Linear Convergence of Gradient Descent for Finite Width Over-parametrized Linear Networks with General Initialization

Ziqing Xu
Johns Hopkins University

Hancheng Min
Johns Hopkins University

Salma Tarmoun
Johns Hopkins University

Enrique Mallada
Johns Hopkins University

René Vidal
University of Pennsylvania

Abstract

Recent theoretical analyses of the convergence of gradient descent (GD) to a global minimum for over-parametrized neural networks make strong assumptions on the step size (infinitesimal), the hidden-layer width (infinite), or the initialization (spectral, balanced). In this work, we relax these assumptions and derive a linear convergence rate for two-layer linear networks trained using GD on the squared loss in the case of finite step size, finite width and general initialization. Despite the generality of our analysis, our rate estimates are significantly tighter than those of prior work. Moreover, we provide a time-varying step size rule that monotonically improves the convergence rate as the loss function decreases to zero. Numerical experiments validate our findings.

1 INTRODUCTION

The empirical success of neural networks on a wide variety of applications, such as natural language processing Vaswani et al. [2017, 2018], computer vision He et al. [2015]; Minaee et al. [2021] and decision making Silver et al. [2016]; Vo et al. [2019], has motivated significant research on understanding theoretically why neural networks work so well in practice. One interesting and puzzling phenomenon is that over-parametrized neural networks trained with gradient descent (GD) enjoy fast convergence even if their loss landscape is non-convex. Much of the recent work in this area has focused on deriving convergence rates for over-parametrized networks. However, existing results

Table 1: Summary of prior work and our contributions.

	step size	width	initialization
Jacot et al. [2018]; Du et al. [2018b]; Lee et al. [2019]; Liu et al. [2022]; Oymak and Soltanolkotabi [2020]	**finite**	very large	sufficiently large
Mei et al. [2018]; Chizat and Bach [2018]; Ding et al. [2022]; Sirignano and Spiliopoulos [2020]	infinitesimal	infinite	**general**
Saxe et al. [2013]; Gidel et al. [2019]; Tarmoun et al. [2021]	infinitesimal	**finite**	spectral
Tarmoun et al. [2021]; Min et al. [2022]	infinitesimal	**finite**	**general**
Arora et al. [2018]; Du et al. [2018a]; Nguegnang et al. [2021]	**finite**	**finite**	large margin and small imbalance
This work	**finite**	**finite**	**general**

require stringent assumptions on the step size (infinitesimally small), the hidden-layer width (infinitely large), or the initialization (spectral, balanced).

Prior work. One line of work Jacot et al. [2018]; Du et al. [2018b]; Lee et al. [2019]; Liu et al. [2022] studies the convergence of GD when the scale of the initialization and the network width are sufficiently large. Under these assumptions, the network weights remain close to their initialization during training, and one can show that GD converges linearly to a global minimum. However, Chizat et al. [2019]; Chen et al. [2022] show that this "lazy training" regime is unrealistic in practice as it limits feature learning. A convergence analysis beyond the so-called lazy regime can be undertaken in the (mean-field) limit of infinitely wide networks Mei et al. [2018]; Rotskoff and Vanden-Eijnden [2018b,a]; Chizat and Bach [2018]; Sirignano and Spiliopoulos [2020]; Ding et al. [2022], where suitable assumptions on the initialization and step size make GD become a Wasserstein flow; a partial differential equation

commonly appearing in optimal transport theory. However, while such analysis can guarantee convergence to the global optimum for a wider range of initializations, it still imposes strong assumptions on the network width (infinite) and step size (infinitesimal).

Another line of work studies the convergence of gradient-based algorithms for over-parametrized networks with finite width. In this finite-width setting, the vast majority of existing results considers linear networks trained using gradient flow (GF). GF can be seen as GD with infinitesimal step size, but its dynamics in this setting are generally easier to analyze. For example, Saxe et al. [2013]; Gidel et al. [2019]; Tarmoun et al. [2021] show that under spectral initialization the dynamics of GF decouple into several scalar dynamics, which allows them to derive a linear convergence rate. For non-spectral initialization, Tarmoun et al. [2021]; Min et al. [2022] show that a large *imbalance* or large *margin* of the initialization can lead to faster convergence of GF, significantly extending the range of initializations from which linear convergence of GF is guaranteed. However, such results require infinitesimal step size. For finite step size, Arora et al. [2018]; Du et al. [2018a]; Nguegnang et al. [2021] prove linear convergence of GD when there is sufficient margin at initialization and the imbalance is small. However, such assumptions rarely hold in practice since commonly used random initializations have a large imbalance.

Paper contributions. In this work, we derive a linear convergence rate for GD in the case of over-parametrized, finite-width, two-layer linear networks with general initialization. Our analysis can be seen as a natural extension of recent results for GF, which cover finite width and imbalanced initializations. However, a key challenge in the case of GD is that quantities such as *imbalance*, which are preserved by GF, are no longer preserved by GD. To address this challenge, we derive quantities that effectively bound the deviation of the discrete dynamics from the continuous dynamics as a function of the step size, thus ensuring sufficient control (via upper and lower bounds) of the level of imbalance throughout training. This leads to a convergence rate that naturally depends on the step size, as well as other quantities, such as the current loss value. Moreover, the dependency of the rate on the step size is a low-degree polynomial, which allows us to easily compute an optimal step size at each iteration of training. Furthermore, we prove that the resulting time-varying step size is *lower-bounded by* the optimal rate of GD for the non-overparametrized problem. Finally, our numerical results show that, despite the generality of our analysis, the step size we derive leads to faster convergence and our Theorem in §3 admits a wider range of step sizes than in Du et al. [2018a]; Arora et al. [2018]. We provide the code to reproduce the simulations at `https://github.com/simonxu97/aistats2023_overparametrization`.

Notation. We use lower case letters a to denote a scalar, and capital letters A and $A^\top$ to denote a matrix and its transpose. We use $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$ to denote the largest and smallest eigenvalues of A, $\sigma_{\max}(A)$ and $\sigma_{\min}(A)$ to denote its largest and smallest singular values, $\|A\|_F$ and $\|A\|_2$ to denote its Frobenius and spectral norms, and $A[i,j]$ to denote its (i,j)-th element. Given two matrices $A \in \mathbb{R}^{n \times m}$ and $B \in \mathbb{R}^{k \times l}$, it will be convenient to use either $\begin{pmatrix} A \\ B \end{pmatrix}$ or (A, B) to represent an element in the product space $\mathbb{R}^{n \times m} \times \mathbb{R}^{k \times l}$, irrespectively of the dimensions. For a function $f(Z)$, we use $\nabla f(Z) := \frac{\partial}{\partial Z} f(Z)$ to denote its gradient, and whenever Z depends on an independent variable t, we use $f(t) := f(Z(t))$ and $\dot{Z}(t) = \frac{d}{dt} Z(t)$, dropping the dependence on t when it is implicit from the context, e.g., $\dot{Z} = \frac{d}{dt} Z$. Finally, we use $\mathcal{N}(\mu, \sigma^2)$ to denote a normal distribution with mean μ and variance σ^2.

2 CONVERGENCE OF GRADIENT FLOW FOR TWO-LAYER LINEAR NETWORKS

In this section, we first consider a linear regression problem and its over-parametrized version, which is equivalent to training a two-layer linear neural network. We then summarize the convergence results for GF in Min et al. [2022], which constitute the starting point of our work. Throughout this section, we thus consider a continuous time $t \in \mathbb{R}$.

Given N training samples $(x_i, y_i)_{i=1}^N$, where $x_i \in \mathbb{R}^n$, $y_i \in \mathbb{R}^m$, we consider the following linear regression problem

$$\min_W \ell(W) = \frac{1}{2}\|Y - XW\|_F^2, \qquad (1)$$

where $W \in \mathbb{R}^{n \times m}$, $X = [x_1, \cdots, x_N]^\top \in \mathbb{R}^{N \times n}$ and $Y = [y_1, \cdots, y_N]^\top \in \mathbb{R}^{N \times m}$. We are interested in solving the optimization problem in equation 1 by solving the following over-parametrized problem

$$\min_{W_1, W_2} L(W_1, W_2) = \frac{1}{2}\|Y - XW_1 W_2\|_F^2, \qquad (2)$$

where $W_1 \in \mathbb{R}^{n \times h}$, $W_2 \in \mathbb{R}^{h \times m}$. This over-parametrized problem corresponds to training a two-layer linear neural network with n inputs, h hidden neurons, m outputs, and weight matrices W_1 and W_2.

To simplify the exposition, we consider the above problems in the under-determined case, i.e., $N \leq n$. We assume that the input data matrix X is full rank, i.e., $\text{rank}(X) = N$.[1] We also assume that $h \geq \min\{n, m\}$. These assumptions imply that the minimum of both problems is zero, i.e., $\min_W \ell(W) = 0$ and $L^* := \min_{W_1, W_2} L(W_1, W_2) = 0$.

---

[1]When X is rank deficient, one can reformulate the problem into one with full-rank input data matrix (see Appendix A for details).

We note, however, that our results generalize the case $N > n$, by properly accounting for a non-zero L^*.

Convergence under GF. Let us consider solving equation 2 via GF

$$\begin{pmatrix} \dot{W}_1 \\ \dot{W}_2 \end{pmatrix} = -\nabla L(W_1, W_2) = - \begin{pmatrix} \nabla \ell(W) W_2^\top \\ W_1^\top \nabla \ell(W) \end{pmatrix}, \quad (3)$$

where $\nabla \ell(W) = X^\top(Y - XW)$. Notice that there exists a linear operator $\gamma(\cdot; W_1, W_2) : \mathbb{R}^{n \times m} \to \mathbb{R}^{n \times h} \times \mathbb{R}^{h \times m}$,

$$\gamma(\nabla \ell(W); W_1, W_2) := \begin{pmatrix} \nabla \ell(W) W_2^\top \\ W_1^\top \nabla \ell(W) \end{pmatrix}, \quad (4)$$

which depends on W_1, W_2, that maps the gradient of the loss $\nabla \ell(W) \in \mathbb{R}^{n \times m}$ to the gradient of the over-parametrized loss $\nabla L(W_1, W_2) \in \mathbb{R}^{n \times h} \times \mathbb{R}^{h \times m}$.

Then, one can show that the evolution of L under GF is

$$\begin{aligned} \dot{L}(W_1, W_2) \\ = \left\langle \frac{\partial L}{\partial W_1}(W_1, W_2), \dot{W}_1 \right\rangle + \left\langle \frac{\partial L}{\partial W_2}(W_1, W_2), \dot{W}_2 \right\rangle \\ = - \langle \gamma(\nabla \ell(W); W_1, W_2), \gamma(\nabla \ell(W); W_1, W_2) \rangle \\ = - \langle \nabla \ell(W), \gamma^* \circ \gamma(\nabla \ell(W); W_1, W_2) \rangle, \end{aligned} \quad (5)$$

where $\gamma^*(\cdot; W_1, W_2)$ is the adjoint of $\gamma(\cdot; W_1, W_2)$. Therefore, the dynamics of L are defined by the following positive semi-definite Hermitian linear operator on $\nabla \ell(W)$:

$$\begin{aligned} \tau(\nabla \ell(W); W_1, W_2) := \gamma^* \circ \gamma(\nabla \ell(W); W_1, W_2) \\ = \nabla \ell(W) W_2^\top W_2 + W_1 W_1^\top \nabla \ell(W). \end{aligned} \quad (6)$$

Then, from equation 5 and the min-max principle of Hermitian operators, we have

$$\dot{L}(t) = -\langle \nabla \ell(t), \tau_t(\nabla \ell(t)) \rangle \le -\lambda_{\min}(\tau_t) \|\nabla \ell(t)\|_F^2, \quad (7)$$

where for simplicity we use $\ell(t)$, $L(t)$ and $\tau_t(\nabla \ell(t))$, resp., as a shorthand for $\ell(W(t))$, $L(W_1(t), W_2(t))$ and $\tau(\nabla \ell(W(t)); W_1(t), W_2(t))$. Similarly, we use $\lambda_{\min}(\tau_t)$ and $\lambda_{\max}(\tau_t)$ as a shorthand for $\lambda_{\min}(\tau(\cdot; W_1(t), W_2(t)))$ and $\lambda_{\max}(\tau(\cdot; W_1(t), W_2(t)))$, respectively.

The core contribution of Min et al. [2022] is to provide a lower bound on $\lambda_{\min}(\tau_t)$ using two quantities: *imbalance*

$$D(t) = W_1^\top(t) W_1(t) - W_2(t) W_2(t)^\top, \quad (8)$$

and *product* $W(t) = W_1(t) W_2(t)$. Specifically, they show there exists an non-negative function $\alpha(D, \sigma_{\min}(W))$ that depends on imbalance and product, such that for all $t \ge 0$,

$$\lambda_{\min}(\tau_t) \ge \alpha(D(t), \sigma_{\min}(W(t))). \quad (9)$$

To find a uniform lower bound on $\alpha(D(t), \sigma_{\min}(W(t)))$ for all $t \ge 0$, they exploit the fact that the imbalance matrix

remains constant along the trajectories of GF Arora et al. [2018]; Du et al. [2018a], i.e., $\dot{D} \equiv 0$ so that $D(t) = D(0)$. As for the product, Min et al. [2022] show (from the fact that the loss $L(t)$ is non-increasing) that

$$\sigma_{\min}(W(t)) \ge p_1(:= margin), \quad (10)$$

where p_1 is defined in Table 2.

Therefore, we can replace the imbalance $D(t)$ in equation 9 by its initial value $D(0)$. Moreover, it can be shown that $\alpha(D, \sigma)$ is a non-decreasing function of the second argument σ, allowing us to use equations 9 and 10 to show that

$$\lambda_{\min}(\tau_t) \ge \alpha(D(t), p_1) = \alpha(D(0), p_1) := \alpha_0, \quad (11)$$

where the expression for α_0 is shown in Table 2. Observe that equation 11 yields a uniform lower bound on $\lambda_{\min}(\tau_t)$. Combining equation 11 with the fact that $\ell(t)$ satisfies the PL condition $\frac{1}{2} \|\nabla \ell(t)\|_F^2 \ge \mu \ell(t)$ with $\mu = \sigma_{\min}^2(X) > 0$, we show that equation 7 can be further upper-bounded by:

$$\begin{aligned} \dot{L}(t) \le -\lambda_{\min}(\tau_t) \|\nabla \ell(t)\|_F^2 \le -\alpha_0 \|\nabla \ell(t)\|_F^2 \\ \le -2\mu\alpha_0 \ell(t) = -2\mu\alpha_0 L(t), \end{aligned} \quad (12)$$

where the third inequality follows from the PL condition. Moreover, if $\alpha_0 > 0$, it follows from Grönwall's inequality that $L(t) \le \exp(-2\mu\alpha_0 t) L(0)$, showing that GF converges exponentially with a rate $2\mu\alpha_0$.

