

CONVERGENCE OF GRADIENT DESCENT FOR LEARNING LINEAR NEURAL NETWORKS

GABIN MAXIME NGUEGNANG¹, HOLGER RAUHUT¹, AND ULRICH TERSTIEGE¹

ABSTRACT. We study the convergence properties of gradient descent for training deep linear neural networks, i.e., deep matrix factorizations, by extending a previous analysis for the related gradient flow. We show that under suitable conditions on the step sizes gradient descent converges to a critical point of the loss function, i.e., the square loss in this article. Furthermore, we demonstrate that for almost all initializations gradient descent converges to a global minimum in the case of two layers. In the case of three or more layers we show that gradient descent converges to a global minimum on the manifold matrices of some fixed rank, where the rank cannot be determined a priori.

1. INTRODUCTION

Deep learning is arguably the most widely used and successful machine learning method, which has led to spectacular breakthroughs in various domains such as image recognition, autonomous driving, machine translation, medical imaging and many more. Despite its widespread use, the understanding of the mathematical principles of deep learning is still in its infancy. A particular widely open question concerns the convergence properties of commonly used (stochastic) gradient descent (SGD) algorithms for learning a deep neural network from training data: Does (S)GD always converge to a critical point of the loss function? Does it converge to a global minimum? Does the network learned via (S)GD generalize well to unseen data?

In order to approach these questions we study gradient descent for learning a deep *linear* network, i.e., a network with activation function being the identity, or in other words, learning a deep matrix factorization. While linear neural networks are not expressive enough for most practical applications, the theoretical study of gradient descent for linear neural networks is highly non-trivial and therefore expected to be very valuable. The difficulty in deriving mathematical convergence guarantees results from the functional being non-convex in terms of the individual matrices in the factorization. We are convinced that the case of linear networks should be well-understood before passing to the more difficult (but more practically relevant) case of nonlinear networks. We expect that some principles (though not all) will transfer to the nonlinear case and the mathematical analysis of the linear case will provide valuable insights.

This article is a continuation of the work started in [5], where a theoretical analysis of the gradient flow related to learning a deep linear network via minimization of the square loss has been studied. Extending earlier contributions [2, 6, 3], it was shown in [5] that gradient flow always converges to a critical point of the square loss. Moreover, for almost all initializations it converges to a global minimizer in the case of two layers. It is conjectured that this result also holds for more than two layers, but currently it is only shown in [5] that for more layers, gradient flow converges to the global minimum of the loss function restricted to the matrices of some fixed rank k for almost all initializations, where unfortunately the result does not allow to determine k a priori. As another interesting discovery, [5] identifies the flow of the product matrix resulting from the gradient flow for the individual matrices in the factorization as a Riemannian gradient flow with respect to a non-trivial and explicitly given Riemannian metric on the manifold of matrices of fixed rank k . This result requires that at initialization the tuple of individual matrices is *balanced*, a term that the authors of [3] introduced. It is important to note that balancedness is preserved by the gradient flow, i.e., this property is related to the natural invariant set of the flow.

¹CHAIR FOR MATHEMATICS OF INFORMATION PROCESSING, RWTH AACHEN UNIVERSITY, PONTDRIESCH 10, 52062 AACHEN, GERMANY

E-mail addresses: nguegnang@mathc.rwth-aachen.de, rauhut@mathc.rwth-aachen.de, terstiege@mathc.rwth-aachen.de.

Date: August 4, 2021.

In this article, we extend the convergence analysis in [5] from gradient flow to gradient descent. Under certain conditions on the step sizes, we show that the gradient descent iterations converge to a critical point of the square loss function. Moreover, for almost all initializations convergence is towards a global minimum in the case of two layers, while for more than two layers we obtain the analogue of the main result in [5] that for almost all initializations the product matrix converges to a global minimum of the square loss restricted to the manifold of rank k matrices for some k .

The extension of the analysis from the gradient flow case to gradient descent turned out to be much more involved than one might initially expect. The reason is that the gradient descent iterations do no longer satisfy exactly the invariance property related to the balancedness. This property of the gradient flow, however, was heavily used in the convergence proof in [5]. In order to circumvent this problem, we develop an induction argument inspired by the article [8], which covers the significantly simpler special case of two layers. The induction proof tracks, in particular, how much the balancedness condition is perturbed during the iterations. In fact, such perturbations stay bounded under suitable assumptions on the step sizes.

Learning linear networks are currently studied also in the context of the so-called implicit bias of gradient descent and gradient flows. It was observed and studied empirically, for instance in [13, 18, 11], that in the context of overparameterization, where more parameters than training samples are used, learned neural networks generalize surprisingly well. This in contrast to the intuition gained from classical statistics, which would predict overfitting of the learned network, i.e., that generalization properties should be poor. In this overparameterized setting there will usually exist many neural networks that exactly interpolate the training examples [18], hence, leading to the global minimum zero for the empirical loss. In particular, a global minimizer is by far not unique. This means that the used learning algorithm, mostly variants of (stochastic) gradient descent, will induce an implicit bias on the learned neural network [11, 18]. As a result, one possible explanation of the phenomenon of the good generalization property of overparameterized learned neural networks is that the implicit bias of (stochastic) gradient descent is towards solutions of low complexity in a suitable sense, resulting in good generalization. While a theoretical analysis of this phenomenon seems difficult for nonlinear networks, first works for linear networks indicate that gradient descent leads to linear networks (factorized matrices) of low rank [2, 7, 10, 11, 16], although many open questions remain. Another important role seems to be played by the random initialization, see e.g. [9].

We expect that the convergence analysis of gradient descent performed in our paper, will be a useful tool also for the detailed analysis of the implicit bias of (stochastic) gradient descent in learning deep overparameterized neural networks.

Our article is structured as follows. Section 2 introduces deep linear networks and gradient descent, recalls the recent results from [5] on gradient flows, and presents our two main results on convergence to a critical point and convergence to a global minimizer for almost all initializations. Section 3 provides the proof of convergence to critical points (in the sense described above), while Section 4 is dedicated to the proof of convergence to global minimizers. Finally, Section 5 presents numerical experiments illustrating our results.

1.1. Notation. The standard ℓ_p -norm on \mathbb{R}^d will be denoted by $\|x\|_p = (\sum_{j=1}^d |x_j|^p)^{1/p}$ for $1 \leq p < \infty$. We write the spectral norm on $\mathbb{R}^{d \times m}$ as $\|A\| = \max_{\|x\|_2=1} \|Ax\|_2 = \sigma_{\max}(A)$, where $\sigma_{\max}(A)$ is the largest singular value of A . Moreover, we let $\sigma_{\min}(A) = \min_{\|x\|_2=1} \|Ax\|_2$ be the smallest singular value of A . The trace of a matrix A is denoted as $\text{tr}(A)$ and its Frobenius norm is defined as $\|A\|_F = \sqrt{\text{tr}(A^T A)} = \sqrt{\sum_{j,k} |A_{j,k}|^2}$. We will often combine matrices W_1, \dots, W_N into a tuple $\vec{W} = (W_1, \dots, W_N)$. We define the Frobenius inner product of two such a tuples \vec{W} and \vec{V} as $\langle \vec{W}, \vec{V} \rangle_F = \sum_{j=1}^N \text{tr}(W_j^T V_j)$ and the corresponding Frobenius norm as $\|\vec{W}\|_F = \sqrt{\langle \vec{W}, \vec{W} \rangle_F} = \left(\sum_{j=1}^N \|W_j\|_F^2 \right)^{1/2}$. The operator norm of a mapping \mathcal{A} acting between tuples of matrices will be denoted as $\|\mathcal{A}\|_{F \rightarrow F} = \max_{\|\vec{W}\|_F=1} \|\mathcal{A}(\vec{W})\|_F$.

2. LINEAR NEURAL NETWORKS AND GRADIENT DESCENT ANALYSIS

A neural network is a function $f : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_y}$ of the form

$$f(x) = f_{W_1, \dots, W_N, b_1, \dots, b_N}(x) = g_N \circ g_{N-1} \circ g_1(x),$$

where the so-called layers $g_j : \mathbb{R}^{d_{j-1}} \rightarrow \mathbb{R}^{d_j}$ are the composition of an affine function with a componentwise activation function, i.e.,

$$g_j(z) = \sigma(W_j z + b_j), \quad \text{for } W_j \in \mathbb{R}^{d_j \times d_{j-1}}, b_j \in \mathbb{R}^{d_j},$$

where $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ applied to a vector $w \in \mathbb{R}^{d_j}$ acts as $(\sigma(w))_k = \sigma(w_k)$, $k \in [d_j]$. Here, $d_0 = d_x$ and $d_N = d_y$, while $d_1, \dots, d_{N-1} \in \mathbb{N}$ are some numbers. Prominent examples for activation functions used in deep learning include $\sigma(t) = \text{ReLU}(t) = \max\{0, t\}$ and $\sigma(t) = \tanh(t)$, but we will simply choose the identity $\sigma(t) = t$ in this article.

Learning a neural network $f = f_{W_1, \dots, W_N, b_1, \dots, b_N}$ consists in adapting the parameters W_j, b_j based on labeled training data, i.e., pairs (x_i, y_i) of input data $x_1, \dots, x_m \in \mathbb{R}^{d_x}$ and output data $y_1, \dots, y_m \in \mathbb{R}^{d_y}$ in a way that $f_{W_1, \dots, W_N, b_1, \dots, b_N}(x_i) \approx y_i$ for $i \in [m]$. (Ideally, the learned neural network f should generalize well to unseen data, i.e., it should predict well the label y corresponding to new input data x . However, we will not discuss this point further in this article.)

The learning process is usually performed via optimization. Given a loss function $\ell : \mathbb{R}^{d_y} \times \mathbb{R}^{d_y} \rightarrow \mathbb{R}_+$ (usually satisfying $\ell(y, y) = 0$), one aims at minimizing the empirical risk function

$$\mathcal{L}(W_1, \dots, W_N, b_1, \dots, b_N) = \sum_{i=1}^m \ell(f_{W_1, \dots, W_N, b_1, \dots, b_N}(x_i), y_i)$$

with respect to the parameters $W_1, \dots, W_N, b_1, \dots, b_N$. Gradient descent and stochastic gradient descent algorithms are most commonly used for this task. A convergence analysis of these algorithms is challenging in general since due to the compositional nature of neural networks, the function \mathcal{L} is not convex in general.

