

TOWARDS RESOLVING THE IMPLICIT BIAS OF GRADIENT DESCENT FOR MATRIX FACTORIZATION: GREEDY LOW-RANK LEARNING

Zhiyuan Li*, Yuping Luo*

Princeton University

{zhiyuanli, yupingl}@cs.princeton.edu

Kaifeng Lyu*

Tsinghua University

vfleaking@gmail.com

ABSTRACT

Matrix factorization is a simple and natural test-bed to investigate the implicit regularization of gradient descent. Gunasekar et al. (2017) conjectured that Gradient Flow with infinitesimal initialization converges to the solution that minimizes the nuclear norm, but a series of recent papers argued that the language of norm minimization is not sufficient to give a full characterization for the implicit regularization. In this work, we provide theoretical and empirical evidence that for depth-2 matrix factorization, gradient flow with infinitesimal initialization is mathematically equivalent to a simple heuristic rank minimization algorithm, Greedy Low-Rank Learning, under some reasonable assumptions. This generalizes the rank minimization view from previous works to a much broader setting and enables us to construct counter-examples to refute the conjecture from Gunasekar et al. (2017). We also extend the results to the case where depth ≥ 3 , and we show that the benefit of being deeper is that the above convergence has a much weaker dependence over initialization magnitude so that this rank minimization is more likely to take effect for initialization with practical scale.

1 INTRODUCTION

There are usually far more learnable parameters in deep neural nets than the number of training data, but still deep learning works well on real-world tasks. Even with explicit regularization, the model complexity of state-of-the-art neural nets is so large that they can fit randomly labeled data easily (Zhang et al., 2017). Towards explaining the mystery of generalization, we must understand what kind of implicit regularization does Gradient Descent (GD) impose during training. Ideally, we are hoping for a nice mathematical characterization of how GD constrains the set of functions that can be expressed by a trained neural net.

As a direct analysis for deep neural nets could be quite hard, a line of works turned to study the implicit regularization on simpler problems to get inspirations, for example, low-rank matrix factorization, a fundamental problem in machine learning and information process. Given a set of observations about an unknown matrix $\mathbf{W}^* \in \mathbb{R}^{d \times d}$ of rank $r^* \ll d$, one needs to find a low-rank solution \mathbf{W} that is compatible with the given observations. Examples include matrix sensing, matrix completion, phase retrieval, robust principal component analysis, just to name a few (see Chi et al. 2019 for a survey). When \mathbf{W}^* is symmetric and positive semidefinite, one way to solve all these problems is to parameterize \mathbf{W} as $\mathbf{W} = \mathbf{U}\mathbf{U}^\top$ for $\mathbf{U} \in \mathbb{R}^{d \times r}$ and optimize $\mathcal{L}(\mathbf{U}) := \frac{1}{2}f(\mathbf{U}\mathbf{U}^\top)$, where $f(\cdot)$ is some empirical risk function depending on the observations, and r is the rank constraint. In theory, if the rank constraint is too loose, the solutions do not have to be low-rank and we may fail to recover \mathbf{W}^* . However, even in the case where the rank is unconstrained (i.e., $r = d$), GD with small initialization can still get good performance in practice. This empirical observation reveals that the implicit regularization of GD exists even in this simple matrix factorization problem, but its mechanism is still on debate. Gunasekar et al. (2017) proved that Gradient Flow (GD with infinitesimal step size, a.k.a., GF) with infinitesimal initialization finds the minimum nuclear norm solution in a special case of matrix sensing, and further conjectured this holds in general.

Conjecture 1.1 (Gunasekar et al. 2017, informal). With sufficiently small initialization, GF converges to the minimum nuclear norm solution of matrix sensing.

*Alphabet ordering.

Subsequently, Arora et al. (2019a) challenged this view by arguing that a simple mathematical norm may not be a sufficient language for characterizing implicit regularization. One example illustrated in Arora et al. (2019a) is regarding matrix sensing with a single observation. They showed that GD with small initialization enhances the growth of large singular values of the solution and attenuates that of smaller ones. This enhancement/attenuation effect encourages low-rank, and it is further intensified with depth in deep matrix factorization (i.e., GD optimizes $f(U_1 \cdots U_L)$ for $L \geq 2$). However, these are not captured by the nuclear norm alone. Gidel et al. (2019); Gissin et al. (2020) further exploited this idea and showed in the special case of full-observation matrix sensing that GF learns solutions with gradually increasing rank. Razin and Cohen (2020) showed in a simple class of matrix completion problems that GF decreases the rank along the trajectory while any norm grows towards infinity. More aggressively, they conjectured that the implicit regularization can be explained by rank minimization rather than norm minimization.

Our Contributions. In this paper, we move one further step towards resolving the implicit regularization in the matrix factorization problem. Our theoretical results show that GD performs rank minimization via a greedy process in a broader setting. Specifically, we provide theoretical evidence that GF with infinitesimal initialization is in general mathematically equivalent to another algorithm called *Greedy Low-Rank Learning* (GLRL). At a high level, GLRL is a greedy algorithm that performs rank-constrained optimization and relaxes the rank constraint by 1 whenever it fails to reach a global minimizer of $f(\cdot)$ with the current rank constraint. As a by-product, we refute Conjecture 1.1 by demonstrating a counterexample (Example 5.9).

We also extend our results to deep matrix factorization Section 6, where we prove that the trajectory of GF with infinitesimal identity initialization converges to a deep version of GLRL, at least in the early stage of the optimization. We also use this result to confirm the intuition achieved on toy models (Gissin et al., 2020), that benefits of depth in matrix factorization is to encourage rank minimization even for initialization with a relatively larger scale, and thus it is more likely to happen in practice. This shows that describing the implicit regularization using GLRL is more expressive than using the language of norm minimization. We validate all our results with experiments in Appendix E.

2 RELATED WORKS

Norm Minimization. The view of norm minimization, or the closely related view of margin maximization, has been explored in different settings. Besides the nuclear norm minimization for matrix factorization (Gunasekar et al., 2017) discussed in the introduction, previous works have also studied the norm minimization/margin maximization for linear regression (Wilson et al., 2017; Soudry et al., 2018a;b; Nacson et al., 2019b;c; Ji and Telgarsky, 2019b), deep linear neural nets (Ji and Telgarsky, 2019a; Gunasekar et al., 2018), homogeneous neural nets (Nacson et al., 2019a; Lyu and Li, 2020), ultra-wide neural nets (Jacot et al., 2018; Arora et al., 2019b; Chizat and Bach, 2020).

Small Initialization and Rank Minimization. The initialization scale can greatly influence the implicit regularization. A sufficiently large initialization can make the training dynamics fall into the lazy training regime defined by Chizat et al. (2019) and diminish test accuracy. Using small initialization is particularly important to bias gradient descent to low-rank solutions for matrix factorization, as empirically observed by Gunasekar et al. (2017). Arora et al. (2019a); Gidel et al. (2019); Gissin et al. (2020); Razin and Cohen (2020) studied how gradient flow with small initialization encourages low-rank in simple settings, as discussed in the introduction. Li et al. (2018) proved recovery guarantees for gradient flow solving matrix sensing under Restricted Isometry Property (RIP), but the proof cannot be generalized easily to the case without RIP. Belabbas (2020) made attempts to prove that gradient flow is approximately rank-1 in the very early phase of training, but it does not exclude the possibility that the approximation error explodes later and gradient flow is not converging to low-rank solutions. Compared to these works, the current paper studies how GF encourages low-rank in a much broader setting.

3 BACKGROUND

Notations. For two matrices A, B , we define $\langle A, B \rangle := \text{Tr}(AB^\top)$ as their inner product. We use $\|A\|_F$, $\|A\|_*$ and $\|A\|_2$ to denote the Frobenius norm, nuclear norm and the largest singular value of A respectively. For a matrix $A \in \mathbb{R}^{d \times d}$, we use $\lambda_1(A), \dots, \lambda_d(A)$ to denote the eigenvalues of A in decreasing order (if they are all reals). We define \mathbb{S}_d as the set of symmetric $d \times d$ matrices and $\mathbb{S}_d^+ \subseteq \mathbb{S}_d$ as the set of positive semidefinite (PSD) matrices. We write $A \succeq B$ or $B \preceq A$ iff $A - B$ is PSD. We use $\mathbb{S}_{d,r}^+$, $\mathbb{S}_{d,\leq r}^+$ to denote the set of $d \times d$ PSD matrices with rank $= r, \leq r$ respectively.

Matrix Factorization. Matrix factorization problem asks one to optimize $\mathcal{L}(\mathbf{U}, \mathbf{V}) := \frac{1}{2}f(\mathbf{UV}^\top)$ among $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{d \times r}$, where $f: \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$ is a convex function and in this paper we assume f is \mathcal{C}^3 -smooth. A notable example is matrix sensing. There is an unknown rank- r^* matrix $\mathbf{W}^* \in \mathbb{R}^{d \times d}$ with $r^* \ll d$. Given m measurements $\mathbf{X}_1, \dots, \mathbf{X}_m \in \mathbb{R}^{d \times d}$, one can observe $y_i := \langle \mathbf{X}_i, \mathbf{W}^* \rangle$ through each measurement. The goal of matrix sensing is to reconstruct \mathbf{W}^* via minimizing $f(\mathbf{W}) := \frac{1}{2} \sum_{i=1}^m (\langle \mathbf{W}, \mathbf{X}_i \rangle - y_i)^2$. Matrix completion is a notable special case of matrix sensing in which every measurement has the form $\mathbf{X}_i = \mathbf{e}_{p_i} \mathbf{e}_{q_i}^\top$, where $\{\mathbf{e}_1, \dots, \mathbf{e}_d\}$ stands for the standard basis (i.e., exactly one entry is observed through each measurement).