As discussed in the introduction, the imbalance matrix $D(t)$ measures the difference of the weights in the two layers, while the margin p_1 depends on the initial error $\|Y - XW_1(0)W_2(0)\|_F$ (the smaller the error, the larger the margin). Min et al. [2022] show that $\alpha_0 > 0$ when there is either 1) sufficient imbalance $\underline{\Delta} > 0$ or 2) sufficient margin $p_1 > 0$, where $\underline{\Delta}$ is defined in Table 2. Moreover, a larger imbalance (as measured by $\underline{\Delta}$) or a larger margin p_1 improves the rate of convergence α_0. In summary, the convergence of GF is completely determined by the initialization $W_1(0), W_2(0)$, and convergence is guaranteed when the initialization satisfies $\alpha_0 > 0$, which is achieved by either being imbalanced or having sufficient margin.

3 CONVERGENCE OF GRADIENT DESCENT FOR TWO-LAYER LINEAR NETWORKS

In this section, we analyze the convergence of GD for over-parametrized two-layer linear networks. We start in §3.1 by highlighting the challenges of analyzing over-parametrized GD when compared to (1) the standard GD algorithm applied to $\ell(W)$ and (2) the GF algorithm applied to $L(W_1, W_2)$ described in the previous section. Alongside, we provide a high-level overview of the overall strategy we use to overcome these challenges. Based on these strategy, we derive in §3.2 a rigorous convergence rate that depends

on not only the imbalance and margin at the initialization but also the step size and condition number of the data. Finally, in §3.3 we propose an adaptive step size scheme that accelerates convergence. Due to the discrete nature of our updates, we thus consider t to be discrete, i.e., $t \in \mathbb{N}$.

3.1 Challenges in the Analysis of Over-parametrized Gradient Descent

Standard GD. We start by deriving the convergence rate of the non-overparametrized regime described in equation 1. Notice that $\ell(t)$ is K-smooth and satisfies the μ-PL condition, where $K = \sigma_{\max}^2(X)$ and $\mu = \sigma_{\min}^2(X)$. Then, the following smoothness inequality holds for any $W(t), W(t+1)$:

$$
\begin{aligned}
\ell(t+1) \leq \ell(t) &+ \langle \nabla \ell(t), W(t+1) - W(t) \rangle \\
&+ \frac{K}{2} \| W(t+1) - W(t) \|_F^2.
\end{aligned} \tag{13}
$$

After substituting the GD update with fixed step size η

$$
W(t+1) = W(t) - \eta \nabla \ell(t) \tag{14}
$$

into the smoothness inequality in equation 13 we obtain

$$
\begin{aligned}
\ell(t+1) &\leq \ell(t) - \eta \| \nabla \ell(t) \|_F^2 + \frac{K}{2} \eta^2 \| \nabla \ell(t) \|_F^2 \\
&= \ell(t) - \eta \left(1 - K \frac{\eta}{2} \right) \| \nabla \ell(t) \|_F^2.
\end{aligned} \tag{15}
$$

Then, if the step size satisfies $\eta < \frac{2}{K}$, then the loss is non-increasing. Moreover, if we apply the PL condition $\frac{1}{2} \| \nabla \ell(t) \|_F^2 \geq \mu \ell(t)$ to equation 15, we obtain

$$
\ell(t+1) \leq (1 - 2\eta\mu + K\mu\eta^2) \ell(t), \tag{16}
$$

which suffices to show the linear convergence of GD, for properly chosen η.

Over-parametrized GD. In the over-parametrized case, we use the chain rule to write the gradient of L with respect to W_1, W_2 in terms of $\nabla \ell(W), W_1, W_2$. The update of weights in GD is

$$
\begin{pmatrix} W_1(t+1) \\ W_2(t+1) \end{pmatrix} = \begin{pmatrix} W_1(t) \\ W_2(t) \end{pmatrix} - \eta \gamma_t (\nabla \ell(t)). \tag{17}
$$

Thus, the update of the product is

$$
\begin{aligned}
W(t+1) &= W_1(t+1) W_2(t+1) \\
&= \left(W_1(t) - \eta \nabla \ell(t) W_2(t)^\top \right) \left(W_2(t) - \eta W_1(t)^\top \nabla \ell(t) \right) \\
&= W(t) - \eta \tau_t (\nabla \ell(t)) + \eta^2 \nabla \ell(t) W(t)^\top \nabla \ell(t). \tag{18}
\end{aligned}
$$

In other words, the update of the product is a polynomial of degree two on the step size η, unlike the update in equation 14, which is a polynomial of degree one. Substituting equation 18 into the smoothness inequality 13, and using the PL condition, we can connect the loss at iteration

$t + 1$ with the loss at iteration t. The following lemma characterizes this property.

Lemma 3.1. *If at the t-th iteration of GD applied to the over-parametrized loss L, the step size η satisfies*

$$
\begin{aligned}
&\lambda_{\min}(\tau_t) - \eta \| \nabla \ell(t) \|_F \sigma_{\max}(W(t)) \\
&- \frac{K\eta}{2} \left[\lambda_{\max}(\tau_t) + \eta \| \nabla \ell(t) \|_F \sigma_{\max}(W(t)) \right]^2 \geq 0,
\end{aligned} \tag{19}
$$

then the following inequality holds

$$
L(t+1) \leq \rho(\eta, t) L(t), \tag{20}
$$

where

$$
\begin{aligned}
\rho(\eta, t) = 1 &- 2\eta\mu\lambda_{\min}(\tau_t) + K\mu\eta^2\lambda_{\max}^2(\tau_t) \\
&+ 2\eta^2\mu\sigma_{\max}(W(t)) \| \nabla \ell(t) \|_F \\
&+ 2\eta^3\mu K\lambda_{\max}(\tau_t)\sigma_{\max}(W(t)) \| \nabla \ell(t) \|_F \\
&+ \eta^4\mu K\sigma_{\max}^2(W(t)) \| \nabla \ell(t) \|_F^2. \tag{21}
\end{aligned}
$$

The proof of the above lemma can be found in Appendix B.

Comparison with non-overparametrized GD. The difference between the inequality we derive in Lemma 3.1 and the one in equation 16 is twofold. Firstly, $\rho(\eta, t)$ in equation 21 includes a quadratic polynomial of η:

$$
1 - 2\eta\mu\lambda_{\min}(\tau_t) + K\mu\eta^2\lambda_{\max}^2(\tau_t) \tag{22}
$$

that resembles the one in equation 16. The only difference is that the second coefficient is now scaled by $\lambda_{\min}(\tau_t)$ and the third coefficient by $\lambda_{\max}^2(\tau_t)$. Equation 22 comes from the term $\eta\tau_t(\nabla \ell(t))$ in the product update in equation 18, which corresponds to moving the weight $W(t)$ along the "skewed gradient direction" $\tau_t(\nabla \ell(t))$ instead of $\nabla \ell(t)$. Secondly, equation 21 has extra second- and higher-order terms in η which come from the other term $\eta^2 \nabla \ell(t) W^\top(t) \nabla \ell(t)$ in equation 18. Overall, compared to equation 16, the over-parametrized GD introduces a more complicated update on the product $W(t)$, leading to the inequality in equation 20 that not only is a polynomial of degree four in η, but also depends on the weights $W_1(t), W_2(t)$ at the current iteration. These differences pose additional challenges in deriving a linear convergence rate for over-parametrized GD.

Towards linear convergence. Lemma 3.1 provides an upper bound on $L(t + 1)$, $\rho(\eta, t)L(t)$, which implicitly depends on $W_1(t)$ and $W_2(t)$ via $\lambda_{\min}(\tau_t)$, $\sigma_{\max}(W(t))$, $\ell(t)$ and $\lambda_{\max}(\tau_t)$. However, it is unclear whether one can find some step size η that can simultaneously satisfy equation 19 and uniformly bound $\rho(\eta, t) \leq \bar{\rho} < 1$, for all t. Only under such conditions Lemma 3.1 would lead to

$$
L(t + 1) < \bar{\rho}L(t) < (\bar{\rho})^{t+1}L(0). \tag{23}
$$

We approach this challenge in a similar spirit as it was done in GF Min et al. [2022].

Step 1. Spectral bounds for τ_t and $W(t)$: First, we seek to find bounds for $\lambda_{\min}(\tau_t)$ and $\lambda_{\max}(\tau_t)$ based on the imbalance $D(t)$ and the singular values of the product, i.e.,

$$\alpha(D(t), \sigma_{\min}(W(t))) \leq \lambda_{\min}(\tau_t)$$
$$\lambda_{\max}(\tau_t) \leq \beta(D(t), \sigma_{\max}(W(t))), \quad (24)$$

where both functions $\alpha(D, \sigma)$ and $\beta(D, \sigma)$ are increasing on the second argument, σ. As a result, if one is able to control $D(t)$ and the singular values of $W(t)$, one can attempt to upper-bound $\rho(\eta, t)$ in equation 21.

For the case of $\sigma_{\min}(W(t))$ and $\sigma_{\max}(W(t))$, a similar monotonicity argument as in GF can be done to obtain

$$p_1 \leq \sigma_{\min}(W(t)) \leq \sigma_{\max}(W(t)) \leq p_2, \quad (25)$$

where p_2 is defined in Table 2.

The additional, non-trivial challenge present in GD is the fact that the imbalance $D(t)$ is no longer preserved, i.e., $D(t) \neq D(0)$, which makes it still difficult control $\lambda_{\min}(\tau_t), \lambda_{\max}(\tau_t)$ by equation 24. Nevertheless, we show in *Theorem 3.1* that if η is sufficiently small, but not infinitesimal, it is possible to control how much the imbalance changes by bounding $\|D(t) - D(0)\|$ for all t, which leads to a uniform bound of the form

$$\alpha_0 c_1 \leq \lambda_{\min}(\tau_t) \leq \lambda_{\max}(\tau_t) \leq \beta_0 c_2, \quad (26)$$

where $\beta_0 := \beta(D(0), p_1)$, and the parameters $0 < c_1 < 1$, $c_2 > 1$ represent an additional level of conservativeness in the bound that is necessary to accommodate the time varying nature of $D(t)$ in GD; see discussion after Theorem 3.1 for more details.

Stage 2. Uniform upper-bound on rate $\rho(\eta, t)$: Once bounds for the spectrum of $W(t)$ and τ_t have been established, one can then proceed to bound $\rho(\eta, t)$ in equation 21. In particular, we will show that $\rho(\eta, t) \leq f(\eta, t)$, where

$$f(\eta, t) := 1 - a_1\eta + a_2(t)\eta^2 + a_3(t)\eta^3 + a_4(t)\eta^4, \quad (27)$$

and the dependency on time is only through $L(t)$, i.e.,

$$a_1 = 2(c_1\alpha_0)\sigma_{\min}^2(X),$$
$$a_2(t) = 2\sqrt{2\kappa L(t)\sigma_{\min}^6(X)}p_2 + \kappa\sigma_{\min}^4(X)(c_2\beta_0)^2,$$
$$a_3(t) = 2\sqrt{2\kappa^3 L(t)\sigma_{\min}^{10}(X)}c_2\beta_0 p_2,$$
$$a_4(t) = 2\kappa^2\sigma_{\min}^6(X)p_2^2 L(t). \quad (28)$$

The above bound for $\rho(\eta, t)$ in equation 27, whose derivation is provided in *Theorem 3.2*, can be then leveraged in multiple ways.

- **Uniform linear rate.** Under mild conditions on the step size, here exists η independent of t such that $f(\eta, t) \leq f(\eta, 0)$ (also in *Theorem 3.2*), leading to

$$L(t) \leq \Pi_{k=0}^t f(\eta, k) L(0) \leq (f(\eta, 0))^t L(0). \quad (29)$$

- **Time-varying step size.** A natural consequence of equations 27 and 29 is the possibility to adaptively choose η_t, using only knowledge of the current loss $L(t)$, so as to improve the convergence rate. This is explored in §3.3; see *Algorithm 1*.

3.2 General bound on linear convergence rate

In this subsection, we derive conditions under which Lemma 3.1 is a descent lemma. Based on this result, we can prove that GD converges linearly to a global minimum of equation 2. We refer the reader to Table 2 for the definition of various quantities appearing in this section.

Before stating our main result, we note that prior work Arora et al. [2018]; Du et al. [2018a] studied optimizing equation 2 via GD, but their results require the initial imbalance to have small Frobenius norm and the initial margin to be sufficiently large. The NTK initialization Du and Hu [2019] does not require small imbalance, but it does require a large hidden-layer width h, and the weights needs to be randomly initialized. To the best of our knowledge, Theorem 3.2 is the first convergence result for GD which provides an explicit convergence rate without making the assumption that the initial imbalance is small or that the width of the network is large.

