
Gradient Descent Finds Global Minima of Deep Neural Networks

Simon S. Du^{*1} Jason D. Lee^{*2} Haochuan Li^{*34} Liwei Wang^{*54} Xiyu Zhai^{*6}

Abstract

Gradient descent finds a global minimum in training deep neural networks despite the objective function being non-convex. The current paper proves gradient descent achieves zero training loss in polynomial time for a deep over-parameterized neural network with residual connections (ResNet). Our analysis relies on the particular structure of the Gram matrix induced by the neural network architecture. This structure allows us to show the Gram matrix is stable throughout the training process and this stability implies the global optimality of the gradient descent algorithm. We further extend our analysis to deep residual convolutional neural networks and obtain a similar convergence result.

1. Introduction

One of the mysteries in deep learning is randomly initialized first-order methods like gradient descent achieve zero training loss, even if the labels are arbitrary (Zhang et al., 2016). Over-parameterization is widely believed to be the main reason for this phenomenon as only if the neural network has a sufficiently large capacity, it is possible for this neural network to fit all the training data. For example, Lu et al. (2017) proved that except for a measure zero set, all functions cannot be approximated by ReLU networks with a width less than the input dimension. In practice, many neural network architectures are highly over-parameterized. For example, Wide Residual Networks have 100x parameters than the number of training data (Zagoruyko & Komodakis, 2016).

^{*}Equal contribution ¹Machine Learning Department, Carnegie Mellon University ²Data Science and Operations Department, University of Southern California ³School of Physics, Peking University ⁴Center for Data Science, Peking University, Beijing Institute of Big Data Research ⁵Key Laboratory of Machine Perception, MOE, School of EECS, Peking University ⁶Department of EECS, Massachusetts Institute of Technology. Correspondence to: Simon S. Du <ssdu@cs.cmu.edu>.

The second mysterious phenomenon in training deep neural networks is “deeper networks are harder to train.” To solve this problem, He et al. (2016) proposed the deep residual network (ResNet) architecture which enables randomly initialized first order method to train neural networks with an order of magnitude more layers. Theoretically, Hardt & Ma (2016) showed that residual links in linear networks prevent gradient vanishing in a large neighborhood of zero, but for neural networks with non-linear activations, the advantages of using residual connections are not well understood.

In this paper, we demystify these two mysterious phenomena. We consider the setting where there are n data points, and the neural network has H layers with width m . We focus on the least-squares loss and assume the activation function is Lipschitz and smooth. This assumption holds for many activation functions including the soft-plus and sigmoid. Our contributions are summarized below.

- As a warm-up, we first consider a fully-connected feedforward network. We show if $m = \Omega(\text{poly}(n)2^{O(H)})^1$, then randomly initialized gradient descent converges to zero training loss at a linear rate.
- Next, we consider the ResNet architecture. We show as long as $m = \Omega(\text{poly}(n, H))$, then randomly initialized gradient descent converges to zero training loss at a linear rate. Comparing with the first result, the dependence on the number of layers improves exponentially for ResNet. This theory demonstrates the advantage of using residual architectures.
- Lastly, we apply the same technique to analyze convolutional ResNet. We show if $m = \text{poly}(n, p, H)$ where p is the number of patches, then randomly initialized gradient descent achieves zero training loss.

Our proof builds on two ideas from previous work on gradient descent for two-layer neural networks. First, we use the observation by (Li & Liang, 2018) that if the neural network is over-parameterized, every weight matrix is close to its initialization. Second, following (Du et al., 2018b),

¹The precise polynomials and data-dependent parameters are stated in Section 5, 6, 7.

we analyze the dynamics of the predictions whose convergence is determined by the least eigenvalue of the Gram matrix induced by the neural network architecture and to lower bound the least eigenvalue, it is sufficient to bound the distance of each weight matrix from its initialization.

Different from these two works, in analyzing deep neural networks, we need to exploit more structural properties of deep neural networks and develop new techniques for analyzing both the initialization and gradient descent dynamics. In Section 4 we give an overview of our proof technique.

1.1. Organization

This paper is organized as follows. In Section 2, we discuss related works. In Section 3, we formally state the problem setup. In Section 4, we present our main analysis techniques. In Section 5, we give a warm-up result for the deep fully-connected neural network. In Section 6, we give our main result for the ResNet. In Section 7, we give our main result for the convolutional ResNet. We conclude in Section 8 and defer all proofs to the appendix.

2. Related Works

Recently, many works try to study the optimization problem in deep learning. Since optimizing a neural network is a non-convex problem, one approach is first to develop a general theory for a class of non-convex problems which satisfy desired geometric properties and then identify that the neural network optimization problem belongs to this class. One promising candidate class is the set of functions that satisfy: a) all local minima are global and b) there exists a negative curvature for every saddle point. For this function class, researchers have shown (perturbed) gradient descent (Jin et al., 2017; Ge et al., 2015; Lee et al., 2016; Du et al., 2017a) can find a global minimum. Many previous works thus try to study the optimization landscape of neural networks with different activation functions (Soudry & Hoffer, 2017; Safran & Shamir, 2018; 2016; Zhou & Liang, 2017; Freeman & Bruna, 2016; Hardt & Ma, 2016; Nguyen & Hein, 2017; Kawaguchi, 2016; Venturi et al., 2018; Soudry & Carmon, 2016; Du & Lee, 2018; Soltanolkotabi et al., 2018; Haeffele & Vidal, 2015). However, even for a three-layer linear network, there exists a saddle point that does not have a negative curvature (Kawaguchi, 2016), so it is unclear whether this geometry-based approach can be used to obtain the global convergence guarantee of first-order methods.

Another way to attack this problem is to study the dynamics of a specific algorithm for a specific neural network architecture. Our paper also belongs to this category. Many previous works put assumptions on the input distribution and assume the label is generated according to a planted neural net-

work. Based on these assumptions, one can obtain global convergence of gradient descent for some shallow neural networks (Tian, 2017; Soltanolkotabi, 2017; Brutzkus & Globerson, 2017; Du et al., 2018a; Li & Yuan, 2017; Du et al., 2017b). Some local convergence results have also been proved (Zhong et al., 2017a;b; Zhang et al., 2018). In comparison, our paper does not try to recover the underlying neural network. Instead, we focus on minimizing the training loss and rigorously prove that randomly initialized gradient descent can achieve zero training loss.

The most related papers are (Li & Liang, 2018; Du et al., 2018b) who observed that when training an over-parameterized two-layer fully-connected neural network, the weights do not change a large amount, which we also use to show the stability of the Gram matrix. They used this observation to obtain the convergence rate of gradient descent on a two-layer over-parameterized neural network for the cross-entropy and least-squares loss. More recently, Allen-Zhu et al. (2018b) generalized ideas from (Li & Liang, 2018) to derive convergence rates of training recurrent neural networks.

Our work extends these previous results in several ways: a) we consider deep networks, b) we generalize to ResNet architectures, and c) we generalize to convolutional networks. To improve the width dependence m on sample size n , we utilize a smooth activation (e.g. smooth ReLU). For example, our results specialized to depth $H = 1$ improve upon (Du et al., 2018b) in the required amount of overparametrization from $m = \Omega(n^6)$ to $m = \Omega(n^4)$. See Theorem 5.1 for the precise statement.

Chizat & Bach (2018b) brought to our attention the paper of Jacot et al. (2018) which proved a similar weight stability phenomenon for deep networks, but only in the asymptotic setting of infinite-width networks and gradient flow run for a finite time. Jacot et al. (2018) do not establish the convergence of gradient flow to a global minimizer. In lieu of their results, our work can be viewed as a generalization of their result to: a) finite width, b) gradient descent as opposed to gradient flow, and c) convergence to a global minimizer.

Mei et al. (2018); Chizat & Bach (2018a); Sirignano & Spiliopoulos (2018); Rotskoff & Vanden-Eijnden (2018); Wei et al. (2018) used optimal transport theory to analyze gradient descent on over-parameterized models. However, their results are limited to two-layer neural networks and may require an exponential amount of over-parametrization.

Daniely (2017) developed the connection between deep neural networks with kernel methods and showed stochastic gradient descent can learn a function that is competitive with the best function in the conjugate kernel space of the network. Andoni et al. (2014) showed that gradient descent can

learn networks that are competitive with polynomial classifiers. However, these results do not imply gradient descent can find a global minimum for the empirical loss minimization problem. Our analysis of the Gram matrices at random initialization is closely related to prior work on the analysis of infinite-width networks as Gaussian Processes (Raghu et al., 2016; Matthews et al., 2018; Lee et al., 2017; Schoenholz et al., 2016). Since we require the initialization analysis for three distinct architectures (ResNet, feed-forward, and convolutional ResNet), we re-derive many of these prior results in a unified fashion in Appendix E.

Finally, in concurrent work, Allen-Zhu et al. (2018c) also analyze gradient descent on deep neural networks. The primary difference between the two papers is that we analyze general smooth activations, and Allen-Zhu et al. (2018c) develop specific analysis for ReLU activation. The two papers also differ significantly on their data assumptions. *We wish to emphasize a fair comparison is not possible due to the difference in setting and data assumptions. We view the two papers as complementary since they address different neural net architectures.*

For ResNet, the primary focus of this manuscript, the required width per layer for Allen-Zhu et al. (2018c) is $m \gtrsim n^{30} H^{30} \log^2 \frac{1}{\epsilon}$ and for this paper’s Theorem 6.1 is $m \gtrsim n^4 H^{2.2}$. Our paper requires a width m that does not depend on the desired accuracy ϵ . As a consequence, Theorem 6.1 guarantees the convergence of gradient descent to a global minimizer. The iteration complexity of Allen-Zhu et al. (2018c) is $T \gtrsim n^6 H^2 \log \frac{1}{\epsilon}$ and of Theorem 6.1 is $T \gtrsim n^2 \log \frac{1}{\epsilon}$.

For fully-connected networks, Allen-Zhu et al. (2018c) requires width $m \gtrsim n^{30} H^{30} \log^2 \frac{1}{\epsilon}$ and iteration complexity $T \gtrsim n^6 H^2 \log \frac{1}{\epsilon}$. Theorem 5.1 requires width $m \gtrsim n^4 2^{O(H)}$ and iteration complexity $T \gtrsim n^2 2^{O(H)} \log \frac{1}{\epsilon}$. The primary difference is for very deep fully-connected networks, Allen-Zhu et al. (2018c) has milder dependence on H , but worse dependence on n . Commonly used fully-connected networks such as VGG are not extremely deep ($H = 16$), yet the dataset size such as ImageNet ($n \sim 10^6$) is very large.

In a second concurrent work, Zou et al. (2018) also analyzed the convergence of gradient descent on fully-connected networks with ReLU activation. The emphasis is on different loss functions (e.g. hinge loss), so the results are not directly comparable. Both Zou et al. (2018) and Allen-Zhu et al. (2018c) train a subset of the layers, instead of all the layers as in this work, but also analyze stochastic gradient.

²In all comparisons, we ignore the polynomial dependency on data-dependent parameters which only depends on the input data and the activation function. The two papers use different measures and are not directly comparable.

3. Preliminaries

3.1. Notations

We Let $[n] = \{1, 2, \dots, n\}$. We use $N(\mathbf{0}, \mathbf{I})$ to denote the standard Gaussian distribution. For a matrix \mathbf{A} , we use \mathbf{A}_{ij} to denote its (i, j) -th entry. We will also use $\mathbf{A}_{i,\cdot}$ to denote the i -th row vector of \mathbf{A} and define $\mathbf{A}_{i,j:k} = (\mathbf{A}_{i,j}, \mathbf{A}_{i,j+1}, \dots, \mathbf{A}_{i,k})$ as part of the vector. Similarly $\mathbf{A}_{\cdot,i}$ is the i -th column vector and $\mathbf{A}_{j:k,i}$ is a part of i -th column vector. For a vector \mathbf{v} , we use $\|\mathbf{v}\|_2$ to denote the Euclidean norm. For a matrix \mathbf{A} we use $\|\mathbf{A}\|_F$ to denote the Frobenius norm and $\|\mathbf{A}\|_2$ to denote the operator norm. If a matrix \mathbf{A} is positive semi-definite, we use $\lambda_{\min}(\mathbf{A})$ to denote its smallest eigenvalue. We use $\langle \cdot, \cdot \rangle$ to denote the standard Euclidean inner product between two vectors or matrices. We let $O(\cdot)$ and $\Omega(\cdot)$ denote standard Big-O and Big-Omega notations, only hiding constants. In this paper we will use C and c to denote constants. The specific value can be different from line to line.

3.2. Activation Function

We use $\sigma(\cdot)$ to denote the activation function. In this paper we impose some technical conditions on the activation function. The guiding example is softplus: $\sigma(z) = \log(1 + \exp(z))$.

Condition 3.1 (Lipschitz and Smooth). *There exists a constant $c > 0$ such that $|\sigma(0)| \leq c$ and for any $z, z' \in \mathbb{R}$,*

$$|\sigma(z) - \sigma(z')| \leq c|z - z'|, \\ \text{and } |\sigma'(z) - \sigma'(z')| \leq c|z - z'|.$$

These two conditions will be used to show the stability of the training process. Note for softplus both Lipschitz constant and smoothness constant are 1. In this paper, we view all activation function related parameters as constants.

Condition 3.2. $\sigma(\cdot)$ is analytic and is not a polynomial function.

This assumption is used to guarantee the positive-definiteness of certain Gram matrices which we will define later. Softplus function satisfies this assumption by definition.

3.3. Problem Setup

In this paper, we focus on the empirical risk minimization problem with the quadratic loss function

$$\min_{\theta} L(\theta) = \frac{1}{2} \sum_{i=1}^n (f(\theta, \mathbf{x}_i) - y_i)^2 \quad (1)$$

where $\{\mathbf{x}_i\}_{i=1}^n$ are the training inputs, $\{y_i\}_{i=1}^n$ are the labels, θ is the parameter we optimize over and f is the prediction

function, which in our case is a neural network. We consider the following architectures.

- **Multilayer fully-connected neural networks:** Let $\mathbf{x} \in \mathbb{R}^d$ be the input, $\mathbf{W}^{(1)} \in \mathbb{R}^{m \times d}$ is the first weight matrix, $\mathbf{W}^{(h)} \in \mathbb{R}^{m \times m}$ is the weight at the h -th layer for $2 \leq h \leq H$, $\mathbf{a} \in \mathbb{R}^m$ is the output layer and $\sigma(\cdot)$ is the activation function.³ We define the prediction function recursively (for simplicity we let $\mathbf{x}^{(0)} = \mathbf{x}$).

$$\begin{aligned} \mathbf{x}^{(h)} &= \sqrt{\frac{c_\sigma}{m}} \sigma\left(\mathbf{W}^{(h)} \mathbf{x}^{(h-1)}\right), 1 \leq h \leq H \\ f(\mathbf{x}, \theta) &= \mathbf{a}^\top \mathbf{x}^{(H)}. \end{aligned} \quad (2)$$

where $c_\sigma = (\mathbb{E}_{x \sim N(0,1)} [\sigma(x)^2])^{-1}$ is a scaling factor to normalize the input in the initialization phase.

- **ResNet⁴:** We use the same notations as the multilayer fully connected neural networks. We define the prediction recursively.

$$\begin{aligned} \mathbf{x}^{(1)} &= \sqrt{\frac{c_\sigma}{m}} \sigma\left(\mathbf{W}^{(1)} \mathbf{x}\right), \\ \mathbf{x}^{(h)} &= \mathbf{x}^{(h-1)} + \frac{c_{res}}{H\sqrt{m}} \sigma\left(\mathbf{W}^{(h)} \mathbf{x}^{(h-1)}\right) \\ &\quad \text{for } 2 \leq h \leq H, \\ f_{res}(\mathbf{x}, \theta) &= \mathbf{a}^\top \mathbf{x}^{(H)} \end{aligned} \quad (3)$$

where $0 < c_{res} < 1$ is a small constant. Note here we use a $\frac{c_{res}}{H\sqrt{m}}$ scaling. This scaling plays an important role in guaranteeing the width per layer only needs to scale polynomially with H . In practice, the small scaling is enforced by a small initialization of the residual connection (Hardt & Ma, 2016; Zhang et al., 2019), which obtains state-of-the-art performance for deep residual networks. We choose to use an explicit scaling, instead of altering the initialization scheme for notational convenience.

- **Convolutional ResNet:** Lastly, we consider the convolutional ResNet architecture. Again we define the prediction function in a recursive way. Let $\mathbf{x}^{(0)} \in \mathbb{R}^{d_0 \times p}$ be the input, where d_0 is the number of input channels and p is the number of pixels. For $h \in [H]$, we let the number of channels be $d_h = m$ and number of

³We assume intermediate layers are square matrices for simplicity. It is not difficult to generalize our analysis to rectangular weight matrices.

⁴We will refer to this architecture as ResNet, although this differs by the standard ResNet architecture since the skip-connections at every layer, instead of every two layers. This architecture was previously studied in (Hardt & Ma, 2016). We study this architecture for the ease of presentation and analysis. It is not hard to generalize our analysis to architectures with skip-connections are every two or more layers.

pixels be p . Given $\mathbf{x}^{(h-1)} \in \mathbb{R}^{d_{h-1} \times p}$ for $h \in [H]$, we first use an operator $\phi_h(\cdot)$ to divide $\mathbf{x}^{(h-1)}$ into p patches. Each patch has size qd_{h-1} and this implies a map $\phi_h(\mathbf{x}^{(h-1)}) \in \mathbb{R}^{qd_{h-1} \times p}$. For example, when the stride is 1 and $q = 3$

$$\begin{aligned} &\phi_h(\mathbf{x}^{(h-1)}) \\ &= \begin{pmatrix} \left(\mathbf{x}_{1,0:2}^{(h-1)}\right)^\top, & \dots, & \left(\mathbf{x}_{1,p-1:p+1}^{(h-1)}\right)^\top \\ \dots, & \dots, & \dots \\ \left(\mathbf{x}_{d_{h-1},0:2}^{(h-1)}\right)^\top, & \dots, & \left(\mathbf{x}_{d_{h-1},p-1:p+1}^{(h-1)}\right)^\top \end{pmatrix} \end{aligned}$$

where we let $\mathbf{x}_{:,0}^{(h-1)} = \mathbf{x}_{:,p+1}^{(h-1)} = \mathbf{0}$, i.e., zero-padding. Note this operator has the property

$$\left\| \mathbf{x}^{(h-1)} \right\|_F \leq \left\| \phi_h(\mathbf{x}^{(h-1)}) \right\|_F \leq \sqrt{q} \left\| \mathbf{x}^{(h-1)} \right\|_F.$$

because each element from $\mathbf{x}^{(h-1)}$ at least appears once and at most appears q times. In practice, q is often small like 3×3 , so throughout the paper we view q as a constant in our theoretical analysis. To proceed, let $\mathbf{W}^{(h)} \in \mathbb{R}^{d_h \times qd_{h-1}}$, we have

$$\begin{aligned} \mathbf{x}^{(1)} &= \sqrt{\frac{c_\sigma}{m}} \sigma\left(\mathbf{W}^{(1)} \phi_1(\mathbf{x})\right) \in \mathbb{R}^{m \times p}, \\ \mathbf{x}^{(h)} &= \mathbf{x}^{(h-1)} + \frac{c_{res}}{H\sqrt{m}} \sigma\left(\mathbf{W}^{(h)} \phi_h(\mathbf{x}^{(h-1)})\right) \in \mathbb{R}^{m \times p} \\ &\quad \text{for } 2 \leq h \leq H, \end{aligned}$$

where $0 < c_{res} < 1$ is a small constant. Finally, for $\mathbf{a} \in \mathbb{R}^{m \times p}$, the output is defined as

$$f_{cnn}(\mathbf{x}, \theta) = \langle \mathbf{a}, \mathbf{x}^{(H)} \rangle.$$

Note here we use the similar scaling $O(\frac{1}{H\sqrt{m}})$ as ResNet.

To learn the deep neural network, we consider the randomly initialized gradient descent algorithm to find the global minimizer of the empirical loss (1). Specifically, we use the following random initialization scheme. For every level $h \in [H]$, each entry is sampled from a standard Gaussian distribution, $\mathbf{W}_{ij}^{(h)} \sim N(0, 1)$ and each entry of the output layer \mathbf{a} is also sampled from $N(0, 1)$. In this paper, we train all layers by gradient descent, for $k = 1, 2, \dots$, and $h \in [H]$

$$\begin{aligned} \mathbf{W}^{(h)}(k) &= \mathbf{W}^{(h)}(k-1) - \eta \frac{\partial L(\theta(k-1))}{\partial \mathbf{W}^{(h)}(k-1)}, \\ \mathbf{a}(k) &= \mathbf{a}(k-1) - \eta \frac{\partial L(\theta(k-1))}{\partial \mathbf{a}(k-1)} \end{aligned}$$

where $\eta > 0$ is the step size.

4. Technique Overview

In this section, we describe our main idea of proving the global convergence of gradient descent. Our proof technique is inspired by [Du et al. \(2018b\)](#) who proposed to study the dynamics of differences between labels and predictions. Here the individual prediction at the k -th iteration is

$$u_i(k) = f(\theta(k), \mathbf{x}_i)$$

and we denote $\mathbf{u}(k) = (u_1(k), \dots, u_n(k))^\top \in \mathbb{R}^n$. [Du et al. \(2018b\)](#) showed that for two-layer fully-connected neural network, the sequence $\{\mathbf{y} - \mathbf{u}(k)\}_{k=0}^\infty$ admits the following dynamics

$$\mathbf{y} - \mathbf{u}(k+1) = (\mathbf{I} - \eta \mathbf{H}(k)) (\mathbf{y} - \mathbf{u}(k))$$

where $\mathbf{H}(k) \in \mathbb{R}^{n \times n}$ is a Gram matrix with⁵

$$\mathbf{H}_{ij}(k) = \left\langle \frac{\partial u_i(k)}{\partial \mathbf{W}^{(1)}(k)}, \frac{\partial u_j(k)}{\partial \mathbf{W}^{(1)}(k)} \right\rangle.$$

The key finding in [\(Du et al., 2018b\)](#) is that if m is sufficiently large, $\mathbf{H}(k) \approx \mathbf{H}^\infty$ for all k where \mathbf{H}^∞ is defined as $\mathbf{H}_{ij}^\infty = \mathbb{E}_{\mathbf{w} \sim N(\mathbf{0}, \mathbf{I})} [\sigma'(\mathbf{w}^\top \mathbf{x}_i) \sigma'(\mathbf{w}^\top \mathbf{x}_j) \mathbf{x}_i^\top \mathbf{x}_j]$. Notably, \mathbf{H}^∞ is a fixed matrix which only depends on the training input, but *does not* depend on neural network parameters θ . As a direct result, in the large m regime, the dynamics of $\{\mathbf{y} - \mathbf{u}(k)\}_{k=0}^\infty$ is approximately *linear*

$$\mathbf{y} - \mathbf{u}(k+1) \approx (\mathbf{I} - \eta \mathbf{H}^\infty) (\mathbf{y} - \mathbf{u}(k)).$$

For this linear dynamics, using standard analysis technique for power method, one can show $\{\mathbf{y} - \mathbf{u}(k)\}_{k=0}^\infty$ converges to $\mathbf{0}$ where the rate is determined by the least eigenvalue of \mathbf{H}^∞ and the step size η .

We leverage this insight to our deep neural network setting. Again we consider the sequence $\{\mathbf{y} - \mathbf{u}(k)\}_{k=0}^\infty$, which admits the dynamics

$$\mathbf{y} - \mathbf{u}(k+1) = (\mathbf{I} - \eta \mathbf{G}(k)) (\mathbf{y} - \mathbf{u}(k))$$

where

$$\begin{aligned} & \mathbf{G}_{ij}(k) \\ &= \left\langle \frac{\partial u_i(k)}{\partial \theta(k)}, \frac{\partial u_j(k)}{\partial \theta(k)} \right\rangle \\ &= \sum_{h=1}^H \left\langle \frac{\partial u_i(k)}{\partial \mathbf{W}^{(h)}(k)}, \frac{\partial u_j(k)}{\partial \mathbf{W}^{(h)}(k)} \right\rangle + \left\langle \frac{\partial u_i(k)}{\partial \mathbf{a}(k)}, \frac{\partial u_j(k)}{\partial \mathbf{a}(k)} \right\rangle \\ &\triangleq \sum_{h=1}^{H+1} \mathbf{G}_{ij}^{(h)}(k). \end{aligned}$$

⁵This formula is for the setting that only the first layer is trained.

Here we define $\mathbf{G}^{(h)} \in \mathbb{R}^{n \times n}$ with $\mathbf{G}_{ij}^{(h)}(k) = \left\langle \frac{\partial u_i(k)}{\partial \mathbf{W}^{(h)}(k)}, \frac{\partial u_j(k)}{\partial \mathbf{W}^{(h)}(k)} \right\rangle$ for $h = 1, \dots, H$ and $\mathbf{G}_{ij}^{(H+1)}(k) = \left\langle \frac{\partial u_i(k)}{\partial \mathbf{a}(k)}, \frac{\partial u_j(k)}{\partial \mathbf{a}(k)} \right\rangle$. Note for all $h \in [H+1]$, each entry of $\mathbf{G}^{(h)}(k)$ is an inner product. Therefore, $\mathbf{G}^{(h)}(k)$ is a positive semi-definite (PSD) matrix for $h \in [H+1]$. Furthermore, if there exists one $h \in [H]$ that $\mathbf{G}^{(h)}(k)$ is strictly positive definite, then if one chooses the step size η to be sufficiently small, the loss decreases at the k -th iteration according the analysis of power method. In this paper we focus on $\mathbf{G}^{(H)}(k)$, the gram matrix induced by the weights from H -th layer for simplicity at the cost of a minor degradation in convergence rate.⁶

We use the similar observation in [\(Du et al., 2018b\)](#) that we show if the width is large enough for all layers, for all $k = 0, 1, \dots$, $\mathbf{G}^{(H)}(k)$ is close to a fixed matrix $\mathbf{K}^{(H)} \in \mathbb{R}^{n \times n}$ which depends on the input data, neural network architecture and the activation but does not depend on neural network parameters θ . According to the analysis of the power method, once we establish this, as long as $\mathbf{K}^{(H)}$ is strictly positive definite, then the gradient descent enjoys a linear convergence rate. We will show for $\mathbf{K}^{(H)}$ is strictly positive definite as long as the training data is not degenerate (c.f. [Proposition F.1](#) and [F.2](#)).

While following the similar high-level analysis framework proposed by [Du et al. \(2018b\)](#), analyzing the convergence of gradient descent for *deep* neural network is significantly more involved and requires new technical tools. To show $\mathbf{G}^{(H)}(k)$ is close to $\mathbf{K}^{(H)}$, we have two steps. First, we show in the initialization phase $\mathbf{G}^{(H)}(0)$ is close to $\mathbf{K}^{(H)}$. Second, we show during training $\mathbf{G}^{(H)}(k)$ is close to $\mathbf{G}^{(H)}(0)$ for $k = 1, 2, \dots$. Below we give overviews of these two steps.

Analysis of Random Initialization Unlike [\(Du et al., 2018b\)](#) in which they showed $\mathbf{H}(0)$ is close to \mathbf{H}^∞ via a simple concentration inequality, showing $\mathbf{G}^{(H)}(0)$ is close to $\mathbf{K}^{(H)}$ requires more subtle calculations. First, as will be clear in the following sections, $\mathbf{K}^{(H)}$ is a recursively defined matrix. Therefore, we need to analyze how the perturbation (due to randomness of initialization and finite m) from lower layers propagates to the H -th layer. Second, this perturbation propagation involves non-linear operations due to the activation function. To quantitatively characterize this perturbation propagation dynamics, we use induction and leverage techniques from Malliavin calculus ([Malliavin, 1995](#)). We derive a general framework that allows us to analyze the initialization behavior for the fully-connected neural network, ResNet, convolutional ResNet and other potential neural network architectures in a unified way.

⁶Using the contribution of all the gram matrices to the minimum eigenvalue can potentially improve the convergence rate.

One important finding in our analysis is that ResNet architecture makes the ‘‘perturbation propagation’’ more stable. The high level intuition is the following. For fully connected neural network, suppose we have some perturbation $\|\mathbf{G}^{(1)}(0) - \mathbf{K}^{(1)}\|_2 \leq \mathcal{E}_1$ in the first layer. This perturbation propagates to the H -th layer admits the form

$$\|\mathbf{G}^{(H)}(0) - \mathbf{K}^{(H)}\|_2 \triangleq \mathcal{E}_H \lesssim 2^{O(H)} \mathcal{E}_1. \quad (4)$$

Therefore, we need to have $\mathcal{E}_1 \leq \frac{1}{2^{O(H)}}$ and this makes m have exponential dependency on H .⁷

On the other hand, for ResNet the perturbation propagation admits the form

$$\mathcal{E}_H \lesssim \left(1 + O\left(\frac{1}{H}\right)\right)^H \epsilon_1 = O(\epsilon_1) \quad (5)$$

Therefore we do not have the exponential explosion problem for ResNet. We refer readers to Section E for details.

Analysis of Perturbation of During Training The next step is to show $\mathbf{G}^{(H)}(k)$ is close to $\mathbf{G}^{(H)}(0)$ for $k = 0, 1, \dots$. Note $\mathbf{G}^{(H)}$ depends on weight matrices from all layers, so to establish that $\mathbf{G}^{(H)}(k)$ is close to $\mathbf{G}^{(H)}(0)$, we need to show $\mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0)$ is small for all $h \in [H]$ and $\mathbf{a}(k) - \mathbf{a}(0)$ is small.

