim \sqrt{\frac{n}{d}}$. It follows that $\|\mathbf{J}_1(\mathbf{W}, \mathbf{v})\| \lesssim \sqrt{\frac{n}{d}} + \frac{\sqrt{n}}{d^{\frac{1}{2}+\frac{\alpha}{7}}} \lesssim \sqrt{\frac{n}{d}}$ and

$\left\| \mathbf{J}_1(\widetilde{\mathbf{W}}, \mathbf{v}) \right\| \lesssim \sqrt{\frac{n}{d}}$. Then we have

$$\begin{aligned}&\left\| \mathbf{J}_1(\mathbf{W}, \mathbf{v}) \mathbf{J}_1(\widetilde{\mathbf{W}}, \mathbf{v})^\top - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}) \mathbf{J}_1(\mathbf{W}(0), \mathbf{v})^\top \right\| \\ &\leq \|\mathbf{J}_1(\mathbf{W}, \mathbf{v})\| \cdot \left\| \mathbf{J}_1(\widetilde{\mathbf{W}}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}) \right\| + \|\mathbf{J}_1(\mathbf{W}(0), \mathbf{v})\| \cdot \|\mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v})\| \\ &\lesssim \sqrt{\frac{n}{d}} \cdot \frac{\sqrt{n}}{d^{\frac{1}{2}+\frac{\alpha}{7}}} + \sqrt{\frac{n}{d}} \cdot \frac{\sqrt{n}}{d^{\frac{1}{2}+\frac{\alpha}{7}}} \\ &\lesssim \frac{n}{d^{1+\frac{\alpha}{7}}}.\end{aligned}$$

Combining the above inequality with Proposition D.2, we obtain

$$\left\| \mathbf{J}_1(\mathbf{W}, \mathbf{v}) \mathbf{J}_1(\widetilde{\mathbf{W}}, \mathbf{v})^\top - \Theta^{\text{lin1}} \right\| \lesssim \frac{n}{d^{1+\frac{\alpha}{7}}} + \frac{n}{d^{1+\alpha}} \lesssim \frac{n}{d^{1+\frac{\alpha}{7}}},$$

completing the proof. \square

Finally, we can instantiate Theorem C.2 to conclude the proof of (7):

Proposition D.7. *There exists a universal constant $c > 0$ such that with high probability, for all $0 \leq t \leq T = c \cdot \frac{d \log d}{\eta_1}$ simultaneously, we have:*

- $\frac{1}{n} \sum_{i=1}^n (f_t^1(\mathbf{x}_i) - f_t^{\text{lin}1}(\mathbf{x}_i))^2 \leq d^{-\frac{\alpha}{4}};$
- $\|\mathbf{W}(t) - \mathbf{W}(0)\|_F \leq \sqrt{d \log d}, \|\boldsymbol{\beta}(t)\| \leq \sqrt{d \log d}.$

Proof. Let $R = \sqrt{d \log d}$ and $\epsilon = C \frac{n}{d^{1+\frac{\alpha}{7}}}$ for a sufficiently large universal constant $C > 0$. From Lemma D.6 we know that Assumption C.1 is satisfied with parameters ϵ and R . (Note that $\epsilon \ll \frac{n}{d} \lesssim \|\Theta^{\text{lin}1}\|$.) Also we have $R^2 \epsilon \ll n$, and $\eta_1 \ll d \lesssim \frac{n}{\|\Theta^{\text{lin}1}\|}$. Therefore, we can apply Theorem C.2 and obtain for all $0 \leq t \leq T$:

$$\sqrt{\sum_{i=1}^n (f_t^1(\mathbf{x}_i) - f_t^{\text{lin}1}(\mathbf{x}_i))^2} \lesssim \frac{\eta_1 t \epsilon}{\sqrt{n}} \lesssim \frac{d \log d \cdot \frac{n}{d^{1+\frac{\alpha}{7}}}}{\sqrt{n}} = \frac{\sqrt{n} \log d}{d^{\frac{\alpha}{7}}} \ll \frac{\sqrt{n}}{d^{\frac{\alpha}{8}}},$$

which implies

$$\frac{1}{n} \sum_{i=1}^n (f_t^1(\mathbf{x}_i) - f_t^{\text{lin}1}(\mathbf{x}_i))^2 \leq d^{-\frac{\alpha}{4}}.$$

Furthermore, Theorem C.2 also tells us $\|\mathbf{W}(t) - \mathbf{W}(0)\|_F \leq \sqrt{d \log d}$ and $\|\boldsymbol{\beta}(t)\| \leq \sqrt{d \log d}$. \square

D.3.3 Agreement on Distribution

Now we prove the second part of Theorem 3.2, (8), which guarantees the agreement between f_t^1 and $f_t^{\text{lin}1}$ on the entire distribution \mathcal{D} . As usual, we will be conditioned on all the high-probability events unless otherwise noted.

Given the initialization $(\mathbf{W}(0), \mathbf{v})$ (recall that $\mathbf{v} = \mathbf{v}(0)$ is always fixed), we define an auxiliary model $f^{\text{aux}1}(\mathbf{x}; \mathbf{W})$ which is the first-order Taylor approximation of the neural network $f(\mathbf{x}; \mathbf{W}, \mathbf{v})$ around $\mathbf{W}(0)$:

$$\begin{aligned} f^{\text{aux}1}(\mathbf{x}; \mathbf{W}) &:= f(\mathbf{x}; \mathbf{W}(0), \mathbf{v}) + \langle \mathbf{W} - \mathbf{W}(0), \nabla_{\mathbf{W}} f(\mathbf{x}; \mathbf{W}(0), \mathbf{v}) \rangle \\ &= \langle \mathbf{W} - \mathbf{W}(0), \nabla_{\mathbf{W}} f(\mathbf{x}; \mathbf{W}(0), \mathbf{v}) \rangle \\ &= \langle \text{vec}(\mathbf{W} - \mathbf{W}(0)), \boldsymbol{\rho}_1(\mathbf{x}) \rangle, \end{aligned}$$

where $\boldsymbol{\rho}_1(\mathbf{x}) := \nabla_{\mathbf{W}} f(\mathbf{x}; \mathbf{W}(0), \mathbf{v})$. Above we have used $f(\mathbf{x}; \mathbf{W}(0), \mathbf{v}) = 0$ according to the symmetric initialization (3). We also denote $f_t^{\text{aux}1}(\mathbf{x}) := f^{\text{aux}1}(\mathbf{x}; \mathbf{W}(t))$ for all t .

For all models, we write their predictions on all training datapoints concisely as $f_t^1(\mathbf{X}), f_t^{\text{lin}1}(\mathbf{X}), f_t^{\text{aux}1}(\mathbf{X}) \in \mathbb{R}^n$. From Proposition D.7 we know that f_t^1 and $f_t^{\text{lin}1}$ make similar predictions on \mathbf{X} (for all $t \leq T$ simultaneously):

$$\|f_t^1(\mathbf{X}) - f_t^{\text{lin}1}(\mathbf{X})\| \leq \frac{\sqrt{n}}{d^{\frac{\alpha}{8}}}. \quad (34)$$

We can also related the predictions of f_t^1 and $f_t^{\text{aux}1}$ by the fundamental theorem for line integrals:

$$\begin{aligned} f_t^1(\mathbf{X}) &= f_t^1(\mathbf{X}) - f_0^1(\mathbf{X}) = \mathbf{J}_1(\mathbf{W}(0) \rightarrow \mathbf{W}(t), \mathbf{v}) \cdot \text{vec}(\mathbf{W}(t) - \mathbf{W}(0)), \\ f_t^{\text{aux}1}(\mathbf{X}) &= f_t^{\text{aux}1}(\mathbf{X}) - f_0^{\text{aux}1}(\mathbf{X}) = \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}) \cdot \text{vec}(\mathbf{W}(t) - \mathbf{W}(0)), \end{aligned} \quad (35)$$

where $\mathbf{J}_1(\mathbf{W}(0) \rightarrow \mathbf{W}(t), \mathbf{v}) := \int_0^1 \mathbf{J}_1(\mathbf{W}(0) + x(\mathbf{W}(t) - \mathbf{W}(0)), \mathbf{v}) dx$. Since $\|\mathbf{W}(t) - \mathbf{W}(0)\|_F \leq \sqrt{d \log d}$ according to Proposition D.7, we can use Lemma D.5 in the same way as in the proof of Lemma D.6 and obtain

$$\|\mathbf{J}_1(\mathbf{W}(0) \rightarrow \mathbf{W}(t), \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v})\| \leq \frac{\sqrt{n}}{d^{\frac{1}{2} + \frac{\alpha}{7}}}.$$

Then it follows from (35) that

$$\begin{aligned}
\|f_t^1(\mathbf{X}) - f_t^{\text{aux1}}(\mathbf{X})\| &= \|(\mathbf{J}_1(\mathbf{W}(0) \rightarrow \mathbf{W}(t), \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v})) \cdot \text{vec}(\mathbf{W}(t) - \mathbf{W}(0))\| \\
&\leq \frac{\sqrt{n}}{d^{\frac{1}{2} + \frac{\alpha}{7}}} \cdot \sqrt{d \log d} \\
&\leq \frac{\sqrt{n}}{d^{\frac{\alpha}{8}}}.
\end{aligned} \tag{36}$$

Combining (34) and (36) we know

$$\|f_t^{\text{aux1}}(\mathbf{X}) - f_t^{\text{lin1}}(\mathbf{X})\| \lesssim \frac{\sqrt{n}}{d^{\frac{\alpha}{8}}}.$$

This implies

$$\frac{1}{n} \sum_{i=1}^n \min \left\{ (f_t^{\text{aux1}}(\mathbf{x}_i) - f_t^{\text{lin1}}(\mathbf{x}_i))^2, 1 \right\} \leq \frac{1}{n} \sum_{i=1}^n (f_t^{\text{aux1}}(\mathbf{x}_i) - f_t^{\text{lin1}}(\mathbf{x}_i))^2 \lesssim d^{-\frac{\alpha}{4}}.$$

Next we will translate these guarantees on the training data to the distribution \mathcal{D} using Rademacher complexity. Note that the model $f_t^{\text{aux1}}(\mathbf{x}) - f_t^{\text{lin1}}(\mathbf{x})$ is by definition linear in the feature $\begin{bmatrix} \boldsymbol{\rho}_1(\mathbf{x}) \\ \boldsymbol{\psi}_1(\mathbf{x}) \end{bmatrix}$, and it belongs to the following function class (for all $t \leq T$):

$$\mathcal{F} := \left\{ \mathbf{x} \mapsto \mathbf{a}^\top \begin{bmatrix} \boldsymbol{\rho}_1(\mathbf{x}) \\ \boldsymbol{\psi}_1(\mathbf{x}) \end{bmatrix} : \|\mathbf{a}\| \leq 2\sqrt{d \log d} \right\}.$$