Table 2: Table of Notation

SYMBOL	DEFINITION
$\ell(t)$	$\ell(W(t))$
$L(t)$	$L(W_1(t), W_2(t))$
$\tau_t(\nabla\ell(t))$	$\tau(\nabla\ell(W(t)); W_1(t), W_2(t))$
$\lambda_{\min}(\tau_t)$	$\lambda_{\min}(\tau(\cdot; W_1(t), W_2(t)))$
$\lambda_{\max}(\tau_t)$	$\lambda_{\max}(\tau(\cdot; W_1(t), W_2(t)))$
$D(t)$	$W_1^\top(t)W_1(t) - W_2(t)W_2(t)^\top$
$W(t)$	$W_1(t)W_2(t)$
$E(t)$	$Y - XW_1(t)W_2(t)$
κ	$\dfrac{\sigma_{\max}^2(X)}{\sigma_{\min}^2(X)}$
p_1	$\dfrac{\max\{\sigma_{\min}(Y) - \|E(0)\|_F, 0\}}{\sigma_{\max}(X)}$
p_2	$\dfrac{\|Y\|_F + \|E(0)\|_F}{\sigma_{\min}(X)}$
Δ_+	$\max(\lambda_{\max}(D(0)), 0) - \max(\lambda_n(D(0)), 0)$
Δ_-	$\max(\lambda_{\max}(-D(0)), 0) - \max(\lambda_m(-D(0)), 0)$
$\underline{\Delta}$	$\max(\lambda_n(D(0)), 0) + \max(\lambda_m(-D(0)), 0)$
λ_+	$\max(\lambda_{\max}(D(0)), 0)$
λ_-	$\max(\lambda_{\max}(-D(0)), 0)$
α_0	$\dfrac{-\Delta_+ + \sqrt{(\Delta_+ + \underline{\Delta})^2 + 4p_1^2}}{2} + \dfrac{-\Delta_- + \sqrt{(\Delta_- + \underline{\Delta})^2 + 4p_1^2}}{2}$
β_0	$\dfrac{\lambda_+ + \sqrt{\lambda_+^2 + 4p_2^2}}{2} + \dfrac{\lambda_- + \sqrt{\lambda_-^2 + 4p_2^2}}{2}$

Theorem 3.1 (Uniform bounds on eigenvalues of τ_t and singular values of $W(t)$). *Assume $\alpha_0 > 0$, and choose $0 < c_1 < 1$, and $c_2 > 1$. Let $\eta_1^{\max}$ and $\eta_2^{\max}$ be, respectively, the unique positive roots of the following two polynomials*

in η

$$a_4(0)\eta^3 + a_3(0)\eta^2 + \left(a_2(0) + \frac{4c_2 L(0)\sigma_{\max}^2(X)}{c_2 - 1}\right)\eta = a_1,$$

$$a_4(0)\eta^3 + a_3(0)\eta^2 + \left(a_2(0) + \frac{8c_2\beta_0 L(0)\sigma_{\max}^2(X)}{(1 - c_1)\alpha_0}\right)\eta = a_1.$$

$$(30)$$

Then, for any $0 < \eta \leq \eta_{\max} := \min\{\eta_1^{\max}, \eta_2^{\max}\}$, the following holds for all $t = 0, 1, \ldots$

$$c_1\alpha_0 \leq \lambda_{\min}(\tau_t) \leq \lambda_{\max}(\tau_t) \leq c_2\beta_0$$
$$p_1 \leq \sigma_{\min}(W(t)) \leq \sigma_{\max}(W(t)) \leq p_2. \qquad (31)$$

The above theorem says that when the step size is small, we can bound the eigenvalues of τ_t and the singular values of $W(t)$ using the initial *imbalance* and *margin*. When $\alpha_0 > 0$, we have $a_1 > 0$, and the LHS of equation 30 is a monotonically increasing function of η, when $\eta > 0$, and is equal to zero, when $\eta = 0$. Therefore, each polynomial has a unique positive root. The condition $\eta \leq \eta_{\max}$ is used to control $\|D(t) - D(0)\|_F$. We use $\lambda_{\min}(\tau_t)$ as an example to illustrate why we need to control $\|D(t) - D(0)\|_F$. In GD, equation 9 still holds. However, since the imbalance is no longer constant, i.e. since $D(t) \neq D(0)$, we no longer have $\alpha(D(t), p_1) = \alpha(D(0), p_1)$. Nonetheless, after careful analysis, we observe that the change of imbalance at each iteration is of order η^2. Moreover, as long as the loss decreases linearly and η is small (see equation 30), we can prove that $\|D(t) - D(0)\|_F \leq O(\eta)$. Thus, we first introduce c_1 to control the change of the eigenvalues of the imbalance matrix. Then, if the step size is bounded, i.e. $\eta \leq \eta_{\max}$, we can show that $\alpha(D(t), p_1) \geq c_1\alpha(D(0), p_1)$. A similar analysis yields the upper bound for $\lambda_{\max}(\tau_t)$. When c_1, c_2 are chosen to be close to one, the change in eigenvalues of imbalance is guaranteed to be small, but it requires a smaller step as $\eta_{\max}$ is small.

Then, based on Theorem 3.1, we can prove the linear convergence of GD.

Theorem 3.2 (Convergence rate of gradient descent on two-layer linear networks). *Under the assumptions in Theorem 3.1, for any $0 < \eta \leq \eta_{\max} := \min\{\eta_1^{\max}, \eta_2^{\max}\}$, the loss function under GD satisfies*

$$L(t+1) \leq f(\eta, t)L(t), \qquad (32)$$

for $f(\eta, t)$ as defined in equation 27, and with

$$0 < f(\eta, t) \leq f(\eta, 0) < 1, \quad \forall t \geq 0. \qquad (33)$$

Thus, the loss converges linearly, i.e.,

$$L(t) \leq \Pi_{k=0}^t f(\eta, k) L(0) \leq f(\eta, 0)^t L(0). \qquad (34)$$

with rate given by $f(\eta, 0)$.

In $f(\eta, t)$, $-a_1\eta$ is an important term that facilitates convergence because it is the only term that is associated with a negative coefficient. Notice that a_1 depends on $p_1, D(0)$ via α_0, and when $\alpha_0 > 0$, i.e., there is either sufficient margin or imbalance, we have $a_1 > 0$. The proof Theorem 3.1 and Theorem 3.2 is presented in Appendix C.

Detailed comparison with SOTA. We compare our results with other works studying the same problem Du et al. [2018a]; Arora et al. [2018]. In both works, the authors make assumptions that the initial imbalance is small. In our work, Theorem 3.2 holds if there is either a sufficient imbalance or sufficient margin at initialization, which is a more general setting. In Du et al. [2018a], they prove the loss decreases, and the imbalance remains small during training, but the paper does not provide an explicit convergence rate. More importantly, a decay in step size is needed to control the difference between $D(t)$ and $D(0)$. In our work, we provide an explicit convergence rate without the need to decrease step size. In Arora et al. [2018], the authors provide an explicit convergence rate. However, their result depends on the property that when step size is small, $\|D(t)\|_F \leq 2\|D(0)\|_F$. We think the *two* used in their proof is an artifact and improve it by introducing c_1 and c_2 and characterize the dependence between step size and c_1, c_2, which is a more general case.

Comparison with non-overparametrized regime. In the GF regime, Min et al. [2022]; Tarmoun et al. [2021] show that if α_0 is sufficiently large, the over-parametrized model can have a faster convergence rate than the non-overparametrized model. However, as shown in the next proposition, such a result does not extend to the GD regime.

Proposition 3.1. *If $\alpha_0 > 0$, for all $0 < \eta \leq \eta_{\max}$ and for all $t = 0, 1, \cdots$, the following inequality holds*

$$f(\eta, t) \geq 1 - \frac{1}{\kappa}, \qquad (35)$$

where $\kappa = \frac{K}{\mu}$ is the condition number of the non-overparametrized Problem 1

In Proposition 3.1, $1 - \frac{1}{\kappa}$ is the theoretical optimal convergence rate of solving Problem 1 via GD (see §3.1 for a derivation of it). As a result, Proposition 3.1 states that the convergence rate derived in Theorem 3.2, i.e., $f(\eta, t)$, for solving the over-parametrized Problem 2 via GD, is always larger. Nevertheless, we point out that Theorem 3.2 only provides an upper bound on the rate, and further study is needed to characterize its tightness.

3.3 Adaptive Step Size Scheme

Motivation. In Theorem 3.2, we used the bound in equation 32 to show that a fixed step size $\eta \leq \eta_{\max}$ guarantees a linear rate of convergence for the loss $L(t)$. It is thus natural to ask whether we can improve upon this rate by using

a tighter bound and an adaptive step size η_t. Specifically, if we can show that there exists a function $h(\eta, t) \geq 0$ such that for all t, $h(\eta, t) \leq f(\eta, t)$, $L(t+1) \leq h(\eta, t)L(t)$ and $\min_\eta h(\eta, t) < 1$, then we can pick a step size η_t that minimizes $h(\eta, t)$, i.e.,

$$\eta_t = \arg\min_\eta h(\eta, t). \tag{36}$$

Such a choice for the step size would theoretically achieve the most decrease in the loss function $L(t)$. As a result, we would obtain the following theoretical upper bound on $L(t)$

$$L(t) \leq \prod_{k=0}^{t-1} h(\eta_k, k)L(0). \tag{37}$$

Choices of $h(\eta, t)$. We already know from equation 20 and equation 32 that $h(\eta, t) = \rho(\eta, t)$ and $h(\eta, t) = f(\eta, t)$, respectively, are valid choices for h. Based on our analysis in §3.1 and §3.2, in the following proposition we present additional choices for $h(\eta, t)$ and the conditions under which $\min_\eta h(\eta, t) < 1$ holds for every t.

Proposition 3.2. *Let $h(\eta, t)$ be one of the following functions $\rho(\eta, t)$, $\hat{\rho}(\eta, t)$ or $f(\eta, t)$, where*

$$\hat{\rho}(\eta, t) = 1 - 2\eta\mu\alpha_t + K\mu\eta^2\beta_t^2 + 2\eta^2\mu p_2\sqrt{2KL(t)} + 2\eta^3\mu K\beta_t p_2\sqrt{2KL(t)} + 2\eta^4\mu K^2 p_2^2 L(t). \tag{38}$$

Then, under the assumptions of Theorem 3.1, whenever, $\eta \leq \eta_{\max}$, the following holds for all choices of $h(\eta, t)$ and every $t = 0, 1, 2, \ldots$

$$L(t+1) \leq h(\eta, t)L(t), \tag{39}$$

with $0 \leq \rho(\eta, t) \leq \hat{\rho}(\eta, t) \leq f(\eta, t) < 1$.

Proposition 3.2 shows that choosing $0 < \eta \leq \eta_{\max}$ guarantees $\min_\eta h(\eta, t) < 1$ holds for every t. Moreover, since one can view $f(\eta, 0)$ as a special choice of $h(\eta, t)$ where the $h(\eta, t) = f(\eta, 0)$ for every t, we can broaden the choice of $h(\eta, t)$, to compactly describe all rate bounds of our paper, i.e.,

$$h(\eta, t) \in \{\rho(\eta, t), \hat{\rho}(\eta, t), f(\eta, t), f(\eta, 0)\}. \tag{40}$$

Algorithm description. As suggested by Proposition 3.2, given a choice of $h(\eta, t)$, one can compute the adaptive step size by solving the following optimization problem

$$\eta_t = \arg\min_{\eta \leq \eta_{\max}} h(\eta, t). \tag{41}$$

Since all choices for $h(\eta, t)$ are fourth-order polynomials, solving equation 41 only requires finding the roots of a third-order polynomial. More specifically, we have:

Claim 3.1. *Suppose $\alpha_0 > 0$. Let η_t' be the unique positive root of the following equation*

$$\frac{dh(\eta, t)}{d\eta} = 0. \tag{42}$$

Then the solution to Problem 41 is $\eta_t = \min(\eta_t', \eta_{\max})$.

The proof is in Appendix E. This suggests that one can find η_t very efficiently at each iteration. The resulting GD algorithm with adaptive step size is summarized in Algorithm 1.

Algorithm 1: GD with Adaptive Step Size

Data: X, Y, and initial $W_1(0), W_2(0)$
Result: W_1^*, W_2^* that minimize $\frac{1}{2}\|Y - XW_1W_2\|_F^2$.
for $t = 0, 1, 2 \cdots$ **do**
  /* adaptive step size   */
  $\eta_t \leftarrow \arg\min_{\eta \leq \eta_{\max}} h(\eta, t)$
  /* GD update with η_t */
  $\begin{pmatrix} W_1(t+1) \\ W_2(t+1) \end{pmatrix} = \begin{pmatrix} W_1(t) \\ W_2(t) \end{pmatrix} - \eta_t\gamma_t(\nabla\ell(t)).$
end

Convergence rate under adaptive step size. Notice that $h(\eta, t)$ depends on the iteration t via the loss function $L(t)$. As the training proceeds, the adaptive step size scheme ensures $h(\eta, t) < 1$ such that the loss $L(t)$ converges to zero. This, in turn, affects the asymptotic expression for $h(\eta, t)$. Specifically, when t is sufficiently large (so that $L(t) \simeq 0$), we have

$$f(\eta, t) \simeq 1 - 2(c_1\alpha_0)\sigma_{\min}^2(X)\eta + \kappa\sigma_{\min}^4(X)(c_2\beta_0)^2\eta^2,$$
$$\hat{\rho}(\eta, t) \simeq 1 - 2\alpha_t\sigma_{\min}^2(X)\eta + \kappa\sigma_{\min}^4(X)\beta_t^2\eta^2,$$
$$\rho(\eta, t) \simeq 1 - 2\lambda_{\min}(\tau_t)\sigma_{\min}^2(X)\eta + \kappa\sigma_{\min}^4(X)\lambda_{\max}(\tau_t)^2\eta^2.$$

Under a proper choice of c_1, c_2 such that

$$\eta_{\max} \geq \frac{c_1\alpha_0}{c_2^2\beta_0^2\kappa\sigma_{\min}^2(X)} := \eta_t^* \tag{43}$$

the adaptive step size schedulers yield a rate

$$\rho(\eta_t^*, t) \leq \hat{\rho}(\eta_t^*, t) \leq f(\eta_t^*, t) \simeq 1 - \frac{(c_1\alpha_0)^2}{(c_2\beta_0)^2}\frac{1}{\kappa}. \tag{44}$$

For adaptive step size scheme in equation 36 under any one of the choices of $h(\eta, t)$ in Proposition 3.2, the asymptotic convergence rate of GD with is upper bounded by a constant depending on both $\frac{\alpha_0}{\beta_0}$ and $\frac{1}{\kappa}$. In Appendix E, we show that there always exists a choice for c_1, c_2 such that equation 43 holds. Our numerical simulations show that GD with all three adaptive step size strategies listed in Proposition 3.2 achieve faster convergence than GD with a fixed step size in Theorem 3.2. Please refer to Section 4.2 for details.

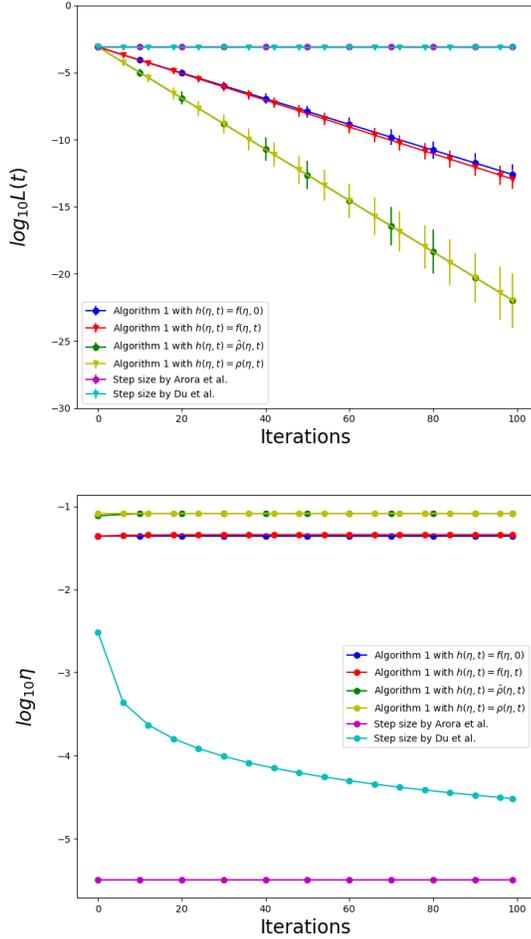linear networks is 1,000. The details of the simulations are presented in Appendix F.