In the two-layer neural network setting (Du et al., 2018b), they are able to show *every* weight vector of the first layer is close to its initialization, i.e., $\|\mathbf{W}^{(1)}(k) - \mathbf{W}^{(1)}(0)\|_{2,\infty}$ is small for $k = 0, 1, \dots$. While establishing this condition for two-layer neural network is not hard, this condition may not hold for multi-layer neural networks. In this paper, we show instead, the averaged Frobenius norm

$$\frac{1}{\sqrt{m}} \|\mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0)\|_F \quad (6)$$

is small for all $k = 0, 1, \dots$

Similar to the analysis in the initialization, showing Equation (6) is small is highly involved because again, we need to analyze how the perturbation propagates. We develop a unified proof strategy for the fully-connected neural network, ResNet and convolutional ResNet. Our analysis in this step again sheds light on the benefit of using ResNet architecture for training. The high-level intuition is similar to Equation (5). See Section B, C, and D for details.

⁷We not mean to imply that fully-connected networks necessarily depend exponentially on H , but simply to illustrate in our analysis why the exponential dependence arises. For specific activations such as ReLU and careful initialization schemes, this exponential dependence may be avoided.

5. Warm Up: Convergence Result of GD for Deep Fully-connected Neural Networks

In this section, as a warm up, we show gradient descent with a constant positive step size converges to the global minimum at a linear rate. As we discussed in Section 4, the convergence rate depends on least eigenvalue of the Gram matrix $\mathbf{K}^{(H)}$.

Definition 5.1. *The Gram matrix $\mathbf{K}^{(H)}$ is recursively defined as follows, for $(i, j) \in [n] \times [n]$, and $h = 1, \dots, H - 1$*

$$\begin{aligned} \mathbf{K}_{ij}^{(0)} &= \langle \mathbf{x}_i, \mathbf{x}_j \rangle, \\ \mathbf{A}_{ij}^{(h)} &= \begin{pmatrix} \mathbf{K}_{ii}^{(h-1)} & \mathbf{K}_{ij}^{(h-1)} \\ \mathbf{K}_{ji}^{(h-1)} & \mathbf{K}_{jj}^{(h-1)} \end{pmatrix}, \\ \mathbf{K}_{ij}^{(h)} &= c_\sigma \mathbb{E}_{(u,v)^\top \sim N(\mathbf{0}, \mathbf{A}_{ij}^{(h)})} [\sigma(u) \sigma(v)], \\ \mathbf{K}_{ij}^{(H)} &= c_\sigma \mathbf{K}_{ij}^{(H-1)} \mathbb{E}_{(u,v)^\top \sim N(\mathbf{0}, \mathbf{A}_{ij}^{(H-1)})} [\sigma'(u) \sigma'(v)]. \end{aligned} \quad (7)$$

The derivation of this Gram matrix is deferred to Section E. The convergence rate and the amount of over-parameterization depends on the least eigenvalue of this Gram matrix. In Section F.1 we show as long as the input training data is not degenerate, then $\lambda_{\min}(\mathbf{K}^{(H)})$ is strictly positive. We remark that if $H = 1$, then $\mathbf{K}^{(H)}$ is the same the Gram matrix defined in (Du et al., 2018b).

Now we are ready to state our main convergence result of gradient descent for deep fully-connected neural networks.

Theorem 5.1 (Convergence Rate of Gradient Descent for Deep Fully-connected Neural Networks). *Assume for all $i \in [n]$, $\|\mathbf{x}_i\|_2 = 1$, $|y_i| = O(1)$ and the number of hidden nodes per layer*

$$m = \Omega \left(2^{O(H)} \max \left\{ \frac{n^4}{\lambda_{\min}^4(\mathbf{K}^{(H)})}, \frac{n}{\delta}, \frac{n^2 \log(\frac{Hn}{\delta})}{\lambda_{\min}^2(\mathbf{K}^{(H)})} \right\} \right)$$

where $\mathbf{K}^{(H)}$ is defined in Equation (7). If we set the step size

$$\eta = O \left(\frac{\lambda_{\min}(\mathbf{K}^{(H)})}{n^2 2^{O(H)}} \right),$$

then with probability at least $1 - \delta$ over the random initialization the loss, for $k = 1, 2, \dots$, the loss at each iteration satisfies

$$L(\theta(k)) \leq \left(1 - \frac{\eta \lambda_{\min}(\mathbf{K}^{(H)})}{2} \right)^k L(\theta(0)).$$

This theorem states that if the width m is large enough and we set step size appropriately then gradient descent converges to the global minimum with zero loss at linear

rate. The main assumption of the theorem is that we need a large enough width of each layer. The width m depends on n , H and $1/\lambda_{\min}(\mathbf{K}^{(H)})$. The dependency on n is only polynomial, which is the same as previous work on shallow neural networks (Du et al., 2018b; Li & Liang, 2018). Similar to (Du et al., 2018b), m also polynomially depends on $1/\lambda_{\min}(\mathbf{K}^{(H)})$. However, the dependency on the number of layers H is exponential. As we discussed in Section B.1, this exponential comes from the instability of the fully-connected architecture (c.f. Equation (4)). In the next section, we show with ResNet architecture, we can reduce the dependency on H from $2^{(H)}$ to $\text{poly}(H)$.

Note the requirement of m has three terms. The first term is used to show the Gram matrix is stable during training. The second term is used to guarantee the output in each layer is approximately normalized at the initialization phase. The third term is used to show the perturbation of Gram matrix at the initialization phase is small. See Section B for proofs.

The convergence rate depends step size η and $\lambda_{\min}(\mathbf{K}^{(H)})$, similar to (Du et al., 2018b). Here we require $\eta = O\left(\frac{\lambda_{\min}(\mathbf{K}^{(H)})}{n^{2^{O(H)}}}\right)$. When $H = 1$, this requirement is the same as the one used in (Du et al., 2018b). However, for deep fully-connected neural network, we require η to be exponentially small in terms of number of layers. The reason is similar to that we require m to be exponentially large. Again, this will be improved in the next section.

6. Convergence Result of GD for ResNet

In this section we consider the convergence of gradient descent for training a ResNet. We will focus on how much over-parameterization is needed to ensure the global convergence of gradient descent and compare it with fully-connected neural networks. Again we first define the key Gram matrix whose least eigenvalue will determine the convergence rate.

Definition 6.1. *The Gram matrix $\mathbf{K}^{(H)}$ is recursively defined as follows, for $(i, j) \in [n] \times [n]$ and $h = 2, \dots, H-1$:*

$$\begin{aligned} \mathbf{K}_{ij}^{(0)} &= \langle \mathbf{x}_i, \mathbf{x}_j \rangle, \\ \mathbf{K}_{ij}^{(1)} &= \mathbb{E}_{(u,v)^\top \sim N\left(\mathbf{0}, \begin{pmatrix} \mathbf{K}_{ii}^{(0)} & \mathbf{K}_{ij}^{(0)} \\ \mathbf{K}_{ji}^{(0)} & \mathbf{K}_{jj}^{(0)} \end{pmatrix}\right)} c_\sigma \sigma(u) \sigma(v), \\ \mathbf{b}_i^{(1)} &= \sqrt{c_\sigma} \mathbb{E}_{u \sim N(0, \mathbf{K}_{ii}^{(0)})} [\sigma(u)], \\ \mathbf{A}_{ij}^{(h)} &= \begin{pmatrix} \mathbf{K}_{ii}^{(h-1)} & \mathbf{K}_{ij}^{(h-1)} \\ \mathbf{K}_{ji}^{(h-1)} & \mathbf{K}_{jj}^{(h-1)} \end{pmatrix} \\ \mathbf{K}_{ij}^{(h)} &= \mathbf{K}_{ij}^{(h-1)} + \mathbb{E}_{(u,v)^\top \sim N(\mathbf{0}, \mathbf{A}_{ij}^{(h)})} \left[\frac{c_{res} \mathbf{b}_i^{(h-1)} \sigma(u)}{H} \right. \end{aligned} \quad (8)$$

$$\begin{aligned} & \left. + \frac{c_{res} \mathbf{b}_j^{(h-1)} \sigma(v)}{H} + \frac{c_{res}^2 \sigma(u) \sigma(v)}{H^2} \right], \\ \mathbf{b}_i^{(h)} &= \mathbf{b}_i^{(h-1)} + \frac{c_{res}}{H} \mathbb{E}_{u \sim N(0, \mathbf{K}_{ii}^{(h-1)})} [\sigma(u)], \\ \mathbf{K}_{ij}^{(H)} &= \frac{c_{res}^2}{H^2} \mathbf{K}_{ij}^{(H-1)} \mathbb{E}_{(u,v)^\top \sim N(\mathbf{0}, \mathbf{A}_{ij}^{(H-1)})} [\sigma'(u) \sigma'(v)]. \end{aligned}$$

Comparing $\mathbf{K}^{(H)}$ of the ResNet and the one of the fully-connect neural network, the definition of $\mathbf{K}^{(H)}$ also depends on a series of $\{\mathbf{b}^{(h)}\}_{h=1}^{H-1}$. This dependency is comes from the skip connection block in the ResNet architecture. See Section E. In Section F.2, we show as long as the input training data is not degenerate, then $\lambda_{\min}(\mathbf{K}^{(H)})$ is strictly positive. Furthermore, $\lambda_{\min}(\mathbf{K}^{(H)})$ does not depend inversely exponentially in H .

Now we are ready to state our main theorem for ResNet.

Theorem 6.1 (Convergence Rate of Gradient Descent for ResNet). *Assume for all $i \in [n]$, $\|\mathbf{x}_i\|_2 = 1$, $|y_i| = O(1)$ and the number of hidden nodes per layer*

$$m = \Omega \left(\max \left\{ \frac{n^4}{\lambda_{\min}^4(\mathbf{K}^{(H)}) H^6}, \frac{n^2}{\lambda_{\min}^2(\mathbf{K}^{(H)}) H^2}, \frac{n}{\delta}, \frac{n^2 \log\left(\frac{Hn}{\delta}\right)}{\lambda_{\min}^2(\mathbf{K}^{(H)})} \right\} \right). \quad (9)$$

If we set the step size $\eta = O\left(\frac{\lambda_{\min}(\mathbf{K}^{(H)}) H^2}{n^2}\right)$, then with probability at least $1 - \delta$ over the random initialization we have for $k = 1, 2, \dots$

$$L(\theta(k)) \leq \left(1 - \frac{\eta \lambda_{\min}(\mathbf{K}^{(H)})}{2}\right)^k L(\theta(0)).$$

In sharp contrast to Theorem 5.1, this theorem is fully polynomial in the sense that both the number of neurons and the convergence rate is polynomially in n and H . Note the amount of over-parameterization depends on $\lambda_{\min}(\mathbf{K}^{(H)})$ which is the smallest eigenvalue of the H -th layer's Gram matrix. The main reason that we do not have any exponential factor here is that the skip connection block makes the overall architecture more stable in both the initialization phase and the training phase.

Note the requirement on m has 4 terms. The first two terms are used to show the Gram matrix stable during training. The third term is used to guarantee the output in each layer is approximately normalized at the initialization phase. The fourth term is used to show bound the size of the perturbation of the Gram matrix at the initialization phase. See Section C for details.

7. Convergence Result of GD for Convolutional ResNet

In this section we generalize the convergence result of gradient descent for ResNet to convolutional ResNet. Again, we focus on how much over-parameterization is needed to ensure the global convergence of gradient descent. Similar to previous sections, we first define the $\mathbf{K}^{(H)}$ for this architecture.

Definition 7.1. *The Gram matrix $\mathbf{K}^{(H)}$ is recursively defined as follows, for $(i, j) \in [n] \times [n]$, $(l, r) \in [p] \times [p]$ and $h = 2, \dots, H - 1$,*

$$\begin{aligned}
 \mathbf{K}_{ij}^{(0)} &= \phi_1(\mathbf{x}_i)^\top \phi_1(\mathbf{x}_j) \in \mathbb{R}^{p \times p}, \\
 \mathbf{K}_{ij}^{(1)} &= \mathbb{E}_{(\mathbf{u}, \mathbf{v}) \sim N\left(\mathbf{0}, \begin{pmatrix} \mathbf{K}_{ii}^{(0)} & \mathbf{K}_{ij}^{(0)} \\ \mathbf{K}_{ji}^{(0)} & \mathbf{K}_{jj}^{(0)} \end{pmatrix}\right)} c_\sigma \sigma(\mathbf{u})^\top \sigma(\mathbf{v}), \\
 \mathbf{b}_i^{(1)} &= \sqrt{c_\sigma} \mathbb{E}_{\mathbf{u} \sim N(\mathbf{0}, \mathbf{K}_{ii}^{(0)})} [\sigma(\mathbf{u})], \\
 \mathbf{A}_{ij}^{(h)} &= \begin{pmatrix} \mathbf{K}_{ii}^{(h-1)} & \mathbf{K}_{ij}^{(h-1)} \\ \mathbf{K}_{ji}^{(h-1)} & \mathbf{K}_{jj}^{(h-1)} \end{pmatrix} \\
 \mathbf{H}_{ij}^{(h)} &= \mathbf{K}_{ij}^{(h-1)} + \\
 &\quad \mathbb{E}_{(\mathbf{u}, \mathbf{v}) \sim N(\mathbf{0}, \mathbf{A}_{ij}^{(h-1)})} \left[\frac{c_{res} \mathbf{b}_i^{(h-1)\top} \sigma(\mathbf{u})}{H} \right. \\
 &\quad \left. + \frac{c_{res} \mathbf{b}_j^{(h-1)\top} \sigma(\mathbf{v})}{H} + \frac{c_{res}^2 \sigma(\mathbf{u})^\top \sigma(\mathbf{v})}{H^2} \right], \\
 \mathbf{K}_{ij,lr}^{(h)} &= \text{tr} \left(\mathbf{H}_{ij, D_l^{(h)} D_r^{(h)}}^{(h)} \right), \\
 \mathbf{b}_i^{(h)} &= \mathbf{b}_i^{(h-1)} + \frac{c_{res}}{H} \mathbb{E}_{\mathbf{u} \sim N(\mathbf{0}, \mathbf{K}_{ii}^{(h-1)})} [\sigma(\mathbf{u})] \\
 \mathbf{M}_{ij,lr}^{(H)} &= \mathbf{K}_{ij,lr}^{(H-1)} \mathbb{E}_{(\mathbf{u}, \mathbf{v}) \sim N(\mathbf{0}, \mathbf{A}_{ij}^{(H-1)})} [\sigma'(u_l) \sigma'(v_r)] \\
 \mathbf{K}_{ij}^{(H)} &= \text{tr}(\mathbf{M}_{ij}^{(H)})
 \end{aligned} \tag{10}$$

where \mathbf{u} and \mathbf{v} are both random row vectors and $D_l^{(h)} \triangleq \{s : \mathbf{x}_{:,s}^{(h-1)} \in \text{the } l^{\text{th}} \text{ patch}\}$.

Note here $\mathbf{K}_{ij}^{(h)}$ has dimension $p \times p$ for $h = 0, \dots, H - 1$ and $\mathbf{K}_{ij,lr}$ denotes the (l, r) -th entry.

Now we state our main convergence theorem for the convolutional ResNet.

Theorem 7.1 (Convergence Rate of Gradient Descent for Convolutional ResNet). *Assume for all $i \in [n]$, $\|\mathbf{x}_i\|_F = 1$, $|y_i| = O(1)$ and the number of hidden nodes per layer*

$$m = \Omega \left(\max \left\{ \frac{n^4}{\lambda_0^4 H^6}, \frac{n^4}{\lambda_0^4 H^2}, \frac{n}{\delta}, \frac{n^2 \log\left(\frac{Hn}{\delta}\right)}{\lambda_0^2} \right\} \text{poly}(p) \right). \tag{11}$$

If we set the step size $\eta = O\left(\frac{\lambda_0 H^2}{n^2 \text{poly}(p)}\right)$, then with probability at least $1 - \delta$ over the random initialization we have for $k = 1, 2, \dots$

$$L(\theta(k)) \leq \left(1 - \frac{\eta \lambda_{\min}(\mathbf{K}^{(H)})}{2} \right)^k L(\theta(0)).$$

This theorem is similar to that of ResNet. The number of neurons required per layer is only polynomial in the depth and the number of data points and step size is only polynomially small. The only extra term is $\text{poly}(p)$ in the requirement of m and η . The analysis is also similar to ResNet and we refer readers to Section D for details.

8. Conclusion

In this paper, we show that gradient descent on deep over-parametrized networks can obtain zero training loss. Our proof builds on a careful analysis of the random initialization scheme and a perturbation analysis which shows that the Gram matrix is increasingly stable under overparametrization. These techniques allow us to show that every step of gradient descent decreases the loss at a geometric rate.

We list some directions for future research:

1. The current paper focuses on the training loss, but does not address the test loss. It would be an important problem to show that gradient descent can also find solutions of low test loss. In particular, existing work only demonstrate that gradient descent works under the same situations as kernel methods and random feature methods (Daniely, 2017; Li & Liang, 2018; Allen-Zhu et al., 2018a; Arora et al., 2019). To further investigate of generalization behavior, we believe some algorithm-dependent analyses may be useful (Hardt et al., 2016; Mou et al., 2018; Chen et al., 2018).
2. The width of the layers m is polynomial in all the parameters for the ResNet architecture, but still very large. Realistic networks have number of parameters, not width, a large constant multiple of n . We consider improving the analysis to cover commonly utilized networks an important open problem.
3. The current analysis is for gradient descent, instead of stochastic gradient descent. We believe the analysis can be extended to stochastic gradient, while maintaining the linear convergence rate.
4. The convergence rate can be potentially improved if the minimum eigenvalue takes into account the contribution of all Gram matrices, but this would considerably complicate the initialization and perturbation analysis.

Acknowledgments

We thank Lijie Chen and Ruosong Wang for useful discussions. SSD acknowledges support from AFRL grant FA8750-17-2-0212 and DARPA D17AP00001. JDL acknowledges support of the ARO under MURI Award W911NF-11-1-0303. This is part of the collaboration between US DOD, UK MOD and UK Engineering and Physical Research Council (EPSRC) under the Multidisciplinary University Research Initiative. HL and LW acknowledge support from National Basic Research Program of China (973 Program) (grant no. 2015CB352502), NSFC (61573026) and BJNSF (L172037). Part of the work is done while SSD was visiting Simons Institute.

References

- Allen-Zhu, Z., Li, Y., and Liang, Y. Learning and generalization in overparameterized neural networks, going beyond two layers. *arXiv preprint arXiv:1811.04918*, 2018a.
- Allen-Zhu, Z., Li, Y., and Song, Z. On the convergence rate of training recurrent neural networks. *arXiv preprint arXiv:1810.12065*, 2018b.
- Allen-Zhu, Z., Li, Y., and Song, Z. A convergence theory for deep learning via over-parameterization. *arXiv preprint arXiv:1811.03962*, 2018c.
- Andoni, A., Panigrahy, R., Valiant, G., and Zhang, L. Learning polynomials with neural networks. In *International Conference on Machine Learning*, pp. 1908–1916, 2014.
- Arora, S., Du, S. S., Hu, W., Li, Z., and Wang, R. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. *arXiv preprint arXiv:1901.08584*, 2019.
- Brutzkus, A. and Globerson, A. Globally optimal gradient descent for a ConvNet with gaussian inputs. In *International Conference on Machine Learning*, pp. 605–614, 2017.
- Chen, Y., Jin, C., and Yu, B. Stability and Convergence Trade-off of Iterative Optimization Algorithms. *arXiv e-prints*, art. arXiv:1804.01619, Apr 2018.
- Chizat, L. and Bach, F. On the global convergence of gradient descent for over-parameterized models using optimal transport. *arXiv preprint arXiv:1805.09545*, 2018a.
- Chizat, L. and Bach, F. A note on lazy training in supervised differentiable programming. *arXiv preprint arXiv:1812.07956*, 2018b.
- Daniely, A. SGD learns the conjugate kernel class of the network. In *Advances in Neural Information Processing Systems*, pp. 2422–2430, 2017.
- Du, S. S. and Lee, J. D. On the power of over-parametrization in neural networks with quadratic activation. *Proceedings of the 35th International Conference on Machine Learning*, pp. 1329–1338, 2018.
- Du, S. S., Jin, C., Lee, J. D., Jordan, M. I., Singh, A., and Póczos, B. Gradient descent can take exponential time to escape saddle points. In *Advances in Neural Information Processing Systems*, pp. 1067–1077, 2017a.
- Du, S. S., Lee, J. D., and Tian, Y. When is a convolutional filter easy to learn? *arXiv preprint arXiv:1709.06129*, 2017b.
- Du, S. S., Lee, J. D., Tian, Y., Póczos, B., and Singh, A. Gradient descent learns one-hidden-layer CNN: Don’t be afraid of spurious local minima. *Proceedings of the 35th International Conference on Machine Learning*, pp. 1339–1348, 2018a.
- Du, S. S., Zhai, X., Póczos, B., and Singh, A. Gradient descent provably optimizes over-parameterized neural networks. *arXiv preprint arXiv:1810.02054*, 2018b.
- Freeman, C. D. and Bruna, J. Topology and geometry of half-rectified network optimization. *arXiv preprint arXiv:1611.01540*, 2016.
- Ge, R., Huang, F., Jin, C., and Yuan, Y. Escaping from saddle points – online stochastic gradient for tensor decomposition. In *Proceedings of The 28th Conference on Learning Theory*, pp. 797–842, 2015.
- Haeffele, B. D. and Vidal, R. Global optimality in tensor factorization, deep learning, and beyond. *arXiv preprint arXiv:1506.07540*, 2015.
- Hardt, M. and Ma, T. Identity matters in deep learning. *arXiv preprint arXiv:1611.04231*, 2016.
- Hardt, M., Recht, B., and Singer, Y. Train faster, generalize better: Stability of stochastic gradient descent. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1225–1234, New York, New York, USA, 20–22 Jun 2016. PMLR. URL <http://proceedings.mlr.press/v48/hardt16.html>.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. *arXiv preprint arXiv:1806.07572*, 2018.

- Jin, C., Ge, R., Netrapalli, P., Kakade, S. M., and Jordan, M. I. How to escape saddle points efficiently. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 1724–1732, 2017.
- Kawaguchi, K. Deep learning without poor local minima. In *Advances In Neural Information Processing Systems*, pp. 586–594, 2016.
- Lee, J., Bahri, Y., Novak, R., Schoenholz, S. S., Pennington, J., and Sohl-Dickstein, J. Deep neural networks as gaussian processes. *arXiv preprint arXiv:1711.00165*, 2017.
- Lee, J. D., Simchowitz, M., Jordan, M. I., and Recht, B. Gradient descent only converges to minimizers. In *Conference on Learning Theory*, pp. 1246–1257, 2016.
- Li, Y. and Liang, Y. Learning overparameterized neural networks via stochastic gradient descent on structured data. *arXiv preprint arXiv:1808.01204*, 2018.
- Li, Y. and Yuan, Y. Convergence analysis of two-layer neural networks with ReLU activation. In *Advances in Neural Information Processing Systems*, pp. 597–607, 2017.
- Lu, Z., Pu, H., Wang, F., Hu, Z., and Wang, L. The expressive power of neural networks: A view from the width. In *Advances in Neural Information Processing Systems 30*, pp. 6231–6239. Curran Associates, Inc., 2017.
- Malliavin, P. Gaussian sobolev spaces and stochastic calculus of variations. 1995.
- Matthews, A. G. d. G., Rowland, M., Hron, J., Turner, R. E., and Ghahramani, Z. Gaussian process behaviour in wide deep neural networks. *arXiv preprint arXiv:1804.11271*, 2018.
- Mei, S., Montanari, A., and Nguyen, P.-M. A mean field view of the landscape of two-layers neural networks. *Proceedings of the National Academy of Sciences*, pp. E7665–E7671, 2018.
- Mou, W., Wang, L., Zhai, X., and Zheng, K. Generalization bounds of sgld for non-convex learning: Two theoretical viewpoints. In Bubeck, S., Perchet, V., and Rigollet, P. (eds.), *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pp. 605–638. PMLR, 06–09 Jul 2018. URL <http://proceedings.mlr.press/v75/mou18a.html>.
- Nguyen, Q. and Hein, M. The loss surface of deep and wide neural networks. In *International Conference on Machine Learning*, pp. 2603–2612, 2017.
- Raghu, M., Poole, B., Kleinberg, J., Ganguli, S., and Sohl-Dickstein, J. On the expressive power of deep neural networks. *arXiv preprint arXiv:1606.05336*, 2016.
- Rotskoff, G. M. and Vanden-Eijnden, E. Neural networks as interacting particle systems: Asymptotic convexity of the loss landscape and universal scaling of the approximation error. *arXiv preprint arXiv:1805.00915*, 2018.
- Safran, I. and Shamir, O. On the quality of the initial basin in overspecified neural networks. In *International Conference on Machine Learning*, pp. 774–782, 2016.
- Safran, I. and Shamir, O. Spurious local minima are common in two-layer ReLU neural networks. In *International Conference on Machine Learning*, pp. 4433–4441, 2018.
- Schoenholz, S. S., Gilmer, J., Ganguli, S., and Sohl-Dickstein, J. Deep information propagation. *arXiv preprint arXiv:1611.01232*, 2016.
- Sirignano, J. and Spiliopoulos, K. Mean field analysis of neural networks. *arXiv preprint arXiv:1805.01053*, 2018.
- Soltanolkotabi, M. Learning ReLUs via gradient descent. In *Advances in Neural Information Processing Systems*, pp. 2007–2017, 2017.
- Soltanolkotabi, M., Javanmard, A., and Lee, J. D. Theoretical insights into the optimization landscape of overparameterized shallow neural networks. *IEEE Transactions on Information Theory*, 2018.
- Soudry, D. and Carmon, Y. No bad local minima: Data independent training error guarantees for multilayer neural networks. *arXiv preprint arXiv:1605.08361*, 2016.
- Soudry, D. and Hoffer, E. Exponentially vanishing sub-optimal local minima in multilayer neural networks. *arXiv preprint arXiv:1702.05777*, 2017.
- Tian, Y. An analytical formula of population gradient for two-layered ReLU network and its applications in convergence and critical point analysis. In *International Conference on Machine Learning*, pp. 3404–3413, 2017.
- Venturi, L., Bandeira, A., and Bruna, J. Neural networks with finite intrinsic dimension have no spurious valleys. *arXiv preprint arXiv:1802.06384*, 2018.
- Vershynin, R. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- Wei, C., Lee, J. D., Liu, Q., and Ma, T. On the margin theory of feedforward neural networks. *arXiv preprint arXiv:1810.05369*, 2018.

- Zagoruyko, S. and Komodakis, N. Wide residual networks. *NIN*, 8:35–67, 2016.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.
- Zhang, H., Dauphin, Y. N., and Ma, T. Residual learning without normalization via better initialization. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=H1gsz30cKX>.
- Zhang, X., Yu, Y., Wang, L., and Gu, Q. Learning one-hidden-layer relu networks via gradient descent. *arXiv preprint arXiv:1806.07808*, 2018.
- Zhong, K., Song, Z., and Dhillon, I. S. Learning non-overlapping convolutional neural networks with multiple kernels. *arXiv preprint arXiv:1711.03440*, 2017a.
- Zhong, K., Song, Z., Jain, P., Bartlett, P. L., and Dhillon, I. S. Recovery guarantees for one-hidden-layer neural networks. *arXiv preprint arXiv:1706.03175*, 2017b.
- Zhou, Y. and Liang, Y. Critical points of neural networks: Analytical forms and landscape properties. *arXiv preprint arXiv:1710.11205*, 2017.
- Zou, D., Cao, Y., Zhou, D., and Gu, Q. Stochastic gradient descent optimizes over-parameterized deep ReLU networks. *arXiv preprint arXiv:1811.08888*, 2018.

Appendix

In the proof we will use the geometric series function $g_\alpha(n) = \sum_{i=0}^{n-1} \alpha^i$ extensively. Some constants we will define below may be different for different network structures, such as c_x , $c_{w,0}$ and $c_{x,0}$. We will also use c to denote a small enough constant, which may be different in different lemmas. For simplicity, we use λ_0 to denote $\lambda_{\min}(\mathbf{K}^{(H)})$ in the proofs.

A. Proof Sketch

Note we can write the loss as

$$L(\theta(k)) = \frac{1}{2} \|\mathbf{y} - \mathbf{u}(k)\|_2^2.$$

Our proof is by induction. Our induction hypothesis is just the following convergence rate of empirical loss.

Condition A.1. *At the k -th iteration, we have*

$$\|\mathbf{y} - \mathbf{u}(k)\|_2^2 \leq \left(1 - \frac{\eta\lambda_0}{2}\right)^k \|\mathbf{y} - \mathbf{u}(0)\|_2^2.$$

Note this condition implies the conclusions we want to prove. To prove Condition A.1, we consider one iteration on the loss function.

$$\begin{aligned} & \|\mathbf{y} - \mathbf{u}(k+1)\|_2^2 \\ &= \|\mathbf{y} - \mathbf{u}(k) - (\mathbf{u}(k+1) - \mathbf{u}(k))\|_2^2 \\ &= \|\mathbf{y} - \mathbf{u}(k)\|_2^2 - 2(\mathbf{y} - \mathbf{u}(k))^\top (\mathbf{u}(k+1) - \mathbf{u}(k)) + \|\mathbf{u}(k+1) - \mathbf{u}(k)\|_2^2. \end{aligned} \quad (12)$$

This equation shows if $2(\mathbf{y} - \mathbf{u}(k))^\top (\mathbf{u}(k+1) - \mathbf{u}(k)) > \|\mathbf{u}(k+1) - \mathbf{u}(k)\|_2^2$, the loss decreases. Note both terms involves $\mathbf{u}(k+1) - \mathbf{u}(k)$, which we will carefully analyze. To simplify notations, we define

$$u'_i(\theta) \triangleq \frac{\partial u_i}{\partial \theta}, \quad u'_i{}^{(h)}(\theta) \triangleq \frac{\partial u_i}{\partial \mathbf{W}^{(h)}}, \quad u'_i{}^{(a)}(\theta) \triangleq \frac{\partial u_i}{\partial \mathbf{a}} \quad \text{and} \quad L'(\theta) = \frac{\partial L(\theta)}{\partial \theta}, \quad L'^{(h)}(\mathbf{W}^{(h)}) = \frac{\partial L(\theta)}{\partial \mathbf{W}^{(h)}}, \quad L'^{(a)}(\theta) \triangleq \frac{\partial L}{\partial \mathbf{a}}.$$

We look one coordinate of $\mathbf{u}(k+1) - \mathbf{u}(k)$.