Due to this difficulty, we reduce to the special case of *linear* neural networks in this article, i.e., we assume that $\sigma(t) = t$ is the identity and that $b_j = 0$ for all j . Consequently, a linear neural networks takes the form

$$f(x) = f_{W_1, \dots, W_N}(x) = W_N \cdots W_1 x = Wx, \quad \text{where } W = W_N \cdot W_{N-1} \cdots W_1.$$

While linear networks may not be expressive enough for many applications, convergence properties of gradient descent applied to learning linear neural networks are still non-trivial to understand. We will concentrate on the square-loss $\ell(z, w) = \frac{1}{2} \|z - w\|_2^2$ here, so that our learning problem consists in minimizing

$$L^N(W_1, \dots, W_N) = \frac{1}{2} \sum_{i=1}^m \|y_i - W_N \cdots W_1 x_i\|_2^2 = \frac{1}{2} \|Y - W_N \cdots W_1 X\|_F^2 = L^1(W_N \cdots W_1)$$

where the data matrix $X \in \mathbb{R}^{d_x \times m}$ contains the data points $x_i \in \mathbb{R}^{d_x}$, $i = 1, \dots, m$ as columns and likewise the matrix $Y \in \mathbb{R}^{d_y \times m}$ contains the label points $y_i \in \mathbb{R}^{d_y}$, $i = 1, \dots, m$. The function L^1 is given by

$$L^1(W) = \frac{1}{2} \|Y - WX\|_F^2.$$

Note that the rank of the matrix $W = W_N \cdots W_1$ is at most $r := \min_{i=0, \dots, N} d_i$, which is strictly smaller than $\min\{d_x, d_y\}$ if one of the ‘‘hidden’’ dimensions d_i is smaller than this number. Hence, we can also view the learning problem as the one of minimizing $L^1(W)$ under the constraint $\text{rank}(W) \leq r$. Instead of directly minimizing over W , we choose an *overparameterized* representation as $W = W_N \cdots W_1$ and consider gradient descent with respect to each factor W_i . While overparameterization seems to be a waste of resources at first sight, it also has certain advantages as it can even accelerate convergence [4] (at least for ℓ_p -losses with $p > 2$) or lead to solutions with better generalization properties [18]. Moreover, we expect that understanding theory for overparameterization in linear neural network will also give insights for overparameterization in nonlinear networks, which is widely used in practice. While speed of convergence or implicit bias are certainly of interest on their own, we will not delve into this, but rather concentrate on mere convergence here.

We consider gradient descent for the loss function L^N with step sizes η_k , i.e.,

$$W_j(k+1) = W_j(k) - \eta_k \nabla_{W_j} L^N(W_1(k), \dots, W_N(k)). \quad (2.1)$$

We further define the matrix W at each iteration k by

$$W(k) = W_N(k) \cdots W_1(k).$$

Before discussing gradient descent itself, let us recall previous results for the related gradient flow, which will guide the intuition for the analysis in this paper.

2.1. Gradient flow analysis. The gradient flow $\vec{W}(t) = (W_1(t), \dots, W_N(t))$, $t \in \mathbb{R}_+$ for the function L^N is defined via the differential equation

$$\frac{d}{dt} W_j(t) = -\nabla_{W_j} L^N(W_1(t), \dots, W_N(t)), \quad W_j(0) = W_{j,0} \quad j = 1, \dots, N, \quad (2.2)$$

for some initial matrices $W_{j,0} \in \mathbb{R}^{d_j \times d_{j-1}}$. This flow represents the continuous analogue of the gradient descent algorithm and has been analyzed in [2, 3, 6, 5].

An important invariance property of the gradient flow (2.2) consists in the fact that the differences

$$W_{j+1}^T(t)W_{j+1}(t) - W_j(t)W_j^T(t), \quad j = 1, \dots, N \quad (2.3)$$

are constant in time, see [2, 3, 6, 5]. This motivates to call a tuple $\vec{W} = (W_1, \dots, W_N)$ *balanced* if

$$W_{j+1}^T W_{j+1} = W_j W_j^T \quad \text{for all } j = 1, \dots, N. \quad (2.4)$$

If $\vec{W}(0) = (W_{1,0}, \dots, W_{N,0})$ is balanced then $\vec{W}(t)$ is balanced for all $t \in \mathbb{R}_+$ as a consequence of the invariance property. Note that by taking the trace on both sides of (2.4), we see that balancedness implies $\|W_j\|_F = \|W_1\|_F$ for all $j = 1, \dots, N$.

It is useful to introduce the “end-to-end” matrix $W(t) = W_N(t) \cdots W_1(t)$, which describes the action of the resulting network and is the object of main interest. It was shown in [3] that if the initial tuple $\vec{W}(0)$ (and hence $\vec{W}(t)$ for any $t \geq 0$) is balanced then the dynamics of $W(t)$ can be described without making use of the individual matrices $W_j(t)$. More precisely, it satisfies the differential equation

$$\frac{d}{dt} W(t) = -\mathcal{A}_{W(t)}(\nabla L^1(W(t))), \quad (2.5)$$

where $\mathcal{A}_W : \mathbb{R}^{d_x \times d_y} \rightarrow \mathbb{R}^{d_x \times d_y}$ is the linear map

$$\mathcal{A}_W(Z) = \sum_{j=1}^N (W W^T)^{\frac{N-j}{N}} \cdot Z \cdot (W^T W)^{\frac{j-1}{N}}.$$

One feature of the flow in (2.5), see [5, Theorem 4.5], is that the rank of $W(t)$ is constant in t , i.e., if $W(0) = W_N(0) \cdots W_1(0)$ has rank r then the $W(t)$ stays in the manifold of rank r matrices for all $t \geq 0$ (but note that the rank may drop in the limit). This property may fail for non-balanced initializations [5, Remark 4.2]. Another interesting observation (which, however, will not be important in our article) is that (2.5) can be interpreted as Riemannian gradient flow with respect to an appropriately defined Riemannian metric on the manifold of rank r matrices, see [5] for all the details.

The convergence properties of the gradient flow (2.2) (in both the unbalanced and balanced case) can be summarized in the following theorems. The first one from [5, Theorem 3.2] significantly generalizes the main result of [6].

Theorem 2.1. *Assume that XX^T has full rank. Then the flow $\vec{W}(t)$ defined by (2.2) is defined and bounded for all $t \geq 0$ and converges to a critical point of L^N as $t \rightarrow \infty$.*

This result is shown via Lojasiewicz’ theorem [1], which requires in turn to show boundedness of all components $W_i(t)$ of $\vec{W}(t)$. While boundedness is straightforward to show for $W(t)$, it is a nontrivial property of the $W_i(t)$. In fact, the proof exploits the invariance of the differences in (2.3).

While convergence to a critical point is nice to have, we would like to obtain more information about the type of critical point; whether it is a global or local minimum or merely a saddle point. Note that the

function L^N built from the square loss has the nice (but rare) property that a local minimum is automatically a global minimum [12, 17]. This means that we only need to single out saddle points. Also observe that we cannot expect to have convergence to a global minimizer for any initialization because the flow will not move when initializing in any critical point, so that we cannot expect convergence to a global minimizer if that critical point is not already a global minimizer. The following result valid for almost all initializations was derived in [5, Theorem 6.12]. In order to state it we need to introduce the matrix

$$Q = YX^T(XX^T)^{-1/2}, \quad (2.6)$$

assuming that XX^T has full rank.

Theorem 2.2. *Assume that XX^T has full rank, let $q = \text{rank}(Q)$, $r = \min_{j=0,\dots,N} d_j$ and $\bar{r} = \min\{q, r\}$.*

- (a) *For almost all initializations $\vec{W}(0)$, the flow (2.2) converges to a critical point $\vec{W}^* = (W_1^*, \dots, W_N^*)$ of L^N such that $W^* := W_N^* \cdots W_1^*$ is a global minimizer of L^1 on the manifold of matrices of fixed rank k for some $0 \leq k \leq \bar{r}$.*
- (b) *If $N = 2$ then for almost all initial values $W_1(0), \dots, W_N(0)$, the flow converges to a global minimizer of L^N on $\mathbb{R}^{d_0 \times d_1} \times \dots \times \mathbb{R}^{d_{N-1} \times d_N}$.*

We conjecture that the statement in part (b) also holds for $N \geq 3$, or in other words, that we can always choose the maximal possible rank $k = \bar{r}$ in (a), but unfortunately, the proof method employed in [5] is not able to deliver this extension without making significant adaptations. In fact, the proof relies on an abstract result, see [14] and [5, Theorem 6.3], which states that for almost all initializations so-called strict saddle points are avoided as limits. Unfortunately, if $N \geq 3$ then minimizers of L^1 restricted to the manifold of matrices of rank $k < \bar{r}$ may correspond to non-strict saddle points of L^N , see [12] and [5, Proposition 6.10], so that the abstract result does not apply to these points.

2.2. Gradient descent analysis. Our main goal is to extend Theorems 2.1 and 2.2 from gradient flow (2.2) to gradient descent (2.1). While the balancedness or more generally the invariance property, see (2.3), does not appear explicitly in the statements of these theorems for gradient flow (although the invariance property is used in the proof of Theorem 2.1), it turns out that it does play a role in the conditions for the step sizes ensuring convergence. Unfortunately, the invariance of the differences in (2.3) does not carry over to the iterations of gradient descent. As a consequence, we cannot simply adapt the proof strategy in [5] for the gradient flow case. However, we will prove that under suitable conditions on the step sizes the differences in (2.3) will stay bounded in norm, which then allows to show boundedness of the components $W_j(k)$ of $\vec{W}(k)$ and to apply Lojasiewicz' theorem to show convergence to a critical point.

In order to state our main results we introduce the following definition.

Definition 2.3. *We say that a tuple $\vec{W} = (W_1, \dots, W_N)$ has **balancedness constant** $\delta \geq 0$ if*

$$\|W_{j+1}^T W_{j+1} - W_j W_j^T\| \leq \delta \quad \text{for all } j = 1, \dots, N-1. \quad (2.7)$$

Obviously, (2.7) quantifies how much the tuple \vec{W} deviates from being balanced, measured in the spectral norm. Note that the authors of [2] introduced a very similar notion and said $\vec{W} = (W_1, \dots, W_N)$ to be δ -balanced if (2.7) holds with the spectral norm replaced by the Frobenius norm.

Theorem 2.4. *Let $X \in \mathbb{R}^{d_x \times m}$, $Y \in \mathbb{R}^{d_y \times m}$ be data matrices such that XX^T is of full rank. Suppose that the initialization $\vec{W}(0)$ of the gradient descent iterations (2.1) has balancedness constant $\alpha\delta$ for some $\delta > 0$ and $\alpha \in [0, 1)$. Assume that the stepsizes $\eta_k > 0$ satisfy $\sum_{k=0}^{\infty} \eta_k = \infty$ and*

$$\eta_k \leq \frac{2(1-\alpha)\delta}{4L^N(\vec{W}(0)) + (1-\alpha)\delta B_\delta} \quad \text{for all } k \in \mathbb{N}_0, \quad (2.8)$$

where

$$B_\delta := 2eNK_\delta^{N-1}\|X\|^2 + \sqrt{e}NK_\delta^{\frac{N}{2}-1}\|XY^T\|, \quad (2.9)$$

$$K_\delta := M^{\frac{2}{N}} + N^2\delta, \quad (2.10)$$

$$M := \frac{\sqrt{2L^N(\vec{W}(0))} + \|Y\|}{\sigma_{\min}(X)} = \frac{\sqrt{2}\|Y - W_N(0) \cdots W_1(0)X\|_F + \|Y\|}{\sigma_{\min}(X)}. \quad (2.11)$$

Then the sequence $\vec{W}(k)$ converges to a critical point of L^N .