Figure 1: Evolution of the loss and of the step size for different choices of the step size schedule in the large margin and balanced spectral initialization regime. We select the two schedules in Arora et al. [2018]; Du et al. [2018a] as well as four other schedules according to Algorithm 1 with $h(\eta, t)$ chosen from equation 40. We run the simulations fifty times. On the top plot, the solid line is the $\log_{10}$ of the mean reconstruction error $L(t)$, and the vertical line is the mean plus and minus one standard deviation.

4 SIMULATIONS

In this section, we first compare the convergence rate of GD using the step sizes presented in §3.3 with those using step size proposed in previous work Arora et al. [2018]; Du et al. [2018a], and our step sizes achieve considerably faster convergence rate. We then present experiments under much more general initializations than in prior work. Our experiments show that that the bounds in equation 37 provide a good characterization of the actual convergence rate. Throughout the experiments, we fix $c_1 = 0.5$ and $c_2 = 1.5$ in our choice of step sizes and the width of the two-layer

4.1 Comparison between different learning rates presented in previous work

In this section, we compare the step sizes proposed in Theorem 2 of Arora et al. [2018] and Theorem 3.1 of Du et al. [2018a] with the step sizes proposed in §3.3. We note that the analyses in Arora et al. [2018]; Du et al. [2018a] assume that the initialization is approximately balanced ($\|D(0)\|_F$ is small). In addition, Arora et al. [2018] requires the initialization to have sufficient margin ($\|Y - XW_1(0)W_2(0)\|_F$ is small). Therefore, we compare our results with Arora et al. [2018]; Du et al. [2018a] using the following simulation setup that initializes the weight matrices via (balanced) spectral initialization. Specifically,

$$
\begin{aligned}
&X = I_{20}, Y = XW(0) + 0.04\varepsilon\,, \\
&W(0) \in \mathbb{R}^{20 \times 1}, W(0)[i, j] \sim \mathcal{N}(0, 1/4)\,, \\
&\varepsilon \in \mathbb{R}^{20 \times 1}, \varepsilon[i, j] \sim \mathcal{N}(0, 1)\,, \\
&W_1(0) = U\Sigma V^\top, W_2(0) = \Sigma^{1/2} V^\top\,.
\end{aligned}
\tag{45}
$$

where $W(0) = U\Sigma V^\top$ is the SVD of $W(0)$. Such data generation setup ensures that $\|Y - XW(0)\|_F$ is small so that there is a sufficiently large margin. The balanced spectral initialization guarantees the initial imbalance is zero. Recall that these conditions of zero balance and large margin are needed by Arora et al. [2018]; Du et al. [2018a] to obtain their convergence results, but they are not needed for our approach. Therefore, we work in this restrictive setting only to facilitate comparison with prior work.

Figure 1 shows that the step size choices proposed in our paper achieve the fastest convergence compared with other SOTA methods Arora et al. [2018]; Du et al. [2018a] (top plot). The bottom plot also shows that the step sizes proposed in this work are larger than the ones proposed in Arora et al. [2018]; Du et al. [2018a]. Moreover, the adaptive step size schemes minimizing $\hat{\rho}(\eta, t), \rho(\eta, t)$ have similar rates and are faster than the one minimizing $f(\eta, t), f(\eta, 0)$. This is because $\hat{\rho}(\eta, t), \rho(\eta, t)$ use $W_1(t), W_2(t)$ to bound the eigenvalue of τ_t at each iteration, which gives tighter bounds on the loss. On the other hand, the coefficients of $f(\eta, t)$ simply use $c_1\alpha_0, c_2\beta_0$ which is conservative.

4.2 Evaluation of the tightness of the theoretical bound on the convergence rate

In this section, we train a two-layer linear network using GD on the squared loss in equation 2. We generate the data matrix as follows: $X \in \mathbb{R}^{20 \times 20}, X[i, j] \sim \mathcal{N}(0, 1)$, and $Y = X\Theta$ where $\Theta \in \mathbb{R}^{20 \times 20}, \Theta[i, j] \sim \mathcal{N}(0, 1)$. The initial weight matrices are generated as $W_1(0) = \sigma U_0, W_2(0) = \frac{1}{\sigma}V_0$, where $U_0 \in \mathbb{R}^{20 \times 1000}, V_0 \in$
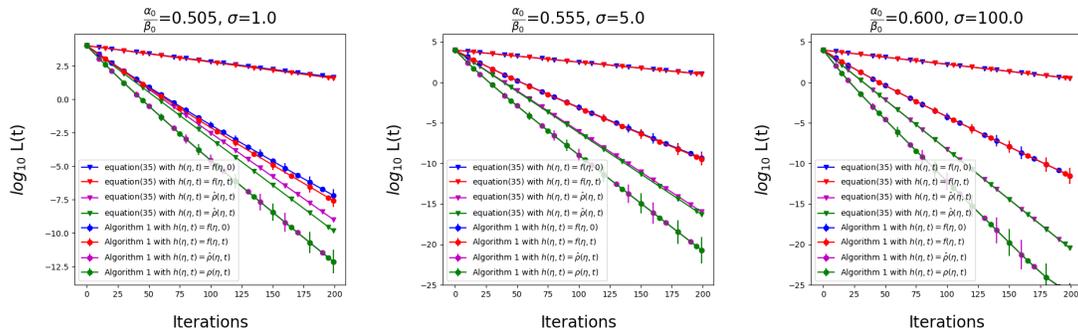
Figure 2: Tightness of the theoretical upper bound in equation 37 with the choice of $h(\eta, t)$ specified in equation 40 versus reconstruction error $L(t)$, shown in different colors. We run the simulations for three different initialization. For each initialization and each choice of step size, we repeat the simulation fifty times. The triangle lines represent the theoretical upper bound on the training loss in equation 37. The circle lines represent the $\log_{10}$ of the mean reconstruction error $L(t)$. The horizontal line is the mean plus and minus one standard deviation.

$\mathbb{R}^{1000 \times 20}$ and have entry-wise i.i.d. samples drawn from a standard Gaussian $\mathcal{N}(0, 1)$. We choose different values of σ to test our convergence rate in different regimes.

Figure 2 compares the convergence rate predicted by equation 37 with the actual convergence rate of $L(t)$, for different values of σ and approximately similar values of $\frac{\alpha_0}{\beta_0}$. In all scenarios, our theoretical bounds in equation 37 follow the empirical results relatively well. Moreover, we see the convergence rate of all methods vary with σ. This is because the rate depends on $\frac{\alpha_0}{\beta_0}$ shown in equation 44. Thus, the larger $\frac{\alpha_0}{\beta_0}$, the faster the convergence. Finally, in this experiment, the initial margin is 0, and there is a large initial imbalance. Those initial conditions violate the assumptions in Arora et al. [2018]; Du et al. [2018a], but still enjoy linear convergence. Thus, our theory applies beyond the regime of Arora et al. [2018]; Du et al. [2018a].

5 CONCLUSIONS

This paper studied the convergence of GD for optimizing two-layer linear networks. In particular, we derived a convergence rate for networks of finite width that are initialized in a non-NTK regime. Our results build upon recent work for GF, which derived convergence rates that depend on the imbalance and margin of the initialization. However, a key challenge in the GD regime is that the imbalance of the weights changes with the iterations of GD. In this paper, we show that when the step size is small, the imbalance at iteration t is close to its value at initialization. Moreover, we show that under this constraint on the step size, the loss is decreasing. In addition, we derive an explicit convergence rate that depends on the margin, imbalance, and condition number of the data matrix. Finally, based on the convergence rate, we propose three adaptive step size schemes that accelerate convergence compared with a con-

stant step size. Empirically, we show the convergence rate derived in our work is tighter than in previous work.

References

Arora, S., Cohen, N., Golowich, N., and Hu, W. (2018). A convergence analysis of gradient descent for deep linear neural networks. *arXiv preprint arXiv:1810.02281*.

Chen, Z., Vanden-Eijnden, E., and Bruna, J. (2022). On feature learning in neural networks with global convergence guarantees. *arXiv preprint arXiv:2204.10782*.

Chizat, L. and Bach, F. (2018). On the global convergence of gradient descent for over-parameterized models using optimal transport. *Advances in neural information processing systems*, 31.

Chizat, L., Oyallon, E., and Bach, F. (2019). On lazy training in differentiable programming. *Advances in Neural Information Processing Systems*, 32.

Ding, Z., Chen, S., Li, Q., and Wright, S. J. (2022). Over-parameterization of deep resnet: Zero loss and mean-field analysis. *J. Mach. Learn. Res.*, 23:48–1.

Du, S. and Hu, W. (2019). Width provably matters in optimization for deep linear neural networks. In *International Conference on Machine Learning*, pages 1655–1664.

Du, S. S., Hu, W., and Lee, J. D. (2018a). Algorithmic regularization in learning deep homogeneous models:

Layers are automatically balanced. *Advances in Neural Information Processing Systems*, 31.

Du, S. S., Zhai, X., Poczos, B., and Singh, A. (2018b). Gradient descent provably optimizes over-parameterized neural networks. *arXiv preprint arXiv:1810.02054*.

Gidel, G., Bach, F., and Lacoste-Julien, S. (2019). Implicit regularization of discrete gradient dynamics in linear neural networks. *Advances in Neural Information Processing Systems*, 32.

He, K., Zhang, X., Ren, S., and Sun, J. (2015). Deep residual learning for image recognition.

Jacot, A., Gabriel, F., and Hongler, C. (2018). Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31.

Lee, J., Xiao, L., Schoenholz, S., Bahri, Y., Novak, R., Sohl-Dickstein, J., and Pennington, J. (2019). Wide neural networks of any depth evolve as linear models under gradient descent. *Advances in neural information processing systems*, 32.

Liu, C., Zhu, L., and Belkin, M. (2022). Loss landscapes and optimization in over-parameterized non-linear systems and neural networks. *Applied and Computational Harmonic Analysis*, 59:85–116.

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.

Min, H., Tarmoun, S., Vidal, R., and Mallada, E. (2022). Convergence and implicit bias of gradient flow on overparametrized linear networks. *arXiv preprint arXiv:2105.06351*.

Minaee, S., Boykov, Y. Y., Porikli, F., Plaza, A. J., Kehtarnavaz, N., and Terzopoulos, D. (2021). Image segmentation using deep learning: A survey. *IEEE transactions on pattern analysis and machine intelligence*.

Nguegnang, G. M., Rauhut, H., and Terstiege, U. (2021). Convergence of gradient descent for learning linear neural networks. *arXiv preprint arXiv:2108.02040*.

Oymak, S. and Soltanolkotabi, M. (2020). Toward moderate overparameterization: Global convergence guarantees for training shallow neural networks. *IEEE Journal on Selected Areas in Information Theory*, 1(1):84–105.

Rotskoff, G. and Vanden-Eijnden, E. (2018a). Parameters as interacting particles: long time convergence and asymptotic error scaling of neural networks. *Advances in neural information processing systems*, 31.

Rotskoff, G. M. and Vanden-Eijnden, E. (2018b). Trainability and accuracy of neural networks: An interacting particle system approach. *arXiv preprint arXiv:1805.00915*.

Saxe, A. M., McClelland, J. L., and Ganguli, S. (2013). Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*.

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., and Hassabis, D. (2016). Mastering the game of go with deep neural networks and tree search. *Nature*, 529:484–503.

Sirignano, J. and Spiliopoulos, K. (2020). Mean field analysis of neural networks: A law of large numbers. *SIAM Journal on Applied Mathematics*, 80(2):725–752.

Tarmoun, S., Franca, G., Haeffele, B. D., and Vidal, R. (2021). Understanding the dynamics of gradient flow in overparameterized linear models. In *International Conference on Machine Learning*, pages 10153–10161. PMLR.

Vaswani, A., Bengio, S., Brevdo, E., Chollet, F., Gomez, A. N., Gouws, S., Jones, L., Kaiser, Ł., Kalchbrenner, N., Parmar, N., et al. (2018). Tensor2tensor for neural machine translation. *arXiv preprint arXiv:1803.07416*.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need.

Vo, N. N., He, X., Liu, S., and Xu, G. (2019). Deep learning for decision making and the optimization of socially responsible investments and portfolio. *Decision Support Systems*, 124:113097.

A CASE WHEN DATA MATRIX IS RANK DEFICIENT

Here, we show for any data matrix X of arbitrary dimensions and rank, the over-parametrized problem

$$\min_{W_1,W_2} L(W_1,W_2) = \frac{1}{2}\|Y - XW_1W_2\|_F^2, \tag{46}$$

can be reparametrized into the following problem

$$\min_{\tilde{W}_1,\tilde{W}_2} L(\tilde{W}_1,\tilde{W}_2) = \frac{1}{2}\|\tilde{Y} - \tilde{X}\tilde{W}_1\tilde{W}_2\|_F^2, \tag{47}$$

where $\tilde{X}$ is a square matrix of full rank.

Let singular value decomposition of X be

$$X = [U_1, U_2] \begin{bmatrix} \Sigma_X & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} V_1^\top \\ V_2^\top \end{bmatrix}, \tag{48}$$

where Σ_X contains all non-zero singular values of X. Then, we have $X = U_1\Sigma_X V_1^\top$. The GD update of W_1, W_2 is

$$W_1(t+1) = W_1(t) + \eta X^\top E(t)W_2^\top(t) = W_1(t) + \eta V_1\Sigma_X U_1^\top E(t)W_2^\top(t),$$
$$W_2(t+1) = W_2(t) + \eta W_1^\top(t)X^\top E(t) = W_2(t) + \eta W_1^\top(t)V_1\Sigma_X U_1^\top E(t). \tag{49}$$

We project W_1 onto the space spanned by V_1, V_2,

$$W_{11} = V_1^\top W_1,$$
$$W_{12} = V_2^\top W_1. \tag{50}$$

Furthermore, we define $\tilde{E}(t) = U_1^\top E(t)$. Based on above, one has

$$W_{11}(t+1) = W_{11}(t) + \eta\Sigma_X\tilde{E}(t)W_2(t),$$
$$W_{12}(t+1) = W_{12}(t),$$
$$W_2(t+1) = W_2(t) + \eta W_{11}^\top(t)\Sigma_X\tilde{E}(t). \tag{51}$$

The update of W_{11}, W_2 is the same to the following problem

$$\min_{W_{11},W_2} L(W_{11},W_2) = \frac{1}{2}\|U_1^\top Y - \Sigma_X W_{11}W_2\|_F^2, \tag{52}$$

where Σ_X is a sqaure matrix of full rank. The above problem takes the same form as equation 47 where $\tilde{Y} = U_1^\top Y, \tilde{X} = \Sigma_X, \tilde{W}_1 = W_{11}, \tilde{W}_2 = W_2$.