Using Taylor expansion, we have

$$\begin{aligned} & u_i(k+1) - u_i(k) \\ &= u_i(\theta(k) - \eta L'(\theta(k))) - u_i(\theta(k)) \\ &= - \int_{s=0}^{\eta} \langle L'(\theta(k)), u'_i(\theta(k) - sL'(\theta(k))) \rangle ds \\ &= - \int_{s=0}^{\eta} \langle L'(\theta(k)), u'_i(\theta(k)) \rangle ds + \int_{s=0}^{\eta} \langle L'(\theta(k)), u'_i(\theta(k)) - u'_i(\theta(k) - sL'(\theta(k))) \rangle ds \\ &\triangleq I_1^i(k) + I_2^i(k). \end{aligned}$$

Denote $\mathbf{I}_1(k) = (I_1^1(k), \dots, I_1^n(k))^\top$ and $\mathbf{I}_2(k) = (I_2^1(k), \dots, I_2^n(k))^\top$ and so $\mathbf{u}(k+1) - \mathbf{u}(k) = \mathbf{I}_1(k) + \mathbf{I}_2(k)$. We will show the $\mathbf{I}_1(k)$ term, which is proportional to η , drives the loss function to decrease and the $\mathbf{I}_2(k)$ term, which is a perturbation term but it is proportional to η^2 so it is small. We further unpack the $I_1^i(k)$ term,

$$\begin{aligned} I_1^i &= -\eta \langle L'(\theta(k)), u'_i(\theta(k)) \rangle \\ &= -\eta \sum_{j=1}^n (u_j - y_j) \langle u'_j(\theta(k)), u'_i(\theta(k)) \rangle \\ &\triangleq -\eta \sum_{j=1}^n (u_j - y_j) \sum_{h=1}^{H+1} \mathbf{G}_{ij}^{(h)}(k) \end{aligned}$$

According to Section 4, we will only look at $\mathbf{G}^{(H)}$ matrix which has the following form

$$\mathbf{G}_{i,j}^{(H)}(k) = (\mathbf{x}_i^{(H-1)}(k))^\top \mathbf{x}_j^{(H-1)}(k) \cdot \frac{c_\sigma}{m} \sum_{r=1}^m a_r^2 \sigma'((\theta_r^{(H)}(k))^\top \mathbf{x}_i^{(H-1)}(k)) \sigma'((\theta_r^{(H)}(k))^\top \mathbf{x}_j^{(H-1)}(k)).$$

Now we analyze $\mathbf{I}_1(k)$. We can write \mathbf{I}_1 in a more compact form with $\mathbf{G}(k)$.

$$\mathbf{I}_1(k) = -\eta \mathbf{G}(k) (\mathbf{u}(k) - \mathbf{y}).$$

Now observe that

$$\begin{aligned} (\mathbf{y} - \mathbf{u}(k))^\top \mathbf{I}_1(k) &= \eta (\mathbf{y} - \mathbf{u}(k))^\top \mathbf{G}(k) (\mathbf{y} - \mathbf{u}(k)) \\ &\geq \lambda_{\min}(\mathbf{G}(k)) \|\mathbf{y} - \mathbf{u}(k)\|_2^2 \\ &\geq \lambda_{\min}(\mathbf{G}^{(H)}(k)) \|\mathbf{y} - \mathbf{u}(k)\|_2^2 \end{aligned}$$

Now recall the progress of loss function in Equation (12):

$$\begin{aligned} &\|\mathbf{y} - \mathbf{u}(k+1)\|_2^2 \\ &= \|\mathbf{y} - \mathbf{u}(k)\|_2^2 - 2(\mathbf{y} - \mathbf{u}(k))^\top \mathbf{I}_1(k) - 2(\mathbf{y} - \mathbf{u}(k))^\top \mathbf{I}_2(k) + \|\mathbf{u}(k+1) - \mathbf{u}(k)\|_2^2 \\ &\leq \left(1 - \eta \lambda_{\min}(\mathbf{G}^{(H)}(k))\right) \|\mathbf{y} - \mathbf{u}(k)\|_2^2 - 2(\mathbf{y} - \mathbf{u}(k))^\top \mathbf{I}_2(k) + \|\mathbf{u}(k+1) - \mathbf{u}(k)\|_2^2. \end{aligned}$$

For the perturbation terms, through standard calculations, we can show both $-2(\mathbf{y} - \mathbf{u}(k))^\top \mathbf{I}_2(k)$ and $\|\mathbf{u}(k+1) - \mathbf{u}(k)\|_2^2$ are proportional to $\eta^2 \|\mathbf{y} - \mathbf{u}(k)\|_2^2$ so if we set η sufficiently small, this term is smaller than $\eta \lambda_{\min}(\mathbf{G}^{(H)}(k)) \|\mathbf{y} - \mathbf{u}(k)\|_2^2$ and thus the loss function decreases with a linear rate.

Therefore, to prove the induction hypothesis, it suffices to prove $\lambda_{\min}(\mathbf{G}^{(H)}(k)) \geq \frac{\lambda_0}{2}$ for $k' = 0, \dots, k$, where λ_0 is independent of m . To analyze the least eigenvalue, we first look at the initialization. Using assumptions of the population Gram matrix and concentration inequalities, we can show at the beginning $\|\mathbf{G}^{(H)}(0) - \mathbf{K}^{(H)}(0)\|_2 \leq \frac{1}{4} \lambda_0$, which implies

$$\lambda_{\min}(\mathbf{G}^{(H)}(0)) \geq \frac{3}{4} \lambda_0.$$

Now for the k -th iteration, by matrix perturbation analysis, we know it is sufficient to show $\|\mathbf{G}^{(H)}(k) - \mathbf{G}^{(H)}(0)\|_2 \leq \frac{1}{4} \lambda_0$. To do this, we use a similar approach as in (Du et al., 2018b). We show as long as m is large enough, every weight matrix is close its initialization in a relative error sense. Ignoring all other parameters except m , $\|\mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0)\|_F \lesssim 1$, and thus the average per-neuron distance from initialization is $\frac{\|\mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0)\|_F}{\sqrt{m}} \lesssim \frac{1}{\sqrt{m}}$ which tends to zero as m increases. See Lemma B.5 for precise statements with all the dependencies.

This fact in turn shows $\|\mathbf{G}^{(H)}(k) - \mathbf{G}^{(H)}(0)\|_2$ is small. The main difference from (Du et al., 2018b) is that we are considering deep neural networks, and when translating the small deviation, $\|\mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0)\|_F$ to $\|\mathbf{G}^{(H)}(k) - \mathbf{G}^{(H)}(0)\|_2$, there is an amplification factor which depends on the neural network architecture.

For deep fully connected neural networks, we show this amplification factor is exponential in H . On the other hand, for ResNet and convolutional ResNet we show this amplification factor is only polynomial in H . We further show the width m required is proportional to this amplification factor.

B. Proofs for Section 5

We first derive the formula of the gradient for the multilayer fully connected neural network

$$\frac{\partial L(\theta)}{\partial \mathbf{W}^{(h)}} = \left(\frac{c_\sigma}{m}\right)^{\frac{H-h+1}{2}} \sum_{i=1}^n (f(\mathbf{x}_i, \theta) - y_i) \mathbf{x}_i^{(h-1)} \mathbf{a}^\top \left(\prod_{k=h+1}^H \mathbf{J}_i^{(k)} \mathbf{W}^{(k)} \right) \mathbf{J}_i^{(h)}$$

where

$$\mathbf{J}^{(h')} \triangleq \text{diag} \left(\sigma' \left((\mathbf{w}_1^{(h')})^\top \mathbf{x}^{(h'-1)} \right), \dots, \sigma' \left((\mathbf{w}_m^{(h')})^\top \mathbf{x}^{(h'-1)} \right) \right) \in \mathbb{R}^{m \times m}$$

are the derivative matrices induced by the activation function and

$$\mathbf{x}^{(h')} = \sqrt{\frac{c_\sigma}{m}} \sigma \left(\mathbf{W}^{(h')} \mathbf{x}^{(h'-1)} \right).$$

is the output of the h' -th layer.

Through standard calculation, we can get the expression of $\mathbf{G}_{i,j}^{(H)}$ of the following form

$$\mathbf{G}_{i,j}^{(H)} = (\mathbf{x}_i^{(H-1)})^\top \mathbf{x}_j^{(H-1)} \cdot \frac{c_\sigma}{m} \sum_{r=1}^m a_r^2 \sigma' \left((\mathbf{w}_r^{(H)})^\top \mathbf{x}_i^{(H-1)} \right) \sigma' \left((\mathbf{w}_r^{(H)})^\top \mathbf{x}_j^{(H-1)} \right). \quad (13)$$

We first present a lemma which shows with high probability the feature of each layer is approximately normalized.

Lemma B.1 (Lemma on Initialization Norms). *If $\sigma(\cdot)$ is L -Lipschitz and $m = \Omega \left(\frac{nHg_C(H)^2}{\delta} \right)$, where $C \triangleq c_\sigma L \left(2|\sigma(0)|\sqrt{\frac{2}{\pi}} + 2L \right)$, then with probability at least $1 - \delta$ over random initialization, for every $h \in [H]$ and $i \in [n]$, we have*

$$\frac{1}{c_{x,0}} \leq \left\| \mathbf{x}_i^{(h)}(0) \right\|_2 \leq c_{x,0}$$

where $c_{x,0} = 2$.

We follow the proof sketch described in Section A. We first analyze the spectral property of $\mathbf{G}^{(H)}(0)$ at the initialization phase. The following lemma lower bounds its least eigenvalue. This lemma is a direct consequence of results in Section E.

Lemma B.2 (Least Eigenvalue at the Initialization). *If $m = \Omega \left(\frac{n^2 \log(Hn/\delta) 2^{O(H)}}{\lambda_0^2} \right)$, we have*

$$\lambda_{\min}(\mathbf{G}^{(H)}(0)) \geq \frac{3}{4} \lambda_0.$$

Now we proceed to analyze the training process. We prove the following lemma which characterizes how the perturbation from weight matrices propagates to the input of each layer. This Lemma is used to prove the subsequent lemmas.

Lemma B.3. *Suppose for every $h \in [H]$, $\left\| \mathbf{W}^{(h)}(0) \right\|_2 \leq c_{w,0} \sqrt{m}$, $\left\| \mathbf{x}^{(h)}(0) \right\|_2 \leq c_{x,0}$ and $\left\| \mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0) \right\|_F \leq \sqrt{m}R$ for some constant $c_{w,0}, c_{x,0} > 0$ and $R \leq c_{w,0}$. If $\sigma(\cdot)$ is L -Lipschitz, we have*

$$\left\| \mathbf{x}^{(h)}(k) - \mathbf{x}^{(h)}(0) \right\|_2 \leq \sqrt{c_\sigma} L c_{x,0} g_{c_x}(h) R$$

where $c_x = 2\sqrt{c_\sigma} L c_{w,0}$.

Here the assumption of $\left\| \mathbf{W}^{(h)}(0) \right\|_2 \leq c_{w,0} \sqrt{m}$ can be shown using Lemma G.2 and taking union bound over $h \in [H]$, where $c_{w,0}$ is a universal constant. Next, we show with high probability over random initialization, perturbation in weight matrices leads to small perturbation in the Gram matrix.

Lemma B.4. *Suppose $\sigma(\cdot)$ is L -Lipschitz and β -smooth. Suppose for $h \in [H]$, $\left\| \mathbf{W}^{(h)}(0) \right\|_2 \leq c_{w,0} \sqrt{m}$, $\left\| \mathbf{a}(0) \right\|_2 \leq a_{2,0} \sqrt{m}$, $\left\| \mathbf{a}(0) \right\|_4 \leq a_{4,0} m^{1/4}$, $\frac{1}{c_{x,0}} \leq \left\| \mathbf{x}^{(h)}(0) \right\|_2 \leq c_{x,0}$, if $\left\| \mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0) \right\|_F, \left\| \mathbf{a}(k) - \mathbf{a}(0) \right\|_2 \leq \sqrt{m}R$ where $R \leq c g_{c_x}(H)^{-1} \lambda_0 n^{-1}$ and $R \leq c g_{c_x}(H)^{-1}$ for some small constant c and $c_x = 2\sqrt{c_\sigma} L c_{w,0}$, we have*

$$\left\| \mathbf{G}^{(H)}(k) - \mathbf{G}^{(H)}(0) \right\|_2 \leq \frac{\lambda_0}{4}.$$

Here the assumption of $\left\| \mathbf{a}(0) \right\|_2 \leq a_{2,0} \sqrt{m}$, $\left\| \mathbf{a}(0) \right\|_4 \leq a_{4,0} m^{1/4}$ can be easily obtained using standard concentration inequalities, where $a_{2,0}$ and $a_{4,0}$ are both universal constants. The following lemma shows if the induction holds, we have every weight matrix close to its initialization.

Lemma B.5. *If Condition A.1 holds for $k' = 1, \dots, k$, we have for any $s = 1, \dots, k + 1$*

$$\begin{aligned} \left\| \mathbf{W}^{(h)}(s) - \mathbf{W}^{(h)}(0) \right\|_F, \|\mathbf{a}(s) - \mathbf{a}(0)\|_2 &\leq R' \sqrt{m} \\ \left\| \mathbf{W}^{(h)}(s) - \mathbf{W}^{(h)}(s-1) \right\|_F, \|\mathbf{a}(s) - \mathbf{a}(s-1)\|_2 &\leq \eta Q'(s-1) \end{aligned}$$

where $R' = \frac{16c_{x,0}a_{2,0}(c_x)^H \sqrt{n} \|\mathbf{y} - \mathbf{u}(0)\|_2}{\lambda_0 \sqrt{m}} \leq cg_{c_x}(H)^{-1}$ for some small constant c with $c_x = \max\{2\sqrt{c_\sigma} Lc_{w,0}, 1\}$ and $Q'(s) = 4c_{x,0}a_{2,0}(c_x)^H \sqrt{n} \|\mathbf{y} - \mathbf{u}(s)\|_2$

Now we proceed to analyze the perturbation terms.

Lemma B.6. *If Condition A.1 holds for $k' = 1, \dots, k$, suppose $\eta \leq c\lambda_0 (n^2 H^2 (c_x)^{3H} g_{2c_x}(H))^{-1}$ for some small constant c , we have*

$$\|\mathbf{I}_2(k)\|_2 \leq \frac{1}{8} \eta \lambda_0 \|\mathbf{y} - \mathbf{u}(k)\|_2.$$

Lemma B.7. *If Condition A.1 holds for $k' = 1, \dots, k$, suppose $\eta \leq c\lambda_0 (n^2 H^2 (c_x)^{2H} g_{2c_x}(H))^{-1}$ for some small constant c , then we have $\|\mathbf{u}(k+1) - \mathbf{u}(k)\|_2^2 \leq \frac{1}{8} \eta \lambda_0 \|\mathbf{y} - \mathbf{u}(k)\|_2^2$.*

We now proceed with the proof of Theorem 5.1. By induction, we assume Condition A.1 for all $k' < k$. Using Lemma B.5, this establishes

$$\begin{aligned} \left\| \mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0) \right\|_F &\leq R' \sqrt{m} \\ &\leq R \sqrt{m} \end{aligned} \quad (\text{using the choice of } m \text{ in the theorem.})$$

By Lemma B.4, this establishes $\lambda_{\min}(\mathbf{G}^{(H)}(k)) \geq \frac{\lambda_0}{2}$.

With these estimates in hand, we are ready to prove the induction hypothesis of Condition A.1.

$$\begin{aligned} &\|\mathbf{y} - \mathbf{u}(k+1)\|_2^2 \\ &= \|\mathbf{y} - \mathbf{u}(k)\|_2^2 - 2\eta (\mathbf{y} - \mathbf{u}(k))^\top \mathbf{G}(k) (\mathbf{y} - \mathbf{u}(k)) - 2(\mathbf{y} - \mathbf{u}(k))^\top \mathbf{I}_2 + \|\mathbf{u}(k+1) - \mathbf{u}(k)\|_2^2 \\ &\leq \|\mathbf{y} - \mathbf{u}(k)\|_2^2 - 2\eta (\mathbf{y} - \mathbf{u}(k))^\top \mathbf{G}^{(H)}(k) (\mathbf{y} - \mathbf{u}(k)) - 2(\mathbf{y} - \mathbf{u}(k))^\top \mathbf{I}_2 + \|\mathbf{u}(k+1) - \mathbf{u}(k)\|_2^2 \\ &\leq (1 - \eta \lambda_0) \|\mathbf{y} - \mathbf{u}(k)\|_2^2 - 2(\mathbf{y} - \mathbf{u}(k))^\top \mathbf{I}_2 + \|\mathbf{u}(k+1) - \mathbf{u}(k)\|_2^2 \\ &\leq \left(1 - \frac{\eta \lambda_0}{2}\right) \|\mathbf{y} - \mathbf{u}(k)\|_2^2. \end{aligned}$$

The first inequality drops the positive terms $(\mathbf{y} - \mathbf{u}(k))^\top \sum_{h \in [H+1], h \neq H} \mathbf{G}^{(h)}(k) (\mathbf{y} - \mathbf{u}(k))$. The second inequality uses the argument above that establishes $\lambda_{\min}(\mathbf{G}^{(H)}(k)) \geq \frac{\lambda_0}{2}$. The third inequality uses Lemmas B.6 and B.7.

B.1. Proofs of Lemmas

Proof of Lemma B.1. We will bound $\left\| \mathbf{x}_i^{(h)}(0) \right\|_2$ by induction on layers. The induction hypothesis is that with probability at least $1 - (h-1) \frac{\delta}{nH}$ over $\mathbf{W}^{(1)}(0), \dots, \mathbf{W}^{(h-1)}(0)$, for every $1 \leq h' \leq h-1$, $\frac{1}{2} \leq 1 - \frac{g_C(h')}{2g_C(H)} \leq \left\| \mathbf{x}_i^{(h')}(0) \right\|_2 \leq 1 + \frac{g_C(h')}{2g_C(H)} \leq 2$. Note that it is true for $h = 1$. We calculate the expectation of $\left\| \mathbf{x}_i^{(h)}(0) \right\|_2^2$ over the randomness from $\mathbf{W}^{(h)}(0)$. Recall

$$\left\| \mathbf{x}_i^{(h)}(0) \right\|_2^2 = \frac{c_\sigma}{m} \sum_{r=1}^m \sigma \left(\mathbf{w}_r^{(h)}(0)^\top \mathbf{x}_i^{(h-1)}(0) \right)^2.$$

Therefore we have

$$\mathbb{E} \left[\left\| \mathbf{x}_i^{(h)}(0) \right\|_2^2 \right] = c_\sigma \mathbb{E} \left[\sigma \left(\mathbf{w}_r^{(h)}(0)^\top \mathbf{x}_i^{(h-1)}(0) \right)^2 \right]$$

$$= c_\sigma \mathbb{E}_{X \sim N(0,1)} \sigma \left(\left\| \mathbf{x}_i^{(h-1)}(0) \right\|_2 X \right)^2.$$

Note that $\sigma(\cdot)$ is L -Lipschitz, for any $\frac{1}{2} \leq \alpha \leq 2$, we have

$$\begin{aligned} & \left| \mathbb{E}_{X \sim N(0,1)} \sigma(\alpha X)^2 - \mathbb{E}_{X \sim N(0,1)} \sigma(X)^2 \right| \\ & \leq \mathbb{E}_{X \sim N(0,1)} \left| \sigma(\alpha X)^2 - \sigma(X)^2 \right| \\ & \leq L |\alpha - 1| \mathbb{E}_{X \sim N(0,1)} |X (\sigma(\alpha X) + \sigma(X))| \\ & \leq L |\alpha - 1| \mathbb{E}_{X \sim N(0,1)} |X| (|2\sigma(0)| + L |(\alpha + 1)X|) \\ & \leq L |\alpha - 1| (2|\sigma(0)| \mathbb{E}_{X \sim N(0,1)} |X| + L |\alpha + 1| \mathbb{E}_{X \sim N(0,1)} X^2) \\ & = L |\alpha - 1| \left(2|\sigma(0)| \sqrt{\frac{2}{\pi}} + L |\alpha + 1| \right) \\ & \leq \frac{C}{c_\sigma} |\alpha - 1|, \end{aligned}$$

where $C \triangleq c_\sigma L \left(2|\sigma(0)| \sqrt{\frac{2}{\pi}} + 2L \right)$, which implies

$$1 - \frac{C g_C(h-1)}{2g_C(H)} \leq \mathbb{E} \left[\left\| \mathbf{x}_i^{(h)}(0) \right\|_2^2 \right] \leq 1 + \frac{C g_C(h-1)}{2g_C(H)}.$$

For the variance we have

$$\begin{aligned} \text{Var} \left[\left\| \mathbf{x}_i^{(h)}(0) \right\|_2^2 \right] &= \frac{c_\sigma^2}{m} \text{Var} \left[\sigma \left(\mathbf{w}_r^{(h)}(0)^\top \mathbf{x}_i^{(h-1)}(0) \right)^2 \right] \\ &\leq \frac{c_\sigma^2}{m} \mathbb{E} \left[\sigma \left(\mathbf{w}_r^{(h)}(0)^\top \mathbf{x}_i^{(h-1)}(0) \right)^4 \right] \\ &\leq \frac{c_\sigma^2}{m} \mathbb{E} \left[\left(|\sigma(0)| + L \left| \mathbf{w}_r^{(h)}(0)^\top \mathbf{x}_i^{(h-1)}(0) \right| \right)^4 \right] \\ &\leq \frac{C_2}{m}. \end{aligned}$$

where $C_2 \triangleq \sigma(0)^4 + 8|\sigma(0)|^3 L \sqrt{2/\pi} + 24\sigma(0)^2 L^2 + 64\sigma(0) L^3 \sqrt{2/\pi} + 512L^4$ and the last inequality we used the formula for the first four absolute moments of Gaussian.

Applying Chebyshev's inequality and plugging in our assumption on m , we have with probability $1 - \frac{\delta}{nH}$ over $\mathbf{W}^{(h)}$,

$$\left| \left\| \mathbf{x}_i^{(h)}(0) \right\|_2^2 - \mathbb{E} \left\| \mathbf{x}_i^{(h)}(0) \right\|_2^2 \right| \leq \frac{1}{2g_C(H)}.$$

Thus with probability $1 - h \frac{\delta}{nH}$ over $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(h)}$,

$$\left| \left\| \mathbf{x}_i^{(h)}(0) \right\|_2 - 1 \right| \leq \left| \left\| \mathbf{x}_i^{(h)}(0) \right\|_2^2 - 1 \right| \leq \frac{C g_C(h-1)}{2g_C(H)} + \frac{1}{2g(H)} = \frac{g_C(h)}{2g_C(H)}.$$

Using union bounds over $[n]$, we prove the lemma. □

Proof of Lemma B.3. We prove this lemma by induction. Our induction hypothesis is

$$\left\| \mathbf{x}^{(h)}(k) - \mathbf{x}^{(h)}(0) \right\|_2 \leq \sqrt{c_\sigma} L R c_{x,0} g_{c_x}(h),$$

where

$$c_x = 2\sqrt{c_\sigma} L c_{w,0}.$$

For $h = 0$, since the input data is fixed, we know the induction hypothesis holds. Now suppose the induction hypothesis holds for $h' = 0, \dots, h-1$, we consider $h' = h$.

$$\begin{aligned}
 \left\| \mathbf{x}^{(h)}(k) - \mathbf{x}^{(h)}(0) \right\|_2 &= \sqrt{\frac{c_\sigma}{m}} \left\| \sigma \left(\mathbf{W}^{(h)}(k) \mathbf{x}^{(h-1)}(k) \right) - \sigma \left(\mathbf{W}^{(h)}(0) \mathbf{x}^{(h-1)}(0) \right) \right\|_2 \\
 &\leq \sqrt{\frac{c_\sigma}{m}} \left\| \sigma \left(\mathbf{W}^{(h)}(k) \mathbf{x}^{(h-1)}(k) \right) - \sigma \left(\mathbf{W}^{(h)}(k) \mathbf{x}^{(h-1)}(0) \right) \right\|_2 \\
 &\quad + \sqrt{\frac{c_\sigma}{m}} \left\| \sigma \left(\mathbf{W}^{(h)}(k) \mathbf{x}^{(h-1)}(0) \right) - \sigma \left(\mathbf{W}^{(h)}(0) \mathbf{x}^{(h-1)}(0) \right) \right\|_2 \\
 &\leq \sqrt{\frac{c_\sigma}{m}} L \left(\left\| \mathbf{W}^{(h)}(0) \right\|_2 + \left\| \mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0) \right\|_F \right) \cdot \left\| \mathbf{x}^{(h-1)}(k) - \mathbf{x}^{(h-1)}(0) \right\|_2 \\
 &\quad + \sqrt{\frac{c_\sigma}{m}} L \left\| \mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0) \right\|_F \left\| \mathbf{x}^{(h-1)}(0) \right\|_2 \\
 &\leq \sqrt{\frac{c_\sigma}{m}} L (c_{w,0} \sqrt{m} + R \sqrt{m}) \sqrt{c_\sigma} L R c_{x,0} g_{c_x} (h-1) + \sqrt{\frac{c_\sigma}{m}} L \sqrt{m} R c_{x,0} \\
 &\leq \sqrt{c_\sigma} L R c_{x,0} (c_x g_{c_x} (h-1) + 1) \\
 &\leq \sqrt{c_\sigma} L R c_{x,0} g_{c_x} (h).
 \end{aligned}$$

□

Proof of Lemma B.4. Because Frobenius-norm of a matrix is bigger than the operator norm, it is sufficient to bound $\left\| \mathbf{G}^{(H)}(k) - \mathbf{G}^{(H)}(0) \right\|_F$. For simplicity define $z_{i,r}(k) = \mathbf{w}_r^{(H)}(k)^\top \mathbf{x}_i^{(H-1)}(k)$, we have

$$\begin{aligned}
 &\left| \mathbf{G}_{i,j}^{(H)}(k) - \mathbf{G}_{i,j}^{(H)}(0) \right| \\
 &= \left| \mathbf{x}_i^{(H-1)}(k)^\top \mathbf{x}_j^{(H-1)}(k) \frac{c_\sigma}{m} \sum_{r=1}^m a_r(k)^2 \sigma'(z_{i,r}(k)) \sigma'(z_{j,r}(k)) \right. \\
 &\quad \left. - \mathbf{x}_i^{(H-1)}(0)^\top \mathbf{x}_j^{(H-1)}(0) \frac{c_\sigma}{m} \sum_{r=1}^m a_r(0)^2 \sigma'(z_{i,r}(0)) \sigma'(z_{j,r}(0)) \right| \\
 &\leq \left| \mathbf{x}_i^{(H-1)}(k)^\top \mathbf{x}_j^{(H-1)}(k) - \mathbf{x}_i^{(H-1)}(0)^\top \mathbf{x}_j^{(H-1)}(0) \right| \frac{c_\sigma}{m} \sum_{r=1}^m a_r(0)^2 \left| \sigma'(z_{i,r}(k)) \sigma'(z_{j,r}(k)) \right| \\
 &\quad + \left| \mathbf{x}_i^{(H-1)}(0)^\top \mathbf{x}_j^{(H-1)}(0) \right| \frac{c_\sigma}{m} \left| \sum_{r=1}^m a_r(0)^2 \left(\sigma'(z_{i,r}(k)) \sigma'(z_{j,r}(k)) - \sigma'(z_{i,r}(0)) \sigma'(z_{j,r}(0)) \right) \right| \\
 &\quad + \left| \mathbf{x}_i^{(H-1)}(k)^\top \mathbf{x}_j^{(H-1)}(k) \right| \frac{c_\sigma}{m} \left| \sum_{r=1}^m (a_r(k)^2 - a_r(0)^2) \sigma'(z_{i,r}(k)) \sigma'(z_{j,r}(k)) \right| \\
 &\leq L^2 c_\sigma a_{2,0}^2 \left| \mathbf{x}_i^{(H-1)}(k)^\top \mathbf{x}_j^{(H-1)}(k) - \mathbf{x}_i^{(H-1)}(0)^\top \mathbf{x}_j^{(H-1)}(0) \right| \\
 &\quad + c_{x,0}^2 \frac{c_\sigma}{m} \left| \sum_{r=1}^m a_r(0)^2 \left(\sigma'(z_{i,r}(k)) \sigma'(z_{j,r}(k)) - \sigma'(z_{i,r}(0)) \sigma'(z_{j,r}(0)) \right) \right| \\
 &\quad + 4L^2 c_{x,0}^2 \frac{c_\sigma}{m} \sum_{r=1}^m |a_r(k)^2 - a_r(0)^2| \\
 &\triangleq I_1^{i,j} + I_2^{i,j} + I_3^{i,j}.
 \end{aligned}$$