Note that matrix factorization in the general case can be reduced to this symmetric case: let $\mathbf{U}' = \begin{bmatrix} \mathbf{U} \\ \mathbf{V} \end{bmatrix} \in \mathbb{R}^{2d \times r}$, $f'(\begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix}) = \frac{1}{2}f(\mathbf{B}) + \frac{1}{2}f(\mathbf{C})$, then $f(\mathbf{UV}^\top) = f'(\mathbf{U}'\mathbf{U}'^\top)$. So in this paper we focus on the symmetric case as in previous works (Gunasekar et al., 2017), i.e., finding a low-rank solution for the convex optimization problem: $\min_{\mathbf{W} \succeq \mathbf{0}} f(\mathbf{W})$. For this, we parameterize \mathbf{W} as $\mathbf{W} = \mathbf{UU}^\top$ for $\mathbf{U} \in \mathbb{R}^{d \times r}$ and optimize $\mathcal{L}(\mathbf{U}) := \frac{1}{2}f(\mathbf{UU}^\top)$. We assume WLOG throughout this paper that $f(\mathbf{W}) = f(\mathbf{W}^\top)$; otherwise, we can set $f'(\mathbf{W}) = \frac{1}{2}(f(\mathbf{W}) + f(\mathbf{W}^\top))$ so that $f'(\mathbf{W}) = f'(\mathbf{W}^\top)$ while $\mathcal{L}(\mathbf{U}) = \frac{1}{2}f'(\mathbf{UU}^\top)$ is unaffected. This assumption makes $\nabla f(\mathbf{W})$ symmetric for every symmetric \mathbf{W} .

Gradient Flow. In this paper, we analyze Gradient Flow (GF) for symmetric matrix factorization, defined as the solution of the following ODE for $\mathbf{U}(t) \in \mathbb{R}^{d \times r}$:

$$\frac{d\mathbf{U}}{dt} = -\nabla \mathcal{L}(\mathbf{U}) = -\nabla f(\mathbf{UU}^\top)\mathbf{U}. \quad (1)$$

Let $\mathbf{W}(t) = \mathbf{U}(t)\mathbf{U}(t)^\top \in \mathbb{R}^{d \times d}$. Then the following end-to-end dynamics holds for $\mathbf{W}(t)$:

$$\frac{d\mathbf{W}}{dt} = -\mathbf{W}\nabla f(\mathbf{W}) - \nabla f(\mathbf{W})\mathbf{W} =: \mathbf{g}(\mathbf{W}). \quad (2)$$

We use $\phi(\mathbf{W}_0, t)$ to denote the matrix $\mathbf{W}(t)$ in (2) when $\mathbf{W}(0) = \mathbf{W}_0 \succeq \mathbf{0}$. Throughout this paper, we assume $\phi(\mathbf{W}_0, t)$ exists for all $t \in \mathbb{R}$, $\mathbf{W}_0 \succeq \mathbf{0}$. It is easy to prove that \mathbf{U} is a stationary point of $\mathcal{L}(\cdot)$ (i.e., $\nabla \mathcal{L}(\mathbf{U}) = \mathbf{0}$) iff $\mathbf{W} = \mathbf{UU}^\top$ is a critical point of (2) (i.e., $\mathbf{g}(\mathbf{W}) = \mathbf{0}$); see Lemma C.1 for a proof. If \mathbf{W} is a minimizer of $f(\cdot)$ in \mathbb{S}_d^+ , then \mathbf{W} is a critical point of (2), but the reverse may not be true, e.g., $\mathbf{g}(\mathbf{0}) = \mathbf{0}$, but $\mathbf{0}$ is not necessarily a minimizer.

In this paper, we particularly focus on the overparameterized case, where $r = d$, to understand the implicit regularization of GF when there is no rank constraint for the matrix \mathbf{W} .

4 WARMUP EXAMPLES

First, we illustrate how GD performs greedy learning using two warmup examples.

Linearization Around the Origin. In general, for a loss function $\mathcal{L}(\mathbf{U}) = \frac{1}{2}f(\mathbf{UU}^\top)$, we can always apply Taylor expansion $f(\mathbf{W}) \approx f(\mathbf{0}) + \langle \mathbf{W}, \nabla f(\mathbf{0}) \rangle$ around the origin to approximate it with a linear function. This motivates us to study the linear case: $f(\mathbf{W}) := f_0 - \langle \mathbf{W}, \mathbf{Q} \rangle$ for some symmetric matrix \mathbf{Q} . In this case, the matrix \mathbf{U} follows the ODE, $\frac{d\mathbf{U}}{dt} = \mathbf{QU}$, which can be understood as a continuous version of the classical power iteration method for solving the top eigenvector. Let $\mathbf{Q} := \sum_{i=1}^d \mu_i \mathbf{v}_i \mathbf{v}_i^\top$ be the eigendecomposition of \mathbf{Q} , where $\mu_1 \geq \mu_2 \geq \dots \geq \mu_d$ and $\mathbf{v}_1, \dots, \mathbf{v}_d$ are orthogonal to each other. Then we can write the solution as:

$$\mathbf{U}(t) = e^{t\mathbf{Q}}\mathbf{U}(0) = \left(\sum_{i=1}^d e^{\mu_i t} \mathbf{v}_i \mathbf{v}_i^\top \right) \mathbf{U}(0). \quad (3)$$

When $\mu_1 > \mu_2$, the ratio between $e^{\mu_1 t}$ and $e^{\mu_i t}$ for $i \neq 1$ increases exponentially fast. As $t \rightarrow +\infty$, $\mathbf{U}(t)$ and $\mathbf{W}(t)$ become approximately rank-1 as long as $\mathbf{v}_1^\top \mathbf{U}(0) \neq \mathbf{0}$, i.e.,

$$\lim_{t \rightarrow \infty} e^{-\mu_1 t} \mathbf{U}(t) = \mathbf{v}_1 \mathbf{v}_1^\top \mathbf{U}(0), \quad \lim_{t \rightarrow \infty} e^{-2\mu_1 t} \mathbf{W}(t) = (\mathbf{v}_1^\top \mathbf{W}(0) \mathbf{v}_1) \mathbf{v}_1 \mathbf{v}_1^\top. \quad (4)$$

The analysis for the simple linear case reveals that GD encourages low-rank through a process similar to power iteration. However, $f(\mathbf{W})$ is non-linear in general, and the linear approximation is close to $f(\mathbf{W})$ only if \mathbf{W} is very small. With sufficiently small initialization, we can imagine that GD still resembles the above power iteration in the early phase of the optimization. But what if $\mathbf{W}(t)$ grows to be so large that the linear approximation is far from the actual $f(\mathbf{W})$?

Full-observation Matrix Sensing. To understand the dynamics of GD when the linearization fails, we now consider a well-studied special case (Gissin et al., 2020): $\mathcal{L}(\mathbf{U}) = \frac{1}{2}f(\mathbf{UU}^\top)$, $f(\mathbf{W}) = \frac{1}{2} \|\mathbf{W} - \mathbf{W}^*\|_{\mathbb{F}}^2$ for some unknown PSD matrix \mathbf{W}^* . GF in this case can be written as:

$$\frac{dU}{dt} = (\mathbf{W}^* - UU^\top)U, \quad \frac{dW}{dt} = (\mathbf{W}^* - W)W + W(\mathbf{W}^* - W). \quad (5)$$

Let $\mathbf{W}^* := \sum_{i=1}^d \mu_i \mathbf{v}_i \mathbf{v}_i^\top$ be the eigendecomposition of \mathbf{W}^* . Our previous analysis shows that the dynamics is approximately $\frac{dU}{dt} = \mathbf{W}^*U$ in the early phase and thus encourages low-rank.

To get a sense for the later phases, we simplify the setting by specifying $U(0) = \sqrt{\alpha}I$ for a small number α . We can write $W(0)$ and \mathbf{W}^* as diagonal matrices $W(0) = \text{diag}(\alpha, \alpha, \dots, \alpha)$, $\mathbf{W}^* = \text{diag}(\mu_1, \mu_2, \dots, \mu_d)$ with respect to the basis $\mathbf{v}_1, \dots, \mathbf{v}_d$. It is easy to see that $W(t)$ is always a diagonal matrix, since the time derivatives of non-diagonal coordinates stay 0 during training. Let $W(t) = \text{diag}(\sigma_1(t), \sigma_2(t), \dots, \sigma_d(t))$, then $\sigma_i(t)$ satisfies the dynamical equation $\frac{d}{dt}\sigma_i(t) = 2\sigma_i(t)(\mu_i - \sigma_i(t))$, and thus $\sigma_i(t) = \frac{\alpha\mu_i}{\alpha + (\mu_i - \alpha)e^{-2\mu_i t}}$. This shows that every $\sigma_i(t)$ increases from α to μ_i over time. As $\alpha \rightarrow 0$, every $\sigma_i(t)$ has a sharp transition from near 0 to near μ_i at time roughly $(\frac{1}{2\mu_i} + o(1)) \log \frac{1}{\alpha}$, which can be seen from the following limit:

$$\lim_{\alpha \rightarrow 0} \sigma_i \left(\left(\frac{1}{2\mu_i} + c \right) \log(1/\alpha) \right) = \lim_{\alpha \rightarrow 0} \frac{\alpha\mu_i}{\alpha + (\mu_i - \alpha)\alpha^{1+2c\mu_i}} = \begin{cases} 0 & c \in (-\frac{1}{2\mu_i}, 0), \\ \mu_i & c \in (0, +\infty). \end{cases}$$

This means for every $q \in (\frac{1}{2\mu_i}, \frac{1}{2\mu_{i+1}})$ for $i = 1, \dots, d-1$ (or $q \in (\frac{1}{2\mu_i}, +\infty)$ for $i = d$), $\lim_{\alpha \rightarrow 0} W(q \log(1/\alpha)) = \text{diag}(\mu_1, \mu_2, \dots, \mu_i, 0, 0, \dots, 0)$. Therefore, when the initialization is sufficiently small, GF learns each component of \mathbf{W}^* one by one, according to the relative order of eigenvalues. At a high level, this shows a greedy nature of GD: GD starts learning with simple models; whenever it underfits, it increases the model complexity (which is rank in our case). This is also called *sequential learning* or *incremental learning* (Gidel et al., 2019; Gissin et al., 2020).