B PROOF OF LEMMA 3.1

In this section, we present detailed proof of Lemma 3.1.

Lemma 3.1. *If at the t-th iteration of GD applied to the over-parametrized loss L, the step size η satisfies*

$$\lambda_{\min}(\tau_t) - \eta\|\nabla\ell(t)\|_F\sigma_{\max}(W(t))$$
$$- \frac{K\eta}{2}\left[\lambda_{\max}(\tau_t) + \eta\|\nabla\ell(t)\|_F\sigma_{\max}(W(t))\right]^2 \geq 0, \tag{53}$$

then the following inequality holds

$$L(t+1) \leq \rho(\eta,t)L(t), \tag{54}$$

where

$$\rho(\eta,t) = 1 - 2\eta\mu\lambda_{\min}(\tau_t) + K\mu\eta^2\lambda_{\max}^2(\tau_t)$$
$$+ 2\eta^2\mu\sigma_{\max}(W(t))\|\nabla\ell(t)\|_F$$
$$+ 2\eta^3\mu K\lambda_{\max}(\tau_t)\sigma_{\max}(W(t))\|\nabla\ell(t)\|_F$$
$$+ \eta^4\mu K\sigma_{\max}^2(W(t))\|\nabla\ell(t)\|_F^2. \tag{55}$$

Proof. Applying smoothness equation 13to the update of the product in equation 18, we get

$$
\begin{aligned}
L(t+1) \leq\ & L(t) - \eta\langle\nabla\ell(t), \tau_t(\nabla\ell(t)) - \eta\nabla\ell(t)W(t)^\top\nabla\ell(t)\rangle \\
& + \frac{K}{2}\eta^2\|\tau_t(\nabla\ell(t)) - \eta\nabla\ell(t)W(t)^\top\nabla\ell(t)\|_F^2 \\
=\ & L(t) - \eta\langle\nabla\ell(t), \tau_t(\nabla\ell(t))\rangle \\
& + \eta^2\big(\langle\nabla\ell(t), \nabla\ell(t)W(t)^\top\nabla\ell(t)\rangle + \tfrac{K}{2}\|\tau_t(\nabla\ell(t))\|_F^2\big) \\
& - \eta^3 K\langle\tau_t(\nabla\ell(t)), \nabla\ell(t)W(t)^\top\nabla\ell(t)\rangle \\
& + \frac{K}{2}\eta^4\|\nabla\ell(t)W(t)^\top\nabla\ell(t)\|_F^2
\end{aligned}
\tag{56}
$$

Then, we upper bound each term in the above inequality separately. First, since τ_t is a positive semi-definite operator, we have

$$
\begin{aligned}
\langle\nabla\ell(t), \tau_t(\nabla\ell(t))\rangle &\geq \lambda_{\min}(\tau_t)\|\nabla\ell(t)\|_F^2 \\
\|\tau_t(\nabla\ell(t))\|_F^2 &\leq \lambda_{\max}^2(\tau_t)\|\nabla\ell(t)\|_F^2
\end{aligned}
\tag{57}
$$

Then, using the sub-multiplicative property of Frobenius norm and Cauchy Schwartz inequality, we can bound the rest terms in equation 56

$$
\begin{aligned}
|\langle\nabla\ell(t), \nabla\ell(t)W(t)^\top\nabla\ell(t)\rangle| &\leq \|\nabla\ell(t)\|_F\|\nabla\ell(t)W(t)^\top\nabla\ell(t)\|_F \leq \|\nabla\ell(t)\|_F^3\sigma_{\max}(W(t)) \\
|\langle\tau_t(\nabla\ell(t)), \nabla\ell(t)W(t)^\top\nabla\ell(t)\rangle| &\leq \|\nabla\ell(t)\|_F\sigma_{\max}(W(t))\langle\nabla\ell(t), \tau_t(\nabla\ell(t))\rangle \leq \lambda_{\max}(\tau_t)\sigma_{\max}(W(t))\|\nabla\ell(t)\|_F^3 \\
\|\nabla\ell(t)W(t)^\top\nabla\ell(t)\|_F^2 &\leq \sigma_{\max}^2(W(t))\|\nabla\ell(t)\|_F^4.
\end{aligned}
\tag{58}
$$

Based on above results, we can further upper bound equation 56

$$
L(t+1) \leq L(t) - \eta\langle\nabla\ell(t), \tau_t(\nabla\ell(t))\rangle \tag{59}
$$
$$
+ \eta^2\big(\langle\nabla\ell(t), \nabla\ell(t)W(t)^\top\nabla\ell(t)\rangle + \tfrac{K}{2}\|\tau_t(\nabla\ell(t))\|_F^2\big) \tag{60}
$$
$$
- \eta^3 K\langle\tau_t(\nabla\ell(t)), \nabla\ell(t)W(t)^\top\nabla\ell(t)\rangle \tag{61}
$$
$$
+ \frac{K}{2}\eta^4\|\nabla\ell(t)W(t)^\top\nabla\ell(t)\|_F^2 \tag{62}
$$
$$
\leq L(t) - \eta\lambda_{\min}(\tau_t)\|\nabla\ell(t)\|_F^2 \tag{63}
$$
$$
+ \eta^2\big(\sigma_{\max}(W(t))\|\nabla\ell(t)\|_F^3 + \frac{K}{2}\lambda_{\max}^2(\tau_t)\|\nabla\ell(t)\|_F^2\big) \tag{64}
$$
$$
+ \eta^3 K\lambda_{\max}(\tau_t)\sigma_{\max}(W(t))\|\nabla\ell(t)\|_F^3 \tag{65}
$$
$$
+ \eta^4\frac{K}{2}\sigma_{\max}^2(W(t))\|\nabla\ell(t)\|_F^4 \tag{66}
$$
$$
= L(t) - \eta\|\nabla\ell(t)\|_F^2 g(\eta) \tag{67}
$$

where

$$
\begin{aligned}
g(\eta) =\ & \lambda_{\min}(\tau_t) - \eta\big(\sigma_{\max}(W(t))\|\nabla\ell(t)\|_F + \frac{K}{2}\lambda_{\max}^2(\tau_t)\big) \\
& - \eta^2 K\lambda_{\max}(\tau_t)\sigma_{\max}(W(t))\|\nabla\ell(t)\|_F - \eta^3\frac{K}{2}\sigma_{\max}^2(W(t))\|\nabla\ell(t)\|_F^2.
\end{aligned}
\tag{68}
$$

When $g(\eta) > 0$, which is assumed in equation 19, we apply PL condition $\frac{1}{2}\|\nabla\ell(t)\|_F^2 \geq \mu\ell(t)$ to the above equation to get

$$
\begin{aligned}
L(t+1) \leq L(t) \times \Big\{ & 1 - 2\eta\mu\lambda_{min}(\tau_t) \\
& + 2\eta^2\mu\big(\sigma_{\max}(W(t))\|\nabla\ell(t)\|_F + \frac{K}{2}\lambda_{\max}^2(\tau_t)\big) \\
& + 2\eta^3\mu K\lambda_{\max}(\tau_t)\sigma_{\max}(W(t))\|\nabla\ell(t)\|_F \\
& + \eta^4\mu K\sigma_{\max}^2(W(t))\|\nabla\ell(t)\|_F^2 \Big\} \\
= \rho(\eta, t)L(t). &
\end{aligned}
\tag{69}
$$

□

C PROOF OF THEOREM 3.1 AND THEOREM 3.2

Here we prove a new Theorem which implies Theorem 3.1 and Theorem 3.2.

Theorem C.1. *Under the assumptions in Theorem 3.1, for any $0 < \eta \leq \eta_{\max} := \min\{\eta_1^{\max}, \eta_2^{\max}\}$, the following four properties hold for all $t = 0, 1, 2, \cdots$.*

- $A_1(t) : L(t) \leq f(\eta, t)L(t-1)$, *where* $f(\eta, t) = 1 - a_1\eta + a_2(t)\eta^2 + a_3(t)\eta^3 + a_4(t)\eta^4 < 1$.

- $A_2(t) : p_1 \leq \sigma_{\min}(W(t)) \leq \sigma_{\max}(W(t)) \leq p_2$.

- $A_3(t) : \|D(t) - D(0)\|_F \leq \frac{2c_2\beta_0\sigma_{\max}^2(X)L(0)\eta^2}{1-f(\eta,0)}$ *when* $\eta < \eta_{\max}$.

- $A_4(t) : c_1\alpha_0 \leq \sigma_{\min}^2(W_1(t)) + \sigma_{\min}^2(W_2(t)) \leq \lambda_{\min}(\tau_t) \leq \lambda_{\max}(\tau_t) \leq \sigma_{\max}^2(W_1(t)) + \sigma_{\max}^2(W_2(t)) \leq c_2\beta_0$.

Notice Theorem 3.1 is Property $A_2(t), A_4(t)$, and Theorem 3.2 is implied by Property $A_1(t)$ because when $L(k) \leq L(0)$ hold for all $k = 0, 1, \cdots, t$, we have $a_2(k) \leq a_2(0), a_3(k) \leq a_3(0), a_4(k) \leq a_4(0)$. Thus, $f(\eta, k) \leq f(\eta, 0)$. As a result, the following inequality holds

$$L(t) \leq f(\eta, t)L(t-1) \leq L(0)\prod_{k=0}^{t-1} f(\eta, k) \leq f(\eta, 0)^t L(0). \tag{70}$$

Before proving Theorem C.1, we first present several **preliminary lemmas**.

Lemma C.1. *For matrix A, B, we have*

$$\sigma_{\min}^2(A)\|B\|_F^2 \leq \|AB\|_F^2 \leq \sigma_{\max}^2(A)\|B\|_F^2$$
$$\sigma_{\min}^2(B)\|A\|_F^2 \leq \|AB\|_F^2 \leq \sigma_{\max}^2(B)\|A\|_F^2. \tag{71}$$

Proof.

$$\begin{aligned}
\|AB\|_F^2 &= \operatorname{tr}\left(ABB^\top A^\top\right) \\
&= \operatorname{tr}\left(A^\top ABB^\top\right) &&\text{use cyclic property of trace} \\
&\leq \lambda_{\max}\left(A^\top A\right)\|B\|_F^2 &&\text{use trace inequality} \\
&= \sigma_{\max}^2(A)\|B\|_F^2.
\end{aligned} \tag{72}$$

For the other way

$$\begin{aligned}
\|AB\|_F^2 &= \operatorname{tr}\left(ABB^\top A^\top\right) \\
&= \operatorname{tr}\left(A^\top ABB^\top\right) \\
&\leq \lambda_{\max}\left(BB^\top\right)\|A\|_F^2 \\
&= \sigma_{\max}^2(B)\|A\|_F^2.
\end{aligned} \tag{73}$$

The lower bound is similar. □

Lemma C.2. *Let $X \in \mathbb{R}^{N \times n}, Y \in \mathbb{R}^{N \times m}$. Assume $N \leq n$ and $rank(X) = N$. For arbitrary $W \in \mathbb{R}^{n \times m}$, the following holds for $\ell(W) = \frac{1}{2}\|Y - XW\|_F^2$*

$$2\sigma_{\min}^2(X)\ell(W) \leq \|\nabla\ell(W)\|_F^2 \leq 2\sigma_{\max}^2(X)\ell(W). \tag{74}$$

Proof. The first inequality is PL inequality. We then prove the second

$$\begin{aligned}
\|\nabla\ell(W)\|_F^2 &= \|X^\top(Y - XW)\|_F^2 &&\text{gradient calculation} \\
&\leq \sigma_{\max}^2(X)\|Y - XW\|_F^2 &&\text{use Lemma C.1} \\
&= 2\sigma_{\max}^2(X)\ell(W).
\end{aligned} \tag{75}$$

□

Lemma C.3. *The difference of the imbalance between iteration $t+1$ and t can be upper bounded by*

$$\|D(t+1) - D(t)\|_F \leq 2\eta^2 \sigma_{\max}^2(X)\big(\sigma_{\max}^2(W_1(t)) + \sigma_{\max}^2(W_2(t))\big)L(t). \tag{76}$$

Proof. Notice the definition of imbalance is $D(t) := W_1^\top(t)W_1(t) - W_2(t)W_2^\top(t)$ and the update of GD is given in equation 4. Thus, using both results, one has

$$
\begin{aligned}
D(t+1) &= \big(W_1(t) - \eta\nabla\ell(t)W_2(t)^\top\big)^\top \big(W_1(t) - \eta\nabla\ell(t)W_2(t)^\top\big) \qquad \text{plug in GD update} \\
&\quad - \big(W_2(t) - \eta W_1(t)^\top\nabla\ell(t)\big)\big(W_2(t) - \eta W_1(t)^\top\nabla\ell(t)\big)^\top \\
&= D(t) + \eta^2\big(W_2(t)\nabla\ell(t)^\top\nabla\ell(t)W_2(t)^\top - W_1(t)^\top\nabla\ell(t)\nabla\ell(t)^\top W_1(t)\big).
\end{aligned} \tag{77}
$$

Then, we can upper bound $\|D(t+1) - D(t)\|_F$ using Lemma C.1 and Lemma C.2

$$
\begin{aligned}
\|D(t+1) - D(t)\|_F &= \eta^2\|W_2(t)\nabla\ell(t)^\top\nabla\ell(t)W_2(t)^\top - W_1(t)^\top\nabla\ell(t)\nabla\ell(t)^\top W_1(t)\|_F \\
&\leq \eta^2\big(\|W_2(t)\nabla\ell(t)^\top\nabla\ell(t)W_2(t)^\top\|_F + \|W_1(t)^\top\nabla\ell(t)\nabla\ell(t)^\top W_1(t)\|_F\big) \\
&\leq \eta^2\big(\|W_2(t)\nabla\ell(t)^\top\|_F^2 + \|W_1(t)^\top\nabla\ell(t)\|_F^2\big) \qquad \text{by Lemma C.1} \\
&\leq \eta^2\big(\sigma_{\max}^2(W_1(t)) + \sigma_{\max}^2(W_2(t))\big)\|\nabla\ell(t)\|_F^2 \qquad \text{by Lemma C.2} \\
&\leq 2\eta^2\sigma_{\max}^2(X)\big(\sigma_{\max}^2(W_1(t)) + \sigma_{\max}^2(W_2(t))\big)L(t).
\end{aligned} \tag{78}
$$

$\square$

Lemma C.4. *Suppose $h > \min\{r, m\}$. Given any $A \in \mathbb{R}^{r\times h}, B \in \mathbb{R}^{h\times m}$ that satisfy $A^\top A - BB^\top = D$, we have*

$$\lambda_m(B^\top B) \geq \frac{-\bar{\lambda} + \underline{\lambda} + \sqrt{(\bar{\lambda} + \underline{\lambda})^2 + 4\sigma_m^2(AB)}}{2} \tag{79}$$

where $\bar{\lambda} = \max\{\lambda_1(D), 0\}$ and $\underline{\lambda} = \max\{\lambda_m(-D), 0\}$.