For $I_1^{i,j}$, using Lemma B.3, we have

$$I_1^{i,j} = L^2 c_\sigma a_{2,0}^2 \left| \mathbf{x}_i^{(H-1)}(k)^\top \mathbf{x}_j^{(H-1)}(k) - \mathbf{x}_i^{(H-1)}(0)^\top \mathbf{x}_j^{(H-1)}(0) \right|$$

$$\begin{aligned}
 &\leq L^2 c_\sigma a_{2,0}^2 \left| (\mathbf{x}_i^{(H-1)}(k) - \mathbf{x}_i^{(H-1)}(0))^\top \mathbf{x}_j^{(H-1)}(k) \right| + L^2 c_\sigma a_{2,0}^2 \left| \mathbf{x}_i^{(H-1)}(0)^\top (\mathbf{x}_j^{(H-1)}(k) - \mathbf{x}_j^{(H-1)}(0)) \right| \\
 &\leq c_\sigma a_{2,0}^2 \sqrt{c_\sigma} L^3 c_{x,0} g_{c_x}(H) R \cdot (c_{x,0} + \sqrt{c_\sigma} L c_{x,0} g_{c_x}(H) R) + c_\sigma \sqrt{c_\sigma} a_{2,0}^2 L^3 c_{x,0} g_{c_x}(H) R c_{x,0} \\
 &\leq 3c_\sigma a_{2,0}^2 c_{x,0}^2 \sqrt{c_\sigma} L^3 g_{c_x}(H) R.
 \end{aligned}$$

For $I_2^{i,j}$, we have

$$\begin{aligned}
 I_2^{i,j} &= c_{x,0}^2 \frac{c_\sigma}{m} \left| \sum_{r=1}^m a_r(0)^2 \sigma'(z_{i,r}(k)) \sigma'(z_{j,r}(k)) - a_r(0)^2 \sigma'(z_{i,r}(0)) \sigma'(z_{j,r}(0)) \right| \\
 &\leq c_{x,0}^2 \frac{c_\sigma}{m} \sum_{r=1}^m a_r(0)^2 \left| (\sigma'(z_{i,r}(k)) - \sigma'(z_{i,r}(0))) \sigma'(z_{j,r}(k)) \right| + a_r(0)^2 \left| (\sigma'(z_{j,r}(k)) - \sigma'(z_{j,r}(0))) \sigma'(z_{i,r}(0)) \right| \\
 &\leq \frac{\beta L c_\sigma c_{x,0}^2}{m} \left(\sum_{r=1}^m a_r(0)^2 |z_{i,r}(k) - z_{i,r}(0)| + a_r(0)^2 |z_{j,r}(k) - z_{j,r}(0)| \right) \\
 &\leq \frac{\beta L c_\sigma a_{4,0}^2 c_{x,0}^2}{\sqrt{m}} \left(\sqrt{\sum_{r=1}^m |z_{i,r}(k) - z_{i,r}(0)|^2} + \sqrt{\sum_{r=1}^m |z_{j,r}(k) - z_{j,r}(0)|^2} \right).
 \end{aligned}$$

Using the same proof for Lemma B.3, it is easy to see

$$\sum_{r=1}^m |z_{i,r}(t) - z_{i,r}(0)|^2 \leq c_{x,0}^2 g_{c_x}(H)^2 m R^2.$$

Thus

$$I_2^{i,j} \leq 2\beta c_\sigma a_{4,0}^2 c_{x,0}^3 L g_{c_x}(H) R.$$

For $I_3^{i,j}$,

$$\begin{aligned}
 I_3^{i,j} &= 4L^2 c_{x,0}^2 \frac{c_\sigma}{m} \sum_{r=1}^m |a_r(k)^2 - a_r(0)^2| \\
 &\leq 4L^2 c_{x,0}^2 \frac{c_\sigma}{m} \sum_{r=1}^m |a_r(k) - a_r(0)| |a_r(k)| + |a_r(k) - a_r(0)| |a_r(0)| \\
 &\leq 12L^2 c_{x,0}^2 c_\sigma a_{2,0} R.
 \end{aligned}$$

Therefore we can bound the perturbation

$$\begin{aligned}
 \left\| \mathbf{G}^{(H)}(t) - \mathbf{G}^{(H)}(0) \right\|_F &= \sqrt{\sum_{(i,j)}^{n,n} \left| \mathbf{G}_{i,j}^{(H)}(t) - \mathbf{G}_{i,j}^{(H)}(0) \right|^2} \\
 &\leq \left[(2\beta c_{x,0} a_{4,0}^2 + 3\sqrt{c_\sigma} L^2) L c_\sigma c_{x,0}^2 a_{2,0} g_{c_x}(H) + 12L^2 c_{x,0}^2 c_\sigma a_{2,0} \right] nR.
 \end{aligned}$$

Plugging in the bound on R , we have the desired result. \square

Proof of Lemma B.5. We will prove this corollary by induction. The induction hypothesis is

$$\left\| \mathbf{W}^{(h)}(s) - \mathbf{W}^{(h)}(0) \right\|_F \leq \sum_{s'=0}^{s-1} \left(1 - \frac{\eta \lambda_0}{2}\right)^{s'/2} \frac{1}{4} \eta \lambda_0 R' \sqrt{m} \leq R' \sqrt{m}, s \in [k+1],$$

$$\|\mathbf{a}(s) - \mathbf{a}(0)\|_2 \leq \sum_{s'=0}^{s-1} \left(1 - \frac{\eta\lambda_0}{2}\right)^{s'/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m} \leq R' \sqrt{m}, s \in [k+1].$$

First it is easy to see it holds for $s' = 0$. Now suppose it holds for $s' = 0, \dots, s$, we consider $s' = s+1$. We have

$$\begin{aligned} & \left\| \mathbf{W}^{(h)}(s+1) - \mathbf{W}^{(h)}(s) \right\|_F \\ &= \eta \left\| \left(\frac{c_\sigma}{m} \right)^{\frac{H-h+1}{2}} \sum_{i=1}^n (y_i - u_i(s)) \mathbf{x}_i^{(h-1)}(s) \left(\mathbf{a}(s)^\top \left(\prod_{k=h+1}^H \mathbf{J}_i^{(k)}(s) \mathbf{W}^{(k)}(s) \right) \mathbf{J}_i^{(h)}(s) \right) \right\|_F \\ &\leq \eta \left(\frac{c_\sigma}{m} \right)^{\frac{H-h+1}{2}} \|\mathbf{a}(s)\|_2 \sum_{i=1}^n |y_i - u_i(s)| \left\| \mathbf{x}_i^{(h-1)}(s) \right\|_2 \prod_{k=h+1}^H \left\| \mathbf{W}^{(k)}(s) \right\|_2 \prod_{k=h}^H \left\| \mathbf{J}^{(k)}(s) \right\|_2, \\ & \|\mathbf{a}(s+1) - \mathbf{a}(s)\|_2 = \eta \left\| \sum_{i=1}^n (y_i - u_i(s)) \mathbf{x}_i^{(H)}(s) \right\|_2. \end{aligned}$$

To bound $\left\| \mathbf{x}_i^{(h-1)}(s) \right\|_2$, we can just apply Lemma B.3 and get

$$\left\| \mathbf{x}_i^{(h-1)}(s) \right\|_2 \leq \sqrt{c_\sigma} L c_{x,0} g_{c_x}(h) R' + c_{x,0} \leq 2c_{x,0}.$$

To bound $\left\| \mathbf{W}^{(k)}(s) \right\|_2$, we use our assumption

$$\begin{aligned} \prod_{k=h+1}^H \left\| \mathbf{W}^{(k)}(s) \right\|_2 &\leq \prod_{k=h+1}^H \left(\left\| \mathbf{W}^{(k)}(0) \right\|_2 + \left\| \mathbf{W}^{(k)}(s) - \mathbf{W}^{(k)}(0) \right\|_2 \right) \\ &\leq \prod_{k=h+1}^H (c_{w,0} \sqrt{m} + R' \sqrt{m}) \\ &= (c_{w,0} + R')^{H-h} m^{\frac{H-h}{2}} \\ &\leq (2c_{w,0})^{H-h} m^{\frac{H-h}{2}}. \end{aligned}$$

Note that $\left\| \mathbf{J}^{(k)}(s) \right\|_2 \leq L$. Plugging in these two bounds back, we obtain

$$\begin{aligned} \left\| \mathbf{W}^{(h)}(s+1) - \mathbf{W}^{(h)}(s) \right\|_F &\leq 4\eta c_{x,0} a_{2,0} c_x^H \sum_{i=1}^n |y_i - u(s)| \\ &\leq 4\eta c_{x,0} a_{2,0} c_x^H \sqrt{n} \|\mathbf{y} - \mathbf{u}(s)\|_2 \\ &= \eta Q'(s) \\ &\leq \left(1 - \frac{\eta\lambda_0}{2}\right)^{s/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m}. \end{aligned}$$

Similarly, we have

$$\begin{aligned} \|\mathbf{a}(s+1) - \mathbf{a}(s)\|_2 &\leq 2\eta c_{x,0} \sum_{i=1}^n |y_i - u(s)| \\ &\leq \eta Q'(s) \\ &\leq \left(1 - \frac{\eta\lambda_0}{2}\right)^{s/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m}. \end{aligned}$$

Thus

$$\left\| \mathbf{W}^{(h)}(s+1) - \mathbf{W}^{(h)}(0) \right\|_F$$

$$\begin{aligned} &\leq \left\| \mathbf{W}^{(h)}(s+1) - \mathbf{W}^{(h)}(s) \right\|_F + \left\| \mathbf{W}^{(h)}(s) - \mathbf{W}^{(h)}(0) \right\|_F \\ &\leq \sum_{s'=0}^s \eta \left(1 - \frac{\eta\lambda_0}{2}\right)^{s'/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m}. \end{aligned}$$

Similarly,

$$\begin{aligned} &\left\| \mathbf{a}(s+1) - \mathbf{a}(0) \right\|_2 \\ &\leq \sum_{s'=0}^s \eta \left(1 - \frac{\eta\lambda_0}{2}\right)^{s'/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m}. \end{aligned}$$

□

Proof of Lemma B.6. Fix $i \in [n]$, we bound

$$|I_2^i(k)| \leq \eta \max_{0 \leq s \leq \eta} \sum_{h=1}^H \left\| L^{(h)}(\theta(k)) \right\|_F \left\| u_i^{(h)}(\theta(k)) - u_i^{(h)}(\theta(k) - sL^{(h)}(\theta(k))) \right\|_F.$$

For the gradient norm, we have

$$\begin{aligned} &\left\| L^{(h)}(\theta(k)) \right\|_F \\ &= \left\| \left(\frac{c_\sigma}{m} \right)^{\frac{H-h+1}{2}} \sum_{i=1}^n (y_i - u_i(k)) \mathbf{x}_i^{(h-1)}(k) \left(\mathbf{a}(k)^\top \left(\prod_{l=h+1}^H \mathbf{J}_i^{(l)}(k) \mathbf{W}^{(l)}(k) \right) \mathbf{J}_i^{(h)}(k) \right) \right\|_F. \end{aligned}$$

Similar to the proof for Lemma B.5, we have

$$\left\| L^{(h)}(\theta(k)) \right\|_F \leq Q'(k).$$

Let $\theta(k, s) = \theta(k) - sL'(\theta(k))$,

$$\begin{aligned} &\left\| u_i^{(h)}(\theta(k)) - u_i^{(h)}(\theta(k, s)) \right\|_F \\ &= \left(\frac{c_\sigma}{m} \right)^{\frac{H-h+1}{2}} \left\| \mathbf{x}_i^{(h-1)}(k) \left(\mathbf{a}(k)^\top \left(\prod_{l=h+1}^H \mathbf{J}_i^{(l)}(k) \mathbf{W}^{(l)}(k) \right) \mathbf{J}_i^{(h)}(k) \right) \right. \\ &\quad \left. - \mathbf{x}_i^{(h-1)}(k, s) \left(\mathbf{a}(k, s)^\top \left(\prod_{l=h+1}^H \mathbf{J}_i^{(l)}(k, s) \mathbf{W}^{(l)}(k, s) \right) \mathbf{J}_i^{(h)}(k, s) \right) \right\|_F \end{aligned}$$

Through standard calculations, we have

$$\begin{aligned} &\left\| \mathbf{W}^{(l)}(k) - \mathbf{W}^{(l)}(k, s) \right\|_F \leq \eta Q'(k), \\ &\left\| \mathbf{a}(k) - \mathbf{a}(k, s) \right\|_2 \leq \eta Q'(k), \\ &\left\| \mathbf{x}_i^{(h-1)}(k) - \mathbf{x}_i^{(h-1)}(k, s) \right\|_F \leq 2\eta \sqrt{c_\sigma} L c_{x,0} g_{2c_x}(H) \frac{Q'(k)}{\sqrt{m}}, \\ &\left\| \mathbf{J}_i^{(l)}(k) - \mathbf{J}_i^{(l)}(k, s) \right\|_F \leq 2\eta \beta \sqrt{c_\sigma} L c_{x,0} g_{2c_x}(H) Q'(k). \end{aligned}$$

According to Lemma G.1, we have

$$\left\| u_i^{(h)}(\mathbf{w}(k)) - u_i^{(h)}(\mathbf{w}(k, s)) \right\|_F$$

$$\begin{aligned} &\leq 4c_{x,0}a_{2,0}c_x^H\eta\frac{Q'(k)}{\sqrt{m}}\left(\frac{H}{2}+\left[\frac{1}{2c_{x,0}}+\frac{H\beta\sqrt{m}}{L}\right]2\sqrt{c_\sigma}Lc_{x,0}g_{2c_x}(H)\right) \\ &\leq 16H\sqrt{c_\sigma}c_{x,0}^2a_{2,0}c_x^Hg_{2c_x}(H)\beta\eta Q'(k). \end{aligned}$$

Thus we have

$$|I_2^i| \leq 16H^2\sqrt{c_\sigma}c_{x,0}^2a_{2,0}c_x^Hg_{2c_x}(H)\beta\eta^2Q'(k)^2.$$

Since this holds for all $i \in [n]$, plugging in η and noting that $\|\mathbf{y} - \mathbf{u}(0)\|_2 = O(\sqrt{n})$, we have

$$\|\mathbf{I}_2(k)\|_2 \leq \frac{1}{8}\eta\lambda_0\|\mathbf{y} - \mathbf{u}(k)\|_2.$$

□

Proof of Lemma B.7.

$$\begin{aligned} \|\mathbf{u}(k+1) - \mathbf{u}(k)\|_2^2 &= \sum_{i=1}^n \left(\mathbf{a}(k+1)^\top \mathbf{x}_i^{(H)}(k+1) - \mathbf{a}(k)^\top \mathbf{x}_i^{(H)}(k) \right)^2 \\ &= \sum_{i=1}^n \left([\mathbf{a}(k+1) - \mathbf{a}(k)]^\top \mathbf{x}_i^{(H)}(k+1) + \mathbf{a}(k)^\top [\mathbf{x}_i^{(H)}(k+1) - \mathbf{x}_i^{(H)}(k)] \right)^2 \\ &\leq 2\|\mathbf{a}(k+1) - \mathbf{a}(k)\|_2^2 \sum_{i=1}^n \left\| \mathbf{x}_i^{(H)}(k+1) \right\|_2^2 + 2\|\mathbf{a}(k)\|_2^2 \sum_{i=1}^n \left\| \mathbf{x}_i^{(H)}(k+1) - \mathbf{x}_i^{(H)}(k) \right\|_2^2 \\ &\leq 8n\eta^2c_{x,0}^2Q'(k)^2 + 4n(2\eta\sqrt{c_\sigma}Lc_{x,0}a_{2,0}g_{2c_x}(H)Q'(k))^2 \\ &\leq \frac{1}{8}\eta\lambda_0\|\mathbf{y} - \mathbf{u}(k)\|_2^2. \end{aligned}$$

□

C. Proofs for Section 6

The gradient for ResNet is

$$\frac{\partial L}{\partial \mathbf{W}^{(h)}} = \frac{c_{res}}{H\sqrt{m}} \sum_{i=1}^n (y_i - u_i) \mathbf{x}_i^{(h-1)} \cdot \left[\mathbf{a}^\top \prod_{l=h+1}^H \left(\mathbf{I} + \frac{c_{res}}{H\sqrt{m}} \mathbf{J}_i^{(l)} \mathbf{W}^{(l)} \right) \mathbf{J}_i^{(h)} \right]$$

For ResNets, $\mathbf{G}^{(H)}$ has the following form:

$$\mathbf{G}_{ij}^{(H)} = \frac{c_{res}^2}{H^2m} (\mathbf{x}_i^{(H-1)})^\top \mathbf{x}_j^{(H-1)} \sum_{r=1}^m a_r^2 \sigma'((\mathbf{w}_r^{(H)})^\top \mathbf{x}_i^{(H-1)}) \sigma'((\mathbf{w}_r^{(H)})^\top \mathbf{x}_j^{(H-1)}). \quad (14)$$

Similar to Lemma B.1, we can show with high probability the feature of each layer is approximately normalized.

Lemma C.1 (Lemma on Initialization Norms). *If $\sigma(\cdot)$ is L -Lipschitz and $m = \Omega(\frac{n}{\delta})$, assuming $\|\mathbf{W}^{(h)}(0)\|_2 \leq c_{w,0}\sqrt{m}$ for $h \in [2, H]$ and $c_{w,0} \approx 2$ for Gaussian initialization. We have with probability at least $1 - \delta$ over random initialization, for every $h \in [H]$ and $i \in [n]$,*

$$\frac{1}{c_{x,0}} \leq \left\| \mathbf{x}_i^{(h)}(0) \right\|_2 \leq c_{x,0}$$

for some universal constant $c_{x,0} > 1$ (only depends on σ).

The following lemma lower bounds $\mathbf{G}^{(H)}(0)$'s least eigenvalue. This lemma is a direct consequence of results in Section E.

Lemma C.2 (Least Eigenvalue at the Initialization). *If $m = \Omega\left(\frac{n^2 \log(Hn/\delta)}{\lambda_0^2}\right)$, we have*

$$\lambda_{\min}(\mathbf{G}^{(H)}(0)) \geq \frac{3}{4}\lambda_0.$$

Next, we characterize how the perturbation on the weight matrices affects the input of each layer.

Lemma C.3. *Suppose $\sigma(\cdot)$ is L -Lipschitz and for $h \in [H]$, $\|\mathbf{W}^{(h)}(0)\|_2 \leq c_{w,0}\sqrt{m}$, $\|\mathbf{x}^{(h)}(0)\|_2 \leq c_{x,0}$ and $\|\mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0)\|_F \leq \sqrt{m}R$ for some constant $c_{w,0}, c_{x,0} > 0$ and $R \leq c_{w,0}$. Then we have*

$$\|\mathbf{x}^{(h)}(k) - \mathbf{x}^{(h)}(0)\|_2 \leq \left(\sqrt{c_\sigma}L + \frac{c_{x,0}}{c_{w,0}}\right) e^{2c_{res}c_{w,0}L} R.$$

Next, we characterize how the perturbation on the weight matrices affect $\mathbf{G}^{(H)}$.

Lemma C.4. *Suppose $\sigma(\cdot)$ is differentiable, L -Lipschitz and β -smooth. Using the same notations in Lemma B.4, if $\|\mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0)\|_F, \|\mathbf{a}(k) - \mathbf{a}(0)\|_2 \leq \sqrt{m}R$ where $R \leq c\lambda_0 H^2 n^{-1}$ and $R \leq c$ for some small constant c , we have*

$$\|\mathbf{G}^{(H)}(k) - \mathbf{G}^{(H)}(0)\|_2 \leq \frac{\lambda_0}{2}.$$

We prove Theorem 6.1 by induction. Our induction hypothesis is just the following convergence rate of empirical loss.

A directly corollary of this condition is the following bound of deviation from the initialization. The proof only involves standard calculations so we defer it to appendix.

Lemma C.5. *If Condition A.1 holds for $k' = 1, \dots, k$, we have for any $s \in [k+1]$*

$$\begin{aligned} \|\mathbf{W}^{(h)}(s) - \mathbf{W}^{(h)}(0)\|_F, \|\mathbf{a}(s) - \mathbf{a}(0)\|_2 &\leq R' \sqrt{m}, \\ \|\mathbf{W}^{(h)}(s) - \mathbf{W}^{(h)}(s-1)\|_F, \|\mathbf{a}(s) - \mathbf{a}(s-1)\|_2 &\leq \eta Q'(s-1), \end{aligned}$$

where $R' = \frac{16c_{res}c_{x,0}a_{2,0}Le^{2c_{res}c_{w,0}L}\sqrt{n}\|\mathbf{y} - \mathbf{u}(0)\|_2}{H\lambda_0\sqrt{m}} < c$ for some small constant c and $Q'(s) = 4c_{res}c_{x,0}a_{2,0}Le^{2c_{res}c_{w,0}L}\sqrt{n}\|\mathbf{y} - \mathbf{u}(s)\|_2/H$.

The next lemma bounds the \mathbf{I}_2 term.

Lemma C.6. *If Condition A.1 holds for $k' = 1, \dots, k$ and $\eta \leq c\lambda_0 H^2 n^{-2}$ for some small constant c , we have*

$$\|\mathbf{I}_2(k)\|_2 \leq \frac{1}{8}\eta\lambda_0 \|\mathbf{y} - \mathbf{u}(k)\|_2.$$

Next we bound the quadratic term.

Lemma C.7. *If Condition A.1 holds for $k' = 1, \dots, k$ and $\eta \leq c\lambda_0 H^2 n^{-2}$ for some small constant c , we have $\|\mathbf{u}(k+1) - \mathbf{u}(k)\|_2^2 \leq \frac{1}{8}\eta\lambda_0 \|\mathbf{y} - \mathbf{u}(k)\|_2^2$.*

Now using the same argument as in the proof for multilayer fully connected neural network, we finish our proof for ResNet.

C.1. Proofs of Lemmas

Proof of Lemma C.1. We will bound $\|\mathbf{x}_i^{(h)}(0)\|_2$ layer by layer. For the first layer, we can calculate

$$\begin{aligned} \mathbb{E} \left[\left\| \mathbf{x}_i^{(1)}(0) \right\|_2^2 \right] &= c_\sigma \mathbb{E} \left[\sigma \left(\mathbf{w}_r^{(1)}(0)^\top \mathbf{x}_i \right)^2 \right] \\ &= c_\sigma \mathbb{E}_{X \sim N(0,1)} \sigma(X)^2 \\ &= 1. \end{aligned}$$

$$\begin{aligned}
 \text{Var} \left[\left\| \mathbf{x}_i^{(1)}(0) \right\|_2^2 \right] &= \frac{c_\sigma^2}{m} \text{Var} \left[\sigma \left(\mathbf{w}_r^{(1)}(0)^\top \mathbf{x}_i(0) \right)^2 \right] \\
 &\leq \frac{c_\sigma^2}{m} \mathbb{E}_{X \sim N(0,1)} \sigma(X)^4 \\
 &\leq \frac{c_\sigma^2}{m} \mathbb{E} \left[\left(|\sigma(0)| + L \left| \mathbf{w}_r^{(1)}(0)^\top \mathbf{x}_i \right| \right)^4 \right] \\
 &\leq \frac{C_2}{m},
 \end{aligned}$$

where $C_2 \triangleq \sigma(0)^4 + 4|\sigma(0)|^3 L\sqrt{2/\pi} + 6\sigma(0)^2 L^2 + 8|\sigma(0)| L^3\sqrt{2/\pi} + 32L^4$. We have with probability at least $1 - \frac{\delta}{n}$,

$$\frac{1}{2} \leq \left\| \mathbf{x}_i^{(1)}(0) \right\|_2 \leq 2.$$

By definition we have for $2 \leq h \leq H$,

$$\begin{aligned}
 \left\| \mathbf{x}_i^{(h-1)}(0) \right\|_2 - \left\| \frac{c_{res}}{H\sqrt{m}} \sigma \left(\mathbf{W}^{(h)}(0) \mathbf{x}_i^{(h-1)}(0) \right) \right\|_2 &\leq \left\| \mathbf{x}^{(h)}(0) \right\|_2 \\
 &\leq \left\| \mathbf{x}_i^{(h-1)}(0) \right\|_2 + \left\| \frac{c_{res}}{H\sqrt{m}} \sigma \left(\mathbf{W}^{(h)}(0) \mathbf{x}^{(h-1)}(0) \right) \right\|_2,
 \end{aligned}$$

where

$$\left\| \frac{c_{res}}{H\sqrt{m}} \sigma \left(\mathbf{W}^{(h)}(0) \mathbf{x}_i^{(h-1)}(0) \right) \right\|_2 \leq \frac{c_{res} c_{w,0} L}{H} \left\| \mathbf{x}_i^{(h-1)}(0) \right\|_2.$$

Thus

$$\left\| \mathbf{x}_i^{(h-1)}(0) \right\|_2 \left(1 - \frac{c_{res} c_{w,0} L}{H} \right) \leq \left\| \mathbf{x}^{(h)}(0) \right\|_2 \leq \left\| \mathbf{x}_i^{(h-1)}(0) \right\|_2 \left(1 + \frac{c_{res} c_{w,0} L}{H} \right),$$

which implies

$$\frac{1}{2} e^{-c_{res} c_{w,0} L} \leq \left\| \mathbf{x}^{(h)}(0) \right\|_2 \leq 2 e^{c_{res} c_{w,0} L}.$$

Choosing $c_{x,0} = 2e^{c_{res} c_{w,0} L}$ and using union bounds over $[n]$, we prove the lemma. □

Proof of Lemma C.3. We prove this lemma by induction. Our induction hypothesis is

$$\left\| \mathbf{x}^{(h)}(k) - \mathbf{x}^{(h)}(0) \right\|_2 \leq g(h),$$

where

$$g(h) = g(h-1) \left[1 + \frac{2c_{res} c_{w,0} L}{H} \right] + \frac{L}{H} R c_{x,0}.$$

For $h = 1$, we have

$$\begin{aligned}
 \left\| \mathbf{x}^{(1)}(k) - \mathbf{x}^{(1)}(0) \right\|_2 &\leq \sqrt{\frac{c_\sigma}{m}} \left\| \sigma \left(\mathbf{W}^{(1)}(k) \mathbf{x} \right) - \sigma \left(\mathbf{W}^{(1)}(0) \mathbf{x} \right) \right\|_2 \\
 &\leq \sqrt{\frac{c_\sigma}{m}} \left\| \mathbf{W}^{(1)}(k) - \mathbf{W}^{(1)}(0) \right\|_F \leq \sqrt{c_\sigma} L R,
 \end{aligned}$$

which implies $g(1) = \sqrt{c_\sigma} L R$, for $2 \leq h \leq H$, we have

$$\begin{aligned}
 \left\| \mathbf{x}^{(h)}(k) - \mathbf{x}^{(h)}(0) \right\|_2 &\leq \frac{c_{res}}{H\sqrt{m}} \left\| \sigma \left(\mathbf{W}^{(h)}(k) \mathbf{x}^{(h-1)}(k) \right) - \sigma \left(\mathbf{W}^{(h)}(0) \mathbf{x}^{(h-1)}(0) \right) \right\|_2 \\
 &+ \left\| \mathbf{x}^{(h-1)}(k) - \mathbf{x}^{(h-1)}(0) \right\|_2
 \end{aligned}$$

$$\begin{aligned}
 &\leq \frac{c_{res}}{H\sqrt{m}} \left\| \sigma \left(\mathbf{W}^{(h)}(k) \mathbf{x}^{(h-1)}(k) \right) - \sigma \left(\mathbf{W}^{(h)}(k) \mathbf{x}^{(h-1)}(0) \right) \right\|_2 \\
 &+ \frac{c_{res}}{H\sqrt{m}} \left\| \sigma \left(\mathbf{W}^{(h)}(k) \mathbf{x}^{(h-1)}(0) \right) - \sigma \left(\mathbf{W}^{(h)}(0) \mathbf{x}^{(h-1)}(0) \right) \right\|_2 \\
 &+ \left\| \mathbf{x}^{(h-1)}(k) - \mathbf{x}^{(h-1)}(0) \right\|_2 \\
 &\leq \frac{c_{res}L}{H\sqrt{m}} \left(\left\| \mathbf{W}^{(h)}(0) \right\|_2 + \left\| \mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0) \right\|_F \right) \cdot \left\| \mathbf{x}^{(h-1)}(k) - \mathbf{x}^{(h-1)}(0) \right\|_2 \\
 &+ \frac{c_{res}L}{H\sqrt{m}} \left\| \mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0) \right\|_F \left\| \mathbf{x}^{(h-1)}(0) \right\|_2 + \left\| \mathbf{x}^{(h-1)}(k) - \mathbf{x}^{(h-1)}(0) \right\|_2 \\
 &\leq \left[1 + \frac{c_{res}L}{H\sqrt{m}} (c_{w,0}\sqrt{m} + R\sqrt{m}) \right] g(h-1) + \frac{c_{res}L}{H\sqrt{m}} \sqrt{m} R c_{x,0} \\
 &\leq \left(1 + \frac{2c_{res}c_{w,0}L}{H} \right) g(h-1) + \frac{c_{res}}{H} L c_{x,0} R.
 \end{aligned}$$

Lastly, simple calculations show $g(h) \leq \left(\sqrt{c_\sigma} L + \frac{c_{x,0}}{c_{w,0}} \right) e^{2c_{res}c_{w,0}L} R$.

