

Generalization Guarantees of Gradient Descent for Shallow Neural Networks

Puyu Wang¹, Yunwen Lei², Di Wang³, Yiming Ying^{4*}, Ding-Xuan Zhou⁴

¹ Hong Kong Baptist University, Hong Kong

² The University of Hong Kong, Hong Kong

³ King Abdullah University of Science and Technology, Saudi Arabia

⁴ The University of Sydney, Australia

Keywords: Learning theory, generalization analysis, algorithmic stability, neural networks, gradient descent

Abstract

Recently, significant progress has been made in understanding the generalization of neural networks (NNs) trained by gradient descent (GD) using the algorithmic stability approach. However, most of the existing research has focused on one-hidden-layer NNs and has not addressed the impact of different network scaling. Here, network scaling corresponds to the normalization of the layers. In this paper, we greatly extend the previous work (Lei et al., 2022; Richards and Kuzborskij, 2021) by conducting a comprehensive stability and generalization analysis of GD for two-layer and three-layer NNs.

For two-layer NNs, our results are established under general network scaling, relaxing previous conditions. In the case of three-layer NNs, our technical contribution lies in demonstrating its nearly co-coercive property by utilizing a novel induction strategy that thoroughly explores the effects of over-parameterization. As a direct application of our general findings, we derive the excess risk rate of $\mathcal{O}(1/\sqrt{n})$ for GD in both two-layer and three-layer NNs. This sheds light on sufficient or necessary conditions for under-parameterized and over-parameterized NNs trained by GD to attain the desired risk rate of $\mathcal{O}(1/\sqrt{n})$. Moreover, we demonstrate that as the scaling factor increases or the network complexity decreases, less over-parameterization is required for GD to achieve the desired error rates. Additionally, under a low-noise condition, we obtain a fast risk rate of $\mathcal{O}(1/n)$ for GD in both two-layer and three-layer NNs.

1 Introduction

Deep neural networks (DNNs) trained by (stochastic) gradient descent (GD) have achieved great success in a wide spectrum of applications such as image recognition (Krizhevsky et al., 2017), speech recognition (Hinton et al., 2012), machine translation (Bahdanau et al., 2014), and reinforcement learning (Silver et al., 2016). In practical applications, most of the deployed DNNs are over-parameterized, i.e., the number of parameters is far larger than the size of the training data. In Zhang et al. (2021), it was empirically demonstrated that over-parameterized NNs trained with SGD can generalize well to the test data while achieving a small training error. This has triggered a surge of theoretical studies on unveiling this generalization mastery of DNNs.

In particular, norm-based generalization bounds are established using the uniform

convergence approach (Bartlett et al., 2017, 2021; Golowich et al., 2018; Long and Sedghi, 2019; Neyshabur et al., 2015, 2018). However, this approach does not take the optimization algorithm and the data distribution into account. Another line of work is to consider the structure of the data distribution and provide algorithm-dependent generalization bounds. In Brutzkus et al. (2018); Li and Liang (2018), it is shown that SGD for over-parameterized two-layer NNs can achieve small generalization error under certain assumptions on the structure of the data. Allen-Zhu et al. (2019a) studies the generalization of SGD for two-layer and three-layer NNs if there exists a true (unknown) NN with low error on the data distribution. The other important line of work is the neural tangent kernel (NTK)-type approach (Arora et al., 2019a; Cao and Gu, 2019; Chizat and Bach, 2018; Nitanda et al., 2019; Kuzborskij and Szepesvári, 2022; Nguyen and Mücke, 2023) which shows that the model trained by GD is well approximated by the tangent space near the initialization and the generalization analysis can be reduced to those of the convex case or kernel methods. However, most of them either require a very high over-parameterization or focus on special function classes.

Recently, the appealing work (Richards and Kuzborskij, 2021) provides an alternative approach in a kernel-free regime, using the concept of algorithmic stability (Bousquet and Elisseeff, 2002; Hardt et al., 2016; Kuzborskij and Lampert, 2018). Specifically, it uses the model-average stability (Lei and Ying, 2020a) to derive generalization bounds of GD for two-layer over-parameterized NNs. Lei et al. (2022) improves this result by deriving generalization bounds for both GD and SGD, and relaxing the over-parameterization requirement. Taheri and Thrampoulidis (2024) derives fast generalization bounds of GD with smooth activation for self-bounded losses under a separable

distribution. However, the above studies only focus on two-layer NNs and a very specific network scaling.

We study the stability and generalization of GD for both two-layer and three-layer NNs with generic network scaling factors. Our contributions are summarized as follows.

- We establish excess risk bounds for GD on both two-layer and three-layer NNs with general network scaling under relaxed over-parameterization conditions. As a direct application of our generalization results, we show that GD can achieve the excess risk rate $\mathcal{O}(1/\sqrt{n})$ when the network width m satisfies certain qualitative conditions related to the scaling parameter c , the size of training data n , and the network complexity measured by the norm of the minimizer of the population risk (refer to Section 3 for further discussions). Further, under a low-noise condition, our excess risk rate can be improved to $\mathcal{O}(1/n)$ for both two-layer and three-layer NNs.

- A crucial technical element in the stability analysis for NNs trained by GD is establishing the almost co-coercivity of the gradient operator. This property naturally holds true for two-layer NNs due to the empirical risks' monotonically decreasing nature which no longer remains valid for three-layer NNs. Our technical contribution lies in demonstrating that the nearly co-coercive property still holds valid throughout the trajectory of GD. To achieve this, we employ a novel induction strategy that fully explores the effects of over-parameterization (refer to Section 4 for further discussions). Furthermore, we are able to eliminate a critical assumption made in Lei et al. (2022) regarding an inequality associated with the population risk of GD's iterates (see Remark 2 for additional details).

- Our results characterize a quantitative condition in terms of network complexity and scaling factor under which GD for two-layer and three-layer NNs can achieve the excess risk rate $\mathcal{O}(1/\sqrt{n})$ in under-parameterization and over-parameterization regimes. Our results shed light on sufficient or necessary conditions for under-parameterized and over-parameterized NNs trained by GD to achieve the desired risk rate $\mathcal{O}(1/\sqrt{n})$. In addition, our results show that as the scaling factor increases or the network complexity decreases, less over-parameterization is needed for GD to achieve the desired rates for two-layer and three-layer NNs.

Organization of the paper: The remaining parts of the paper are organized as follows. Section 2 illustrates the formulation of the problem. We present the stability and generalization results of GD for two-layer and three-layer NNs in Section 3 and sketch the main idea of the proof in Section 4. Section 5 concludes the paper.

2 Problem Formulation

Let P be a probability measure defined on a sample space $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, where $\mathcal{X} \subseteq \mathbb{R}^d$ and $\mathcal{Y} \subseteq \mathbb{R}$. Let $S = \{z_i = (\mathbf{x}_i, y_i)\}_{i=1}^n$ be a training dataset drawn from P . One aims to build a prediction model $f_{\mathbf{W}} : \mathcal{X} \mapsto \mathbb{R}$ parameterized by \mathbf{W} in some parameter space \mathcal{W} based on S . The performance of $f_{\mathbf{W}}$ can be measured by the population risk defined as $L(f_{\mathbf{W}}) = \frac{1}{2} \iint_{\mathcal{X} \times \mathcal{Y}} (f_{\mathbf{W}}(\mathbf{x}) - y)^2 dP(\mathbf{x}, y)$. The corresponding empirical risk is defined as

$$L_S(f_{\mathbf{W}}) = \frac{1}{2n} \sum_{i=1}^n (f_{\mathbf{W}}(\mathbf{x}_i) - y_i)^2. \quad (1)$$

We denote by $\ell(\mathbf{W}; z) = \frac{1}{2}(f_{\mathbf{W}}(\mathbf{x}) - y)^2$ the loss function of \mathbf{W} on a data point $z = (\mathbf{x}, y)$. The best possible model is the regression function f^* defined as $f^*(\mathbf{x}) = \mathbb{E}[y|\mathbf{x}]$,

where $\mathbb{E}[\cdot|\mathbf{x}]$ is the conditional expectation given \mathbf{x} . In this paper, we consider the prediction model $f_{\mathbf{W}}$ with a neural network structure. In particular, we are interested in two-layer and three-layer fully-connected NNs.

Two-layer NNs. A two-layer NN of width $m > 0$ and scaling parameter $c \in [1/2, 1]$ takes the form

$$f_{\mathbf{W}}(\mathbf{x}) = \frac{1}{m^c} \sum_{k=1}^m a_k \sigma(\mathbf{w}_k \mathbf{x}),$$

where $\sigma : \mathbb{R} \mapsto \mathbb{R}$ is an activation function, $\mathbf{W} = [\mathbf{w}_1^\top, \dots, \mathbf{w}_m^\top]^\top \in \mathcal{W}$ with $\mathcal{W} = \mathbb{R}^{m \times d}$ is the weight matrix of the first layer, and a fixed $\mathbf{a} = [a_1, \dots, a_m]$ with $a_k \in \{-1, +1\}$ is the weight of the output layer. In the above formulation, $\mathbf{w}_k \in \mathbb{R}^{1 \times d}$ denotes the weight of the edge connecting the input to the k -th hidden node, and a_k denotes the weight of the edge connecting the k -th hidden node to the output node. The output of the network is scaled by a factor m^{-c} that is decreasing with the network width m . Two popular choices of scaling are Neural Tangent Kernel (NTK) (Allen-Zhu et al., 2019b; Arora et al., 2019a,b; Jacot et al., 2018; Du et al., 2018, 2019) with $c = 1/2$ and mean field (Chizat and Bach, 2018; Chizat et al., 2019; Mei et al., 2019, 2018) with $c = 1$. Richards and Rabbat (2021) studied two-layer NNs with $c \in [1/2, 1]$ and discussed the influence of the scaling by trading off it with the network complexity. We also focus on the scaling $c \in [1/2, 1]$ to ensure that meaningful generalization bounds can be obtained. We fix the output layer weight \mathbf{a} and only optimize the first layer weight \mathbf{W} in this setting.

Three-layer NNs. For a matrix \mathbf{W} , let $\mathbf{W}_{s:}$ and \mathbf{W}_{is} denote the s -th row and the (i, s) -th entry of \mathbf{W} . For an input $\mathbf{x} \in \mathbb{R}^d$, a three-layer fully-connected NN with width

$m > 0$ and scaling $c \in (1/2, 1]$ is formulated as follows

$$f_{\mathbf{W}}(\mathbf{x}) = \frac{1}{m^c} \sum_{i=1}^m a_i \sigma \left(\frac{1}{m^c} \sum_{s=1}^m \mathbf{W}_{is}^{(2)} \sigma(\mathbf{W}_s^{(1)} \mathbf{x}) \right),$$

where $\sigma : \mathbb{R} \mapsto \mathbb{R}$ is an activation function, $\mathbf{W} = [\mathbf{W}^{(1)}, \mathbf{W}^{(2)}] \in \mathcal{W} = \mathbb{R}^{m \times (d+m)}$ is the weight matrix of the neural network, and $\mathbf{a} = [a_1, \dots, a_m]$ with $a_i \in \{-1, +1\}$ is the fixed output layer weight. Here, $\mathbf{W}^{(1)} \in \mathbb{R}^{m \times d}$ and $\mathbf{W}^{(2)} \in \mathbb{R}^{m \times m}$ are the weights at the first and the second layer, respectively. For simplicity, we assume that the widths of each hidden layer are the same. Our findings can readily extend to scenarios where the widths and scaling factors vary across layers. We consider the setting of optimizing the weights in both the first and the second layers.

We consider the gradient descent to solve the minimization problem: $\min_{\mathbf{W}} L_S(f_{\mathbf{W}})$.

For simplicity, let $L(\mathbf{W}) = L(f_{\mathbf{W}})$ and $L_S(\mathbf{W}) = L_S(f_{\mathbf{W}})$.

Definition 1 (Gradient Descent). Let $\mathbf{W}_0 \in \mathcal{W}$ be an initialization point, and $\{\eta_t : t \in \mathbb{N}\}$ be a sequence of step sizes. Let ∇ denote the gradient operator. At iteration t , the update rule of GD is

$$\mathbf{W}_{t+1} = \mathbf{W}_t - \eta_t \nabla L_S(\mathbf{W}_t). \quad (2)$$

We say the gradient operator $\nabla L_S(\mathbf{W})$ is ϵ -almost co-coercivity with $\epsilon > 0$ and $\alpha \geq 0$ if $\langle \mathbf{W} - \mathbf{W}', \nabla L_S(\mathbf{W}) - \nabla L_S(\mathbf{W}') \rangle \geq \alpha \|\nabla L_S(\mathbf{W}) - \nabla L_S(\mathbf{W}')\|_2^2 - \epsilon \|\mathbf{W} - \mathbf{W}' - \eta_t(\nabla L_S(\mathbf{W}) - \nabla L_S(\mathbf{W}'))\|_2^2$.

Target of Analysis. Let $\mathbf{W}^* = \arg \min_{\mathbf{W} \in \mathcal{W}} L(\mathbf{W})$ where the minimizer is chosen to be the one enjoying the smallest norm. For a randomized algorithm \mathcal{A} to solve (1), let $\mathcal{A}(S)$ be the output of \mathcal{A} applied to the dataset S . The generalization performance of $\mathcal{A}(S)$ is measured by its *excess population risk*, i.e., $L(\mathcal{A}(S)) - L(\mathbf{W}^*)$. In this paper,

we are interested in studying the excess population risk of models trained by GD for both two-layer and three-layer NNs. For any $\lambda > 0$, let $\mathbf{W}_\lambda^* = \operatorname{argmin}_{\mathbf{W}} \{L(\mathbf{W}) + \frac{\lambda}{2} \|\mathbf{W} - \mathbf{W}_0\|_2^2\}$. We use the following error decomposition

$$\begin{aligned} \mathbb{E}[L(\mathcal{A}(S))] - L(\mathbf{W}^*) &= \mathbb{E}[L(\mathcal{A}(S)) - L_S(\mathcal{A}(S))] \\ &+ \mathbb{E}[L_S(\mathcal{A}(S)) - L_S(\mathbf{W}_\lambda^*) - \frac{\lambda}{2} \|\mathbf{W}_\lambda^* - \mathbf{W}_0\|_2^2] + [L(\mathbf{W}_\lambda^*) + \frac{\lambda}{2} \|\mathbf{W}_\lambda^* - \mathbf{W}_0\|_2^2 - L(\mathbf{W}^*)], \end{aligned} \quad (3)$$

where we use $\mathbb{E}[L_S(\mathbf{W}_\lambda^*)] = L(\mathbf{W}_\lambda^*)$. Here, we denote $\|\cdot\|_2$ the standard Euclidean norm of a vector or a matrix. The first term in (3) is called the *generalization error*, which can be controlled by stability analysis. The second term is the *optimization error*, which can be estimated by tools from optimization theory. The third term, denoted by $\Lambda_\lambda = L(\mathbf{W}_\lambda^*) + \frac{\lambda}{2} \|\mathbf{W}_\lambda^* - \mathbf{W}_0\|_2^2 - L(\mathbf{W}^*)$, is the *approximation error* which will be estimated by introducing an assumption on the complexity of the NN (Assumption 3 in Section 3.1).

We use the following on-average argument stability (Lei and Ying, 2020a) to study generalization error.

Definition 2 (On-average argument stability). Let $S = \{z_i\}_{i=1}^n$ and $S' = \{z'_i\}_{i=1}^n$ be drawn i.i.d. from an unknown distribution P . For any $i \in [n] =: \{1, \dots, n\}$, define $S^{(i)} = \{z_1, \dots, z_{i-1}, z'_i, z_{i+1}, \dots, z_n\}$ as the set formed from S by replacing the i -th element with z'_i . We say \mathcal{A} is on-average argument ϵ -stable if

$$\mathbb{E}_{S, S', \mathcal{A}} \left[\frac{1}{n} \sum_{i=1}^n \|\mathcal{A}(S) - \mathcal{A}(S^{(i)})\|_2^2 \right] \leq \epsilon^2.$$

Our analysis requires the following standard assumptions (Lei et al., 2022; Richards and Kuzborskij, 2021). Both sigmoid and hyperbolic tangent activation functions satisfy Assumption 1.

Assumption 1 (Activation function). *The activation function $a \mapsto \sigma(a)$ is continuous and twice differentiable with $|\sigma(a)| \leq B_\sigma$, $|\sigma'(a)| \leq B_{\sigma'}$ and $|\sigma''(a)| \leq B_{\sigma''}$, where $B_\sigma, B_{\sigma'}, B_{\sigma''} > 0$.*

Assumption 2 (Inputs, labels and loss). *There exist constants $c_x, c_y, c_0 > 0$ such that $\|\mathbf{x}\|_2 \leq c_x$, $|y| \leq c_y$ and $\max\{\ell(\mathbf{0}; z), \ell(\mathbf{W}_0; z)\} \leq c_0$ for any $\mathbf{x} \in \mathcal{X}$, $y \in \mathcal{Y}$, $z \in \mathcal{Z}$.*

Assumption 2 requires that the loss is uniformly bounded when evaluated under the initialization. One can verify that this requirement holds true for some widely used activations including sigmoid activation and hyperbolic tangent activation for any initialization.

3 Main Results

In this section, we present our main results on the excess population risk of NNs trained by GD.

3.1 Two-layer Neural Networks with Scaling Parameters

We say a function $\ell : \mathcal{W} \times \mathcal{Z} \rightarrow \mathbb{R}$ is ρ -smooth if for any $\mathbf{W}, \mathbf{W}' \in \mathcal{W}$ and $z \in \mathcal{Z}$, there holds $\ell(\mathbf{W}; z) - \ell(\mathbf{W}'; z) \leq \langle \nabla \ell(\mathbf{W}'; z), \mathbf{W} - \mathbf{W}' \rangle + \frac{\rho}{2} \|\mathbf{W} - \mathbf{W}'\|_2^2$. We denote $B \asymp B'$ if there exist some universal constants $c_1, c_2 > 0$ such that $c_1 B \geq B' \geq c_2 B$, and denote $B \gtrsim B'$ ($B \lesssim B'$) if there exists a constant $c > 0$ such that $B \geq cB'$ ($cB \leq B'$). Let $\rho = c_x^2 \left(\frac{B_{\sigma'}^2 + B_\sigma B_{\sigma''}}{m^{2c-1}} + \frac{B_{\sigma''} c_y}{m^c} \right)$. The following theorem to be proved in Appendix A.1 presents generalization error bounds of GD for two-layer NNs.

Theorem 1 (Generalization error). *Suppose Assumptions 1 and 2 hold. Let $\{\mathbf{W}_t\}$ be*

produced by (2) with $\eta_t \equiv \eta \leq 1/(2\rho)$ based on S . Assume

$$m \gtrsim ((\eta T)^2(1 + \eta\rho)\sqrt{\rho(\rho\eta T + 2)}/n)^{\frac{2}{4c-1}} + (\eta T)^{\frac{3}{4c-1}} + (\eta T)^{\frac{1}{c}}. \quad (4)$$

Then, for $t \in [T]$, there holds

$$\mathbb{E}[L(\mathbf{W}_t) - L_S(\mathbf{W}_t)] \leq \left(\frac{4(e\eta\rho)^2 t}{n^2} + \frac{4e\eta\rho}{n}\right) \sum_{j=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_j)].$$

Remark 1. Theorem 1 establishes the first generalization results for general $c \in [1/2, 1]$, which shows that the requirement on m becomes weaker as c becomes larger. Indeed, with a typical choice $\eta T \asymp n^{\frac{c}{6\mu+2c-3}}$ (as shown in Corollary 4 below), the assumption (4) becomes $m \gtrsim n^{\frac{1}{6\mu+2c-3}(1+\frac{51-96\mu}{8(8c-3)})}$. This requirement becomes milder as c increases for any $\mu \in [0, 1]$, which implies that large scaling reduces the requirement on the network width. In particular, our assumption only requires $m \gtrsim \eta T$ when $c = 1$. Lei et al. (2022) provided a similar generalization bound with an assumption $m \gtrsim (\eta T)^5/n^2 + (\eta T)^2$, and our result with $c = 1/2$ is consistent with their result.

The following theorem to be proved in Appendix A.2 develops the optimization error bounds of GD.

Theorem 2 (Optimization error). *Suppose Assumptions 1 and 2 hold. Let $\{\mathbf{W}_t\}$ be produced by (2) with $\eta_t \equiv \eta \leq 1/(2\rho)$ based on S . Let $\tilde{\mathbf{b}} = c_{\mathbf{x}}^2 B_{\sigma''}(\frac{2B_{\sigma'}c_{\mathbf{x}}}{m^{c-1/2}} + \sqrt{2c_0})$ and assume (4) holds. Assume*

$$m \gtrsim \left(T(\sqrt{\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)\left(\frac{\eta^3 \rho^2 T^2}{n^2} + \frac{\eta^2 T \rho}{n} + 1\right)\right)^{\frac{1}{c}}. \quad (5)$$

Then we have

$$\begin{aligned}
& \mathbb{E}[L_S(\mathbf{W}_T)] - L_S(\mathbf{W}_{\frac{1}{\eta T}}^*) - \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 \\
& \leq \frac{2\tilde{b}}{m^c} (\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) \left(\frac{8e^2\eta^3\rho^2T^2}{n^2} + \frac{8e\eta^2T\rho}{n} \right) \sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] \\
& \quad + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + \eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)].
\end{aligned}$$

We can combine generalization and optimization error bounds together to derive our main result for two-layer NNs. It is worth mentioning that our excess population bound is dimension-independent, which is mainly due to that stability analysis focuses on the optimization process trajectory, as opposed to the uniform convergence approach which involves the complexity of function space and thus is often dimension-dependent. In addition, the following theorem holds for both random initialization and fixed initialization settings. Without loss of generality, we assume $\eta T \geq 1$.

Theorem 3 (Excess population risk). *Suppose Assumptions 1 and 2 hold. Let $\{\mathbf{W}_t\}$ be produced by (2) with $\eta_t \equiv \eta \leq 1/(2\rho)$. Assume (4) and (5) hold. For any $c \in [1/2, 1]$, if $\eta T m^{1-2c} = \mathcal{O}(n)$ and $m \gtrsim (\eta T (\sqrt{\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2))^{1/c}$, then*

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left(\frac{\eta T m^{1-2c}}{n} L(\mathbf{W}^*) + \Lambda_{\frac{1}{\eta T}}\right).$$

Remark 2. Theorem 3 provides the first excess risk bounds of GD for two-layer NNs with general scaling parameter $c \in [1/2, 1]$ which recovers the previous work (Richards and Kuzborskij, 2021; Lei et al., 2022) with $c = 1/2$. Specifically, Richards and Kuzborskij (2021) derived the excess risk bound $\mathcal{O}\left(\frac{\eta T}{n} L(\mathbf{W}^*) + \Lambda_{\frac{1}{\eta T}}\right)$ with $m \gtrsim (\eta T)^5$, we relax this condition to $m \gtrsim (\eta T)^3$ by providing a better estimation of the smallest eigenvalue of a Hessian matrix of the empirical risk. This bound was obtained in Lei

et al. (2022) under a crucial condition $\mathbb{E}[L(\mathbf{W}_s)] \geq L(\mathbf{W}_{\frac{1}{\eta T}}^*)$ for any $s \in [T]$, which is difficult to verify in practice. Here, we are able to remove this condition by using $L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}_s) \leq L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*) \leq \Lambda_{\frac{1}{\eta T}}$ since $L(\mathbf{W}_s) \geq L(\mathbf{W}^*)$ when controlling the bound of GD iterates (see Lemma 15 in Appendix A for details). Further, if we ignore the effect of $\mathbf{W}_{\frac{1}{\eta T}}^*$ and $L(\mathbf{W}^*)$, Theorem 3 implies the larger the scaling parameter c is, the better the excess risk bound is. The reason could be that the smoothness parameter related to the objective function along the GD trajectory becomes smaller for a larger c (see Lemma 11 in Appendix A for details).

As a direct application of Theorem 3, we can derive the excess risk rates of GD for two-layer NNs by properly trade-offing the generalization, optimization and approximation errors. To this end, we introduce the following assumption to control the approximation error which can interpret the complexity of the network, i.e., how well the network approximates the least population risk. To achieve guarantees on the excess population risk, the approximation error should be sufficiently small, which can be interpreted as a complexity assumption of the learning problem.

Assumption 3 (Richards and Rabbat (2021)). *There is $\mu \in [0, 1]$ and population risk minimizer \mathbf{W}^* with $\|\mathbf{W}^*\|_2 \leq m^{\frac{1}{2}-\mu}$.*

Indeed, the approximation error $L(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 - L(\mathbf{W}^*) \leq L(\mathbf{W}^*) + \frac{1}{2\eta T} \|\mathbf{W}^* - \mathbf{W}_0\|_2^2 - L(\mathbf{W}^*) = \frac{1}{2\eta T} \|\mathbf{W}^* - \mathbf{W}_0\|_2^2$. It implies that controlling the approximation error requires placing assumptions on the weights $\|\mathbf{W}^*\|_2$. Hence, the norm $\|\mathbf{W}^*\|_2$ serves as an indicator of the network's complexity. Assumption 3 states that the population risk minimizer \mathbf{W}^* satisfies $\|\mathbf{W}^*\|_2 \leq m^{\frac{1}{2}-\mu}$. Therefore, a large μ aligns with a stronger assumption which corresponds to a simpler network, as the minimum

risk can be attained by a set of weights with a smaller norm.