Remark 2.5. (a) If $\vec{W}(0)$ is balanced, i.e., has balancedness constant 0, we can choose $\alpha = 0$ above. Then, for any $\delta > 0$, choosing the stepsizes η_k such that (3.14) below is satisfied ensures convergence to a critical point and that all the iterates $\vec{W}(k)$, $k \in \mathbb{N}$, have balancedness constant δ , see Proposition 3.3. This latter property will be a crucial ingredient for the proof of the theorem.

(b) Intuitively, the step sizes η_k should be chosen as large as possible in order to have fast convergence in practice, while it does not seem to be crucial to have the balancedness constant δ as small as possible during the iterations. This suggests to maximize the right hand side of (3.14) with respect to δ in order to make the condition on the stepsizes as weak as possible. While the analytical maximization seems difficult, this may be done numerically in practice. A reasonably good choice for δ seems to be

$$\delta = \frac{1}{N^3}M^{\frac{2}{N}}.$$

Then $K_\delta = (1 + \frac{1}{N})M^{\frac{2}{N}}$ so that $K_\delta^N \leq eM^2$. Since $2L^N(\vec{W}(0)) \leq \sigma_{\min}^2(X)M^2$, Condition (2.8) is then satisfied if

$$\eta_k \leq \frac{2}{2(1-\alpha)^{-1}N^3M^{2-\frac{2}{N}}\sigma_{\min}^2(X) + 2e^{2-\frac{1}{N}}NM^{2-\frac{2}{N}}\|X\|^2 + e^{1-\frac{1}{N}}NM^{1-\frac{2}{N}}\|XY^T\|}.$$

In particular, the required bound does not decrease exponentially with N .

(c) The stepsizes η_k in the theorem can be chosen a priori, for instance $\eta_k = \eta$ (constant step size), or $\eta_k = ck^{-\alpha}$ for some $\alpha \in [0, 1)$, or adaptively, i.e., depending on the current iterate $\vec{W}(k)$, as long as the step size condition (2.8) is satisfied. In practice, it seems that a large constant step size leads to best performance in terms of convergence speed.

Of course, more information on the type of critical point to which $\vec{W}(k)$ converges is desirable. Our next theorem states the analogue of Theorem 2.2 that essentially convergence is towards global minimizers for almost all initializations. Since Condition (3.14) on the stepsizes η_k ensuring mere convergence to a critical point depends on the initialization $\vec{W}(0)$, we can only expect to state a result for almost all initializations for sets of tuples \vec{W} of matrices for which the balancedness constant δ and M in (3.14) have a uniform upper bound. Consequently, we choose $\mathcal{B} \subset \mathbb{R}^{d_0 \times d_1} \times \cdots \times \mathbb{R}^{d_{N-1} \times d_N}$ to be bounded and let

$$\delta_{\mathcal{B}} = \sup_{\vec{W} \in \mathcal{B}} \max_{j=1, \dots, N-1} \|W_{j+1}^T W_{j+1} - W_j W_j^T\|, \quad (2.12)$$

$$L_{\mathcal{B}} = \sup_{\vec{W} \in \mathcal{B}} L^N(\vec{W}), \quad M_{\mathcal{B}} = \left(\sqrt{2L_{\mathcal{B}}} + \|Y\| \right) \sigma_{\min}^{-1}(X). \quad (2.13)$$

Note that $\delta_{\mathcal{B}}$ and $M_{\mathcal{B}}$ are finite (assuming XX^T has full rank) since L^N is continuous. Let us also recall the definition of the matrix $Q = YX^T(XX^T)^{-1/2}$ in (2.6).

Theorem 2.6. Let $\mathcal{B} \subset \mathbb{R}^{d_0 \times d_1} \times \cdots \times \mathbb{R}^{d_{N-1} \times d_N}$ be a bounded set with constants $\delta_{\mathcal{B}} \leq \alpha\delta$ as in (2.12) for some $\delta > 0$ and $\alpha \in [0, 1)$ and $L_{\mathcal{B}}$, $M_{\mathcal{B}}$ defined by (2.13). Let $q = \text{rank}(Q)$, $r = \min\{d_0, \dots, d_N\}$ and $\bar{r} = \min\{q, r\}$ and let $(\eta_k)_{k \in \mathbb{N}_0}$ be a sequence of positive stepsizes such that

$$\eta_k \leq \frac{2(1-\alpha)\delta}{4L_{\mathcal{B}} + (1-\alpha)\delta B_\delta} \quad \text{for all } k \in \mathbb{N}_0, \quad (2.14)$$

where

$$K_\delta := M_{\mathcal{B}}^{\frac{2}{N}} + N^2\delta, \quad B_\delta := 2eNK_\delta^{N-1}\|X\|^2 + \sqrt{e}NK_\delta^{\frac{N}{2}-1}\|XY^T\|.$$

Assume that additionally one of the following conditions is satisfied.

- (1) The sequence (η_k) is constant, i.e., $\eta_k = \eta$ for some $\eta > 0$ for all $k \in \mathbb{N}$.
- (2) It holds

$$\eta_k \geq C\frac{1}{k} \quad \text{for some } C > 0 \quad \text{and} \quad \lim_{k \rightarrow \infty} \eta_k = 0.$$

Then the following statements hold.

- (a) For almost all initializations $\vec{W}(0) = (W_1(0), \dots, W_N(0)) \in \mathcal{B}$, gradient descent (2.1) with step sizes η_k converges to a critical point \vec{W} of L^N such that $W = W_N \cdots W_1$ is a global minimum of L^1 on the manifold \mathcal{M}_k of matrices of rank $k = \text{rank}(W) \in \{0, 1, \dots, \bar{r}\}$ on $\mathbb{R}^{d_N \times d_0}$.
- (b) For $N = 2$, gradient descent (2.1) converges to a global minimum of L^N on $\mathbb{R}^{d_0 \times d_1} \times \mathbb{R}^{d_1 \times d_2}$ for almost all $\vec{W}(0) = (W_1(0), W_2(0)) \in \mathcal{B}$.

Similar as for Theorem 2.2, we conjecture that part (b) extends to $N \geq 3$ or equivalently that part (a) holds with $k = \bar{r}$. As for Theorem 2.2, the current proof method based on a strict saddle point analysis cannot be extended to show this conjecture.

It is currently not clear, whether the theorem holds under more general assumptions on the step sizes η_k , i.e., whether it is necessary that one of the two additional conditions on η_k holds. The current proof can only handle those two cases, for corresponding abstract results are available, see [14], [15]. It seems crucial for these general results that the stepsizes are chosen a priori and independently of the choice of $\vec{W}(0)$ (or the further iterates). In particular, adaptive step size choices are not covered by our theorem.

3. CONVERGENCE TO CRITICAL POINTS

We will prove Theorem 2.4 in this section. For $\vec{W} = (W_1, \dots, W_N)$ will always denote the corresponding product matrix by

$$W = W_N \cdots W_1,$$

and similarly, we denote by $W(k) = W_N(k) \cdots W_1(k)$ the sequence of product matrices associated to a sequence $\vec{W}(k) = (W_1(k), \dots, W_N(k))$, $k \in \mathbb{N}_0$. We recall from [3, 2, 6, 5] that

$$\nabla L^1(W) = WXX^T - YX^T, \tag{3.1}$$

$$\nabla_{W_j} L^N(W_1, \dots, W_N) = W_{j+1}^T \cdots W_N^T \nabla L^1(W) W_1^T \cdots W_{j-1}^T. \tag{3.2}$$

3.1. Auxiliary bounds. We start with a useful bound for $\|W\|$ in terms of $L^1(W)$.

Lemma 3.1. *Assume that XX^T has full rank. Then $W \in \mathbb{R}^{d_x \times d_y}$ satisfies*

$$\|W\| \leq (\|Y - WX\| + \|Y\|)\sigma_{\min}^{-1}(X) \leq \left(\sqrt{2L^1(W)} + \|Y\|\right)\sigma_{\min}^{-1}(X). \tag{3.3}$$

Consequently, if $L^N(\vec{W}(j+1)) \leq L^N(\vec{W}(j))$ for all $j \in \{0, \dots, k-1\}$, then

$$\|W(j)\| = \|W_N(j) \cdots W_1(j)\| \leq \left(\sqrt{2L^N(\vec{W}(0))} + \|Y\|\right)\sigma_{\min}^{-1}(X) \quad \text{for all } j \in \{0, \dots, k\}.$$

Furthermore,

$$\|\nabla L^1(W)\| \leq \|WX - Y\|\|X\| \leq \sqrt{2L^1(W)}\|X\|. \tag{3.4}$$

Proof. Arguing similarly as in the proof of [5, Theorem 3.2] gives

$$\begin{aligned} \|W\| &= \|WXX^T(XX^T)^{-1}\| \leq \|WX\|\|X^T(XX^T)^{-1}\| \leq (\|Y - WX\| + \|Y\|)\sigma_{\min}^{-1}(X) \\ &\leq (\|Y - WX\|_F + \|Y\|)\sigma_{\min}^{-1}(X) = \left(\sqrt{2L^1(W)} + \|Y\|\right)\sigma_{\min}^{-1}(X). \end{aligned}$$

The second claim follows then as an easy consequence recalling that $L^1(W(k)) = L^N(\vec{W}(k))$.

For the third claim we use the explicit formula (3.1) for the gradient of L^1 to conclude that

$$\|\nabla L^1(W)\| = \|WXX^T - YX^T\| \leq \|WX - Y\| \|X^T\| \leq \|WX - Y\|_F \|X\| = \sqrt{2L^1(W)} \|X\|.$$

This completes the proof. \square

A crucial ingredient in our proof is to show boundedness of all matrices $W_j(k)$, $k \in \mathbb{N}_0$. While boundedness for the product $W(k) = W_N(k) \cdots W_1(k)$ follows easily from the previous lemma, it does not immediately imply boundedness of all the factors $W_j(k)$. For instance, multiplying one factor $W_j(k)$ by a constant $\alpha > 0$ and another factor $W_\ell(k)$ by α^{-1} leaves the product $W(k)$ invariant but changes the norm of $W_j(k)$ and $W_\ell(k)$. In particular, letting $\alpha \rightarrow \infty$ shows that a bound for $W(k)$ alone does not imply boundedness for $W_j(k)$, $k \in \mathbb{N}_0$. This is where the balancedness comes in. In particular, if a tuple $\vec{W} = (W_1, \dots, W_N)$ has balancedness constant $\delta \geq 0$, then we can bound $\|W_j\|$, $j = 1, \dots, N$, by an expression (continuously) depending on $\|W\|$. This is the essence of the next statement.