This is because we have $\|\text{vec}(\mathbf{W}(t) - \mathbf{W}(0))\| \leq \sqrt{d \log d}$ and $\|\boldsymbol{\beta}(t)\| \leq \sqrt{d \log d}$ for all $t \leq T$. Using the well-known bound on the empirical Rademacher complexity of a linear function class with bounded ℓ_2 norm (see e.g. [Bartlett and Mendelson \[2002\]](#)), we can bound the empirical Rademacher complexity of the function class \mathcal{F} :

$$\begin{aligned}
\hat{\mathcal{R}}_{\mathbf{X}}(\mathcal{F}) &:= \frac{1}{n} \mathbb{E}_{\varepsilon_1, \dots, \varepsilon_n \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(\{\pm 1\})} \left[\sup_{h \in \mathcal{F}} \sum_{i=1}^n \varepsilon_i h(\mathbf{x}_i) \right] \\
&\lesssim \frac{\sqrt{d \log d}}{n} \sqrt{\sum_{i=1}^n (\|\boldsymbol{\rho}_1(\mathbf{x}_i)\|^2 + \|\boldsymbol{\psi}_1(\mathbf{x}_i)\|^2)} \\
&= \frac{\sqrt{d \log d}}{n} \sqrt{\text{Tr}[\boldsymbol{\Theta}_1(\mathbf{W}(0), \mathbf{v})] + \text{Tr}[\boldsymbol{\Theta}^{\text{lin1}}]}.
\end{aligned} \tag{37}$$

Since ϕ' is bounded and $\frac{\|\mathbf{x}_i\|^2}{d} = O(1)$ ($\forall i \in [n]$), we can bound

$$\text{Tr}[\boldsymbol{\Theta}_1(\mathbf{W}(0), \mathbf{v})] = \sum_{i=1}^n \frac{1}{m} \sum_{r=1}^m \phi' \left(\mathbf{w}_r(0)^\top \mathbf{x}_i / \sqrt{d} \right)^2 \cdot \frac{\|\mathbf{x}_i\|^2}{d} \lesssim n,$$

and

$$\text{Tr}[\boldsymbol{\Theta}^{\text{lin1}}] = \sum_{i=1}^n \left(\zeta^2 \frac{\|\mathbf{x}_i\|^2}{d} + \frac{\nu^2}{d} \right) \lesssim n.$$

Therefore we have

$$\hat{\mathcal{R}}_{\mathbf{X}}(\mathcal{F}) \lesssim \frac{\sqrt{d \log d}}{n} \sqrt{n} = \sqrt{\frac{d \log d}{n}}.$$

Now using the standard generalization bound via Rademacher complexity (see e.g. [Mohri et al. \[2012\]](#)), and noticing that the function $z \mapsto \min\{z^2, 1\}$ is 2-Lipschitz and bounded in $[0, 1]$, we have with high probability, for all $t \leq T$ simultaneously,

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\min \left\{ (f_t^{\text{aux1}}(\mathbf{x}) - f_t^{\text{lin1}}(\mathbf{x}))^2, 1 \right\} \right]$$

$$\begin{aligned}
&\leq \frac{1}{n} \sum_{i=1}^n \min \left\{ (f_t^{\text{aux1}}(\mathbf{x}_i) - f_t^{\text{lin1}}(\mathbf{x}_i))^2, 1 \right\} + O\left(\sqrt{\frac{d \log d}{n}}\right) + O\left(\frac{1}{\sqrt{n}}\right) \\
&\lesssim d^{-\frac{\alpha}{4}} + \sqrt{\frac{d \log d}{d^{1+\alpha}}} \quad (n \gtrsim d^{1+\alpha}) \\
&\lesssim d^{-\frac{\alpha}{4}}. \tag{38}
\end{aligned}$$

Therefore we have shown that f_t^{aux1} and f_t^{lin1} are close on the distribution \mathcal{D} for all $t \leq T$. To complete the proof, we need to show that f_t^1 and f_t^{aux1} are close on \mathcal{D} . For this, we take an imaginary set of test datapoints $\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}$, which are independent of the training samples. Let $\tilde{\mathbf{X}} \in \mathbb{R}^{n \times d}$ be the corresponding test data matrix. Since the test data are from the same distribution \mathcal{D} , the concentration properties in Claim 3.1 still hold, and the Jacobian perturbation bounds in Lemma D.5 hold as well. Hence we can apply the exact same arguments in (36) and obtain with high probability for all $t \leq T$,

$$\left\| f_t^1(\tilde{\mathbf{X}}) - f_t^{\text{aux1}}(\tilde{\mathbf{X}}) \right\| \leq \frac{\sqrt{n}}{d^{\frac{\alpha}{8}}},$$

which implies

$$\frac{1}{n} \sum_{i=1}^n \min \left\{ (f_t^1(\tilde{\mathbf{x}}_i) - f_t^{\text{aux1}}(\tilde{\mathbf{x}}_i))^2, 1 \right\} \leq d^{-\frac{\alpha}{4}}.$$

Now notice that f_t^1 and f_t^{aux1} are independent of $\tilde{\mathbf{X}}$. Thus, by Hoeffding inequality, for each t , with probability at least $1 - \delta$ we have

$$\begin{aligned}
&\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\min \left\{ (f_t^1(\mathbf{x}) - f_t^{\text{aux1}}(\mathbf{x}))^2, 1 \right\} \right] \\
&\leq \frac{1}{n} \sum_{i=1}^n \min \left\{ (f_t^1(\tilde{\mathbf{x}}_i) - f_t^{\text{aux1}}(\tilde{\mathbf{x}}_i))^2, 1 \right\} + O\left(\sqrt{\frac{\log \frac{1}{\delta}}{n}}\right) \\
&\lesssim d^{-\frac{\alpha}{4}} + \sqrt{\frac{\log \frac{1}{\delta}}{n}}.
\end{aligned}$$

Then letting $\delta = \frac{1}{100T}$ and taking a union bound over $t \leq T$, we obtain that with high probability, for all $t \leq T$ simultaneously,

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\min \left\{ (f_t^1(\mathbf{x}) - f_t^{\text{aux1}}(\mathbf{x}))^2, 1 \right\} \right] \lesssim d^{-\frac{\alpha}{4}} + \sqrt{\frac{\log T}{n}}. \tag{39}$$

Therefore we have proved that f_t^1 and f_t^{aux1} are close on \mathcal{D} . Finally, combining (38) and (39), we know that with high probability, for all $t \leq T$,

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\min \left\{ (f_t^1(\mathbf{x}) - f_t^{\text{lin1}}(\mathbf{x}))^2, 1 \right\} \right] \lesssim d^{-\frac{\alpha}{4}} + \sqrt{\frac{\log T}{n}}.$$

Here we have used $\min\{(a+b)^2, 1\} \leq 2(\min\{a^2, 1\} + \min\{b^2, 1\})$ ($\forall a, b \in \mathbb{R}$). Therefore we have finished the proof of (8). The proof of Theorem 3.2 is done.

D.4 Proof of Corollary 3.3 (Training the First Layer, Well-Conditioned Data)

Proof of Corollary 3.3. We continue to adopt the notation in Appendix D.3.3 to use $f_t^1(\mathbf{X})$, $f_t^{\text{lin1}}(\mathbf{X})$, etc. to represent the predictions of a model on all n training datapoints. Given Theorem 3.2, it suffices to prove that f_T^{lin1} and f_*^{lin1} are close in the following sense:

$$\frac{1}{n} \sum_{i=1}^n (f_T^{\text{lin1}}(\mathbf{x}_i) - f_*^{\text{lin1}}(\mathbf{x}_i))^2 \lesssim d^{-\Omega(\alpha)}, \tag{40}$$

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\min \left\{ (f_T^{\text{lin1}}(\mathbf{x}) - f_*^{\text{lin1}}(\mathbf{x}))^2, 1 \right\} \right] \lesssim d^{-\Omega(\alpha)}. \tag{41}$$

According to the linear dynamics (6), we have the following relation (see Claim C.1):

$$\begin{aligned} f_T^{\text{lin1}}(\mathbf{X}) - \mathbf{y} &= \left(\mathbf{I} - \frac{1}{n}\eta_1 \Theta^{\text{lin1}}\right)^T (-\mathbf{y}), \\ f_*^{\text{lin1}}(\mathbf{X}) - \mathbf{y} &= \lim_{t \rightarrow \infty} \left(\mathbf{I} - \frac{1}{n}\eta_1 \Theta^{\text{lin1}}\right)^t (-\mathbf{y}) =: \left(\mathbf{I} - \frac{1}{n}\eta_1 \Theta^{\text{lin1}}\right)^\infty (-\mathbf{y}). \end{aligned}$$

From the well-conditioned data assumption, it is easy to see that Θ^{lin1} 's non-zero eigenvalues are all $\Omega(\frac{n}{d})$ with high probability. As a consequence, in all the non-zero eigen-directions of Θ^{lin1} , the corresponding eigenvalues of $\left(\mathbf{I} - \frac{1}{n}\eta_1 \Theta^{\text{lin1}}\right)^T$ are at most $\left(1 - \frac{1}{n}\eta_1 \cdot \Omega(\frac{n}{d})\right)^T \leq \exp\left(-\Omega\left(\frac{\eta_1 T}{d}\right)\right) = \exp(-\Omega(\log d)) = d^{-\Omega(1)}$. This implies

$$\left\| f_T^{\text{lin1}}(\mathbf{X}) - f_*^{\text{lin1}}(\mathbf{X}) \right\| \leq \left\| \left(\mathbf{I} - \frac{1}{n}\eta_1 \Theta^{\text{lin1}}\right)^T - \left(\mathbf{I} - \frac{1}{n}\eta_1 \Theta^{\text{lin1}}\right)^\infty \right\| \cdot \|\mathbf{y}\| \lesssim d^{-\Omega(1)} \sqrt{n},$$

which completes the proof of (40).

To prove (41), we further apply the standard Rademacher complexity argument (similar to Appendix D.3.3). For this we just need to bound the ℓ_2 norm of the parameters, $\|\beta(T)\|$ and $\|\beta_*\|$. From Proposition D.7, we already have $\|\beta(T)\| \leq \sqrt{d \log d}$. Regarding β_* , we can directly write down its expression

$$\beta_* = (\Psi_1^\top \Psi_1)^\dagger \Psi_1^\top \mathbf{y}.$$

Here Ψ_1 is the feature matrix defined in (19), and \dagger stands for the Moore–Penrose pseudo-inverse. Recall that $\Theta^{\text{lin1}} = \Psi_1 \Psi_1^\top$. Notice that every non-zero singular value of $(\Psi_1^\top \Psi_1)^\dagger \Psi_1^\top$ is the inverse of a non-zero singular value of Ψ_1 , and that every non-zero singular value of Ψ_1 is $\Omega(\sqrt{\frac{n}{d}})$. This implies $\|(\Psi_1^\top \Psi_1)^\dagger \Psi_1^\top\| \lesssim \sqrt{\frac{d}{n}}$. Hence we have

$$\|\beta_*\| \lesssim \sqrt{\frac{d}{n}} \sqrt{n} = \sqrt{d}.$$

Therefore we can apply the standard Rademacher complexity argument and conclude the proof of (41). \square

D.5 Proof of Theorem 3.5 (Training the Second Layer)

Since the first layer is kept fixed in this case, we let $\mathbf{W} = \mathbf{W}(0)$ for notational convenience. Similar to the proof of Theorem 3.2 in Appendix D.3, we still divide the proof into 3 parts: analyzing the NTK at initialization (which is also the NTK throughout training in this case), proving the agreement on training data, and proving the agreement on the distribution.