Corollary 4. *Let assumptions in Theorem 3 hold and Assumption 3 hold. Suppose*

$$\frac{c}{3} + \mu > \frac{1}{2}.$$

(a) *If $c \in [1/2, 3/4)$ and $c + \mu \geq 1$, we can choose $m \asymp (\eta T)^{\frac{3}{2c}}$ and ηT such that*

$$n^{\frac{c}{6\mu+2c-3}} \lesssim \eta T \lesssim n^{\frac{c}{3-4c}}. \text{ If } c \in [3/4, 1], \text{ we can choose } m \asymp (\eta T)^{\frac{3}{2c}} \text{ and } \eta T \text{ such}$$

that $\eta T \gtrsim n^{\frac{c}{6\mu+2c-3}}$. Then there holds

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}(1/\sqrt{n}).$$

(b) *Assume $L(\mathbf{W}^*) = 0$. We can choose $m \asymp (\eta T)^{\frac{3}{2c}}$ and ηT such that $\eta T \gtrsim n^{\frac{2c}{6\mu+2c-3}}$,*

and get

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}(1/n).$$

Corollary 4 shows that for a fixed μ , the lower bound of m becomes smaller as c increases. It means that large c relaxes the requirement to the width m for GD to achieve the excess risk rate $\mathcal{O}(1/\sqrt{n})$. The difference made over $c = 3/4$ in part (a) is proof artifacts. More details are given in Appendix A.3.

Interpretation of the Results: Corollary 4 indicates a quantitative condition via the network complexity and scaling factor where GD for two-layer NNs can achieve the excess risk rate $\mathcal{O}(1/\sqrt{n})$ in under-parameterization and over-parameterization regimes. Here, the under-parameterization regime and over-parameterization regime correspond to the region that $m \lesssim n$ and $m \gtrsim n$, respectively.

Let us first explain the meanings of the regions and lines in Figure 1. Specifically, the *blue regions* (with or without dots) correspond to the conditions $\frac{c}{3} + \mu > \frac{1}{2}$ and $c + \mu \geq 1$ in part (a) of Corollary 4, while our results do not hold in the *pink region*




Figure 1: Scaling parameter c versus Network complexity parameter μ for Part (a) in Corollary 4. *Blue Region without dots*: values of c and μ where over-parameterization is necessary to achieve error bound $\mathcal{O}(1/\sqrt{n})$. *Blue Region with dots*: values of c and μ where under-parameterization is sufficient to achieve error bound $\mathcal{O}(1/\sqrt{n})$. *Pink Region*: the desired bound cannot be guaranteed.

which violates the conditions in part (a). Furthermore, under conditions $\frac{c}{3} + \mu > \frac{1}{2}$ and $c + \mu \geq 1$, that the desired bound in part (a) can be achieved by choosing $m \asymp (\eta T)^{\frac{3}{2c}}$ for any ηT satisfying $\eta T \gtrsim n^{\frac{c}{6\mu+2c-3}}$ if $c \geq 3/4$ and $n^{\frac{c}{3-4c}} \gtrsim \eta T \gtrsim n^{\frac{c}{6\mu+2c-3}}$ if $c \in [1/2, 3/4)$, which further implies that GD with any width $m \gtrsim n^{\frac{3}{2(6\mu+2c-3)}}$ if $c \in [3/4, 1]$ and $n^{\frac{3}{2(3-4c)}} \gtrsim m \gtrsim n^{\frac{3}{2(6\mu+2c-3)}}$ if $c \in [1/2, 3/4)$ with suitable iterations can achieve the error rate $\mathcal{O}(1/\sqrt{n})$. This observation tells us that **the smallest width** for guaranteeing our results in part (a) is $m \asymp n^{\frac{3}{2(6\mu+2c-3)}}$ for any $c \in [1/2, 1]$. The dotted line $\mu = 3/4 - c/3$ in the figure corresponds to the setting $m \asymp n$, i.e., $\frac{3}{2(6\mu+2c-3)} = 1$. Correspondingly, we use the *dotted blue region* above the dotted line to indicate the under-parameterization region and *the blue region without dots* below the dotted line

for the over-parameterization region. With the above explanations, we can interpret Figure 1 as follows.

- Firstly, from the figure, we know that, if values of c and μ are located above the dotted line $\mu \geq 3/4 - c/3$, i.e., the blue region with dots, under-parameterization is *sufficient* for GD to achieve the error rate $\mathcal{O}(1/\sqrt{n})$. It implies that the sufficient condition for under-parameterized NNs trained by GD achieving the desired rate is $\mu \geq 3/4 - c/3$. The potential reason is that the population risk minimizer \mathbf{W}^* is well-behaved in terms of its norm being relatively small with μ being relatively large there. In particular, when $\mu > 1/2$, $\|\mathbf{W}^*\|_2 \leq m^{1/2-\mu}$ tends to 0 as m tends to infinity. Hence, it is expected that under-parameterized NNs can learn this relatively simple \mathbf{W}^* well. However, it is worthy of mentioning that over-parameterization can also achieve the rate $\mathcal{O}(1/\sqrt{n})$ since the dotted line only indicates the smallest width required for achieving such an error rate.

- Secondly, from the figure, we see that, if c and μ belong to the blue region without dots which is between the solid lines and the dotted line, then over-parameterization is *necessary* for achieving the error rate $\mathcal{O}(1/\sqrt{n})$. This is because, in the blue region without dots, that the conditions of choosing $m \gtrsim n^{\frac{3}{2(6\mu+2c-3)}}$ in part (a) of Corollary 4 which will always indicate the over-parameterization region, i.e., $m \gtrsim n$. Furthermore, from the above discussions, our theoretical results indicate that the over-parameterization does bring benefit for GD to achieve good generalization in the sense that GD can achieve excess risk rate $\mathcal{O}(1/\sqrt{n})$ when c and μ is in the whole blue region (with or without dots) while under-parameterization can only do so for the blue region with dots where the network complexity is relatively simple, i.e., μ is relatively large.

- Thirdly, our results do not hold for GD when values of c and μ are in the pink region in the figure. In particular, when $\mu < 1/6$, our bounds do not hold for any $c \in [1/2, 1]$. We suspect that this is due to the artifacts of our analysis tools and it remains an open question to us whether we can get a generalization error bound $\mathcal{O}(1/\sqrt{n})$ when $\mu < 1/6$. In addition, our results in Corollary 4 also indicate that the requirement on m becomes weaker as c and μ become larger. It implies that networks with larger scaling and simpler network complexity are biased to weaken the over-parameterization for GD to achieve the desired error rates for two-layer NNs.

Remark 3. Note Lemma 11 in Appendix A shows that $f_{\mathbf{W}}$ is $B_{\sigma'} c_{\mathbf{x}} m^{\frac{1}{2}-c}$ -Lipschitz. Combining this result with Assumption 3 we know $|f_{\mathbf{W}^*}(\mathbf{x}) - f_0(\mathbf{x})|^2 = \mathcal{O}(m^{2(1-c-\mu)})$. In order for $|f_{\mathbf{W}^*}(\mathbf{x}) - f_0(\mathbf{x})|^2$ to not vanish as m tends to infinity, one needs $c + \mu \leq 1$. In Corollary 4, we also need $\frac{c}{3} + \mu > \frac{1}{2}$ to ensure the excess risk bounds vanish. Combining these two conditions together implies that c can not be larger than $3/4$. That is, for the range $c \in (3/4, 1]$, the conditions in Corollary 4 restrict the class of functions the networks can represent as m tends to infinity. However, we want to emphasize that even for the simplest case that $|f_{\mathbf{W}^*}(\mathbf{x}) - f_0(\mathbf{x})|^2$ tends to 0 as m tends to infinity, our results still imply that over-parameterization does bring benefit for GD to achieve optimal excess risk rate $\mathcal{O}(1/\sqrt{n})$. Besides, our corollary mainly discusses the conditions for achieving the excess risk rate $\mathcal{O}(1/\sqrt{n})$ and $\mathcal{O}(1/n)$. The above-mentioned conditions will be milder if we consider the slower excess risk rates. Then the restriction on c will be weaker. Furthermore, our main result (i.e., Theorem 3) does not rely on Assumption 3, and it holds for any setting.

Comparison with the Existing Work: Part (b) in Corollary 4 shows fast rate $\mathcal{O}(1/n)$

can be derived under a low-noise condition $L(\mathbf{W}^*) = 0$ which is equivalent to the fact that there is a true network such that $y = f_{\mathbf{W}^*}(\mathbf{x})$ almost surely. Similar to part (a), large c and large μ also help weaken the requirement on m in this case. For a special case $\mu = 1/2$ and $c = 1/2$, Lei et al. (2022) proved that GD for two-layer NNs achieves the excess risk rate $\mathcal{O}(1/\sqrt{n})$ with $m \asymp n^{3/2}$ and $\eta T \asymp \sqrt{n}$ in the general case, which is further improved to $\mathcal{O}(1/n)$ with $m \asymp n^3$ and $\eta T \asymp n$ in a low-noise case. Corollary 4 recovers their results with the same conditions on m and ηT for this setting.

Richards and Rabbat (2021) studied GD with weakly convex losses and showed that the excess population risk is controlled by $\mathcal{O}\left(\frac{\eta T L^2}{n} + \frac{\|\mathbf{W}_\epsilon - \mathbf{W}_0\|_2^2}{\eta T} + \epsilon(\eta T + \|\mathbf{W}^* - \mathbf{W}_0\|_2)\right)$ if $2\eta\epsilon < 1/T$ when the empirical risk is ϵ -weakly convex and L -Lipschitz continuous, where $\mathbf{W}_\epsilon = \arg \min_{\mathbf{W}} L_S(\mathbf{W}) + \epsilon\|\mathbf{W} - \mathbf{W}_0\|_2^2$. If the approximation error is small enough, then the $\mathcal{O}(1/\sqrt{n})$ bound can be achieved by choosing $\eta T = \sqrt{n}$ if $\|\mathbf{W}_\epsilon\|_2 = \mathcal{O}(1)$. Indeed, their excess risk bound will not converge for the general case. Specifically, note that $L_S(\mathbf{W}_\epsilon) + \epsilon\|\mathbf{W}_\epsilon - \mathbf{W}_0\|_2^2 \leq L_S(0) + \epsilon\|\mathbf{W}_0\|_2^2$, then there holds $\|\mathbf{W}_\epsilon - \mathbf{W}_0\|_2^2 = \mathcal{O}(1/\epsilon)$. The simultaneous appearance of $\frac{1}{\eta T \epsilon}$ and $\eta T \epsilon$ causes the non-vanishing error bound. Richards and Rabbat (2021) also investigated the weak convexity of two-layer NNs with a smooth activation function. Under the assumption that the derivative of the loss function is uniformly bounded by a constant, they proved that the weak convexity parameter is controlled by $\mathcal{O}(d/m^c)$. We provide a dimension-independent weak convexity parameter which further yields a dimension-independent excess risk rate $\mathcal{O}(1/\sqrt{n})$. More discussion can be found in Appendix A.3.

3.2 Three-layer Neural Networks with Scaling Parameters

Now, we present our results for three-layer NNs. Let $\hat{\rho} = 4B_2(1 + 2B_1)$ and $\mathcal{B}_T = \sqrt{\eta T} + \|\mathbf{W}_0\|_2$, where $B_1, B_2 > 0$ are constants depending on $c_x, B_{\sigma'}, B_{\sigma''}$ and c_0 , whose specific forms are given in Appendix B.1. We first present the generalization bounds for three-layer NNs.

Theorem 5 (Generalization error). *Suppose Assumptions 1 and 2 hold. Let $\{\mathbf{W}_t\}$ be produced by (2) with $\eta_t \equiv \eta \leq 1/(8\hat{\rho})$ based on S . Assume*

$$m \gtrsim (\eta T)^4 + (\eta T)^{\frac{1}{4c-2}} + \|\mathbf{W}_0\|_2^{\frac{4}{8c-3}} + \|\mathbf{W}_0\|_2^{\frac{1}{6c-3}} + ((\eta T \mathcal{B}_T)^2 + \frac{(\eta T)^{\frac{7}{2}} \mathcal{B}_T}{n})^{\frac{1}{5c-\frac{1}{2}}} + ((\eta T)^{\frac{3}{2}} \mathcal{B}_T^2 + \frac{(\eta T)^3 \mathcal{B}_T}{n})^{\frac{1}{4c-1}} + ((\eta T)^2 \mathcal{B}_T + \frac{(\eta T)^{\frac{7}{2}}}{n})^{\frac{1}{5c-1}} + ((\eta T)^{\frac{3}{2}} \mathcal{B}_T + \frac{(\eta T)^3}{n})^{\frac{1}{4c-\frac{3}{2}}}. \quad (6)$$

Then, for any $t \in [T]$,

$$\mathbb{E}[L(\mathbf{W}_t) - L_S(\mathbf{W}_t)] \leq \left(\frac{4e^2 \eta^2 \hat{\rho}^2 t}{n^2} + \frac{4e\eta \hat{\rho}}{n} \right) \sum_{j=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_j)].$$

Similar to Theorem 1, Theorem 5 also implies that a larger scaling c relaxes the requirement on m .

Remark 4. As compared to two-layer NNs, the analysis of three-layer NNs is more challenging since we can only show that $\lambda_{\max}(\nabla^2 \ell(\mathbf{W}; z)) \leq \rho_{\mathbf{W}}$ where $\rho_{\mathbf{W}}$ depends on $\|\mathbf{W}\|_2$, i.e., the smoothness parameter of ℓ relies on the upper bound of \mathbf{W} , while that of two-layer NNs is uniformly bounded. In this way, three-layer NNs do not enjoy almost co-coercivity, which is the key step to controlling the stability of GD. To handle this problem, we first establish a crude estimate $\|\mathbf{W}_t - \mathbf{W}_0\|_2 \leq \eta t m^{2c-1}$ for any $c > 1/2$ by induction strategy. Once having the estimate of $\|\mathbf{W}_t\|_2$, we can upper bound $\rho_{\mathbf{W}}$ by a constant $\hat{\rho}$ if m satisfies (6). It implies that the loss is $\hat{\rho}$ -smooth along the trajectory of

GD. Finally, by assuming $\eta \leq 1/(2\hat{\rho})$ we build the upper bound $\|\mathbf{W}_t - \mathbf{W}_0\|_2 \leq \sqrt{2c_0\eta t}$. However, for the case $c = 1/2$, we cannot get a similar bound due to the condition $m^{2-4c}\|\mathbf{W}_t - \mathbf{W}_0\|_2^2 = \mathcal{O}(m^{2c-1})$. Specifically, the upper bound of $\|\mathbf{W}_t - \mathbf{W}_0\|_2$ in this case contains a worse term 2^t , which is not easy to control. Therefore, we only consider $c \in (1/2, 1]$ for three-layer NNs. The estimate of $\|\mathbf{W}_t - \mathbf{W}_0\|_2$ when $c = 1/2$ remains an open problem. The detailed proof of the theorem is given in Appendix B.1.

Let $\mathcal{C}_T^n = \eta T + \eta^3 T^2/n^2$, $\hat{\mathcal{C}}_{\mathbf{W}} = 4B_1(m^{-3c}(\|\mathbf{W}\|_2^2 + 2c_0\eta T) + m^{\frac{1}{2}-2c}(\|\mathbf{W}\|_2 + \sqrt{2c_0\eta T}))$ and $\hat{B}_{\mathbf{W}} = (\frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-1/2}}(\sqrt{2c_0\eta T} + \|\mathbf{W}\|_2) + \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}})(2\sqrt{2\eta T c_0} + \|\mathbf{W}\|_2) + \sqrt{2c_0}$. The following theorem gives optimization error bounds for three-layer NNs. The proof is given in Appendix B.2.

Theorem 6 (Optimization error). *Let Assumptions 1, 2 hold. Let $\{\mathbf{W}_t\}$ be produced by (2) with $\eta_t \equiv \eta \leq 1/(8\hat{\rho})$. Assume m is large enough such that (6) and*

$$m \gtrsim (\mathcal{C}_T^n(\mathcal{B}_T + \|\mathbf{W}_{\frac{1}{\eta T}}^*\|_2)^4)^{\frac{1}{5c-\frac{1}{2}}} + (\mathcal{C}_T^n(\mathcal{B}_T + \|\mathbf{W}_{\frac{1}{\eta T}}^*\|_2)^3)^{\frac{1}{4c-1}} + (\mathcal{C}_T^n(\mathcal{B}_T + \|\mathbf{W}_{\frac{1}{\eta T}}^*\|_2)^2)^{\frac{1}{4c-\frac{3}{2}}} + (\mathcal{C}_T^n(\mathcal{B}_T + \|\mathbf{W}_{\frac{1}{\eta T}}^*\|_2))^{\frac{1}{2c-\frac{1}{2}}}. \quad (7)$$

Then we have

$$\begin{aligned} & \mathbb{E}[L_S(\mathbf{W}_T)] - L_S(\mathbf{W}_{\frac{1}{\eta T}}^*) - \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 \\ & \leq \hat{\mathcal{C}}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \hat{B}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \left(\left(\frac{4e^2\eta^3 T^2 \hat{\rho}^2}{n^2} + \frac{4e\eta^2 T \hat{\rho}}{n} \right) \sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 \right. \\ & \quad \left. + \eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)] \right). \end{aligned}$$

Now, we develop excess risk bounds of GD for three-layer NNs by combining Theorem 5 and Theorem 6 together. The proof is given in Appendix B.3.

Theorem 7 (Excess population risk). *Suppose Assumptions 1 and 2 hold. Let $\{\mathbf{W}_t\}$ be produced by (2) with $\eta \leq 1/(8\hat{\rho})$. Assume (6) and (7) hold. For any $c \in (1/2, 1]$, if $n \gtrsim \eta T$, then there holds*

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left(\frac{\eta T}{n} L(\mathbf{W}^*) + \Lambda_{\frac{1}{\eta T}}\right).$$

Finally, we establish excess risk bounds of GD for three-layer NNs by assuming Assumption 3 holds.

Corollary 8. *Let assumptions in Theorem 7 and Assumption 3 hold. The following statements hold.*

(a) *Assume $\mu \geq 1/2$. If $c \in [9/16, 1]$, we can choose $m \asymp (\eta T)^4$ and ηT such that $n^{\frac{1}{2(8\mu-3)}} \lesssim \eta T \lesssim \sqrt{n}$. If $c \in (1/2, 9/16)$, we can choose $m \asymp (\eta T)^{\frac{1}{4c-2}}$ and ηT such that $n^{\frac{2c-1}{2\mu+4c-3}} \lesssim \eta T \lesssim \sqrt{n}$. Then*

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}(1/\sqrt{n}).$$

(b) *Assume $L(\mathbf{W}^*) = 0$ and $\mu \geq 1/2$. If $c \in [9/16, 1]$, we can choose $m \asymp (\eta T)^4$ and $\eta T \gtrsim n^{\frac{1}{8\mu-3}}$. If $c \in (1/2, 9/16)$, we can choose $m \asymp (\eta T)^{\frac{1}{4c-2}}$ and ηT such that $\eta T \gtrsim n^{\frac{4c-2}{4c+2\mu-3}}$. Then*

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}(1/n).$$

Discussion of the Results: Part (a) in Corollary 8 shows GD for three-layer NNs can achieve the excess risk rate $\mathcal{O}(1/\sqrt{n})$ with $m \asymp (\eta T)^4$ and $n^{\frac{1}{2(8\mu-3)}} \lesssim \eta T \lesssim \sqrt{n}$ for the case $c \in [9/16, 1]$ and $m \asymp (\eta T)^{\frac{1}{4c-2}}$ and $n^{\frac{2c-1}{2\mu+4c-3}} \lesssim \eta T \lesssim \sqrt{n}$ for the case $c \in (1/2, 9/16)$, respectively. Note that there is an additional assumption $\mu \geq 1/2$ in part (a). Combining this assumption with $\|\mathbf{W}^*\|_2 \leq m^{\frac{1}{2}-\mu}$ together, we know

that the population risk minimizer \mathbf{W}^* cannot be too large to reach the power of the exponent of m . The potential reason is that we use a constant to bound the smoothness parameter in the analysis for three-layer NNs. Part (a) also indicates a quantitative condition in terms of m and c where GD for three-layer NNs can achieve the excess risk rate $\mathcal{O}(1/\sqrt{n})$ in under-parameterization and over-parameterization regimes, which is interpreted in Figure 2. The results in part (a) tell us that the smallest width for guaranteeing the desired bounds are $m \asymp n^{\frac{2}{8\mu-3}}$ for $c \in [9/16, 1]$ and $m \asymp n^{\frac{1}{2(2\mu+4c-3)}}$ for $c \in (1/2, 9/16)$. Similar to Figure 1, the dotted lines $\mu = \frac{5}{8}$ and $\mu = \frac{7}{4} - 2c$ in Figure 2 correspond to the setting $m \asymp n$, i.e., $\frac{2}{8\mu-3} = 1$ and $\frac{1}{2(2\mu+4c-3)} = 1$. Hence, when c and μ belong to the blue region with dots, under-parameterization is *sufficient* to achieve the desired rate. When c and μ are located in the blue region without dots which is between the solid lines and the dotted line, over-parameterization is *necessary* for GD to achieve the rate $\mathcal{O}(1/\sqrt{n})$. Our results for three-layer NNs also imply that the over-parameterization does bring benefit for GD to achieve good generalization in the sense that GD can achieve excess risk rate $\mathcal{O}(1/\sqrt{n})$. Under a low-noise condition, part (b) implies that the excess risk rate can be improved to $\mathcal{O}(1/n)$ with suitable choices of m and ηT . These results imply that the larger the scaling parameter is, the less over-parameterization is needed for GD to achieve the desired error rate for both the general case and the low-noise case.

Comparison with the Existing Work: Richards and Rabbat (2021) studied the minimal eigenvalue of the empirical risk Hessian for a three-layer NN with a linear activation in the first layer for Lipschitz and convex losses (e.g., logistic loss), while we focus on NNs with more general activation functions for least square loss. See more detailed




Figure 2: Scaling parameter c versus Network complexity parameter μ for Part (a) in Corollary 8. *Blue Region without dots*: values of c and μ where over-parameterization is necessary to achieve error bound $\mathcal{O}(1/\sqrt{n})$. *Blue Region with dots*: values of c and μ where under-parameterization is sufficient to achieve error bound $\mathcal{O}(1/\sqrt{n})$. *Pink Region*: the desired bound cannot be guaranteed.

discussion in Appendix B.4. Ju et al. (2022) studied the generalization performance of overparameterized three-layer NTK models with absolute loss and ReLU activation. They showed that the generalization error is in the order of $\mathcal{O}(1/\sqrt{n})$ when there are infinitely many neurons. They only trained the middle-layer weights of the networks. To the best of our knowledge, our work is the first study on the stability and generalization of GD to train both the first and the second layers of the network in the kernel-free regime.

Remark 5. Extending our results to multi-layer NNs would be interesting but challenging. The key difficulty lies in estimating the maximum and minimum eigenvalues of a Hessian matrix of the empirical risk, which rely on the norm of the coupled weights of

different layers. Once having the estimates of eigenvalues, we can derive the stability bounds and further the generalization bounds by establishing the almost co-coercivity of the gradient operator. However, it's not easy to control the coupled weights of different layers by using the existing technique. We acknowledge the importance of considering the applicability of our findings to networks with multiple layers. Since the existing stability analysis mainly focused on two-layer NNs, we regard our work as an important starting point to study multiple layers for future research directions.

4 Main Idea of the Proof

In this section, we present the main ideas for proving the main results in Section 3.

4.1 Two-layer Neural Networks

Eq.(3) decomposes the excess population risk into three terms: generalization error, optimization error and approximation error. We estimate these three terms separately.

Generalization error. From Lemma 10 we know the generalization error can be upper bounded by the on-average argument stability of GD. Hence, it remains to figure out the on-average argument stability of GD. A key step in our stability analysis is to show that the loss is strongly smooth and weakly convex, which can be obtained by the following results given in Lemma 11: $\lambda_{\max}(\nabla^2\ell(\mathbf{W}; z)) \leq \rho$ and

$$\lambda_{\min}(\nabla^2\ell(\mathbf{W}; z)) \geq - (c_{\mathbf{x}}^3 B_{\sigma'} B_{\sigma''} m^{\frac{1}{2}-2c} \|\mathbf{W} - \mathbf{W}_0\|_2 + c_{\mathbf{x}}^2 B_{\sigma} \sqrt{2c_0} m^{-c}).$$

Here $\rho = \mathcal{O}(m^{1-2c})$, $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$ denote the largest and the smallest eigenvalue of A , respectively. The lower bound of $\lambda_{\min}(\nabla^2\ell(\mathbf{W}; z))$ is related to $\|\mathbf{W} - \mathbf{W}_0\|_2$. We further show it is uniformly bounded (see the proof of Theorem 14). Then

the smoothness and weak convexity of the loss scale with m^{1-2c} and m^c , which implies that the loss becomes more smooth and more convex for wider networks with a larger scaling.