Proposition 3.2. *Let $\vec{W} = (W_1, \dots, W_N) \in \mathbb{R}^{d_0 \times d_1} \times \dots \times \mathbb{R}^{d_{N-1} \times d_N}$ with balancedness constant $\delta \geq 0$ and let $W = W_N \cdots W_1$. Then*

$$\|W_j\|^2 \leq \|W\|^{\frac{2}{N}} + N^2 \delta \quad \text{for all } j = 1, \dots, N.$$

Proof. We will first prove that

$$\|W_1\|^{2N} \leq \|W\|^2 + Q_{N,\delta}(\|W_1\|^2), \quad (3.5)$$

where $Q_{N,\delta}$ is the polynomial of degree $N - 1$ defined as

$$Q_{N,\delta}(x) = x(x + \delta)(x + 2\delta) \cdots (x + (N - 1)\delta) - x^N.$$

In order to prove this claim, we let $D_j := W_{j-1}W_{j-1}^T - W_j^T W_j$ for $j = 2, \dots, N$ and note that $\|D_j\| \leq \delta$ by assumption. Moreover,

$$\|W_j\|^2 = \|W_j^T W_j\| = \|W_{j-1}W_{j-1}^T - D_j\| \leq \|W_{j-1}\|^2 + \delta, \quad \text{for all } j = 2, \dots, N, \quad (3.6)$$

and consequently

$$\|W_j\|^2 \leq \|W_1\|^2 + (j - 1)\delta \quad \text{for } j = 1, \dots, N. \quad (3.7)$$

We observe that by basic properties of the spectral norm

$$\begin{aligned} \|W_1\|^{2N} &= \|(W_1^T W_1)^N\| = \|W_1^T (W_1 W_1^T)^{N-1} W_1\| = \|W_1^T (W_2^T W_2 + D_2)^{N-1} W_1\| \\ &\leq \|W_1^T (W_2^T W_2)^{N-1} W_1\| + \sum_{k=0}^{N-2} \binom{N-1}{k} \|W_1\| \|W_2^T W_2\|^k \|D_2\|^{N-k-1} \|W_1\| \\ &\leq \|W_1^T (W_2^T W_2)^{N-1} W_1\| + \|W_1\|^2 \left(\sum_{k=0}^{N-1} \binom{N-1}{k} \|W_2\|^{2k} \delta^{N-k-1} - \|W_2\|^{2(N-1)} \right) \\ &= \|W_1^T W_2^T (W_2 W_2^T)^{N-2} W_2 W_1\| + \|W_1\|^2 \left((\|W_2\|^2 + \delta)^{N-1} - \|W_2\|^{2(N-1)} \right). \end{aligned} \quad (3.8)$$

In the first inequality, we expanded $(W_2^T W_2 + D_2)^{N-1}$ as a (matrix) polynomial in $W_2^T W_2$ and D_2 , observing that the highest degree term is $(W_2^T W_2)^{N-1}$. Applying the triangle inequality separates this term from the rest of the polynomial. Applying the submultiplicativity of the spectral norm to all the summands and collecting terms (which now consist of commuting scalars, i.e, the spectral norms $\|W_1\|$, $\|W_2^T W_2\|$ and $\|D_2\|$) gives the sum in (3.8), where the index $k = N - 1$ is left out as it was already taken care of in the first term in (3.8).

We continue in this way, replacing $(W_2 W_2^T)^{N-2}$ by $(W_3^T W_3 + D_3)^{N-2}$ and so on. Using also (3.7), we observe that similarly as above, for $j = 2, \dots, N - 1$,

$$\begin{aligned} &\|W_1^T \cdots W_j^T (W_j W_j^T)^{N-j} W_j \cdots W_1\| \\ &\leq \|W_1^T \cdots W_{j+1}^T (W_{j+1} W_{j+1}^T)^{N-j-1} W_{j+1} \cdots W_1\| + \|W_j\|^2 \cdots \|W_1\|^2 \left((\|W_{j+1}\|^2 + \delta)^{N-j} - \|W_{j+1}\|^{2(N-j)} \right) \\ &\leq \|W_1^T \cdots W_{j+1}^T (W_{j+1} W_{j+1}^T)^{N-j-1} W_{j+1} \cdots W_1\| \\ &\quad + \|W_1\|^2 (\|W_1\|^2 + \delta) \cdots (\|W_1\|^2 + (j - 1)\delta) \left((\|W_1\|^2 + j\delta)^{N-j} - (\|W_1\|^2 + (j - 1)\delta)^{N-j} \right). \end{aligned}$$

Hereby, we have also used that the function $x \mapsto (x + \delta)^{N-j} - x^{N-j}$ is monotonically increasing in $x \geq 0$. With this estimate we obtain, noting below that the sum in the second line is telescoping,

$$\begin{aligned} \|W_1\|^{2N} &\leq \|W_1^T \cdots W_N^T W_N \cdots W_1\| \\ &+ \sum_{j=1}^{N-1} \|W_1\|^2 (\|W_1\|^2 + \delta) \cdots (\|W_1\|^2 + (j-1)\delta) ((\|W_1\|^2 + j\delta)^{N-j} - (\|W_1\|^2 + (j-1)\delta)^{N-j}) \\ &= \|W_N \cdots W_1\|^2 + \|W_1\|^2 (\|W_1\|^2 + \delta) \cdots (\|W_1\|^2 + (N-1)\delta) - \|W_1\|^{2N} \\ &= \|W\|^2 + Q_{N,\delta}(\|W_1\|^2). \end{aligned}$$

This proves the claimed inequality (3.5).

The fact that for all $z, \alpha \in \mathbb{R}$ it holds $z(z + \alpha) \leq (z + \frac{\alpha}{2})^2$ implies that

$$x(x + \delta) \cdots (x + (N-1)\delta) = (x(x + (N-1)\delta)) \cdot (x + \delta)(x + (N-2)\delta) \cdots \leq \left(x + \frac{N-1}{2}\delta\right)^N. \quad (3.9)$$

Setting $x = \|W_1\|^2$, $a = \|W\|^2$ and $b = \frac{N-1}{2}\delta$ and combining inequality (3.5) and the definition of $Q_{N,\delta}$ with (3.9), leads to

$$x^N \leq a + (x + b)^N - x^N. \quad (3.10)$$

The mean-value theorem applied to the map $x \mapsto X^N$ gives

$$(x + b)^N = x^N + N\xi^{N-1}b \quad \text{for some } \xi \in [x, x + b].$$

Hence,

$$x^N \leq a + N\xi^{N-1}b \leq a + N(x + b)^{N-1}b.$$

We assume now that $a > 0$ and will comment on the case $a = 0$ below. Then the previous inequality implies

$$\frac{x^N}{a} \leq 1 + N \frac{(x + b)^{N-1}b}{a},$$

which is equivalent to

$$\left(\frac{x}{a^{\frac{1}{N}}}\right)^N \leq 1 + N \left(\frac{x}{a^{\frac{1}{N}}} + \frac{b}{a^{\frac{1}{N}}}\right)^{N-1} \frac{b}{a^{\frac{1}{N}}}. \quad (3.11)$$

Setting $z = a^{-\frac{1}{N}}x$ and $c = a^{-\frac{1}{N}}b$, we obtain

$$z^N \leq 1 + Nc(z + c)^{N-1}.$$

We claim that $z \leq 1 + 2Nc$. Assume on the contrary that $z > 1 + 2Nc$. Then (3.11) gives

$$z \leq \frac{1}{z^{N-1}} + Nc \left(1 + \frac{c}{z}\right)^{N-1} < 1 + Nc \left(1 + \frac{c}{1 + 2Nc}\right)^{N-1} \leq 1 + Nc \left(1 + \frac{1/2}{N}\right)^N \leq 1 + Nce^{\frac{1}{2}}.$$

The last inequality implies $z \leq 1 + 2Nc$, which is a contradiction. Thus, we showed the claim that $z \leq 1 + 2Nc$, that is, $xa^{-\frac{1}{N}} \leq 1 + 2Na^{-\frac{1}{N}}b$, which is equivalent to

$$x \leq a^{\frac{1}{N}} + 2Nb. \quad (3.12)$$

The last inequality also holds in the case $a = 0$, since for $a = 0$ inequality (3.10) remains true if we replace a by any positive number ε and then by our reasoning above $x \leq \varepsilon^{\frac{1}{N}} + 2Nb$. Since this is true for any $\varepsilon > 0$, it follows that for $a = 0$ we have $x \leq 2Nb = a^{\frac{1}{N}} + 2Nb$, thus (3.12) also holds for $a = 0$.

Using the definitions of a, b and x , we obtain from (3.12) that

$$\|W_1\|^2 \leq \|W\|^{\frac{2}{N}} + N(N-1)\delta.$$

For any $j = 1, \dots, N$, (3.7) implies then that

$$\|W_j\|^2 \leq \|W_1\|^2 + (j-1)\delta \leq \|W\|^{\frac{2}{N}} + N(N-1)\delta + (j-1)\delta \leq \|W\|^{\frac{2}{N}} + N^2\delta.$$

This completes the proof. \square

3.2. Preservation of approximate balancedness. The key ingredient to the proof of Theorem 2.4 is the following proposition. It is a highly non-trivial extension of [8, Lemma 3.1] from $N = 2$ layers to an arbitrary number of layers.

Proposition 3.3. *Assume that XX^T has full rank and that $\vec{W}(0) = (W_1(0), \dots, W_N(0))$ has balancedness constant $\alpha\delta$ for some $\delta > 0$ and $\alpha \in [0, 1)$. Assume that the positive stepsizes η_k satisfy (3.14). Then the gradient descent iterates $\vec{W}(k) = (W_1(k), \dots, W_N(k))$ defined by (2.1) satisfy, for all $k \in \mathbb{N}_0$:*

(1) $\vec{W}(k)$ has balancedness constant δ , i.e.,

$$\|W_{j+1}^T(k)W_{j+1}(k) - W_j(k)W_j^T(k)\| \leq \delta \quad \text{for all } j = 1, \dots, N-1; \quad (3.13)$$

(2) $L^N(\vec{W}(k)) \leq L^N(\vec{W}(0))$;

(3) $\|W_j(k)\|^2 \leq K_\delta = M^{\frac{2}{N}} + N^2\delta$ for $j = 1, \dots, N$;

(4) $L^N(\vec{W}(k)) - L^N(\vec{W}(k+1)) \geq \sigma\eta_k \left\| \nabla L^N(\vec{W}(k)) \right\|_F^2$.