It is easy to see from the definition of Θ^{lin2} in (20) and Claim 3.1 that if $\vartheta_0 \neq 0$, then $\|\Theta^{\text{lin2}}\| = O(n)$ with high probability, and if $\vartheta_0 = 0$, then $\|\Theta^{\text{lin2}}\| = O(\frac{n \log n}{d})$ with high probability. As we will see in the proof, this is why we distinguish these two cases in Theorem 3.5.

D.5.1 The NTK at Initialization

Proposition D.8. *With high probability over the random initialization \mathbf{W} and the training data \mathbf{X} , we have*

$$\|\Theta_2(\mathbf{W}) - \Theta^{\text{lin2}}\| \lesssim \frac{n}{d^{1+\frac{\alpha}{3}}}.$$

To prove Proposition D.8, we will prove $\Theta_2(\mathbf{W})$ is close to its expectation Θ_2^* (defined in (18)), and then prove Θ_2^* is close to Θ^{lin2} . We do these steps in the next two propositions.

Proposition D.9. *With high probability over the training data \mathbf{X} , we have*

$$\|\Theta_2^* - \Theta^{\text{lin2}}\| \lesssim \frac{n}{d^{1+\alpha}}.$$

Proof. We will be conditioned on the high probability events stated in Claim 3.1.

By the definition of Θ_2^* , we know

$$[\Theta_2^*]_{i,j} = \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\phi(\mathbf{w}^\top \mathbf{x}_i / \sqrt{d}) \phi(\mathbf{w}^\top \mathbf{x}_j / \sqrt{d}) \right], \quad i, j \in [n].$$

We define

$$\Gamma(a, b, c) := \mathbb{E}_{(z_1, z_2) \sim \mathcal{N}(\mathbf{0}, \mathbf{\Lambda})} [\phi(z_1)\phi(z_2)], \text{ where } \mathbf{\Lambda} = \begin{pmatrix} a^2 & c \\ c & b^2 \end{pmatrix}, \quad a \geq 0, b \geq 0, |c| \leq ab.$$

Then we can write

$$[\Theta_2^*]_{i,j} = \Gamma\left(\frac{\|\mathbf{x}_i\|}{\sqrt{d}}, \frac{\|\mathbf{x}_j\|}{\sqrt{d}}, \frac{\mathbf{x}_i^\top \mathbf{x}_j}{d}\right).$$

Denote $e_i := \frac{\|\mathbf{x}_i\|}{\sqrt{d}} - 1$ and $s_{i,j} := \frac{\mathbf{x}_i^\top \mathbf{x}_j}{d}$. Below we consider the diagonal and off-diagonal entries of Θ_2^* separately.

For $i \neq j$, we do a Taylor expansion of Γ around $(1, 1, 0)$:

$$\begin{aligned} & [\Theta_2^*]_{i,j} \\ &= \Gamma(1, 1, 0) + \nabla \Gamma(1, 1, 0)^\top \begin{bmatrix} e_i \\ e_j \\ s_{i,j} \end{bmatrix} + \frac{1}{2} [e_i, e_j, s_{i,j}] \cdot \nabla^2 \Gamma(1, 1, 0) \cdot \begin{bmatrix} e_i \\ e_j \\ s_{i,j} \end{bmatrix} + O(|e_i|^3 + |e_j|^3 + |s_{i,j}|^3) \\ &= \vartheta_0^2 + \vartheta_0 \vartheta_1 (e_i + e_j) + \zeta^2 s_{i,j} + \vartheta_0 \vartheta_2 (e_i^2 + e_j^2) + \vartheta_1^2 e_i e_j + \frac{1}{2} \vartheta_1^2 s_{i,j}^2 + \gamma s_{i,j} (e_i + e_j) \pm \tilde{O}\left(\frac{1}{d^{3/2}}\right) \\ &= (\vartheta_0 + \vartheta_1 e_i + \vartheta_2 e_i^2)(\vartheta_0 + \vartheta_1 e_j + \vartheta_2 e_j^2) - \vartheta_1 \vartheta_2 (e_i e_j^2 + e_i^2 e_j) - \vartheta_2^2 e_i^2 e_j^2 + \zeta^2 s_{i,j} + \frac{1}{2} \vartheta_1^2 s_{i,j}^2 \\ &\quad + \gamma s_{i,j} (e_i + e_j) \pm \tilde{O}\left(\frac{1}{d^{3/2}}\right) \\ &= [\mathbf{q}]_i [\mathbf{q}]_j \pm \tilde{O}\left(\frac{1}{d^{3/2}}\right) \pm \tilde{O}\left(\frac{1}{d^2}\right) + \zeta^2 s_{i,j} + \frac{1}{2} \vartheta_1^2 s_{i,j}^2 + \gamma s_{i,j} (e_i + e_j) \pm \tilde{O}\left(\frac{1}{d^{3/2}}\right) \\ &= [\mathbf{q}]_i [\mathbf{q}]_j + \zeta^2 s_{i,j} + \frac{1}{2} \vartheta_1^2 s_{i,j}^2 + \gamma s_{i,j} (e_i + e_j) \pm \tilde{O}\left(\frac{1}{d^{3/2}}\right). \end{aligned}$$

Here $\zeta, \vartheta_0, \vartheta_1, \vartheta_2$ are defined in (9), and γ is the $(1, 3)$ -th entry in the Hessian $\nabla^2 \Gamma(1, 1, 0)$ whose specific value is not important to us. Recall that $[\mathbf{q}]_i = \vartheta_0 + \vartheta_1 e_i + \vartheta_2 e_i^2$.

On the other hand, by the definition (20) we have

$$[\Theta^{\text{lin}2}]_{i,j} = \zeta^2 s_{i,j} + \frac{\nu^2}{2d} + [\mathbf{q}]_i [\mathbf{q}]_j.$$

It follows that

$$\begin{aligned} [\Theta_2^* - \Theta^{\text{lin}2}]_{i,j} &= \frac{1}{2} \vartheta_1^2 s_{i,j}^2 - \frac{\nu^2}{2d} + \gamma s_{i,j} (e_i + e_j) \pm \tilde{O}\left(\frac{1}{d^{3/2}}\right) \\ &= \frac{1}{2} \vartheta_1^2 \left(s_{i,j}^2 - \frac{\text{Tr}[\Sigma^2]}{d^2} \right) + \gamma s_{i,j} (e_i + e_j) \pm \tilde{O}\left(\frac{1}{d^{3/2}}\right). \end{aligned}$$

Here we have used the definition of ν in (5). In the proof of Proposition D.4, we have proved that all the error terms above contribute to at most $\tilde{O}\left(\frac{n}{d^{1.25}}\right)$ in spectral norm. Using the analysis there we get

$$\|(\Theta_2^* - \Theta^{\text{lin}2})_{\text{off}}\| = \tilde{O}\left(\frac{n}{d^{1.25}}\right).$$

Regarding the diagonal entries, it is easy to see that all the diagonal entries in Θ_2^* and $\Theta^{\text{lin}2}$ are $O(1)$, which implies

$$\|(\Theta_2^* - \Theta^{\text{lin}2})_{\text{diag}}\| = O(1).$$

Therefore we have

$$\|\Theta_2^* - \Theta^{\text{lin}2}\| = \tilde{O}\left(\frac{n}{d^{1.25}}\right) + O(1) = O\left(\frac{n}{d^{1+\alpha}}\right),$$

since $n \gtrsim d^{1+\alpha}$ ($0 < \alpha < \frac{1}{4}$). □

Proposition D.10. *With high probability over the random initialization \mathbf{W} and the training data \mathbf{X} , we have*

$$\|\Theta_2(\mathbf{W}) - \Theta_2^*\| \lesssim \frac{n}{d^{1+\frac{\alpha}{3}}}.$$

Proof. For convenience we denote $\Theta_2 = \Theta_2(\mathbf{W})$ in the proof. We will be conditioned on \mathbf{X} and on Claim 3.1, and only consider the randomness in \mathbf{W} . From Proposition D.9 we know that

$$\|\Theta_2^*\| = \begin{cases} \tilde{O}(n/d), & \text{if } \vartheta_0 = \mathbb{E}[\phi(g)] = 0 \\ O(n), & \text{otherwise} \end{cases}.$$

Define $\Theta_2^{(r)} := \phi(\mathbf{X}\mathbf{w}_r/\sqrt{d})\phi(\mathbf{X}\mathbf{w}_r/\sqrt{d})^\top$ for each $r \in [m]$. We have $\Theta_2 = \frac{1}{m} \sum_{r=1}^m \Theta_2^{(r)}$. According to the initialization scheme (3), we know that $\Theta_2^{(1)}, \Theta_2^{(2)}, \dots, \Theta_2^{(m/2)}$ are independent, $\Theta_2^{(m/2+1)}, \Theta_2^{(m/2+2)}, \dots, \Theta_2^{(m)}$ are independent, and $\mathbb{E}[\Theta_2^{(r)}] = \Theta_2^*$ for all $r \in [m]$.

Since the matrices $\Theta_2^{(r)}$ are possibly unbounded, we will use a variant of the matrix Bernstein inequality for unbounded matrices, which can be found as Proposition 4.1 in [Klochkov and Zhivotovskiy \[2020\]](#). There are two main steps in order to use this inequality: (i) showing that $\|\Theta_2^{(r)} - \Theta_2^*\|$ is a sub-exponential random variable for each r and bounding its sub-exponential norm; (ii) bounding the variance $\left\| \sum_{r=1}^{m/2} \mathbb{E}[(\Theta_2^{(r)} - \Theta_2^*)^2] \right\|$. For the first step, we have

$$\begin{aligned} \|\Theta_2^{(r)} - \Theta_2^*\| &\leq \|\Theta_2^{(r)}\| + \|\Theta_2^*\| \\ &= \left\| \phi(\mathbf{X}\mathbf{w}_r/\sqrt{d}) \right\|^2 + O(n) \\ &\lesssim \|\phi(\mathbf{0}_n)\|^2 + \left\| \mathbf{X}\mathbf{w}_r/\sqrt{d} \right\|^2 + n && (\phi \text{ is Lipschitz}) \\ &\lesssim n + \frac{\|\mathbf{X}\|^2 \|\mathbf{w}_r\|^2}{d} \\ &\lesssim n + \frac{n}{d} \|\mathbf{w}_r\|^2. \end{aligned}$$

Since $\|\mathbf{w}_r\|^2$ is a χ^2 random variable with d degrees of freedom, it has sub-exponential norm $O(d)$, which implies that the random variable $\|\Theta_2^{(r)} - \Theta_2^*\|$ has sub-exponential norm $O(n)$.