Based on the above results, we derive the following uniform stability bounds $\|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2 \leq 2\eta eT \sqrt{2c_0\rho(\rho\eta T + 2)}/n$, where $\mathbf{W}_t^{(i)}$ is the iterates derived by applying SGD to $S^{(i)}$. Here, the weak convexity of the loss plays an important role in presenting the almost co-coercivity of the gradient operator, which helps us establish the recursive relationship of $\|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2$. Note $\rho = \mathcal{O}(m^{1-2c}) = \mathcal{O}(1)$. Then we know $\|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2 = \mathcal{O}(\frac{m^{1-2c}}{n}(\eta T)^{3/2})$, which implies that GD is more stable for a wider neural network with a large scaling.

Optimization error. A key step is to use the smoothness and weak convexity to control

$$\begin{aligned} \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_t\|_2^2] &\leq \left(\frac{8e^2\eta^3\rho^2t^2}{n^2} + \frac{8e\eta^2t\rho}{n}\right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] + 2\mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2] \\ &\quad + 2\eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)] \end{aligned}$$

in Lemma 11. We use $L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}_s) \leq L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*) \leq \Lambda_{\frac{1}{\eta T}}$ to remove the condition $\mathbb{E}[L(\mathbf{W}_s)] \geq L(\mathbf{W}_{\frac{1}{\eta T}}^*)$ for any $s \in [T]$ in Lei et al. (2022). Then the optimization error can be controlled by the monotonically decreasing of $\{L_S(\mathbf{W}_t)\}$.

The proofs are given in Appendix A.2.

Excess population risk. Combining stability bounds, optimization bounds and approximation error together, and noting that $L(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 \leq L(\mathbf{W}^*) + \frac{1}{2\eta T} \|\mathbf{W}^* - \mathbf{W}_0\|_2^2$, one can get the final error bound. The detailed proof can be found in Appendix A.3.

4.2 Three-layer Neural Networks

One of our technical contributions is to show that this almost co-coercivity still holds true along the GD's trajectory by a novel induction strategy fully exploring the over-parameterization. In particular, we develop stability bounds and control the generalization error by estimating the smoothness and curvature of the loss function:

$$\lambda_{\max}(\nabla^2 \ell(\mathbf{W}; z)) \leq \rho_{\mathbf{W}},$$

$$\lambda_{\min}(\nabla^2 \ell(\mathbf{W}; z)) \geq -C_{\mathbf{W}} \left(\frac{2B_{\sigma'} B_{\sigma}}{m^{2c-1}} \|\mathbf{W}^{(2)}\|_2 + \sqrt{2c_0} \right),$$

where $\rho_{\mathbf{W}}$ and $C_{\mathbf{W}}$ depend on $\|\mathbf{W}\|_2$. The specific forms of $\rho_{\mathbf{W}}$ and $C_{\mathbf{W}}$ are given in Appendix B.1. As mentioned in Remark 4, it is not easy to estimate $\|\mathbf{W}_t\|_2$ since the smoothness of the loss relies on the norm of \mathbf{W} in this case. We address this difficulty by first giving a rough bound of $\|\mathbf{W}_t - \mathbf{W}_0\|_2$ by induction, i.e., $\|\mathbf{W}_t - \mathbf{W}_0\|_2 \leq \eta t m^{2c-1}$. Then, for any \mathbf{W} produced by GD iterates, we can control $\rho_{\mathbf{W}}$ by a constant $\hat{\rho}$ if m is large enough. Finally, by assuming $\eta \hat{\rho} \leq 1/2$, we prove that $\|\mathbf{W}_t - \mathbf{W}_0\|_2 \leq \sqrt{2c_0 \eta t}$ and further get $\|\mathbf{W}_t\|_2 \leq \|\mathbf{W}_0\|_2 + \sqrt{2c_0 \eta t}$.

After estimating $\|\mathbf{W}_t\|_2$, we can develop the following almost co-coercivity of the gradient operator:

$$\begin{aligned} \langle \mathbf{W}_t - \mathbf{W}_t^{(i)}, \nabla L_{S^i}(\mathbf{W}_t) - \nabla L_{S^i}(\mathbf{W}_t^{(i)}) \rangle &\geq 2\eta(1 - 4\eta\hat{\rho}) \|\nabla L_{S^i}(\mathbf{W}_t) - \nabla L_{S^i}(\mathbf{W}_t^{(i)})\|_2^2 \\ &- \tilde{\epsilon}_t \|\mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta(\nabla L_{S^i}(\mathbf{W}_t) - \nabla L_{S^i}(\mathbf{W}_t^{(i)}))\|_2^2. \end{aligned}$$

It helps to establish the uniform stability bound $\|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2 \leq 2e\eta T \sqrt{2c_0 \hat{\rho} (\hat{\rho} \eta T + 2)}/n$.

Based on stability bounds, we get generalization bounds by Lemma 10 in Appendix A.1.

The proofs are given in Appendix B.1.

Similar to two-layer NNs, to estimate the optimization error for three-layer NNs, we first control $\mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_t\|_2^2]$ by using the smoothness and weak convexity of the loss. Then the desired bound can be obtained by the monotonically decreasing property of $\{L_S(\mathbf{W}_t)\}$. The final error bounds (Theorem 7 and Corollary 8) can be derived by plugging generalization and optimization bounds back into (3). The detailed proof can be found in Appendix B.3.

5 Conclusion

In this paper, we extend the previous work (Lei et al., 2022; Richards and Kuzborskij, 2021) by establishing the stability and generalization of GD for NNs with generic network scaling factors. Our results provide a quantitative condition that relates the scaling factor and network complexity to the achievement of the desired excess risk rate in two-layer and three-layer NNs.

There are several remaining questions for further study. Firstly, can our analysis of GD for NNs be extended to SGD with reduced computation cost? The main challenge lies in the non-monotonicity of objective functions in SGD, which is crucial for the stability analysis. Secondly, our analysis for three-layer NNs does not apply to the case of $c = 1/2$. It would be interesting to develop a result specifically for this setting. Lastly, extending our results to multiple layers would be interesting. The main difficulty lies in estimating the maximum and minimum eigenvalues of a Hessian matrix of the empirical risk, which rely on the norm of the coupled weights of different layers.

Acknowledgments

The corresponding author is Yiming Ying. The work by Yunwen Lei is partially supported by the Research Grants Council of Hong Kong [Project No. 22303723]. The work of Di Wang is supported in part by the baseline funding BAS/1/1689-01-01, funding from the CRG grand URF/1/4663-01-01, and URF/1/5508-01-01 of King Abdullah University of Science and Technology (KAUST). The work of Ding-Xuan Zhou is partially supported by Discovery Project (DP240101919) of the Australian Research Council.

References

- Allen-Zhu, Z., Li, Y., and Liang, Y. (2019a). Learning and generalization in over-parameterized neural networks, going beyond two layers. In *Advances in Neural Information Processing Systems*, volume 32.
- Allen-Zhu, Z., Li, Y., and Song, Z. (2019b). A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pages 242–252. PMLR.
- Arora, S., Du, S., Hu, W., Li, Z., and Wang, R. (2019a). Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pages 322–332. PMLR.
- Arora, S., Du, S. S., Hu, W., Li, Z., Salakhutdinov, R. R., and Wang, R. (2019b). On exact computation with an infinitely wide neural net. *Advances in Neural Information Processing Systems*, 32.

- Bahdanau, D., Cho, K., and Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Bartlett, P. L., Foster, D. J., and Telgarsky, M. J. (2017). Spectrally-normalized margin bounds for neural networks. *Advances in Neural Information Processing Systems*, 30.
- Bartlett, P. L., Montanari, A., and Rakhlin, A. (2021). Deep learning: a statistical viewpoint. *Acta numerica*, 30:87–201.
- Bousquet, O. and Elisseeff, A. (2002). Stability and generalization. *The Journal of Machine Learning Research*, 2:499–526.
- Brutzkus, A., Globerson, A., Malach, E., and Shalev-Shwartz, S. (2018). Sgd learns over-parameterized networks that provably generalize on linearly separable data. In *International Conference on Learning Representations*.
- Cao, Y. and Gu, Q. (2019). Generalization bounds of stochastic gradient descent for wide and deep neural networks. In *Advances in Neural Information Processing Systems*, volume 32.
- Charles, Z. and Papailiopoulos, D. (2018). Stability and generalization of learning algorithms that converge to global optima. In *International Conference on Machine Learning*, pages 745–754. PMLR.
- Chizat, L. and Bach, F. (2018). On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Advances in Neural Information Processing Systems*, volume 31.

- Chizat, L., Oyallon, E., and Bach, F. (2019). On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems*, volume 32.
- Du, S., Lee, J., Li, H., Wang, L., and Zhai, X. (2019). Gradient descent finds global minima of deep neural networks. In *International conference on machine learning*, pages 1675–1685. PMLR.
- Du, S. S., Zhai, X., Póczos, B., and Singh, A. (2018). Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*.
- Golowich, N., Rakhlin, A., and Shamir, O. (2018). Size-independent sample complexity of neural networks. In *Conference On Learning Theory*, pages 297–299. PMLR.
- Hardt, M., Recht, B., and Singer, Y. (2016). Train faster, generalize better: Stability of stochastic gradient descent. In *International Conference on Machine Learning*, pages 1225–1234. PMLR.
- Hinton, G., Deng, L., Yu, D., Dahl, G. E., Mohamed, A.-r., Jaitly, N., Senior, A., Vanhoucke, V., Nguyen, P., Sainath, T. N., et al. (2012). Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal processing magazine*, 29(6):82–97.
- Jacot, A., Gabriel, F., and Hongler, C. (2018). Neural tangent kernel: Convergence and generalization in neural networks. *Advances in Neural Information Processing Systems*, 31.
- Ju, P., Lin, X., and Shroff, N. (2022). On the generalization power of the overfitted

- three-layer neural tangent kernel model. *Advances in Neural Information Processing Systems*, 35:26135–26146.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2017). Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90.
- Kuzborskij, I. and Lampert, C. (2018). Data-dependent stability of stochastic gradient descent. In *International Conference on Machine Learning*, pages 2820–2829.
- Kuzborskij, I. and Szepesvári, C. (2022). Learning lipschitz functions by gd-trained shallow overparameterized relu neural networks. *arXiv preprint arXiv:2212.13848*.
- Lei, Y., Jin, R., and Ying, Y. (2022). Stability and generalization analysis of gradient methods for shallow neural networks. In *Advances in Neural Information Processing Systems*, volume 35.
- Lei, Y. and Ying, Y. (2020a). Fine-grained analysis of stability and generalization for stochastic gradient descent. In *International Conference on Machine Learning*, pages 5809–5819.
- Lei, Y. and Ying, Y. (2020b). Sharper generalization bounds for learning with gradient-dominated objective functions. In *International Conference on Learning Representations*.
- Li, Y. and Liang, Y. (2018). Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*, volume 31.

- Long, P. M. and Sedghi, H. (2019). Generalization bounds for deep convolutional neural networks. *arXiv preprint arXiv:1905.12600*.
- Mei, S., Misiakiewicz, T., and Montanari, A. (2019). Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. In *Conference on Learning Theory*, pages 2388–2464. PMLR.
- Mei, S., Montanari, A., and Nguyen, P.-M. (2018). A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671.
- Mou, W., Wang, L., Zhai, X., and Zheng, K. (2018). Generalization bounds of sgld for non-convex learning: Two theoretical viewpoints. In *Conference on Learning Theory*, pages 605–638.
- Neyshabur, B., Bhojanapalli, S., and Srebro, N. (2018). A pac-bayesian approach to spectrally-normalized margin bounds for neural networks. In *International Conference on Learning Representations*.
- Neyshabur, B., Tomioka, R., and Srebro, N. (2015). Norm-based capacity control in neural networks. In *Conference on Learning Theory*, pages 1376–1401. PMLR.
- Nguyen, M. and Mücke, N. (2023). How many neurons do we need? a refined analysis for shallow networks trained with gradient descent. *arXiv preprint arXiv:2309.08044*.
- Nitanda, A., Chinot, G., and Suzuki, T. (2019). Gradient descent can learn less over-parameterized two-layer neural networks on classification problems. *arXiv preprint arXiv:1905.09870*.

- Richards, D. and Kuzborskij, I. (2021). Stability & generalisation of gradient descent for shallow neural networks without the neural tangent kernel. In *Advances in Neural Information Processing Systems*, volume 34. PMLR.
- Richards, D. and Rabbat, M. (2021). Learning with gradient descent and weakly convex losses. In *International Conference on Artificial Intelligence and Statistics*, pages 1990–1998. PMLR.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489.
- Srebro, N., Sridharan, K., and Tewari, A. (2010). Smoothness, low noise and fast rates. In *Advances in Neural Information Processing Systems*, pages 2199–2207.
- Taheri, H. and Thrampoulidis, C. (2024). Generalization and stability of interpolating neural networks with minimal width. *Journal of Machine Learning Research*, 25(156):1–41.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2021). Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115.
- Zhou, Y., Liang, Y., and Zhang, H. (2022). Understanding generalization error of sgd in nonconvex optimization. *Machine Learning*, 111(1):345–375.

Appendix

A Proofs of Two-layer Neural Networks

A.1 Proofs of Generalization Bounds

We first introduce the self-bounding property of smooth functions (Srebro et al., 2010).

Lemma 9 (Self-bounding property). *Suppose for all $z \in \mathcal{Z}$, the function $\mathbf{W} \mapsto \ell(\mathbf{W}; z)$ is nonnegative and ρ -smooth. Then $\|\nabla \ell(\mathbf{W}; z)\|_2^2 \leq 2\rho \ell(\mathbf{W}; z)$.*

The following lemma establishes the connection (Lei and Ying, 2020a) between the on-average argument stability and the generalization error.

Lemma 10. *Let \mathcal{A} be an algorithm. If for any z , the map $\mathbf{W} \mapsto \ell(\mathbf{W}; z)$ is ρ -smooth, then*

$$\begin{aligned} \mathbb{E}[L(\mathcal{A}(S)) - L_S(\mathcal{A}(S))] &\leq \frac{\rho}{2n} \sum_{i=1}^n \mathbb{E}[\|\mathcal{A}(S) - \mathcal{A}(S^{(i)})\|_2^2] \\ &\quad + \left(\frac{2\rho \mathbb{E}[L_S(\mathcal{A}(S))]}{n} \sum_{i=1}^n \mathbb{E}[\|\mathcal{A}(S) - \mathcal{A}(S^{(i)})\|_2^2] \right)^{\frac{1}{2}}. \end{aligned}$$

We work with vectorized quantities so $\mathbf{W} \in \mathbb{R}^{md}$. Then $\nabla f_{\mathbf{W}}(\mathbf{x}) \in \mathbb{R}^{md}$ and $\nabla^2 f_{\mathbf{W}}(\mathbf{x}) \in \mathbb{R}^{md \times md}$. Denote by $\|\mathbf{W}\|_{op}$ the spectral norm of a matrix \mathbf{W} . We first introduce the following lemma, which shows that the loss function is smooth and weakly convex.

Lemma 11 (Smoothness and Curvature). *Suppose Assumptions 1 and 2 hold. Let \mathbf{W}_0 be the initial point of GD. For any fixed $\mathbf{W} \in \mathbb{R}^{m \times d}$ and any $z \in \mathcal{Z}$, there holds*

$$\lambda_{\max}(\nabla^2 \ell(\mathbf{W}; z)) \leq \rho \text{ with } \rho = c_{\mathbf{x}}^2 \left(\frac{B_{\sigma'}^2 + B_{\sigma} B_{\sigma''}}{m^{2c-1}} + \frac{B_{\sigma''} c_y}{m^c} \right),$$

$$\lambda_{\min}(\nabla^2 \ell(\mathbf{W}; z)) \geq -\left(\frac{c_{\mathbf{x}}^3 B_{\sigma'} B_{\sigma''}}{m^{2c-\frac{1}{2}}} \|\mathbf{W} - \mathbf{W}_0\|_2 + \frac{c_{\mathbf{x}}^2 B_{\sigma} \sqrt{2c_0}}{m^c}\right).$$

Proof. Recall that $f_{\mathbf{W}}(\mathbf{x}) = \frac{1}{m^c} \sum_{k=1}^m a_k \sigma(\mathbf{w}_k \mathbf{x})$. Let $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_m) \in \mathbb{R}^{dm}$ with $\mathbf{v}_k \in \mathbb{R}^d$ where $k \in [m]$. According to Assumption 1,2 and noting that $|a_k| = 1$, we can give the following estimations:

$$\begin{aligned} \|\nabla^2 f_{\mathbf{W}}(\mathbf{x})\|_{op} &= \max_{\mathbf{v}: \|\mathbf{v}\|_2 \leq 1} \sum_{k=1}^m \frac{a_k}{m^c} \langle \mathbf{v}_k, \mathbf{x} \rangle^2 \sigma''(\mathbf{w}_k \mathbf{x}) \\ &\leq \frac{\|\mathbf{x}\|_2^2 B_{\sigma''}}{m^c} \max_{\mathbf{v}: \|\mathbf{v}\|_2 \leq 1} \sum_{k=1}^m \|\mathbf{v}_k\|_2^2 \\ &\leq \frac{c_{\mathbf{x}}^2 B_{\sigma''}}{m^c}, \end{aligned} \quad (8)$$

$$\|\nabla f_{\mathbf{W}}(\mathbf{x})\|_2^2 = \sum_{k=1}^m \left\| \frac{a_k}{m^c} \mathbf{x} \sigma'(\mathbf{w}_k \mathbf{x}) \right\|_2^2 \leq \frac{B_{\sigma'}^2 c_{\mathbf{x}}^2}{m^{2c-1}} \quad (9)$$

and

$$|f_{\mathbf{W}}(\mathbf{x}) - y| \leq |f_{\mathbf{W}}(\mathbf{x})| + c_y \leq m^{1-c} B_{\sigma} + c_y. \quad (10)$$

Note

$$\nabla^2 \ell(\mathbf{W}; z) = \nabla f_{\mathbf{W}}(\mathbf{x}) \nabla f_{\mathbf{W}}(\mathbf{x})^{\top} + \nabla^2 f_{\mathbf{W}}(\mathbf{x}) (f_{\mathbf{W}}(\mathbf{x}) - y). \quad (11)$$

Then for any $\mathbf{W} \in \mathbb{R}^{md}$, we can upper bound the maximum eigenvalue of the Hessian as

$$\begin{aligned} \lambda_{\max}(\nabla^2 \ell(\mathbf{W}; z)) &\leq \|\nabla f_{\mathbf{W}}(\mathbf{x})\|_2^2 + \|\nabla^2 f_{\mathbf{W}}(\mathbf{x})\|_{op} |f_{\mathbf{W}}(\mathbf{x}) - y| \\ &\leq \frac{B_{\sigma'}^2 c_{\mathbf{x}}^2}{m^{2c-1}} + \frac{c_{\mathbf{x}}^2 B_{\sigma''}}{m^c} (m^{1-c} B_{\sigma} + c_y) \\ &= c_{\mathbf{x}}^2 \left(\frac{B_{\sigma'}^2 + B_{\sigma} B_{\sigma''}}{m^{2c-1}} + \frac{B_{\sigma''} c_y}{m^c} \right), \end{aligned}$$

which implies that the loss function is ρ -smooth with $\rho = c_{\mathbf{x}}^2 \left(\frac{B_{\sigma'}^2 + B_{\sigma} B_{\sigma''}}{m^{2c-1}} + \frac{B_{\sigma''} c_y}{m^c} \right)$.

For any $\mathbf{W}, \mathbf{W}' \in \mathbb{R}^{md}$, from Assumption 1 we know

$$\begin{aligned} |f_{\mathbf{W}}(\mathbf{x}) - f_{\mathbf{W}'}(\mathbf{x})| &\leq \frac{1}{m^c} \sum_{k=1}^m |\sigma(\mathbf{w}_k \mathbf{x}) - \sigma(\mathbf{w}'_k \mathbf{x})| \leq \frac{B_{\sigma'}}{m^c} \sum_{k=1}^m |(\mathbf{w}_k - \mathbf{w}'_k) \mathbf{x}| \\ &\leq \frac{c_{\mathbf{x}} B_{\sigma'}}{m^{c-1/2}} \|\mathbf{W} - \mathbf{W}'\|_2. \end{aligned} \quad (12)$$

Combining (11) with the fact that $\nabla f_{\mathbf{W}}(\mathbf{x}) \nabla f_{\mathbf{W}}(\mathbf{x})^\top$ is positive semi-definite together,

we obtain

$$\begin{aligned} \lambda_{\min}(\nabla^2 \ell(\mathbf{W}; z)) &\geq -\|\nabla^2 f_{\mathbf{W}}(\mathbf{x})\|_{op} |f_{\mathbf{W}}(\mathbf{x}) - y| \\ &\geq -\frac{c_{\mathbf{x}}^2 B_{\sigma''}}{m^c} \left(|f_{\mathbf{W}}(\mathbf{x}) - f_{\mathbf{W}_0}(\mathbf{x})| + |f_{\mathbf{W}_0}(\mathbf{x}) - y| \right) \\ &\geq -\frac{c_{\mathbf{x}}^2 B_{\sigma''}}{m^c} \left(\frac{c_{\mathbf{x}} B_{\sigma'}}{m^{c-1/2}} \|\mathbf{W} - \mathbf{W}_0\|_2 + \sqrt{2\ell(\mathbf{W}_0; z)} \right) \\ &\geq -\frac{c_{\mathbf{x}}^2 B_{\sigma''}}{m^c} \left(\frac{c_{\mathbf{x}} B_{\sigma'}}{m^{c-1/2}} \|\mathbf{W} - \mathbf{W}_0\|_2 + \sqrt{2c_0} \right), \end{aligned} \quad (13)$$

where in the third inequality we used (12) with $\mathbf{W}' = \mathbf{W}_0$ and in the last inequality we used Assumption 2. The proof is completed. \square

To give an upper bound of the uniform stability, we need the following lemma which shows how the GD iterate will deviate from the initial point.

Lemma 12 (Richards and Kuzborskij (2021)). *Suppose the loss is ρ -smooth and $\eta \leq 1/(2\rho)$. Then for any $t \geq 0, i \in [n]$,*

$$\|\mathbf{W}_t - \mathbf{W}_0\|_2 \leq \sqrt{2\eta t L_S(\mathbf{W}_0)},$$

$$\|\mathbf{W}_t^{(i)} - \mathbf{W}_0\|_2 \leq \sqrt{2\eta t L_{S^{(i)}}(\mathbf{W}_0)}.$$

The following lemma shows an almost co-coercivity of the gradient operator associated with shallow neural networks. For any $i \in [n]$, define $S^{(i)} = \{z_1, \dots, z_{i-1}, z'_i, z_{i+1}, \dots, z_n\}$

as the set formed from S by replacing the i -th element with z'_i . For any $\mathbf{W} \in \mathcal{W}$,

$$L_{S \setminus i}(\mathbf{W}) = L_S(\mathbf{W}) - \frac{1}{n} \ell(\mathbf{W}; z_i) = L_{S^{(i)}}(\mathbf{W}) - \frac{1}{n} \ell(\mathbf{W}; z'_i).$$

Let $\{\mathbf{W}_t\}$ and $\{\mathbf{W}_t^{(i)}\}$ be the sequence produced by GD based on S and $S^{(i)}$, respectively.