However, it is unclear how and why this sequential learning/incremental learning can occur in general. Through the first warmup example, we may understand why GD learns a rank-1 matrix in the early phase, but does GD always learn solutions with rank 2, 3, 4, ... sequentially? If true, what is the mechanism behind this? The current paper answers the questions by providing both theoretical and empirical evidence that the greedy learning behavior does occur in general with a similar reason as for the first warmup example.

5 GREEDY LOW-RANK LEARNING (GLRL)

In this section, we present a trajectory-based analysis for the implicit bias of GF on matrix factorization. Our main result is that GF with infinitesimal initialization is generically the same as that of a simple greedy algorithm, *Greedy Low-Rank Learning* (GLRL, Algorithm 1). See Appendix A for a comparison with existing greedy algorithms for rank-constrained optimization.

The GLRL algorithm consists of several phases, numbered from 1. In phase r , GLRL increases the rank constraint to r and optimizes $\mathcal{L}(U_r) := \frac{1}{2}f(U_r U_r^\top)$ among $U_r \in \mathbb{R}^{d \times r}$ via GD until it reaches a stationary point $U_r(\infty)$, i.e., $\nabla \mathcal{L}(U_r(\infty)) = \mathbf{0}$. At convergence, $W_r := U_r(\infty)U_r^\top(\infty)$ is a critical point of (2), and we call it the r -th *critical point* of GLRL. If W_r is further a minimizer of $f(\cdot)$ in \mathbb{S}_d^+ , or equivalently, $\lambda_1(-\nabla f(W_r)) \leq 0$ (see Lemma C.2), GLRL returns W_r ; otherwise GLRL enters phase $r+1$.

To set the initial point of GD in phase r , GLRL appends a small column vector $\delta_r \in \mathbb{R}^d$ to the resulting stationary point $U_{r-1}(\infty)$ from the last phase, i.e., $U_r(0) \leftarrow [U_{r-1}(\infty) \ \delta_r] \in \mathbb{R}^{d \times r}$ (in the case of $r=1$, $U_1(0) \leftarrow [\delta_1] \in \mathbb{R}^{d \times 1}$). In this way, $U_r(0)U_r^\top(0) = W_{r-1} + \delta_r \delta_r^\top$ is perturbed away from the $(r-1)$ -th critical point. In GLRL, we set $\delta_r = \sqrt{\epsilon} \mathbf{u}_r$, where \mathbf{u}_r is the top eigenvector of $-\nabla f(W_r)$ with unit norm $\|\mathbf{u}_r\|_2 = 1$, and $\epsilon > 0$ is a parameter controlling the magnitude of perturbation (preferably very small). Note that it is guaranteed that $\lambda_1(-\nabla f(W_{r-1})) > 0$; otherwise W_{r-1} is a minimizer of the convex function $f(\cdot)$ in \mathbb{S}_d^+ and GLRL exits before phase r .

Algorithm 1: Greedy Low-Rank Learning

parameter: step size $\eta > 0$; small $\epsilon > 0$

$r \leftarrow 0, W_0 \leftarrow \mathbf{0} \in \mathbb{R}^{d \times d}$, and

$U_0(\infty) \in \mathbb{R}^{d \times 0}$

while $\lambda_1(-\nabla f(W_r)) > 0$ **do**

$r \leftarrow r + 1$

$\mathbf{u}_r \leftarrow$ unit top eigenvector of $-\nabla f(W_{r-1})$

$U_r(0) \leftarrow [U_{r-1}(\infty) \ \sqrt{\epsilon} \mathbf{u}_r] \in \mathbb{R}^{d \times r}$

for $t = 0, 1, \dots$ **do**

$U_r(t+1) \leftarrow U_r(t) - \eta \nabla \mathcal{L}(U_r(t))$

$W_r \leftarrow U_r(\infty)U_r^\top(\infty)$ ^a

return W_r

^aIn practice, we approximate the infinite time limit by running sufficiently many steps.

Trajectory of GLRL. We define the (limiting) trajectory of GLRL by taking the learning rate $\eta \rightarrow 0$. The goal is to show that the trajectory of GLRL is close to that of GF with infinitesimal initialization. Recall that $\phi(\mathbf{W}_0, t)$ stands for the solution $\mathbf{W}(t)$ in (2) when $\mathbf{W}(0) = \mathbf{W}_0$.

Definition 5.1 (Trajectory of GLRL). Let $\overline{\mathbf{W}}_{0,\epsilon} := \mathbf{0}$ be the 0th critical point of GLRL. For every $r \geq 1$, if the $(r-1)$ -th critical point $\overline{\mathbf{W}}_{r-1,\epsilon}$ exists and is not a minimizer of $f(\cdot)$ in \mathbb{S}_d^+ , we define $\mathbf{W}_{r,\epsilon}^G(t) := \phi(\overline{\mathbf{W}}_{r-1,\epsilon} + \epsilon \mathbf{u}_{r,\epsilon} \mathbf{u}_{r,\epsilon}^\top, t)$, where $\mathbf{u}_{r,\epsilon}$ is a top eigenvector of $\nabla f(\overline{\mathbf{W}}_{r-1,\epsilon})$ with unit norm, $\|\mathbf{u}_{r,\epsilon}\|_2 = 1$. We define $\overline{\mathbf{W}}_{r,\epsilon} := \lim_{t \rightarrow +\infty} \mathbf{W}_{r,\epsilon}^G(t)$ to be the r -th critical point of GLRL if the limit exists.

Throughout this paper, we always focus on the case where the top eigenvalue of every $\nabla f(\overline{\mathbf{W}}_{r-1,\epsilon})$ is unique. In this case, the trajectory of GLRL is unique for every $\epsilon > 0$, since the normalized top eigenvectors can only be $\pm \mathbf{u}_{r,\epsilon}$, and both of them lead to the same $\mathbf{W}_{r,\epsilon}^G(t)$.

5.1 THE LIMITING TRAJECTORY: A GENERAL THEOREM FOR DYNAMICAL SYSTEM

To prove the equivalence between GF and GLRL, we first introduce our high-level idea by analyzing the behavior of a more general dynamical system around its critical point, say $\mathbf{0}$. A specific example is (2) if we set $\boldsymbol{\theta}$ to be the vectorization of \mathbf{W} .

$$\frac{d\boldsymbol{\theta}}{dt} = \mathbf{g}(\boldsymbol{\theta}), \quad \text{where } \mathbf{g}(\mathbf{0}) = \mathbf{0}. \quad (6)$$

We use $\phi(\boldsymbol{\theta}_0, t)$ to denote the value of $\boldsymbol{\theta}(t)$ in the case of $\boldsymbol{\theta}(0) = \boldsymbol{\theta}_0$. We assume that $\mathbf{g}(\boldsymbol{\theta})$ is \mathcal{C}^2 -smooth with $\mathbf{J}(\boldsymbol{\theta})$ being the Jacobian matrix and $\phi(\boldsymbol{\theta}_0, t)$ exists for all $\boldsymbol{\theta}_0$ and t . For ease of presentation, in the main text we assume $\mathbf{J}(\mathbf{0})$ is diagonalizable over \mathbb{R} and defer the same result for the general case into Appendix G.3. Let $\mathbf{J}(\mathbf{0}) = \tilde{\mathbf{V}} \tilde{\mathbf{D}} \tilde{\mathbf{V}}^{-1}$ be the eigendecomposition, where $\tilde{\mathbf{V}}$ is an invertible matrix and $\tilde{\mathbf{D}} = \text{diag}(\tilde{\mu}_1, \dots, \tilde{\mu}_d)$ is the diagonal matrix consisting of the eigenvalues $\tilde{\mu}_1 \geq \tilde{\mu}_2 \geq \dots \geq \tilde{\mu}_d$. Let $\tilde{\mathbf{V}} = (\tilde{\mathbf{v}}_1, \dots, \tilde{\mathbf{v}}_d)$ and $\tilde{\mathbf{V}}^{-1} = (\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_d)^\top$, then $\tilde{\mathbf{u}}_i, \tilde{\mathbf{v}}_i$ are left and right eigenvectors associated with $\tilde{\mu}_i$ and $\tilde{\mathbf{u}}_i^\top \tilde{\mathbf{v}}_j = \delta_{ij}$. We can rewrite the eigendecomposition as $\mathbf{J}(\mathbf{0}) = \sum_{i=1}^d \tilde{\mu}_i \tilde{\mathbf{v}}_i \tilde{\mathbf{u}}_i^\top$. We also assume the top eigenvalue $\tilde{\mu}_1$ is positive and unique. Note $\tilde{\mu}_1 > 0$ means the critical point $\boldsymbol{\theta} = \mathbf{0}$ is unstable, and in matrix factorization it means $\mathbf{0}$ is a strict saddle point of $\mathcal{L}(\cdot)$. The key observation is that if the initialization is infinitesimal, the trajectory is almost uniquely determined. To be more precise, we need the following definition:

Definition 5.2. For any $\boldsymbol{\theta}_0 \in \mathbb{R}^d$ and $\mathbf{u} \in \mathbb{R}^d$, we say that $\{\boldsymbol{\theta}_\alpha\}_{\alpha \in (0,1)}$ converges to $\boldsymbol{\theta}_0$ with positive alignment with \mathbf{u} if $\lim_{\alpha \rightarrow 0} \boldsymbol{\theta}_\alpha = \boldsymbol{\theta}_0$ and $\liminf_{\alpha \rightarrow 0} \left\langle \frac{\boldsymbol{\theta}_\alpha - \boldsymbol{\theta}_0}{\|\boldsymbol{\theta}_\alpha - \boldsymbol{\theta}_0\|_2}, \mathbf{u} \right\rangle > 0$.