Next we bound the variance. Let $B > 0$ be a threshold to be determined. We have:

$$\begin{aligned} &\left\| \mathbb{E}[(\Theta_2^{(r)} - \Theta_2^*)^2] \right\| \\ &= \left\| \mathbb{E}[(\Theta_2^{(r)})^2] - (\Theta_2^*)^2 \right\| \\ &\leq \left\| \mathbb{E}[(\Theta_2^{(r)})^2] \right\| + \|\Theta_2^*\|^2 \\ &= \left\| \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\|^2 \phi(\mathbf{X}\mathbf{w}/\sqrt{d})\phi(\mathbf{X}\mathbf{w}/\sqrt{d})^\top \right] \right\| + \|\Theta_2^*\|^2 \\ &\leq \left\| \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\mathbb{1}_{\{\|\phi(\mathbf{X}\mathbf{w}/\sqrt{d})\| \leq B\}} \left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\|^2 \phi(\mathbf{X}\mathbf{w}/\sqrt{d})\phi(\mathbf{X}\mathbf{w}/\sqrt{d})^\top \right] \right\| \\ &\quad + \left\| \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\mathbb{1}_{\{\|\phi(\mathbf{X}\mathbf{w}/\sqrt{d})\| > B\}} \left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\|^2 \phi(\mathbf{X}\mathbf{w}/\sqrt{d})\phi(\mathbf{X}\mathbf{w}/\sqrt{d})^\top \right] \right\| + \|\Theta_2^*\|^2 \\ &\leq B^2 \left\| \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\phi(\mathbf{X}\mathbf{w}/\sqrt{d})\phi(\mathbf{X}\mathbf{w}/\sqrt{d})^\top \right] \right\| \\ &\quad + \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\mathbb{1}_{\{\|\phi(\mathbf{X}\mathbf{w}/\sqrt{d})\| > B\}} \left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\|^4 \right] + \|\Theta_2^*\|^2 \\ &= B^2 \|\Theta_2^*\| + \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\mathbb{1}_{\{\|\phi(\mathbf{X}\mathbf{w}/\sqrt{d})\| > B\}} \left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\|^4 \right] + \|\Theta_2^*\|^2 \end{aligned}$$

$$\begin{aligned}
&\leq B^2 \|\Theta_2^*\| + \|\Theta_2^*\|^2 + \sqrt{\mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\mathbb{1}_{\{\|\phi(\mathbf{X}\mathbf{w}/\sqrt{d})\| > B\}} \right] \cdot \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\|^8 \right]} \\
&\hspace{15em} \text{(Cauchy-Schwarz inequality)} \\
&= B^2 \|\Theta_2^*\| + \|\Theta_2^*\|^2 + \sqrt{\Pr_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\| > B \right] \cdot \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\|^8 \right]}.
\end{aligned}$$

Note that $\left| \left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\| - \left\| \phi(\mathbf{X}\mathbf{w}'/\sqrt{d}) \right\| \right| \lesssim \left\| \mathbf{X}\mathbf{w}/\sqrt{d} - \mathbf{X}\mathbf{w}'/\sqrt{d} \right\| \leq \frac{\|\mathbf{X}\|}{\sqrt{d}} \|\mathbf{w} - \mathbf{w}'\| \lesssim \sqrt{\frac{n}{d}} \|\mathbf{w} - \mathbf{w}'\|$ for all $\mathbf{w}, \mathbf{w}' \in \mathbb{R}^d$. Then by the standard Lipschitz concentration bound for Gaussian variables (see e.g. [Wainwright \[2019\]](#)) we know that for any $s > 0$:

$$\Pr_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\| > M + s \right] \leq e^{-\Omega\left(\frac{s^2}{n/d}\right)},$$

where $M := \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\| \right]$ which can be bounded as

$$\begin{aligned}
M^2 &\leq \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\|^2 \right] \\
&\lesssim \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\|\phi(\mathbf{0}_n)\|^2 + \|\mathbf{X}\mathbf{w}/\sqrt{d}\|^2 \right] \\
&\lesssim n + \frac{n}{d} \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\|\mathbf{w}\|^2 \right] \\
&\lesssim n.
\end{aligned}$$

Thus, letting $\frac{s^2}{n/d} = C \log n$ for a sufficiently large universal constant $C > 0$, we know that with probability at least $1 - n^{-10}$ over $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$,

$$\left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\| \leq M + s \lesssim \sqrt{n} + \sqrt{\frac{n}{d} \log n} \lesssim \sqrt{n}.$$

Hence we pick the threshold $B = C' \sqrt{n}$ which is the upper bound above, where $C' > 0$ is a universal constant.

We can also bound

$$\begin{aligned}
&\mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\|^8 \right] \\
&= \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left(\sum_{i=1}^n \phi(\mathbf{x}_i^\top \mathbf{w}/\sqrt{d})^2 \right)^4 \right] \\
&\lesssim \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left(\sum_{i=1}^n \left(\phi(0)^2 + (\mathbf{x}_i^\top \mathbf{w}/\sqrt{d})^2 \right) \right)^4 \right] \\
&\hspace{15em} (\phi \text{ is Lipschitz \& Cauchy-Schwartz inequality}) \\
&\lesssim \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left(n + \sum_{i=1}^n (\mathbf{x}_i^\top \mathbf{w}/\sqrt{d})^2 \right)^4 \right] \hspace{5em} (|\phi(0)| = O(1)) \\
&\lesssim n^4 + \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left(\sum_{i=1}^n (\mathbf{x}_i^\top \mathbf{w}/\sqrt{d})^2 \right)^4 \right] \hspace{5em} \text{(Jensen's inequality)} \\
&= n^4 + n^4 \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left(\frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i^\top \mathbf{w}/\sqrt{d})^2 \right)^4 \right] \\
&\leq n^4 + n^4 \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i^\top \mathbf{w}/\sqrt{d})^8 \right] \hspace{5em} \text{(Jensen's inequality)}
\end{aligned}$$

$$\begin{aligned}
&= n^4 + n^3 \sum_{i=1}^n \mathbb{E}_{x \sim \mathcal{N}(0, \|\mathbf{x}_i\|^2/d)} [x^8] \\
&\lesssim n^4. \tag{\|\mathbf{x}_i\|^2/d = O(1)}
\end{aligned}$$

Combining all the above, we get

$$\begin{aligned}
&\left\| \mathbb{E}[(\Theta_2^{(r)} - \Theta_2^*)^2] \right\| \\
&\leq B^2 \|\Theta_2^*\| + \|\Theta_2^*\|^2 + \sqrt{\Pr_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\| > B \right] \cdot \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \left[\left\| \phi(\mathbf{X}\mathbf{w}/\sqrt{d}) \right\|^8 \right]} \\
&\lesssim n \|\Theta_2^*\| + \|\Theta_2^*\|^2 + \sqrt{n^{-10} \cdot n^4} \\
&= n \|\Theta_2^*\| + \|\Theta_2^*\|^2 + n^{-3}.
\end{aligned}$$

We will discuss two cases separately.

Case 1: $\vartheta_0 \neq 0$. Recall that in this case Theorem 3.5 assumes $m \gtrsim d^{2+\alpha}$.

Since $\|\Theta_2^*\| = O(n)$, we have $\left\| \mathbb{E}[(\Theta_2^{(r)} - \Theta_2^*)^2] \right\| \lesssim n^2$ which implies

$$\left\| \sum_{r=1}^{m/2} \mathbb{E}[(\Theta_2^{(r)} - \Theta_2^*)^2] \right\| \lesssim mn^2.$$

Applying Proposition 4.1 in Klochkov and Zhivotovskiy [2020], we know that for any $u \gg \max\{n \log m, n\sqrt{m}\} = n\sqrt{m}$,

$$\Pr \left[\left\| \sum_{r=1}^{m/2} (\Theta_2^{(r)} - \Theta_2^*) \right\| > u \right] \lesssim n \cdot \exp \left(-\Omega \left(\min \left\{ \frac{u^2}{mn^2}, \frac{u}{n \log m} \right\} \right) \right).$$

Let $u = m \cdot \frac{n}{d^{1+\frac{\alpha}{3}}}$. We can verify $u \gg n\sqrt{m}$ since $m \gtrsim d^{2+\alpha}$. Then we have

$$\Pr \left[\left\| \sum_{r=1}^{m/2} (\Theta_2^{(r)} - \Theta_2^*) \right\| > m \cdot \frac{n}{d^{1+\frac{\alpha}{3}}} \right] \lesssim n \cdot \exp \left(-\Omega \left(\min \left\{ \frac{m}{d^{2+\frac{2\alpha}{3}}}, \frac{m}{d^{1+\frac{\alpha}{3}} \log m} \right\} \right) \right) \ll 1.$$

Similarly, for the second half of the neurons we also have $\left\| \sum_{r=m/2+1}^m (\Theta_2^{(r)} - \Theta_2^*) \right\| \leq m \cdot \frac{n}{d^{1+\frac{\alpha}{3}}}$ with high probability. Therefore we have with high probability,

$$\|\Theta_2 - \Theta_2^*\| \lesssim \frac{n}{d^{1+\frac{\alpha}{3}}}.$$

Case 2: $\vartheta_0 = 0$. Recall that in this case Theorem 3.5 assumes $m \gtrsim d^{1+\alpha}$.

Since $\|\Theta_2^*\| = \tilde{O}(n/d)$, we have $\left\| \mathbb{E}[(\Theta_2^{(r)} - \Theta_2^*)^2] \right\| \lesssim n \cdot \tilde{O}(n/d) + \tilde{O}((n/d)^2) + n^{-3} = \tilde{O}(n^2/d)$ which implies

$$\left\| \sum_{r=1}^{m/2} \mathbb{E}[(\Theta_2^{(r)} - \Theta_2^*)^2] \right\| \leq \tilde{O}(mn^2/d) \lesssim \frac{mn^2}{d^{1-\frac{\alpha}{10}}}.$$

Applying Proposition 4.1 in Klochkov and Zhivotovskiy [2020], we know that for any $u \gg \max\{n \log m, n\sqrt{m/d^{1-\frac{\alpha}{10}}}\} = n\sqrt{m/d^{1-\frac{\alpha}{10}}}$,

$$\Pr \left[\left\| \sum_{r=1}^{m/2} (\Theta_2^{(r)} - \Theta_2^*) \right\| > u \right] \lesssim n \cdot \exp \left(-\Omega \left(\min \left\{ \frac{u^2}{mn^2/d^{1-\frac{\alpha}{10}}}, \frac{u}{n \log m} \right\} \right) \right).$$

Let $u = m \cdot \frac{n}{d^{1+\frac{\alpha}{3}}}$. We can verify $u \gg n\sqrt{m/d^{1-\frac{\alpha}{10}}}$ since $m \gtrsim d^{1+\alpha}$. Then we have

$$\Pr \left[\left\| \sum_{r=1}^{m/2} (\Theta_2^{(r)} - \Theta_2^*) \right\| > m \cdot \frac{n}{d^{1+\frac{\alpha}{3}}} \right] \lesssim n \cdot \exp \left(-\Omega \left(\min \left\{ \frac{m}{d^{1+0.77\alpha}}, \frac{m}{d^{1+\frac{\alpha}{3}} \log m} \right\} \right) \right) \ll 1.$$

Similarly, for the second half of the neurons we also have $\left\| \sum_{r=m/2+1}^m (\Theta_2^{(r)} - \Theta_2^*) \right\| \leq m \cdot \frac{n}{d^{1+\frac{\alpha}{3}}}$ with high probability. Therefore we have with high probability,

$$\|\Theta_2 - \Theta_2^*\| \lesssim \frac{n}{d^{1+\frac{\alpha}{3}}}.$$

The proof is completed. \square

Combining Propositions D.9 and D.10 directly gives Proposition D.8.