Lemma 13 (Almost Co-coercivity of the Gradient Operator). *Suppose the loss is ρ -smooth and $\eta \leq 1/(2\rho)$. Then*

$$\begin{aligned} \langle \mathbf{W}_t - \mathbf{W}_t^{(i)}, \nabla L_{S \setminus i}(\mathbf{W}_t) - \nabla L_{S \setminus i}(\mathbf{W}_t^{(i)}) \rangle &\geq 2\eta \left(1 - \frac{\eta\rho}{2}\right) \|\nabla L_{S \setminus i}(\mathbf{W}_t) - \nabla L_{S \setminus i}(\mathbf{W}_t^{(i)})\|_2^2 \\ &\quad - \epsilon_t \|\mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta(\nabla L_{S \setminus i}(\mathbf{W}_t) - \nabla L_{S \setminus i}(\mathbf{W}_t^{(i)}))\|_2^2, \end{aligned}$$

$$\text{where } \epsilon_t = \frac{c_{\mathbf{x}}^2 B_{\sigma''}}{m^c} \left(\frac{c_{\mathbf{x}} B_{\sigma'}}{m^{c-1/2}} (1 + \eta\rho) \|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2 + \frac{c_{\mathbf{x}} B_{\sigma'} \sqrt{2\eta T c_0}}{m^{c-1/2}} + \sqrt{2c_0} \right).$$

Proof. This lemma can be proved in a similar way as Lemma 5 in Richards and Kuzborskij (2021) except the estimation of the eigenvalue of Hessian matrix. Specifically, for $\alpha \in [0, 1]$, let $\mathbf{W}(\alpha) = \alpha \mathbf{W}_t + (1 - \alpha) \mathbf{W}_t^{(i)} - \alpha \eta (\nabla \ell(\mathbf{W}_t; z) - \nabla \ell(\mathbf{W}_t^{(i)}; z))$. According to (8) and (11), for any $\mathbf{W} \in \mathcal{W}$, we know

$$\lambda_{\min}(\nabla^2 L_{S \setminus i}(\mathbf{W})) \geq -\frac{c_{\mathbf{x}}^2 B_{\sigma''}}{m^c} \left(\frac{1}{n} \sum_{j \in [n], j \neq i} |f_{\mathbf{W}}(\mathbf{x}_j) - y_j| \right).$$

Let $\mathbf{W} = \mathbf{W}(\alpha)$. Note Lemma 11 shows that the loss is ρ -smooth with $\rho = c_{\mathbf{x}}^2 \left(\frac{B_{\sigma'}^2 + B_{\sigma} B_{\sigma''}}{m^{2c-1}} + \right.$

$\frac{B_{\sigma''}c_y}{m^c}$). Then from (12) and the smoothness of ℓ we can get

$$\begin{aligned}
& \frac{1}{n} \sum_{j \in [n], j \neq i} |f_{\mathbf{W}(\alpha)}(\mathbf{x}_j) - y_j| \\
& \leq \frac{1}{n} \sum_{j \in [n], j \neq i} |(f_{\mathbf{W}(\alpha)}(\mathbf{x}_j) - f_{\mathbf{W}_t^{(i)}}(\mathbf{x}_j)) + (f_{\mathbf{W}_t^{(i)}}(\mathbf{x}_j) - f_{\mathbf{W}_0}(\mathbf{x}_j)) + (f_{\mathbf{W}_0}(\mathbf{x}_j) - y_j)| \\
& \leq \frac{c_{\mathbf{x}}B_{\sigma'}}{m^{c-1/2}} \|\mathbf{W}(\alpha) - \mathbf{W}_t^{(i)}\|_2 + \frac{c_{\mathbf{x}}B_{\sigma'}}{m^{c-1/2}} \|\mathbf{W}_t^{(i)} - \mathbf{W}_0\|_2 + \sqrt{2L_{S^i}(\mathbf{W}_0)} \\
& \leq \frac{c_{\mathbf{x}}B_{\sigma'}\alpha}{m^{c-1/2}} (\|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2 + \eta \|\nabla \ell(\mathbf{W}_t; z) - \nabla \ell(\mathbf{W}_t^{(i)}; z)\|_2) + \frac{c_{\mathbf{x}}B_{\sigma'}\sqrt{2\eta T c_0}}{m^{c-1/2}} + \sqrt{2c_0} \\
& \leq \frac{c_{\mathbf{x}}B_{\sigma'}(1 + \eta\rho)}{m^{c-1/2}} \|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2 + \frac{c_{\mathbf{x}}B_{\sigma'}\sqrt{2\eta T c_0}}{m^{c-1/2}} + \sqrt{2c_0},
\end{aligned}$$

where in the third inequality we used Lemma 12 and $\ell(\mathbf{W}_0; z) \leq c_0$.

Combining the above two inequalities together, we get

$$\lambda_{\min}(\nabla^2 L_{S^i}(\mathbf{W}(\alpha))) \geq -\frac{c_{\mathbf{x}}^2 B_{\sigma''}}{m^c} \left(\frac{c_{\mathbf{x}}B_{\sigma'}(1 + \eta\rho)}{m^{c-1/2}} \|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2 + \frac{c_{\mathbf{x}}B_{\sigma'}\sqrt{2\eta T c_0}}{m^{c-1/2}} + \sqrt{2c_0} \right).$$

Similarly, let $\widetilde{\mathbf{W}}(\alpha) = \alpha \mathbf{W}_t^{(i)} + (1 - \alpha) \mathbf{W}_t - \alpha \eta (\nabla \ell(\mathbf{W}_t^{(i)}; z) - \nabla \ell(\mathbf{W}_t; z))$, we

can prove that

$$\lambda_{\min}(\nabla^2 L_{S^i}(\widetilde{\mathbf{W}}(\alpha))) \geq -\frac{c_{\mathbf{x}}^2 B_{\sigma''}}{m^c} \left(\frac{c_{\mathbf{x}}B_{\sigma'}(1 + \eta\rho)}{m^{c-1/2}} \|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2 + \frac{c_{\mathbf{x}}B_{\sigma'}\sqrt{2\eta T c_0}}{m^{c-1/2}} + \sqrt{2c_0} \right).$$

The remaining arguments in proving the lemma are the same as Lemma 5 in Richards and Kuzborskij (2021). We omit the proof for simplicity. \square

Based on the almost co-coercivity property of the gradient operator, we give the following uniform stability theorem.

Theorem 14 (Uniform Stability). *Suppose Assumptions 1 and 2 hold. Let $S, S^{(i)}$ be constructed in Definition 2. Let $\{\mathbf{W}_t\}$ and $\{\mathbf{W}_t^{(i)}\}$ be produced by (2) with $\eta \leq 1/(2\rho)$ based on S and $S^{(i)}$, respectively. Assume (4) holds. For any $t \in [T]$, there holds*

$$\|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2 \leq \frac{2\eta e T \sqrt{2c_0 \rho (\rho \eta T + 2)}}{n}.$$

Proof. Recall that

$$L_{S^{\setminus i}}(\mathbf{W}) = L_S(\mathbf{W}) - \frac{1}{n}\ell(\mathbf{W}; z_i) = L_{S^{(i)}}(\mathbf{W}) - \frac{1}{n}\ell(\mathbf{W}; z'_i).$$

Note $\mathcal{W} = \mathbb{R}^{m \times d}$. Then by the update rule $\mathbf{W}_{t+1} = \mathbf{W}_t - \eta \nabla L_S(\mathbf{W}_t)$, there holds

$$\begin{aligned} & \|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2^2 \\ &= \|\mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta(\nabla L_{S^{\setminus i}}(\mathbf{W}_t) - \nabla L_{S^{\setminus i}}(\mathbf{W}_t^{(i)})) - \frac{\eta}{n}(\nabla \ell(\mathbf{W}_t; z_i) - \nabla \ell(\mathbf{W}_t^{(i)}; z'_i))\|_2^2 \\ &\leq (1+p)\|\mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta(\nabla L_{S^{\setminus i}}(\mathbf{W}_t) - \nabla L_{S^{\setminus i}}(\mathbf{W}_t^{(i)}))\|_2^2 \\ &\quad + \frac{\eta^2(1+1/p)}{n^2}\|\nabla \ell(\mathbf{W}_t; z_i) - \nabla \ell(\mathbf{W}_t^{(i)}; z'_i)\|_2^2 \\ &\leq (1+p)\|\mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta(\nabla L_{S^{\setminus i}}(\mathbf{W}_t) - \nabla L_{S^{\setminus i}}(\mathbf{W}_t^{(i)}))\|_2^2 \\ &\quad + \frac{2\eta^2(1+1/p)}{n^2}\left(\|\nabla \ell(\mathbf{W}_t; z_i)\|_2^2 + \|\nabla \ell(\mathbf{W}_t^{(i)}; z'_i)\|_2^2\right), \end{aligned} \tag{14}$$

where in the first inequality we used $(a+b)^2 \leq (1+p)a^2 + (1+1/p)b^2$.

According to Lemma 13 we can get

$$\begin{aligned} & \|\mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta(\nabla L_{S^{\setminus i}}(\mathbf{W}_t) - \nabla L_{S^{\setminus i}}(\mathbf{W}_t^{(i)}))\|_2^2 \\ &= \|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2^2 + \eta^2\|\nabla L_{S^{\setminus i}}(\mathbf{W}_t) - \nabla L_{S^{\setminus i}}(\mathbf{W}_t^{(i)})\|_2^2 \\ &\quad - 2\eta\langle \mathbf{W}_t - \mathbf{W}_t^{(i)}, \nabla L_{S^{\setminus i}}(\mathbf{W}_t) - \nabla L_{S^{\setminus i}}(\mathbf{W}_t^{(i)}) \rangle \\ &\leq \|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2^2 + \eta^2\|\nabla L_{S^{\setminus i}}(\mathbf{W}_t) - \nabla L_{S^{\setminus i}}(\mathbf{W}_t^{(i)})\|_2^2 \\ &\quad - 4\eta^2\left(1 - \frac{\eta\rho}{2}\right)\|\nabla L_{S^{\setminus i}}(\mathbf{W}_t) - \nabla L_{S^{\setminus i}}(\mathbf{W}_t^{(i)})\|_2^2 \\ &\quad + 2\eta\epsilon_t\|\mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta(\nabla L_{S^{\setminus i}}(\mathbf{W}_t) - \nabla L_{S^{\setminus i}}(\mathbf{W}_t^{(i)}))\|_2^2, \end{aligned}$$

where $\epsilon_t = \frac{c_{\mathbf{x}}B_{\sigma'}\mu}{m^c} \left(\frac{c_{\mathbf{x}}B_{\sigma'}}{m^{c-1/2}}(1+\eta\rho)\|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2 + \frac{c_{\mathbf{x}}B_{\sigma'}\sqrt{2\eta T c_0}}{m^{c-1/2}} + \sqrt{2c_0} \right)$.

Rearranging the above inequality and noting that $\eta\rho \leq 1/2$, we obtain

$$\begin{aligned} & (1 - 2\eta\epsilon_t) \left\| \mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta(\nabla L_{S \setminus i}(\mathbf{W}_t) - \nabla L_{S \setminus i}(\mathbf{W}_t^{(i)})) \right\|_2^2 \\ & \leq \left\| \mathbf{W}_t - \mathbf{W}_t^{(i)} \right\|_2^2 + \eta^2(2\eta\rho - 3) \left\| \nabla L_{S \setminus i}(\mathbf{W}_t) - \nabla L_{S \setminus i}(\mathbf{W}_t^{(i)}) \right\|_2^2 \leq \left\| \mathbf{W}_t - \mathbf{W}_t^{(i)} \right\|_2^2. \end{aligned}$$

We can choose m large enough to ensure $2\eta\epsilon_t < 1$ holds for any $t \in [T]$. Indeed,

$2\eta\epsilon_t < 1$ holds as long as condition (4) holds. We will discuss it at the end of the proof.

Now, plugging the above inequality back into (14) yields

$$\begin{aligned} & \left\| \mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)} \right\|_2^2 \\ & \leq \frac{1+p}{1-2\eta\epsilon_t} \left\| \mathbf{W}_t - \mathbf{W}_t^{(i)} \right\|_2^2 + \frac{2\eta^2(1+1/p)}{n^2} \left(\left\| \nabla \ell(\mathbf{W}_t; z_i) \right\|_2^2 + \left\| \nabla \ell(\mathbf{W}_t^{(i)}; z'_i) \right\|_2^2 \right). \end{aligned} \quad (15)$$

We can apply (15) recursively and derive

$$\begin{aligned} & \left\| \mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)} \right\|_2^2 \\ & \leq \frac{2\eta^2(1+1/p)}{n^2} \sum_{j=0}^t \left(\left\| \nabla \ell(\mathbf{W}_j; z_i) \right\|_2^2 + \left\| \nabla \ell(\mathbf{W}_j^{(i)}; z'_i) \right\|_2^2 \right) \prod_{\tilde{j}=j+1}^t \frac{1+p}{1-2\eta\epsilon_{\tilde{j}}}, \end{aligned} \quad (16)$$

where we used $\mathbf{W}_0 = \mathbf{W}_0^{(i)}$.

According to Lemma 9 and Lemma 12, we know

$$\begin{aligned} \left\| \nabla \ell(\mathbf{W}_j; z) \right\|_2^2 & \leq 2 \left\| \nabla \ell(\mathbf{W}_j; z) - \nabla \ell(\mathbf{W}_0; z) \right\|_2^2 + 2 \left\| \nabla \ell(\mathbf{W}_0; z) \right\|_2^2 \\ & \leq 2\rho^2 \left\| \mathbf{W}_j - \mathbf{W}_0 \right\|_2^2 + 4\rho \ell(\mathbf{W}_0; z) \leq 4\rho^2 \eta j L_S(\mathbf{W}_0) + 4\rho \ell(\mathbf{W}_0; z). \end{aligned}$$

Similarly, we can show that

$$\left\| \nabla \ell(\mathbf{W}_j^{(i)}; z) \right\|_2^2 \leq 4\rho^2 \eta j L_{S^{(i)}}(\mathbf{W}_0) + 4\rho \ell(\mathbf{W}_0; z).$$

Combining the above three inequalities together, we get

$$\begin{aligned}
& \|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2^2 \\
& \leq \frac{8\rho\eta^2(1+1/p)}{n^2} \sum_{j=0}^t \left(\rho\eta j L_S(\mathbf{W}_0) + \ell(\mathbf{W}_0; z_i) + \rho\eta j L_{S^{(i)}}(\mathbf{W}_0) + \ell(\mathbf{W}_0; z'_i) \right) \prod_{\tilde{j}=j+1}^t \frac{1+p}{1-2\eta\epsilon_{\tilde{j}}} \\
& \leq \frac{8\rho\eta^2(1+1/p)}{n^2} \prod_{\tilde{j}=1}^t \frac{1+p}{1-2\eta\epsilon_{\tilde{j}}} \sum_{j=0}^t \left(\rho\eta j L_S(\mathbf{W}_0) + \ell(\mathbf{W}_0; z_i) + \rho\eta j L_{S^{(i)}}(\mathbf{W}_0) + \ell(\mathbf{W}_0; z'_i) \right) \\
& = \frac{4\rho\eta^2(1+1/p)}{n^2} \prod_{\tilde{j}=1}^t \frac{1+p}{1-2\eta\epsilon_{\tilde{j}}} \left(\rho\eta t(t+1)(L_S(\mathbf{W}_0) + L_{S^{(i)}}(\mathbf{W}_0)) + 2(t+1)(\ell(\mathbf{W}_0; z_i) \right. \\
& \quad \left. + \ell(\mathbf{W}_0; z'_i)) \right) \\
& \leq \frac{8\rho\eta^2 c_0(1+1/p)(1+t)(\rho\eta t+2)}{n^2} \prod_{\tilde{j}=1}^t \frac{1+p}{1-2\eta\epsilon_{\tilde{j}}},
\end{aligned}$$

where we used $\ell(\mathbf{W}_0; z) \leq c_0$ for any $z \in \mathcal{Z}$. If we further choose $p = 1/t$, then there holds

$$\|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2^2 \leq \frac{8\rho\eta^2 c_0 e(1+t)^2(\rho\eta t+2)}{n^2} \prod_{\tilde{j}=1}^t \frac{1}{1-2\eta\epsilon_{\tilde{j}}}, \quad (17)$$

where we used $(1+1/t)^t \leq e$.

Now, we prove by induction to show

$$\|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2 \leq \frac{2\eta e T \sqrt{2c_0\rho(\rho\eta T+2)}}{n}. \quad (18)$$

Note (18) with $k = 0$ holds trivially. Assume (18) holds with all $k \leq t$, i.e., for all $k \leq t$

$$\|\mathbf{W}_k - \mathbf{W}_k^{(i)}\|_2 \leq \frac{2\eta e T \sqrt{2c_0\rho(\rho\eta T+2)}}{n}. \quad (19)$$

and we want to show it holds with $k = t+1 \leq T$. Recall that $\epsilon_k = \frac{c_{\mathbf{x}} B_{\sigma''}}{m^c} \left(\frac{c_{\mathbf{x}} B_{\sigma'}}{m^{c-1/2}} (1 + \eta\rho) \|\mathbf{W}_k - \mathbf{W}_k^{(i)}\|_2 + \frac{c_{\mathbf{x}} B_{\sigma'} \sqrt{2\eta T c_0}}{m^{c-1/2}} + \sqrt{2c_0} \right)$. From (19), for any $\tilde{j} \leq t$, we know

$$\epsilon_{\tilde{j}} \leq \epsilon' := \frac{c_{\mathbf{x}} B_{\sigma''}}{m^c} \left(\frac{2\sqrt{2c_0\rho(\rho\eta T+2)}\eta e T(1+\eta\rho)c_{\mathbf{x}} B_{\sigma'}}{nm^{c-1/2}} + \frac{c_{\mathbf{x}} B_{\sigma'} \sqrt{2\eta T c_0}}{m^{c-1/2}} + \sqrt{2c_0} \right).$$

Putting the above inequality back into (17), we get

$$\|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2^2 \leq \frac{8\rho\eta^2 c_0 e(1+t)^2(\rho\eta t + 2)}{n^2} \left(\frac{1}{1-2\eta\epsilon'}\right)^t.$$

If m is large enough such that $2\eta\epsilon' \leq 1/(t+1)$, then we can show

$$\left(\frac{1}{1-2\eta\epsilon'}\right)^t \leq \left(\frac{1}{1-1/(t+1)}\right)^t \leq e. \quad (20)$$

Then there holds

$$\|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2 \leq \frac{2\eta e(1+t)\sqrt{2\rho c_0(\rho\eta t + 2)}}{n} \leq \frac{2\eta e T \sqrt{2\rho c_0(\rho\eta T + 2)}}{n}. \quad (21)$$

Now, we discuss the conditions on m . Suppose m satisfies the following conditions

$$m \geq C_1 \left(\frac{(\eta T)^2(1+\eta\rho)\sqrt{\rho(\rho\eta T + 2)}}{n}\right)^{\frac{2}{4c-1}}, m \geq C_2(\eta T)^{\frac{3}{4c-1}} \text{ and } m \geq C_3(\eta T)^{\frac{1}{c}},$$

where $C_1 = (8ec_x^2 B_{\sigma'} B_{\sigma''} \sqrt{2c_0})^{\frac{2}{4c-1}}$, $C_2 = (4\sqrt{2c_0} c_x^2 B_{\sigma'} B_{\sigma''})^{\frac{3}{4c-1}}$, $C_3 = (8\sqrt{2c_0} c_x B_{\sigma''})^{1/c}$.

Then it is easy to verify that

$$\frac{2\eta c_x B_{\sigma''}}{m^c} \left(\frac{2\sqrt{2c_0\rho(\rho\eta T + 2)}\eta e T(1+\eta\rho)c_x B_{\sigma'}}{nm^{c-1/2}} + \frac{c_x B_{\sigma'}\sqrt{2\eta T c_0}}{m^{c-1/2}} + \sqrt{2c_0}\right) \leq \frac{1}{T} \leq \frac{1}{1+t},$$

which ensures that $2\eta\epsilon' \leq 1/(t+1)$, and then (21) holds. The proof is completed. \square

We can combine Theorem 14 and Lemma 10 together to get the upper bound of the generalization error.

Proof of Theorem 1. Eq.(16) with $p = 1/t$ and Eq.(20) imply

$$\begin{aligned} \|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2^2 &\leq \frac{2e^2\eta^2(1+t)}{n^2} \sum_{j=0}^t \left(\|\nabla\ell(\mathbf{W}_j; z_i)\|_2^2 + \|\nabla\ell(\mathbf{W}_j^{(i)}; z'_i)\|_2^2 \right) \\ &\leq \frac{4e^2\eta^2\rho(1+t)}{n^2} \sum_{j=0}^t \left(\ell(\mathbf{W}_j; z_i) + \ell(\mathbf{W}_j^{(i)}; z'_i) \right), \end{aligned}$$

where in the last inequality we used the self-bounding property of the smooth loss (Lemma 9). Now, taking an average over $i \in [n]$ and using $\mathbb{E}[\ell(\mathbf{W}_j; z_i)] = \mathbb{E}[\ell(\mathbf{W}_j^{(i)}; z'_i)]$, we have

$$\begin{aligned}
& \frac{1}{n} \sum_{i=1}^n \mathbb{E} \|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2^2 \\
& \leq \frac{4e^2\eta^2\rho(1+t)}{n^3} \sum_{j=0}^t \left(\sum_{i=1}^n \mathbb{E}[\ell(\mathbf{W}_j; z_i)] + \mathbb{E}[\ell(\mathbf{W}_j^{(i)}; z'_i)] \right) \\
& = \frac{8e^2\eta^2\rho(1+t)}{n^3} \sum_{j=0}^t \sum_{i=1}^n \mathbb{E}[\ell(\mathbf{W}_j; z_i)] \\
& = \frac{8e^2\eta^2\rho(1+t)}{n^2} \sum_{j=0}^t \mathbb{E}[L_S(\mathbf{W}_j)].
\end{aligned}$$

Combining the above stability bounds with Lemma 10 together, we get

$$\begin{aligned}
& \mathbb{E}[L(\mathbf{W}_t) - L_S(\mathbf{W}_t)] \\
& \leq \frac{4e^2\eta^2\rho^2t}{n^2} \sum_{j=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_j)] + \left(\frac{16e^2\eta^2\rho^2t\mathbb{E}[L_S(\mathbf{W}_t)]}{n^2} \sum_{j=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_j)] \right)^{\frac{1}{2}} \\
& \leq \frac{4e^2\eta^2\rho^2t}{n^2} \sum_{j=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_j)] + \frac{4e\eta\rho}{n} \sum_{j=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_j)],
\end{aligned}$$

where in the last inequality we used $L_S(\mathbf{W}_t) \leq \frac{1}{t} \sum_{j=1}^{t-1} L_S(\mathbf{W}_j)$ (Richards and Kuzborskij, 2021). The proof is completed. \square

A.2 Proofs of Optimization Bounds

Before giving the proofs of optimization error bounds, we first introduce the following lemma on the bound of GD iterates.

Lemma 15. *Suppose Assumptions 1 and 2 hold, and $\eta \leq 1/(2\rho)$. Assume (4) and (5)*

hold. Then for any $t \in [T]$, there holds

$$\begin{aligned} 1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_t\|_2^2] &\leq \left(\frac{8e^2\eta^3\rho^2t^2}{n^2} + \frac{8e\eta^2t\rho}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] + 2\mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2] \\ &\quad + 2\eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)]. \end{aligned}$$

Proof. For any $\mathbf{W}, \widetilde{\mathbf{W}} \in \mathbb{R}^{md}$ and $\alpha \in [0, 1]$, define $\mathbf{W}(\alpha) := \widetilde{\mathbf{W}} + \alpha(\mathbf{W} - \widetilde{\mathbf{W}})$. Note

that

$$\begin{aligned} &\lambda_{\min}(\nabla^2 L_S(\mathbf{W}(\alpha))) \\ &\geq -\max_i \{\|\nabla^2 f(\mathbf{x}_i)\|_2\} \left(\frac{1}{n} \sum_{i=1}^n |f_{\mathbf{W}(\alpha)}(\mathbf{x}_i) - y_i| \right) \\ &\geq -\frac{c_x^2 B_{\sigma''}}{m^c} \frac{1}{n} \left(\sum_{i=1}^n (|f_{\mathbf{W}(\alpha)}(\mathbf{x}_i) - f_{\widetilde{\mathbf{W}}}(\mathbf{x}_i)| + |f_{\widetilde{\mathbf{W}}}(\mathbf{x}_i) - f_{\mathbf{W}_0}(\mathbf{x}_i)| + |f_{\mathbf{W}_0}(\mathbf{x}_i) - y_i|) \right) \\ &\geq -\frac{c_x^2 B_{\sigma''}}{m^c} \left(\frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \|\mathbf{W}(\alpha) - \widetilde{\mathbf{W}}\|_2 + \frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \|\widetilde{\mathbf{W}} - \mathbf{W}_0\|_2 + \sqrt{2L_S(\mathbf{W}_0)} \right) \\ &\geq -\frac{c_x^2 B_{\sigma''}}{m^c} \left(\frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \|\mathbf{W} - \widetilde{\mathbf{W}}\|_2 + \frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \|\widetilde{\mathbf{W}} - \mathbf{W}_0\|_2 + \sqrt{2c_0} \right). \end{aligned}$$

Then for any $t \in [T]$, let $\widetilde{\mathbf{W}} = \mathbf{W}_t$, and define

$$\begin{aligned} g(\alpha) &:= L_S(\mathbf{W}(\alpha)) + \frac{c_x^2 B_{\sigma''}}{m^c} \frac{\alpha^2}{2} \left(\frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \|\mathbf{W} - \mathbf{W}_t\|_2 + \frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \|\mathbf{W}_t - \mathbf{W}_0\|_2 + \sqrt{2c_0} \right) \\ &\quad \times (1 \vee \mathbb{E}[\|\mathbf{W} - \mathbf{W}_t\|_2^2]). \end{aligned}$$

It is obvious that $g''(\alpha) \geq 0$. Then $g(\alpha)$ is convex in $\alpha \in [0, 1]$. Now, by convexity we know

$$\begin{aligned} &g(1) - g(0) \\ &= L_S(\mathbf{W}) + \frac{c_x^2 B_{\sigma''}}{2m^c} \left(\frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \|\mathbf{W} - \mathbf{W}_t\|_2 + \frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \|\mathbf{W}_t - \mathbf{W}_0\|_2 + \sqrt{2c_0} \right) \times (1 \vee \mathbb{E}[\|\mathbf{W} - \mathbf{W}_t\|_2^2]) \\ &\quad - L_S(\mathbf{W}_t) \\ &\geq \langle \mathbf{W} - \mathbf{W}_t, \nabla L_S(\mathbf{W}_t) \rangle = g'(0). \end{aligned}$$

Rearranging the above inequality we get

$$\begin{aligned}
L_S(\mathbf{W}_t) &\leq L_S(\mathbf{W}) + \frac{c_{\mathbf{x}}^2 B_{\sigma''}}{2m^c} \left(\frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \|\mathbf{W} - \mathbf{W}_t\|_2 + \frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \|\mathbf{W}_t - \mathbf{W}_0\|_2 + \sqrt{2c_0} \right) \\
&\quad \times (1 \vee \mathbb{E}[\|\mathbf{W} - \mathbf{W}_t\|_2^2]) - \langle \mathbf{W} - \mathbf{W}_t, \nabla L_S(\mathbf{W}_t) \rangle. \tag{22}
\end{aligned}$$

Combining (22) with the smoothness of the loss we can get

$$\begin{aligned}
&L_S(\mathbf{W}_{t+1}) \\
&\leq L_S(\mathbf{W}_t) + \langle \nabla L_S(\mathbf{W}_t), \mathbf{W}_{t+1} - \mathbf{W}_t \rangle + \frac{\rho}{2} \|\mathbf{W}_{t+1} - \mathbf{W}_t\|_2^2 \\
&\leq L_S(\mathbf{W}_t) - \eta \langle \nabla L_S(\mathbf{W}_t), \nabla L_S(\mathbf{W}_t) \rangle + \frac{\rho}{2} \|\mathbf{W}_{t+1} - \mathbf{W}_t\|_2^2 \\
&\leq L_S(\mathbf{W}_t) - \eta \left(1 - \frac{\eta\rho}{2}\right) \|\nabla L_S(\mathbf{W}_t)\|_2^2 \\
&\leq L_S(\mathbf{W}_t) - \frac{\eta}{2} \|\nabla L_S(\mathbf{W}_t)\|_2^2 \\
&\leq L_S(\mathbf{W}) + \frac{c_{\mathbf{x}}^2 B_{\sigma''}}{2m^c} \left(\frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \|\mathbf{W} - \mathbf{W}_t\|_2 + \frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \|\mathbf{W}_t - \mathbf{W}_0\|_2 + \sqrt{2c_0} \right) \times (1 \vee \mathbb{E}[\|\mathbf{W} - \mathbf{W}_t\|_2^2]) \\
&\quad - \langle \mathbf{W} - \mathbf{W}_t, \nabla L_S(\mathbf{W}_t) \rangle - \frac{\eta}{2} \|\nabla L_S(\mathbf{W}_t)\|_2^2,
\end{aligned}$$

where in the third inequality we used the update rule (2) and $\eta\rho \leq 1$.