A special case is that the direction of $\boldsymbol{\theta}_\alpha - \boldsymbol{\theta}_0$ converges, i.e., $\bar{\boldsymbol{\theta}} := \lim_{\alpha \rightarrow 0} \frac{\boldsymbol{\theta}_\alpha - \boldsymbol{\theta}_0}{\|\boldsymbol{\theta}_\alpha - \boldsymbol{\theta}_0\|_2}$ exists. In this case, $\{\boldsymbol{\theta}_\alpha\}$ has positive alignment with either \mathbf{u} or $-\mathbf{u}$ except for a zero-measure subset of $\bar{\boldsymbol{\theta}}$. This means any convergent sequence generically falls into either of these two categories.

The following theorem shows that if the initial point $\boldsymbol{\theta}_\alpha$ converges to $\mathbf{0}$ with positive alignment with $\tilde{\mathbf{u}}_1$ as $\alpha \rightarrow 0$, the trajectory starting with $\boldsymbol{\theta}_\alpha$ converges to a unique trajectory $\mathbf{z}(t) := \phi(\alpha \tilde{\mathbf{v}}_1, t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha})$. By symmetry, there is another unique trajectory for sequences $\{\boldsymbol{\theta}_\alpha\}$ with positive alignment to $-\tilde{\mathbf{u}}_1$, which is $\mathbf{z}'(t) := \phi(-\alpha \tilde{\mathbf{v}}_1, t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha})$. This is somewhat surprising: different initial points should lead to very different trajectories, but our analysis shows that generically there are only two limiting trajectories for infinitesimal initialization. We will soon see how this theorem helps in our analysis for matrix factorization in Sections 5.2 and 5.3.

Theorem 5.3. Let $\mathbf{z}_\alpha(t) := \phi(\alpha \tilde{\mathbf{v}}_1, t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha})$ for every $\alpha > 0$, then $\mathbf{z}(t) := \lim_{\alpha \rightarrow 0} \mathbf{z}_\alpha(t)$ exists and is also a solution of (6), i.e., $\mathbf{z}(t) = \phi(\mathbf{z}(0), t)$. If $\boldsymbol{\delta}_\alpha$ converges to $\mathbf{0}$ with positive alignment with $\tilde{\mathbf{u}}_1$ as $\alpha \rightarrow 0$, then $\forall t \in \mathbb{R}$, there is a constant $C > 0$ such that

$$\left\| \phi\left(\boldsymbol{\delta}_\alpha, t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\langle \boldsymbol{\delta}_\alpha, \tilde{\mathbf{u}}_1 \rangle}\right) - \mathbf{z}(t) \right\|_2 \leq C \cdot \|\boldsymbol{\delta}_\alpha\|_2^{\frac{\tilde{\gamma}}{\tilde{\mu}_1 + \tilde{\gamma}}}, \quad (7)$$

for every sufficiently small α , where $\tilde{\gamma} := \tilde{\mu}_1 - \tilde{\mu}_2 > 0$ is the eigenvalue gap.

Proof sketch. The main idea is to linearize the dynamics near origin as we have done for the first warmup example. For sufficiently small $\boldsymbol{\theta}$, by Taylor expansion of $\mathbf{g}(\boldsymbol{\theta})$, the dynamics is approximately $\frac{d\boldsymbol{\theta}}{dt} \approx \mathbf{J}(\mathbf{0})\boldsymbol{\theta}$, which can be understood as a continuous version of power iteration. If the linear approximation is exact, then $\boldsymbol{\theta}(t) = e^{t\mathbf{J}(\mathbf{0})}\boldsymbol{\theta}(0)$. For large enough t_0 ,

$e^{t_0 \mathbf{J}(0)} = \sum_{i=1}^d e^{\tilde{\mu}_i t_0} \tilde{\mathbf{v}}_i \tilde{\mathbf{u}}_i^\top = e^{\tilde{\mu}_1 t_0} \tilde{\mathbf{v}}_1 \tilde{\mathbf{u}}_1^\top + O(e^{\tilde{\mu}_2 t_0})$. Therefore, as long as the initial point $\boldsymbol{\theta}(0)$ has a positive inner product with $\tilde{\mathbf{u}}_1$, $\boldsymbol{\theta}(t_0)$ should be very close to $\epsilon \tilde{\mathbf{v}}_1$ for some $\epsilon > 0$, and the rest of the trajectory after t_0 should be close to the trajectory starting from $\epsilon \tilde{\mathbf{v}}_1$. However, here is a tradeoff: we should choose t_0 to be large enough so that the power iteration takes effect; but if t_0 is so large that the norm of $\boldsymbol{\theta}(t_0)$ reaches a constant scale, then the linearization fails unavoidably. Nevertheless, if the initialization scale is sufficiently small, we show via a careful error analysis that there is always a suitable choice of t_0 such that $\boldsymbol{\theta}(t_0)$ is well approximated by $\epsilon \tilde{\mathbf{v}}_1$ and the difference between $\boldsymbol{\theta}(t_0 + t)$ and $\phi(\epsilon \tilde{\mathbf{v}}_1, t)$ is bounded as well. We defer the details to Appendix G. \square

5.2 EQUIVALENCE BETWEEN GD AND GLRL: RANK-ONE CASE

Now we establish the equivalence between GF and GLRL in the first phase. The main idea is to apply Theorem 5.3 on (2). For this, we need the following lemma on the eigenvalues and eigenvectors.

Lemma 5.4. *Let $\mathbf{g}(\mathbf{W}) := -\mathbf{W} \nabla f(\mathbf{W}) - \nabla f(\mathbf{W}) \mathbf{W}$ and $\mathbf{J}(\mathbf{W})$ be its Jacobian. Then $\mathbf{J}(\mathbf{0})$ is symmetric and thus diagonalizable. Let $-\nabla f(\mathbf{0}) = \sum_{i=1}^d \mu_i \mathbf{u}_{1[i]} \mathbf{u}_{1[i]}^\top$ be the eigendecomposition of the symmetric matrix $-\nabla f(\mathbf{0})$, where $\mu_1 \geq \mu_2 \geq \dots \geq \mu_d$. Then $\mathbf{J}(\mathbf{0})$ has the form:*

$$\mathbf{J}(\mathbf{0})[\Delta] = \sum_{i=1}^d \sum_{j=1}^d (\mu_i + \mu_j) \left\langle \Delta, \mathbf{u}_{1[i]} \mathbf{u}_{1[j]}^\top \right\rangle \mathbf{u}_{1[i]} \mathbf{u}_{1[j]}^\top, \quad (8)$$

where $\mathbf{J}(\mathbf{0})[\Delta]$ stands for the resulting matrix produced by left-multiplying $\mathbf{J}(\mathbf{0})$ to the vectorization of Δ . For every pair of $1 \leq i \leq j \leq d$, $\mu_i + \mu_j$ is an eigenvalue of $\mathbf{J}(\mathbf{0})$ and $\mathbf{u}_{1[i]} \mathbf{u}_{1[j]}^\top + \mathbf{u}_{1[j]} \mathbf{u}_{1[i]}^\top$ is a corresponding eigenvector. All the other eigenvalues are 0.

We simplify the notation by letting $\mathbf{u}_1 := \mathbf{u}_{1[1]}$. A direct corollary of Lemma 5.4 is that $\mathbf{u}_1 \mathbf{u}_1^\top$ is the top eigenvector of $\mathbf{J}(\mathbf{0})$. According to Theorem 5.3, now there are only two types of trajectories, which correspond to infinitesimal initialization $\mathbf{W}_\alpha \rightarrow \mathbf{0}$ with positive alignment with $\mathbf{u}_1 \mathbf{u}_1^\top$ or $-\mathbf{u}_1 \mathbf{u}_1^\top$. As the initialization must be PSD, $\mathbf{W}_\alpha \rightarrow \mathbf{0}$ cannot have positive alignment with $-\mathbf{u}_1 \mathbf{u}_1^\top$. For the former case, Theorem 5.6 below states that, for every fixed time t , the GF solution $\phi(\mathbf{W}_\alpha, T(\mathbf{W}_\alpha) + t)$ after shifting by a time offset $T(\mathbf{W}_\alpha) := \frac{1}{2\mu_1} \log(\langle \mathbf{W}_\alpha, \mathbf{u}_1 \mathbf{u}_1^\top \rangle^{-1})$ converges to the GLRL solution $\mathbf{W}_1^G(t)$ as $\mathbf{W}_\alpha \rightarrow \mathbf{0}$. The only assumption for this result is that $\mathbf{0}$ is not a minimizer of $f(\cdot)$ in \mathbb{S}_d^+ (which is equivalent to $\lambda_1(-\nabla f(\mathbf{0})) > 0$) and $-\nabla f(\mathbf{0})$ has an eigenvalue gap. In the full observation case, this assumption is satisfied easily if the ground-truth matrix has a unique top eigenvalue. The proof for Theorem 5.6 is deferred to Appendix I.1.