D.5.2 Agreement on Training Data

To prove the agreement between f_t^2 and $f_t^{\text{lin}2}$ on training data for all $t \leq T = c \cdot \frac{d \log d}{\eta_2}$, we still apply Theorem C.2. This case is much easier than training the first layer (Appendix D.3.2), since the Jacobian for the second layer does not change during training, and thus Proposition D.8 already verifies Assumption C.1. Therefore we can directly instantiate Theorem C.2 with $\epsilon = C \frac{n}{d^{1+\frac{\alpha}{3}}}$ (for a sufficiently large constant C) and $R = \sqrt{d \log d}$, which gives (notice that the choice of η_2 in Theorem 3.5 also satisfies the condition in Theorem C.2)

$$\sqrt{\sum_{i=1}^n (f_t^2(\mathbf{x}_i) - f_t^{\text{lin}2}(\mathbf{x}_i))^2} \lesssim \frac{\eta_2 t \epsilon}{\sqrt{n}} \lesssim \frac{d \log d \cdot \frac{n}{d^{1+\frac{\alpha}{3}}}}{\sqrt{n}} = \frac{\sqrt{n} \log d}{d^{\frac{\alpha}{3}}} \ll \frac{\sqrt{n}}{d^{\frac{\alpha}{4}}},$$

i.e.,

$$\frac{1}{n} \sum_{i=1}^n (f_t^2(\mathbf{x}_i) - f_t^{\text{lin}2}(\mathbf{x}_i))^2 \leq d^{-\frac{\alpha}{2}}.$$

This proves the first part in Theorem 3.5.

Note that Theorem C.2 also tells us $\|\mathbf{v}(t) - \mathbf{v}(0)\| \leq \sqrt{d \log d}$ and $\|\gamma(t)\| \leq \sqrt{d \log d}$, which will be useful for proving the guarantee on the distribution \mathcal{D} .

D.5.3 Agreement on Distribution

Now we prove the second part in Theorem 3.5, which is the agreement between f_t^2 and $f_t^{\text{lin}2}$ on the distribution \mathcal{D} . The proof is similar to the case of training the first layer (Appendix D.3.3), but our case here is again simpler. In particular, we do not need to define an auxiliary model anymore because $f(\mathbf{x}; \mathbf{W}, \mathbf{v})$ is already linear in the parameters \mathbf{v} . Now that $f_t^2 - f_t^{\text{lin}2}$ is a linear model (in some feature space) with bounded parameters, we can bound the Rademacher complexity of the linear function class it belongs to, similar to Appendix D.3.3. Similar to (37), we can bound the Rademacher complexity by

$$\frac{\sqrt{d \log d}}{n} \sqrt{\text{Tr}[\Theta_2(\mathbf{W})] + \text{Tr}[\Theta^{\text{lin}2}]}$$

Next we bound the above two traces. First, we have

$$\begin{aligned} [\Theta_2(\mathbf{W})]_{i,i} &= \frac{1}{m} \sum_{r=1}^m \phi(\mathbf{x}_i^\top \mathbf{w}_r / \sqrt{d})^2 \lesssim \frac{1}{m} \sum_{r=1}^m \left(\phi(0)^2 + (\mathbf{x}_i^\top \mathbf{w}_r / \sqrt{d})^2 \right) \\ &= 1 + \frac{1}{dm} \sum_{r=1}^m (\mathbf{x}_i^\top \mathbf{w}_r)^2 \lesssim 1 + \frac{1}{dm} (dm + \log n) \lesssim 1 \end{aligned}$$

with high probability for all $i \in [n]$ together. Here we have used the standard tail bound for χ^2 random variables and a union bound over $i \in [n]$. Hence we have $\text{Tr}[\Theta_2(\mathbf{W})] \lesssim n$. For the second trace, we have

$$\text{Tr}[\Theta^{\text{lin}2}] = \sum_{i=1}^n \left(\zeta^{-2} \frac{\|\mathbf{x}_i\|^2}{d} + \frac{\nu^2}{2d} + [\mathbf{q}]_i^2 \right) \lesssim n$$

with high probability. Therefore we can bound the Rademacher complexity by $\sqrt{\frac{d \log d}{n}}$. Then we can conclude the agreement guarantee on the distribution \mathcal{D} , i.e., for all $t \leq T$ simultaneously,

$$\begin{aligned} & \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \left[\min \left\{ (f_t^2(\mathbf{x}) - f_t^{\text{lin}2}(\mathbf{x}))^2, 1 \right\} \right] \\ & \leq \frac{1}{n} \sum_{i=1}^n \min \left\{ (f_t^2(\mathbf{x}_i) - f_t^{\text{lin}2}(\mathbf{x}_i))^2, 1 \right\} + O \left(\sqrt{\frac{d \log d}{n}} \right) + O \left(\frac{1}{\sqrt{n}} \right) \\ & \lesssim d^{-\frac{\alpha}{2}} + \sqrt{\frac{d \log d}{d^{1+\alpha}}} \quad (n \gtrsim d^{1+\alpha}) \\ & \lesssim d^{-\frac{\alpha}{2}}. \end{aligned}$$

This completes the proof of Theorem 3.5.

D.6 Proof of Theorem D.1 (Training Both Layers)

The proof for training both layers follows the same ideas in the proofs for training the first layer only and the second layer only. In fact, most technical components needed in the proof were already developed in the previous proofs. The only new component is a Jacobian perturbation bound for the case of training both layers, Lemma D.12 (analog of Lemma D.5 for training the first layer).

As before, we proceed in three steps.

D.6.1 The NTK at Initialization

Proposition D.11. *With high probability over the random initialization $(\mathbf{W}(0), \mathbf{v}(0))$ and the training data \mathbf{X} , we have*

$$\|\Theta(\mathbf{W}(0), \mathbf{v}(0)) - \Theta^{\text{lin}}\| \lesssim \frac{n}{d^{1+\frac{\alpha}{3}}}.$$

Proof. This is a direct corollary of Propositions D.2 and D.8, given that $\Theta(\mathbf{W}(0), \mathbf{v}(0)) = \Theta_1(\mathbf{W}(0), \mathbf{v}(0)) + \Theta(\mathbf{W}(0))$ ((17)) and $\Theta^{\text{lin}} = \Theta^{\text{lin}1} + \Theta^{\text{lin}2}$ ((20)). \square

D.6.2 Agreement on Training Data

The proof for the agreement on training data is similar to the case of training the first layer only (Appendix D.3.2). We will again apply Theorem C.2. For this we need a new Jacobian perturbation lemma to replace Lemma D.5, since both layers are allowed to move now.

Lemma D.12 (Jacobian perturbation for both layers). *If ϕ is a smooth activation as in Assumption 3.2, then with high probability over the training data \mathbf{X} , we have*

$$\|\mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0))\| \lesssim \sqrt{\frac{n}{md}} \|\mathbf{W} - \mathbf{W}(0)\|_F + \sqrt{\frac{n}{m}} \|\mathbf{v} - \mathbf{v}(0)\|, \quad \forall \mathbf{W}, \mathbf{v}. \quad (42)$$

If ϕ is a piece-wise linear activation as in Assumption 3.2, then with high probability over the random initialization $\mathbf{W}(0)$ and the training data \mathbf{X} , we have

$$\begin{aligned} \|\mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0))\| & \lesssim \sqrt{\frac{n}{d}} \left(\frac{\|\mathbf{W} - \mathbf{W}(0)\|^{1/3}}{m^{1/6}} + \left(\frac{\log n}{m} \right)^{1/4} \right) + \sqrt{\frac{n}{md}} \|\mathbf{v} - \mathbf{v}(0)\|, \\ & \quad \forall \mathbf{W}, \mathbf{v}. \end{aligned} \quad (43)$$

Furthermore, with high probability over the training data \mathbf{X} , we have

$$\|\mathbf{J}_2(\mathbf{W}) - \mathbf{J}_2(\widetilde{\mathbf{W}})\| \lesssim \sqrt{\frac{n}{md}} \|\mathbf{W} - \widetilde{\mathbf{W}}\|_F, \quad \forall \mathbf{W}, \widetilde{\mathbf{W}}. \quad (44)$$

Proof. We will be conditioned on \mathbf{X} and on the high-probability events in Claim 3.1.

We first consider the first-layer Jacobian. By the definition of $\mathbf{J}_1(\mathbf{W}, \mathbf{v})$ in (15), we have

$$\begin{aligned} & (\mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0)))(\mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0)))^\top \\ &= \frac{1}{md} \left(\left(\phi' \left(\mathbf{X} \mathbf{W}^\top / \sqrt{d} \right) \text{diag}(\mathbf{v}) - \phi' \left(\mathbf{X} \mathbf{W}(0)^\top / \sqrt{d} \right) \text{diag}(\mathbf{v}(0)) \right) \right. \\ & \quad \left. \cdot \left(\phi' \left(\mathbf{X} \mathbf{W}^\top / \sqrt{d} \right) \text{diag}(\mathbf{v}) - \phi' \left(\mathbf{X} \mathbf{W}(0)^\top / \sqrt{d} \right) \text{diag}(\mathbf{v}(0)) \right)^\top \right) \odot (\mathbf{X} \mathbf{X}^\top). \end{aligned} \quad (45)$$