According to the equality $2\langle x - y, x - z \rangle = \|x - y\|_2^2 + \|x - z\|_2^2 - \|y - z\|_2^2$, we

know

$$\begin{aligned}
& - \langle \mathbf{W} - \mathbf{W}_t, \nabla L_S(\mathbf{W}_t) \rangle - \frac{\eta}{2} \|\nabla L_S(\mathbf{W}_t)\|_2^2 \\
&= \frac{1}{\eta} \langle \mathbf{W} - \mathbf{W}_t, \mathbf{W}_{t+1} - \mathbf{W}_t \rangle - \frac{1}{2\eta} \|\mathbf{W}_{t+1} - \mathbf{W}_t\|_2^2 \\
&= \frac{1}{2\eta} (\|\mathbf{W} - \mathbf{W}_t\|_2^2 - \|\mathbf{W}_{t+1} - \mathbf{W}_t\|_2^2).
\end{aligned}$$

Then there holds

$$\begin{aligned}
& L_S(\mathbf{W}_{t+1}) \\
& \leq L_S(\mathbf{W}) + \frac{c_{\mathbf{x}}^2 B_{\sigma''}}{2m^c} \left(\frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \|\mathbf{W} - \mathbf{W}_t\|_2 + \frac{B_{\sigma'} c_{\mathbf{x}} \sqrt{2\eta T c_0}}{m^{c-1/2}} + \sqrt{2c_0} \right) \times (1 \vee \mathbb{E}[\|\mathbf{W} - \mathbf{W}_t\|_2^2]) \\
& \quad + \frac{1}{2\eta} \left(\|\mathbf{W} - \mathbf{W}_t\|_2^2 - \|\mathbf{W}_{t+1} - \mathbf{W}\|_2^2 \right). \tag{23}
\end{aligned}$$

The above inequality with $\mathbf{W} = \mathbf{W}_{\frac{1}{\eta T}}^*$ implies

$$\begin{aligned}
& \frac{1}{t} \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] + \frac{\mathbb{E}[\|\mathbf{W}_t - \mathbf{W}_{\frac{1}{\eta T}}^*\|_2^2]}{2\eta t} \\
& \leq \mathbb{E}[L(\mathbf{W}_{\frac{1}{\eta T}}^*)] + \frac{\mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2]}{2\eta t} + \frac{c_{\mathbf{x}}^2 B_{\sigma''}}{2m^c t} \sum_{s=0}^{t-1} \left(\frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2] \right. \\
& \quad \left. + \frac{B_{\sigma'} c_{\mathbf{x}} \sqrt{2\eta T c_0}}{m^{c-1/2}} + \sqrt{2c_0} \right) \times (1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2]).
\end{aligned}$$

Combined the above inequality with Theorem 1 implies

$$\begin{aligned}
& \frac{\mathbb{E}[\|\mathbf{W}_t - \mathbf{W}_{\frac{1}{\eta T}}^*\|_2^2]}{2\eta t} \\
& \leq \frac{1}{t} \sum_{s=0}^{t-1} [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - \mathbb{E}[L(\mathbf{W}_s)]] + \frac{\mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2]}{2\eta t} + \left(\frac{4e^2 \eta^2 \rho^2 t}{n^2} + \frac{4e\eta\rho}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] \\
& \quad + \frac{c_{\mathbf{x}}^2 B_{\sigma''}}{2m^c t} \sum_{s=0}^{t-1} \left(\frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2] + \frac{B_{\sigma'} c_{\mathbf{x}} \sqrt{2\eta T c_0}}{m^{c-1/2}} + \sqrt{2c_0} \right) (1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2]) \\
& \leq L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*) + \frac{\mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2]}{2\eta t} + \left(\frac{4e^2 \eta^2 \rho^2 t}{n^2} + \frac{4e\eta\rho}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] \\
& \quad + \frac{c_{\mathbf{x}}^2 B_{\sigma''}}{2m^c t} \sum_{s=0}^t \left(\frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2] + \frac{B_{\sigma'} c_{\mathbf{x}} \sqrt{2\eta T c_0}}{m^{c-1/2}} + \sqrt{2c_0} \right) (1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2]), \tag{24}
\end{aligned}$$

where in the second inequality we used $L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}_s) \leq L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)$ since

$L(\mathbf{W}_s) \geq L(\mathbf{W}^*)$ for any $s \in [t-1]$.

On the other hand, using Lemma 12 we can obtain

$$\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2 \leq \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2 + \|\mathbf{W}_s - \mathbf{W}_0\|_2 \leq \sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2. \quad (25)$$

Then we know

$$\begin{aligned} & \left(\frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2] + \frac{B_{\sigma'} c_{\mathbf{x}} \sqrt{2\eta T c_0}}{m^{c-1/2}} + \sqrt{2c_0} \right) (1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2]) \\ & \leq \left(\frac{B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} (2\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) + \sqrt{2c_0} \right) (1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2]). \end{aligned}$$

Plugging the above inequality back into (24) yields

$$\begin{aligned} & \frac{\mathbb{E}[\|\mathbf{W}_t - \mathbf{W}_{\frac{1}{\eta T}}^*\|_2^2]}{2\eta t} \\ & \leq \frac{\mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2]}{2\eta t} + \frac{\tilde{b}(\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c t} \sum_{s=0}^t (1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2]) \\ & \quad + \left(\frac{4e^2 \eta^2 \rho^2 t}{n^2} + \frac{4e\eta\rho}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] + L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*), \end{aligned}$$

where $\tilde{b} = \frac{c_{\mathbf{x}}^2 B_{\sigma''}}{2} \left(\frac{2B_{\sigma'} c_{\mathbf{x}}}{m^{c-1/2}} + \sqrt{2c_0} \right)$.

Multiplying both sides by $2\eta t$ yields

$$\begin{aligned} & \mathbb{E}[\|\mathbf{W}_t - \mathbf{W}_{\frac{1}{\eta T}}^*\|_2^2] \\ & \leq \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2] + \frac{2\tilde{b}\eta(\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c} \sum_{s=0}^t (1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2]) \\ & \quad + \left(\frac{8e^2 \eta^3 \rho^2 t^2}{n^2} + \frac{8e\eta^2 \rho t}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] + 2\eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)]. \end{aligned}$$

Let $x = \max_{s \in [T]} \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2] \vee 1$. Then the above inequality implies

$$\begin{aligned} x & \leq \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2] + \frac{2\tilde{b}\eta T(\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c} x \\ & \quad + \left(\frac{8e^2 \eta^3 \rho^2 t^2}{n^2} + \frac{8e\eta^2 \rho t}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] + 2\eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)]. \end{aligned}$$

Without loss of generality, we assume $\eta \leq 1$. Condition (5) implies $m \geq (4\tilde{b}\eta T(\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2))^{\frac{1}{c}}$. Then there holds $\frac{2\tilde{b}\eta T(\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c} \leq \frac{1}{2}$. Hence

$$\begin{aligned} x \leq & \left(\frac{16e^2\eta^3\rho^2t^2}{n^2} + \frac{16e\eta^2t\rho}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] + 2\mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2] \\ & + 2\eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)]. \end{aligned}$$

It then follows that

$$\begin{aligned} 1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_t\|_2^2] \leq & \left(\frac{16e^2\eta^3\rho^2t^2}{n^2} + \frac{16e\eta^2t\rho}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] + 2\mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2] \\ & + 2\eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)]. \end{aligned}$$

This completes the proof. \square

Now, we can give the proof of Theorem 2.

Proof of Theorem 2. Recall that $\tilde{b} = \frac{c_x^2 B_{\sigma''}}{2} \left(\frac{2B_{\sigma'} c_x}{m^{c-1/2}} + \sqrt{2c_0} \right)$. Eq. (23) with $\mathbf{W} = \mathbf{W}_{\frac{1}{\eta T}}^*$

implies

$$\begin{aligned} & \frac{1}{T} \sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] \\ & \leq \mathbb{E}[L_S(\mathbf{W}_{\frac{1}{\eta T}}^*)] + \frac{\tilde{b}(\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c T} \sum_{s=0}^{T-1} 1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2] \\ & \quad + \frac{\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2}{2\eta T}, \end{aligned} \tag{26}$$

where in the last inequality we used (25).

Further, by monotonic decrease of $\{L_S(\mathbf{W}_t)\}$, we know

$$\begin{aligned} \mathbb{E}[L_S(\mathbf{W}_T)] \leq & \mathbb{E}[L_S(\mathbf{W}_{\frac{1}{\eta T}}^*)] + \frac{\tilde{b}(\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c T} \sum_{s=0}^{T-1} 1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2] \\ & + \frac{\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2}{2\eta T}. \end{aligned}$$

Note that Lemma 15 shows

$$1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_t\|_2^2] \leq \left(\frac{16e^2\eta^3\rho^2t^2}{n^2} + \frac{16e\eta^2t\rho}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] + 2\mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2] \\ + 2\eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)].$$

Combining the above two inequalities together, we get

$$\mathbb{E}[L_S(\mathbf{W}_T)] \leq \mathbb{E}[L_S(\mathbf{W}_{\frac{1}{\eta T}}^*)] + \frac{2\tilde{b}(\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c} \left(\left(\frac{8e^2\eta^3\rho^2T^2}{n^2} + \frac{8e\eta^2T\rho}{n} \right) \right. \\ \left. \times \sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + \eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)] \right) + \frac{\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2}{2\eta T}.$$

The theorem is proved. \square

Lemma 16. *Suppose Assumptions 1 and 2 hold. Let $\{\mathbf{W}_t\}$ be produced by (2) with $\eta \leq 1/(2\rho)$. Assume (4) and (5) hold. Then*

$$\sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] \leq 4TL(\mathbf{W}_{\frac{1}{\eta T}}^*) - 2\eta TL(\mathbf{W}^*) \\ + \left(\frac{2\tilde{b}T(\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c} + \frac{1}{2\eta} \right) \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2.$$

Proof. Multiplying T over both sides of (26) and using Lemma 15 we get

$$\sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] \\ \leq TL(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{\tilde{b}(\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c} \sum_{s=0}^{T-1} 1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2] + \frac{\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2}{2\eta} \\ \leq TL(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{\tilde{b}T(\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c} \left(\left(\frac{16e^2\eta^3\rho^2T^2}{n^2} + \frac{16e\eta^2T\rho}{n} \right) \sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] \right. \\ \left. + 2\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + 2\eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)] \right) + \frac{\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2}{2\eta}.$$

Condition (5) implies $m \geq (2\tilde{b}T(\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) \left(\frac{16e^2\eta^3\rho^2T^2}{n^2} + \frac{16e\eta^2T\rho}{n} \right))^{1/c}$

and $m \geq (2\tilde{b}T(\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2))^{1/c}$, there holds

$$\begin{aligned} & \sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] \\ & \leq 4TL(\mathbf{W}_{\frac{1}{\eta T}}^*) - 2\eta TL(\mathbf{W}^*) + \left(\frac{2\tilde{b}T(\sqrt{2\eta T c_0} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c} + \frac{1}{2\eta} \right) \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2, \end{aligned}$$

which completes the proof. \square

A.3 Proofs of Excess Risk Bounds

Proof of Theorem 3. According to Lemma 16 and noting that $\tilde{b} = \mathcal{O}(1)$, we know

$$\sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] = \mathcal{O} \left(TL(\mathbf{W}_{\frac{1}{\eta T}}^*) + \left(\frac{T(\sqrt{\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c} + \frac{1}{\eta} \right) \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 \right). \quad (27)$$

The upper bound of the generalization error can be controlled by plugging (27) back into Theorem 1

$$\begin{aligned} & \mathbb{E}[L(\mathbf{W}_T) - L_S(\mathbf{W}_T)] \\ & = \mathcal{O} \left(\left(\frac{\eta^2 \rho^2 T}{n^2} + \frac{\eta \rho}{n} \right) \sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] \right) \\ & = \mathcal{O} \left(\left(\frac{\eta^2 \rho^2 T^2}{n^2} + \frac{\eta T \rho}{n} \right) \left(L(\mathbf{W}_{\frac{1}{\eta T}}^*) + \left(\frac{\sqrt{\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2}{m^c} + \frac{1}{\eta T} \right) \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 \right) \right). \end{aligned} \quad (28)$$

The estimation of the optimization error is given by plugging (27) back into Theorem 2

$$\begin{aligned}
& \mathbb{E}[L_S(\mathbf{W}_T) - L_S(\mathbf{W}_{\frac{1}{\eta T}}^*) - \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2] \\
&= \mathcal{O}\left(\frac{(\sqrt{\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c} \left[\left(\frac{\eta^3 \rho^2 T^2}{n^2} + \frac{\eta^2 T \rho}{n}\right) \sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 \right. \right. \\
&\quad \left. \left. + \eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)] \right] \right) \\
&= \mathcal{O}\left(\frac{(\sqrt{\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c} \left(\frac{\eta^3 \rho^2 T^2}{n^2} + \frac{\eta^2 T \rho}{n}\right) \left(\frac{T(\sqrt{\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c} + \frac{1}{\eta}\right) \right. \\
&\quad \times \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + \frac{\eta T(\sqrt{\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c} \left(\frac{\eta^2 \rho^2 T^2}{n^2} + \frac{\eta T \rho}{n}\right) L(\mathbf{W}_{\frac{1}{\eta T}}^*) \\
&\quad \left. + \frac{(\sqrt{\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c} (\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + \eta T \Lambda_{\frac{1}{\eta T}})\right), \tag{29}
\end{aligned}$$

where we used the fact that $L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*) \leq \Lambda_{\frac{1}{\eta T}}$.

Combining (28) and (29) together and noting that the approximation error $\Lambda_{\frac{1}{\eta T}} = L(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 - L(\mathbf{W}^*)$ we get

$$\begin{aligned}
& \mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] \\
&= \left[\mathbb{E}[L(\mathbf{W}_T) - L_S(\mathbf{W}_T)] + \mathbb{E}\left[L_S(\mathbf{W}_T) - \left(L_S(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2\right)\right] \right] \\
&\quad + \left[L(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 - L(\mathbf{W}^*) \right] \\
&= \mathcal{O}\left(\frac{\eta T \rho}{n} \left(\frac{\eta \rho T}{n} + 1\right) \left(1 + \frac{\eta T(\sqrt{\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c}\right) \left[L(\mathbf{W}_{\frac{1}{\eta T}}^*) \right. \right. \\
&\quad \left. \left. + \left(\frac{1}{2\eta T} + \frac{\sqrt{\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2}{m^c}\right) \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 \right] \right. \\
&\quad \left. + \frac{(\sqrt{\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)}{m^c} (\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + \eta T \Lambda_{\frac{1}{\eta T}}) + \Lambda_{\frac{1}{\eta T}}\right).
\end{aligned}$$

Recalling that $\rho = \mathcal{O}(m^{1-2c})$. If $\eta T m^{1-2c} = \mathcal{O}(n)$ and $\eta T(\sqrt{\eta T} + \|\mathbf{W}^* - \mathbf{W}_0\|_2) = \mathcal{O}(m^c)$, there holds $\eta T \rho = \mathcal{O}(n)$ and $\eta T(\sqrt{\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2)/m^c = \mathcal{O}(1)$. Then

from the above bound we can get

$$\begin{aligned} & \mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] \\ &= \mathcal{O}\left(\frac{\eta T \rho}{n} \left[L(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 \right] + \frac{1}{\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + \Lambda_{\frac{1}{\eta T}}\right). \end{aligned}$$

Combining the above bound with the facts $L(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 = L(\mathbf{W}^*) + \Lambda_{\frac{1}{\eta T}}$ and $\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2 \leq \sqrt{\eta T \Lambda_{\frac{1}{\eta T}}}$ together we get

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left(\frac{\eta T \rho}{n} L(\mathbf{W}^*) + \Lambda_{\frac{1}{\eta T}}\right).$$

The proof is completed. \square

Proof of Corollary 4. Part (a). Case 1. From the definition of the approximation error $\Lambda_{\frac{1}{\eta T}}$, we know that $\Lambda_{\frac{1}{\eta T}} \leq \frac{1}{2\eta T} \|\mathbf{W}^* - \mathbf{W}_0\|_2^2$. Combining this with Theorem 3, we have

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left(\frac{\eta T \rho}{n} L(\mathbf{W}^*) + \frac{1}{\eta T} \|\mathbf{W}^* - \mathbf{W}_0\|_2^2\right).$$

Without loss of generality, we consider $\|\mathbf{W}_0\|_2$ as a constant. To obtain the excess risk rate, we discuss the following two cases: $2c + 6\mu - 3 > 0$ and $2c + 6\mu - 3 \leq 0$.

For the case $2c + 6\mu - 3 > 0$, to ensure conditions (4), (5) and $\eta T(\sqrt{\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) = \mathcal{O}(m^c)$ hold, we set $m \asymp (\eta T)^{\frac{3}{2c}}$ for this case. Then according to Theorem 3 and Assumption 3 we know

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left((\eta T)^{\frac{3-4c}{2c}} n^{-1} + (\eta T)^{\frac{3-6\mu-2c}{2c}}\right).$$

If $c < 3/4$, under the condition $\eta T \lesssim n^{\frac{c}{3-4c}}$ and $n^{\frac{c}{6\mu+2c-3}} \lesssim \eta T$, there holds $(\eta T)^{\frac{3-4c}{2c}} n^{-1} + (\eta T)^{\frac{3-6\mu-2c}{2c}} = \mathcal{O}(1/\sqrt{n})$. To ensure the above-mentioned conditions hold simultaneously, we further require $c + \mu \geq 1$ such that $n^{\frac{c}{3-4c}} \gtrsim n^{\frac{c}{6\mu+2c-3}}$. Therefore, if $c \in [1 - \mu, 3/4)$ and $c + 3\mu > 3/2$, we can obtain

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left(\frac{1}{\sqrt{n}}\right)$$

with $m \asymp (\eta T)^{\frac{3}{2c}}$ and $n^{\frac{c}{6\mu+2c-3}} \lesssim \eta T \lesssim n^{\frac{c}{3-4c}}$.

If $c \geq 3/4$, for any $\eta T \geq 1$ and $n \geq 1$, there holds $(\eta T)^{\frac{3-4c}{2c}} n^{-1} = \mathcal{O}(1/\sqrt{n})$.

Similar to before, if $\eta T \gtrsim n^{\frac{c}{6\mu+2c-3}}$, there holds $(\eta T)^{\frac{3-6\mu-2c}{2c}} = \mathcal{O}(1/\sqrt{n})$. Then we can obtain the excess population bound $\mathcal{O}(1/\sqrt{n})$ with $m \asymp (\eta T)^{\frac{3}{2c}}$ and $\eta T \gtrsim n^{\frac{c}{6\mu+2c-3}}$.

Case 2. For the case $2c + 6\mu - 3 \leq 0$, we can choose $m \asymp (\eta T)^{\frac{1}{c+\mu-1/2}}$ to ensure conditions (4), (5) and $\eta T(\sqrt{\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) = \mathcal{O}(m^c)$ hold. From Theorem 3 we know

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left((\eta T)^{\frac{1-2c+2\mu}{2c+2\mu-1}} n^{-1} + (\eta T)^{\frac{3-6\mu-2c}{2c+2\mu-1}}\right).$$

Note $3 - 6\mu - 2c \geq 0$ and $2c + 2\mu - 1 > 0$. Then the term $(\eta T)^{\frac{3-6\mu-2c}{2c+2\mu-1}}$ will not converge for any choice of ηT in this case. The proof of Part (a) is completed.

Part (b). Now, we consider the low noise case, i.e., $L(\mathbf{W}^*) = 0$. Combining the fact $\Lambda_{\frac{1}{\eta T}} \leq \frac{1}{2\eta T} \|\mathbf{W}^* - \mathbf{W}_0\|_2^2$, Assumption 3 and Theorem 3 with $L(\mathbf{W}^*) = 0$, we can get

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left(\frac{m^{1-2\mu}}{\eta T}\right).$$

Similar to part (a), we can set $m \asymp (\eta T)^{\frac{3}{2c}}$, $\eta T \gtrsim n^{\frac{2c}{6\mu+2c-3}}$ and obtain

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left(\frac{1}{n}\right).$$

We can check that this choice of m and ηT satisfies conditions (4) and (5). The proof is completed. \square

Remark 6. Several works (Charles and Papailiopoulos, 2018; Hardt et al., 2016; Lei and Ying, 2020b; Zhou et al., 2022; Mou et al., 2018) studied the stability behavior of stochastic gradient methods for non-convex losses, which can be applied to two-layer networks. Specifically, to obtain meaningful stability bounds, Hardt et al. (2016)

required a time-dependent step size $\eta_t = 1/t$, which is insufficient to get a good convergence rate for optimization error. Charles and Papailiopoulos (2018); Lei and Ying (2020b); Zhou et al. (2022) established generalization bounds by introducing the Polyak-Łojasiewicz condition, which depends on a problem-dependent number. This number might be large in practice and results in a worse generalization bound. It is hard to provide a direct comparison with their results since the learning settings are different.

B Proofs of Three-layer Neural Networks

B.1 Proofs of Generalization Bounds

For a matrix \mathbf{W} , let $\mathbf{W}_{s\cdot}$ and \mathbf{W}_{is} denote the s -th row and the (i, s) -th entry of \mathbf{W} , respectively.

Lemma 17 (Smoothness and Curvature). *Suppose Assumptions 1 and 2 hold. For any fixed $\mathbf{W} = (\mathbf{W}^{(1)}, \mathbf{W}^{(2)}) \in \mathbb{R}^{m \times d} \times \mathbb{R}^{m \times m}$ and any $z \in \mathcal{Z}$, there holds*

$$\lambda_{\max}(\nabla^2 \ell(\mathbf{W}; z)) \leq \rho_{\mathbf{W}},$$

$$\lambda_{\min}(\nabla^2 \ell(\mathbf{W}; z)) \geq -C_{\mathbf{W}} \left(\frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \|\mathbf{W}^{(2)}\|_2 + \sqrt{2c_0} \right),$$

where $\rho_{\mathbf{W}} = \frac{B_{\sigma'}^4 c_{\mathbf{x}}^2}{m^{4c-1}} \|\mathbf{W}^{(2)}\|_2^2 + \frac{B_{\sigma'}^2 B_{\sigma}^2}{m^{4c-2}} + C_{\mathbf{W}} \left(\frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \|\mathbf{W}^{(2)}\|_2 + \sqrt{2c_0} \right)$, and $C_{\mathbf{W}} = \frac{B_{\sigma'}^2 B_{\sigma''} c_{\mathbf{x}}^2}{m^{3c}} \|\mathbf{W}^{(2)}\|_2^2 + \left(\frac{B_{\sigma'} B_{\sigma''} c_{\mathbf{x}}^2}{m^{2c-\frac{1}{2}}} + \frac{2B_{\sigma''} B_{\sigma'} B_{\sigma} c_{\mathbf{x}}}{m^{3c-\frac{1}{2}}} \right) \|\mathbf{W}^{(2)}\|_2 + \frac{B_{\sigma''} B_{\sigma}^2}{m^{3c-1}} + \frac{2B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-\frac{1}{2}}}$.