□

Proof of Lemma C.4. Similar to the proof of Lemma B.4, we can obtain

$$\left| \mathbf{G}_{i,j}^{(H)}(k) - \mathbf{G}_{i,j}^{(H)}(0) \right| \leq \frac{c_{res}^2}{H^2} \left(I_1^{i,j} + I_2^{i,j} + I_3^{i,j} \right).$$

For $I_1^{i,j}$, using Lemma C.3, we have

$$\begin{aligned}
 I_1^{i,j} &= L^2 a_{2,0}^2 \left| \mathbf{x}_i^{(H-1)}(k)^\top \mathbf{x}_j^{(H-1)}(k) - \mathbf{x}_i^{(H-1)}(0)^\top \mathbf{x}_j^{(H-1)}(0) \right| \\
 &\leq L^2 a_{2,0}^2 \left| (\mathbf{x}_i^{(H-1)}(k) - \mathbf{x}_i^{(H-1)}(0))^\top \mathbf{x}_j^{(H-1)}(k) \right| + L^2 a_{2,0}^2 \left| \mathbf{x}_i^{(H-1)}(0)^\top (\mathbf{x}_i^{(H-1)}(k) - \mathbf{x}_i^{(H-1)}(0)) \right| \\
 &\leq c_x L^2 a_{2,0}^2 R \cdot (c_{x,0} + c_x R) + c_{x,0} c_x L^2 a_{2,0}^2 R \\
 &\leq 3c_{x,0} c_x L^2 a_{2,0}^2 R,
 \end{aligned}$$

where $c_x \triangleq \left(\sqrt{c_\sigma} L + \frac{c_{x,0}}{c_{w,0}} \right) e^{2c_{res}c_{w,0}L}$. To bound $I_2^{i,j}$, we have

$$\begin{aligned}
 I_2^{i,j} &= c_{x,0}^2 \frac{1}{m} \left| \sum_{r=1}^m a_r(0)^2 \sigma'(z_{i,r}(k)) \sigma'(z_{j,r}(k)) - a_r(0)^2 \sigma'(z_{i,r}(0)) \sigma'(z_{j,r}(0)) \right| \\
 &\leq c_{x,0}^2 \frac{1}{m} \sum_{r=1}^m a_r(0)^2 \left| (\sigma'(z_{i,r}(k)) - \sigma'(z_{i,r}(0))) \sigma'(z_{j,r}(k)) \right| + a_r(0)^2 \left| (\sigma'(z_{j,r}(k)) - \sigma'(z_{j,r}(0))) \sigma'(z_{i,r}(0)) \right| \\
 &\leq \frac{\beta L c_{x,0}^2}{m} \left(\sum_{r=1}^m a_r(0)^2 |z_{i,r}(k) - z_{i,r}(0)| + a_r(0)^2 |z_{j,r}(k) - z_{j,r}(0)| \right) \\
 &\leq \frac{\beta L a_{4,0}^2 c_{x,0}^2}{\sqrt{m}} \left(\sqrt{\sum_{r=1}^m |z_{i,r}(k) - z_{i,r}(0)|^2} + \sqrt{\sum_{r=1}^m |z_{j,r}(k) - z_{j,r}(0)|^2} \right).
 \end{aligned}$$

Using the same proof for Lemma C.3, it is easy to see

$$\sum_{r=1}^m |z_{i,r}(k) - z_{i,r}(0)|^2 \leq (2c_x c_{w,0} + c_{x,0})^2 L^2 m R^2.$$

Thus

$$I_2^{i,j} \leq 2\beta c_{x,0}^2 (2c_x c_{w,0} + c_{x,0}) L^2 R.$$

The bound of $I_3^{i,j}$ is similar to that in Lemma B.4,

$$I_3^{i,j} \leq 12L^2 c_{x,0}^2 a_{2,0} R.$$

Therefore we can bound the perturbation

$$\begin{aligned} \left\| \mathbf{G}^{(H)}(k) - \mathbf{G}^{(H)}(0) \right\|_F &= \sqrt{\sum_{(i,j)}^{n,n} \left| \mathbf{G}_{i,j}^{(H)}(k) - \mathbf{G}_{i,j}^{(H)}(0) \right|^2} \\ &\leq \frac{c_{res}^2 L^2 n R}{H^2} \left[3c_{x,0} c_x a_{2,0}^2 + 2\beta c_{x,0}^2 (2c_x c_w + c_{x,0}) a_{4,0}^2 + 12c_{x,0}^2 a_{2,0} \right]. \end{aligned}$$

Plugging in the bound on R , we have the desired result. □

Proof of Lemma C.5. We will prove this corollary by induction. The induction hypothesis is

$$\begin{aligned} \left\| \mathbf{W}^{(h)}(s) - \mathbf{W}^{(h)}(0) \right\|_F &\leq \sum_{s'=0}^{s-1} \left(1 - \frac{\eta\lambda_0}{2}\right)^{s'/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m} \leq R' \sqrt{m}, s \in [k+1], \\ \left\| \mathbf{a}(s) - \mathbf{a}(0) \right\|_2 &\leq \sum_{s'=0}^{s-1} \left(1 - \frac{\eta\lambda_0}{2}\right)^{s'/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m} \leq R' \sqrt{m}, s \in [k+1]. \end{aligned}$$

First it is easy to see it holds for $s' = 0$. Now suppose it holds for $s' = 0, \dots, s$, we consider $s' = s + 1$. Similar to Lemma B.5, we have

$$\begin{aligned} &\left\| \mathbf{W}^{(h)}(s+1) - \mathbf{W}^{(h)}(s) \right\|_F \\ &\leq \eta \frac{L c_{res}}{H \sqrt{m}} \left\| \mathbf{a} \right\|_2 \sum_{i=1}^n |y_i - u_i(s)| \left\| \mathbf{x}_i^{(h-1)}(s) \right\|_2 \prod_{k=h+1}^H \left\| \mathbf{I} + \frac{c_{res} \lambda^{3/2}}{H \sqrt{m}} \mathbf{J}_i^{(k)}(s) \mathbf{W}^{(k)}(s) \right\|_2 \\ &\leq 2\eta c_{res} c_{x,0} L a_{2,0} e^{2c_{res} c_w L} \sqrt{n} \left\| \mathbf{y} - \mathbf{u}(s) \right\|_2 / H \\ &= \eta Q'(s) \\ &\leq \left(1 - \frac{\eta\lambda_0}{2}\right)^{s/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m}, \end{aligned}$$

Similarly, we have

$$\begin{aligned} \left\| \mathbf{a}(s+1) - \mathbf{a}(s) \right\|_2 &\leq 2\eta c_{x,0} \sum_{i=1}^n |y_i - u(s)| \\ &\leq \eta Q'(s) \\ &\leq \left(1 - \frac{\eta\lambda_0}{2}\right)^{s/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m}. \end{aligned}$$

Thus

$$\begin{aligned} &\left\| \mathbf{W}^{(h)}(s+1) - \mathbf{W}^{(h)}(0) \right\|_F \\ &\leq \left\| \mathbf{W}^{(h)}(s+1) - \mathbf{W}^{(h)}(s) \right\|_F + \left\| \mathbf{W}^{(h)}(s) - \mathbf{W}^{(h)}(0) \right\|_F \\ &\leq \sum_{s'=0}^s \eta \left(1 - \frac{\eta\lambda_0}{2}\right)^{s'/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m}. \end{aligned}$$

Similarly,

$$\begin{aligned} & \|\mathbf{a}(s+1) - \mathbf{a}(0)\|_2 \\ & \leq \sum_{s'=0}^s \eta \left(1 - \frac{\eta\lambda_0}{2}\right)^{s'/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m}. \end{aligned}$$

□

Proof of Lemma C.6. Similar to Lemma B.6, we first bound the gradient norm.

$$\begin{aligned} & \left\| L^{(h)}(\mathbf{w}(k)) \right\|_F \\ & = \left\| \frac{c_{res}}{H\sqrt{m}} \sum_{i=1}^n (y_i - u_i(k)) \mathbf{x}_i^{(h-1)}(k) \cdot \left[\mathbf{a}(k)^\top \prod_{l=h+1}^H \left(\mathbf{I} + \frac{c_{res}}{H\sqrt{m}} \mathbf{J}_i^{(l)}(k) \mathbf{W}^{(l)}(k) \right) \mathbf{J}_i^{(h)}(k) \right] \right\|_F \\ & \leq \frac{c_{res}L}{H\sqrt{m}} \|\mathbf{a}(k)\|_2 \sum_{i=1}^n |y_i - u_i(k)| \left\| \mathbf{x}_i^{(h-1)}(k) \right\|_2 \prod_{k=h+1}^H \left\| \mathbf{I} + \frac{c_{res}}{H\sqrt{m}} \mathbf{J}_i^{(k)}(k) \mathbf{W}^{(k)}(k) \right\|_2. \end{aligned}$$

We have bounded the RHS in the proof for Lemma C.5, thus

$$\left\| L^{(h)}(\theta(k)) \right\|_F \leq \lambda_0 Q'(k).$$

Let $\theta(k, s) = \theta(k) - sL'(\theta(k))$, we have

$$\begin{aligned} & \left\| u_i^{(h)}(\theta(k)) - u_i^{(h)}(\theta(k, s)) \right\|_F = \\ & \frac{c_{res}}{H\sqrt{m}} \left\| \mathbf{x}_i^{(h-1)}(k) \mathbf{a}(k)^\top \prod_{l=h+1}^H \left(\mathbf{I} + \frac{c_{res}}{H\sqrt{m}} \mathbf{J}_i^{(l)}(k) \mathbf{W}^{(l)}(k) \right) \mathbf{J}_i^{(h)}(k) \right. \\ & \quad \left. - \mathbf{x}_i^{(h-1)}(k, s) \mathbf{a}(k, s)^\top \prod_{l=h+1}^H \left(\mathbf{I} + \frac{c_{res}}{H\sqrt{m}} \mathbf{J}_i^{(l)}(k, s) \mathbf{W}^{(l)}(k, s) \right) \mathbf{J}_i^{(h)}(k, s) \right\|_F. \end{aligned}$$

Through standard calculations, we have

$$\begin{aligned} & \left\| \mathbf{W}^{(l)}(k) - \mathbf{W}^{(l)}(k, s) \right\|_F \leq \eta Q'(k), \\ & \left\| \mathbf{a}(k) - \mathbf{a}(k, s) \right\|_F \leq \eta Q'(k), \\ & \left\| \mathbf{x}_i^{(h-1)}(k) - \mathbf{x}_i^{(h-1)}(k, s) \right\|_F \leq \eta c_x \frac{Q'(k)}{\sqrt{m}}, \\ & \left\| \mathbf{J}^{(l)}(k) - \mathbf{J}^{(l)}(k, s) \right\|_F \leq 2(c_{x,0} + c_{w,0}c_x) \eta \beta Q'(k), \end{aligned}$$

where $c_x \triangleq \left(\sqrt{c_\sigma} L + \frac{c_{x,0}}{c_{w,0}} \right) e^{3c_{res}c_{w,0}L}$. According to Lemma G.1, we have

$$\begin{aligned} & \left\| u_i^{(h)}(\theta(k)) - u_i^{(h)}(\theta(k, s)) \right\|_F \\ & \leq \frac{4}{H} c_{res} c_{x,0} L a_{2,0} e^{2Lc_{w,0}} \eta \frac{Q'(k)}{\sqrt{m}} \left(\frac{c_x}{c_{x,0}} + \frac{2}{L} (c_{x,0} + c_{w,0}c_x) \beta \sqrt{m} + 4c_{w,0} (c_{x,0} + c_{w,0}c_x) \beta + L + 1 \right) \\ & \leq \frac{32}{H} c_{res} c_{x,0} a_{2,0} e^{2Lc_{w,0}} (c_{x,0} + c_{w,0}c_x) \beta \eta Q'(k). \end{aligned}$$

Thus we have

$$|I_2^i| \leq 32c_{res}c_{x,0}a_{2,0}e^{2Lc_{w,0}}(c_{x,0} + c_{w,0}c_x)\beta\eta^2Q'(k)^2 \leq \frac{1}{8}\eta\lambda_0\|\mathbf{y} - \mathbf{u}(k)\|_2,$$

where we used the bound of η and that $\|\mathbf{y} - \mathbf{u}(0)\|_2 = O(\sqrt{n})$. \square

Proof of Lemma C.7.

$$\begin{aligned} \|\mathbf{u}(k+1) - \mathbf{u}(k)\|_2^2 &= \sum_{i=1}^n \left(\mathbf{a}(k+1)^\top \mathbf{x}_i^{(H)}(k+1) - \mathbf{a}(k)^\top \mathbf{x}_i^{(H)}(k) \right)^2 \\ &= \sum_{i=1}^n \left([\mathbf{a}(k+1) - \mathbf{a}(k)]^\top \mathbf{x}_i^{(H)}(k+1) + \mathbf{a}(k)^\top [\mathbf{x}_i^{(H)}(k+1) - \mathbf{x}_i^{(H)}(k)] \right)^2 \\ &\leq 2\|\mathbf{a}(k+1) - \mathbf{a}(k)\|_2^2 \sum_{i=1}^n \|\mathbf{x}_i^{(H)}(k+1)\|_2^2 + 2\|\mathbf{a}(k)\|_2^2 \sum_{i=1}^n \|\mathbf{x}_i^{(H)}(k+1) - \mathbf{x}_i^{(H)}(k)\|_2^2 \\ &\leq 8n\eta^2c_{x,0}^2Q'(k)^2 + 4n(\eta a_{2,0}c_xQ'(k))^2 \\ &\leq \frac{1}{8}\eta\lambda_0\|\mathbf{y} - \mathbf{u}(k)\|_2^2. \end{aligned}$$

\square

D. Proofs for Section 7

For CNN, denote $\mathbf{x}_{i,l} = \phi(\mathbf{x}_{i,l})_{:,l}$, $\mathbf{G}^{(H)}$ has the following form:

$$\mathbf{G}_{ij}^{(H)} = \frac{c_{res}^2}{H^2m} \sum_{r=1}^m \left[\sum_{l=1}^p a_{l,r} \mathbf{x}_{i,l}^{(H-1)} \sigma' \left(\left(\mathbf{w}_r^{(H)} \right)^\top \mathbf{x}_{i,l}^{(H-1)} \right) \right]^\top \left[\sum_{k=1}^p a_{k,r} \mathbf{x}_{j,k}^{(H-1)} \sigma' \left(\left(\mathbf{w}_r^{(H)} \right)^\top \mathbf{x}_{j,k}^{(H-1)} \right) \right]. \quad (15)$$

We define a constant $c_{\sigma,c_0} = \left(\min_{c_0 \leq \alpha \leq 1} \mathbb{E}_{X \sim N(0,1)} \sigma(\alpha X)^2 \right)^{-1} > 0$, where $0 < c_0 \leq 1$. In particular, it is easy to see for smooth ReLU, $c_{\sigma, \frac{1}{\sqrt{p}}} = \text{poly}(p)$.

Similar to Lemma B.1, we can show with high probability the feature of each layer is approximately normalized.

Lemma D.1 (Lemma on Initialization Norms). *If $\sigma(\cdot)$ is L -Lipschitz and $m = \Omega\left(\frac{p^2n}{c_{\sigma, \frac{1}{\sqrt{p}}}^2\delta}\right)$, assuming $\|\mathbf{W}^{(h)}(0)\|_2 \leq c_{w,0}\sqrt{m}$ for $h \in [H]$, we have with probability at least $1 - \delta$ over random initialization, for every $h \in [H]$ and $i \in [n]$,*

$$\frac{1}{c_{x,0}} \leq \|\mathbf{x}_i^{(h)}(0)\|_F \leq c_{x,0}$$

for some constant $c_{x,0} = \text{poly}(p) > 1$.

The following lemma lower bounds $\mathbf{G}^{(H)}(0)$'s least eigenvalue. This lemma is a direct consequence of results in Section E.

Lemma D.2 (Least Eigenvalue at the Initialization). *If $m = \Omega\left(\frac{n^2p^2 \log(Hn/\delta)}{\lambda_0^2}\right)$, we have*

$$\lambda_{\min}(\mathbf{G}^{(H)}(0)) \geq \frac{3}{4}\lambda_0.$$

Next, we prove the following lemma which characterizes how the perturbation from weight matrices propagates to the input of each layer.

Lemma D.3. *Suppose $\sigma(\cdot)$ is L -Lipschitz and for $h \in [H]$, $\|\mathbf{W}^{(h)}(0)\|_2 \leq c_{w,0}\sqrt{m}$, $\|\mathbf{x}^{(h)}(0)\|_F \leq c_{x,0}$ and $\|\mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0)\|_F \leq \sqrt{m}R$ for some constant $c_{w,0}, c_{x,0} > 1$ and $R \leq c_{w,0}$. Then we have*

$$\|\mathbf{x}^{(h)}(k) - \mathbf{x}^{(h)}(0)\|_F \leq \left(\sqrt{c_\sigma}L\sqrt{q} + \frac{c_{x,0}}{c_{w,0}} \right) e^{2c_{w,0}L\sqrt{q}c_{res}}R.$$

Next, we show with high probability over random initialization, perturbation in weight matrices leads to small perturbation in the Gram matrix.

Lemma D.4. *Suppose $\sigma(\cdot)$ is differentiable, L -Lipschitz and β -smooth. Using the same notations in Lemma B.4, if $\|\mathbf{a}_{:,i}\|_2 \leq a_{2,0}\sqrt{m}$ and $\|\mathbf{a}_{:,i}\|_4 \leq a_{4,0}m^{1/4}$ for any $i \in [p]$, $\|\mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0)\|_F, \|\mathbf{a}(k) - \mathbf{a}(0)\|_F \leq \sqrt{m}R$ where $R \leq c\lambda_0 H^2 (n)^{-1} \text{poly}(p)^{-1}$ for some small constant c , we have*

$$\left\| \mathbf{G}^{(H)}(k) - \mathbf{G}^{(H)}(0) \right\|_2 \leq \frac{\lambda_0}{2}.$$

Lemma D.5. *If Condition A.1 holds for $k' = 1, \dots, k$, we have for any $s \in [k+1]$*

$$\begin{aligned} \left\| \mathbf{W}^{(h)}(s) - \mathbf{W}^{(h)}(0) \right\|_F, \|\mathbf{a}(s) - \mathbf{a}(0)\|_F &\leq R' \sqrt{m}, \\ \left\| \mathbf{W}^{(h)}(s) - \mathbf{W}^{(h)}(s-1) \right\|_F, \|\mathbf{a}(s) - \mathbf{a}(s-1)\|_F &\leq \eta Q'(s-1), \end{aligned}$$

where $R' = \frac{16c_{res}c_{x,0}L\sqrt{pq}e^{2c_{res}c_w,0L}a_{2,0}\sqrt{q}\sqrt{n}\|\mathbf{y} - \mathbf{u}(0)\|_2}{H\lambda_0\sqrt{m}} < c$ for some small constant c and

$$Q'(s) = 4c_{res}c_{x,0}La_{2,0}\sqrt{pq}e^{2c_{res}c_w,0L}\sqrt{q}\sqrt{n}\|\mathbf{y} - \mathbf{u}(s)\|_2 / H.$$

The follow lemma bounds the norm of \mathbf{I}_2 .

Lemma D.6. *If Condition A.1 holds for $k' = 1, \dots, k$ and $\eta \leq c\lambda_0 H^2 n^{-2} \text{poly}(1/p)$ for some small constant c , we have*

$$\|\mathbf{I}_2(k)\|_2 \leq \frac{1}{8}\eta\lambda_0 \|\mathbf{y} - \mathbf{u}(k)\|_2.$$

Next we also bound the quadratic term.

Lemma D.7. *If Condition A.1 holds for $k' = 1, \dots, k$ and $\eta \leq c\lambda_0 H^2 n^{-2} \text{poly}(1/p)$ for some small constant c , we have $\|\mathbf{u}(k+1) - \mathbf{u}(k)\|_2^2 \leq \frac{1}{8}\eta\lambda_0 \|\mathbf{y} - \mathbf{u}(k)\|_2^2$.*

Now using the same argument as in the proof for multilayer fully connected neural network, we finish our proof for CNN.

D.1. Proofs of Lemmas

Proof of Lemma D.1. We will bound $\left\| \mathbf{x}_i^{(h)}(0) \right\|_F$ layer by layer. For the first layer, we can calculate

$$\begin{aligned} \mathbb{E} \left[\left\| \mathbf{x}_i^{(1)}(0) \right\|_F^2 \right] &= c_\sigma \sum_{l=1}^{p_1} \mathbb{E} \left[\sigma \left(\mathbf{w}_r^{(1)}(0)^\top \mathbf{x}_{i,l} \right)^2 \right] \\ &\geq \frac{c_\sigma}{c_{\sigma, \frac{1}{\sqrt{p}}}}, \end{aligned}$$

where the inequality we use the definition of $c_{\sigma, \frac{1}{\sqrt{p}}}$ and the fact that there must exist $l' \in [p]$ such that $\|\mathbf{x}_{i,l'}\|_2^2 \geq \frac{1}{p_1} \geq \frac{1}{p}$. For the variance,

$$\begin{aligned} \text{Var} \left[\left\| \mathbf{x}_i^{(1)}(0) \right\|_F^2 \right] &= \frac{c_\sigma^2}{m} \text{Var} \left[\sum_{l=1}^{p_1} \sigma \left(\mathbf{w}_r^{(1)}(0)^\top \mathbf{x}_{i,l} \right)^2 \right] \\ &\leq \frac{c_\sigma^2}{m} \mathbb{E} \left[\left(\sum_{l=1}^{p_1} \left(|\sigma(0)| + L \left| \mathbf{w}_r^{(1)}(0)^\top \mathbf{x}_{i,l} \right| \right)^2 \right)^2 \right] \\ &\leq \frac{p^2 C_2}{m}, \end{aligned}$$

where $C_2 \triangleq \sigma(0)^4 + 4|\sigma(0)|^3 L\sqrt{2/\pi} + 6\sigma(0)^2 L^2 + 8|\sigma(0)| L^3\sqrt{2/\pi} + 32L^4$. We have with probability at least $1 - \frac{\delta}{n}$,

$$\left\| \mathbf{x}_i^{(1)}(0) \right\|_F^2 \geq \frac{c_\sigma}{2c_{\sigma, \frac{1}{\sqrt{p}}}}.$$

It is easy to get its upper bound

$$\left\| \mathbf{x}_i^{(1)}(0) \right\|_F^2 = \frac{c_\sigma}{m} \left\| \sigma \left(\mathbf{W}^{(1)} \phi(\mathbf{x}_i) \right) \right\|_F^2 \leq qL^2 c_\sigma c_{w,0}^2.$$

By definition we have for $2 \leq h \leq H$

$$\begin{aligned} \left\| \mathbf{x}_i^{(h-1)}(0) \right\|_F - \left\| \frac{c_{res}}{H\sqrt{m}} \sigma \left(\mathbf{W}^{(h)}(0) \phi \left(\mathbf{x}_i^{(h-1)}(0) \right) \right) \right\|_F &\leq \left\| \mathbf{x}_i^{(h)}(0) \right\|_F \\ &\leq \left\| \mathbf{x}_i^{(h-1)}(0) \right\|_F + \left\| \frac{c_{res}}{H\sqrt{m}} \sigma \left(\mathbf{W}^{(h)}(0) \phi \left(\mathbf{x}_i^{(h-1)}(0) \right) \right) \right\|_F, \end{aligned}$$

where

$$\left\| \frac{c_{res}}{H\sqrt{m}} \sigma \left(\mathbf{W}^{(h)}(0) \phi \left(\mathbf{x}_i^{(h-1)}(0) \right) \right) \right\|_F \leq \frac{\sqrt{q} c_{res} c_{w,0} L}{H} \left\| \mathbf{x}_i^{(h-1)}(0) \right\|_F.$$

Thus

$$\left\| \mathbf{x}_i^{(h-1)}(0) \right\|_F \left(1 - \frac{\sqrt{q} c_{res} c_{w,0} L}{H} \right) \leq \left\| \mathbf{x}^{(h)}(0) \right\|_F \leq \left\| \mathbf{x}_i^{(h-1)}(0) \right\|_F \left(1 + \frac{\sqrt{q} c_{res} c_{w,0} L}{H} \right),$$

which implies

$$\sqrt{\frac{c_\sigma}{2c_{\sigma, \frac{1}{\sqrt{p}}}}} e^{-\sqrt{q} c_{res} c_{w,0} L} \leq \left\| \mathbf{x}^{(h)}(0) \right\|_F \leq \sqrt{qL^2 c_\sigma c_{w,0}^2} e^{\sqrt{q} c_{res} c_{w,0} L}.$$

Choosing $c_{x,0} = \max\left\{ \sqrt{qL^2 c_\sigma c_{w,0}^2}, \sqrt{\frac{2c_{\sigma, \frac{1}{\sqrt{p}}}}{c_\sigma}} \right\} e^{\sqrt{q} c_{res} c_{w,0} L}$ and using union bounds over $[n]$, we prove the lemma. \square

Proof of Lemma D.3. We prove this lemma by induction. Our induction hypothesis is

$$\left\| \mathbf{x}^{(h)}(k) - \mathbf{x}^{(h)}(0) \right\|_F \leq g(h),$$

where

$$g(h) = g(h-1) \left[1 + \frac{2c_{res} c_{w,0} L \sqrt{q}}{H} \right] + \frac{c_{res} L \sqrt{q}}{H} R c_{x,0}.$$

For $h = 1$, we have

$$\begin{aligned} \left\| \mathbf{x}^{(1)}(k) - \mathbf{x}^{(1)}(0) \right\|_F &\leq \sqrt{\frac{c_\sigma}{m}} \left\| \sigma \left(\mathbf{W}^{(1)}(k) \phi_1(\mathbf{x}) \right) - \sigma \left(\mathbf{W}^{(1)}(0) \phi_1(\mathbf{x}) \right) \right\|_F \\ &\leq \sqrt{\frac{c_\sigma}{m}} L \sqrt{q} \left\| \mathbf{W}^{(1)}(k) - \mathbf{W}^{(1)}(0) \right\|_F \leq \sqrt{c_\sigma} L \sqrt{q} R, \end{aligned}$$

which implies $g(1) = \sqrt{c_\sigma} L \sqrt{q} R$, for $2 \leq h \leq H$, we have

$$\begin{aligned} &\left\| \mathbf{x}^{(h)}(k) - \mathbf{x}^{(h)}(0) \right\|_F \\ &\leq \frac{c_{res}}{H\sqrt{m}} \left\| \sigma \left(\mathbf{W}^{(h)}(k) \phi_h \left(\mathbf{x}^{(h-1)}(k) \right) \right) - \sigma \left(\mathbf{W}^{(h)}(0) \phi_h \left(\mathbf{x}^{(h-1)}(0) \right) \right) \right\|_F + \left\| \mathbf{x}^{(h-1)}(k) - \mathbf{x}^{(h-1)}(0) \right\|_F \\ &\leq \frac{c_{res}}{H\sqrt{m}} \left\| \sigma \left(\mathbf{W}^{(h)}(k) \phi_h \left(\mathbf{x}^{(h-1)}(k) \right) \right) - \sigma \left(\mathbf{W}^{(h)}(k) \phi_h \left(\mathbf{x}^{(h-1)}(0) \right) \right) \right\|_F \end{aligned}$$

$$\begin{aligned}
 & + \frac{c_{res}}{H\sqrt{m}} \left\| \sigma \left(\mathbf{W}^{(h)}(k) \phi_h \left(\mathbf{x}^{(h-1)}(0) \right) \right) - \sigma \left(\mathbf{W}^{(h)}(0) \phi_h \left(\mathbf{x}^{(h-1)}(0) \right) \right) \right\|_F + \left\| \mathbf{x}^{(h-1)}(k) - \mathbf{x}^{(h-1)}(0) \right\|_F \\
 & \leq \frac{L\sqrt{q}c_{res}}{H\sqrt{m}} \left(\left\| \mathbf{W}^{(h)}(0) \right\|_2 + \left\| \mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0) \right\|_F \right) \cdot \left\| \mathbf{x}^{(h-1)}(k) - \mathbf{x}^{(h-1)}(0) \right\|_F \\
 & \quad + \frac{L\sqrt{q}c_{res}}{H\sqrt{m}} \left\| \mathbf{W}^{(h)}(k) - \mathbf{W}^{(h)}(0) \right\|_F \left\| \mathbf{x}^{(h-1)}(0) \right\|_F + \left\| \mathbf{x}^{(h-1)}(k) - \mathbf{x}^{(h-1)}(0) \right\|_F \\
 & \leq \left[1 + \frac{L\sqrt{q}c_{res}}{H\sqrt{m}} (c_{w,0}\sqrt{m} + R\sqrt{m}) \right] g(h-1) + \frac{L\sqrt{q}c_{res}}{H\sqrt{m}} \sqrt{m} R c_{x,0} \\
 & \leq \left(1 + \frac{2c_{w,0}L\sqrt{q}c_{res}}{H} \right) g(h-1) + \frac{1}{H} L\sqrt{q}c_{res} c_{x,0} R.
 \end{aligned}$$

Lastly, simple calculations show $g(h) \leq \left(\sqrt{c_\sigma} L\sqrt{q} + \frac{c_{x,0}}{c_{w,0}} \right) e^{2c_{w,0}L\sqrt{q}c_{res}} R$.