Assumption 5.5. $\mu_1 > \max\{\mu_2, 0\}$, where $\mu_1 := \lambda_1(-\nabla f(\mathbf{0}))$, $\mu_2 := \lambda_2(-\nabla f(\mathbf{0}))$.

Theorem 5.6. *Under Assumption 5.5, the following limit $\mathbf{W}_1^G(t)$ exists and is a solution of (2).*

$$\mathbf{W}_1^G(t) := \lim_{\epsilon \rightarrow 0} \mathbf{W}_{1,\epsilon}^G \left(\frac{1}{2\mu_1} \log \frac{1}{\epsilon} + t \right) = \lim_{\epsilon \rightarrow 0} \phi \left(\epsilon \mathbf{u}_1 \mathbf{u}_1^\top, \frac{1}{2\mu_1} \log \frac{1}{\epsilon} + t \right). \quad (9)$$

Let $\{\mathbf{W}_\alpha\} \subseteq \mathbb{S}_d^+$ be PSD matrices converging to $\mathbf{0}$ with positive alignment with $\mathbf{u}_1 \mathbf{u}_1^\top$ as $\alpha \rightarrow 0$, that is, $\lim_{\alpha \rightarrow 0} \mathbf{W}_\alpha = \mathbf{0}$ and $\exists \alpha_0, q > 0$ such that $\langle \mathbf{W}_\alpha, \mathbf{u}_1 \mathbf{u}_1^\top \rangle \geq q \|\mathbf{W}_\alpha\|_F$ for all $\alpha < \alpha_0$. Then $\forall t \in \mathbb{R}$, there is a constant $C > 0$ such that

$$\left\| \phi \left(\mathbf{W}_\alpha, \frac{1}{2\mu_1} \log \frac{1}{\langle \mathbf{W}_\alpha, \mathbf{u}_1 \mathbf{u}_1^\top \rangle} + t \right) - \mathbf{W}_1^G(t) \right\|_F \leq C \|\mathbf{W}_\alpha\|_F^{\frac{\tilde{\gamma}}{2\mu_1 + \tilde{\gamma}}} \quad (10)$$

for every sufficiently small α , where $\tilde{\gamma} := 2\mu_1 - (\mu_1 + \mu_2) = \mu_1 - \mu_2$.

It is worth to note that $\mathbf{W}_1^G(t)$ has rank ≤ 1 for any $t \in \mathbb{R}$, since every $\mathbf{W}_{1,\epsilon}^G(t)$ has rank ≤ 1 and the set $\mathbb{S}_{d,\leq 1}^+$ is closed. This matches with the first warmup example: GD does start learning with rank-1 solutions. Interestingly, in the case where the limit $\overline{\mathbf{W}}_1 := \lim_{t \rightarrow +\infty} \mathbf{W}_1^G(t)$ happens to be a minimizer of $f(\cdot)$ in \mathbb{S}_d^+ , GLRL should exit with the rank-1 solution $\overline{\mathbf{W}}_1$ after the first phase, and the following theorem shows that this is also the solution found by GF.

Assumption 5.7. $f(\mathbf{W})$ is locally analytic at each point.

Theorem 5.8. *Under Assumptions 5.5 and 5.7, if $\|\mathbf{W}_1^G(t)\|_F$ is bounded for all $t \geq 0$, then the limit $\overline{\mathbf{W}}_1 := \lim_{t \rightarrow +\infty} \mathbf{W}_1^G(t)$ exists. Further, if $\overline{\mathbf{W}}_1$ is a minimizer of $f(\cdot)$ in \mathbb{S}_d^+ , then for PSD matrices $\{\mathbf{W}_\alpha\} \subseteq \mathbb{S}_d^+$ converging to $\mathbf{0}$ with positive alignment with $\mathbf{u}_1 \mathbf{u}_1^\top$ as $\alpha \rightarrow 0$, it holds that $\lim_{\alpha \rightarrow 0} \lim_{t \rightarrow +\infty} \phi(\mathbf{W}_\alpha, t) = \overline{\mathbf{W}}_1$.*

Assumption 5.7 is a natural assumption, since $f(\cdot)$ in most cases of matrix factorization is a quadratic or polynomial function (e.g., matrix sensing, matrix completion). In general, it is unlikely for a gradient-based optimization process to get stuck at saddle points (Lee et al., 2017; Panageas et al., 2019). Thus, we should expect to see in general that GLRL finds the rank-1 solution if the problem is feasible with rank-1 matrices. This means at least for this subclass of problems, the implicit regularization of GD is rather unrelated to norm minimization. Below is a concrete example:

Example 5.9 (Counter-example of Conjecture 1.1, Gunasekar et al. 2017). Theorem 5.8 enables us to construct counterexamples of the *implicit nuclear norm regularization* conjecture in (Gunasekar et al., 2017). The idea is to construct a loss $\mathcal{L} : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$ where every rank-1 stationary point of $\mathcal{L}(\mathbf{U})$ attains the global minimum but none of them is minimizing the nuclear norm. Below we give a concrete matrix completion problem that meets the above requirement. Let \mathbf{M} be a partially observed matrix to be recovered, where the entries in $\Omega = \{(1, 3), (1, 4), (2, 3), (3, 1), (3, 2), (4, 1)\}$ are observed and the others (marked with “?”) are unobserved. The optimization problem is defined formally by $\mathcal{L}(\mathbf{U}) = \frac{1}{2}f(\mathbf{U}\mathbf{U}^\top)$, $f(\mathbf{W}) = \frac{1}{2} \sum_{(i,j) \in \Omega} (W_{ij} - M_{ij})^2$.

$$\mathbf{M} = \begin{bmatrix} ? & ? & 1 & R \\ ? & ? & R & ? \\ 1 & R & ? & ? \\ R & ? & ? & ? \end{bmatrix}, \mathbf{M}_{\text{norm}} = \begin{bmatrix} R & 1 & 1 & R \\ 1 & R & R & 1 \\ 1 & R & R & 1 \\ R & 1 & 1 & R \end{bmatrix}, \mathbf{M}_{\text{rank}} = \begin{bmatrix} 1 & R & 1 & R \\ R & R^2 & R & R^2 \\ 1 & R & 1 & R \\ R & R^2 & R & R^2 \end{bmatrix}.$$

Here $R > 1$ is a large constant, e.g., $R = 100$. The minimum nuclear norm solution is the rank-2 matrix \mathbf{M}_{norm} , which has $\|\mathbf{M}_{\text{norm}}\|_* = 4R$ (which is 400 when $R = 100$). \mathbf{M}_{rank} is a rank-1 solution with much larger nuclear norm, $\|\mathbf{M}_{\text{rank}}\|_* = 2R^2 + 2$ (which is 20002 when $R = 100$). We can verify that $f(\cdot)$ satisfies Assumptions 5.5 and 5.7 and $\mathbf{W}_1^G(t)$ converges to the rank-1 solution \mathbf{M}_{rank} . Therefore, GF with infinitesimal initialization converges to \mathbf{M}_{rank} rather than \mathbf{M}_{norm} , which refutes the conjecture in (Gunasekar et al., 2017). See Appendix D for a formal statement.

5.3 EQUIVALENCE BETWEEN GD AND GLRL: GENERAL CASE

Theorem 5.6 shows that for any *fixed* time t , the trajectory of GLRL in the first phase approximates GF with infinitesimal initialization, i.e., $\mathbf{W}_1^G(t) = \lim_{\alpha \rightarrow 0} \widehat{\mathbf{W}}_\alpha(t)$, where $\widehat{\mathbf{W}}_\alpha(t) := \phi(\mathbf{W}_\alpha, \frac{1}{2\mu_1} \log((\mathbf{W}_\alpha, \mathbf{u}_1 \mathbf{u}_1^\top)^{-1}) + t)$. However, $\mathbf{W}_1^G(\infty) \neq \lim_{\alpha \rightarrow 0} \widehat{\mathbf{W}}_\alpha(\infty)$ does not hold in general, unless the prerequisite in Theorem 5.8 is satisfied, i.e., unless $\overline{\mathbf{W}}_1 = \mathbf{W}_1^G(\infty)$ is a minimizer of $f(\cdot)$ in \mathbb{S}_d^+ . This is because of the well-known result that GD converges to local minimizers (Lee et al., 2016; 2017). We adapt Theorem 2 of Lee et al. (2017) to the setting of GF (Theorem I.5) and obtain the following result (Theorem 5.10); see Appendix I.4 for the proof.