Then if ϕ is a smooth activation, we have with high probability,

$$\begin{aligned} & \|\mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0))\|^2 \\ & \leq \frac{1}{md} \left\| \phi' \left(\mathbf{X} \mathbf{W}^\top / \sqrt{d} \right) \text{diag}(\mathbf{v}) - \phi' \left(\mathbf{X} \mathbf{W}(0)^\top / \sqrt{d} \right) \text{diag}(\mathbf{v}(0)) \right\|^2 \cdot \max_{i \in [n]} \|\mathbf{x}_i\|^2 \\ & \quad \text{((45) and Lemma B.3)} \\ & \lesssim \frac{1}{m} \left\| \phi' \left(\mathbf{X} \mathbf{W}^\top / \sqrt{d} \right) \text{diag}(\mathbf{v}) - \phi' \left(\mathbf{X} \mathbf{W}(0)^\top / \sqrt{d} \right) \text{diag}(\mathbf{v}(0)) \right\|^2 \quad \text{(Claim 3.1)} \\ & \lesssim \frac{1}{m} \left\| \phi' \left(\mathbf{X} \mathbf{W}^\top / \sqrt{d} \right) \text{diag}(\mathbf{v}(0)) - \phi' \left(\mathbf{X} \mathbf{W}(0)^\top / \sqrt{d} \right) \text{diag}(\mathbf{v}(0)) \right\|^2 \\ & \quad + \frac{1}{m} \left\| \phi' \left(\mathbf{X} \mathbf{W}^\top / \sqrt{d} \right) \text{diag}(\mathbf{v}) - \phi' \left(\mathbf{X} \mathbf{W}^\top / \sqrt{d} \right) \text{diag}(\mathbf{v}(0)) \right\|^2 \\ & \leq \frac{1}{m} \left\| \phi' \left(\mathbf{X} \mathbf{W}^\top / \sqrt{d} \right) - \phi' \left(\mathbf{X} \mathbf{W}(0)^\top / \sqrt{d} \right) \right\|_F^2 \cdot \|\text{diag}(\mathbf{v}(0))\|^2 \\ & \quad + \frac{1}{m} \left\| \phi' \left(\mathbf{X} \mathbf{W}^\top / \sqrt{d} \right) \text{diag}(\mathbf{v} - \mathbf{v}(0)) \right\|_F^2 \\ & \leq \frac{n}{md} \|\mathbf{W} - \mathbf{W}(0)\|_F^2 + \frac{1}{m} \left\| \phi' \left(\mathbf{X} \mathbf{W}^\top / \sqrt{d} \right) \text{diag}(\mathbf{v} - \mathbf{v}(0)) \right\|_F^2 \\ & \quad \text{(using the proof of Lemma D.5, and } \|\text{diag}(\mathbf{v}(0))\| = 1) \\ & \leq \frac{n}{md} \|\mathbf{W} - \mathbf{W}(0)\|_F^2 + \frac{n}{m} \|\mathbf{v} - \mathbf{v}(0)\|^2. \quad (\phi' \text{ is bounded}) \end{aligned}$$

This proves (42).

If ϕ is a piece-wise linear activation, then with high probability,

$$\begin{aligned} & \|\mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0))\|^2 \\ & \leq \frac{1}{md} \|\mathbf{X} \mathbf{X}^\top\| \cdot \max_{i \in [n]} \left\| \text{diag}(\mathbf{v}) \phi'(\mathbf{W} \mathbf{x}_i / \sqrt{d}) - \text{diag}(\mathbf{v}(0)) \phi'(\mathbf{W}(0) \mathbf{x}_i / \sqrt{d}) \right\|^2 \\ & \quad \text{((45) and Lemma B.3)} \\ & \leq \frac{n}{md} \cdot \max_{i \in [n]} \left\| \text{diag}(\mathbf{v}) \phi'(\mathbf{W} \mathbf{x}_i / \sqrt{d}) - \text{diag}(\mathbf{v}(0)) \phi'(\mathbf{W}(0) \mathbf{x}_i / \sqrt{d}) \right\|^2 \quad \text{(Claim 3.1)} \\ & \lesssim \frac{n}{md} \cdot \max_{i \in [n]} \left\| \text{diag}(\mathbf{v}(0)) \phi'(\mathbf{W} \mathbf{x}_i / \sqrt{d}) - \text{diag}(\mathbf{v}(0)) \phi'(\mathbf{W}(0) \mathbf{x}_i / \sqrt{d}) \right\|^2 \\ & \quad + \frac{n}{md} \cdot \max_{i \in [n]} \left\| \text{diag}(\mathbf{v}) \phi'(\mathbf{W} \mathbf{x}_i / \sqrt{d}) - \text{diag}(\mathbf{v}(0)) \phi'(\mathbf{W} \mathbf{x}_i / \sqrt{d}) \right\|^2 \\ & \leq \frac{n}{md} \cdot \max_{i \in [n]} \left\| \phi'(\mathbf{W} \mathbf{x}_i / \sqrt{d}) - \phi'(\mathbf{W}(0) \mathbf{x}_i / \sqrt{d}) \right\|^2 + \frac{n}{md} \cdot \max_{i \in [n]} \left\| \text{diag}(\mathbf{v} - \mathbf{v}(0)) \phi'(\mathbf{W} \mathbf{x}_i / \sqrt{d}) \right\|^2 \\ & \quad (\|\text{diag}(\mathbf{v}(0))\| = 1) \\ & \lesssim \frac{n}{d} \left(\frac{\|\mathbf{W} - \mathbf{W}(0)\|^{2/3}}{m^{1/3}} + \sqrt{\frac{\log n}{m}} \right) + \frac{n}{md} \|\mathbf{v} - \mathbf{v}(0)\|^2. \\ & \quad \text{(using the proof of Lemma D.5, and } \phi' \text{ is bounded)} \end{aligned}$$

This proves (43).

For the second-layer Jacobian, we have with high probability,

$$\begin{aligned}
\left\| \mathbf{J}_2(\mathbf{W}) - \mathbf{J}_2(\widetilde{\mathbf{W}}) \right\| &= \frac{1}{\sqrt{m}} \left\| \phi(\mathbf{X}\mathbf{W}^\top/\sqrt{d}) - \phi(\mathbf{X}\widetilde{\mathbf{W}}^\top/\sqrt{d}) \right\| \\
&\leq \frac{1}{\sqrt{m}} \left\| \mathbf{X}(\mathbf{W} - \widetilde{\mathbf{W}})^\top/\sqrt{d} \right\|_F \quad (\phi' \text{ is bounded}) \\
&\leq \frac{\|\mathbf{X}\|}{\sqrt{md}} \left\| \mathbf{W} - \widetilde{\mathbf{W}} \right\|_F \\
&\leq \sqrt{\frac{n}{md}} \left\| \mathbf{W} - \widetilde{\mathbf{W}} \right\|_F,
\end{aligned}$$

completing the proof of (44). \square

Based on Lemma D.12, we can now verify Assumption C.1 for the case of training both layers:

Lemma D.13. *Let $R = \sqrt{d \log d}$. With high probability over the random initialization and the training data, for all (\mathbf{W}, \mathbf{v}) and $(\widetilde{\mathbf{W}}, \tilde{\mathbf{v}})$ such that $\|\mathbf{W} - \mathbf{W}(0)\|_F \leq R$, $\|\widetilde{\mathbf{W}} - \mathbf{W}(0)\|_F \leq R$, $\|\mathbf{v} - \mathbf{v}(0)\| \leq R$ and $\|\tilde{\mathbf{v}} - \mathbf{v}(0)\| \leq R$, we have*

$$\left\| \mathbf{J}(\mathbf{W}, \mathbf{v}) \mathbf{J}(\widetilde{\mathbf{W}}, \tilde{\mathbf{v}})^\top - \Theta^{\text{lin}} \right\| \lesssim \frac{n}{d^{1+\frac{\alpha}{3}}}.$$

Proof. This proof is conditioned on all the high-probability events we have shown.

Now consider (\mathbf{W}, \mathbf{v}) and $(\widetilde{\mathbf{W}}, \tilde{\mathbf{v}})$ which satisfy the conditions stated in the lemma.

If ϕ is a smooth activation, from Lemma D.12 we know

$$\begin{aligned}
\left\| \mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0)) \right\| &\lesssim \sqrt{\frac{n}{md}} \left\| \mathbf{W} - \mathbf{W}(0) \right\|_F + \sqrt{\frac{n}{m}} \left\| \mathbf{v} - \mathbf{v}(0) \right\| \\
&\leq \sqrt{\frac{n}{md}} \cdot \sqrt{d \log d} + \sqrt{\frac{n}{m}} \cdot \sqrt{d \log d} \\
&\lesssim \sqrt{\frac{nd \log d}{m}} \\
&\lesssim \sqrt{\frac{n \log d}{d^{1+\alpha}}} \\
&\ll \sqrt{\frac{n}{d^{1+\frac{2\alpha}{3}}}},
\end{aligned}$$

where we have used $m \gtrsim d^{2+\alpha}$. If ϕ is a piece-wise linear activation, from Lemma D.12 we have

$$\begin{aligned}
\left\| \mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0)) \right\| &\lesssim \sqrt{\frac{n}{d}} \left(\frac{\|\mathbf{W} - \mathbf{W}(0)\|^{1/3}}{m^{1/6}} + \left(\frac{\log n}{m} \right)^{1/4} \right) + \sqrt{\frac{n}{md}} \left\| \mathbf{v} - \mathbf{v}(0) \right\| \\
&\leq \sqrt{\frac{n}{d}} \left(\frac{(d \log d)^{1/6}}{m^{1/6}} + \left(\frac{\log n}{m} \right)^{1/4} \right) + \sqrt{\frac{n \log d}{m}} \\
&\lesssim \sqrt{\frac{n}{d}} \cdot \frac{(d \log d)^{1/6}}{d^{1/3+\alpha/6}} + \sqrt{\frac{n \log d}{d^{2+\alpha}}} \\
&\ll \frac{\sqrt{n}}{d^{\frac{2}{3}}}.
\end{aligned}$$

Hence in either case have $\left\| \mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0)) \right\| \leq \frac{\sqrt{n}}{d^{\frac{1}{2}+\frac{\alpha}{3}}}$. Similarly, we have

$$\left\| \mathbf{J}_1(\widetilde{\mathbf{W}}, \tilde{\mathbf{v}}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0)) \right\| \leq \frac{\sqrt{n}}{d^{\frac{1}{2}+\frac{\alpha}{3}}}.$$

Also, we know from Proposition D.2 that $\|\mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0))\| \lesssim \sqrt{\frac{n}{d}}$. It follows that $\|\mathbf{J}_1(\mathbf{W}, \mathbf{v})\| \lesssim \sqrt{\frac{n}{d}}$ and $\|\mathbf{J}_1(\widetilde{\mathbf{W}}, \widetilde{\mathbf{v}})\| \lesssim \sqrt{\frac{n}{d}}$. Then we have

$$\begin{aligned} & \left\| \mathbf{J}_1(\mathbf{W}, \mathbf{v})\mathbf{J}_1(\widetilde{\mathbf{W}}, \widetilde{\mathbf{v}})^\top - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0))\mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0))^\top \right\| \\ & \leq \|\mathbf{J}_1(\mathbf{W}, \mathbf{v})\| \cdot \left\| \mathbf{J}_1(\widetilde{\mathbf{W}}, \widetilde{\mathbf{v}}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0)) \right\| + \|\mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0))\| \cdot \|\mathbf{J}_1(\mathbf{W}, \mathbf{v}) - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0))\| \\ & \lesssim \sqrt{\frac{n}{d}} \cdot \frac{\sqrt{n}}{d^{\frac{1}{2} + \frac{\alpha}{3}}} + \sqrt{\frac{n}{d}} \cdot \frac{\sqrt{n}}{d^{\frac{1}{2} + \frac{\alpha}{3}}} \\ & \lesssim \frac{n}{d^{1 + \frac{\alpha}{3}}}. \end{aligned}$$