Lemma C.4 is cited from Min et al. [2022] and the proof can be found in Min et al. [2022] Lemma 8.

Lemma C.5. *Suppose $h > \min\{r, m\}$. Given any $A \in \mathbb{R}^{r\times h}, B \in \mathbb{R}^{h\times m}$ that satisfy $A^\top A - BB^\top = D$, we have*

$$\lambda_{\max}(B^\top B) \leq \frac{\max(\lambda_{\max}(-D), 0) + \sqrt{\max(\lambda_{\max}(-D), 0)^2 + 4\sigma_{\max}^2(AB)}}{2} \tag{80}$$

Proof. We first choose $z \in \mathbb{R}^m$ with $\|z\|_2 = 1$ s.t.

$$z^\top B^\top Bz = \lambda_{\max}(B^\top B). \tag{81}$$

Then, we have

$$
\begin{aligned}
\lambda_{\max}^2(B^\top B) - z^\top B^\top A^\top ABz &= z^\top B^\top BB^\top Bz - z^\top B^\top A^\top ABz \\
&= z^\top (B^\top BB^\top B - B^\top A^\top AB)z \\
&= z^\top B^\top (BB^\top - A^\top A)Bz \\
&= z^\top B^\top (-D)Bz.
\end{aligned} \tag{82}
$$

Notice

$$
\begin{aligned}
\lambda_{\max}^2(B^\top B) - z^\top B^\top A^\top ABz &\geq \lambda_{\max}^2(B^\top B) - \sigma_{\max}^2(AB) \\
z^\top B^\top (-D)Bz &\leq \max(\lambda_{\max}(-D), 0)\|Bz\|_2^2 \leq \max(\lambda_{\max}(-D), 0)\lambda_{\max}(B^\top B).
\end{aligned} \tag{83}
$$

Thus, we have

$$\lambda_{\max}(B^\top B)^2 - \sigma_{\max}^2(AB) \leq \max(\lambda_{\max}(-D), 0)\lambda_{\max}(B^\top B). \tag{84}$$

The solution to the above inequality gives us the results. $\square$

Then, we begin the proof of Theorem C.1.

Proof. Assume $A_1(k), A_2(k), A_3(k), A_4(k)$ hold at iteration $k = 1, 2, \cdots, t$, then we prove they all hold at iteration $t+1$. First, **we prove** $A_1(t+1)$ **hold**. According to Lemma 3.1, we have

$$
\begin{aligned}
L(t+1) \leq L(t) \times \Big\{ &1 - 2\eta\mu\lambda_{min}(\tau_t) \\
&+ 2\eta^2\mu\big(\sigma_{\max}(W(t))\|\nabla\ell(t)\|_F + \frac{K}{2}\lambda_{\max}^2(\tau_t)\big) \\
&+ 2\eta^3\mu K\lambda_{\max}(\tau_t)\sigma_{\max}(W(t))\|\nabla\ell(t)\|_F \\
&+ \eta^4\mu K\sigma_{\max}^2(W(t))\|\nabla\ell(t)\|_F^2 \Big\}
\end{aligned}
\tag{85}
$$

Since $A_2(t), A_4(t)$ hold, we can further upper bound the above inequality

$$
\begin{aligned}
L(t+1) \leq L(t) \times \Big\{ &1 - 2\eta\mu c_1\alpha_0 + 2\eta^2\mu\big(p_2\|\nabla\ell(t)\|_F + \frac{K}{2}(c_1\beta_0)^2\big) \\
&+ 2\eta^3\mu K c_1\beta_0 p_2\|\nabla\ell(t)\|_F + \eta^4\mu K p_2^2\|\nabla\ell(t)\|_F^2 \Big\}
\end{aligned}
\tag{86}
$$

Apply Lemma C.2

$$
\begin{aligned}
L(t+1) &\leq L(t) \times \Big\{ 1 - 2\eta\mu c_1\alpha_0 + 2\eta^2\mu\big(p_2\sqrt{2\sigma_{\max}^2(X)L(t)} + \frac{\sigma_{\max}^2(X)}{2}(c_2\beta_0)^2\big) \\
&\quad + 2\eta^3\mu\sigma_{\max}^2(X)c_2\beta_0 p_2\sqrt{2\sigma_{\max}^2(X)L(t)} + 2\eta^4\mu\sigma_{\max}^4(X)p_2^2 L(t) \Big\} \\
&= L(t) \times \Big\{ 1 - 2\eta\sigma_{\min}^2(X)c_1\alpha_0 + 2\eta^2\big(p_2\sqrt{2\kappa\sigma_{\min}^6(X)L(t)} + \frac{\kappa\sigma_{\min}^4(X)}{2}(c_2\beta_0)^2\big) \\
&\quad + 2\eta^3 c_2\beta_0 p_2\sqrt{2\kappa^3\sigma_{\min}^{10}(X)L(t)} + 2\eta^4 p_2^2\kappa^2\sigma_{\min}^6(X)L(t) \Big\} \\
&= L(t) \times [1 - a_1\eta + a_2(t)\eta^2 + a_3(t)\eta^3 + a_4(t)\eta^4]
\end{aligned}
\tag{87}
$$

Finally, we show when $0 < \eta \leq \eta_{\max}, f(\eta, t) < 1$. Notice $f(\eta, t)$ is a decreasing functions in t, it suffices to show $f(\eta, 0) < 1$

$$
f(\eta, 0) < 1 \iff a_4(0)\eta^3 + a_3(0)\eta^2 + a_2(0)\eta < a_1.
\tag{88}
$$

Compare the above inequality with equation 30, one has

$$
\begin{aligned}
a_4(0)\eta^3 + a_3(0)\eta^2 + a_2(0)\eta &< a_4(0)\eta^3 + a_3(0)\eta^2 + \big(a_2(0) + \frac{4c_2 L(0)\sigma_{\max}^2(X)}{c_2 - 1}\big)\eta \\
a_4(0)\eta^3 + a_3(0)\eta^2 + a_2(0)\eta &< a_4(0)\eta^3 + a_3(0)\eta^2 + \big(a_2(0) + \frac{8c_2\beta_0 L(0)\sigma_{\max}^2(X)}{(1 - c_1)\alpha_0}\big)\eta.
\end{aligned}
\tag{89}
$$

Thus, when $0 < \eta \leq \eta_{\max}$, we have

$$
\begin{aligned}
a_4(0)\eta^3 + a_3(0)\eta^2 + a_2(0)\eta &< a_4(0)\eta^3 + a_3(0)\eta^2 + \big(a_2(0) + \frac{4c_2 L(0)\sigma_{\max}^2(X)}{c_2 - 1}\big)\eta \leq a_1 \\
a_4(0)\eta^3 + a_3(0)\eta^2 + a_2(0)\eta &< a_4(0)\eta^3 + a_3(0)\eta^2 + \big(a_2(0) + \frac{8c_2\beta_0 L(0)\sigma_{\max}^2(X)}{(1 - c_1)\alpha_0}\big)\eta \leq a_1.
\end{aligned}
\tag{90}
$$

which is equivalent to $f(\eta, 0) < 1$. Thus, $A_1(t+1)$ is proved.

Then, **we prove** $A_2(t+1)$ **hold**. Since loss is decreasing, i.e. $L(t+1) \leq L(t) \leq L(0)$, we have

$$
\|Y - XW(t+1)\|_F \leq \|E(0)\|_F.
\tag{91}
$$

equation 91 is equivalent to

$$\|Y\|_F - \|E(0)\|_F \leq \|XW(t+1)\|_F \leq \|Y\|_F + \|E(0)\|_F. \tag{92}$$

In Min et al. [2022], Theorem 3, the lower bound is proved. For the upper bound,

$$\sigma_{\max}(W(t+1))\sigma_{\min}(X) \leq \|W(t+1)\|_F\sigma_{\min}(X) \leq \|XW(t+1)\|_F \leq \|Y\|_F + \|E(0)\|_F, \tag{93}$$

Thus,

$$\sigma_{\max}(W(t+1)) \leq \frac{\|Y\|_F + \|E(0)\|_F}{\sigma_{\min}(X)} =: p_2. \tag{94}$$

Then, **we prove** $A_3(t+1)$ **hold**.

$$
\begin{aligned}
\|D(t+1) - D(0)\|_F &\leq \sum_{k=0}^{t}\|D(K+1) - D(k)\|_F \qquad \text{use Lemma C.3} \\
&\leq \sum_{k=0}^{t} 2\eta^2\sigma_{\max}^2(X)\big(\sigma_{\max}^2(W_1(k)) + \sigma_{\max}^2(W_2(k))\big)L(k) \qquad \text{use } A_4(k) \\
&\leq 2\eta^2\sigma_{\max}^2(X)c_2\beta_0 \sum_{k=0}^{t} L(k) \qquad \text{use } A_1(k) \\
&\leq 2\eta^2 c_2\beta_0\sigma_{\max}^2(X) \sum_{k=0}^{t}(1 - a_1\eta + a_2(k)\eta^2 + a_3(k)\eta^3 + a_4(k)\eta^4)^k L(0) \qquad (95) \\
&\leq 2\eta^2 c_2\beta_0\sigma_{\max}^2(X) \sum_{k=0}^{t}(1 - a_1\eta + a_2(0)\eta^2 + a_3(0)\eta^3 + a_4(0)\eta^4)^k L(0) \\
&\leq \frac{2\eta^2 c_2\beta_0\sigma_{\max}^2(X)L(0)}{1 - f(\eta,0)}. \qquad (96)
\end{aligned}
$$

where we upper bound $a_i(k)$ by $a_i(0)$ in equation 95 for $i = 1, 2, 3, 4$.

Finally, **we prove** $A_4(t+1)$ **hold**. $\lambda_{\min}(\tau) \leq \lambda_{\max}(\tau)$ is obvious. We begin with the second inequality

$$
\begin{aligned}
\lambda_{\min}(\tau) &= \min_{\|W\|_F=1} \langle W, WW_2^\top W_2 + W_1 W_1^\top W \rangle \qquad \text{definition of operator norm} \\
&\geq \min_{\|W\|_F=1} \langle W, WW_2^\top W_2 \rangle + \min_{\|W\|_F=1} \langle W, W_1 W_1^\top W \rangle \\
&= \sigma_{\min}^2(W_1) + \sigma_{\min}^2(W_2). \qquad (97)
\end{aligned}
$$

The fourth inequality can be proved similarly

$$
\begin{aligned}
\lambda_{\max}(\tau) &= \max_{\|W\|_F=1} \langle W, WW_2^\top W_2 + W_1 W_1^\top W \rangle \\
&\leq \max_{\|W\|_F=1} \langle W, WW_2^\top W_2 \rangle + \max_{\|W\|_F=1} \langle W, W_1 W_1^\top W \rangle \\
&= \sigma_{\max}^2(W_1) + \sigma_{\max}^2(W_2) \qquad (98)
\end{aligned}
$$

Then, we prove the first inequality and last inequality holds. According to Lemma C.4, we have

$$\sigma_{\min}^2(W_1(t+1)) \geq \frac{-\bar{\lambda}_-(t+1) + \underline{\lambda}_+(t+1) + \sqrt{\big(\bar{\lambda}_-(t+1) + \underline{\lambda}_+(t+1)\big)^2 + 4\sigma_n^2(W(t+1))}}{2}.$$

$$\sigma_{\min}^2(W_2(t+1)) \geq \frac{-\bar{\lambda}_+(t+1) + \underline{\lambda}_-(t+1) + \sqrt{\big(\bar{\lambda}_+(t+1) + \underline{\lambda}_-(t+1)\big)^2 + 4\sigma_m^2(W(t+1))}}{2}. \tag{99}$$

where

$$
\begin{aligned}
\bar{\lambda}_+(t) &= \max\left(\lambda_1(D(t)), 0\right) \\
\underline{\lambda}_-(t) &= \max\left(\lambda_m(-D(t)), 0\right) \\
\bar{\lambda}_-(t) &= \max\left(\lambda_1(-D(t)), 0\right) \\
\underline{\lambda}_+(t) &= \max\left(\lambda_n(D(t)), 0\right)
\end{aligned}
\tag{100}
$$

We define

$$
\begin{aligned}
h_1(\Delta_1, \Delta_2) &:= \frac{-\bar{\lambda}_-(0) + \Delta_1 + \underline{\lambda}_+(0) + \Delta_2 + \sqrt{\left(\bar{\lambda}_-(0) + \underline{\lambda}_+(0) + \Delta_1 + \Delta_2\right)^2 + 4p_1^2}}{2} \\
h_2(\Delta_3, \Delta_4) &= \frac{-\bar{\lambda}_+(0) + \Delta_1 + \underline{\lambda}_-(0) + \Delta_2 + \sqrt{\left(\bar{\lambda}_+(0) + \underline{\lambda}_-(0) + \Delta_1 + \Delta_2\right)^2 + 4p_1^2}}{2}
\end{aligned}
\tag{101}
$$

where

$$
\begin{aligned}
\Delta_1 &= \bar{\lambda}_-(t+1) - \bar{\lambda}_-(0) \\
\Delta_2 &= \underline{\lambda}_+(t+1) - \underline{\lambda}_+(0) \\
\Delta_3 &= \bar{\lambda}_+(t+1) - \bar{\lambda}_+(0) \\
\Delta_4 &= \underline{\lambda}_-(t+1) - \underline{\lambda}_-(0).
\end{aligned}
\tag{102}
$$

Then, we use $\sigma_{\min}(W(t+1)) \geq p_1$ to lower bound equation 99

$$
\begin{aligned}
\sigma_{\min}^2(W_1(t+1)) &\geq \frac{-\bar{\lambda}_-(t+1) + \underline{\lambda}_+(t+1) + \sqrt{\left(\bar{\lambda}_-(t+1) + \underline{\lambda}_+(t+1)\right)^2 + 4\sigma_n^2\left(W(t+1)\right)}}{2}. \\
&\geq \frac{-\bar{\lambda}_-(t+1) + \underline{\lambda}_+(t+1) + \sqrt{\left(\bar{\lambda}_-(t+1) + \underline{\lambda}_+(t+1)\right)^2 + 4p_1^2}}{2}. \\
&:= h_1(\Delta_1, \Delta_2).
\end{aligned}
\tag{103}
$$

Similarly, we have

$$
\sigma_{\min}^2(W_2(t+1)) \geq h_2(\Delta_3, \Delta_4).
\tag{104}
$$

Notice $h_1(0,0) + h_2(0,0) = \alpha_0$ which is independent of t. Our goal is to lower bound $h_1(\Delta_1, \Delta_2) + h_2(\Delta_3, \Delta_4)$ using $h_1(0,0) + h_2(0,0)$. A natural solution is that if we can quantify how large $|\Delta_k|, k = 1, 2, 3, 4$ is, i.e. $|\Delta_k| \leq \Delta_h$, and if we can show $h_1(\cdot, \cdot), h_2(\cdot, \cdot)$ are both L_h-Lipschitz continuous. Using these two ingredients, one can show

$$
\begin{aligned}
&|h_1(\Delta_1, \Delta_2) - h_1(0,0)| \leq L_h \sqrt{\Delta_1^2 + \Delta_2^2} \\
&\Rightarrow h_1(\Delta_1, \Delta_2) \geq h_1(0,0) - L_h \sqrt{\Delta_1^2 + \Delta_2^2} \geq h_1(0,0) - \sqrt{2} L_h \Delta_h.
\end{aligned}
\tag{105}
$$

Similarly, we have

$$
h_2(\Delta_3, \Delta_4) \geq h_2(0,0) - \sqrt{2} L_h \Delta_h.
\tag{106}
$$

Based on above two equations, one has

$$
h_1(\Delta_1, \Delta_2) + h_2(\Delta_3, \Delta_4) \geq h_1(0,0) + h_2(0,0) - 2\sqrt{2} L_h \Delta_h.
\tag{107}
$$

Next, we show the above two assumptions hold

1. $h_1(\cdot, \cdot), h_2(\cdot, \cdot)$ are both L_h-Lipschitz continuous.