Proof. Let $A_{\mathbf{W}^{(1)}} = [\sigma(\mathbf{W}_{1\cdot}^{(1)}; \mathbf{x}), \dots, \sigma(\mathbf{W}_{m\cdot}^{(1)}; \mathbf{x})]^\top \in \mathbb{R}^m$. Let $\tilde{\mathbf{w}} = \mathbf{w}^\top$. We first estimate the upper bound of $\|\nabla f_{\mathbf{W}}(\mathbf{x})\|_2$. Note that for any $k = 1, \dots, m$

$$\nabla_{\tilde{\mathbf{w}}_k^{(1)}} f_{\mathbf{W}}(\mathbf{x}) = \frac{\partial f_{\mathbf{W}}(\mathbf{x})}{\partial \tilde{\mathbf{w}}_k^{(1)}} = \frac{1}{m^{2c}} \sum_{i=1}^m a_i \sigma' \left(\frac{1}{m^c} \sum_{s=1}^m \mathbf{W}_{is}^{(2)} \sigma(\mathbf{W}_{s\cdot}^{(1)}; \mathbf{x}) \right) \mathbf{W}_{ik}^{(2)} \sigma'(\mathbf{W}_{k\cdot}^{(1)}; \mathbf{x}) \mathbf{x}$$

and

$$\nabla_{\tilde{\mathbf{w}}_k^{(2)}} f_{\mathbf{W}}(\mathbf{x}) = \frac{\partial f_{\mathbf{W}}(\mathbf{x})}{\partial \tilde{\mathbf{w}}_k^{(2)}} = \frac{a_k}{m^{2c}} \sigma' \left(\frac{1}{m^c} \sum_{s=1}^m \mathbf{W}_{ks}^{(2)} \sigma(\mathbf{W}_{s;\cdot}^{(1)} \mathbf{x}) \right) A_{\mathbf{W}^{(1)}}.$$

According to Assumptions 1 and 2, the upper bound of the gradient can be controlled as follows

$$\begin{aligned} \|\nabla f_{\mathbf{W}}(\mathbf{x})\|_2^2 &= \sum_{k=1}^m \left(\|\nabla_{\tilde{\mathbf{w}}_k^{(1)}} f_{\mathbf{W}}(\mathbf{x})\|_2^2 + \|\nabla_{\tilde{\mathbf{w}}_k^{(2)}} f_{\mathbf{W}}(\mathbf{x})\|_2^2 \right) \\ &\leq \frac{B_{\sigma'}^4 c_{\mathbf{x}}^2}{m^{4c-1}} \sum_{k=1}^m \sum_{i=1}^m |\mathbf{W}_{ik}^{(2)}|^2 + \frac{B_{\sigma'}^2 B_{\sigma}^2}{m^{4c-2}} \\ &\leq \frac{B_{\sigma'}^4 c_{\mathbf{x}}^2}{m^{4c-1}} \|\mathbf{W}^{(2)}\|_2^2 + \frac{B_{\sigma'}^2 B_{\sigma}^2}{m^{4c-2}}. \end{aligned} \quad (30)$$

For any $k, j \in [m]$, we know

$$\begin{aligned} \frac{\partial^2 f_{\mathbf{W}}(\mathbf{x})}{(\partial \tilde{\mathbf{w}}_k^{(1)})^2} &= \frac{1}{m^{3c}} \sum_{i=1}^m a_i \sigma'' \left(\frac{1}{m^c} \sum_{s=1}^m \mathbf{w}_{is}^2 \sigma(\mathbf{W}_{s;\cdot}^{(1)} \mathbf{x}) \right) (\mathbf{W}_{ik}^{(2)})^2 (\sigma'(\mathbf{W}_{k;\cdot}^{(1)} \mathbf{x}))^2 \mathbf{x} \mathbf{x}^\top \\ &\quad + \frac{1}{m^{2c}} \sum_{i=1}^m a_i \sigma' \left(\frac{1}{m^c} \sum_{s=1}^m \mathbf{W}_{is}^2 \sigma(\mathbf{W}_{s;\cdot}^{(1)} \mathbf{x}) \right) \mathbf{W}_{ik}^{(2)} \sigma''(\mathbf{W}_{k;\cdot}^{(1)} \mathbf{x}) \mathbf{x} \mathbf{x}^\top, \\ \frac{\partial^2 f_{\mathbf{W}}(\mathbf{x})}{(\partial \tilde{\mathbf{w}}_k^{(2)})^2} &= \frac{a_k}{m^{3c}} \sigma'' \left(\frac{1}{m^c} \sum_{s=1}^m \mathbf{W}_{ks}^{(2)} \sigma(\mathbf{W}_{s;\cdot}^{(1)} \mathbf{x}) \right) A_{\mathbf{W}^{(1)}} A_{\mathbf{W}^{(1)}}^\top \end{aligned}$$

and

$$\begin{aligned} \frac{\partial^2 f_{\mathbf{W}}(\mathbf{x})}{\partial \tilde{\mathbf{w}}_k^{(1)} \partial \tilde{\mathbf{w}}_j^{(2)}} &= \frac{1}{m^{3c}} a_j \sigma'' \left(\frac{1}{m^c} \sum_{s=1}^m \mathbf{W}_{js}^{(2)} \sigma(\mathbf{W}_{s;\cdot}^{(1)} \mathbf{x}) \right) \mathbf{W}_{jk}^{(2)} \sigma'(\mathbf{W}_{k;\cdot}^{(1)} \mathbf{x}) \mathbf{x} A_{\mathbf{W}^{(1)}}^\top \\ &\quad + \frac{1}{m^{2c}} a_j \sigma' \left(\frac{1}{m^c} \sum_{s=1}^m \mathbf{W}_{js}^{(2)} \sigma(\mathbf{W}_{s;\cdot}^{(1)} \mathbf{x}) \right) \sigma'(\mathbf{W}_{k;\cdot}^{(1)} \mathbf{x}) \mathbf{x} B_k, \end{aligned}$$

where $B_k \in \mathbb{R}^{1 \times m}$ with k -th element is 1 and others are 0. Let the vector $\mathbf{u} \in \mathbb{R}^{md+m^2}$ have unit norm $\|\mathbf{u}\|_2 = 1$ and be composed in a manner matching the parameter $\mathbf{W} = (\mathbf{W}^{(1)}, \mathbf{W}^{(2)})$ so that $\mathbf{u} = (\mathbf{u}^{(1)}, \mathbf{u}^{(2)})$, where $\mathbf{u}^{(1)} \in \mathbb{R}^{m \times d}$ and $\mathbf{u}^{(2)} \in \mathbb{R}^{m \times m}$ have been

vectorised in a row-major manner with $\mathbf{u}_k^{(1)} \in \mathbb{R}^{1 \times d}$ and $\mathbf{u}_k^{(2)} \in \mathbb{R}^{1 \times m}$. Then

$$\begin{aligned} \mathbf{u}^\top \nabla^2 f_{\mathbf{W}}(\mathbf{x}) \mathbf{u} &= \sum_{k=1}^m \left(\mathbf{u}_k^{(1)} \frac{\partial^2 f_{\mathbf{W}}(\mathbf{x})}{(\partial \tilde{\mathbf{w}}_k^{(1)})^2} (\mathbf{u}_k^{(1)})^\top + \mathbf{u}_k^{(2)} \frac{\partial^2 f_{\mathbf{W}}(\mathbf{x})}{(\partial \tilde{\mathbf{w}}_k^{(2)})^2} (\mathbf{u}_k^{(2)})^\top \right. \\ &\quad \left. + 2 \sum_{j=1}^m \mathbf{u}_k^{(1)} \frac{\partial^2 f_{\mathbf{W}}(\mathbf{x})}{\partial \tilde{\mathbf{w}}_k^{(1)} \partial \tilde{\mathbf{w}}_j^{(2)}} (\mathbf{u}_j^{(2)})^\top \right). \end{aligned} \quad (31)$$

We estimate the above three terms separately. Let $\mathbf{W}_{\cdot k}^{(2)}$ denote the k -th column of $\mathbf{W}^{(2)}$.

$$\begin{aligned} &\sum_{k=1}^m \mathbf{u}_k^{(1)} \frac{\partial^2 f_{\mathbf{W}}(\mathbf{x})}{(\partial \tilde{\mathbf{w}}_k^{(1)})^2} (\mathbf{u}_k^{(1)})^\top \\ &\leq \frac{B_{\sigma''} B_{\sigma'}^2}{m^{3c}} \sum_{k=1}^m \mathbf{u}_k^{(1)} \sum_{i=1}^m (\mathbf{W}_{ik}^{(2)})^2 \mathbf{x} \mathbf{x}^\top (\mathbf{u}_k^{(1)})^\top + \frac{B_{\sigma'} B_{\sigma''}}{m^{2c}} \sum_{k=1}^m \mathbf{u}_k^{(1)} \sum_{i=1}^m \mathbf{W}_{ik}^{(2)} \mathbf{x} \mathbf{x}^\top (\mathbf{u}_k^{(1)})^\top \\ &\leq \frac{B_{\sigma''} B_{\sigma'}^2}{m^{3c}} \sum_{i=1}^m \sum_{k=1}^m (\mathbf{W}_{ik}^{(2)})^2 (\mathbf{u}_k^{(1)} \mathbf{x})^2 + \frac{B_{\sigma'} B_{\sigma''}}{m^{2c}} \sqrt{\sum_{k=1}^m \sum_{i=1}^m (\mathbf{W}_{ik}^{(2)})^2} \sqrt{\sum_{k=1}^m (\mathbf{u}_k^{(1)} \mathbf{x} \mathbf{x}^\top (\mathbf{u}_k^{(1)})^\top)^2} \\ &\leq \frac{B_{\sigma''} B_{\sigma'}^2 c_{\mathbf{x}}^2}{m^{3c}} \sum_{k=1}^m \sum_{i=1}^m (\mathbf{W}_{ik}^{(2)})^2 \|\mathbf{u}_k^{(1)}\|_2^2 + \frac{B_{\sigma'} B_{\sigma''}}{m^{2c}} \sqrt{m \sum_{k=1}^m \sum_{i=1}^m (\mathbf{W}_{ik}^{(2)})^2} \sum_{k=1}^m (\mathbf{u}_k^{(1)} \mathbf{x} \mathbf{x}^\top (\mathbf{u}_k^{(1)})^\top) \\ &\leq \frac{B_{\sigma''} B_{\sigma'}^2 c_{\mathbf{x}}^2}{m^{3c}} \|\mathbf{W}^{(2)}\|_2^2 + \frac{B_{\sigma'} B_{\sigma''} c_{\mathbf{x}}^2}{m^{2c-\frac{1}{2}}} \|\mathbf{W}^{(2)}\|_2, \end{aligned} \quad (32)$$

where in the third inequality we used $\sum_{k=1}^m (\mathbf{u}_k^{(1)} \mathbf{x} \mathbf{x}^\top (\mathbf{u}_k^{(1)})^\top)^2 \leq (\sum_{k=1}^m \mathbf{u}_k^{(1)} \mathbf{x} \mathbf{x}^\top (\mathbf{u}_k^{(1)})^\top)^2$ and $(\sum_{i=1}^m \mathbf{W}_{ik}^{(2)})^2 \leq m \sum_{i=1}^m (\mathbf{W}_{ik}^{(2)})^2$, the last inequality follows from $\mathbf{u}_k^{(1)} \mathbf{x} \mathbf{x}^\top (\mathbf{u}_k^{(1)})^\top = (\mathbf{u}_k^{(1)} \mathbf{x})^2 \leq c_{\mathbf{x}}^2 \|\mathbf{u}_k^{(1)}\|_2^2$ and $\|\mathbf{u}_k^{(1)}\|_2^2 \leq 1$.

For the second term in (31), we control it by

$$\sum_{k=1}^m \mathbf{u}_k^{(2)} \frac{\partial^2 f_{\mathbf{W}}(\mathbf{x})}{(\partial \tilde{\mathbf{w}}_k^{(2)})^2} (\mathbf{u}_k^{(2)})^\top \leq \frac{B_{\sigma''}}{m^{3c}} \sum_{k=1}^m (\mathbf{u}_k^{(2)} A_{\mathbf{W}^{(1)}})^2 \leq \frac{B_{\sigma''} B_{\sigma}}{m^{3c-1}} \|\mathbf{u}^{(2)}\|_2^2 \leq \frac{B_{\sigma''} B_{\sigma}}{m^{3c-1}}. \quad (33)$$

Further, according to Cauchy-Schwarz inequality, we can get

$$\begin{aligned}
& \sum_{k=1}^m \sum_{j=1}^m \mathbf{u}_k^{(1)} \frac{\partial^2 f_{\mathbf{W}}(\mathbf{x})}{\partial \tilde{\mathbf{W}}_k^{(1)} \partial \tilde{\mathbf{W}}_j^{(2)}} (\mathbf{u}_j^{(2)})^\top \\
& \leq \frac{B_{\sigma''} B_{\sigma'} B_{\sigma}}{m^{3c}} \sum_{j=1}^m \mathbf{W}_{j;\cdot}^{(2)} \mathbf{u}^{(1)} \mathbf{x} \|\mathbf{u}_j^{(2)}\|_1 + \frac{B_{\sigma'}^2}{m^{2c}} \sum_{k=1}^m \mathbf{u}_k^{(1)} \mathbf{x} \left(\sum_{j=1}^m \mathbf{u}_{jk}^{(2)} \right) \\
& \leq \frac{B_{\sigma''} B_{\sigma'} B_{\sigma}}{m^{3c}} \sqrt{\sum_{j=1}^m \|\mathbf{u}_j^{(2)}\|_1^2} \sqrt{\sum_{j=1}^m \left(\mathbf{W}_{j;\cdot}^{(2)} \mathbf{u}^{(1)} \mathbf{x} \right)^2} + \frac{B_{\sigma'}^2}{m^{2c}} \sqrt{\sum_{k=1}^m \left(\mathbf{u}_k^{(1)} \mathbf{x} \right)^2} \sqrt{\sum_{k=1}^m \|\mathbf{u}_{\cdot k}^{(2)}\|_1^2} \\
& \leq \frac{B_{\sigma''} B_{\sigma'} B_{\sigma}}{m^{3c}} \sqrt{m \sum_{j=1}^m \|\mathbf{u}_j^{(2)}\|_2^2} \sqrt{\sum_{j=1}^m \|\mathbf{W}_{j;\cdot}^{(2)}\|_2^2 \|\mathbf{u}^{(1)} \mathbf{x}\|_2^2} + \frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c}} \sqrt{\sum_{k=1}^m \|\mathbf{u}_k^{(1)}\|_2^2} \sqrt{m \sum_{k=1}^m \|\mathbf{u}_{\cdot k}^{(2)}\|_2^2} \\
& \leq \frac{B_{\sigma''} B_{\sigma'} B_{\sigma} c_{\mathbf{x}}}{m^{3c-\frac{1}{2}}} \|\mathbf{u}^{(1)}\|_2 \|\mathbf{u}^{(2)}\|_2 \|\mathbf{W}^{(2)}\|_2 + \frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-\frac{1}{2}}} \|\mathbf{u}^{(1)}\|_2 \|\mathbf{u}^{(2)}\|_2 \\
& \leq \frac{B_{\sigma''} B_{\sigma'} B_{\sigma} c_{\mathbf{x}}}{m^{3c-\frac{1}{2}}} \|\mathbf{W}^{(2)}\|_2 + \frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-\frac{1}{2}}}, \tag{34}
\end{aligned}$$

where in the first equality we used $\sum_{k=1}^m \mathbf{u}_k^{(1)} \mathbf{W}_{jk}^{(2)} \mathbf{x} = \mathbf{W}_{j;\cdot}^{(2)} \mathbf{u}^{(1)} \mathbf{x}$, the second inequality follows from $\left(\sum_{j=1}^m \mathbf{u}_{jk}^{(2)} \right) \leq \|\mathbf{u}_{\cdot k}^{(2)}\|_1$. Here $\mathbf{u}_{\cdot k}^{(2)}$ denotes the k -th column of $\mathbf{u}^{(2)}$.

Plugging (32), (33) and (34) back into (31) we can get

$$\begin{aligned}
\|\nabla^2 f_{\mathbf{W}}(\mathbf{x})\|_{op} & \leq \frac{B_{\sigma'}^2 B_{\sigma''} c_{\mathbf{x}}^2}{m^{3c}} \|\mathbf{W}^{(2)}\|_2^2 + \left(\frac{B_{\sigma'} B_{\sigma''} c_{\mathbf{x}}^2}{m^{2c-\frac{1}{2}}} + \frac{2B_{\sigma''} B_{\sigma'} B_{\sigma} c_{\mathbf{x}}}{m^{3c-\frac{1}{2}}} \right) \|\mathbf{W}^{(2)}\|_2 \\
& \quad + \frac{B_{\sigma''} B_{\sigma}^2}{m^{3c-1}} + \frac{2B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-\frac{1}{2}}} := C_{\mathbf{W}}. \tag{35}
\end{aligned}$$

For any $\mathbf{W}, \tilde{\mathbf{W}}$, according to Assumptions 1 and 2 we can get

$$\begin{aligned}
& |f_{\mathbf{W}}(\mathbf{x}) - f_{\tilde{\mathbf{W}}}(\mathbf{x})| \\
& \leq \frac{1}{m^c} \sum_{k=1}^m \left| \sigma \left(\frac{1}{m^c} \sum_{s=1}^m \mathbf{W}_{ks}^{(2)} \sigma(\mathbf{W}_s^{(1)} \mathbf{x}) \right) - \sigma \left(\frac{1}{m^c} \sum_{s=1}^m \tilde{\mathbf{W}}_{ks}^{(2)} \sigma(\tilde{\mathbf{W}}_{s;\cdot}^{(1)} \mathbf{x}) \right) \right| \\
& \leq \frac{B_{\sigma'}}{m^{2c}} \sum_{k=1}^m \sum_{s=1}^m \left| \mathbf{W}_{ks}^{(2)} \sigma(\mathbf{W}_{s;\cdot}^{(1)} \mathbf{x}) - \tilde{\mathbf{W}}_{ks}^{(2)} \sigma(\mathbf{W}_{s;\cdot}^{(1)} \mathbf{x}) + \tilde{\mathbf{W}}_{ks}^{(2)} \sigma(\mathbf{W}_{s;\cdot}^{(1)} \mathbf{x}) - \tilde{\mathbf{W}}_{ks}^{(2)} \sigma(\tilde{\mathbf{W}}_{s;\cdot}^{(1)} \mathbf{x}) \right| \\
& \leq \frac{B_{\sigma'} B_{\sigma}}{m^{2c}} \sum_{k=1}^m \sum_{s=1}^m \left| \mathbf{W}_{ks}^{(2)} - \tilde{\mathbf{W}}_{ks}^{(2)} \right| + \frac{B_{\sigma'}^2}{m^{2c}} \sum_{k=1}^m \sum_{s=1}^m \left| \tilde{\mathbf{W}}_{ks}^{(2)} \right| \cdot \left| (\mathbf{W}_{s;\cdot}^{(1)} - \tilde{\mathbf{W}}_{s;\cdot}^{(1)}) \mathbf{x} \right| \\
& \leq \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \|\mathbf{W}^{(2)} - \tilde{\mathbf{W}}^{(2)}\|_2 + \frac{B_{\sigma'}^2 c_{\mathbf{x}} \|\tilde{\mathbf{W}}^{(2)}\|_{\infty}}{m^{2c-\frac{3}{2}}} \|\mathbf{W}^{(1)} - \tilde{\mathbf{W}}^{(1)}\|_2. \tag{36}
\end{aligned}$$

Since

$$\nabla^2 \ell(\mathbf{W}; z) = \nabla f_{\mathbf{W}}(\mathbf{x}) \nabla f_{\mathbf{W}}(\mathbf{x})^\top + \nabla^2 f_{\mathbf{W}}(\mathbf{x}) (f_{\mathbf{W}}(\mathbf{x}) - y). \quad (37)$$

Then for any $\mathbf{W} \in \mathbb{R}^{md+m^2}$, we can upper bound the maximal eigenvalue of the Hessian by combining (30), (35) and (36) with $\widetilde{\mathbf{W}} = \mathbf{0}$ together

$$\begin{aligned} & \lambda_{\max}(\nabla^2 \ell(\mathbf{W}; z)) \\ & \leq \|\nabla f_{\mathbf{W}}(\mathbf{x})\|_2^2 + \|\nabla^2 f_{\mathbf{W}}(\mathbf{x})\|_{op} |f_{\mathbf{W}}(\mathbf{x}) - y| \\ & \leq \frac{B_{\sigma'}^4 c_{\mathbf{x}}^2}{m^{4c-1}} \|\mathbf{W}^{(2)}\|_2^2 + \frac{B_{\sigma'}^2 B_{\sigma}^2}{m^{4c-2}} + C_{\mathbf{W}} (|f_{\mathbf{W}}(\mathbf{x}) - f_{\mathbf{0}}(\mathbf{x})| + |f_{\mathbf{0}}(\mathbf{x}) - y|) \\ & \leq \frac{B_{\sigma'}^4 c_{\mathbf{x}}^2}{m^{4c-1}} \|\mathbf{W}^{(2)}\|_2^2 + \frac{B_{\sigma'}^2 B_{\sigma}^2}{m^{4c-2}} + C_{\mathbf{W}} \left(\frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \|\mathbf{W}^{(2)}\|_2 + \sqrt{2\ell(\mathbf{0}; z)} \right) \\ & \leq \frac{B_{\sigma'}^4 c_{\mathbf{x}}^2}{m^{4c-1}} \|\mathbf{W}^{(2)}\|_2^2 + \frac{B_{\sigma'}^2 B_{\sigma}^2}{m^{4c-2}} + C_{\mathbf{W}} \left(\frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \|\mathbf{W}^{(2)}\|_2 + \sqrt{2c_0} \right). \end{aligned}$$

Note that $\nabla f_{\mathbf{W}}(\mathbf{x}) \nabla f_{\mathbf{W}}(\mathbf{x})^\top$ is positive semi-definite, then from (35), (37) and (36) with $\widetilde{\mathbf{W}} = \mathbf{0}$ we can get

$$\begin{aligned} \lambda_{\min}(\nabla^2 \ell(\mathbf{W}; z)) & \geq -\|\nabla^2 f_{\mathbf{W}}(\mathbf{x})\|_{op} |f_{\mathbf{W}}(\mathbf{x}) - y| \\ & \geq -C_{\mathbf{W}} (|f_{\mathbf{W}}(\mathbf{x}) - f_{\mathbf{0}}(\mathbf{x})| + |f_{\mathbf{0}}(\mathbf{x}) - y|) \\ & \geq -C_{\mathbf{W}} \left(\frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \|\mathbf{W}^{(2)}\|_2 + \sqrt{2\ell(\mathbf{W}_0; z)} \right) \\ & \geq -C_{\mathbf{W}} \left(\frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \|\mathbf{W}^{(2)}\|_2 + \sqrt{2c_0} \right). \end{aligned}$$

The proof is completed. □

Let

$$B_1 = \max \{ B_{\sigma'}^2 B_{\sigma''} c_{\mathbf{x}}^2, B_{\sigma'} B_{\sigma''} c_{\mathbf{x}}^2, 2B_{\sigma''} B_{\sigma} c_{\mathbf{x}}, B_{\sigma''} B_{\sigma}^2, 2B_{\sigma''}^2 c_{\mathbf{x}} \},$$

$$B_2 = \max \{ B_{\sigma'}^4 c_{\mathbf{x}}^2, B_{\sigma'}^2 B_{\sigma''}^2, B_{\sigma''} B_{\sigma}, \sqrt{2c_0} \}.$$

Lemma 18. *Suppose Assumptions 1 and 2 hold. Let $\{\mathbf{W}_t\}$ and $\{\mathbf{W}'_t\}$ be produced by GD iterates with T iterations based on S and S' , respectively. Let $C > 0$ be a constant. Assume $\eta \leq 1/(2\hat{\rho})$ and (6) hold. Then for any $c \in (1/2, 1]$ and any $t = 0, \dots, T$ there holds*

$$\|\mathbf{W}_t - \mathbf{W}_0\|_2 \leq \sqrt{2c_0\eta t}$$

and

$$\|\nabla\ell(\mathbf{W}_t; z) - \nabla\ell(\mathbf{W}'_t; z)\|_2 \leq \hat{\rho}\|\mathbf{W}_t - \mathbf{W}'_t\|_2.$$

Proof. We will prove by induction to show $\|\mathbf{W}_t - \mathbf{W}_0\|_2 \leq \eta tm^{2c-1}$. Further, we can show that $\rho_{\mathbf{W}} = \mathcal{O}(1)$ for any \mathbf{W} produced by GD iterates if m satisfies (6). Then by assuming $\eta \leq 1/(2\hat{\rho})$ we can prove that $\|\mathbf{W}_t - \mathbf{W}_0\|_2 \leq \sqrt{2c_0\eta t}$.