Next we look at the second-layer Jacobian. From Lemma D.12 we know $\|\mathbf{J}_2(\mathbf{W}) - \mathbf{J}_2(\mathbf{W}(0))\| \lesssim \sqrt{\frac{n}{md}} \cdot \sqrt{d \log d} \lesssim \sqrt{\frac{n \log d}{d^{2 + \alpha}}} \ll \frac{\sqrt{n}}{d^{1 + \frac{\alpha}{3}}}$. Similarly we have $\left\| \mathbf{J}_2(\widetilde{\mathbf{W}}) - \mathbf{J}_2(\mathbf{W}(0)) \right\| \ll \frac{\sqrt{n}}{d^{1 + \frac{\alpha}{3}}}$. Also, from Proposition D.8 we know $\|\mathbf{J}_2(\mathbf{W}(0))\| \lesssim \sqrt{n}$, which implies $\|\mathbf{J}_2(\mathbf{W})\| \lesssim \sqrt{n}$ and $\|\mathbf{J}_2(\widetilde{\mathbf{W}})\| \lesssim \sqrt{n}$. It follows that

$$\begin{aligned} & \left\| \mathbf{J}_2(\mathbf{W})\mathbf{J}_2(\widetilde{\mathbf{W}})^\top - \mathbf{J}_2(\mathbf{W}(0))\mathbf{J}_2(\mathbf{W}(0))^\top \right\| \\ & \leq \|\mathbf{J}_2(\mathbf{W})\| \cdot \left\| \mathbf{J}_2(\widetilde{\mathbf{W}}) - \mathbf{J}_2(\mathbf{W}(0)) \right\| + \|\mathbf{J}_2(\mathbf{W}(0))\| \cdot \|\mathbf{J}_2(\mathbf{W}) - \mathbf{J}_2(\mathbf{W}(0))\| \\ & \lesssim \sqrt{n} \cdot \frac{\sqrt{n}}{d^{1 + \frac{\alpha}{3}}} + \sqrt{n} \cdot \frac{\sqrt{n}}{d^{1 + \frac{\alpha}{3}}} \\ & \lesssim \frac{n}{d^{1 + \frac{\alpha}{3}}}. \end{aligned}$$

Combining the above arguments for two layers, we obtain

$$\begin{aligned} & \left\| \mathbf{J}(\mathbf{W}, \mathbf{v})\mathbf{J}(\widetilde{\mathbf{W}}, \widetilde{\mathbf{v}})^\top - \mathbf{J}(\mathbf{W}(0), \mathbf{v}(0))\mathbf{J}(\mathbf{W}(0), \mathbf{v}(0))^\top \right\| \\ & = \left\| \mathbf{J}_1(\mathbf{W}, \mathbf{v})\mathbf{J}_1(\widetilde{\mathbf{W}}, \widetilde{\mathbf{v}})^\top + \mathbf{J}_2(\mathbf{W})\mathbf{J}_2(\widetilde{\mathbf{W}})^\top \right. \\ & \quad \left. - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0))\mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0))^\top - \mathbf{J}_2(\mathbf{W}(0))\mathbf{J}_2(\mathbf{W}(0))^\top \right\| \\ & \leq \left\| \mathbf{J}_1(\mathbf{W}, \mathbf{v})\mathbf{J}_1(\widetilde{\mathbf{W}}, \widetilde{\mathbf{v}})^\top - \mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0))\mathbf{J}_1(\mathbf{W}(0), \mathbf{v}(0))^\top \right\| \\ & \quad + \left\| \mathbf{J}_2(\mathbf{W})\mathbf{J}_2(\widetilde{\mathbf{W}})^\top - \mathbf{J}_2(\mathbf{W}(0))\mathbf{J}_2(\mathbf{W}(0))^\top \right\| \\ & \lesssim \frac{n}{d^{1 + \frac{\alpha}{3}}} + \frac{n}{d^{1 + \frac{\alpha}{3}}} \\ & \lesssim \frac{n}{d^{1 + \frac{\alpha}{3}}}. \end{aligned}$$

Combining the above inequality with Proposition D.11, the proof is finished. \square

Finally, we can apply Theorem C.2 with $R = \sqrt{d \log d}$ and $\epsilon = O\left(\frac{n}{d^{1 + \frac{\alpha}{3}}}\right)$, and obtain that for all $t \leq T$:

$$\sqrt{\sum_{i=1}^n (f_t(\mathbf{x}_i) - f_t^{\text{lin}}(\mathbf{x}_i))^2} \lesssim \frac{\eta t \epsilon}{\sqrt{n}} \lesssim \frac{d \log d \cdot \frac{n}{d^{1 + \frac{\alpha}{3}}}}{\sqrt{n}} = \frac{\sqrt{n} \log d}{d^{\frac{\alpha}{3}}} \ll \frac{\sqrt{n}}{d^{\frac{\alpha}{4}}},$$

i.e.,

$$\frac{1}{n} \sum_{i=1}^n (f_t(\mathbf{x}_i) - f_t^{\text{lin}}(\mathbf{x}_i))^2 \leq d^{-\frac{\alpha}{2}}.$$

This proves the first part in Theorem D.1.

Note that Theorem C.2 also tells us $\|\mathbf{W}(t) - \mathbf{W}(0)\| \leq \sqrt{d \log d}$, $\|\mathbf{v}(t) - \mathbf{v}(0)\| \leq \sqrt{d \log d}$ and $\|\delta(t)\| \leq \sqrt{d \log d}$, which will be useful for proving the guarantee on the distribution \mathcal{D} .

D.6.3 Agreement on Distribution

The proof for the second part of Theorem D.1 is basically identical to the case of training the first layer (Appendix D.3.3), so we will only sketch the differences here to avoid repetition.

Recall that in Appendix D.3.3 we define an auxiliary model which is the first-order approximation of the network around initialization. Here since we are training both layers, we need to modify the definition of the auxiliary model to incorporate deviation from initialization in both layers:

$$f^{\text{aux}}(\mathbf{x}; \mathbf{W}, \mathbf{v}) := \langle \mathbf{W} - \mathbf{W}(0), \nabla_{\mathbf{W}} f(\mathbf{x}; \mathbf{W}(0), \mathbf{v}(0)) \rangle + \langle \mathbf{v} - \mathbf{v}(0), \nabla_{\mathbf{v}} f(\mathbf{x}; \mathbf{W}(0), \mathbf{v}(0)) \rangle.$$

Then we denote $f_t^{\text{aux}}(\mathbf{x}) := f^{\text{aux}}(\mathbf{x}; \mathbf{W}(t), \mathbf{v}(t))$.

There are two more minor changes to Appendix D.3.3:

1. When proving f_t and f_t^{aux} are close on both training data and imaginary test data, we need to bound a Jacobian perturbation. In Appendix D.3.3 this step is done using Lemma D.5. Now we simply need to use Lemma D.12 instead and note that $\|\mathbf{W}(t) - \mathbf{W}(0)\| \leq \sqrt{d \log d}$ and $\|\mathbf{v}(t) - \mathbf{v}(0)\| \leq \sqrt{d \log d}$.
2. Instead of (37), the empirical Rademacher complexity of the function class that each $f_t^{\text{aux}} - f_t^{\text{lin}}$ lies in will be

$$\begin{aligned} & \frac{\sqrt{d \log d}}{n} \sqrt{\text{Tr}[\Theta(\mathbf{W}(0), \mathbf{v}(0))] + \text{Tr}[\Theta^{\text{lin}}]} \\ &= \frac{\sqrt{d \log d}}{n} \sqrt{\text{Tr}[\Theta_1(\mathbf{W}(0), \mathbf{v}(0))] + \text{Tr}[\Theta_2(\mathbf{W}(0))] + \text{Tr}[\Theta^{\text{lin}1}] + \text{Tr}[\Theta^{\text{lin}2}]}. \end{aligned}$$

In Appendices D.3.3 and D.5.3, we have shown that the above 4 traces are all $O(n)$ with high probability. Hence we get the same Rademacher complexity bound as before.

Modulo these differences, the proof proceeds the same as Appendix D.3.3. Therefore we conclude the proof of Theorem D.1.

D.7 Proof of Claim 3.1

Proof of Claim 3.1. According to Assumption 3.1, we have $\mathbf{x}_i = \Sigma^{1/2} \bar{\mathbf{x}}_i$ where $\mathbb{E}[\bar{\mathbf{x}}_i] = \mathbf{0}$, $\mathbb{E}[\bar{\mathbf{x}}_i \bar{\mathbf{x}}_i^\top] = \mathbf{I}$, and $\bar{\mathbf{x}}_i$'s entries are independent and $O(1)$ -subgaussian.

By Hanson-Wright inequality (specifically, Theorem 2.1 in Rudelson and Vershynin [2013]), we have for any $t \geq 0$,

$$\Pr \left[\left| \|\Sigma^{1/2} \bar{\mathbf{x}}_i\| - \|\Sigma^{1/2}\|_F \right| > t \right] \leq 2 \exp \left(-\Omega \left(\frac{t^2}{\|\Sigma^{1/2}\|^2} \right) \right),$$

i.e.,

$$\Pr \left[\left| \|\mathbf{x}_i\| - \sqrt{d} \right| > t \right] \leq 2 \exp \left(-\Omega(t^2) \right).$$

Let $t = C\sqrt{\log n}$ for a sufficiently large constant $C > 0$. Taking a union bound over all $i \in [n]$, we obtain that with high probability, $\|\mathbf{x}_i\| = \sqrt{d} \pm O(\sqrt{\log n})$ for all $i \in [n]$ simultaneously. This proves the first property in Claim 3.1.

For $i \neq j$, we have $\langle \mathbf{x}_i, \mathbf{x}_j \rangle = \bar{\mathbf{x}}_i^\top \Sigma \bar{\mathbf{x}}_j$. Conditioned on $\bar{\mathbf{x}}_j$, we know that $\bar{\mathbf{x}}_i^\top \Sigma \bar{\mathbf{x}}_j$ is zero-mean and $O(\|\Sigma \bar{\mathbf{x}}_j\|^2)$ -subgaussian, which means for any $t \geq 0$,

$$\Pr \left[\left| \bar{\mathbf{x}}_i^\top \Sigma \bar{\mathbf{x}}_j \right| > t \mid \bar{\mathbf{x}}_j \right] \leq 2 \exp \left(-\frac{t^2}{\|\Sigma \bar{\mathbf{x}}_j\|^2} \right).$$

Since we have shown that $\|\Sigma \bar{\mathbf{x}}_j\|^2 \lesssim \|\mathbf{x}_j\|^2 \lesssim \sqrt{d} + \sqrt{\log n} \lesssim \sqrt{d}$ with probability at least $1 - n^{-10}$, we have

$$\Pr \left[\left| \bar{\mathbf{x}}_i^\top \Sigma \bar{\mathbf{x}}_j \right| > t \right] \leq n^{-10} + 2 \exp \left(-\Omega \left(\frac{t^2}{d} \right) \right).$$

Then we can take $t = C\sqrt{d\log n}$ and apply a union bound over i, j , which gives $|\langle \mathbf{x}_i, \mathbf{x}_j \rangle| \lesssim \sqrt{d\log n}$ for all $i \neq j$ with high probability. This completes the proof of the second statement in Claim 3.1.