Theorem 5.10. *Let $f : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$ be a convex \mathcal{C}^2 -smooth function. (1). All stationary points of $\mathcal{L} : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$, $\mathcal{L}(\mathbf{U}) = \frac{1}{2}f(\mathbf{U}\mathbf{U}^\top)$ are either strict saddles or global minimizers; (2). For any random initialization, GF (1) converges to strict saddles of $\mathcal{L}(\mathbf{U})$ with probability 0.*

Therefore, for convex $f(\cdot)$ such as matrix sensing and completion, suppose $f(\cdot)$ has no rank-1 PSD minimizer, then no matter how small α is, $\widehat{\mathbf{W}}_\alpha(\infty)$ (if exists) is a minimizer of $f(\cdot)$ with a higher rank and thus away from the rank-1 matrix $\overline{\mathbf{W}}_1$. In other words, $\mathbf{W}_1^G(t)$ only describes the limiting trajectory of GF in the first phase, i.e., when GF goes from near $\mathbf{0}$ to near $\overline{\mathbf{W}}_1$. After a sufficiently long time (depending on α), GF escapes the critical point $\overline{\mathbf{W}}_1$, but this is not described by $\mathbf{W}_1^G(t)$.

To understand how GF escapes $\overline{\mathbf{W}}_1$, a priori, we need to know how GF approaches $\overline{\mathbf{W}}_1$. Using a similar argument for Theorem 5.3, Theorem 5.11 shows that generically GF only escapes in the direction of $\mathbf{v}_1 \mathbf{v}_1^\top$, where \mathbf{v}_1 is the (unique) top eigenvector of $-\nabla f(\overline{\mathbf{W}}_1)$, and thus the limiting trajectory exactly matches with that of GLRL in the second phase until GF gets close to another critical point $\overline{\mathbf{W}}_2 \in \mathbb{S}_{d, \leq 2}^+$. If $\overline{\mathbf{W}}_2$ is still not a minimizer of $f(\cdot)$ in \mathbb{S}_d^+ (but it is a local minimizer in $\mathbb{S}_{d, \leq 2}^+$ generically), then GF escapes $\overline{\mathbf{W}}_2$ and the above process repeats until $\overline{\mathbf{W}}_K$ is a minimizer in \mathbb{S}_d^+ for some K . Here by “generically” we hide some technical assumptions and we elaborate on them in Appendix J. See Figure 1 and Figure 2 for experimental verification of the equivalence between GD and GLRL. We end this section with the following characterization of GF:

Theorem 5.11 (Theorem I.2, informal). *Let $\overline{\mathbf{W}}$ be a critical point of (2) satisfying that $\overline{\mathbf{W}}$ is a local minimizer of $f(\cdot)$ in $\mathbb{S}_{d, \leq r}^+$ for some $r \geq 1$ but not a minimizer in \mathbb{S}_d^+ . Let $-\nabla f(\overline{\mathbf{W}}) =$*

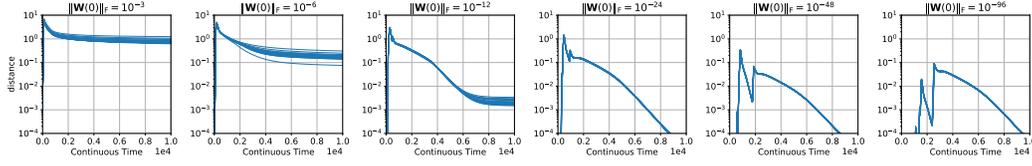


Figure 1: The trajectory of depth-2 GD, $\mathbf{W}_{\text{GD}}(t)$, converges to the trajectory of GLRL, $\mathbf{W}_{\text{GLRL}}(t)$, as the initialization scale goes to 0. We plot $\text{dist}(t) = \min_{t' \in \mathcal{T}} \|\mathbf{W}_{\text{GD}}(t) - \mathbf{W}_{\text{GLRL}}(t')\|_{\text{F}}$ for different initialization scale $\|\mathbf{W}(0)\|_{\text{F}}$, where \mathcal{T} is a discrete subset of \mathbb{R} that δ -covers the entire trajectory of GLRL: $\max_t \min_{t' \in \mathcal{T}} \|\mathbf{W}_{\text{GLRL}}(t) - \mathbf{W}_{\text{GLRL}}(t')\|_{\text{F}} \leq \delta$ for $\delta \approx 0.00042$. For each $\|\mathbf{W}(0)\|_{\text{F}}$, we run 20 random seeds and plot them separately. The ground truth $\mathbf{W}^* \in \mathbb{R}^{20 \times 20}$ is a randomly generated rank-3 matrix with $\|\mathbf{W}^*\|_{\text{F}} = 20$. 30% entries are observed. See more in Appendix E.1.

$\sum_{i=1}^d \mu_i \mathbf{v}_i \mathbf{v}_i^{\top}$ be the eigendecomposition of $-\nabla f(\overline{\mathbf{W}})$. If $\mu_1 > \mu_2$ and if there exists time $T_\alpha \in \mathbb{R}$ for every α so that $\phi(\mathbf{W}_\alpha, T_\alpha)$ converges to $\overline{\mathbf{W}}$ with positive alignment with the top principal component $\mathbf{v}_1 \mathbf{v}_1^{\top}$ as $\alpha \rightarrow 0$, then for every fixed t , $\lim_{\alpha \rightarrow 0} \phi(\mathbf{W}_\alpha, T_\alpha + \frac{1}{2\mu_1} \log \frac{1}{\langle \phi(\mathbf{W}_\alpha, T_\alpha), \mathbf{v}_1 \mathbf{v}_1^{\top} \rangle} + t)$ exists and is equal to $\mathbf{W}^{\text{G}}(t) := \lim_{\epsilon \rightarrow 0} \phi(\overline{\mathbf{W}} + \epsilon \mathbf{v}_1 \mathbf{v}_1^{\top}, \frac{1}{2\mu_1} \log \frac{1}{\epsilon} + t)$.

Characterization of the trajectory of GF. Generically, the trajectory of GF with small initialization can be split into K phases by $K + 1$ critical points of (2), $\{\overline{\mathbf{W}}_r\}_{r=0}^K$ ($\overline{\mathbf{W}}_0 = \mathbf{0}$), where in phase r GF escapes from $\overline{\mathbf{W}}_{r-1}$ in the direction of the top principal component of $-\nabla f(\overline{\mathbf{W}}_{r-1})$ and gets close to $\overline{\mathbf{W}}_r$. Each $\overline{\mathbf{W}}_r$ is a local minimizer of $f(\cdot)$ in $\mathbb{S}_{d, \leq r}^+$, but none of them is a minimizer of $f(\cdot)$ in \mathbb{S}_d^+ except $\overline{\mathbf{W}}_K$. The smaller the initialization is, the longer GF stays around each $\overline{\mathbf{W}}_r$. Moreover, $\{\overline{\mathbf{W}}_r\}_{r=0}^K$ corresponds to $\{\overline{\mathbf{W}}_{r, \epsilon}\}_{r=0}^K$ in Definition 5.1 with infinitesimal $\epsilon > 0$.

6 BENEFITS OF DEPTH: A VIEW FROM GLRL

In this section, we consider matrix factorization problems with depth $L \geq 3$. Our goal is to understand the effect of the depth- L parametrization $\mathbf{W} = \mathbf{U}_1 \mathbf{U}_2 \cdots \mathbf{U}_L$ on the implicit bias — how does depth encourage GF to find low rank solutions? We take the standard assumption in existing analysis for the end-to-end dynamics that the weight matrices have a balanced initialization, i.e. $\mathbf{U}_i^{\top}(0) \mathbf{U}_i(0) = \mathbf{U}_{i+1}(0) \mathbf{U}_{i+1}^{\top}(0)$, $\forall 1 \leq i \leq L - 1$. Arora et al. (2018) showed that if $\{\mathbf{U}_i\}_{i=1}^L$ is balanced at initialization, then we have the following end-to-end dynamics. Similar to the depth-2 case, we use $\phi(\mathbf{W}(0), t)$ to denote $\mathbf{W}(t)$, where

$$\frac{d\mathbf{W}}{dt} = - \sum_{i=0}^{L-1} (\mathbf{W} \mathbf{W}^{\top})^{\frac{i}{L}} \nabla f(\mathbf{W}) (\mathbf{W}^{\top} \mathbf{W})^{1 - \frac{i+1}{L}}. \quad (11)$$

The lemma below is the foundation of our analysis for the deep case, which greatly simplifies (11). Due to the space limit, we defer its derivations and applications into Appendix K.

Lemma 6.1. For $M(t) := \mathbf{W}(t)^{2/L}$, we have $\frac{dM}{dt} = -\nabla f(M^{L/2}) M^{L/2} - M^{L/2} \nabla f(M^{L/2})$.

Our main result, Theorem 6.2, gives a characterization of the limiting trajectory for deep matrix factorization with infinitesimal identity initialization. Here $\overline{\mathbf{W}}(t) := \lim_{\alpha \rightarrow 0} \mathbf{W}_\alpha^{\text{G}}(t)$ is the trajectory of deep GLRL, where $\mathbf{W}_\alpha^{\text{G}}(t) := \phi(\alpha \mathbf{e}_1 \mathbf{e}_1^{\top}, \frac{\alpha^{-(1-1/P)}}{2\mu_1(P-1)} + t)$ (see Algorithm 2). The dynamics for general initialization is more complicated. Please see discussions in Appendix L.