It's obvious that $\|\mathbf{W}_t - \mathbf{W}_0\|_2 \leq \eta tm^{2c-1}$ with $t = 0$ holds trivially. Assume $\|\mathbf{W}_t - \mathbf{W}_0\|_2 \leq \eta tm^{2c-1}$. According to the update rule (2) we know

$$\begin{aligned} & \|\mathbf{W}_{t+1} - \mathbf{W}_0\|_2 \\ & \leq \|\mathbf{W}_t - \mathbf{W}_0\|_2 + \eta \|\nabla L_S(\mathbf{W}_t)\|_2 \\ & \leq \|\mathbf{W}_t - \mathbf{W}_0\|_2 + \eta \max_{i \in [n]} \|\nabla f_{\mathbf{W}_t}(\mathbf{x}_i)\|_2 |f_{\mathbf{W}_t}(\mathbf{x}_i) - y_i| \\ & \leq \|\mathbf{W}_t - \mathbf{W}_0\|_2 + \eta \left(\frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-\frac{1}{2}}} (\|\mathbf{W}_t - \mathbf{W}_0\|_2 + \|\mathbf{W}_0\|_2) + \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \right) (|f_{\mathbf{W}_t}(\mathbf{x}_i) - f_{\mathbf{0}}(\mathbf{W}_t)| \\ & \quad + |f_{\mathbf{0}}(\mathbf{W}_t) - y_i|) \\ & \leq \|\mathbf{W}_t - \mathbf{W}_0\|_2 + \eta \left(\frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-\frac{1}{2}}} (\|\mathbf{W}_t - \mathbf{W}_0\|_2 + \|\mathbf{W}_0\|_2) + \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \right) \left(\frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} (\|\mathbf{W}_t - \mathbf{W}_0\|_2 \right. \\ & \quad \left. + \|\mathbf{W}_0\|_2) + \sqrt{2c_0} \right), \end{aligned} \tag{38}$$

where in the third inequality we used (30), the last inequality used (36) with $\widetilde{\mathbf{W}} = \mathbf{0}$. If

m is large enough such that

$$\begin{aligned} & \left(\frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-\frac{1}{2}}} (\|\mathbf{W}_t - \mathbf{W}_0\|_2 + \|\mathbf{W}_0\|_2) + \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \right) \left(\frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} (\|\mathbf{W}_t - \mathbf{W}_0\|_2 + \|\mathbf{W}_0\|_2) + \sqrt{2c_0} \right) \\ & \leq m^{2c-1}, \end{aligned} \quad (39)$$

then from (38) we know that $\|\mathbf{W}_{t+1} - \mathbf{W}_0\|_2 \leq \eta(t+1)m^{2c-1}$. The first part of the lemma can be proved. Now, we discuss the conditions on m such that (39) holds. Let $x = \|\mathbf{W}_t - \mathbf{W}_0\|_2 + \|\mathbf{W}_0\|_2$. To guarantee (39), it suffices that the following three inequalities hold

$$\frac{B_{\sigma'}^3 B_{\sigma} c_{\mathbf{x}}}{m^{4c-\frac{3}{2}}} x^2 \leq \frac{m^{2c-1}}{3}, \left(\frac{B_{\sigma'}^2 B_{\sigma}^2}{m^{4c-2}} + \frac{\sqrt{2c_0} B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-\frac{1}{2}}} \right) x \leq \frac{m^{2c-1}}{3}, \frac{\sqrt{2c_0} B_{\sigma'} B_{\sigma}}{m^{2c-1}} \leq \frac{m^{2c-1}}{3}. \quad (40)$$

It's easy to verify that (40) holds if $m \gtrsim \max\{(\eta T)^{\frac{4}{4c-1}}, (\eta T)^{\frac{1}{4c-2}}, \|\mathbf{W}_0\|_2^{\frac{1}{3c-\frac{5}{4}}}, \|\mathbf{W}_0\|_2^{\frac{1}{6c-3}}\}$, which can be ensured by (6). Hence, if $c \in (1/2, 1]$ and (6) holds, we have $\|\mathbf{W}_t - \mathbf{W}_0\|_2 \leq \eta t m^{2c-1}$ for all $t = 0, 1, \dots, T$.

Recall that

$$B_1 = \max \{ B_{\sigma'}^2 B_{\sigma''} c_{\mathbf{x}}^2, B_{\sigma'} B_{\sigma''} c_{\mathbf{x}}^2, 2B_{\sigma''} B_{\sigma} c_{\mathbf{x}}, B_{s\sigma''} B_{\sigma}^2, 2B_{s\sigma''}^2 c_{\mathbf{x}} \}$$

and

$$B_2 = \max \{ B_{\sigma'}^4 c_{\mathbf{x}}^2, B_{\sigma'}^2 B_{\sigma''}^2, B_{s\sigma''} B_{\sigma}, \sqrt{2c_0} \}.$$

Then from Lemma 17 we know

$$C_{\mathbf{W}} \leq B_1 \left(m^{-3c} \|\mathbf{W}^{(2)}\|_2^2 + (m^{\frac{1}{2}-2c} + m^{\frac{1}{2}-3c}) \|\mathbf{W}^{(2)}\|_2 + m^{1-3c} + m^{\frac{1}{2}-2c} \right)$$

and

$$\rho_{\mathbf{W}} \leq B_2 \left(m^{1-4c} \|\mathbf{W}\|_2^2 + m^{2-4c} + C_{\mathbf{W}} (m^{1-2c} \|\mathbf{W}\|_2 + 1) \right).$$

Note that (6) implies $m \gtrsim (\eta T)^4 + \|\mathbf{W}_0\|_2^{\frac{4}{8c-3}}$. By using $\|\mathbf{W}_t\|_2 \leq \eta t m^{2c-1} + \|\mathbf{W}_0\|_2$ we can verify that

$$\rho_{\mathbf{W}_t} \leq 4B_2(1 + 2B_1) := \hat{\rho} \quad \text{for any } t = 0, \dots, T.$$

Hence, we know that ℓ is $\hat{\rho}$ -smooth when the parameter space is the trajectory of GD.

Then for any $t = 0, \dots, T$ and any $\mathbf{W}_t, \mathbf{W}'_t$ produced by GD iterates, there holds

$$\|\nabla \ell(\mathbf{W}_t; z) - \nabla \ell(\mathbf{W}'_t; z)\|_2 \leq \hat{\rho} \|\mathbf{W}_t - \mathbf{W}'_t\|_2. \quad (41)$$

In addition, by the smoothness of ℓ we can get for any $j = 0, \dots, T - 1$

$$L_S(\mathbf{W}_{j+1}) \leq L_S(\mathbf{W}_j) - \eta \left(1 - \frac{\eta \hat{\rho}}{2}\right) \|\nabla L_S(\mathbf{W}_j)\|_2^2.$$

Rearranging and summing over j yields

$$\eta \left(1 - \frac{\eta \hat{\rho}}{2}\right) \sum_{j=0}^t \|\nabla L_S(\mathbf{W}_j)\|_2^2 \leq \sum_{j=0}^t L_S(\mathbf{W}_j) - L_S(\mathbf{W}_{j+1}) \leq L_S(\mathbf{W}_0).$$

Note that the update rule of GD (2) implies

$$\mathbf{W}_{t+1} = \mathbf{W}_0 - \eta \sum_{j=0}^t \nabla L_S(\mathbf{W}_j).$$

Combining the above two equations together and noting that $\eta \hat{\rho} \leq 1/2$, we obtain

$$\|\mathbf{W}_{t+1} - \mathbf{W}_0\|_2^2 = \eta^2 \left\| \sum_{j=0}^t \nabla L_S(\mathbf{W}_j) \right\|_2^2 \leq \eta^2 t \sum_{j=0}^t \|\nabla L_S(\mathbf{W}_j)\|_2^2 \leq 2\eta t L_S(\mathbf{W}_0) \leq 2c_0 \eta t.$$

The proof is completed. \square

The almost co-coercivity of the gradient operator for three-layer neural networks is given as follows. Recall that $S^{(i)} = \{z_1, \dots, z_{i-1}, z'_i, z_{i+1}, z_n\}$ is the set formed from S by replacing the i -th element with z'_i . For any \mathbf{W} , we have

$$L_{S \setminus i}(\mathbf{W}) = L_S(\mathbf{W}) - \frac{1}{n} \ell(\mathbf{W}; z_i) = L_{S^{(i)}}(\mathbf{W}) - \frac{1}{n} \ell(\mathbf{W}; z'_i).$$

Let $\{\mathbf{W}_t\}$ and $\{\mathbf{W}_t^{(i)}\}$ be the sequence produced by GD based on S and $S^{(i)}$, respectively. Let $C_{3,T} = 4B_1(2c_0\eta T m^{-3c} + (m^{\frac{1}{2}-2c} + m^{\frac{1}{2}-3c})\sqrt{2c_0\eta T} + m^{1-3c} + m^{\frac{1}{2}-2c})$.

Lemma 19. *Suppose $\eta \leq 1/(8\hat{\rho})$ where $\hat{\rho} = 4B_2(1 + 2B_1)$. Assume (6) holds. Then*

$$\begin{aligned} \langle \mathbf{W}_t - \mathbf{W}_t^{(i)}, \nabla L_{S^{(i)}}(\mathbf{W}_t) - \nabla L_{S^{(i)}}(\mathbf{W}_t^{(i)}) \rangle &\geq 2\eta \left(1 - 4\eta\hat{\rho}\right) \left\| \nabla L_{S^{(i)}}(\mathbf{W}_t) - \nabla L_{S^{(i)}}(\mathbf{W}_t^{(i)}) \right\|_2^2 \\ &\quad - \tilde{\epsilon}_t \left\| \mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta(\nabla L_{S^{(i)}}(\mathbf{W}_t) - \nabla L_{S^{(i)}}(\mathbf{W}_t^{(i)})) \right\|_2^2, \end{aligned}$$

where $\tilde{\epsilon}_t = C_{3,T} \left(B_3 \left(\frac{\sqrt{\eta T} + \|\mathbf{W}_0\|_2}{m^{2c-\frac{1}{2}}} + m^{1-2c} \right) \left((1 + \eta\hat{\rho}) \|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2 + 2(\sqrt{2c_0\eta T} + \|\mathbf{W}_0\|_2) + \sqrt{2c_0} \right) \right)$.

Proof. For any $\mathbf{W} \in \mathcal{B}(0, 2(\sqrt{2c_0\eta T} + \|\mathbf{W}_0\|_2))$, defining the following two functions

$$G_1(\mathbf{W}) = L_{S^{(i)}}(\mathbf{W}) - \langle \nabla L_{S^{(i)}}(\mathbf{W}_t^{(i)}), \mathbf{W} \rangle, \quad G_2(\mathbf{W}) = L_{S^{(i)}}(\mathbf{W}) - \langle \nabla L_{S^{(i)}}(\mathbf{W}_t), \mathbf{W} \rangle.$$

Note that

$$\begin{aligned} \langle \mathbf{W}_t - \mathbf{W}_t^{(i)}, \nabla L_{S^{(i)}}(\mathbf{W}_t) - \nabla L_{S^{(i)}}(\mathbf{W}_t^{(i)}) \rangle &= \\ &= (G_1(\mathbf{W}_t) - G_1(\mathbf{W}_t^{(i)})) + (G_2(\mathbf{W}_t^{(i)}) - G_2(\mathbf{W}_t)). \end{aligned} \quad (42)$$

Hence, it is enough to lower bound $G_1(\mathbf{W}_t) - G_1(\mathbf{W}_t^{(i)})$ and $G_2(\mathbf{W}_t^{(i)}) - G_2(\mathbf{W}_t)$.

Note that for any \mathbf{W}_t , Lemma 17 implies that $\|\mathbf{W}_t\|_2 \leq \sqrt{2c_0\eta t} + \|\mathbf{W}_0\|_2$. Then there holds

$$\begin{aligned} \left\| \mathbf{W}_t - \eta \nabla G_1(\mathbf{W}_t) \right\|_2 &\leq \left\| \mathbf{W}_t \right\|_2 + \eta \left\| \nabla L_{S^{(i)}}(\mathbf{W}_t) - \nabla L_{S^{(i)}}(\mathbf{W}_t^{(i)}) \right\|_2 \\ &\leq \left\| \mathbf{W}_t \right\|_2 + \eta\hat{\rho} \left\| \mathbf{W}_t - \mathbf{W}_t^{(i)} \right\|_2 \\ &\leq \sqrt{2c_0\eta t} + \|\mathbf{W}_0\|_2 + 2\eta\hat{\rho}(\sqrt{2c_0\eta t} + \|\mathbf{W}_0\|_2) \\ &\leq 2(\sqrt{2c_0\eta t} + \|\mathbf{W}_0\|_2), \end{aligned}$$

where in the last inequality we used $\eta\hat{\rho} \leq 1/8$. Similarly, we can show that $\|\mathbf{W}_t^{(i)} - \eta\nabla G_2(\mathbf{W}_t^{(i)})\|_2 \leq 2(\sqrt{2c_0\eta T} + \|\mathbf{W}_0\|_2)$. Hence, we know $\mathbf{W}_t - \eta\nabla G_1(\mathbf{W}_t) \in \mathcal{B}(0, 2(\sqrt{2c_0\eta T} + \|\mathbf{W}_0\|_2))$ and $\mathbf{W}_t^{(i)} - \eta\nabla G_2(\mathbf{W}_t^{(i)}) \in \mathcal{B}(0, 2(\sqrt{2c_0\eta T} + \|\mathbf{W}_0\|_2))$. On the other hand, similar to Lemma 18, we can show that $G_1(\mathbf{W})$ and $G_2(\mathbf{W})$ is $8\hat{\rho}$ -smooth for any $\mathbf{W} \in \mathcal{B}(0, 2(\sqrt{2c_0\eta T} + \|\mathbf{W}_0\|_2))$. Combining the above results, we can get

$$G_1(\mathbf{W}_t - \eta\nabla G_1(\mathbf{W}_t)) \leq G_1(\mathbf{W}_t) - \eta(1 - 4\eta\hat{\rho})\|\nabla G_1(\mathbf{W}_t)\|_2^2, \quad (43)$$

$$G_2(\mathbf{W}_t^{(i)} - \eta\nabla G_2(\mathbf{W}_t^{(i)})) \leq G_2(\mathbf{W}_t^{(i)}) - \eta(1 - 4\eta\hat{\rho})\|\nabla G_2(\mathbf{W}_t^{(i)})\|_2^2. \quad (44)$$

If we can further show that

$$G_1(\mathbf{W}_t - \eta\nabla G_1(\mathbf{W}_t)) \geq G_1(\mathbf{W}_t^{(i)}) - \frac{\tilde{\epsilon}_t}{2}\|\mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta\nabla G_1(\mathbf{W}_t)\|_2^2, \quad (45)$$

$$G_2(\mathbf{W}_t^{(i)} - \eta\nabla G_2(\mathbf{W}_t^{(i)})) \geq G_2(\mathbf{W}_t) - \frac{\tilde{\epsilon}_t}{2}\|\mathbf{W}_t^{(i)} - \mathbf{W}_t - \eta\nabla G_2(\mathbf{W}_t)\|_2^2, \quad (46)$$

with $\tilde{\epsilon}_t = C_{3,T}(B_3(\frac{\sqrt{\eta T} + \|\mathbf{W}_0\|_2}{m^{2c-\frac{1}{2}}} + m^{1-2c})((1 + \eta\hat{\rho})\|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2 + 2(\sqrt{2c_0\eta T} + \|\mathbf{W}_0\|_2) + \sqrt{2c_0}))$, where $C_{3,T} = 4B_1(2c_0\eta T m^{-3c} + (m^{\frac{1}{2}-2c} + m^{\frac{1}{2}-3c})\sqrt{2c_0\eta T} + m^{1-3c} + m^{\frac{1}{2}-2c})$. Then combining (43), (44), (45) and (46) together yields

$$G_1(\mathbf{W}_t) - G_1(\mathbf{W}_t^{(i)}) \geq \eta(1 - 4\eta\hat{\rho})\|\nabla G_1(\mathbf{W}_t)\|_2^2 - \frac{\tilde{\epsilon}_t}{2}\|\mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta\nabla G_1(\mathbf{W}_t)\|_2^2,$$

$$G_2(\mathbf{W}_t^{(i)}) - G_2(\mathbf{W}_t) \geq \eta(1 - 4\eta\hat{\rho})\|\nabla G_2(\mathbf{W}_t)\|_2^2 - \frac{\tilde{\epsilon}_t}{2}\|\mathbf{W}_t^{(i)} - \mathbf{W}_t - \eta\nabla G_2(\mathbf{W}_t)\|_2^2.$$

Plugging the above two inequalities back into (42) yields

$$\begin{aligned} & \langle \mathbf{W}_t - \mathbf{W}_t^{(i)}, \nabla L_{S^i}(\mathbf{W}_t) - \nabla L_{S^i}(\mathbf{W}_t^{(i)}) \rangle \\ &= G_1(\mathbf{W}_t) - G_1(\mathbf{W}_t^{(i)}) + G_2(\mathbf{W}_t^{(i)}) - G_2(\mathbf{W}_t) \\ &\geq 2\eta(1 - 4\eta\hat{\rho})\|\nabla L_{S^i}(\mathbf{W}_t) - \nabla L_{S^i}(\mathbf{W}_t^{(i)})\|_2^2 - \tilde{\epsilon}_t\|\mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta(\nabla L_{S^i}(\mathbf{W}_t) \\ &\quad - \nabla L_{S^i}(\mathbf{W}_t^{(i)}))\|_2^2. \end{aligned}$$

The desired result has been proved.

Now, we give the proof of (45) and (46). For $\alpha \in [0, 1]$, let $\mathbf{W}(\alpha) = \alpha \mathbf{W}_t + (1 - \alpha) \mathbf{W}_t^{(i)} - \alpha \eta (\nabla L_{S^i}(\mathbf{W}_t) - \nabla L_{S^i}(\mathbf{W}_t^{(i)}))$. For any $\alpha \in [0, 1]$, it's obvious that $\|\mathbf{W}(\alpha)\|_2 \leq 2(\sqrt{2c_0\eta t} + \|\mathbf{W}_0\|_2)$ by using Lemma 18. Combining this observation with (37) we can obtain

$$\begin{aligned}
& \lambda_{\min}(\nabla^2 L_{S^i}(\mathbf{W}(\alpha))) \\
& \geq -\max_{j \in [n]} \|\nabla^2 f_{\mathbf{W}(\alpha)}(\mathbf{x}_j)\|_{op} |f_{\mathbf{W}(\alpha)}(\mathbf{x}_j) - y_j| \\
& \geq -C_{\mathbf{W}(\alpha)} \max_{j \in [n]} \left(|f_{\mathbf{W}(\alpha)}(\mathbf{x}_j) - f_{\mathbf{W}_t^{(i)}}(\mathbf{x}_j)| + |f_{\mathbf{W}_t^{(i)}}(\mathbf{x}_j) - f_0(\mathbf{x}_j)| + |f_0(\mathbf{x}_j) - y_j| \right) \\
& \geq -C_{\mathbf{W}(\alpha)} \left(\left(\frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-\frac{1}{2}}} \|\mathbf{W}(\alpha)\|_2 + \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \right) (\|\mathbf{W}(\alpha) - \mathbf{W}_t^{(i)}\|_2 + \|\mathbf{W}_t^{(i)}\|_2) + \sqrt{2c_0} \right) \\
& \geq -C_{\mathbf{W}(\alpha)} \left(B_3 \left(\frac{\sqrt{\eta T} + \|\mathbf{W}_0\|_2}{m^{2c-\frac{1}{2}}} + m^{1-2c} \right) (\|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2 + \eta \|\nabla \ell(\mathbf{W}_t) - \nabla \ell(\mathbf{W}_t^{(i)})\|_2 \right. \\
& \quad \left. + \|\mathbf{W}_t^{(i)}\|_2) + \sqrt{2c_0} \right) \\
& \geq -\tilde{\epsilon}_t, \tag{47}
\end{aligned}$$

where the second inequality is due to (35), the third inequality is according to (30) and in the last inequality we used Lemma 18. Similarly, let $\widetilde{\mathbf{W}}(\alpha) = \alpha \mathbf{W}_t^{(i)} + (1 - \alpha) \mathbf{W}_t - \alpha \eta (\nabla L_{S^i}(\mathbf{W}_t^{(i)}) - \nabla L_{S^i}(\mathbf{W}_t))$, we can also control $\lambda_{\min}(\nabla^2 L_{S^i}(\widetilde{\mathbf{W}}(\alpha)))$ by $-\tilde{\epsilon}_t$.

Let $\Delta = \|\mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta (\nabla L_{S^i}(\mathbf{W}_t) - \nabla L_{S^i}(\mathbf{W}_t^{(i)}))\|_2^2$. We define

$$g_1(\alpha) = G_1(\mathbf{W}(\alpha)) + \frac{\tilde{\epsilon}_t \alpha^2}{2} \Delta, \quad g_2(\alpha) = G_2(\widetilde{\mathbf{W}}(\alpha)) + \frac{\tilde{\epsilon}_t \alpha^2}{2} \Delta.$$

From (47) we know that $g_1''(\alpha) \geq 0$ for any $\alpha \in [0, 1]$. Hence, g_1 is convex on $[0, 1]$.

Then from the convexity of g_1 we can get that

$$0 = g_1'(0) \leq g_1(1) - g_1(0) \leq G_1(\mathbf{W}_t - \eta \nabla G_1(\mathbf{W}_t)) + \frac{\tilde{\epsilon}_t}{2} \Delta - G_1(\mathbf{W}_t^{(i)}),$$

which completes the proof of (45). We can also show $g_2(\alpha)$ is convex on $[0, 1]$ and prove (46) in a similar way. The proof is completed. \square

Based on Lemma 17, Lemma 18 and Lemma 19, we can establish the following uniform stability bounds for three-layer neural networks.

Theorem 20 (Uniform Stability). *Suppose Assumptions 1 and 2 hold. Let S and $S^{(i)}$ be constructed in Definition 2. Let $\{\mathbf{W}_t\}$ and $\{\mathbf{W}_t^{(i)}\}$ be produced by (2) with $\eta \leq 1/(8\hat{\rho})$ based on S and $S^{(i)}$, respectively. Assume (6) holds. Then, for any $t \in [T]$, there holds*

$$\|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2 \leq \frac{2e\eta T \sqrt{2c_0\hat{\rho}(\hat{\rho}\eta T + 2)}}{n}.$$

Proof. Similar to (14), by the update rule $\mathbf{W}_{t+1} = \mathbf{W}_t - \eta \nabla L_S(\mathbf{W}_t)$ we know

$$\begin{aligned} \|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2^2 &\leq (1+p) \|\mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta(\nabla L_{S \setminus i}(\mathbf{W}_t) - \nabla L_{S \setminus i}(\mathbf{W}_t^{(i)}))\|_2^2 \\ &\quad + \frac{2\eta^2(1+1/p)}{n^2} (\|\nabla \ell(\mathbf{W}_t; z_i)\|_2^2 + \|\nabla \ell(\mathbf{W}_t^{(i)}; z_i')\|_2^2). \end{aligned} \quad (48)$$

From Lemma 19 we know

$$\begin{aligned} &\|\mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta(\nabla L_{S \setminus i}(\mathbf{W}_t) - \nabla L_{S \setminus i}(\mathbf{W}_t^{(i)}))\|_2^2 \\ &= \|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2^2 + \eta^2 \|\nabla L_{S \setminus i}(\mathbf{W}_t) - \nabla L_{S \setminus i}(\mathbf{W}_t^{(i)})\|_2^2 - 2\eta \langle \mathbf{W}_t - \mathbf{W}_t^{(i)}, \nabla L_{S \setminus i}(\mathbf{W}_t) \\ &\quad - \nabla L_{S \setminus i}(\mathbf{W}_t^{(i)}) \rangle \\ &\leq \|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2^2 + \eta^2 \|\nabla L_{S \setminus i}(\mathbf{W}_t) - \nabla L_{S \setminus i}(\mathbf{W}_t^{(i)})\|_2^2 - 4\eta^2 (1 - 4\eta\hat{\rho}) \|\nabla L_{S \setminus i}(\mathbf{W}_t) \\ &\quad - \nabla L_{S \setminus i}(\mathbf{W}_t^{(i)})\|_2^2 + 2\eta\tilde{\epsilon}_t \|\mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta(\nabla L_{S \setminus i}(\mathbf{W}_t) - \nabla L_{S \setminus i}(\mathbf{W}_t^{(i)}))\|_2^2. \end{aligned}$$

Note $4\eta\hat{\rho} \leq 1/2$ implies $1 - 4(1 - 4\eta\hat{\rho}) < 0$ and condition (6) ensures that $2\eta\tilde{\epsilon}_t < 1$ for any $t \in [T]$. Then from the above inequality we can get

$$(1 - 2\eta\tilde{\epsilon}_t) \|\mathbf{W}_t - \mathbf{W}_t^{(i)} - \eta(\nabla L_{S \setminus i}(\mathbf{W}_t) - \nabla L_{S \setminus i}(\mathbf{W}_t^{(i)}))\|_2^2 \leq \|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2^2. \quad (49)$$

Now, plugging (49) back into (48) we have

$$\begin{aligned} & \|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2^2 \\ & \leq \frac{1+p}{1-2\eta\tilde{\epsilon}_t} \|\mathbf{W}_t - \mathbf{W}_t^{(i)}\|_2^2 + \frac{2\eta^2(1+1/p)}{n^2} \left(\|\nabla\ell(\mathbf{W}_t; z_i)\|_2^2 + \|\nabla\ell(\mathbf{W}_t^{(i)}; z'_i)\|_2^2 \right). \end{aligned}$$