Finally, for $\mathbf{X}\mathbf{X}^\top$, we can use standard covariance concentration (see, e.g., Lemma A.6 in Du et al. [2020]) to obtain $0.9\Sigma \preceq \frac{1}{n}\mathbf{X}^\top\mathbf{X} \preceq 1.1\Sigma$ with high probability. This implies $\|\mathbf{X}\mathbf{X}^\top\| = \|\mathbf{X}^\top\mathbf{X}\| = \Theta(n)$. \square

E Omitted Details in Section 4

Proof of Proposition 4.1. For an input $\mathbf{x} \in \mathbb{R}^d$ and an index $k \in [d]$, we let $[\mathbf{x}]_{k:k+q}$ be the patch of size q starting from index k , i.e., $[\mathbf{x}]_{k:k+q} := [[\mathbf{x}]_k, [\mathbf{x}]_{k+1}, \dots, [\mathbf{x}]_{k+q-1}]^\top \in \mathbb{R}^q$.

For two datapoints \mathbf{x}_i and \mathbf{x}_j ($i, j \in [n]$) and a location $k \in [d]$, we define

$$\rho_{i,j,k} := \frac{\langle [\mathbf{x}_i]_{k:k+q}, [\mathbf{x}_j]_{k:k+q} \rangle}{q}$$

which is a local correlation between \mathbf{x}_i and \mathbf{x}_j .

Now we calculate the infinite-width NTK matrix Θ_{CNN} , which is also the expectation of a finite-width NTK matrix with respect to the randomly initialized weights (\mathbf{W}, \mathbf{V}) . We divide the NTK matrix into two components corresponding to two layers: $\Theta_{\text{CNN}} = \Theta_{\text{CNN}}^{(1)} + \Theta_{\text{CNN}}^{(2)}$, and consider the two layers separately.

Step 1: the second-layer NTK. Since the CNN model (12) is linear in the second layer weights, it is easy to derive the formula for the second-layer NTK:

$$\begin{aligned} [\Theta_{\text{CNN}}^{(2)}]_{i,j} &= \frac{1}{d} \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}_q, \mathbf{I}_q)} [\phi(\mathbf{w} * \mathbf{x}_i / \sqrt{q})^\top \phi(\mathbf{w} * \mathbf{x}_j / \sqrt{q})] \\ &= \frac{1}{d} \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}_q, \mathbf{I}_q)} \left[\sum_{k=1}^d \phi([\mathbf{w} * \mathbf{x}_i]_k / \sqrt{q}) \phi([\mathbf{w} * \mathbf{x}_j]_k / \sqrt{q}) \right] \\ &= \frac{1}{d} \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}_q, \mathbf{I}_q)} \left[\sum_{k=1}^d \phi\left(\frac{\langle \mathbf{w}, [\mathbf{x}_i]_{k:k+q} \rangle}{\sqrt{q}}\right) \phi\left(\frac{\langle \mathbf{w}, [\mathbf{x}_j]_{k:k+q} \rangle}{\sqrt{q}}\right) \right] \\ &= \frac{1}{d} \sum_{k=1}^d P(\rho_{i,j,k}), \end{aligned}$$

where

$$P(\rho) := \mathbb{E}_{(z_1, z_2) \sim \mathcal{N}(\mathbf{0}, \mathbf{\Lambda})} [\phi(z_1)\phi(z_2)], \text{ where } \mathbf{\Lambda} = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}, \quad |\rho| \leq 1.$$

Note that we have used the property $\|[\mathbf{x}_j]_{k:k+q}\| = \|[\mathbf{x}_j]_{k:k+q}\| = \sqrt{q}$ since the data are from the hypercube $\{\pm 1\}^d$.

For $i \neq j$, we can do a Taylor expansion of P around 0: $P(\rho) = \zeta^2 \rho \pm O(|\rho|^3)$. Here since $\phi = \text{erf}$ is an odd function, all the even-order terms in the expansion vanish. Therefore we have

$$[\Theta_{\text{CNN}}^{(2)}]_{i,j} = \frac{1}{d} \sum_{k=1}^d (\zeta^2 \rho_{i,j,k} \pm O(|\rho_{i,j,k}|^3)) = \frac{1}{d} \zeta^2 \mathbf{x}_i^\top \mathbf{x}_j \pm \frac{1}{d} \sum_{k=1}^d O(|\rho_{i,j,k}|^3).$$

Next we bound the error term $\frac{1}{d} \sum_{k=1}^d |\rho_{i,j,k}|^3$ for all $i \neq j$. For each i, j, k ($i \neq j$), since $\mathbf{x}_i, \mathbf{x}_j \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(\{\pm 1\}^d)$, by Hoeffding's inequality we know that with probability $1 - \delta$, we have $|\rho_{i,j,k}| \lesssim \sqrt{\frac{\log \frac{1}{\delta}}{q}}$. Taking a union bound, we know that with high probability, for all i, j, k ($i \neq j$)

we have $|\rho_{i,j,k}| = \tilde{O}(q^{-1/2})$. Now we will be conditioned on this happening. Then we write

$$\sum_{k=1}^d |\rho_{i,j,k}|^3 = \sum_{k=1}^q |\rho_{i,j,k}|^3 + \sum_{k=q+1}^{2q} |\rho_{i,j,k}|^3 + \dots,$$

i.e., we divide the sum into $\lceil d/q \rceil$ groups each containing no more than q terms. By the definition of $\rho_{i,j,k}$, it is easy to see that the groups are independent. Also, we have shown that the sum in each group is at most $q \cdot \tilde{O}(q^{-3/2}) = \tilde{O}(q^{-1/2})$. Therefore, using another Hoeffding's inequality among the groups, and applying a union bound over all i, j , we know that with high probability for all i, j ($i \neq j$),

$$\frac{1}{d} \sum_{k=1}^d |\rho_{i,j,k}|^3 \leq \frac{1}{d} \tilde{O}(q^{-1/2}) \cdot \tilde{O}(\sqrt{d/q}) = \tilde{O}\left(\frac{1}{q\sqrt{d}}\right).$$

Therefore we have shown that with high probability, for all $i \neq j$,

$$\left| \left[\Theta_{\text{CNN}}^{(2)} - \zeta^2 \mathbf{X} \mathbf{X}^\top / d \right]_{i,j} \right| = \tilde{O}\left(\frac{1}{q\sqrt{d}}\right).$$

This implies

$$\begin{aligned} \left\| \left(\Theta_{\text{CNN}}^{(2)} - \zeta^2 \mathbf{X} \mathbf{X}^\top / d \right)_{\text{off}} \right\| &\leq \left\| \left(\Theta_{\text{CNN}}^{(2)} - \zeta^2 \mathbf{X} \mathbf{X}^\top / d \right)_{\text{off}} \right\|_F = \tilde{O}\left(\frac{n}{q\sqrt{d}}\right) \\ &= \tilde{O}\left(\frac{n}{d^{\frac{1}{2}+2\alpha}\sqrt{d}}\right) = O\left(\frac{n}{d^{1+\alpha}}\right). \end{aligned}$$

For the diagonal entries, we can easily see

$$\left\| \left(\Theta_{\text{CNN}}^{(2)} - \zeta^2 \mathbf{X} \mathbf{X}^\top / d \right)_{\text{diag}} \right\| = O(1) = O\left(\frac{n}{d^{1+\alpha}}\right).$$

Combining the above two equations, we obtain

$$\left\| \Theta_{\text{CNN}}^{(2)} - \zeta^2 \mathbf{X} \mathbf{X}^\top / d \right\| = O\left(\frac{n}{d^{1+\alpha}}\right).$$

Step 2: The first-layer NTK. We calculate the derivative of the output of the CNN with respect to the first-layer weights as:

$$\nabla_{\mathbf{w}_r} f_{\text{CNN}}(\mathbf{x}; \mathbf{W}, \mathbf{V}) = \frac{1}{\sqrt{md}} \sum_{k=1}^d [\mathbf{v}_r]_k \phi' \left(\langle \mathbf{w}_r, [\mathbf{x}]_{k:k+q} \rangle / \sqrt{q} \right) [\mathbf{x}]_{k:k+q} / \sqrt{q}.$$

Therefore, the entries in the first-layer NTK matrix are

$$\begin{aligned} \left[\Theta_{\text{CNN}}^{(2)} \right]_{i,j} &= \mathbb{E}_{\mathbf{W}, \mathbf{V}} \left[\sum_{r=1}^m \langle \nabla_{\mathbf{w}_r} f_{\text{CNN}}(\mathbf{x}_i; \mathbf{W}, \mathbf{V}), \nabla_{\mathbf{w}_r} f_{\text{CNN}}(\mathbf{x}_j; \mathbf{W}, \mathbf{V}) \rangle \right] \\ &= \mathbb{E}_{\mathbf{W}} \left[\frac{1}{md} \sum_{r=1}^m \sum_{k=1}^d \phi' \left(\langle \mathbf{w}_r, [\mathbf{x}_i]_{k:k+q} \rangle / \sqrt{q} \right) \phi' \left(\langle \mathbf{w}_r, [\mathbf{x}_j]_{k:k+q} \rangle / \sqrt{q} \right) \rho_{i,j,k} \right] \\ &= \mathbb{E}_{\mathbf{w} \sim \mathcal{N}(\mathbf{0}_q, \mathbf{I}_q)} \left[\frac{1}{d} \sum_{k=1}^d \phi' \left(\langle \mathbf{w}, [\mathbf{x}_i]_{k:k+q} \rangle / \sqrt{q} \right) \phi' \left(\langle \mathbf{w}, [\mathbf{x}_j]_{k:k+q} \rangle / \sqrt{q} \right) \rho_{i,j,k} \right] \\ &= \frac{1}{d} \sum_{k=1}^d Q(\rho_{i,j,k}) \cdot \rho_{i,j,k}, \end{aligned}$$

where

$$Q(\rho) := \mathbb{E}_{(z_1, z_2) \sim \mathcal{N}(\mathbf{0}, \mathbf{\Lambda})} [\phi'(z_1) \phi'(z_2)], \text{ where } \mathbf{\Lambda} = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}, \quad |\rho| \leq 1.$$

For $i \neq j$, we can do a Taylor expansion of Q around 0: $Q(\rho) = \zeta^2 \pm O(\rho^2)$. Here since $\phi' = \text{erf}'$ is an even function, all the odd-order terms in the expansion vanish. Therefore we have

$$\left[\Theta_{\text{CNN}}^{(1)} \right]_{i,j} = \frac{1}{d} \sum_{k=1}^d (\zeta^2 \pm O(\rho_{i,j,k}^2)) \rho_{i,j,k} = \frac{1}{d} \zeta^2 \mathbf{x}_i^\top \mathbf{x}_j \pm \frac{1}{d} \sum_{k=1}^d O(|\rho_{i,j,k}|^3).$$

Then, using the exact same analysis for the second-layer NTK, we know that with high probability,

$$\left\| \Theta_{\text{CNN}}^{(1)} - \zeta^2 \mathbf{X} \mathbf{X}^\top / d \right\| = O\left(\frac{n}{d^{1+\alpha}}\right).$$

Finally, combining the results for two layers, we conclude the proof of Proposition 4.1. \square