Proof. We will show statements (1), (2) and (3) by induction under the condition that

$$\eta_k \leq \min \left\{ \frac{2(1-\sigma)}{B_\delta}, \frac{\sigma(1-\alpha)\delta}{2L^N(\vec{W}(0))} \right\} \quad \text{for all } k \in \mathbb{N}, \quad (3.14)$$

holds for some $\sigma \in (0, 1)$. The choice

$$\sigma = \frac{4L^N(\vec{W}(0))}{4L^N(\vec{W}(0)) + (1-\alpha)\delta B_\delta}$$

reduces (3.14) to (2.8). In the induction step for (2), we will show that if (3) holds for k , then (4) holds for k as well. Below we will always denote $W(k) = W_N(k) \cdots W_1(k)$.

Since $\vec{W}(0)$ has balancedness constant $\alpha\delta < \delta$ by assumption, (3.13) is clearly satisfied for $k = 0$. Statement (2) is trivial for $k = 0$. The bound in (3) follows from a direct combination of Proposition 3.2 with Lemma 3.1, i.e., for $j = 1, \dots, N$,

$$\|W_j(0)\|^2 \leq \|W(0)\|^{\frac{2}{N}} + N^2\delta \leq \left(\frac{\sqrt{2L^N(\vec{W}(0))} + \|Y\|}{\sigma_{\min}^{-1}(X)} \right)^{\frac{N}{2}} + N^2\delta = M^{\frac{2}{N}} + N^2\delta,$$

using also the definition of M in (3.2).

For the induction step, we assume that (1), (2) and (3) hold for $0, 1, \dots, k$ and prove that these three properties hold for $k+1$ as well.

Step 1: We first prove statement (2) for $k+1$. To do so, we will show that if statement (3) holds for k then statement (4) holds for k as well. This also proves (4) once the induction for (1), (2) and (3) is completed.

We consider the Taylor expansion

$$\begin{aligned} L^N(\vec{W}(k+1)) &= L^N(\vec{W}(k)) + \left\langle \nabla_{\vec{W}} L^N(\vec{W}(k)), \vec{W}(k+1) - \vec{W}(k) \right\rangle \\ &\quad + \frac{1}{2} \left\langle \left(\vec{W}(k+1) - \vec{W}(k) \right)^T \nabla^2 L^N(\vec{A}_\xi), \vec{W}(k+1) - \vec{W}(k) \right\rangle, \end{aligned}$$

where $\nabla L^N(\vec{W}(k)) = \begin{pmatrix} \nabla_{W_1} L^N(\vec{W}(k)) \\ \vdots \\ \nabla_{W_N} L^N(\vec{W}(k)) \end{pmatrix}$ and $\vec{A}_\xi = (A_\xi^1, \dots, A_\xi^N)$ with

$$A_\xi^i = W_i(k) + \xi(W_i(k+1) - W_i(k)) \quad \text{for some } \xi \in [0, 1], \quad i = 1, \dots, N.$$

By definition of $\vec{W}(k+1)$, this Taylor expansion can be written as

$$\begin{aligned} L^N(\vec{W}(k+1)) &= L^N(\vec{W}(k)) - \eta_k \langle \nabla L^N(\vec{W}(k)), \nabla L^N(\vec{W}(k)) \rangle_F \\ &\quad + \frac{1}{2} \eta_k^2 \langle \nabla L^N(\vec{W}(k)), \nabla^2 L^N(\vec{A}_\xi) \nabla L^N(\vec{W}(k)) \rangle_F. \end{aligned}$$

By the Cauchy-Schwarz inequality, we obtain

$$\begin{aligned} L^N(\vec{W}(k)) - L^N(\vec{W}(k+1)) &\geq \eta_k \left\| \nabla L^N(\vec{W}(k)) \right\|_F^2 - \frac{1}{2} \eta_k^2 \left\| \nabla L^N(\vec{W}(k)) \right\|_F^2 \left\| \nabla^2 L^N(\vec{A}_\xi) \right\|_{F \rightarrow F} \\ &\geq \left(1 - \frac{1}{2} \eta_k \left\| \nabla^2 L^N(\vec{A}_\xi) \right\|_{F \rightarrow F} \right) \eta_k \left\| \nabla L^N(\vec{W}(k)) \right\|_F^2. \end{aligned} \quad (3.15)$$

The crucial point now is to show that $\left\| \nabla^2 L^N(\vec{A}_\xi) \right\|_{F \rightarrow F}$ is bounded by the constant B_δ defined in (2.9). By setting $\vec{\Delta} = (\Delta_1, \dots, \Delta_N)$, $\Delta_j \in \mathbb{R}^{d_j \times d_{j-1}}$, and writing $\nabla^2 L^N(\vec{W})(\vec{\Delta}, \vec{\Delta})$ for $\langle \vec{\Delta}, \nabla^2 L^N(\vec{W}) \vec{\Delta} \rangle$, the quadratic form $\nabla^2 L^N(\vec{W})(\vec{\Delta}, \vec{\Delta})$ defined by the Hessian can be written as

$$\begin{aligned} \nabla^2 L^N(\vec{W})(\vec{\Delta}, \vec{\Delta}) &= \sum_{j=1}^N \sum_{i=1}^N \left\langle \Delta_j, \frac{\partial^2 L^N(\vec{W})}{\partial W_i \partial W_j} \Delta_i \right\rangle \\ &= \sum_{i=1}^N \left\langle \Delta_i, \frac{\partial^2 L^N(\vec{W})}{\partial W_i^2} \Delta_i \right\rangle + \sum_{j=1}^N \sum_{\substack{i=1 \\ i \neq j}}^N \left\langle \Delta_j, \frac{\partial^2 L^N(\vec{W})}{\partial W_i \partial W_j} \Delta_i \right\rangle. \end{aligned}$$

In order to compute mixed second derivatives we introduce the notation

$$\begin{aligned} Q_i(\vec{W}, \Delta_i) &= W_N \cdots W_{i+1} \Delta_i W_{i-1} \cdots W_1 X \\ P_{i,j}(\vec{W}, \Delta_i, \Delta_j) &= \begin{cases} W_N \cdots W_{j+1} \Delta_j W_{j-1} \cdots W_{i+1} \Delta_i W_{i-1} \cdots W_1 & \text{if } j > i, \\ W_N \cdots W_{i+1} \Delta_i W_{i-1} \cdots W_{j+1} \Delta_j W_{j-1} \cdots W_1 & \text{if } j < i, \end{cases} \end{aligned}$$

with the understanding that $W_{i-1} \cdot W_1 = \text{Id}$ for $i = 1$ and $W_N \cdots W_{i+1} = \text{Id}$ for $i = N$. Using the first partial derivatives of L^N , cf. (3.2), we obtain, for $i = 1, \dots, N$,

$$\left\langle \Delta_i, \frac{\partial^2 L^N(\vec{W})}{\partial W_i^2} \Delta_i \right\rangle = \langle Q_i(\vec{W}, \Delta_i), Q_i(\vec{W}, \Delta_i) \rangle = \|Q_i(\vec{W}, \Delta_i)\|_F^2.$$

The mixed second order derivatives are given, for $i \neq j$, by

$$\left\langle \Delta_i, \frac{\partial^2 L^N(\vec{W})}{\partial W_i \partial W_j} \Delta_j \right\rangle = \langle Q_i(\vec{W}, \Delta_i), Q_j(\vec{W}, \Delta_j) \rangle + \langle L^N(\vec{W}), P_{i,j}(\vec{W}, \Delta_i, \Delta_j) \rangle.$$

This implies that

$$\begin{aligned} \nabla^2 L^N(\vec{A}_\xi)(\vec{\Delta}, \vec{\Delta}) &= \sum_{i=1}^N \|Q_i(\vec{A}_\xi, \Delta_i)\|_F^2 + \sum_{\substack{i,j=1 \\ i \neq j}}^N \langle Q_i(\vec{A}_\xi, \Delta_i), Q_j(\vec{A}_\xi, \Delta_j) \rangle \\ &\quad + \sum_{\substack{i,j=1 \\ i \neq j}}^N \langle A_\xi X X^T - Y X^T, P_{i,j}(\vec{A}_\xi, \Delta_i, \Delta_j) \rangle, \end{aligned}$$

where $A_\xi = A_\xi^N \cdots A_\xi^1$ as usual. The Cauchy-Schwarz inequality for the trace inner product together with $\|AB\|_F \leq \|A\| \|B\|_F$ for any matrices A, B of matching dimensions gives, for $i > j$,

$$\begin{aligned} & \left| \left\langle A_\xi X X^T - Y X^T, P_{i,j}(\vec{A}_\xi, \Delta_i, \Delta_j) \right\rangle \right| \\ &= \left| \text{tr}((A_\xi X X^T - Y X^T)^T A_\xi^N \cdots A_\xi^{i+1} \Delta_i A_\xi^{i-1} \cdots A_\xi^{j+1} \Delta_j A_\xi^{j-1} \cdots A_\xi^1) \right| \\ &\leq \|(A_\xi X X^T - Y X^T)^T A_\xi^N \cdots A_\xi^{i+1} \Delta_i\|_F \|A_\xi^{i-1} \cdots A_\xi^{j+1} \Delta_j A_\xi^{j-1} \cdots A_\xi^1\|_F \\ &\leq \|A_\xi X X^T - Y X^T\| \|A_\xi^N\| \cdots \|A_\xi^{i+1}\| \|\Delta_i\|_F \|A_\xi^{i-1}\| \cdots \|A_\xi^{j+1}\| \|\Delta_j\|_F \|A_\xi^{j-1}\| \cdots \|A_\xi^1\|, \end{aligned}$$

and similarly, for $i < j$. Another application of the Cauchy-Schwarz inequality gives $|\langle Q_i(\vec{A}_\xi, \Delta_i), Q_j(\vec{A}_\xi, \Delta_j) \rangle| \leq \|Q_i(\vec{A}_\xi, \Delta_i)\|_F \|Q_j(\vec{A}_\xi, \Delta_j)\|_F$. Consequently

$$\begin{aligned} \left| \nabla^2 L^N(\vec{A}_\xi)(\vec{\Delta}, \vec{\Delta}) \right| &\leq \sum_{i,j=1}^N \|Q_i(\vec{A}_\xi, \Delta_i)\|_F \|Q_j(\vec{A}_\xi, \Delta_j)\|_F \\ &\quad + \sum_{\substack{i,j=1 \\ i \neq j}}^N \|A_\xi X X^T - Y X^T\| \|\Delta_i\|_F \|\Delta_j\|_F \prod_{\substack{k=1 \\ k \neq i,j}}^N \|A_\xi^k\| \\ &\leq \|X\|^2 \sum_{i,j=1}^N \|\Delta_i\|_F \|\Delta_j\|_F \left(\prod_{\substack{k=1 \\ k \neq i}}^N \|A_\xi^k\| \right) \left(\prod_{\substack{k=1 \\ k \neq j}}^N \|A_\xi^k\| \right) \\ &\quad + \sum_{\substack{i,j=1 \\ i \neq j}}^N \|\Delta_i\|_F \|\Delta_j\|_F \left(\|X\|^2 \prod_{k=1}^N \|A_\xi^k\| + \|Y X^T\| \right) \prod_{\substack{k=1 \\ k \neq i,j}}^N \|A_\xi^k\|. \end{aligned} \quad (3.16)$$