2. $|\Delta_k| \leq \Delta_h$ hold for all $k = 1, 2, 3, 4$.

For the first one, using Weyl's inequality and Property $A_3(t+1)$, we can upper bound $|\Delta_k|$

$$
\begin{aligned}
|\Delta_1| &= \left|\max(\lambda_1(-D(t+1)), 0) - \max(\lambda_1(-D(0)), 0)\right| \\
&\leq \left|\lambda_1(-D(t+1)) - \lambda_1(-D(0)), 0\right| \qquad \text{use Weyl's inequality} \\
&\leq \|D(t+1) - D(0)\|_F \qquad \text{use Lemma C.3} \\
&\leq \frac{2\eta^2 c_2 \beta_0 \sigma_{\max}^2(X) L(0)}{1 - f(\eta, 0)}.
\end{aligned}
\tag{108}
$$

Similarly, we have

$$
|\Delta_2|, |\Delta_3|, |\Delta_4| \leq \|D(t+1) - D(0)\|_F \leq \frac{2\eta^2 c_2 \beta_0 \sigma_{\max}^2(X) L(0)}{1 - f(\eta, 0)}.
\tag{109}
$$

What's more,

$$
\begin{aligned}
\left|\frac{dh_1(x, y)}{dx}\right| &= \left|-\frac{1}{2} + \frac{x + y + \bar{\lambda}_-(0) + \underline{\lambda}_+(0)}{2\sqrt{\left(\bar{\lambda}_-(0) + \underline{\lambda}_+(0) + x + y\right)^2 + 4p_1^2}}\right| \\
&\leq \frac{1}{2} + \left|\frac{x + y + \bar{\lambda}_-(0) + \underline{\lambda}_+(0)}{2\sqrt{\left(\bar{\lambda}_-(0) + \underline{\lambda}_+(0) + x + y\right)^2 + 4p_1^2}}\right| \\
&\leq \frac{1}{2} + \frac{1}{2} \\
&\leq 1.
\end{aligned}
\tag{110}
$$

Similarly, we have $\left|\frac{dh_1(x,y)}{dy}\right|, \left|\frac{dh_2(x,y)}{dx}\right|, \left|\frac{dh_2(x,y)}{dy}\right| \leq 1$. Combine with equation 110, we have $h_1(\cdot, \cdot), h_2(\cdot, \cdot)$ are $\sqrt{2}$-Lipschitz continuous. Thus, we have

$$
\begin{aligned}
\sigma_{\min}^2(W_1(t)) + \sigma_{\min}^2(W_2(t)) &\geq h_1(\Delta_1, \Delta_2) + h_2(\Delta_3, \Delta_4) \\
&\geq \alpha_0 - 2L_h \sqrt{2} \|D(t+1) - D(0)\|_F \qquad L_h = \sqrt{2} \\
&\geq \alpha_0 - \frac{8\eta^2 c_2 \beta_0 \sigma_{\max}^2(X) L(0)}{1 - f(\eta, 0)}.
\end{aligned}
\tag{111}
$$

Although the above lower bound is smaller than α_0, it is close to α_0 when η is small. This motivates us to introduce $0 < c_1 < 1$ so that when η is small, the above inequality is lower bounded by $c_1 \alpha_0$. To derive the upper bound on η, it is equivalent to ensure

$$
\begin{aligned}
&\alpha_0 - \frac{8\eta^2 c_2 \beta_0 \sigma_{\max}^2(X) L(0)}{1 - f(\eta, 0)} \geq c_1 \alpha_0 \\
\iff\ & (1 - c_1)\alpha_0 \geq \frac{8\eta^2 c_2 \beta_0 \sigma_{\max}^2(X) L(0)}{1 - f(\eta, 0)} \\
\iff\ & (1 - c_1)\alpha_0 \geq \frac{8\eta c_2 \beta_0 \sigma_{\max}^2(X) L(0)}{a_1 - a_2(0)\eta - a_3(0)\eta^2 - a_4(0)\eta^3} \\
\iff\ & a_4(0)\eta^3 + a_3(0)\eta^2 + \left(a_2(0) + \frac{8c_2 \beta_0 L(0)\sigma_{\max}^2(X)}{(1 - c_1)\alpha_0}\right)\eta \leq a_1
\end{aligned}
\tag{112}
$$

which is ensured when $0 < \eta < \eta_{\max}$.

The proof for the fourth inequality $\sigma_{\max}^2(W_1(t+1)) + \sigma_{\max}^2(W_2(t+1)) \leq c_2 \beta_0$ in $A_4(t+1)$ is similar. According to Lemma C.5, we have

$$
\begin{aligned}
&\sigma_{\max}^2(W_1(t+1)) + \sigma_{\max}^2(W_2(t+1)) \\
&\leq \frac{\max(\lambda_{\max}(D(0)), 0) + \Delta_3 + \sqrt{4\sigma_{\max}^2(W(t+1)) + [\max(\lambda_{\max}(D(0)), 0) + \Delta_3]^2}}{2} \\
&\quad + \frac{\max(\lambda_{\max}(-D(0)), 0) + \Delta_4 + \sqrt{4\sigma_{\max}^2(W(t+1) + [\max(\lambda_{\max}(-D(0)), 0) + \Delta_4]^2}}{2}
\end{aligned}
$$

$$\leq \frac{\max(\lambda_{\max}(D(0)), 0) + \Delta_3 + \sqrt{4p_2^2 + [\max(\lambda_{\max}(D(0)), 0) + \Delta_3]^2}}{2}$$

$$+ \frac{\max(\lambda_{\max}(-D(0)), 0) + \Delta_4 + \sqrt{4p_2^2 + [\max(\lambda_{\max}(-D(0)), 0) + \Delta_4]^2}}{2}$$

$$:= h_3(\Delta_5, \Delta_6), \tag{113}$$

where

$$\Delta_5 = \max(\lambda_{\max}(D(t+1)), 0) - \max(\lambda_{\max}(D(0)), 0)$$
$$\Delta_6 = \max(\lambda_{\max}(-D(t+1)), 0) - \max(\lambda_{\max}(-D(0)), 0). \tag{114}$$

Since

$$\left|\frac{dh_3(x,y)}{dx}\right| = \left|\frac{1}{2} + \frac{x + \max(\lambda_{\max}(D(t+1)), 0)}{2\sqrt{4p_2^2 + [\max(\lambda_{\max}(D(t+1)), 0) + \Delta_5]^2}}\right| \leq 1. \tag{115}$$

Similarly, $\left|\frac{dh_3(x,y)}{dy}\right| \leq 1$. What's more, Weyl's inequality gives us

$$|\Delta_5| = |\max(\lambda_{\max}(D(t+1)), 0) - \max(\lambda_{\max}(D(0)), 0)|$$
$$\leq |\lambda_{\max}(D(t+1)) - \lambda_{\max}(D(0))|$$
$$\leq \|D(t+1) - D(0)\|_F \tag{116}$$

Similarly, we have $|\Delta_6| \leq \|D(t+1) - D(0)\|_F$. Thus, we have

$$\sigma_{\max}^2(W_1(t+1)) + \sigma_{\max}^2(W_2(t+1)) = h_3(\Delta_5, \Delta_6)$$
$$\leq h_3(0,0) + \sqrt{2}\sqrt{\Delta_5^2 + \Delta_6^2}$$
$$\leq \beta_0 + \frac{4\eta^2 c_2 \beta_0 \sigma_{\max}^2(X) L(0)}{1 - f(\eta, 0)} \tag{117}$$
$$\leq \beta_0 c_2$$

where the last inequality holds if and only if

$$a_4 \eta^3 + a_3 \eta^2 + \left(a_2 + \frac{4c_2 L(0)\sigma_{\max}^2(X)}{c_2 - 1}\right)\eta \leq a_1. \tag{118}$$

$\square$

D   PROOF OF PROPOSITION 3.1

Proposition 3.1. *If $\alpha_0 > 0$, for all $0 < \eta \leq \eta_{\max}$ and for all $t = 0, 1, \cdots$, the following inequality holds*

$$f(\eta, t) \geq 1 - \frac{1}{\kappa} \tag{119}$$

where $\kappa = \frac{K}{\mu}$ is the condition number of the non-overparametrized Problem 1

Proof. The theoretical optimal convergence rate for non-overparametrized regime is $1 - \frac{1}{\kappa}$. Then

$$f(\eta, t) - (1 - \frac{1}{\kappa}) = \frac{1}{\kappa} - a_1 \eta + a_2(t)\eta^2 + a_3(t)\eta^3 + a_4(t)\eta^4 \qquad \text{drop last two terms which are non-negative}$$

$$\geq \frac{1}{\kappa} - 2c_1 \alpha_0 \sigma_{\min}^2(X)\eta + \left(2\sqrt{2\kappa L(t)\sigma_{\min}^6(X)}p_2 + \kappa\mu^2 c_2^2 \beta_0^2\right)\eta^2$$

$$\geq \frac{1}{\kappa} - 2c_1 \alpha_0 \sigma_{\min}^2(X)\eta + \kappa\sigma_{\min}^4(X)c_2^2 \beta_0^2 \eta^2 \qquad \text{use } \beta_0 \geq \alpha_0 \text{ to lower bound last term}$$

$$\geq \frac{1}{\kappa} - 2c_1 \alpha_0 \sigma_{\min}^2(X)\eta + \kappa\sigma_{\min}^4(X)c_2^2 \alpha_0^2 \eta^2$$

$$= \left(\frac{1}{\sqrt{\kappa}} - \sqrt{\kappa}\sigma_{\min}^2(X)c_2 \alpha_0 \eta\right)^2$$

$$\geq 0. \tag{120}$$

Thus, the results are proved. $\square$

E   PROOF OF CLAIM 3.1

Claim 3.1. *Suppose $\alpha_0 > 0$. Let η_t' be the unique positive root of the following equation*

$$\frac{dh(\eta, t)}{d\eta} = 0. \tag{121}$$

Then the solution to Problem 41 is $\eta_t = \min(\eta_t', \eta_{\max})$.

Proof. We first observe $h(\eta, t)$ takes the following general form

$$h(\eta, t) = 1 - h_1(t)\eta + h_2(t)\eta^2 + h_3(t)\eta^3 + h_4(t)\eta^4 \tag{122}$$

where $h_1(t), h_2(t), h_3(t), h_4(t) > 0$. Then, $\frac{dh(\eta,t)}{d\eta} = -h_1 + 2h_2\eta + 3h_3\eta^2 + 4h_4\eta^3$. Notice the derivative of $h(\eta, t)$ with respect to η is monotonically increasing when $\eta > 0$ and $\frac{d^2h(\eta,t)}{d\eta^2} > 0$. Thus, if $\eta_t' \leq \eta_{\max}$, the minimizer of Problem 41 is $\eta_{\max}$. If $\eta_t' \geq \eta_{\max}$, since $\frac{dh(\eta,t)}{d\eta}$ is negative when $0 < \eta \leq \eta_{\max} \leq \eta_t'$, $h(\eta, t)$ is decreasing in the same range. Thus, the minimizer is $\eta_{\max}$. Combing the above two cases, the minimizer of Problem 41 is

$$\eta_t = \min(\eta_t', \eta_{\max}). \tag{123}$$

$\square$

Claim E.1. *Given some $0 < c_1 < \frac{2}{3}$, pick any*

$$c_2 \geq \max\left\{ \frac{M + \frac{16L(0)}{\beta_0}}{c_1\alpha_0\sigma_{\min}^2(X)}, \sqrt{\frac{M + \frac{8\alpha_0 L(0)}{\beta_0^2}}{\alpha_0\sigma_{\min}^2(X)}}, 2 \right\}, \tag{124}$$

where $M = \frac{2\alpha_0^3 p_2^2 L(0)}{\beta_0^6 \kappa} + \frac{2\sqrt{2\sigma_{\min}^2(X)L(0)}p_2\alpha_0^2}{\sqrt{\kappa}\beta_0^3} + \frac{2\sqrt{2L(0)\sigma_{\min}^2(X)}p_2\alpha_0}{\beta_0^2\sqrt{\kappa}}$.

Such choice of c_1, c_2 ensures $\eta_{\max} \geq \eta_t^*$ for all $t = 0, 1, 2, \cdots$.