Applying the above inequality recursively and noting that $\mathbf{W}_0 = \mathbf{W}_0^{(i)}$ we get

$$\begin{aligned} & \|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2^2 \\ & \leq \frac{2\eta^2(1+1/p)}{n^2} \sum_{j=0}^t \left(\|\nabla\ell(\mathbf{W}_j; z_i)\|_2^2 + \|\nabla\ell(\mathbf{W}_j^{(i)}; z'_i)\|_2^2 \right) \prod_{\tilde{j}=j+1}^t \frac{1+p}{1-2\eta\tilde{\epsilon}_{\tilde{j}}} \\ & \leq \frac{2\eta^2(1+1/p)(1+p)^t}{n^2(1-2\eta\tilde{\epsilon}_j)^t} \sum_{j=0}^t \left(\|\nabla\ell(\mathbf{W}_j; z_i)\|_2^2 + \|\nabla\ell(\mathbf{W}_j^{(i)}; z'_i)\|_2^2 \right). \end{aligned}$$

Let $p = 1/t$ and note that $(1 + 1/t)^t \leq e$, we have

$$\begin{aligned} & \|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2^2 \\ & \leq \frac{2e\eta^2(1+t)}{n^2} \sum_{j=0}^t \left(\|\nabla\ell(\mathbf{W}_j; z_i)\|_2^2 + \|\nabla\ell(\mathbf{W}_j^{(i)}; z'_i)\|_2^2 \right) \prod_{\tilde{j}=j+1}^t \frac{1}{1-2\eta\tilde{\epsilon}_{\tilde{j}}}. \quad (50) \end{aligned}$$

According to Lemma 17 and Lemma 9, Assumption 2 and noting that $\|\mathbf{W}_j - \mathbf{W}_0\|_2 \leq \sqrt{2c_0\eta j}$ for any $j \leq t$, we know

$$\begin{aligned} \|\nabla\ell(\mathbf{W}_j; z)\|_2^2 & \leq 2\|\nabla\ell(\mathbf{W}_j; z) - \nabla\ell(\mathbf{W}_0; z)\|_2^2 + 2\|\nabla\ell(\mathbf{W}_0; z)\|_2^2 \\ & \leq 2\hat{\rho}^2\|\mathbf{W}_j - \mathbf{W}_0\|_2^2 + 4\hat{\rho}\ell(\mathbf{W}_0; z) \leq 4c_0\hat{\rho}(\hat{\rho}\eta j + 1). \end{aligned}$$

Similarly, we have

$$\|\nabla\ell(\mathbf{W}_j^{(i)}; z)\|_2^2 \leq 4c_0\hat{\rho}(\hat{\rho}\eta j + 1).$$

Combining the above three inequalities together, we get

$$\|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2^2 \leq \frac{8c_0e\eta^2(1+t)^2\hat{\rho}(\hat{\rho}\eta t + 2)}{n^2} \prod_{\tilde{j}=j+1}^t \frac{1}{1-2\eta\tilde{\epsilon}_{\tilde{j}}}.$$

Similar to the proof of Theorem 14, we can derive the following stability result by induction

$$\|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2 \leq \frac{2e\eta T \sqrt{2c_0 \hat{\rho}(\hat{\rho}\eta T + 2)}}{n}.$$

Here, the condition $\frac{1}{(1-2\eta\tilde{\epsilon}_j)^t} \leq \left(\frac{1}{1-1/(t+1)}\right)^t \leq e$ is ensured by condition (6), i.e.,

$$\begin{aligned} m \gtrsim & ((\eta T \mathcal{B}_T)^2 + \frac{(\eta T)^{\frac{7}{2}} \mathcal{B}_T}{n})^{\frac{1}{5c-\frac{1}{2}}} + ((\eta T)^{\frac{3}{2}} \mathcal{B}_T^2 + \frac{(\eta T)^3 \mathcal{B}_T}{n})^{\frac{1}{4c-1}} + ((\eta T)^2 \mathcal{B}_T + \frac{(\eta T)^{\frac{7}{2}}}{n})^{\frac{1}{5c-1}} \\ & + ((\eta T)^{\frac{3}{2}} \mathcal{B}_T + \frac{(\eta T)^3}{n})^{\frac{1}{4c-\frac{3}{2}}} \end{aligned}$$

with $\mathcal{B}_T = \sqrt{\eta T} + \|\mathbf{W}_0\|_2$. The proof is the same as Theorem 14. We omit it for simplicity. \square

We can combine Theorem 20 and Lemma 10 together to get the upper bound of the generalization error.

Proof of Theorem 5. The proof is similar to that of Theorem 1. From (50) we know that

$$\|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2^2 \leq \frac{4e^2 \eta^2 \hat{\rho}(1+t)}{n^2} \sum_{j=0}^t \left(\ell(\mathbf{W}_j; z_i) + \ell(\mathbf{W}_j^{(i)}; z_i') \right),$$

where we used the self-bounding property of the smooth loss (Lemma 9).

Then, taking an average over $i \in [n]$ and noting that $\mathbb{E}[\ell(\mathbf{W}_j; z_i)] = \mathbb{E}[\ell(\mathbf{W}_j^{(i)}; z_i')]$,

we have

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E} \|\mathbf{W}_{t+1} - \mathbf{W}_{t+1}^{(i)}\|_2^2 \leq \frac{8e^2 \eta^2 \hat{\rho}(1+t)}{n^2} \sum_{j=0}^t \mathbb{E}[L_S(\mathbf{W}_j)].$$

Combining the above stability bounds with Lemma 10 together and noting that $L_S(\mathbf{W}_t) \leq \frac{1}{t} \sum_{j=1}^{t-1} L_S(\mathbf{W}_j)$ (Richards and Kuzborskij, 2021), the desired result is obtained. \square

B.2 Proofs of Optimization Bounds

To show optimization error bounds, we first introduce the following lemma on the bound of GD iterates.

Lemma 21. *Suppose Assumptions 1 and 2 hold, and $\eta \leq 1/(8\hat{\rho})$. Assume (6) and (7)*

hold. Then for any $t \in [T]$, there holds

$$1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_t\|_2^2] \leq \left(\frac{16e^2\eta^3 t^2 \hat{\rho}^2}{n^2} + \frac{16e\eta^2 t \hat{\rho}}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] + 2\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + 2\eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)].$$

Proof. For any $\mathbf{W} \in \mathbb{R}^{md}$ and $\alpha \in [0, 1]$, define $\mathbf{W}(\alpha) := \mathbf{W}_t + \alpha(\mathbf{W} - \mathbf{W}_t)$. Similar to (47), according to Lemma 17 we can show that

$$\begin{aligned} & \lambda_{\min}(\nabla^2 L_S(\mathbf{W}(\alpha))) \\ & \geq -C_{\mathbf{W}(\alpha)} \left(\left(\frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-1/2}} \|\mathbf{W}(\alpha)\|_2 + \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} (\|\mathbf{W} - \mathbf{W}_t\|_2 + \|\mathbf{W}_t\|_2) + \sqrt{2c_0} \right), \right. \end{aligned}$$

where $C_{\mathbf{W}(\alpha)} = \frac{B_{\sigma'}^2 B_{\sigma''} c_{\mathbf{x}}^2}{m^{3c}} \|\mathbf{W}(\alpha)^{(2)}\|_2^2 + \left(\frac{B_{\sigma'} B_{\sigma''} c_{\mathbf{x}}^2}{m^{2c-\frac{1}{2}}} + \frac{2B_{\sigma''} B_{\sigma'} B_{\sigma} c_{\mathbf{x}}}{m^{3c-\frac{1}{2}}} \right) \|\mathbf{W}(\alpha)^{(2)}\|_2 + \frac{B_{\sigma''} B_{\sigma}^2}{m^{3c-1}} + \frac{2B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-\frac{1}{2}}}$. Let $\hat{C}_{\mathbf{W}} = 4B_1 (m^{-3c} (\|\mathbf{W}\|_2^2 + 4c_0 \eta T + 2\|\mathbf{W}_0\|_2^2) + m^{\frac{1}{2}-2c} (\|\mathbf{W}\|_2 + \sqrt{2c_0 \eta T} + \|\mathbf{W}_0\|_2))$. According to Lemma 18, we can verify that $C_{\mathbf{W}(\alpha)} \leq \hat{C}_{\mathbf{W}}$ for any $\alpha \in [0, 1]$.

Now, let

$$\begin{aligned} g(\alpha) := & L_S(\mathbf{W}(\alpha)) + \frac{\alpha^2 \hat{C}_{\mathbf{W}}}{2} \left(\left(\frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-1/2}} \|\mathbf{W}(\alpha)\|_2 + \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} (\|\mathbf{W} - \mathbf{W}_t\|_2 + \|\mathbf{W}_t\|_2) \right. \right. \\ & \left. \left. + \sqrt{2c_0} \right) \times (1 \vee \mathbb{E}[\|\mathbf{W} - \mathbf{W}_t\|_2^2]). \end{aligned}$$

It is obvious that $g(\alpha)$ is convex in $\alpha \in [0, 1]$. Similar to the proof of Lemma 15, by

convexity of g and smoothness of the loss we can show that

$$\begin{aligned}
& \frac{1}{t} \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] + \frac{\mathbb{E}[\|\mathbf{W}_t - \mathbf{W}\|_2^2]}{2\eta t} \\
& \leq \mathbb{E}[L_S(\mathbf{W})] + \frac{\mathbb{E}[\|\mathbf{W} - \mathbf{W}_0\|_2^2]}{2\eta t} + \frac{\hat{C}_{\mathbf{W}}}{2t} \sum_{s=0}^{t-1} \left(\left(\frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-1/2}} \|\mathbf{W}(\alpha)\|_2 + \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \right) \right. \\
& \quad \left. \times (\mathbb{E}[\|\mathbf{W} - \mathbf{W}_s\|_2] + \sqrt{2c_0 \eta T} + \|\mathbf{W}_0\|_2) + \sqrt{2c_0} \right) (1 \vee \mathbb{E}[\|\mathbf{W} - \mathbf{W}_s\|_2^2]). \quad (51)
\end{aligned}$$

Combining the above inequality with Theorem 5 and let $\mathbf{W} = \mathbf{W}_{\frac{1}{\eta T}}^*$, we have

$$\begin{aligned}
& \frac{\mathbb{E}[\|\mathbf{W}_t - \mathbf{W}_{\frac{1}{\eta T}}^*\|_2^2]}{2\eta t} \\
& \leq \frac{1}{t} \sum_{s=1}^{t-1} [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - \mathbb{E}[L(\mathbf{W}_s)]] + \frac{\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2}{2\eta t} + \left(\frac{4e^2 \eta^2 t \hat{\rho}^2}{n^2} + \frac{4e\eta \hat{\rho}}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] \\
& \quad + \frac{1}{2t} \sum_{s=0}^{t-1} \hat{C}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \left(\left(\frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-1/2}} \|\mathbf{W}(\alpha)\|_2 + \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \right) (\mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2] + \sqrt{2c_0 \eta T} + \|\mathbf{W}_0\|_2) \right. \\
& \quad \left. + \sqrt{2c_0} \right) \times (1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2]) \\
& \leq L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*) + \frac{\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2}{2\eta t} + \left(\frac{4e^2 \eta^2 t \hat{\rho}^2}{n^2} + \frac{4e\eta \hat{\rho}}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] \\
& \quad + \frac{1}{2t} \sum_{s=0}^{t-1} \hat{C}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \left(\left(\frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-1/2}} \|\mathbf{W}(\alpha)\|_2 + \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \right) (\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2 + \sqrt{2c_0 \eta T} + \|\mathbf{W}_0\|_2) \right. \\
& \quad \left. + \sqrt{2c_0} \right) \times (1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2]), \quad (52)
\end{aligned}$$

where in the second inequality we used $L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}_s) \leq L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)$ since $L(\mathbf{W}_s) \geq L(\mathbf{W}^*)$ for any $s \in [t-1]$.

According to Lemma 18 we can get

$$\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2 \leq \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2 + \|\mathbf{W}_s - \mathbf{W}_0\|_2 \leq \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2 + \sqrt{2c_0 \eta T}. \quad (53)$$

Then there holds

$$\begin{aligned}
& \left(\frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-1/2}} \|\mathbf{W}(\alpha)\|_2 + \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \right) (\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2 + \sqrt{2c_0 \eta T} + \|\mathbf{W}_0\|_2) + \sqrt{2c_0} \\
& \leq \left(\frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-1/2}} (2\sqrt{2c_0 \eta T} + 2\|\mathbf{W}_0\|_2 + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) + \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \right) (2\sqrt{2c_0 \eta T} + \|\mathbf{W}_0\|_2 \\
& \quad + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) + \sqrt{2c_0}. \tag{54}
\end{aligned}$$

Plugging the above inequality back into (52) and multiplying both sides by $2\eta t$ yields

$$\begin{aligned}
& \mathbb{E}[\|\mathbf{W}_t - \mathbf{W}_{\frac{1}{\eta T}}^*\|_2^2] \\
& \leq \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + \eta \hat{C}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \left(\left(\frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-1/2}} (2\sqrt{2c_0 \eta T} + \|\mathbf{W}_0\|_2 + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) + \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \right) \right. \\
& \quad \times (2\sqrt{2\eta T c_0} + \|\mathbf{W}_0\|_2 + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) + \sqrt{2c_0} \left. \right) \sum_{s=0}^{t-1} (1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2]) \\
& \quad + \left(\frac{8e^2 \eta^3 t^2 \hat{\rho}^2}{n^2} + \frac{8e\eta^2 t \hat{\rho}}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] + 2\eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)].
\end{aligned}$$

Let $x = \max_{s \in [T]} \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2] \vee 1$. Then the above inequality implies

$$\begin{aligned}
x & \leq \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + \eta T \hat{C}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \left(\left(\frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-1/2}} (2\sqrt{2c_0 \eta T} + \|\mathbf{W}_0\|_2 + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) + \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \right) \right. \\
& \quad \times (2\sqrt{2\eta T c_0} + \|\mathbf{W}_0\|_2 + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) + \sqrt{2c_0} \left. \right) x \\
& \quad + \left(\frac{8e^2 \eta^3 t^2 \hat{\rho}^2}{n^2} + \frac{8e\eta^2 t \hat{\rho}}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] + 2\eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)].
\end{aligned}$$

Note that condition (7) implies that $\eta T \hat{C}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \left(\left(\frac{B_{\sigma'}^2 c_{\mathbf{x}}}{m^{2c-1/2}} (2\sqrt{2c_0 \eta T} + \|\mathbf{W}_0\|_2 + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) + \frac{B_{\sigma'} B_{\sigma}}{m^{2c-1}} \right) (2\sqrt{2\eta T c_0} + \|\mathbf{W}_0\|_2 + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) + \sqrt{2c_0} \right) \leq \frac{1}{2}$. Then there holds

$$x \leq \left(\frac{16e^2 \eta^3 t^2 \hat{\rho}^2}{n^2} + \frac{16e\eta^2 t \hat{\rho}}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] + 2\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + 2\eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)].$$

It then follows that

$$\begin{aligned}
& 1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_t\|_2^2] \\
& \leq \left(\frac{16e^2 \eta^3 t^2 \hat{\rho}^2}{n^2} + \frac{16e\eta^2 t \hat{\rho}}{n} \right) \sum_{s=0}^{t-1} \mathbb{E}[L_S(\mathbf{W}_s)] + 2\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + 2\eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)].
\end{aligned}$$

This completes the proof. \square

Proof of Theorem 6. Combining (54) and (51) with $\mathbf{W} = \mathbf{W}_{\frac{1}{\eta T}}^*$ together yields

$$\begin{aligned}
& \frac{1}{T} \sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] \\
& \leq \mathbb{E}[L_S(\mathbf{W}_{\frac{1}{\eta T}}^*)] + \frac{\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2}{2\eta T} + \frac{\hat{C}_{\mathbf{W}_{\frac{1}{\eta T}}^*}}{2T} \sum_{s=0}^{T-1} \left(\left(\frac{B_{\sigma'}^2 c_x}{m^{2c-1/2}} (2\sqrt{2c_0\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) \right. \right. \\
& \quad \left. \left. + \|\mathbf{W}_0\|_2 + \frac{B_{\sigma'} B_\sigma}{m^{2c-1}} \right) (\mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2] + 2\sqrt{2c_0\eta T} + \|\mathbf{W}_0\|_2) + \sqrt{2c_0} \right) (1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2]) \\
& \leq \mathbb{E}[L_S(\mathbf{W}_{\frac{1}{\eta T}}^*)] + \frac{1}{2T} \hat{C}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \hat{B}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \sum_{s=0}^{T-1} 1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2] + \frac{\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2}{2\eta T}, \quad (55)
\end{aligned}$$

where $\hat{B}_{\mathbf{W}_{\frac{1}{\eta T}}^*} = \left(\frac{B_{\sigma'}^2 c_x}{m^{2c-1/2}} (2\sqrt{2c_0\eta T} + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) + \frac{B_{\sigma'} B_\sigma}{m^{2c-1}} \right) (2\sqrt{2\eta T} c_0 + \|\mathbf{W}_0\|_2 + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2) + \sqrt{2c_0}$ and in the last inequality we used (53).

By monotonic decrease of $\{L_S(\mathbf{W}_t)\}$ and Lemma 21, we further know

$$\begin{aligned}
& \mathbb{E}[L_S(\mathbf{W}_T)] \\
& \leq \mathbb{E}[L_S(\mathbf{W}_{\frac{1}{\eta T}}^*)] + \frac{\hat{C}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \hat{B}_{\mathbf{W}_{\frac{1}{\eta T}}^*}}{2T} \sum_{s=0}^{T-1} 1 \vee \mathbb{E}[\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_s\|_2^2] + \frac{\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2}{2\eta T} \\
& \leq \mathbb{E}[L_S(\mathbf{W}_{\frac{1}{\eta T}}^*)] + \hat{C}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \hat{B}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \left(\left(\frac{8e^2\eta^3 T^2 \hat{\rho}^2}{n^2} + \frac{8e\eta^2 T \hat{\rho}}{n} \right) \sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 \right. \\
& \quad \left. + \eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)] \right) + \frac{\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2}{2\eta T},
\end{aligned}$$

which completes the proof. \square

Recall that $\mathcal{C}_{T,n} = \eta T + \eta^3 T^2 / n^2$.

Lemma 22. *Suppose Assumptions 1 and 2 hold. Let $\{\mathbf{W}_t\}$ be produced by (2) with*

$\eta \leq 1/(8\hat{\rho})$. Assume (6) and (7) hold. Then

$$\sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] \leq 2TL(\mathbf{W}_{\frac{1}{\eta T}}^*) - \eta TL(\mathbf{W}^*) + \left(1 + \frac{1}{2\eta}\right) \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2.$$

Proof. Multiplying T over both sides of (55) and using Lemma 21 we get

$$\begin{aligned} \sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] &\leq TL(\mathbf{W}_{\frac{1}{\eta T}}^*) + \hat{C}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \hat{B}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \left(\left(\frac{8e^2 \eta^3 T^2 \hat{\rho}^2}{n^2} + \frac{8e \eta^2 T \hat{\rho}}{n} \right) \sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] \right. \\ &\quad \left. + \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + \eta T [L(\mathbf{W}_{\frac{1}{\eta T}}^*) - L(\mathbf{W}^*)] \right) + \frac{\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2}{2\eta}. \end{aligned} \quad (56)$$

Condition (7) implies $\hat{C}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \hat{B}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \left(\frac{4e^2 \eta^3 T^2 \hat{\rho}^2}{n^2} + \frac{4e \eta^2 T \hat{\rho}}{n} \right) \leq 1/2$ and $\hat{C}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \hat{B}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \leq 1$.

Then there holds

$$\sum_{s=0}^{T-1} \mathbb{E}[L_S(\mathbf{W}_s)] \leq 2TL(\mathbf{W}_{\frac{1}{\eta T}}^*) - \eta TL(\mathbf{W}^*) + \left(1 + \frac{1}{2\eta}\right) \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2,$$

which completes the proof. \square

B.3 Proofs of Excess Population Bounds

Proof of Theorem 7. Note $\hat{\rho} = \mathcal{O}(1)$. According to Theorem 5 and Lemma 22 we know

$$\mathbb{E}[L(\mathbf{W}_T) - L_S(\mathbf{W}_T)] = \mathcal{O}\left(\left(\frac{\eta^2 T^2}{n^2} + \frac{\eta T}{n}\right) \left(L(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{1}{\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2\right)\right). \quad (57)$$

The estimation of the optimization error is given by combining Lemma 22 and Theorem 6 together

$$\begin{aligned} &\mathbb{E}[L_S(\mathbf{W}_T) - L_S(\mathbf{W}_{\frac{1}{\eta T}}^*) - \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2] \\ &= \mathcal{O}\left(\left(\frac{\eta^2 T^2}{n^2} + \frac{\eta T}{n}\right) \left(L(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2\right) + \frac{1}{\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + \Lambda_{\frac{1}{\eta T}}\right), \end{aligned}$$

where we used the fact that $\hat{C}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \hat{B}_{\mathbf{W}_{\frac{1}{\eta T}}^*} \leq 1/\eta T$ implied by condition (7) and $L(\mathbf{W}_{\frac{1}{\eta T}}^*) -$

$$L(\mathbf{W}^*) \leq \Lambda_{\frac{1}{\eta T}}.$$

Combining the above two inequalities together we get

$$\begin{aligned}
& \mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] \\
&= \left[\mathbb{E}[L(\mathbf{W}_T) - L_S(\mathbf{W}_T)] + \mathbb{E}\left[L_S(\mathbf{W}_T) - \left(L_S(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 \right) \right] \right] \\
&\quad + \left[L(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 - L(\mathbf{W}^*) \right] \\
&= \mathcal{O}\left(\frac{\eta T}{n} \left(\frac{\eta T}{n} + 1 \right) \left[L(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 \right] + \frac{1}{\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + \Lambda_{\frac{1}{\eta T}} \right).
\end{aligned}$$

Since $n \gtrsim \eta T$, we further have

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left(\frac{\eta T}{n} \left[L(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 \right] + \frac{1}{\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 + \Lambda_{\frac{1}{\eta T}} \right).$$

Finally, note that $L(\mathbf{W}_{\frac{1}{\eta T}}^*) + \frac{1}{2\eta T} \|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 = L(\mathbf{W}^*) + \Lambda_{\frac{1}{\eta T}}$ and $\|\mathbf{W}_{\frac{1}{\eta T}}^* - \mathbf{W}_0\|_2^2 \leq \eta T \Lambda_{\frac{1}{\eta T}}$, we have

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left(\frac{\eta T}{n} L(\mathbf{W}^*) + \Lambda_{\frac{1}{\eta T}} \right).$$

The proof is completed. □

Proof of Corollary 8. Part (a). From the definition of the approximation error $\Lambda_{\frac{1}{\eta T}}$ and Theorem 7 we can get

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left(\frac{\eta T}{n} L(\mathbf{W}^*) + \frac{1}{\eta T} \|\mathbf{W}^* - \mathbf{W}_0\|_2^2 \right).$$

For the case $c \in [9/16, 1]$, to ensure conditions (6) and (7) hold, we choose $m \asymp (\eta T)^4$ for this case. Then according to Theorem 7 and Assumption 3, there holds

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left(\frac{\eta T}{n} + (\eta T)^{3-8\mu} \right).$$

Here, the condition $\mu \geq 1/2$ ensures that $3 - 8\mu < 0$. Hence, the bound will vanish as ηT tends to 0. Further, if $n^{\frac{1}{2(8\mu-3)}} \lesssim \eta T \lesssim \sqrt{n}$ (the existence of ηT is ensured by

$\mu \geq 1/2$), then there holds $\eta T/n = \mathcal{O}(n^{-1/2})$ and $(\eta T)^{3-8\mu} = \mathcal{O}(n^{-1/2})$. That is

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left(\frac{1}{\sqrt{n}}\right).$$

For the case $c \in (1/2, 9/16)$, we choose $m \asymp (\eta T)^{\frac{1}{4c-2}}$, and $n^{\frac{2c-1}{2\mu+4c-3}} \lesssim \eta T \lesssim \sqrt{n}$.

From Theorem 7 and Assumption 3 we have

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left(\frac{1}{\sqrt{n}}\right).$$

The first part of the theorem is proved.

Part (b). For the case $c \in [9/16, 1]$, by choosing $m \asymp (\eta T)^4$ and $\eta T \gtrsim n^{\frac{1}{8\mu-3}}$, from Theorem 7 and Assumption 3 we have

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left(\frac{1}{n}\right).$$

For the case $c \in (1/2, 9/16)$, by choosing $m \asymp (\eta T)^{\frac{1}{4c-2}}$ and $\eta T \gtrsim n^{\frac{4c-2}{4c+2\mu-3}}$, from Theorem 7 and Assumption 3 we have

$$\mathbb{E}[L(\mathbf{W}_T) - L(\mathbf{W}^*)] = \mathcal{O}\left(\frac{1}{n}\right).$$

The proof is completed. □

B.4 More Discussion on Related Works

Richards and Rabbat (2021) derived a lower bound on the minimum eigenvalue of the Hessian for the three-layer NN, where the first layer activation is linear, and the second activation is smooth. They proved the weak convexity of the empirical risk scales with $m^{\frac{1}{2}-2c}$ when optimizing the first and third layers of weights with Lipschitz and convex losses. We train the first and the second layers of the network with general smooth activation functions for both layers. Our result shows that the weak convexity of the least square loss scales with $m^{\frac{1}{2}-2c}$.