Using the recursive definition of $W_i(k+1)$ and that $\xi \in [0, 1]$ we further obtain, for $i = 1, \dots, N$,

$$\begin{aligned} \|A_\xi^i\| &= \|W_i(k) + \xi(W_i(k+1) - W_i(k))\| \leq \|W_i(k)\| + \|W_i(k+1) - W_i(k)\| \\ &= \|W_i(k)\| + \left\| \eta_k \nabla_{W_i} L^N(\vec{W}(k)) \right\| \\ &= \|W_i(k)\| + \eta_k \|W_{i+1}^T(k) \cdots W_N^T(k) \nabla L^1(W(k)) W_1(k)^T \cdots W_{i-1}(k)\|. \end{aligned}$$

It follows from (3.4) and the induction hypothesis (2) for k that

$$\|\nabla L^1(W(k))\| \leq \sqrt{2L^N(\vec{W}(k))} \|X\| \leq \sqrt{2L^N(\vec{W}(0))} \|X\|. \quad (3.17)$$

Using the induction hypothesis (3) for k this gives

$$\|A_\xi^i\| \leq \|W_i(k)\| + \eta_k \sqrt{2L^N(\vec{W}(0))} \|X\| \left(\prod_{\substack{j=1 \\ j \neq i}}^N \|W_j(k)\| \right) \leq K_\delta^{1/2} + \eta_k \sqrt{2L^N(\vec{W}(0))} \|X\| K_\delta^{\frac{N-1}{2}}.$$

By the assumption (3.14) on the stepsize η_k and the definitions of K_δ and B_δ we have

$$\begin{aligned} \eta_k \sqrt{2L^N(\vec{W}(0))} \|X\| K_\delta^{\frac{N-1}{2}} &\leq \frac{2(1-\sigma)}{B_\delta} \sqrt{2L^N(\vec{W}(0))} \|X\| K_\delta^{\frac{N-1}{2}} \leq \frac{2\sqrt{2L^N(\vec{W}(0))} \|X\| K_\delta^{\frac{N-1}{2}}}{2eN K_\delta^{N-1} \|X\|^2} \\ &\leq \frac{M \sigma_{\min}(X) \|X\| K_\delta^{\frac{N-1}{2}}}{eN K_\delta^{N-1} \|X\|^2} \leq \frac{(M^{\frac{2}{N}} + N^2 \delta)^{\frac{N}{2}} K_\delta^{\frac{N-1}{2}}}{eN K_\delta^{N-1}} \leq \frac{1}{2N} K_\delta^{\frac{1}{2}}. \end{aligned}$$

hence,

$$\|A_\xi^i\| \leq \left(1 + \frac{1}{2N} \right) K_\delta^{1/2}.$$

Substituting this bound into (3.16), we obtain

$$\begin{aligned}
& \left| \nabla^2 L^N \left(\vec{A}_\xi \right) \left(\vec{\Delta}, \vec{\Delta} \right) \right| \\
& \leq \left(1 + \frac{1}{2N} \right)^{2N-2} K_\delta^{N-1} \|X\|^2 \sum_{i,j=1}^N \|\Delta_j\|_F \|\Delta_W\|_F \\
& + \left(\left(1 + \frac{1}{2N} \right)^{2N-2} K_\delta^{N-1} \|X\|^2 + \left(1 + \frac{1}{2N} \right)^{N-2} K_\delta^{N/2-1} \|XY^T\| \right) \sum_{j=1}^N \sum_{\substack{i=1 \\ i \neq j}}^N \|\Delta_j\|_F \|\Delta_{W_i}\|_F \\
& \leq e K_\delta^{N-1} \|X\|^2 \left(\sum_{j=1}^N \|\Delta_j\| \right)^2 + \left(e K_\delta^{N-1} \|X\|^2 + e^{1/2} K_\delta^{N/2-1} \|XY^T\| \right) \left(\sum_{j=1}^N \|\Delta_j\| \right)^2 \\
& \leq \left[2eN K_\delta^{N-1} \|X\|^2 + \sqrt{e} N K_\delta^{N/2-1} \|XY^T\| \right] \|\vec{\Delta}\|_F^2,
\end{aligned}$$

where we have used that $(1 + 1/(2N))^{2N} \leq e$ and that $\sum_{j=1}^N \|\Delta_{W_j}\|_F \leq \sqrt{N} \|\vec{\Delta}\|_F$. Hence, we derived that

$$\left\| \nabla^2 L^N \left(\vec{A}_\xi \right) \right\|_{F \rightarrow F} \leq 2eN K_\delta^{N-1} \|X\|^2 + \sqrt{e} N K_\delta^{\frac{N}{2}-1} \|XY^T\| = B_\delta.$$

Substituting this estimate into (3.15) and using that the step sizes satisfy gives (3.14)

$$L^N \left(\vec{W}(k) \right) - L^N \left(\vec{W}(k+1) \right) \geq \left(1 - \frac{1}{2} \eta_k B_\delta \right) \eta_k \left\| \nabla L^N \left(\vec{W}(k) \right) \right\|_F^2 \geq \sigma \eta_k \left\| \nabla L^N \left(\vec{W}(k) \right) \right\|_F^2 \geq 0. \quad (3.18)$$

This shows the statement (4) for k . It follows by the induction hypothesis (2) for k that

$$L^N \left(\vec{W}(0) \right) \geq L^N \left(\vec{W}(k) \right) \geq L^N \left(\vec{W}(k+1) \right).$$

This shows statement (2) for $k+1$.

Step 2: Let us now show that statement (1) holds at iteration $k+1$. For $j = 1, \dots, N-1$ we obtain

$$\begin{aligned}
& \left\| W_{j+1}^T(k+1) W_{j+1}(k+1) - W_j(k+1) W_j^T(k+1) \right\| \\
& = \left\| \left(W_{j+1}^T(k) - \eta_k \nabla_{W_{j+1}}^T L^N(\vec{W}(k)) \right) \left(W_{j+1}(k) - \eta_k \nabla_{W_{j+1}} L^N(\vec{W}(k)) \right) \right. \\
& \quad \left. - \left(W_j(k) - \eta_k \nabla_{W_j} L^N(\vec{W}(k)) \right) \left(W_j^T(k) - \eta_k \nabla_{W_j}^T L^N(\vec{W}(k)) \right) \right\| \\
& = \left\| W_{j+1}^T(k) W_{j+1}(k) - W_j(k) W_j^T(k) \right. \\
& \quad + \eta_k \left(-W_{j+1}^T(k) W_{j+2}^T(k) \dots W_N^T(k) \nabla L^1(W(k)) W_1^T(k) \dots W_j^T(k) \right. \\
& \quad \quad \left. - W_j(k) \dots W_1(k) \nabla^T L^1(W(k)) W_N(k) \dots W_{j+2}(k) W_{j+1}(k) \right) \\
& \quad + W_j(k) W_{j-1}(k) \dots W_1(k) \nabla^T L^1(W(k)) W_N(k) \dots W_{j+2}(k) W_{j+1}(k) \\
& \quad \left. + W_{j+1}^T(k) W_{j+2}^T(k) \dots W_N^T(k) \nabla L^1(W(k)) W_1^T(k) \dots W_{j-1}^T(k) W_j^T(k) \right) \\
& \quad + \eta_k^2 \left(\nabla_{W_{j+1}}^T L^N(\vec{W}(k)) \nabla_{W_{j+1}} L^N(\vec{W}(k)) - \nabla_{W_j} L^N(\vec{W}(k)) \nabla_{W_j}^T L^N(\vec{W}(k)) \right) \left\| \right. \\
& \leq \left\| W_{j+1}^T(k) W_{j+1}(k) - W_j(k) W_j^T(k) \right\| + \eta_k^2 \left(\left\| \nabla_{W_{j+1}} L^N(\vec{W}(k)) \right\|^2 + \left\| \nabla_{W_j} L^N(\vec{W}(k)) \right\|^2 \right).
\end{aligned}$$

Applying this inequality repeatedly, we obtain

$$\begin{aligned}
& \|W_{j+1}^T(k+1)W_{j+1}(k+1) - W_j(k+1)W_j^T(k+1)\| \\
& \leq \|W_{j+1}^T(0)W_{j+1}(0) - W_j(0)W_j^T(0)\| + \sum_{\ell=0}^k \eta_\ell^2 \left(\|\nabla_{W_{j+1}} L^N(\vec{W}(\ell))\|^2 + \|\nabla_{W_j} L^N(\vec{W}(\ell))\|^2 \right) \\
& \leq \alpha\delta + 2 \left(\max_{\ell=0, \dots, k} \eta_\ell \right) \sum_{\ell=0}^k \eta_\ell \|\nabla L^N(\vec{W}(\ell))\|_{F \rightarrow F}^2, \tag{3.19}
\end{aligned}$$

where we have used that $\vec{W}(0)$ has balancedness constant $\alpha\delta$ by assumption and that

$$\|\nabla L^N(\vec{W}(k))\|_{F \rightarrow F}^2 \geq \max_{\ell=1, \dots, N} \|\nabla_{W_\ell} L^N(\vec{W}(k))\|^2 \geq \frac{1}{2} \left(\|\nabla_{W_j} L^N(\vec{W}(k))\|^2 + \|\nabla_{W_{j+1}} L^N(\vec{W}(k))\|^2 \right).$$

The inequality (3.18) from the previous step gives

$$L^N(\vec{W}(0)) - L^N(\vec{W}(k+1)) = \sum_{j=0}^k \left(L^N(\vec{W}(j)) - L^N(\vec{W}(j+1)) \right) \geq \sigma \sum_{j=0}^k \eta_j \|\nabla L^N(\vec{W}(j))\|_{F \rightarrow F}^2. \tag{3.20}$$

Combining inequalities (3.19) and (3.20) yields

$$\begin{aligned}
& \|W_{j+1}^T(k+1)W_{j+1}(k+1) - W_j(k+1)W_j^T(k+1)\| \leq \alpha\delta + \frac{2}{\sigma} \left(\max_{\ell=0, \dots, k} \eta_\ell \right) \left(L^N(\vec{W}(0)) - L^N(\vec{W}(k+1)) \right) \\
& \leq \alpha\delta + \frac{2}{\sigma} \left(\max_{\ell=0, \dots, k} \eta_\ell \right) L^N(\vec{W}(0)) \leq \alpha\delta + (1-\alpha)\delta = \delta,
\end{aligned}$$

where we have used Condition (3.14) on the stepsizes. This proves statement (1) for $k+1$.