□

Proof of Lemma D.4. Similar to Lemma C.4, define $z_{i,l,r} = \left(\mathbf{w}_r^{(H)} \right)^\top \mathbf{x}_{i,l}^{(H-1)}$, we have

$$\begin{aligned}
 & \left| \mathbf{G}_{i,j}^{(H)}(k) - \mathbf{G}_{i,j}^{(H)}(0) \right| \\
 & = \frac{c_{res}^2}{H^2} \left| \sum_{l=1}^p \sum_{k=1}^p \mathbf{x}_{i,l}^{(H-1)}(k)^\top \mathbf{x}_{j,k}^{(H-1)}(k) \frac{1}{m} \sum_{r=1}^m a_{r,l}(k) a_{r,k}(k) \sigma'(z_{i,l,r}(k)) \sigma'(z_{j,k,r}(k)) \right. \\
 & \quad \left. - \sum_{l=1}^p \sum_{k=1}^p \mathbf{x}_{i,l}^{(H-1)}(0)^\top \mathbf{x}_{j,k}^{(H-1)}(0) \frac{1}{m} \sum_{r=1}^m a_{r,l}(0) a_{r,k}(0) \sigma'(z_{i,l,r}(0)) \sigma'(z_{j,k,r}(0)) \right| \\
 & \leq \frac{c_{res}^2 L^2 a_{2,0}^2}{H^2} \left| \sum_{l=1}^p \sum_{k=1}^p \mathbf{x}_{i,l}^{(H-1)}(k)^\top \mathbf{x}_{j,k}^{(H-1)}(k) - \mathbf{x}_{i,l}^{(H-1)}(0)^\top \mathbf{x}_{j,k}^{(H-1)}(0) \right| \\
 & \quad + \frac{c_{res}^2}{H^2} \sum_{l=1}^p \sum_{k=1}^p \left| \mathbf{x}_{i,l}^{(H-1)}(0)^\top \mathbf{x}_{j,k}^{(H-1)}(0) \right| \frac{1}{m} \sum_{r=1}^m |a_{r,l}(0) a_{r,k}(0)| \left| \sigma'(z_{i,l,r}(k)) \sigma'(z_{j,k,r}(k)) - \sigma'(z_{i,l,r}(0)) \sigma'(z_{j,k,r}(0)) \right| \\
 & \quad + \frac{c_{res}^2}{H^2} L^2 \sum_{l=1}^p \sum_{k=1}^p \left| \mathbf{x}_{i,l}^{(H-1)}(k)^\top \mathbf{x}_{j,k}^{(H-1)}(k) \right| \frac{1}{m} \sum_{r=1}^m |a_{r,l}(k) a_{r,k}(k) - a_{r,l}(0) a_{r,k}(0)| \\
 & \triangleq \frac{c_{res}^2}{H^2} \left(I_1^{i,j} + I_2^{i,j} + I_3^{i,j} \right).
 \end{aligned}$$

For $I_1^{i,j}$, using Lemma D.3, we have

$$\begin{aligned}
 I_1^{i,j} & = L^2 a_{2,0}^2 \left| \sum_{l=1}^p \sum_{k=1}^p \mathbf{x}_{i,l}^{(H-1)}(k)^\top \mathbf{x}_{j,k}^{(H-1)}(k) - \mathbf{x}_{i,l}^{(H-1)}(0)^\top \mathbf{x}_{j,k}^{(H-1)}(0) \right| \\
 & \leq L^2 a_{2,0}^2 \sum_{l=1}^p \sum_{k=1}^p \left| \left(\mathbf{x}_{i,l}^{(H-1)}(k) - \mathbf{x}_{i,l}^{(H-1)}(0) \right)^\top \mathbf{x}_{j,k}^{(H-1)}(k) \right| + L^2 a_{2,0}^2 \sum_{l=1}^p \sum_{k=1}^p \left| \mathbf{x}_{i,l}^{(H-1)}(0)^\top \left(\mathbf{x}_{j,k}^{(H-1)}(k) - \mathbf{x}_{j,k}^{(H-1)}(0) \right) \right| \\
 & \leq L^2 a_{2,0}^2 \sqrt{\sum_{l=1}^p \sum_{k=1}^p \left\| \mathbf{x}_{i,l}^{(H-1)}(k) - \mathbf{x}_{i,l}^{(H-1)}(0) \right\|_2^2} \sqrt{\sum_{l=1}^p \sum_{k=1}^p \left\| \mathbf{x}_{j,k}^{(H-1)}(k) \right\|_2^2} \\
 & \quad + L^2 a_{2,0}^2 \sqrt{\sum_{l=1}^p \sum_{k=1}^p \left\| \mathbf{x}_{i,l}^{(H-1)}(0) \right\|_2^2} \sqrt{\sum_{l=1}^p \sum_{k=1}^p \left\| \mathbf{x}_{j,k}^{(H-1)}(k) - \mathbf{x}_{j,k}^{(H-1)}(0) \right\|_2^2} \\
 & \leq L^2 a_{2,0}^2 p \left\| \mathbf{x}_i^{(H-1)}(k) - \mathbf{x}_i^{(H-1)}(0) \right\|_F \left\| \mathbf{x}_j^{(H-1)}(k) \right\|_F + L^2 a_{2,0}^2 p \left\| \mathbf{x}_i^{(H-1)}(0) \right\|_F \left\| \mathbf{x}_j^{(H-1)}(k) - \mathbf{x}_j^{(H-1)}(0) \right\|_F \\
 & \leq 3c_{x,0} c_x L^2 a_{2,0}^2 p R,
 \end{aligned}$$

where $c_x \triangleq \left(\sqrt{c_\sigma} L \sqrt{q} + \frac{c_{x,0}}{c_{w,0}} \right) e^{2c_{res} c_{w,0} L \sqrt{q}}$. To bound $I_2^{i,j}$, we have

$$\begin{aligned}
 I_2^{i,j} &= \sum_{l=1}^p \sum_{k=1}^p \left| \mathbf{x}_{i,l}^{(H-1)}(0)^\top \mathbf{x}_{j,k}^{(H-1)}(0) \right| \frac{1}{m} \sum_{r=1}^m |a_{r,l}(0) a_{r,k}(0)| |\sigma'(z_{i,l,r}(k)) \sigma'(z_{j,k,r}(k)) - \sigma'(z_{i,l,r}(0)) \sigma'(z_{j,k,r}(0))| \\
 &\leq \sum_{l=1}^p \sum_{k=1}^p \left| \mathbf{x}_{i,l}^{(H-1)}(0)^\top \mathbf{x}_{j,k}^{(H-1)}(0) \right| \frac{\beta L}{m} \left(\sum_{r=1}^m |a_{r,l}(0) a_{r,k}(0)| (|z_{i,l,r}(k) - z_{i,l,r}(0)| + |z_{j,k,r}(k) - z_{j,k,r}(0)|) \right) \\
 &\leq \frac{\beta L}{m} \sqrt{\sum_{l=1}^p \sum_{k=1}^p \left\| \mathbf{x}_{i,l}^{(H-1)}(0) \right\|_2^2 \left\| \mathbf{x}_{j,k}^{(H-1)}(0) \right\|_2^2} \\
 &\quad \left(\sqrt{\sum_{l=1}^p \sum_{k=1}^p \left(\sum_{r=1}^m |a_{r,l}(0) a_{r,k}(0)| |z_{i,l,r}(k) - z_{i,l,r}(0)| \right)^2} + \sqrt{\sum_{l=1}^p \sum_{k=1}^p \left(\sum_{r=1}^m |a_{r,l}(0) a_{r,k}(0)| |z_{j,k,r}(k) - z_{j,k,r}(0)| \right)^2} \right) \\
 &\leq \frac{\beta L c_{x,0}^2 a_{4,0}^2}{m} \left(\sqrt{m \sum_{l=1}^p \sum_{k=1}^p \sum_{r=1}^m |z_{i,l,r}(k) - z_{i,l,r}(0)|^2} + \sqrt{m \sum_{l=1}^p \sum_{k=1}^p \sum_{r=1}^m |z_{j,k,r}(k) - z_{j,k,r}(0)|^2} \right) \\
 &\leq \frac{\beta L a_{4,0}^2 \sqrt{p} c_{x,0}^2}{\sqrt{m}} (\|\mathbf{z}_i\|_F + \|\mathbf{z}_j\|_F).
 \end{aligned}$$

Using the same proof for Lemma D.3, it is easy to see

$$\|\mathbf{z}_i\|_F \leq (2c_x c_{w,0} \sqrt{q} + c_{x,0}) R \sqrt{m}.$$

Thus

$$I_2^{i,j} \leq 2\beta L a_{4,0}^2 \sqrt{p} c_{x,0}^2 (2c_x c_{w,0} \sqrt{q} + c_{x,0}) R.$$

Similarly for $I_3^{i,j}$, we have

$$\begin{aligned}
 I_3^{i,j} &= \frac{c_{res}^2}{H^2} L^2 \sum_{l=1}^p \sum_{k=1}^p \left| \mathbf{x}_{i,l}^{(H-1)}(k)^\top \mathbf{x}_{j,k}^{(H-1)}(k) \right| \frac{1}{m} \sum_{r=1}^m |a_{r,l}(k) a_{r,k}(k) - a_{r,l}(0) a_{r,k}(0)| \\
 &\leq \frac{c_{res}^2}{H^2} L^2 \sum_{l=1}^p \sum_{k=1}^p \left| \mathbf{x}_{i,l}^{(H-1)}(k)^\top \mathbf{x}_{j,k}^{(H-1)}(k) \right| \frac{1}{m} \sum_{r=1}^m (|a_{r,l}(k) - a_{r,l}(0)| |a_{r,k}(k)| + |a_{r,k}(k) - a_{r,k}(0)| |a_{r,l}(0)|) \\
 &\leq \frac{c_{res}^2}{H^2} L^2 \sum_{l=1}^p \sum_{k=1}^p \left| \mathbf{x}_{i,l}^{(H-1)}(k)^\top \mathbf{x}_{j,k}^{(H-1)}(k) \right| \frac{1}{m} (\|\mathbf{a}_{:,l}(k) - \mathbf{a}_{:,l}(0)\|_2 \|\mathbf{a}_{:,k}(k)\|_2 + \|\mathbf{a}_{:,k}(k) - \mathbf{a}_{:,k}(0)\|_2 \|\mathbf{a}_{:,l}(0)\|_2) \\
 &\leq \frac{c_{res}^2}{H^2 m} L^2 \sqrt{\sum_{l=1}^p \sum_{k=1}^p \left\| \mathbf{x}_{i,l}^{(H-1)}(k) \right\|_2^2 \left\| \mathbf{x}_{j,k}^{(H-1)}(k) \right\|_2^2} \\
 &\quad \left(\sqrt{\sum_{l=1}^p \sum_{k=1}^p \|\mathbf{a}_{:,l}(k) - \mathbf{a}_{:,l}(0)\|_2^2 \|\mathbf{a}_{:,k}(k)\|_2^2} + \sqrt{\sum_{l=1}^p \sum_{k=1}^p \|\mathbf{a}_{:,k}(k) - \mathbf{a}_{:,k}(0)\|_2^2 \|\mathbf{a}_{:,l}(0)\|_2^2} \right) \\
 &\leq \frac{c_{res}^2}{H^2 m} L^2 \left\| \mathbf{x}_i^{(H-1)}(k) \right\|_F \left\| \mathbf{x}_j^{(H-1)}(k) \right\|_F (\|\mathbf{a}(k) - \mathbf{a}(0)\|_F \|\mathbf{a}(k)\|_F + (\|\mathbf{a}(k) - \mathbf{a}(0)\|_F \|\mathbf{a}(0)\|_F)) \\
 &\leq \frac{12 a_{2,0} c_{res}^2 c_{x,0}^2 L^2 \sqrt{p} R}{H^2}.
 \end{aligned}$$

Therefore we can bound the perturbation

$$\left\| \mathbf{G}^{(H)}(k) - \mathbf{G}^{(H)}(0) \right\|_2 \leq \left\| \mathbf{G}^{(H)}(k) - \mathbf{G}^{(H)}(0) \right\|_F$$

$$\begin{aligned}
 &= \sqrt{\sum_{(i,j)}^{n,n} \left| \mathbf{G}_{i,j}^{(H)}(k) - \mathbf{G}_{i,j}^{(H)}(0) \right|^2} \\
 &\leq \frac{c_{res}^2}{H^2} \left[3c_{x,0}c_xLa_{2,0}^2p + 2\beta c_{x,0}^2a_{4,0}^2\sqrt{p}(2c_xc_{w,0}\sqrt{q} + c_{x,0}) + 12c_{x,0}^2La_{2,0}\sqrt{p} \right] LnR.
 \end{aligned}$$

Plugging in the bound on R , we have the desired result. \square

Proof of Lemma D.5. We will prove this corollary by induction. The induction hypothesis is

$$\begin{aligned}
 \left\| \mathbf{W}^{(h)}(s) - \mathbf{W}^{(h)}(0) \right\|_F &\leq \sum_{s'=0}^{s-1} \left(1 - \frac{\eta\lambda_0}{2}\right)^{s'/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m} \leq R' \sqrt{m}, s \in [k+1], \\
 \|\mathbf{a}(s) - \mathbf{a}(0)\|_F &\leq \sum_{s'=0}^{s-1} \left(1 - \frac{\eta\lambda_0}{2}\right)^{s'/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m} \leq R' \sqrt{m}, s \in [k+1].
 \end{aligned}$$

First it is easy to see it holds for $s' = 0$. Now suppose it holds for $s' = 0, \dots, s$, we consider $s' = s + 1$. Similar to Lemma B.5, we have

$$\begin{aligned}
 &\left\| \mathbf{W}^{(h)}(s+1) - \mathbf{W}^{(h)}(s) \right\|_F \\
 &\leq \eta \frac{c_{res}L}{H\sqrt{m}} \|\mathbf{a}\|_F \sum_{i=1}^n |y_i - u(s)| \left\| \phi_h \left(\mathbf{x}^{(h-1)}(s) \right) \right\|_F \prod_{k=h+1}^H \left\| \mathbf{I} + \frac{c_{res}}{H\sqrt{m}} \mathbf{W}^{(k)}(s) \phi_k \right\|_{op} \\
 &\leq 2\eta c_{res} c_{x,0} La_{2,0} \sqrt{pq} e^{2c_{res}c_{w,0}L\sqrt{q}} \sqrt{n} \|\mathbf{y} - \mathbf{u}(s)\|_2 / H \\
 &= \eta Q'(s) \\
 &\leq \left(1 - \frac{\eta\lambda_0}{2}\right)^{s/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m},
 \end{aligned}$$

where $\|\cdot\|_{op}$ denotes the operator norm. Similarly, we have

$$\begin{aligned}
 \|\mathbf{a}(s+1) - \mathbf{a}(s)\|_2 &\leq 2\eta c_{x,0} \sum_{i=1}^n |y_i - u(s)| \\
 &\leq \eta Q'(s) \\
 &\leq \left(1 - \frac{\eta\lambda_0}{2}\right)^{s/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m}.
 \end{aligned}$$

Thus

$$\begin{aligned}
 &\left\| \mathbf{W}^{(h)}(s+1) - \mathbf{W}^{(h)}(0) \right\|_F \\
 &\leq \left\| \mathbf{W}^{(h)}(s+1) - \mathbf{W}^{(h)}(s) \right\|_F + \left\| \mathbf{W}^{(h)}(s) - \mathbf{W}^{(h)}(0) \right\|_F \\
 &\leq \sum_{s'=0}^s \eta \left(1 - \frac{\eta\lambda_0}{2}\right)^{s'/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m}.
 \end{aligned}$$

Similarly,

$$\begin{aligned}
 &\|\mathbf{a}(s+1) - \mathbf{a}(0)\|_2 \\
 &\leq \sum_{s'=0}^s \eta \left(1 - \frac{\eta\lambda_0}{2}\right)^{s'/2} \frac{1}{4} \eta\lambda_0 R' \sqrt{m}.
 \end{aligned}$$

\square

Proof of Lemma D.6.

$$|I_2^i| \leq \eta \max_{0 \leq s \leq \eta} \sum_{h=1}^H \left\| L'^{(h)}(\theta(k)) \right\|_F \left\| u_i'^{(h)}(\theta(k)) - u_i'^{(h)}(\theta(k) - sL'^{(h)}(\theta(k))) \right\|_F.$$

For the gradient norm, we have

$$\begin{aligned} & \left\| L'^{(h)}(\theta(k)) \right\|_F \\ & \leq \frac{Lc_{res}}{H\sqrt{m}} \|\mathbf{a}(k)\|_F \sum_{i=1}^n |y_i - u_i(k)| \left\| \phi_h(\mathbf{x}_i^{(h-1)}(k)) \right\|_F \prod_{k=h+1}^H \left\| \mathbf{I} + \frac{c_{res}}{H\sqrt{m}} \mathbf{J}_i^{(k)}(k) \mathbf{W}^{(k)}(k) \phi_k \right\|_{op}, \end{aligned}$$

which we have bounded in Lemma D.5, thus

$$\left\| L'^{(h)}(\theta(k)) \right\|_F \leq Q'(k).$$

Let $\theta(k, s) = \theta(k) - sL'(\theta(k))$, similar to the proof of Lemma B.6, we have

$$\begin{aligned} & \left\| u_i'^{(h)}(\theta(k)) - u_i'^{(h)}(\theta(k, s)) \right\|_F \\ & \leq \frac{2}{H} c_{res} c_{x,0} L a_{2,0} \sqrt{q} e^{2c_{res} L c_{w,0} \sqrt{q}} \eta \frac{Q'(k)}{\sqrt{m}} \left(\frac{c_x}{c_{x,0}} + \frac{2}{L} (c_{x,0} + c_{w,0} c_x) \beta \sqrt{m} + 4\sqrt{q} c_{w,0} (c_{x,0} + c_{w,0} c_x) \beta \sqrt{m} + (L+1)\sqrt{q} \right) \\ & \leq \frac{24}{H} c_{res} c_{x,0} L a_{2,0} \sqrt{q} c_{w,0} e^{2c_{res} L c_{w,0} \sqrt{q}} (c_{x,0} + c_{w,0} c_x) \beta \eta Q'(k). \end{aligned}$$

Thus

$$|I_2^i| \leq 24 c_{res} c_{x,0} L a_{2,0} \sqrt{q} c_{w,0} e^{2c_{res} L c_{w,0} \sqrt{q}} (c_{x,0} + c_{w,0} c_x) \beta \eta^2 \lambda_0 \sqrt{m} Q'(k) R' \leq \frac{1}{8} \eta \lambda_0 \|\mathbf{y} - \mathbf{u}(k)\|_2.$$

where we used the bound of η and that $\|\mathbf{y} - \mathbf{u}(0)\|_2 = O(\sqrt{n})$. \square

Proof of Lemma D.7.

$$\begin{aligned} \|\mathbf{u}(k+1) - \mathbf{u}(k)\|_2^2 &= \sum_{i=1}^n \left(\langle \mathbf{a}(k+1), \mathbf{x}_i^{(H)}(k+1) \rangle - \langle \mathbf{a}(k), \mathbf{x}_i^{(H)}(k+1) \rangle \right)^2 \\ &\leq \sum_{i=1}^n \left(\langle \mathbf{a}(k+1) - \mathbf{a}(k), \mathbf{x}_i^{(H)}(k+1) \rangle + \langle \mathbf{a}(k), \mathbf{x}_i^{(H)}(k+1) - \mathbf{x}_i^{(H)}(k) \rangle \right)^2 \\ &\leq 2 \|\mathbf{a}(k+1) - \mathbf{a}(k)\|_F^2 \sum_{i=1}^n \left\| \mathbf{x}_i^{(H)}(k+1) \right\|_F^2 + 2 \|\mathbf{a}(k)\|_F^2 \sum_{i=1}^n \left\| \mathbf{x}_i^{(H)}(k+1) - \mathbf{x}_i^{(H)}(k) \right\|_F^2 \\ &\leq 8n\eta^2 c_{x,0}^2 Q'(k)^2 + 4np(\eta a_{2,0} c_x Q'(k))^2 \\ &\leq \frac{1}{8} \eta \lambda_0 \|\mathbf{y} - \mathbf{u}(k)\|_2^2. \end{aligned}$$

\square

E. Analysis of Random Initialization

E.1. A General Framework for Analyzing Random Initialization in First $(H-1)$ Layers

In this section we provide a self-contained framework to analyze the Gram matrix at the initialization phase. There are two main objectives. First, we provide the expression of the Gram matrix as $m \rightarrow \infty$, i.e., the population Gram matrix. Second, we quantitatively study how much over-parameterization is needed to ensure the Gram matrix generated by the random initialization. The bound will depend on number of samples n and properties of the activation function. This analysis

framework is fully general that it can explain fully connected neural network, ResNet, convolutional neural considered in this paper and other neural network architectures that satisfy the general setup defined below.

We begin with some notations. Suppose that we have a sequence of real vector spaces

$$\mathbb{R}^{p^{(0)}} \rightarrow \mathbb{R}^{p^{(1)}} \rightarrow \dots \rightarrow \mathbb{R}^{p^{(H)}}.$$

Remark E.1. For fully-connected neural network and ResNet, $p^{(0)} = p^{(1)} = \dots = p^{(H)} = 1$. For convolutional neural network, $p^{(h)}$ is the number of patches of the h -th layer.

For each pair $(\mathbb{R}^{p^{(h-1)}}, \mathbb{R}^{p^{(h)}})$, let $\mathcal{W} \subset \mathcal{L}(\mathbb{R}^{p^{(h-1)}}, \mathbb{R}^{p^{(h)}}) = \mathbb{R}^{p^{(h)} \times p^{(h-1)}}$ be a linear subspace.

Remark E.2. For convolutional neural network, the dimension of \mathcal{W} is the filter size.

In this section, by Gaussian distribution \mathcal{P} over a q -dimensional subspace \mathcal{W} , we mean that for a basis $\{\mathbf{e}_1, \dots, \mathbf{e}_q\}$ of \mathcal{W} and $(v_1, \dots, v_q) \sim N(\mathbf{0}, \mathbf{I})$ such that $\sum_{i=1}^q v_i \mathbf{e}_i \sim \mathcal{P}$. In this section, we equip one Gaussian distribution $\mathcal{P}^{(h)}$ with each linear subspace $\mathcal{W}^{(h)}$. By an abuse of notation, we also use \mathcal{W} to denote a transformation. For $\mathbf{K} \in \mathbb{R}^{p^{(h-1)} \times p^{(h-1)}}$, we let

$$\mathcal{W}^{(h)}(\mathbf{K}) = \mathbb{E}_{\mathbf{W} \sim \mathcal{P}^{(h)}} [\mathbf{W} \mathbf{K} \mathbf{W}^T].$$

We also consider a deterministic linear mapping $\mathcal{D}^{(h)} : \mathbb{R}^{n^{(h-1)}} \rightarrow \mathbb{R}^{n^{(h)}}$. For this section, we denote $\mathcal{D}^{(1)} = \mathbf{0}$, i.e., the zero mapping.

Remark E.3. For full-connected neural networks, we take $\mathcal{D}^{(h)}$ to be the zero mapping. For ResNet and convolutional ResNet, we take $\mathcal{D}^{(h)}$ to be the identity mapping.

Let $\rho^{(1)}, \dots, \rho^{(H)}$ be a sequence of activation functions over \mathbb{R} . Note here we use ρ instead of σ to denote the activation function because we will incorporate the scaling in ρ for the ease of presentation and the full generality.

Now we recursively define the output of each layer in this setup. In the following, we use $h \in [H]$ to index layers, $i \in [n]$ to index data points, $\alpha, \beta, \gamma \in [m]$ or $[d]$ to index channels (for CNN) or weight vectors (for fully connected neural networks or ResNet).

Remark E.4. $d = 1$ for fully connected neural network and ResNet and $d \geq 1$ for convolutional neural network because d represents the number of input channels.

We denote $\mathbf{X}_i^{(h),[\alpha]}$ an $p^{(h)}$ -dimensional vector which is the output at $(h-1)$ -th layer. We have the following recursive formula

$$\begin{aligned} \mathbf{X}_i^{(1),(\alpha)} &= \rho^{(1)} \left(\sum_{\beta} \mathbf{W}_{(\beta)}^{(1),(\alpha)} \mathbf{X}_i^{(0),(\beta)} \right) \\ \mathbf{X}_i^{(h),(\alpha)} &= \mathcal{D}^{(h)}(\mathbf{X}_i^{(h-1),(\alpha)}) + \rho^{(h)} \left(\frac{\sum_{\beta} \mathbf{W}_{(\beta)}^{(h),(\alpha)} \mathbf{X}_i^{(h-1),(\beta)}}{\sqrt{m}} \right) \end{aligned}$$

where $\mathbf{W}_{(\beta)}^{(h),(\alpha)}$ is $p^{(h)} \times p^{(h-1)}$ matrix generated according to the following rule

- for $h = 1$, $\mathbf{W}_{[\beta]}^{(h),(\alpha)}$ is defined for $1 \leq \alpha \leq m$ and $1 \leq \beta \leq d$; for $h > 1$, $\mathbf{W}_{(\beta)}^{(h),(\alpha)}$ is defined for $1 \leq \alpha \leq m$ and $1 \leq \beta \leq m$;
- the set of random variables $\{\mathbf{W}_{(\beta)}^{(h),(\alpha)}\}_{h,\alpha,\beta}$ are independently generated;
- for fixed h, α, β , $\mathbf{W}_{(\beta)}^{(h),(\alpha)} \sim \mathcal{P}^{(h)}$.

Remark E.5. Choosing $\rho^{(h)}(z)$ to be $\sigma(z)$ and $\mathcal{D}^{(h)}$ to be the zero mapping, we recover the fully-connected architecture. Choosing $\rho^{(h)}(z)$ to be $\frac{\text{cres}}{H} \sigma(z)$ and $\mathcal{D}^{(h)}$ to be the identity mapping, we recover ResNet architecture.

Remark E.6. Note here $\mathbf{X}_i^{(h)} = \mathbf{x}_i^{(h)} \sqrt{m}$ for $h \geq 1$ and $\mathbf{X}_i^{(h)} = \mathbf{x}_i^{(h)}$ for $h = 0$ in the main text. We change the scaling here to simplify the calculation of expectation and the covariance in this section.

With these notations, we first define the population Gram matrices recursively.

Definition E.1. We fix $(i, j) \in [n] \times [n]$, for $h = 1, \dots, H$. The population Gram matrices are defined according to the following formula

$$\begin{aligned} \mathbf{K}_{ij}^{(0)} &= \sum_{\gamma} (\mathbf{X}_i^{(0),[\gamma]})^\top \mathbf{X}_j^{(0),[\gamma]}, \\ \mathbf{b}_i^{(0)} &= \mathbf{0}, \\ \mathbf{K}_{ij}^{(h)} &= \mathcal{D}^{(h)} \mathbf{K}_{ij}^{(h-1)} \mathcal{D}^{(h)\top} + \mathbb{E}_{(\mathbf{U}, \mathbf{V})} \left(\rho(\mathbf{U}) \mathcal{D}^{(h)} (\mathbf{b}_j^{(h-1)})^\top + (\mathcal{D}^{(h)} (\mathbf{b}_i^{(h-1)})) \rho(\mathbf{V})^\top + \rho(\mathbf{U}) \rho(\mathbf{V})^\top \right), \\ \mathbf{b}_i^{(h)} &= \mathcal{D}^{(h)} (\mathbf{b}_i^{(h-1)}) + \mathbb{E}_{\mathbf{U}} \rho^{(h)}(\mathbf{U}), \end{aligned} \quad (16)$$

where

$$(\mathbf{U}, \mathbf{V}) \sim N \left(\mathbf{0}, \begin{pmatrix} \mathcal{W} \begin{pmatrix} \mathbf{K}_{ii}^{(h-1)} \\ \mathbf{K}_{ji}^{(h-1)} \end{pmatrix} & \mathcal{W} \begin{pmatrix} \mathbf{K}_{ij}^{(h-1)} \\ \mathbf{K}_{jj}^{(h-1)} \end{pmatrix} \end{pmatrix} \right). \quad (17)$$

Notice that the Gram matrix of the next layer $\mathbf{K}^{(h)}$ not only depends on the previous layer's Gram matrix $\mathbf{K}^{(h-1)}$ but also depends on a ‘‘bias’’ term $\mathbf{b}^{(h-1)}$.

Given the population Gram matrices defined in Equation (16) and (17), we derive the following quantitative bounds which characterizes how much over-parameterization, i.e., how large m is needed to ensure the randomly generated Gram matrices is close to the population Gram matrices.

Theorem E.1. With probability $1 - \delta$ over the $\{\mathbf{W}_{(\beta)}^{(h),(\alpha)}\}_{h,\alpha,\beta}$, for any $1 \leq h \leq H - 1, 1 \leq i, j \leq n$,

$$\left\| \frac{1}{m} \sum_{\alpha=1}^m (\mathbf{X}_i^{(h),(\alpha)})^\top \mathbf{X}_j^{(h),(\alpha)} - \mathbf{K}_{ij}^{(h)} \right\|_{\infty} \leq \mathcal{E} \sqrt{\frac{\log(Hn \max_h p^{(h)})/\delta}{m}} \quad (18)$$

and any $h \in [H - 1], \forall 1 \leq i \leq n$,

$$\left\| \frac{1}{m} \sum_{\alpha=1}^m \mathbf{X}_i^{(h),(\alpha)} - \mathbf{b}_i^{(h)} \right\|_{\infty} \leq \mathcal{E} \sqrt{\frac{\log(Hn \max_h p^{(h)})/\delta}{m}} \quad (19)$$

The error constant \mathcal{E} satisfies there exists an absolute constant $C > 0$ such that

$$\mathcal{E} \leq C \left(\prod_{h=2}^{H-1} \left(A_{(h)} + \Lambda_{(h)} \mathfrak{W} + C_{(h)} A_{(h)} B \mathfrak{W} + C_{(h)} A_{(h)} \sqrt{\mathfrak{W}_{(h)} M} \right) \right) \times \max\{\mathfrak{W} \sqrt{(1 + C_{(1)}^2) M^2}, \sqrt{C_{(1)}^2 M}\}$$

where $M, B, \Lambda_{(h)}, C_{(h)}, A_{(h)}, \mathfrak{W}_{(h)}$ are defined by:

- $M = 1 + 100 \max_{i,j,p,q,h} |\mathcal{W}^{(h)}(\mathbf{K}_{ij}^{(h-1)})_{pq}|$,
- $A_{(h)} = 1 + \max\{\|\mathcal{D}^{(h)}\|_{L^\infty \rightarrow L^\infty}, \|\mathcal{D}^{(h)}(\cdot) \mathcal{D}^{(h)\top}\|_{L^\infty \rightarrow L^\infty}\}$,
- $B = 1 + 100 \max_{i,p,h} |\mathbf{b}_{ip}^{(h)}|$,
- $C_{(h)} = |\rho(0)| + \sup_{x \in \mathbb{R}} |\rho'(x)|$,
- $\Lambda_{(h)}$ is a constant that only depends on $\rho^{(h)}$,
- $\mathfrak{W}_{(h)} = 1 + \|\mathcal{W}^{(h)}\|_{L^\infty \rightarrow L^\infty}$.