Theorem 6.2. Let $P = \frac{L}{2}$, $L \geq 3$. Suppose $\|\nabla f(\mathbf{0})\|_2 = \lambda_1(-\nabla f(\mathbf{0})) > \max\{\lambda_2(-\nabla f(\mathbf{0})), 0\}$,¹

$$\text{for every fixed } t \in \mathbb{R}, \quad \left\| \phi\left(\alpha \mathbf{I}, \frac{\alpha^{-(1-1/P)}}{2\mu_1(P-1)} + t\right) - \overline{\mathbf{W}}(t) \right\|_{\text{F}} = O(\alpha^{\frac{1}{P(P+1)}}), \quad (12)$$

and for any $2 \leq k \leq d$,

$$\text{for every fixed } t \in \mathbb{R}, \quad \lambda_k\left(\phi\left(\alpha \mathbf{I}, \frac{\alpha^{-(1-1/P)}}{2\mu_1(P-1)} + t\right)\right) = O(\alpha). \quad (13)$$

So how does depth encourage GF to find low-rank solutions? When the ground truth is low-rank, say rank- k , our experiments (Figure 2) suggest that GF with small initialization finds solutions with smaller k -low-rankness compared to the depth-2 case, thus achieving better generalization. At first glance, this is contradictory to what Theorem 6.2 suggests, i.e., the convergence rate of deep GLRL at a constant time gets slower as the depth increases. However, it turns out the uniform upper bound for the distance between GF and GLRL is not the ideal metric for the eventual k -low-rankness of learned solution. Below we will illustrate why the r -low-rankness of GF within each phase r is a better metric and how they are different.

Definition 6.3 (r -low-rankness). For matrix $M \in \mathbb{R}^{d \times d}$, we define the r -low-rankness of M as $\sqrt{\sum_{i=r+1}^d \sigma_i^2(M)}$, where $\sigma_i(M)$ is the i -th largest singular value of M .

¹We believe assumption $\|\nabla f(\mathbf{0})\|_2 = \lambda_1(-\nabla f(\mathbf{0}))$ could be removed with a more refined analysis.

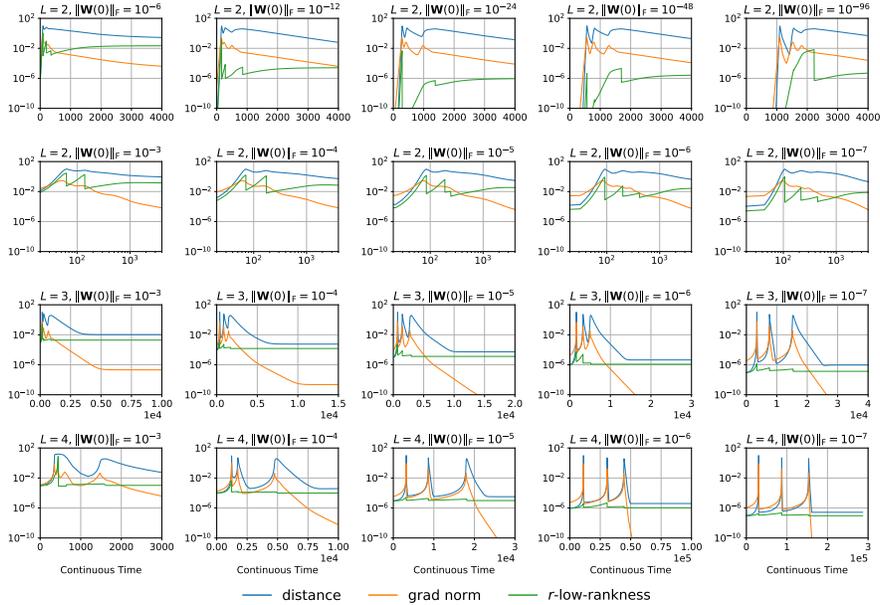


Figure 2: GD passes by the same set of critical points as GLRL when the initialization scale is small, and gets much closer to the critical points when $L \geq 3$. Depth-2 GD requires a much smaller initialization scale to maintain small low-rankness. Here the ground truth matrix $\mathbf{W}^* \in \mathbb{R}^{20 \times 20}$ is of rank 3 as stated in Appendix E.1. In this case, GLRL has 3 phases and 4 critical points $\{\overline{\mathbf{W}}_r\}_{r=0}^3$, where $\overline{\mathbf{W}}_0 = \mathbf{0}$ and $\overline{\mathbf{W}}_3 = \mathbf{W}^*$. For each depth L and initialization scale $\|\mathbf{W}(0)\|_F$, we plot the distance between the current step of GD and the closest critical point of GLRL, $\|\mathbf{W}_{\text{GD}}(t) - \overline{\mathbf{W}}_r\|_F$, the norm of full gradient, $\|\nabla_{U_{1:L}} \mathcal{L}(U_{1:L})\|_F$ and the $(r+1)$ -low-rankness of $\mathbf{W}_{\text{GD}}(t)$ with $r := \arg \min_{0 \leq i \leq 3} \|\mathbf{W}_{\text{GD}}(t) - \overline{\mathbf{W}}_i\|_F$.

Suppose $f(\cdot)$ admits a unique minimizer \mathbf{W}_0 in $\mathbb{S}_{d,1}^+$, and we run GF from $\alpha \mathbf{I}$ for both depth-2 and depth- L cases. Intuitively, the 1-low-rankness of the depth-2 solution is $\Omega(\alpha^{1-\mu_2/\mu_1})$, which can be seen from the second warmup example in Section 4. For the depth- L solution, though it may diverge from the trajectory of deep GLRL more than the depth-2 solution does, its 1-low-rankness is only $O(\alpha)$, as shown in Theorem 6.4. The key idea is to show that there is a basin in the manifold of rank-1 matrices around \mathbf{W}_0 such that any GF starting within the basin converges to \mathbf{W}_0 . Based on this, we can prove that starting from any matrix $O(\alpha)$ -close to the basin, GF converges to a solution $O(\alpha)$ -close to \mathbf{W}_0 . See Appendix M for more details.

Theorem 6.4. *In the same settings as Theorem 6.2, if $\overline{\mathbf{W}}(\infty)$ exists and is a minimizer of $f(\cdot)$ in $\mathbb{S}_{d,\leq 1}^+$, under regularity assumption M.1, we have $\inf_{t \in \mathbb{R}} \|\phi(\alpha \mathbf{I}, t) - \overline{\mathbf{W}}(\infty)\|_F = O(\alpha)$.*

Interpretation for the advantage of depth with multiple phases. For depth-2 GLRL, the low-rankness is raised to some power less than 1 per phase (depending on the eigengap). For deep GLRL, we show the low-rankness is only multiplied by some constant for the first phase and speculate it to be true for later phases. This conjecture is supported by our experiments; see Figure 2. Interestingly, our theory and experiments (Figure 5) suggest that while **being deep is good for generalization**, **being much deeper may not be much better**: once $L \geq 3$, increasing the depth does not improve the order of low-rankness significantly. While this theoretical result is only for identity initialization, Theorem F.1 and Corollary F.2 further show that the dynamics of GF (11) with any initialization pointwise converges as $L \rightarrow \infty$, under a suitable time rescaling. See Figure 6 for experimental verification.

7 CONCLUSION AND FUTURE DIRECTIONS

In this work, we connect gradient descent to Greedy Low-Rank Learning (GLRL) to explain the success of using gradient descent to find low-rank solutions in the matrix factorization problem. This enables us to construct counterexamples to the implicit nuclear norm conjecture in (Gunasekar et al., 2017). Taking the view of GLRL can also help us understand the benefits of depth.

ACKNOWLEDGMENTS

The authors thank Sanjeev Arora and Jason D. Lee for helpful discussions. The authors also thank Runzhe Wang for useful suggestions on writing. ZL and YL acknowledge support from NSF, ONR, Simons Foundation, Schmidt Foundation, Mozilla Research, Amazon Research, DARPA and SRC. ZL is also supported by Microsoft PhD Fellowship.

REFERENCES

- Akshay Agrawal, Robin Verschueren, Steven Diamond, and Stephen Boyd. A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1):42–60, 2018.
- Sanjeev Arora, N Cohen, and Elad Hazan. On the optimization of deep networks: Implicit acceleration by overparameterization. In *35th International Conference on Machine Learning*, 2018.
- Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo. Implicit regularization in deep matrix factorization. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 7411–7422. Curran Associates, Inc., 2019a.
- Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, Russ R Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8139–8148. Curran Associates, Inc., 2019b.
- Mohamed Ali Belabbas. On implicit regularization: Morse functions and applications to matrix factorization. *arXiv preprint arXiv:2001.04264*, 2020.
- Jeff Bezanson, Stefan Karpinski, Viral B Shah, and Alan Edelman. Julia: A fast dynamic language for technical computing. *arXiv preprint arXiv:1209.5145*, 2012.
- Yuejie Chi, Yue M Lu, and Yuxin Chen. Nonconvex optimization meets low-rank matrix factorization: An overview. *IEEE Transactions on Signal Processing*, 67(20):5239–5269, 2019.
- Lénaïc Chizat and Francis Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. volume 125 of *Proceedings of Machine Learning Research*, pages 1305–1338. PMLR, 09–12 Jul 2020.
- Lénaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 2937–2947. Curran Associates, Inc., 2019.
- Francis H. Clarke, Yuri S. Ledyev, Ronald J. Stern, and Peter R. Wolenski. *Nonsmooth analysis and control theory*, volume 178. Springer Science & Business Media, 2008.
- Frank H. Clarke. Generalized gradients and applications. *Transactions of the American Mathematical Society*, 205:247–262, 1975.
- Frank H Clarke. *Optimization and Nonsmooth Analysis*. Society for Industrial and Applied Mathematics, 1990. doi: 10.1137/1.9781611971309.
- Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.
- Simon Du and Jason Lee. On the power of over-parametrization in neural networks with quadratic activation. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1329–1338, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.
- Gauthier Gidel, Francis Bach, and Simon Lacoste-Julien. Implicit regularization of discrete gradient dynamics in linear neural networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 3196–3206. Curran Associates, Inc., 2019.
- Daniel Gissin, Shai Shalev-Shwartz, and Amit Daniely. The implicit bias of depth: How incremental learning drives generalization. In *International Conference on Learning Representations*, 2020.
- Suriya Gunasekar, Blake E Woodworth, Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Implicit regularization in matrix factorization. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6151–6159. Curran Associates, Inc., 2017.