Step 3: For the proof of statement (3) for $k+1$ we use that we have already shown that (1) and (2) hold for $k+1$. It follows from Proposition 3.2 and Lemma 3.1 that

$$\|W_j(k+1)\|^2 \leq \|W(k+1)\|^{\frac{2}{N}} + N^2\delta \leq \left(\frac{\sqrt{2L^N(\vec{W}(0))} + \|Y\|}{\sigma_{\min}(X)} \right)^{\frac{2}{N}} + N^2\delta = K_\delta.$$

This shows (3) for $k+1$ and completes the proof of the proposition. \square

3.3. Convergence of gradient descent to a critical point. We will use a result from [1] to prove Theorem 2.4, which is based on the following definition.

Definition 3.4 (Strong descent conditions [1]). *We say that a sequence $x_k \in \mathbb{R}^n$ satisfies the strong descent conditions (for a differentiable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$) if*

$$f(x_k) - f(x_{k+1}) \geq \sigma \|\nabla f(x_k)\| \|x_{k+1} - x_k\| \tag{3.21}$$

$$\text{and } f(x_{k+1}) = f(x_k) \implies x_{k+1} = x_k \tag{3.22}$$

hold for some $\sigma > 0$ and for all k larger than some K .

The next theorem is essentially an extension of Lojasiewicz' theorem to discrete variants of gradient flows.

Theorem 3.5. [1, Theorem 3.2] *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be an analytic cost function. Let the sequence $\{x_k\}_{k=1,2,\dots}$ satisfy the strong descent conditions (Definition 3.4). Then, either $\lim_{k \rightarrow \infty} \|x_k\| = +\infty$, or there exists a single point $x^* \in \mathbb{R}$ such that*

$$\lim_{k \rightarrow \infty} x_k = x^*$$

Now we are ready to prove Theorem 2.4.

Proof. By point (4) of Proposition 3.3 and since $\|\vec{W}(k+1) - \vec{W}(k)\|_F = \|\eta_k \nabla L^N(\vec{W}(k))\|_F$ for all $k \in \mathbb{N}_0$, we have

$$L^N(\vec{W}(k)) - L^N(\vec{W}(k+1)) \geq \sigma \|\nabla L^N(\vec{W}(k))\|_F \|\vec{W}(k+1) - \vec{W}(k)\|_F, \quad (3.23)$$

which means that the first part (3.21) of the strong descent condition holds. This implies then that also the second part (3.22) of the strong descent condition holds, since if $L^N(\vec{W}(k+1)) = L^N(\vec{W}(k))$, it follows that

$$\|\nabla L^N(\vec{W}(k))\|_F \|\vec{W}(k+1) - \vec{W}(k)\|_F = 0,$$

hence $\vec{W}(k+1) = \vec{W}(k)$ or $\nabla L^N(\vec{W}(k)) = 0$, but the latter again implies $\vec{W}(k+1) = \vec{W}(k)$. Thus indeed $\vec{W}(k+1) = \vec{W}(k)$ if $L^N(\vec{W}(k+1)) = L^N(\vec{W}(k))$.

Since by Proposition 3.3, the sequence $(\vec{W}(k))_{k \in \mathbb{N}_0}$ is bounded and L^N is analytic, it follows from Theorem 3.5 that there exists \vec{W}^* such that

$$\lim_{k \rightarrow \infty} \vec{W}(k) = \vec{W}^*.$$

It remains to show that \vec{W}^* is a critical point of L^N . Since $\nabla L^N(\vec{W})$ is continuous in \vec{W} , it follows that $\nabla L^N(\vec{W}^*) = \lim_{k \rightarrow \infty} \nabla L^N(\vec{W}(k))$ and that $\|\nabla L^N(\vec{W}^*)\|_F = \lim_{k \rightarrow \infty} \|\nabla L^N(\vec{W}(k))\|_F := c$. In order to show that \vec{W}^* is a critical point, it suffices to show that $c = 0$. A repeated application of point (4) of Proposition 3.3 gives

$$L^N(\vec{W}(0)) - L^N(\vec{W}(k+1)) \geq \sigma \sum_{j=0}^k \eta_j \|\nabla L^N(\vec{W}(j))\|_F^2 \quad \text{for any } k \in \mathbb{N},$$

hence, taking the limit,

$$L^N(\vec{W}(0)) \geq \sigma \sum_{k=0}^{\infty} \eta_k \|\nabla L^N(\vec{W}(k))\|_F^2.$$

Assume now that $c \neq 0$. Then $c > 0$ and there exists $k_0 \in \mathbb{N}$ such that

$$\|\nabla L^N(\vec{W}(k))\|_F \geq \frac{c}{2} \quad \forall k \geq k_0.$$

But then

$$L^N(\vec{W}(0)) \geq \sigma \sum_{k=k_0}^{\infty} \eta_k \|\nabla L^N(\vec{W}(k))\|_F^2 \geq \frac{c^2}{4} \sigma \sum_{k=k_0}^{\infty} \eta_k,$$

which by $\sigma > 0$ contradicts our assumption that $\sum_{k=0}^{\infty} \eta_k = \infty$. Thus indeed $c = 0$ and \vec{W}^* is a critical point of L^N . \square

4. CONVERGENCE TO A GLOBAL MINIMUM FOR ALMOST ALL INITIALIZATIONS

Let us now transfer [5, Theorem 6.12] to our situation of the gradient descent method by showing Theorem 2.6. Our proof is based on the following abstract theorem, which basically states that gradient descent schemes avoid strict saddle points for almost all initializations. The case of constant step sizes (condition (1)) was shown in [14, Proposition 1], while the one for step sizes converging to zero was proven in [15, Theorem 5.1]. We call a critical point z^* of a twice continuously differentiable function f a strict saddle point if the Hessian $\nabla^2 f(z^*)$ has at least one negative eigenvalue. Note that an analysis of the strict saddle points of L^N has been performed in [5], extending [12, 17].

Theorem 4.1. *Let $f : \mathbb{R}^p \rightarrow \mathbb{R}$ be a twice continuously differentiable function and consider the gradient descent scheme*

$$z(k+1) = z(k) - \eta_k \nabla f(z(k)),$$

where (η_k) satisfies one of the following conditions.

- (1) The sequence (η_k) is constant, i.e., $\eta_k = \eta$ for some $\eta > 0$ for all $k \in \mathbb{N}$.

(2) It holds

$$\eta_k \geq C \frac{1}{k} \quad \text{for some } C > 0 \quad \text{and} \quad \lim_{k \rightarrow \infty} \eta_k = 0.$$

Then the set of initializations $z(0) \in \mathbb{R}^p$ such that $(z(k))_k$ converges to a strict saddle point of f has measure zero.

Now we are ready to prove Theorem 2.6.

Proof. Due to definitions (2.12), (2.13) of the constants $\delta_{\mathcal{B}}$, $L_{\mathcal{B}}$ and $M_{\mathcal{B}}$ together with condition (2.14) on the step sizes η_k , the conditions of Theorem 2.4 are satisfied for each initialization $\vec{W}(0) \in \mathcal{B}$. Hence, $\vec{W}(k)$ converges to a critical point of L^N for all $\vec{W}(0) \in \mathcal{B}$. By Theorem 4.1 convergence of gradient descent with initial values in \mathcal{B} and with step sizes η_k to a strict saddle point occurs only for subset of \mathcal{B} that has measure zero.

The rest of the proof is the same as the corresponding reasoning in the proof of [5, Theorem 6.12]. Let us repeat only the main aspects from [5]. We denote by $\vec{W} = (W_1, \dots, W_N)$ the limit of $\vec{W}(k)$, $W = W_N \cdots W_1$ and $k = \text{rank}(W)$. Then $k \leq r$ and W is a critical point of L^1 restricted to manifold \mathcal{M}_k of rank k matrices [5, Proposition 6.8(a)]. Then [5, Proposition 6.6(1)] implies that $k \leq q$. If W is not a global minimizer of L^1 restricted to \mathcal{M}_k then W is a strict saddle point of L^N by [5, Proposition 6.9]. As argued above, the set of initializations converging to such a point has measure zero, showing part (a). (Note that for $N \geq 3$ and $k < \min\{r, q\}$ a global minimizer of L^1 restricted to \mathcal{M}_k may correspond to a non-strict saddle point \vec{W} of L^N , see [5, Proposition 6.10].) If $N = 2$, then by [5, Proposition 6.11] any critical point $\vec{W} = (W_1, W_2)$ such that $W = W_2 W_1$ is a global minimum of L^1 restricted to \mathcal{M}_k for some $k < \bar{r}$ is a strict saddle point of L^2 , which shows part (b) of the theorem. \square

5. NUMERICAL EXPERIMENTS

In this section we illustrate our theoretical results with numerical experiments. In particular, we test convergence of gradient descent for various choices of constant and decreasing step sizes and with $N = 2$, $N = 3$ and $N = 5$ layers.

The sample size is chosen as $m = 3 \cdot d$ with $d = 70$. For our experiments we generate our dataset $X \in \mathbb{R}^{d_x \times m}$ randomly with entries drawn from a mean zero Gaussian distribution with variance $\sigma^2 = 1/d$, where $d_x = d$. The data matrix $Y \in \mathbb{R}^{d \times m}$ is a random matrix of rank $r = 2$, which is generated as described below. We initialize the weight matrices $W_j \in \mathbb{R}^{d_j \times d_{j-1}}$ such that $\vec{W}(0) = (W_1, \dots, W_N)$ is balanced, i.e., has balancedness constant 0 so that $\alpha = 0$ in Theorems 2.4 and 2.6, in the following way. The rank parameter is chosen as $r = 2$ and the dimensions d_j as

$$d_0 = d, \quad d_1 = r, \quad d_j = \text{round} \left(r + (j-1) \frac{d-r}{N-1} \right), \quad j = 2, \dots, N,$$

where $\text{round}(z)$ rounds a real number z to the nearest integer. We randomly generate orthogonal matrices $U_1 \in \mathbb{R}^{d \times d}$, $V_j \in \mathbb{R}^{d_j \times d_j}$, $j = 1, \dots, N$, according to the uniform distribution on the corresponding unitary groups and let $U_j \in \mathbb{R}^{d_j \times d_1}$, $j = 2, \dots, N$ be the matrix composed of the first d_1 columns of V_{j-1} . We then set

$$W_j = V_j I_{d_j, d_1} U_j^T,$$

where for any $n_1, n_2 \in \mathbb{N}$ the matrix $I_{n_1, n_2} \in \mathbb{R}^{n_1 \times n_2}$ is a rectangular diagonal matrix with ones on the diagonal. By orthogonality and construction of U_{j+1} , it follows that for all $j = 1, \dots, N-1$, we have

$$W_{j+1}^T W_{j+1} = U_{j+1} I_{d_1, d_{j+1}} V_{j+1}^T V_{j+1} I_{d_{j+1}, d_1} U_{j+1}^T = U_{j+1} U_{j+1}^T = V_j I_{d_j, d_1} U_j^T U_j I_{d_1, d_j} V_j^T = W_j W_j^T$$

so that the tuple (W_1, \dots, W_N) is balanced.