Remark E.7. For fully-connected neural networks, we have $M = O(1)$, $A_{(h)} = 0$, $B = O(1)$, $C_{(h)} = O(1)$, $\Lambda_{(h)} = O(1)$, $\mathfrak{W}_{(h)} = O(1)$, so we need $m = \Omega\left(\frac{n^2 \log(Hn/\delta) 2^{O(H)}}{\lambda_0^2}\right)$. For ResNet, we have $M = O(1)$, $A_{(h)} = 1$, $B = O(1)$, $C_{(h)} = O(\frac{1}{H})$, $\Lambda_{(h)} = O(\frac{1}{H})$, $\mathfrak{W}_{(h)} = O(1)$, so we need $m = \Omega\left(\frac{n^2 \log(Hn/\delta)}{\lambda_0^2}\right)$. The convolutional ResNet has the same parameters as ResNet but because the Gram matrix is $np \times np$, so we need $m = \Omega\left(\frac{n^2 p^2 \log(Hnp/\delta)}{\lambda_0^2}\right)$.

Proof of Theorem E.1. The proof is by induction. For the base case, $h = 1$, recall

$$\mathbf{X}_i^{(1),[\alpha]} = \rho^{(1)}\left(\sum_{\beta} \mathbf{W}_{(\beta)}^{(1),(\alpha)} \mathbf{X}_i^{(0),(\beta)}\right).$$

We define

$$\mathbf{U}_i^{(1),(\alpha)} = \sum_{\beta} W_{(\beta)}^{(1),(\alpha)} \mathbf{X}_i^{(0),(\beta)}.$$

By our generating process of $\left\{\mathbf{W}_{(\beta)}^{(h),(\alpha)}\right\}_{h,\alpha,\beta}$, the collection $\left\{\mathbf{U}_i^{(1),(\beta)}\right\}_{1 \leq i \leq n, 1 \leq \beta \leq m}$ is a mean-zero Gaussian variable with covariance matrix:

$$\begin{aligned} & \mathbb{E} \mathbf{U}_i^{(1),(\alpha)} \left(\mathbf{U}_j^{(1),(\beta)}\right)^{\top} \\ &= \mathbb{E} \sum_{\gamma, \gamma'} \mathbf{W}_{(\gamma)}^{(1),(\alpha)} \mathbf{X}_i^{(0),(\gamma)} \left(\mathbf{X}_j^{(0),(\gamma')} \right)^{\top} \left(\mathbf{W}_{(\gamma')}^{(1),(\beta)}\right)^{\top} \\ &= \delta_{\alpha\beta} \mathcal{W}^{(1)} \left(\sum_{\gamma} \left(\mathbf{X}_i^{(0),(\gamma)} \mathbf{X}_j^{(0),(\gamma)}\right)^{\top}\right) \\ &= \delta_{\alpha\beta} \mathcal{W}^{(1)}(\mathbf{K}_{ij}^{(0)}) \end{aligned}$$

Therefore, we have

$$\begin{aligned} \mathbb{E} \left[\frac{1}{m} \sum_{i=1}^m \mathbf{X}_i^{(1),(\alpha)} \mathbf{X}_j^{(1),(\alpha)\top} \right] &= \mathbf{K}_{ij}^{(1)} \\ \mathbb{E} \left[\frac{1}{m} \sum_{i=1}^m \mathbf{X}_i^{(1),(\alpha)} \right] &= \mathbf{b}_i^{(1)}. \end{aligned}$$

Now we have calculated the expectation. Note since inside the expectation is an average, we can apply standard Bernstein bounds and Hoeffding bound and obtain the following concentration inequalities. With probability at least $1 - \frac{\delta}{H}$, we have

$$\begin{aligned} \max_{i,j} \left\| \frac{1}{m} \sum_{i=1}^m \mathbf{X}_i^{(1),(\alpha)} \mathbf{X}_j^{(1),(\alpha)\top} - \mathbf{K}_{ij}^{(1)} \right\|_{\infty} &\leq \sqrt{\frac{16(1 + 2C_{(1)}^2/\sqrt{\pi})M^2 \log(4Hn^2(p^{(1)})^2/\delta)}{m}}, \\ \max_{i,p} \left| \frac{1}{m} \sum_{\alpha=1}^m \mathbf{X}_{ip}^{(1),(\alpha)} - \mathbf{b}_{ip}^{(1)} \right| &\leq \sqrt{\frac{2C_{(1)}^2 M \log(2np^{(1)}H/\delta)}{m}} \end{aligned}$$

Now we prove the induction step. Define for $1 \leq h \leq H$

$$\begin{aligned} \hat{\mathbf{K}}_{ij}^{(h)} &= \frac{1}{m} \sum_{\gamma} \mathbf{X}_i^{(h),(\gamma)} \left(\mathbf{X}_j^{(h),(\gamma)}\right)^{\top} \\ \hat{\mathbf{b}}_i^{(h)} &= \frac{1}{m} \sum_{\gamma} \mathbf{X}_i^{(1),(\gamma)} \end{aligned}$$

In the following, by $\mathbb{E}^{(h)}$ we mean taking expectation conditioned on first $(h - 1)$ layers.

Now suppose that Equation (18) and (19) hold for $1 \leq l \leq h$ with probability at least $1 - \frac{h}{H} \delta$, now we want to show the equations holds for $h + 1$ with probability at least $1 - \delta/H$ conditioned on previous layers satisfying Equation (18) and (19). Let $l = h + 1$. recall

$$\mathbf{X}_i^{(l),(\alpha)} = \mathcal{D}^{(l)}(\mathbf{X}^{(l-1)}) + \rho^{(l)} \left(\frac{\sum_{\beta} \mathbf{W}_{(\beta)}^{(l),(\alpha)} \mathbf{X}_i^{(l-1),(\beta)}}{\sqrt{m}} \right).$$

Similar to the base case, denote

$$\mathbf{U}_i^{(l),(\alpha)} = \frac{\sum_{\beta} \mathbf{W}_{(\beta)}^{(l),(\alpha)} \mathbf{X}_i^{(l-1),(\beta)}}{\sqrt{m}}.$$

Again note that $\{\mathbf{U}_i^{(l),(\beta)}\}_{1 \leq i \leq n, 1 \leq \beta \leq m}$ is a collection of mean-zero Gaussian variables with covariance matrix:

$$\mathbb{E} \left[\mathbf{U}_i^{(1),(\alpha)} \left(\mathbf{U}_j^{(1),(\beta)} \right)^{\top} \right] = \delta_{\alpha\beta} \mathcal{W}^{(l)}(\hat{\mathbf{K}}_{ij}^{(l-1)})$$

Now we get the following formula for the expectation:

$$\begin{aligned} \mathbb{E}^{(l)}[\hat{\mathbf{K}}_{ij}^{(l)}] &= \mathcal{D}^{(l)} \hat{\mathbf{K}}_{ij}^{(l-1)} \left(\mathcal{D}^{(l)} \right)^{\top} + \mathbb{E}_{(\mathbf{U}, \mathbf{V})} \left(\rho^{(l)}(\mathbf{U})^{\top} \mathcal{D}^{(l)}(\hat{\mathbf{b}}_j^{(l-1)}) + (\mathcal{D}^{(l)}(\hat{\mathbf{b}}_i^{(l-1)}))^{\top} \rho^{(l)}(\mathbf{V}) + \rho^{(l)}(\mathbf{U})^{\top} \rho^{(l)}(\mathbf{V}) \right) \\ \mathbb{E}^{(l)} \hat{\mathbf{b}}_i^{(l)} &= \mathcal{D}^{(l)}(\hat{\mathbf{b}}_i^{(l-1)}) + \mathbb{E}_{\mathbf{U}} \rho^{(l)}(\mathbf{U}) \end{aligned}$$

with

$$(\mathbf{U}, \mathbf{V}) \sim N \left(\mathbf{0}, \begin{pmatrix} \mathcal{W}^{(l)}(\hat{\mathbf{K}}_{ii}^{(l-1)}) & \mathcal{W}^{(l)}(\hat{\mathbf{K}}_{ij}^{(l-1)}) \\ \mathcal{W}^{(l)}(\hat{\mathbf{K}}_{ji}^{(l-1)}) & \mathcal{W}^{(l)}(\hat{\mathbf{K}}_{jj}^{(l-1)}) \end{pmatrix} \right)$$

Same as the base case, applying concentration inequalities, we have with probability at least $1 - \delta/H$,

$$\begin{aligned} \max_{ij} \|\mathbb{E}^{(l)} \hat{\mathbf{K}}_{ij}^{(l)} - \hat{\mathbf{K}}_{ij}^{(l)}\|_{\infty} &\leq \sqrt{\frac{16(1 + 2C_{(l)}^2/\sqrt{\pi})M^2 \log(4Hn^2(p^{(l)})^2/\delta)}{m}}, \\ \max_i \|\mathbb{E}^{(l)} \hat{\mathbf{b}}_i^{(l)} - \hat{\mathbf{b}}_i^{(l)}\|_{\infty} &\leq \sqrt{\frac{2C_{(l)}^2 M \log(2np^{(1)}H/\delta)}{m}} \end{aligned}$$

Now it remains to bound the differences

$$\max_{ij} \left\| \mathbb{E}^{(l)} \hat{\mathbf{K}}_{ij}^{(l)} - \mathbf{K}_{ij}^{(l)} \right\|_{\infty} \quad \text{and} \quad \max_i \left\| \mathbb{E}^{(l)} \hat{\mathbf{b}}_i^{(l)} - \mathbf{b}_{ij}^{(l)} \right\|_{\infty}$$

which determine how the error propagates through layers.

We analyze the error directly.

$$\begin{aligned} &\left\| \mathbb{E}^{(l)} \hat{\mathbf{K}}_{ij}^{(l)} - \mathbf{K}_{ij}^{(l)} \right\|_{\infty} \\ &\leq \left\| \mathcal{D}^{(l)} \hat{\mathbf{K}}_{ij}^{(l-1)} \mathcal{D}^{(l)\top} - \mathcal{D}^{(l)} \mathbf{K}_{ij}^{(l-1)} \mathcal{D}^{(l)\top} \right\|_{\infty} \\ &\quad + \left\| \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \hat{\mathbf{A}}} \rho^{(l)}(\mathbf{U})^{\top} \mathcal{D}^{(l)}(\hat{\mathbf{b}}_j^{(l-1)}) - \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \mathbf{A}} \rho^{(l)}(\mathbf{U})^{\top} \mathcal{D}^{(l)}(\mathbf{b}_j^{(l-1)}) \right\|_{\infty} \\ &\quad + \left\| \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \hat{\mathbf{A}}} (\mathcal{D}^{(l)}(\hat{\mathbf{b}}_i^{(l-1)}))^{\top} \rho^{(l)}(\mathbf{V}) - \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \mathbf{A}} (\mathcal{D}^{(l)}(\mathbf{b}_i^{(l-1)}))^{\top} \rho^{(l)}(\mathbf{V}) \right\|_{\infty} \\ &\quad + \left\| \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \hat{\mathbf{A}}} \rho^{(l)}(\mathbf{U})^{\top} \rho^{(l)}(\mathbf{V}) - \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \mathbf{A}} \rho^{(l)}(\mathbf{U})^{\top} \rho^{(l)}(\mathbf{V}) \right\|_{\infty} \end{aligned}$$

$$\begin{aligned}
 &\leq \left\| \mathcal{D}^{(l)} \hat{\mathbf{K}}_{ij}^{(l-1)} \mathcal{D}^{(l)\top} - \mathcal{D}^{(l)} \mathbf{K}_{ij}^{(l-1)} \mathcal{D}^{(l)\top} \right\|_{\infty} \\
 &+ \left\| \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \hat{\mathbf{A}}} \rho^{(l)}(\mathbf{U})^{\top} \mathcal{D}^{(l)}(\hat{\mathbf{b}}_j^{(l-1)}) - \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \hat{\mathbf{A}}} \rho^{(l)}(\mathbf{U})^{\top} \mathcal{D}^{(l)}(\mathbf{b}_j^{(l-1)}) \right\|_{\infty} \\
 &+ \left\| \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \hat{\mathbf{A}}} \rho^{(l)}(\mathbf{U})^{\top} \mathcal{D}^{(l)}(\mathbf{b}_j^{(l-1)}) - \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \mathbf{A}} \rho^{(l)}(\mathbf{U})^{\top} \mathcal{D}^{(l)}(\mathbf{b}_j^{(l-1)}) \right\|_{\infty} \\
 &+ \left\| \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \hat{\mathbf{A}}} (\mathbf{a}^{(l)}(\hat{\mathbf{b}}_i^{(l-1)}))^{\top} \rho^{(l)}(\mathbf{V}) - \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \hat{\mathbf{A}}} (\mathcal{D}^{(l)}(\mathbf{b}_i^{(l-1)}))^{\top} \rho^{(l)}(\mathbf{V}) \right\|_{\infty} \\
 &+ \left\| \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \hat{\mathbf{A}}} (\mathcal{D}^{(l)}(\mathbf{b}_i^{(l-1)}))^{\top} \rho^{(l)}(\mathbf{V}) - \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \mathbf{A}} (\mathcal{D}^{(l)}(\mathbf{b}_i^{(l-1)}))^{\top} \rho^{(l)}(\mathbf{V}) \right\|_{\infty} \\
 &+ \left\| \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \hat{\mathbf{A}}} \rho^{(l)}(\mathbf{U})^{\top} \rho^{(l)}(\mathbf{V}) - \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \mathbf{A}} \rho^{(l)}(\mathbf{U})^{\top} \rho^{(l)}(\mathbf{V}) \right\|_{\infty}
 \end{aligned}$$

where we define

$$\hat{\mathbf{A}} = \begin{pmatrix} \mathcal{W}^{(l)}(\hat{\mathbf{K}}_{ii}^{(l-1)}) & \mathcal{W}^{(l)}(\hat{\mathbf{K}}_{ij}^{(l-1)}) \\ \mathcal{W}^{(l)}(\hat{\mathbf{K}}_{ji}^{(l-1)}) & \mathcal{W}^{(l)}(\hat{\mathbf{K}}_{jj}^{(l-1)}) \end{pmatrix} \text{ and } \mathbf{A} = \begin{pmatrix} \mathcal{W}^{(l)}(\mathbf{K}_{ii}^{(l-1)}) & \mathcal{W}^{(l)}(\mathbf{K}_{ij}^{(l-1)}) \\ \mathcal{W}^{(l)}(\mathbf{K}_{ji}^{(l-1)}) & \mathcal{W}^{(l)}(\mathbf{K}_{jj}^{(l-1)}) \end{pmatrix}$$

By definition, we have

$$\begin{aligned}
 \|\mathbf{A} - \hat{\mathbf{A}}\|_{\infty} &\leq \mathfrak{W} \max_{ij} \|\hat{\mathbf{K}}_{ij}^{(l-1)} - \mathbf{K}_{ij}^{(l-1)}\|_{\infty} \text{ and} \\
 \left\| \mathcal{D}^{(l)} \hat{\mathbf{K}}_{ij}^{(l-1)} \mathcal{D}^{(l)\top} - \mathcal{D}^{(l)} \mathbf{K}_{ij}^{(l-1)} \mathcal{D}^{(l)\top} \right\|_{\infty} &\leq A^{(l)} \max_{ij} \|\hat{\mathbf{K}}_{ij}^{(l-1)} - \mathbf{K}_{ij}^{(l-1)}\|_{\infty}.
 \end{aligned}$$

We can also estimate other terms

$$\begin{aligned}
 &\left\| \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \hat{\mathbf{A}}} \rho^{(l)}(\mathbf{U})^{\top} \mathcal{D}^{(l)}(\hat{\mathbf{b}}_j^{(l-1)}) - \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \hat{\mathbf{A}}} \rho^{(l)}(\mathbf{U})^{\top} \mathcal{D}^{(l)}(\mathbf{b}_j^{(l-1)}) \right\|_{\infty} \\
 &\leq \left\| \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \hat{\mathbf{A}}} \rho^{(l)}(\mathbf{U})^{\top} \mathcal{D}^{(l)}(\hat{\mathbf{b}}_j^{(l-1)} - \mathbf{b}_j^{(l-1)}) \right\|_{\infty} \\
 &\leq C^{(l)} A^{(l)} \sqrt{\mathfrak{W} \max_{ij} \|\hat{\mathbf{K}}_{ij}^{(l)}\|_{\infty}} \max_i \|\hat{\mathbf{b}}_i^{(l-1)} - \mathbf{b}_i^{(l-1)}\|_{\infty} \\
 &\leq C^{(l)} A^{(l)} \sqrt{\mathfrak{W}^{(l)} M} \max_i \|\hat{\mathbf{b}}_i^{(l-1)} - \mathbf{b}_i^{(l-1)}\|_{\infty},
 \end{aligned}$$

$$\begin{aligned}
 &\left\| \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \hat{\mathbf{A}}} \rho^{(l)}(\mathbf{U})^{\top} \mathcal{D}^{(l)}(\mathbf{b}_j^{(l-1)}) - \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \mathbf{A}} \rho^{(l)}(\mathbf{U})^{\top} \mathcal{D}^{(l)}(\mathbf{b}_j^{(l-1)}) \right\|_{\infty} \\
 &\leq A^{(l)} B C^{(l)} \|\mathbf{A} - \hat{\mathbf{A}}\|_{\infty} \\
 &\leq A^{(l)} B C^{(l)} \mathfrak{W} \max_{ij} \|\hat{\mathbf{K}}_{ij}^{(l-1)} - \mathbf{K}_{ij}^{(l-1)}\|_{\infty},
 \end{aligned}$$

and

$$\begin{aligned}
 &\left\| \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \hat{\mathbf{A}}} \rho^{(l)}(\mathbf{U})^{\top} \rho^{(l)}(\mathbf{V}) - \mathbb{E}_{(\mathbf{U}, \mathbf{V}) \sim \mathbf{A}} \rho^{(l)}(\mathbf{U})^{\top} \rho^{(l)}(\mathbf{V}) \right\|_{\infty} \\
 &\leq \Lambda^{(l)} \|\mathbf{A} - \hat{\mathbf{A}}\|_{\infty} \\
 &\leq \Lambda^{(l)} \mathfrak{W} \max_{ij} \|\hat{\mathbf{K}}_{ij}^{(l-1)} - \mathbf{K}_{ij}^{(l-1)}\|_{\infty}.
 \end{aligned}$$

where we have used Lemma G.5.

Putting these estimates together, we have

$$\max_{ij} \|\mathbb{E}^{(l)} \hat{\mathbf{K}}_{ij}^{(l)} - \mathbf{K}_{ij}^{(l)}\|_{\infty}$$

$$\begin{aligned} &\leq (A_{(l)} + \Lambda_{(l)}\mathfrak{W} + 2C_{(l)}A_{(l)}B\mathfrak{W}) \max_{ij} \|\hat{\mathbf{K}}_{ij}^{(l-1)} - \mathbf{K}_{ij}^{(l-1)}\|_\infty + 2C_{(l)}A_{(l)}\sqrt{\mathfrak{W}_{(l)}M} \max_i \|\hat{\mathbf{b}}_{ij}^{(l-1)} - \mathbf{b}_{ij}^{(l-1)}\|_\infty \\ &\leq \left(A_{(l)} + \Lambda_{(l)}\mathfrak{W} + 2C_{(l)}A_{(l)}B\mathfrak{W} + 2C_{(l)}A_{(l)}\sqrt{\mathfrak{W}_{(l)}M} \right) \left(\max_{ij} \|\hat{\mathbf{K}}_{ij}^{(l-1)} - \mathbf{K}_{ij}^{(l-1)}\|_\infty \vee \max_i \|\hat{\mathbf{b}}_{ij}^{(l-1)} - \mathbf{b}_{ij}^{(l-1)}\|_\infty \right) \end{aligned}$$

and

$$\begin{aligned} \max_i \left\| \mathbb{E}^{(l)} \hat{\mathbf{b}}_{ij}^{(l)} - \mathbf{b}_{ij}^{(l)} \right\|_\infty &\leq \Lambda_{(l)}\mathfrak{W} \max_{ij} \left\| \hat{\mathbf{K}}_{ij}^{(l-1)} - \mathbf{K}_{ij}^{(l-1)} \right\|_\infty + A_{(l)} \max_i \left\| \hat{\mathbf{b}}_{ij}^{(l-1)} - \mathbf{b}_{ij}^{(l-1)} \right\|_\infty \\ &\leq (A_{(l)} + \Lambda_{(l)}\mathfrak{W}) \left(\max_{ij} \|\hat{\mathbf{K}}_{ij}^{(l-1)} - \mathbf{K}_{ij}^{(l-1)}\|_\infty \vee \max_i \|\hat{\mathbf{b}}_{ij}^{(l-1)} - \mathbf{b}_{ij}^{(l-1)}\|_\infty \right). \end{aligned}$$

These two bounds imply the theorem. \square

E.2. From $\mathbf{K}^{(H-1)}$ to $\mathbf{K}^{(H)}$

Recall $\mathbf{K}^{(H)}$ defined in Equation (7), (8) and (10). Note the definition of $\mathbf{K}^{(H)}$ is qualitatively different from that of $\mathbf{K}^{(h)}$ for $h = 1, \dots, H-1$ because $\mathbf{K}^{(H)}$ depends on $\mathbf{K}^{(H)}$ and $\sigma'(\cdot)$ instead of $\sigma(\cdot)$. Therefore, we take special care of $\mathbf{K}^{(H)}$. Further note $\mathbf{K}^{(H)}$ for our three architectures have the same form and only differ in scaling and dimension, so we will only prove the bound for the fully-connected architecture. The generalization to ResNet and convolutional ResNet is straightforward.

Lemma E.1. For $(i, j) \in [n] \times [n]$, define

$$\hat{\mathbf{K}}_{ij}^{(H-1)} = \hat{\mathbf{K}}_{ij}^{(H-1)} \mathbb{E}_{\mathbf{w} \sim N(\mathbf{0}, \mathbf{I})} \left[\sigma'(\mathbf{w}^\top \mathbf{x}_i^{(H-1)}(0)) \sigma'(\mathbf{w}^\top \mathbf{x}_j^{(H-1)}(0)) \right].$$

and suppose $\left| \hat{\mathbf{K}}_{ij}^{(H-1)} - \mathbf{K}_{ij}^{(H-1)} \right| \leq \frac{c\lambda_0}{n^2}$ for some small constant $c > 0$. Then if $m = \Omega\left(\frac{n^2 \log(n/\delta)}{\lambda_0^2}\right)$, we have with probability at least $1 - \delta$ over $\{\mathbf{w}_r^{(H)}(0)\}_{r=1}^m$ and $\{a_r(0)\}_{r=1}^m$, we have $\|\mathbf{G}^{(H)}(0) - \mathbf{K}^{(H)}\|_{op} \leq \frac{\lambda_0}{4}$.

Proof of Lemma E.1. We decompose

$$\mathbf{G}^{(H)}(0) - \mathbf{K}^{(H)} = \left(\mathbf{G}^{(H)}(0) - \hat{\mathbf{K}}^{(H)} \right) + \left(\hat{\mathbf{K}}^{(H)} - \mathbf{K}^{(H)} \right).$$

Recall $\mathbf{G}^{(H)}$ defined in Equation (13). Based on its expression, it is straightforward to use concentration inequality to show if $m = \Omega\left(\frac{n^2 \log(n/\delta)}{\lambda_0^2}\right)$, we have

$$\left\| \mathbf{G}^{(H)}(0) - \hat{\mathbf{K}}^{(H)} \right\|_{op} \leq \frac{\lambda_0}{8}.$$

For the other

$$\text{Recall } \mathbf{A}_{ij}^{(H)} = \begin{pmatrix} \mathbf{K}_{ii}^{(H-1)} & \mathbf{K}_{ij}^{(H-1)} \\ \mathbf{K}_{ji}^{(H-1)} & \mathbf{K}_{jj}^{(H-1)} \end{pmatrix} \text{ and let } \hat{\mathbf{A}}_{ij}^{(H)} = \begin{pmatrix} \hat{\mathbf{K}}_{ii}^{(H-1)} & \hat{\mathbf{K}}_{ij}^{(H-1)} \\ \hat{\mathbf{K}}_{ji}^{(H-1)} & \hat{\mathbf{K}}_{jj}^{(H-1)} \end{pmatrix}.$$

According to Lemma G.4 (viewing $\sigma'(\cdot)$ as the $\sigma(\cdot)$ in Lemma G.4), we know

$$\left| \mathbb{E}_{(u,v) \sim \hat{\mathbf{A}}_{ij}} [\sigma'(u)\sigma'(v)] - \mathbb{E}_{(u,v) \sim \mathbf{A}_{ij}} [\sigma'(u)\sigma'(v)] \right| \leq C \left| \hat{\mathbf{A}}_{ij} - \mathbf{A}_{ij} \right|$$

for some constant $C > 0$. Since c is small enough, we directly have

$$\left\| \hat{\mathbf{K}}^{(H)} - \mathbf{K}^{(H)} \right\|_{op} \leq \frac{\lambda_0}{8}$$

\square

Remark E.8. Combing Theorem E.1, Lemma E.1 and standard matrix perturbation bound directly have Lemma B.2. Similarly we can prove Lemma C.2 and Lemma D.2.

F. Full Rankness of $\mathbf{K}^{(h)}$

F.1. Full Rankness of $\mathbf{K}^{(h)}$ for the Fully-connected Neural Network

In this section we show as long as no two input vectors are parallel, then $\mathbf{K}^{(H)}$ defined in Equation (8) is strictly positive definite.

Proposition F.1. *Assume $\sigma(\cdot)$ satisfies Condition 3.2 and for any $i, j \in [n], i \neq j, \mathbf{x}_i \not\parallel \mathbf{x}_j$. Then we have $\lambda_{\min}(\mathbf{K}^{(H)}) > 0$ where $\lambda_{\min}(\mathbf{K}^{(H)})$ is defined in Equation (7).*

Proof of Proposition F.1. By our assumption on the data point and using Lemma F.1 we know $\mathbf{K}^{(1)}$ is strictly positive definite.

By letting $\mathbf{Z} = \mathbf{D}^{1/2}\mathbf{U}^\top$, where $\mathbf{U}\mathbf{D}\mathbf{U}^\top = \mathbf{K}^h$. We then use Lemma F.1 inductively for $(H - 2)$ times to conclude $\mathbf{K}^{(H-1)}$ is strictly positive definite. Lastly we use Lemma F.2 to finish the proof. \square

Lemma F.1. *Assume $\sigma(\cdot)$ is analytic and not a polynomial function. Consider data $Z = \{\mathbf{z}_i\}_{i \in [n]}$ of n non-parallel points (meaning $\mathbf{z}_i \notin \text{span}(\mathbf{z}_j)$ for all $i \neq j$). Define*

$$\mathbf{G}(Z)_{ij} = \mathbb{E}_{\mathbf{w} \sim N(\mathbf{0}, \mathbf{I})} [\sigma(\mathbf{w}^\top \mathbf{z}_i) \sigma(\mathbf{w}^\top \mathbf{z}_j)].$$

Then $\lambda_{\min}(\mathbf{G}(Z)) > 0$.

Proof of Lemma F.1. The feature map induced by the kernel \mathbf{G} is given by $\phi_{\mathbf{z}}(\mathbf{w}) = \sigma(\mathbf{w}^\top \mathbf{z})\mathbf{z}$. To show that $\mathbf{G}(Z)$ is strictly positive definite, we need to show $\phi_{\mathbf{z}_1}(\mathbf{w}), \dots, \phi_{\mathbf{z}_n}(\mathbf{w})$ are linearly independent functions. Assume that there are a_i such that

$$0 = \sum_i a_i \phi_{\mathbf{z}_i} = \sum_i a_i \sigma(\mathbf{w}^\top \mathbf{z}_i) \mathbf{z}_i.$$

We wish to show that $a_i = 0$. Differentiating the above equation $(n - 2)$ times with respect to \mathbf{w} , we have

$$0 = \sum_i \left(a_i \sigma^{(n-1)}(\mathbf{w}^\top \mathbf{z}_i) \right) \mathbf{z}_i^{\otimes (n-1)}.$$

Using Lemma G.6, we know $\left\{ \mathbf{z}_i^{\otimes (n-1)} \right\}_{i=1}^n$ are linearly independent. Therefore, we must have $a_i \sigma^{(n-1)}(\mathbf{w}^\top \mathbf{z}_i) = 0$ for all i . Now choosing a \mathbf{w} such that $\sigma^{(n-1)}(\mathbf{w}^\top \mathbf{z}_i) \neq 0$ for all $i \in [n]$ (such \mathbf{w} exists because of our assumption on σ), we have $a_i = 0$ for all $i \in [n]$. \square

Lemma F.2. *Assume $\sigma(\cdot)$ is analytic and not a polynomial function. Consider data $Z = \{\mathbf{z}_i\}_{i \in [n]}$ of n non-parallel points (meaning $\mathbf{z}_i \notin \text{span}(\mathbf{z}_j)$ for all $i \neq j$). Define*

$$\mathbf{G}(Z)_{ij} = \mathbb{E}_{\mathbf{w} \sim N(\mathbf{0}, \mathbf{I})} [\sigma'(\mathbf{w}^\top \mathbf{z}_i) \sigma'(\mathbf{w}^\top \mathbf{z}_j) (\mathbf{z}_i^\top \mathbf{z}_j)].$$

Then $\lambda_{\min}(\mathbf{G}(Z)) > 0$.