What's more, we have

$$\rho(\eta_t^*, t) \leq \hat{\rho}(\eta_t^*, t) \leq f(\eta_t^*, t) \simeq 1 - \frac{(c_1\alpha_0)^2}{(c_2\beta_0)^2}\frac{1}{\kappa}. \tag{125}$$

Remark E.1. *Claim E.1 implies for proper choice of c_1, c_2, one has $\eta_{\max} \geq \eta_t^*$ for all $t = 0, 1, 2, \cdots$. In the limiting case when $t \to \infty$, one has*

$$\lim_{t\to\infty} \rho(\eta, t) \leq \lim_{t\to\infty} \hat{\rho}(\eta, t) \leq \lim_{t\to\infty} f(\eta, t) = 1 - 2(c_1\alpha_0)\sigma_{\min}^2(X)\eta + \kappa\sigma_{\min}^4(X)(c_2\beta_0)^2\eta^2 \tag{126}$$

With the choice of c_1, c_2 specified, we have $\eta_\infty^* \leq \eta_{\max}$. Thus, the asymptotic convergence rate is

$$\rho(\eta_\infty^*, \infty) \leq \hat{\rho}(\eta_\infty^*, \infty) \leq f(\eta_\infty^*, \infty) = 1 - \frac{(c_1\alpha_0)^2}{(c_2\beta_0)^2}\frac{1}{\kappa} \tag{127}$$

The asymptotic convergence rate is determined by $\frac{c_1\alpha_0}{c_2\beta_0}$ and condition number κ. The smaller κ is, the faster convergence rate is. What's more, since $\frac{\lambda_{\min}(\tau_t)}{\lambda_{\max}(\tau_t)} \geq \frac{c_1\alpha_0}{c_2\beta_0}$, we can view $\frac{c_1\alpha_0}{c_2\beta_0}$ as a lower bound on the condition number of the operator τ_t. The more ill-conditioned τ_t is, i.e. $\frac{c_1\alpha_0}{c_2\beta_0}$ is small, the slower the convergence rate is.

Proof. Notice $a_2(t), a_3(t), a_4(t)$ depends on $L(t)$ and $L(t)$ decreases as t increases, so $a_2(t), a_3(t), a_4(t)$ decrease as t increase. From equation 122, we can see η_t^* increases as t increases. Thus, to prove $\eta_t^* \leq \eta_{\max}$, it suffices to show

$$\lim_{t\to\infty} \eta_t^* = \frac{c_1\alpha_0}{c_2^2\beta_0^2\kappa\sigma_{\min}^2(X)} \leq \eta_{\max}. \tag{128}$$

which is equivalent to the following inequalities

$$a_4(0)\left(\frac{c_1\alpha_0}{c_2^2\beta_0^2\kappa\sigma_{\min}^2(X)}\right)^3+a_3(0)\left(\frac{c_1\alpha_0}{c_2^2\beta_0^2\kappa\sigma_{\min}^2(X)}\right)^2+\left(a_2(0)+\frac{4c_2L(0)\sigma_{\max}^2(X)}{c_2-1}\right)\frac{c_1\alpha_0}{c_2^2\beta_0^2\kappa\sigma_{\min}^2(X)}\leq a_1, \qquad (129)$$

$$a_4(0)\left(\frac{c_1\alpha_0}{c_2^2\beta_0^2\kappa\sigma_{\min}^2(X)}\right)^3+a_3(0)\left(\frac{c_1\alpha_0}{c_2^2\beta_0^2\kappa\sigma_{\min}^2(X)}\right)^2+\left(a_2(0)+\frac{8c_2\beta_0L(0)\sigma_{\max}^2(X)}{(1-c_1)\alpha_0}\right)\frac{c_1\alpha_0}{c_2^2\beta_0^2\kappa\sigma_{\min}^2(X)}\leq a_1. \qquad (130)$$

For equation 129 to hold, we study its LHS

$$\begin{aligned}
\text{LHS of equation 129} &= \frac{2c_1^3\alpha_0^3p_2^2L(0)}{c_2^6\beta_0^6\kappa}+\frac{2\sqrt{2L(0)\sigma_{\min}^2(X)}p_2c_1^2\alpha_0^2}{\sqrt{\kappa}c_2^3\beta_0^3}\\
&\quad +\frac{4c_1\alpha_0L(0)}{(c_2-1)c_2\beta_0^2}+\frac{2\sqrt{2L(0)\sigma_{\min}^2(X)}p_2c_1\alpha_0}{c_2^2\beta_0^2\sqrt{\kappa}}+c_1\alpha_0\sigma_{\min}^2(X)\\
&= \frac{2c_1^3\alpha_0^3p_2^2L(0)}{c_2^6\beta_0^6\kappa}+\frac{Pc_1^2\alpha_0}{c_2^3\beta_0}+\frac{4c_1\alpha_0L(0)}{(c_2-1)c_2\beta_0^2}+\frac{Pc_1}{c_2^2}+c_1\alpha_0\sigma_{\min}^2(X)
\end{aligned}$$

where

$$P=\frac{2\sqrt{2L(0)\sigma_{\min}^2(X)}p_2\alpha_0}{\sqrt{\kappa}\beta_0^2}. \qquad (131)$$

Since $c_2\geq 2$, so $c_2-1\geq\frac{c_2}{2}$. Then, we upper bound the above equality by substituting higher order terms of $c_1^k, k\geq 2$ with c_1 in the numerator by one except for the last term and replace higher order terms of $c_2^k, k\geq 3$ with c_2^2,

$$\begin{aligned}
\text{LHS of equation 129} &= \frac{2c_1^3\alpha_0^3p_2^2L(0)}{c_2^6\beta_0^6\kappa}+\frac{Pc_1^2\alpha_0}{c_2^3\beta_0}+\frac{4c_1\alpha_0L(0)}{(c_2-1)c_2\beta_0^2}+\frac{Pc_1}{c_2^2}+c_1\alpha_0\sigma_{\min}^2(X)\\
&\leq \frac{2c_1^3\alpha_0^3p_2^2L(0)}{c_2^6\beta_0^6\kappa}+\frac{Pc_1^2\alpha_0}{c_2^3\beta_0}+\frac{8c_1\alpha_0L(0)}{c_2^2\beta_0^2}+\frac{Pc_1}{c_2^2}+c_1\alpha_0\sigma_{\min}^2(X) \qquad \text{use } c_2-1\geq\frac{c_2}{2} \text{ in the first term}\\
&\leq \frac{2c_1\alpha_0^3p_2^2L(0)}{c_2^2\beta_0^6\kappa}+\frac{Pc_1\alpha_0}{c_2^2\beta_0}+\frac{8c_1\alpha_0L(0)}{c_2^2\beta_0^2}+\frac{Pc_1}{c_2^2}+c_1\alpha_0\sigma_{\min}^2(X)\\
&\qquad \text{use } c_1\geq c_1^k, k\geq 2 \text{ in the numerator and } c_2^2\leq c_2^k, k\geq 3 \text{ in denominator}\\
&= \frac{c_1}{c_2^2}\left[\frac{2\alpha_0^3p_2^2L(0)}{\beta_0^6\kappa}+\frac{P\alpha_0}{\beta_0}+\frac{8\alpha_0L(0)}{\beta_0^2}+P\right]+c_1\alpha_0\sigma_{\min}^2(X)\\
&= \frac{c_1}{c_2^2}\left(M+\frac{8\alpha_0L(0)}{\beta_0^2}\right)+c_1\alpha_0\sigma_{\min}^2(X) \qquad \text{use second condition in equation 124}\\
&\leq c_1\alpha_0\sigma_{\min}^2(X)+c_1\alpha_0\sigma_{\min}^2(X)=a_1. \qquad (132)
\end{aligned}$$

For equation 130 to hold, we study its LHS

$$\text{LHS of equation 130}=\frac{2c_1^3\alpha_0^3p_2^2L(0)}{c_2^6\beta_0^6\kappa}+\frac{Pc_1^2\alpha_0}{c_2^3\beta_0}+\frac{8L(0)c_1}{(1-c_1)c_2\beta_0}+\frac{Pc_1}{c_2^2}+c_1\alpha_0\sigma_{\min}^2(X). \qquad (133)$$

Since $0<c_1<\frac{2}{3}$, we have $1-c_1\geq\frac{c_1}{2}$. Then, we upper bound the above equality by substituting c_1 with 1 in the numerator by one except for the last term and replace higher order terms of $c_2^k, k\geq 2$ with c_2,

$$\begin{aligned}
\text{LHS of equation 130} &= \frac{2c_1^3\alpha_0^3p_2^2L(0)}{c_2^6\beta_0^6\kappa}+\frac{Pc_1^2\alpha_0}{c_2^3\beta_0}+\frac{8L(0)c_1}{(1-c_1)c_2\beta_0}+\frac{Pc_1}{c_2^2}+c_1\alpha_0\sigma_{\min}^2(X)\\
&\leq \frac{2c_1^3\alpha_0^3p_2^2L(0)}{c_2^6\beta_0^6\kappa}+\frac{Pc_1^2\alpha_0}{c_2^3\beta_0}+\frac{16L(0)}{c_2\beta_0}+\frac{Pc_1}{c_2^2}+c_1\alpha_0\sigma_{\min}^2(X) \qquad \text{use } 1-c_1\geq\frac{c_1}{2}\\
&\leq \frac{2\alpha_0^3p_2^2L(0)}{c_2\beta_0^6\kappa}+\frac{P\alpha_0}{c_2\beta_0}+\frac{16L(0)}{c_2\beta_0}+\frac{P}{c_2}+c_1\alpha_0\sigma_{\min}^2(X)\\
&\qquad \text{use } c_1\leq 1 \text{ in the numerator and } c_2\geq c_2^k, k\geq 1 \text{ in the numerator}\\
&= \frac{1}{c_2}\left[\frac{2\alpha_0^3p_2^2L(0)}{\beta_0^6\kappa}+\frac{P\alpha_0}{\beta_0}+\frac{16L(0)}{\beta_0}+P\right]+c_1\alpha_0\sigma_{\min}^2
\end{aligned}$$

$$= \frac{1}{c_2}\Big[M + \frac{16L(0)}{\beta_0}\Big] + c_1\alpha_0\sigma_{\min}^2 \qquad \text{use first condition in equation 124}$$

$$\leq c_1\alpha_0\sigma_{\min}^2(X) + c_1\alpha_0\sigma_{\min}^2(X) = a_1. \tag{134}$$

Then, we prove

$$\rho(\eta_t^*, t) \leq \hat{\rho}(\eta_t^*, t) \leq f(\eta_t^*, t) \simeq 1 - \frac{(c_1\alpha_0)^2}{(c_2\beta_0)^2}\frac{1}{\kappa}. \tag{135}$$

This is true because

$$\hat{\rho}(\eta_\infty^*, \infty) = 1 - 2\alpha_\infty\sigma_{\min}^2(X)\frac{c_1\alpha_0}{c_2^2\beta_0^2\kappa\sigma_{\min}^2(X)} + \kappa\sigma_{\min}^2(X)\beta_\infty^2\Big(\frac{c_1\alpha_0}{c_2^2\beta_0^2\kappa\sigma_{\min}^2(X)}\Big)^2$$

$$= 1 - \frac{1}{\kappa}\Big(\frac{2c_1\alpha_0\alpha_\infty}{c_2^2\beta_0^2\sigma_{\min}^2(X)} - \frac{c_1^2\alpha_0^2\beta_\infty^2}{c_4^2\beta_0^4}\Big) \tag{136}$$

Since $\alpha_\infty \geq c_1\alpha_0, \beta_\infty \leq c_2\beta_0$, the above can be further upper bounded

$$\hat{\rho}(\eta_\infty^*, \infty) = 1 - \frac{1}{\kappa}\Big(\frac{2c_1\alpha_0\alpha_\infty}{c_2^2\beta_0^2\sigma_{\min}^2(X)} - \frac{c_1^2\alpha_0^2\beta_\infty^2}{c_4^2\beta_0^4}\Big) \tag{137}$$

$$\leq 1 - \frac{1}{\kappa}\Big(\frac{2c_1\alpha_0c_1\alpha_0}{c_2^2\beta_0^2\sigma_{\min}^2(X)} - \frac{c_1^2\alpha_0^2(c_2\beta_0)^2}{c_4^2\beta_0^4}\Big) \tag{138}$$

$$= 1 - \Big(\frac{c_1\alpha_0}{c_2\beta_0}\Big)^2\frac{1}{\kappa}. \tag{139}$$

Similarly, we can prove

$$\rho(\eta_\infty^*, \infty) \leq 1 - \Big(\frac{c_1\alpha_0}{c_2\beta_0}\Big)^2\frac{1}{\kappa}. \tag{140}$$

$\square$

F SIMULATIONS

In Section 4.1, we compare the step sizes proposed in Arora et al. [2018]; Du et al. [2018a], Theorem 3.2 and Algorithm 1. In Du et al. [2018a], they choose an adaptive step size

$$\eta_t = \frac{\sqrt{\epsilon/r}}{100(t+1)\|Y\|_F^{\frac{3}{2}}}, \tag{141}$$

where $0 < \epsilon < \|Y\|_F$ is the final precision we want to achieve, r is the rank of Y. When comparing, we set $\epsilon = \|Y\|_F$ to select the largest step size possible in their work.

In Arora et al. [2018], they choose constant step size which satisfies

$$\eta \leq \frac{p_1^3}{6144 \times 2^3 \times \|Y\|_F^4}, \tag{142}$$

When comparing, we select the largest step size possible, i.e. $\eta = \frac{p_1^3}{6144 \times 2^3 \times \|Y\|_F^4}$.

In Arora et al. [2018]; Du et al. [2018a], the authors make assumptions that there is sufficient margin and zero imbalance at initialization. What's more, they both choose the setting of matrix factorization and claim it's equivalent to linear networks. To make fair comparison, we generate X using identity matrix. For initialization of the network, we follow Proposition F.1 in Arora et al. [2018] to create a balanced initialization. The magnitude 0.04 of noise added to Y is a hyperparameter which ensures there is sufficient margin at initialization. The procedure to ensure there is zero imbalance at initialization is given below

Proposition F.1 (Spectral Balanced Initialization). *Given $d_0, d_1, \ldots, d_N \in \mathbb{N}$ such that $\min\{d_1, \ldots, d_{N-1}\} \geq \min\{d_0, d_N\}$ and a distribution $\mathcal{D}$ over $d_N \times d_0$ matrices, a balanced initialization of $W_j \in \mathbb{R}^{d_j \times d_{j-1}}$, $j=1, \ldots, N$, assigns these weights as follows:*

1. *Sample $A \in \mathbb{R}^{d_N \times d_0}$ according to $\mathcal{D}$.*

2. *Take singular value decomposition $A = U\Sigma V^\top$, where $U \in \mathbb{R}^{d_N \times \min\{d_0, d_N\}}$, $V \in \mathbb{R}^{d_0 \times \min\{d_0, d_N\}}$ have orthonormal columns, and $\Sigma \in \mathbb{R}^{\min\{d_0, d_N\} \times \min\{d_0, d_N\}}$ is diagonal and holds the singular values of A.*

3. *Set $W_N \simeq U\Sigma^{1/N}, W_{N-1} \simeq \Sigma^{1/N}, \ldots, W_2 \simeq \Sigma^{1/N}, W_1 \simeq \Sigma^{1/N}V^\top$, where the symbol "$\simeq$" stands for equality up to zero-valued padding.*