The random matrix $Y \in \mathbb{R}^{d \times m}$ of rank 2 is generated as $Y = \widetilde{W}_N \cdots \widetilde{W}_1 X$ with matrices \widetilde{W}_j generated in the same way as the matrices W_j . (We decided to choose a matrix Y of rank 2 so that the global minimizer of L^1 is also of rank 2 and convergence to it means that L^N converges to zero, which is simple to check.)


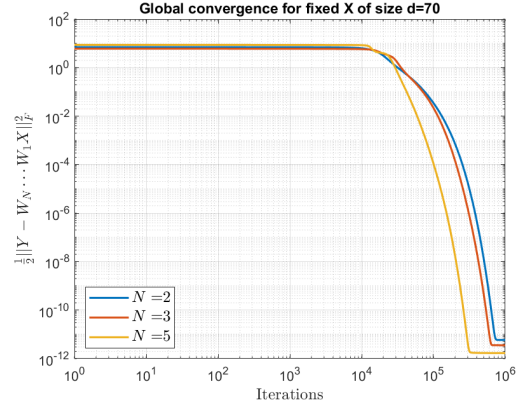

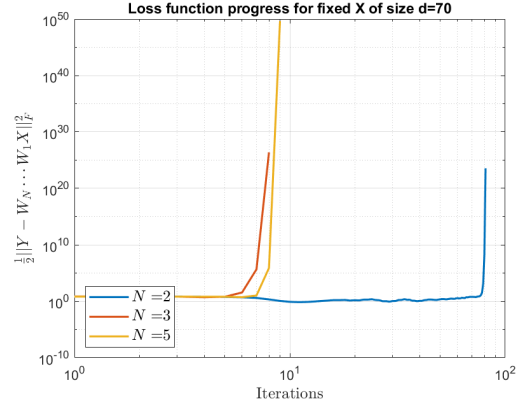
(A) Constant step size η meeting upper bound (5.1)(B) Constant step size $\eta = 10^{-5}$ (C) Constant step size $\eta = 10^{-3}$ (D) Constant step size $\eta = 0.05$

FIGURE 1. Progress of loss function L^N for training linear networks via gradient descent for various values of the constant step size

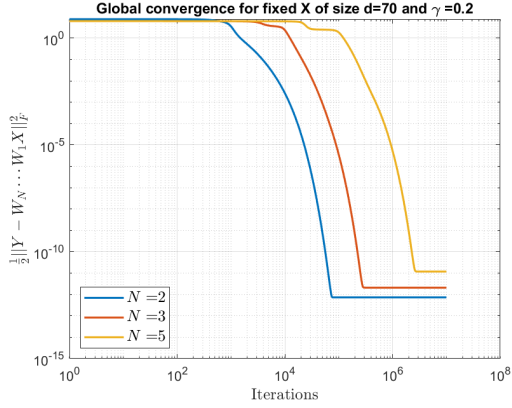



In our first set of experiments, we use a constant step size, i.e., $\eta_k = \eta$. Using $\alpha = 0$, the sufficient condition in Theorem 2.4 reads

$$\eta \leq \frac{2\delta}{4L^N(\vec{W}(0)) + \delta B_\delta}, \quad (5.1)$$

with B_δ in (2.9). As suggested in Remark 2.5(b) we choose

$$\delta = \frac{M^{\frac{2}{N}}}{N^3}.$$

In Figure 1 $L^N(\vec{W}(k))$ is plotted versus the iteration number. For the plot 1a the stepsize is chosen to exactly meet the upper bound in (5.1) (with $\delta = M^{2/N}/N^3$), resulting for this experiment in the values $\eta = 7.73 \cdot 10^{-4}$, $\eta = 1.29 \cdot 10^{-4}$ and $\eta = 3.91 \cdot 10^{-5}$ for depth 2, 3 and 5, respectively. For the plot 1b, the step size η is chosen somewhat smaller than the upper bound in (5.1), while for plots 1c and 1d the bound (5.1) is not satisfied. Since we observe convergence in plot 1c, this suggests that the bound of Theorem 2.4 may not be entirely sharp. But increasing the step size beyond a certain value leads to divergence as suggested by plot 2d, so that some bound on the step size is necessary (see also [7, Lemma A.1] for a necessary condition in a special case).

(A) Step sizes η_k in (5.2) with $\gamma = 0.2$ and a_1, a_2 in (5.3)(B) Step sizes η_k in (5.2) with $\gamma = 0.4$ and a_1, a_2 in (5.3)(C) Step sizes η_k in (5.2) with $\gamma = 0.2$, $a_1 = 0.1$, $a_2 = 0.01$ (D) Step sizes η_k in (5.2) with $\gamma = 0.2$, $a_1 = 0.05$, $a_2 = 0.1$ FIGURE 2. Gradient descent with decreasing step sizes η_k as in (5.2)

In our second set of experiments we use a sequence of step sizes η_k that converges to zero at various speeds. For some decay rate $\gamma \geq 0$ and some constants a_1, a_2 we set

$$\eta_k = \min \left\{ a_1, \frac{a_2}{(k+1)^\gamma} \right\} \quad \gamma \geq 0, \quad \text{for all } k \in \mathbb{N}. \quad (5.2)$$

The upper bound of Theorem 2.4 is satisfied for (see also the beginning of the proof of Proposition 3.3)

$$a_1 = a_2 = \frac{2(1-\sigma)}{B_\delta}, \quad \sigma = \frac{4L^N(\vec{W}(0))}{4L^N(\vec{W}(0)) + \delta B_\delta}. \quad (5.3)$$

Again, as suggested in Remark 2.5(b) we choose $\delta = \frac{1}{N^3} M^{\frac{2}{N}}$, when testing with these values for a_1 and a_2 .

The plots in Figure 2 illustrate the convergence behavior for various choices of the constants a_1, a_2 and decay rate γ in (5.2), for $N = 2, 3, 5$. Plot 2a and 2b both show convergence for the choices a_1, a_2 in (5.2) and for $\gamma = 0.2$ and $\gamma = 0.4$ respectively, leading to step sizes satisfying the condition of Theorem 2.4. (In these experiments, the resulting values of $a_1 = a_2$ are $a_1 = 7.73 \cdot 10^{-4}$ for $N = 2$, $a_1 = 1.29 \cdot 10^{-4}$ for $N = 3$ and $a_1 = 3.91 \cdot 10^{-5}$ for $N = 5$.) Comparing the two plots, as well as with the plots for constant step size in Figure 1, shows that fast decay of the step size leads to slower convergence of gradient descent, as expected. Note that we observe that larger values of γ are possible, but will further slow down convergence, so that we decided to omit the corresponding experiments here.

Plot 2c shows convergence for a decay rate of $y = 0.2$ even though the constants a_1 and a_2 are such that η_k does not satisfy the bound of Theorem 2.4 for all k , while further increasing the value of a_2 leads to divergence as illustrated in Plot 2d.

ACKNOWLEDGEMENTS

All authors acknowledge funding by DAAD (German Foreign Exchange Service) through the project *Understanding stochastic gradient descent in deep learning* (grant no: 57417829).

REFERENCES

- [1] P.-A. Absil, R. Mahony, and B. Andrews. Convergence of the iterates of descent methods for analytic cost functions. *SIAM Journal on Optimization*, 16(2):531–547, 2005.
- [2] S. Arora, N. Cohen, N. Golowich, and W. Hu. A convergence analysis of gradient descent for deep linear neural networks. In *International Conference on Learning Representations*, 2019.
- [3] S. Arora, N. Cohen, and E. Hazan. On the optimization of deep networks: Implicit acceleration by overparameterization. In *International Conference on Machine Learning*, 2018.
- [4] S. Arora, N. Cohen, W. Hu, and Y. Luo. Implicit regularization in deep matrix factorization. In *Advances in Neural Information Processing Systems*, pages 7413–7424, 2019.
- [5] B. Bah, H. Rauhut, U. Terstiege, and M. Westdickenberg. Learning deep linear neural networks: Riemannian gradient flows and convergence to global minimizers. *Information and Inference: A Journal of the IMA*, to appear. DOI:10.1093/imaiai/iaaa039.
- [6] Y. Chitour, Z. Liao, and R. Couillet. A geometric approach of gradient descent algorithms in neural networks. *Preprint arXiv:1811.03568*, 2018.
- [7] H. Chou, C. Gieshoff, J. Maly, and H. Rauhut. Gradient descent for deep matrix factorization: Dynamics and implicit bias towards low rank. *Preprint arXiv:2011.13772*, 2020.
- [8] S. S. Du, W. Hu, and J. D. Lee. Algorithmic regularization in learning deep homogeneous models: Layers are automatically balanced. In *ICML 2018 Workshop on Nonconvex Optimization*, 2018.
- [9] S. S. Du, X. Zhai, B. Póczos, and A. Singh. Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*, 2019.
- [10] K. Geyer, A. Kyrillidis, and A. Kalev. Low-rank regularization and solution uniqueness in over-parameterized matrix sensing. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, pages 930–940, 2020.
- [11] S. Gunasekar, B. E. Woodworth, S. Bhojanapalli, B. Neyshabur, and N. Srebro. Implicit regularization in matrix factorization. In *Advances in Neural Information Processing Systems*, pages 6151–6159, 2017.
- [12] K. Kawaguchi. Deep learning without poor local minima. *Advances in Neural Information Processing Systems 29*, pages 586–594, 2016.
- [13] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang. On large-batch training for deep learning: Generalization gap and sharp minima. In *International Conference on Learning Representations*, 2017.
- [14] J. D. Lee, I. Panageas, G. Piliouras, M. Simchowitz, M. I. Jordan, and B. Recht. First-order methods almost always avoid strict saddle points. *Mathematical Programming*, 176(1):311–337, 2019.
- [15] I. Panageas, G. Piliouras, and X. Wang. First-order methods almost always avoid saddle points: The case of vanishing step-sizes. In *Conference on Neural Information Processing Systems*, 2019.
- [16] N. Razin and N. Cohen. Implicit regularization in deep learning may not be explainable by norms. In *Conference on Neural Information Processing Systems*, 2020.
- [17] M. Trager, K. Kohn, and J. Bruna. Pure and spurious critical points: a geometric study of linear networks. In *International Conference on Learning Representations*, 2020.
- [18] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017.