Proof of Lemma F.2. The feature map induced by the kernel \mathbf{G} is given by $\phi_{\mathbf{z}}(\mathbf{w}) = \sigma'(\mathbf{w}^\top \mathbf{z})\mathbf{z}$. To show that $\mathbf{G}(Z)$ is strictly positive definite, we need to show $\phi_{\mathbf{z}_1}(\mathbf{w}), \dots, \phi_{\mathbf{z}_n}(\mathbf{w})$ are linearly independent functions. Assume that there are a_i such that

$$0 = \sum_i a_i \phi_{\mathbf{z}_i} = \sum_i a_i \sigma'(\mathbf{w}^\top \mathbf{z}_i) \mathbf{z}_i.$$

We wish to show that $a_i = 0$. Differentiating the above equation $(n - 2)$ times with respect to \mathbf{w} , we have

$$0 = \sum_i \left(a_i \sigma^{(n)}(\mathbf{w}^\top \mathbf{z}_i) \right) \mathbf{z}_i^{\otimes (n-1)}.$$

Using Lemma G.6, we know $\left\{ \mathbf{z}_i^{\otimes (n-1)} \right\}_{i=1}^n$ are linearly independent. Therefore, we must have $a_i \sigma^{(n)}(\mathbf{w}^\top \mathbf{z}_i) = 0$ for all i . Now choosing a \mathbf{w} such that $\sigma^{(n)}(\mathbf{w}^\top \mathbf{z}_i) \neq 0$ for all $i \in [n]$ (such \mathbf{w} exists because of our assumption on σ), we have $a_i = 0$ for all $i \in [n]$. \square

F.2. Full Rankness of $\mathbf{K}^{(h)}$ for ResNet

In this section we show as long as no two input vectors are parallel, then $\mathbf{K}^{(H)}$ defined in Equation (8) is strictly positive definite. Furthermore, $\lambda_{\min}(\mathbf{K}^{(H)})$ does not depend inverse exponentially in H .

Proposition F.2. *Assume $\sigma(\cdot)$ satisfies Condition 3.2 and for any $i, j \in [n], i \neq j, \mathbf{x}_i \not\parallel \mathbf{x}_j$. Recall that in Equation (8), we define*

$$\mathbf{K}_{ij}^{(H)} = c_H \mathbf{K}_{ij}^{(H-1)} \cdot \mathbb{E}_{(u,v)^\top \sim N\left(\mathbf{0}, \begin{pmatrix} \mathbf{K}_{ii}^{(H-1)} & \mathbf{K}_{ij}^{(H-1)} \\ \mathbf{K}_{ji}^{(H-1)} & \mathbf{K}_{jj}^{(H-1)} \end{pmatrix}\right)} [\sigma'(u)\sigma'(v)],$$

where $c_H \sim \frac{1}{H^2}$. Then we have $\lambda_{\min}(\mathbf{K}^{(H)}) \geq c_H \kappa$, where κ is a constant that only depends on the activation σ and the input data. In particular, κ does not depend on the depth.

Proof of Proposition F.2. First note $\mathbf{K}_{ii}^{(H-1)} \in [1/c_{x,0}^2, c_{x,0}^2]$ for all H , so it is in a bounded range that does not depend on the depth (c.f. Lemma C.1). Define a function

$$\mathbf{G} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$$

such that $\mathbf{G}(\mathbf{K})_{ij} = \mathbf{K}_{ij} \mathbb{E}_{(u,v)^\top \sim N\left(\mathbf{0}, \begin{pmatrix} \mathbf{K}_{ii} & \mathbf{K}_{ij} \\ \mathbf{K}_{ji} & \mathbf{K}_{jj} \end{pmatrix}\right)} [\sigma'(u)\sigma'(v)]$. Now define a scalar function

$$g(\lambda) = \min_{\mathbf{K} : \mathbf{K} \succ 0, \frac{1}{c_{x,0}^2} \leq \mathbf{K}_{ii} \leq c_{x,0}, \lambda(\mathbf{K}) \geq \lambda} \lambda_{\min}(\mathbf{G}(\mathbf{K}))$$

with

$$\lambda(\mathbf{K}) = \min_{ij} \begin{pmatrix} \mathbf{K}_{ii} & \mathbf{K}_{ij} \\ \mathbf{K}_{ji} & \mathbf{K}_{jj} \end{pmatrix}.$$

By Lemma F.3, we know $\lambda(\mathbf{K}^{(H-1)}) \geq c_H \lambda(\mathbf{K}^{(0)})$.

Next, let $\mathbf{U}\mathbf{D}\mathbf{U}^\top = \mathbf{K}^{(H-1)}$ be the eigen-decomposition of \mathbf{K} , and $\mathbf{Z} = \mathbf{D}^{1/2}\mathbf{U}^\top$ be the feature embedding into \mathbb{R}^n . Since $\begin{pmatrix} \mathbf{z}_i^\top \mathbf{z}_i & \mathbf{z}_i^\top \mathbf{z}_j \\ \mathbf{z}_j^\top \mathbf{z}_i & \mathbf{z}_j^\top \mathbf{z}_j \end{pmatrix}$ is full rank, then $\mathbf{z}_i \notin \text{span}(\mathbf{z}_j)$. Then using Lemma F.2, we know $g(\lambda(\mathbf{K}^{(0)})) > 0$. Thus we have established that $\lambda_{\min}(\mathbf{K}^{(H)}) \geq c_H g(\lambda(\mathbf{K}^{(0)}))$, where $g(\lambda(\mathbf{K}^{(0)}))$ only depends on the input data and activation σ . In particular, it is independent of the depth. \square

Lemma F.3. *If $\mathcal{D}^{(h)}$ is the identity mapping defined in Section E, then $\lambda(\mathbf{K}^{(H)}) \geq \min_{(i,j) \in [n] \times [n]} \lambda_{\min} \begin{pmatrix} \mathbf{K}_{ii}^{(0)} & \mathbf{K}_{ij}^{(0)} \\ \mathbf{K}_{ji}^{(0)} & \mathbf{K}_{jj}^{(0)} \end{pmatrix}$.*

Proof of Lemma F.3. First recall

$$(\mathbf{U}, \mathbf{V}) \sim N\left(\mathbf{0}, \begin{pmatrix} \mathcal{W}^{(h)}(\mathbf{K}_{ii}^{(h-1)}) & \mathcal{W}^{(h)}(\mathbf{K}_{ij}^{(h-1)}) \\ \mathcal{W}^{(h)}(\mathbf{K}_{ji}^{(h-1)}) & \mathcal{W}^{(h)}(\mathbf{K}_{jj}^{(h-1)}) \end{pmatrix}\right)$$

Then we compute

$$\begin{aligned} & \mathbf{K}_{ij}^{(h)} - \mathbf{b}_i^{(h)} \mathbf{b}_j^{(h)\top} \\ &= \mathcal{D}^{(h)} \mathbf{K}_{ij}^{(h-1)} \mathcal{D}^{(h)\top} + \mathbb{E}_{(\mathbf{U}, \mathbf{V})} \left(\rho(\mathbf{U}) \mathcal{D}^{(h)} (\mathbf{b}_j^{(h-1)})^\top + (\mathcal{D}^{(h)} (\mathbf{b}_i^{(h-1)})) \rho(\mathbf{V})^\top + \rho(\mathbf{U}) \rho(\mathbf{V})^\top \right) \\ & - \left(\mathcal{D}^{(h)} (\mathbf{b}_i^{(h-1)}) + \mathbb{E}_{\mathbf{U}} \rho^{(h)}(\mathbf{U}) \right) \left(\mathcal{D}^{(h)} (\mathbf{b}_j^{(h-1)}) + \mathbb{E}_{\mathbf{V}} \rho^{(h)}(\mathbf{V}) \right)^\top \\ &= \mathcal{D}^{(h)} \left(\mathbf{K}_{ij}^{(h-1)} - \mathbf{b}_i^{(h-1)} \mathbf{b}_j^{(h-1)\top} \right) \mathcal{D}^{(h)\top} + \mathbb{E}_{(\mathbf{U}, \mathbf{V})} \left(\rho(\mathbf{U}) \rho(\mathbf{V})^\top \right) - \left(\mathbb{E}_{\mathbf{U}} \rho^{(h)}(\mathbf{U}) \right) \left(\mathbb{E}_{\mathbf{V}} \rho^{(h)}(\mathbf{V}) \right)^\top \end{aligned}$$

For ResNet, $\mathcal{D}^{(h)}$ is the identity mapping so we have

$$\begin{aligned} & \mathbf{K}_{ij}^{(h)} - \mathbf{b}_i^{(h)} \mathbf{b}_j^{(h)\top} \\ &= \mathbf{K}_{ij}^{(h-1)} - \mathbf{b}_i^{(h-1)} \mathbf{b}_j^{(h-1)\top} + \mathbb{E}_{(\mathbf{U}, \mathbf{V})} \left(\rho(\mathbf{U}) \rho(\mathbf{V})^\top \right) - \left(\mathbb{E}_{\mathbf{U}} \rho^{(h)}(\mathbf{U}) \right) \left(\mathbb{E}_{\mathbf{V}} \rho^{(h)}(\mathbf{V}) \right)^\top. \end{aligned}$$

To proceed, we calculate

$$\begin{aligned} & \begin{pmatrix} \mathbf{K}_{ii}^{(h)} & \mathbf{K}_{ij}^{(h)} \\ \mathbf{K}_{ji}^{(h)} & \mathbf{K}_{jj}^{(h)} \end{pmatrix} - \begin{pmatrix} \mathbf{b}_i^{(h)} \\ \mathbf{b}_j^{(h)} \end{pmatrix} \begin{pmatrix} \mathbf{b}_i^{(h)\top} & \mathbf{b}_j^{(h)\top} \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{K}_{ii}^{(h-1)} & \mathbf{K}_{ij}^{(h-1)} \\ \mathbf{K}_{ji}^{(h-1)} & \mathbf{K}_{jj}^{(h-1)} \end{pmatrix} - \begin{pmatrix} \mathbf{b}_i^{(h-1)} \\ \mathbf{b}_j^{(h-1)} \end{pmatrix} \begin{pmatrix} \mathbf{b}_i^{(h-1)\top} & \mathbf{b}_j^{(h-1)\top} \end{pmatrix} \\ &+ \left(\mathbb{E}_{\mathbf{U}, \mathbf{V}} \begin{pmatrix} \rho^{(h)}(\mathbf{U}) \rho^{(h)}(\mathbf{U})^\top & \rho^{(h)}(\mathbf{U}) \rho^{(h)}(\mathbf{V})^\top \\ \rho^{(h)}(\mathbf{V}) \rho^{(h)}(\mathbf{U})^\top & \rho^{(h)}(\mathbf{V}) \rho^{(h)}(\mathbf{V})^\top \end{pmatrix} - \mathbb{E}_{\mathbf{U}, \mathbf{V}} \begin{pmatrix} \rho(\mathbf{U}) \\ \rho(\mathbf{V}) \end{pmatrix} \mathbb{E}_{\mathbf{U}, \mathbf{V}} \begin{pmatrix} \rho(\mathbf{U})^\top & \rho(\mathbf{V})^\top \end{pmatrix} \right) \\ &\geq \begin{pmatrix} \mathbf{K}_{ii}^{(h-1)} & \mathbf{K}_{ij}^{(h-1)} \\ \mathbf{K}_{ji}^{(h-1)} & \mathbf{K}_{jj}^{(h-1)} \end{pmatrix} - \begin{pmatrix} \mathbf{b}_i^{(h-1)} \\ \mathbf{b}_j^{(h-1)} \end{pmatrix} \begin{pmatrix} \mathbf{b}_i^{(h-1)\top} & \mathbf{b}_j^{(h-1)\top} \end{pmatrix} \end{aligned}$$

As a result, we have

$$\begin{aligned} & \lambda_{\min} \begin{pmatrix} \mathbf{K}_{ii}^{(h)} & \mathbf{K}_{ij}^{(h)} \\ \mathbf{K}_{ji}^{(h)} & \mathbf{K}_{jj}^{(h)} \end{pmatrix} \\ &\geq \lambda_{\min} \begin{pmatrix} \mathbf{K}_{ii}^{(h)} & \mathbf{K}_{ij}^{(h)} \\ \mathbf{K}_{ji}^{(h)} & \mathbf{K}_{jj}^{(h)} \end{pmatrix} - \begin{pmatrix} \mathbf{b}_i^{(h)} \\ \mathbf{b}_j^{(h)} \end{pmatrix} \begin{pmatrix} \mathbf{b}_i^{(h)\top} & \mathbf{b}_j^{(h)\top} \end{pmatrix} \\ &\geq \min \lambda_{\min} \begin{pmatrix} \mathbf{K}_{ii}^{(h-1)} & \mathbf{K}_{ij}^{(h-1)} \\ \mathbf{K}_{ji}^{(h-1)} & \mathbf{K}_{jj}^{(h-1)} \end{pmatrix} - \begin{pmatrix} \mathbf{b}_i^{(h-1)} \\ \mathbf{b}_j^{(h-1)} \end{pmatrix} \begin{pmatrix} \mathbf{b}_i^{(h-1)\top} & \mathbf{b}_j^{(h-1)\top} \end{pmatrix} \\ &\geq \dots \\ &\geq \lambda_{\min} \begin{pmatrix} \mathbf{K}_{ii}^{(0)} & \mathbf{K}_{ij}^{(0)} \\ \mathbf{K}_{ji}^{(0)} & \mathbf{K}_{jj}^{(0)} \end{pmatrix} - \begin{pmatrix} \mathbf{b}_i^{(0)} \\ \mathbf{b}_j^{(0)} \end{pmatrix} \begin{pmatrix} \mathbf{b}_i^{(0)\top} & \mathbf{b}_j^{(0)\top} \end{pmatrix} \\ &= \lambda_{\min} \begin{pmatrix} \mathbf{K}_{ii}^{(0)} & \mathbf{K}_{ij}^{(0)} \\ \mathbf{K}_{ji}^{(0)} & \mathbf{K}_{jj}^{(0)} \end{pmatrix}. \end{aligned} \tag{20}$$

We now prove the theorem. □

G. Useful Technical Lemmas

Lemma G.1. *Given a set of matrices $\{\mathbf{A}_i, \mathbf{B}_i : i \in [n]\}$, if $\|\mathbf{A}_i\|_2 \leq M_i$, $\|\mathbf{B}_i\|_2 \leq M_i$ and $\|\mathbf{A}_i - \mathbf{B}_i\|_F \leq \alpha_i M_i$, we have*

$$\left\| \prod_{i=1}^n \mathbf{A}_i - \prod_{i=1}^n \mathbf{B}_i \right\|_F \leq \left(\sum_{i=1}^n \alpha_i \right) \prod_{i=1}^n M_i.$$

Proof of Lemma G.1.

$$\left\| \prod_{i=1}^n \mathbf{A}_i - \prod_{i=1}^n \mathbf{B}_i \right\|_F$$

$$\begin{aligned}
 &= \left\| \sum_{i=1}^n \left(\prod_{j=1}^{i-1} \mathbf{A}_j \right) (\mathbf{A}_i - \mathbf{B}_i) \left(\prod_{k=i+1}^n \mathbf{B}_k \right) \right\|_F \\
 &\leq \sum_{i=1}^n \left\| \left(\prod_{j=1}^{i-1} \mathbf{A}_j \right) (\mathbf{A}_i - \mathbf{B}_i) \left(\prod_{k=i+1}^n \mathbf{B}_k \right) \right\|_F \\
 &\leq \left(\sum_{i=1}^n \alpha_i \right) \prod_{i=1}^n M_i.
 \end{aligned}$$

□

Lemma G.2. Given a matrix $\mathbf{W} \in \mathbb{R}^{m \times cm}$ with $\mathbf{W}_{i,j} \sim N(0, 1)$, where c is a constant. We have with probability at least $1 - \exp\left(-\frac{(c_{w,0} - \sqrt{c} - 1)^2 m}{2}\right)$

$$\|\mathbf{W}\|_2 \leq c_{w,0} \sqrt{m},$$

where $c_{w,0} > \sqrt{c} + 1$ is a constant.

Proof of Lemma G.2. The lemma is a consequence of well-known deviations bounds concerning the singular values of Gaussian random matrices (Vershynin, 2010)

$$P(\lambda_{\max}(\mathbf{W}) > \sqrt{m} + \sqrt{cm} + t) \leq e^{t^2/2}.$$

Choosing $t = (c_{w,0} - \sqrt{c} - 1) \sqrt{m}$, we prove the lemma. □

Lemma G.3. Assume $\sigma(\cdot)$ satisfies Condition 3.1. For $a, b \in \mathbb{R}$ with $\frac{1}{c} < \min(a, b)$, $\max(a, b) < c$ for some constant $c > 0$, we have

$$|\mathbb{E}_{z \sim N(0,a)}[\sigma(z)] - \mathbb{E}_{z \sim N(0,b)}[\sigma(z)]| \leq C |a - b|.$$

for some constant $C > 0$ that depends only on c and the constants in Condition 3.1.

Proof of Lemma G.3. We compute for any $\min(a, b) \leq \alpha \leq \max(a, b)$

$$\left| \frac{d\mathbb{E}_{z \sim N(0,\alpha)}[\sigma(z)]}{d\alpha} \right| = \left| \frac{d\mathbb{E}_{z \sim N(0,1)}[\sigma(\alpha z)]}{d\alpha} \right| = |\mathbb{E}_{z \sim N(0,1)}[z\sigma'(\alpha z)]| \leq C.$$

Applying Taylor's Theorem we finish the proof. □

Lemma G.4. Assume $\sigma(\cdot)$ satisfies Condition 3.1. Suppose that there exists some constant $c > 0$ such that $\mathbf{A} = \begin{bmatrix} a_1^2 & \rho a_1 b_1 \\ \rho a_1 b_1 & b_1^2 \end{bmatrix}$, $\frac{1}{c} \leq \min(a_1, b_1)$, $\max(a_1, b_1) \leq c$, $\mathbf{B} = \begin{bmatrix} a_2^2 & \rho a_2 b_2 \\ \rho a_2 b_2 & b_2^2 \end{bmatrix}$, $\frac{1}{c} \leq \min(a_2, b_2)$, $\max(a_2, b_2) \leq c$

and $\mathbf{A}, \mathbf{B} \succ 0$. Define $F(\mathbf{A}) = \mathbb{E}_{(u,v) \sim N(\mathbf{0}, \mathbf{A})} \sigma(u) \sigma(v)$. Then, we have

$$|F(\mathbf{A}) - F(\mathbf{B})| \leq C \|\mathbf{A} - \mathbf{B}\|_F \leq 2C \|\mathbf{A} - \mathbf{B}\|_\infty.$$

for some constant $C > 0$ that depends only on c and the constants in Condition 3.1.

Proof. Let $\mathbf{A}' = \begin{bmatrix} a^2 & \rho ab \\ \rho ab & b^2 \end{bmatrix} \succ 0$ with $\min(a_1, a_2) \leq a \leq \max(a_1, a_2)$, $\min(b_1, b_2) \leq b \leq \max(b_1, b_2)$ and $\min(\rho_1, \rho_2) \leq \rho \leq \max(\rho_1, \rho_2)$. We can express

$$F(\mathbf{A}') = \mathbb{E}_{(z_1, z_2) \sim \mathcal{N}(0, \mathbf{C})} \sigma(az_1) \sigma(bz_2) \text{ with } \mathbf{C} = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}.$$

Recall $\mathcal{L}^2 = \{f : \int f(z) e^{-z^2/2} dz < \infty\}$ is the Gaussian function space. We compute

$$\frac{dF}{da} = \mathbb{E}[\sigma'(az_1) \sigma(bz_2) z_1]$$

$$\begin{aligned} \left| \frac{dF}{da} \right| &\leq \|\sigma'(az_1)z_1\|_{L^2} \|\sigma(bz_2)\|_{L^2} && (\|f\|_{L^2} := (\mathbb{E}f(z)^2)^{1/2}, \text{Cauchy}) \\ &< \infty && \text{(by Condition 3.1)} \end{aligned}$$

By the same argument, we have

$$\left| \frac{dF}{db} \right| < \infty$$

Next, let $\sigma_a(z) := \sigma(az)$ with Hermite expansion $\sigma_a(z) = \sum_{i=0}^{\infty} \alpha_i h_i(z)$ and similarly $\sigma_b(z) = \sum_i \beta_i h_i(z)$. Using the orthonormality that $\mathbb{E}[h_i(z)h_j(z)] = 1_{i=j}$,

$$F(A) = \sum_{i=0}^{\infty} \alpha_i \beta_i \rho^i.$$

Differentiating, we have

$$\begin{aligned} \left| \frac{dF}{d\rho} \right| &= \left| \sum_{i=1}^{\infty} \alpha_i \beta_i i \rho^{i-1} \right| \\ &< \left(\sum_{i=1}^{\infty} \alpha_i^2 i \right)^{1/2} \left(\sum_{i=1}^{\infty} \beta_i^2 i \right)^{1/2} && (\rho = 1 \text{ and Cauchy}) \\ &< \infty && \text{(Condition 3.1)} \end{aligned}$$

Note by Condition 3.1 we know there exists B_ρ , B_a and B_b such that $\left| \frac{dF}{d\rho} \right| \leq B_\rho$, $\left| \frac{dF}{da} \right| \leq B_a$, and $\left| \frac{dF}{db} \right| \leq B_b$.

Next, we bound $\nabla_{\mathbf{A}'} F(\mathbf{A}')$.

We see that

$$\begin{aligned} \left| \frac{dF}{dA'_{11}} \right| &\leq \left| \frac{dF}{da} \right| \left| \frac{da}{dA'_{11}} \right| \\ &\leq B_a \frac{1}{2\sqrt{A'_{11}}} && \text{(since } a = \sqrt{A'_{11}} \text{)} \\ &\leq \frac{1}{2} B_a / c \\ \left| \frac{dF}{dA'_{11}} \right| &\leq \frac{1}{2} B_b / c && \text{(analogous argument as above.)} \end{aligned}$$

Using the change of variables, let

$$g(A'_{11}, A'_{22}, A'_{12}) = [\sqrt{A'_{11}}, \sqrt{A'_{22}}, A'_{12}/\sqrt{A'_{11}A'_{22}}] = [a, b, \rho].$$

By chain rule, we know

$$\frac{\partial F}{\partial A'_{12}} = \frac{\partial F}{\partial a} \frac{\partial a}{\partial A'_{12}} + \frac{\partial F}{\partial b} \frac{\partial b}{\partial A'_{12}} + \frac{\partial F}{\partial \rho} \frac{\partial \rho}{\partial A'_{12}} = \frac{\partial F}{\partial \rho} \frac{\partial \rho}{\partial A'_{12}}.$$

We can easily verify that $\left| \frac{\partial \rho}{\partial A'_{12}} \right| \leq 1/c^2$, and so we have

$$\left| \frac{\partial F}{\partial A'_{12}} \right| \leq \frac{B_\rho}{c^2}.$$

Similarly, we have

$$\left| \frac{\partial F}{\partial A'_{11}} \right| \leq \frac{B_a}{c^2}$$

$$\left| \frac{\partial F}{\partial A'_{22}} \right| \leq \frac{B_b}{c^2}$$

Define $B_\sigma = \max(B_a, B_b, B_\rho)$. This establishes $\|\nabla F(\mathbf{A}')\|_F \leq 2B_\sigma/c^2 \leq C$ for some constant $C > 0$. Thus by Taylor's Theorem, we have

$$|F(\mathbf{A}) - F(\mathbf{B})| \leq C\|\mathbf{A} - \mathbf{B}\|_F \leq 2C\|\mathbf{A} - \mathbf{B}\|_\infty.$$

□

With Lemma G.3 and G.4, we can prove the following useful lemma.

Lemma G.5. *Suppose $\sigma(\cdot)$ satisfies Condition 3.1 For a positive definite matrix $\mathbf{A} \in \mathbb{R}^{2p \times 2p}$, define*

$$\begin{aligned} \mathbf{F}(\mathbf{A}) &= \mathbb{E}_{\mathbf{U} \sim N(\mathbf{0}, \mathbf{A})} \left[\sigma(\mathbf{U}) \sigma(\mathbf{U})^\top \right], \\ \mathbf{G}(\mathbf{A}) &= \mathbb{E}_{\mathbf{U} \sim N(\mathbf{0}, \mathbf{A})} [\sigma(\mathbf{U})]. \end{aligned}$$

Then for any two positive definite matrices \mathbf{A}, \mathbf{B} with $\frac{1}{c} \leq \mathbf{A}_{ii}, \mathbf{B}_{ii} \leq c$ for some constant $c > 0$, we have

$$\|\mathbf{G}(\mathbf{A}) - \mathbf{G}(\mathbf{B})\|_\infty \vee \|\mathbf{F}(\mathbf{A}) - \mathbf{F}(\mathbf{B})\|_\infty \leq C\|\mathbf{A} - \mathbf{B}\|_\infty$$

for some constant $C > 0$.

Proof of Lemma G.5. The result follows by applying Lemma G.3 to all coordinates and applying Lemma G.4 to all 2×2 submatrices. □

Lemma G.6. *If $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^d$ satisfy that $\|\mathbf{v}_i\|_2 = 1$ and non-parallel (meaning $\mathbf{v}_i \notin \text{span}(\mathbf{v}_j)$ for $i \neq j$), then the matrix $[\text{vec}(\mathbf{v}_1^{\otimes n}), \dots, \text{vec}(\mathbf{v}_n^{\otimes n})] \in \mathbb{R}^{d^n \times n}$ has rank- n .*

Proof of Lemma G.6. We prove by induction. For $n = 2$, $v_1 v_1^\top, v_2 v_2^\top$ are linearly independent under the non-parallel assumption. By induction suppose $\{\text{vec}(\mathbf{v}_1^{\otimes(n-1)}), \dots, \text{vec}(\mathbf{v}_{n-1}^{\otimes(n-1)})\}$ are linearly independent. Suppose the conclusion does not hold, then there exists $\alpha_1, \dots, \alpha_n \in \mathbb{R}$ not identically 0, such that

$$\sum_{i=1}^n \alpha_i \text{vec}(\mathbf{v}_i^{\otimes n}) = 0,$$

which implies for $p = 1, \dots, d$

$$\sum_{i=1}^n (\alpha_i \mathbf{v}_{i,p}) \text{vec}(\mathbf{v}_i^{\otimes(n-1)}) = 0.$$

Note by induction hypothesis any size $(n-1)$ subset of

$\{\text{vec}(\mathbf{v}_1^{\otimes(n-1)}), \dots, \text{vec}(\mathbf{v}_n^{\otimes(n-1)})\}$ is linearly independent. This implies if $\alpha_i \mathbf{v}_{i,p} = 0$ for some $i \in [n]$ and $p \in [d]$, then we must have $\alpha_j \mathbf{v}_{j,p} = 0$ for all $j \in [n]$. Combining this observation with the assumption that every \mathbf{v}_i is non-zero, there must exist $p \in [d]$ such that $\mathbf{v}_{i,p} \neq 0$ for all $i \in [n]$. Without loss of generality, we assume $\mathbf{v}_{i,1} \neq 0$ for all $i \in [n]$.

Next, note if there exists $\alpha_i = 0$, then we have $\alpha_j = 0$ for all $j \in [n]$ because $\mathbf{v}_{j,p} \neq 0$ for all $j \in [n]$ and the linear independence induction hypothesis. Therefore from now on we assume $\alpha_i \neq 0$ for all $i \in [n]$.

For any $p \in [d]$, we have

$$\sum_{i=1}^n (\alpha_i \mathbf{v}_{i,p}) \text{vec}(\mathbf{v}_i^{\otimes(n-1)}) = 0 \text{ and } \sum_{i=1}^n (\alpha_i \mathbf{v}_{i,1}) \text{vec}(\mathbf{v}_i^{\otimes(n-1)}) = 0.$$

By multiplying the second equation by $\frac{\mathbf{v}_{1,p}}{\mathbf{v}_{1,1}}$ and subtracting,

$$\sum_{i=2}^n (\alpha_i \mathbf{v}_{i,p} - \alpha_i \frac{\mathbf{v}_{1,p}}{\mathbf{v}_{1,1}} \mathbf{v}_{i,1}) \text{vec} \left(\mathbf{v}_i^{\otimes (n-1)} \right) = 0.$$

Using the linear independence induction hypothesis, we know for $i = 2, \dots, n$:

$$\frac{\mathbf{v}_{i,p}}{\mathbf{v}_{1,1}} = \frac{\mathbf{v}_{1,p}}{\mathbf{v}_{1,1}}.$$

Therefore we know

$$\frac{\mathbf{v}_{1,p}}{\mathbf{v}_{1,1}} = \dots = \frac{\mathbf{v}_{n,p}}{\mathbf{v}_{n,1}}.$$

Thus there exists $c_2, \dots, c_d \in \mathbb{R}^d$ such that

$$\mathbf{v}_{i,p} = c_p \mathbf{v}_{i,1} \text{ for all } i \in [n].$$

Note this implies all $\mathbf{v}_i, i \in [n]$ are on the same line. This contradicts with the non-parallel assumption. □