- Suriya Gunasekar, Jason D Lee, Daniel Soudry, and Nati Srebro. Implicit bias of gradient descent on linear convolutional networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 9482–9491. Curran Associates, Inc., 2018.
- Benjamin D. Haeffele and René Vidal. Structured Low-Rank Matrix Factorization: Global Optimality, Algorithms, and Applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 42(6):1468–1482, 2019.
- Jean-Baptiste Hiriart-Urruty and A. S. Lewis. The clarke and michel-penot subdifferentials of the eigenvalues of a symmetric matrix. *Comput. Optim. Appl.*, 13(1-3):13–23, 1999. doi: 10.1023/A:1008644520093.
- Arthur Jacot, Franck Gabriel, and Clement Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 8571–8580. Curran Associates, Inc., 2018.
- Ziwei Ji and Matus Telgarsky. Gradient descent aligns the layers of deep linear networks. In *International Conference on Learning Representations*, 2019a.
- Ziwei Ji and Matus Telgarsky. A refined primal-dual analysis of the implicit bias. *arXiv preprint arXiv:1906.04540*, 2019b.
- Rajiv Khanna, Ethan Elenberg, Alexandros G Dimakis, and Sahand Negahban. On approximation guarantees for greedy low rank optimization. *arXiv preprint arXiv:1703.02721*, 2017.
- Jason D Lee, Max Simchowitz, Michael I Jordan, and Benjamin Recht. Gradient descent only converges to minimizers. In *Conference on learning theory*, pages 1246–1257, 2016.
- Jason D Lee, Ioannis Panageas, Georgios Piliouras, Max Simchowitz, Michael I Jordan, and Benjamin Recht. First-order methods almost always avoid saddle points. *arXiv preprint arXiv:1710.07406*, 2017.
- Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet, editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 2–47. PMLR, 06–09 Jul 2018.
- Kaifeng Lyu and Jian Li. Gradient descent maximizes the margin of homogeneous neural networks. In *International Conference on Learning Representations*, 2020.
- Mor Shpigel Nacson, Suriya Gunasekar, Jason Lee, Nathan Srebro, and Daniel Soudry. Lexicographic and depth-sensitive margins in homogeneous and non-homogeneous deep models. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 4683–4692, Long Beach, California, USA, 09–15 Jun 2019a. PMLR.
- Mor Shpigel Nacson, Jason Lee, Suriya Gunasekar, Pedro Henrique Pamplona Savarese, Nathan Srebro, and Daniel Soudry. Convergence of gradient descent on separable data. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 3420–3428. PMLR, 16–18 Apr 2019b.
- Mor Shpigel Nacson, Nathan Srebro, and Daniel Soudry. Stochastic gradient descent on separable data: Exact convergence with a fixed learning rate. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 3051–3059. PMLR, 16–18 Apr 2019c.
- Ioannis Panageas, Georgios Piliouras, and Xiao Wang. First-order methods almost always avoid saddle points: The case of vanishing step-sizes. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 6474–6483. Curran Associates, Inc., 2019.

- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- Lawrence Perko. *Differential equations and dynamical systems*, volume 7. Springer Science & Business Media, 2013.
- Noam Razin and Nadav Cohen. Implicit regularization in deep learning may not be explainable by norms. *arXiv preprint arXiv:2005.06398*, 2020.
- Shai Shalev-Shwartz and Yoram Singer. On the equivalence of weak learnability and linear separability: New relaxations and efficient boosting algorithms. *Machine learning*, 80(2-3):141–163, 2010.
- Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *Journal of Machine Learning Research*, 19(70):1–57, 2018a.
- Daniel Soudry, Elad Hoffer, and Nathan Srebro. The implicit bias of gradient descent on separable data. In *International Conference on Learning Representations*, 2018b.
- Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2):26–31, 2012.
- Zheng Wang, Ming-Jun Lai, Zhaosong Lu, Wei Fan, Hasan Davulcu, and Jieping Ye. Rank-one matrix pursuit for matrix completion. In *International Conference on Machine Learning*, pages 91–99, 2014.
- Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, and Benjamin Recht. The marginal value of adaptive gradient methods in machine learning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4148–4158. Curran Associates, Inc., 2017.
- Quanming Yao and James Tin Yau Kwok. Greedy learning of generalized low-rank models. In *IJCAI International Joint Conference on Artificial Intelligence*, 2016.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*, 2017.
- Stanisław Łojasiewicz. Ensembles semi-analytiques. *IHES notes*, 1965.

A COMPARISON TO EXISTING GREEDY ALGORITHMS FOR RANK-CONSTRAINED OPTIMIZATION

The most related one to GLRL (Algorithm 1) is probably *Rank-1 Matrix Pursuit* (RIMP) proposed by Wang et al. (2014) for matrix completion, which was later generalized to general convex loss in (Yao and Kwok, 2016). RIMP maintains a set of rank-1 matrices as the basis, and in phase r , RIMP adds the same $\mathbf{u}_r \mathbf{u}_r^\top$ as defined in Algorithm 1 into its basis and solve $\min_{\alpha} f(\sum_{i=1}^r \alpha_i \mathbf{u}_i \mathbf{u}_i^\top)$ for rank- r estimation. The main difference between RIMP and GLRL is that the optimization in each phase of RIMP is performed on the coefficients α , while the entire \mathbf{U}_r evolves with GD in each phase of GLRL. In Figure 3, we provide empirical evidence that GLRL generalizes better than RIMP when ground truth is low-rank, although GLRL may have a higher computational cost depending on η, ϵ .

Similar to RIMP, Greedy Efficient Component Optimization (GECO, Shalev-Shwartz and Singer 2010) also chooses the r -th component of its basis as the top eigenvector of $-\nabla f(\mathbf{W}_r)$, while it solves $\min_{\beta} f(\sum_{1 \leq i, j \leq r} \beta_{ij} \mathbf{u}_i \mathbf{u}_j^\top)$ for the rank- r estimation. Khanna et al. (2017) provided convergence guarantee for GECO assuming strong convexity. Haeffele and Vidal (2019) proposed a local-descent meta algorithm, of which GLRL can be viewed as a specific realization.

B DEEP GLRL ALGORITHM

Algorithm 2: Deep Greedy Low-Rank Learning (Deep GLRL)

parameter : step size $\eta > 0$; small $\epsilon > 0$
 $\epsilon' \leftarrow \epsilon^{1/L}$, $\mathcal{L}(\mathbf{U}_1, \dots, \mathbf{U}_L) := f(\mathbf{W}_1 \cdots \mathbf{W}_L)$.
 $\mathbf{W}_0 \leftarrow \mathbf{0} \in \mathbb{R}^{d \times d}$, and $\mathbf{U}_{0,1}(\infty), \dots, \mathbf{U}_{0,L}(\infty) \in \mathbb{R}^{d \times 0}$ are empty matrices
while $\lambda_1(-\nabla f(\mathbf{W}_r)) > 0$ **do**
 $r \leftarrow r + 1$
 let \mathbf{u}_r be a top (unit) eigenvector of $-\nabla f(\mathbf{W}_{r-1})$
 $\mathbf{U}_{r,1}(0) \leftarrow [\mathbf{U}_{r-1,1}(\infty) \quad \epsilon' \mathbf{u}_r] \in \mathbb{R}^{d \times r}$
 $\mathbf{U}_{r,k}(0) \leftarrow \begin{bmatrix} \mathbf{U}_{r-1,k}(\infty) & \mathbf{0} \\ \mathbf{0} & \epsilon' \end{bmatrix} \in \mathbb{R}^{r \times r}$ for all $2 \leq k \leq L-1$
 $\mathbf{U}_{r,L}(0) \leftarrow \begin{bmatrix} \mathbf{U}_{r-1,L}(\infty) \\ \epsilon' \mathbf{u}_r^\top \end{bmatrix} \in \mathbb{R}^{r \times d}$
 for $t = 0, 1, \dots$ **do**
 $\lfloor \mathbf{U}_{r,i}(t+1) \leftarrow \mathbf{U}_{r,i}(t) - \eta \nabla_{\mathbf{U}_i} \mathcal{L}(\mathbf{U}_{r,1}(t), \dots, \mathbf{U}_{r,L}(t)), \forall 1 \leq i \leq L.$
 $\mathbf{W}_r \leftarrow \mathbf{U}_{r,1}(\infty) \cdots \mathbf{U}_{r,L}(\infty)$
return \mathbf{W}_r

C PRELIMINARY LEMMAS

Lemma C.1. For $\mathbf{U}_0 \in \mathbb{R}^{d \times r}$ and $\mathbf{W}_0 := \mathbf{U}_0 \mathbf{U}_0^\top$, the following statements are equivalent:

- (1). \mathbf{U}_0 is a stationary point of $\mathcal{L}(\mathbf{U}) = \frac{1}{2} f(\mathbf{U} \mathbf{U}^\top)$;
- (2). $\nabla f(\mathbf{W}_0) \mathbf{W}_0 = \mathbf{0}$;
- (3). $\mathbf{W}_0 := \mathbf{U}_0 \mathbf{U}_0^\top$ is a critical point of (2).

Proof. (2) \Rightarrow (3) is trivial. We only prove (1) \Rightarrow (2), (3) \Rightarrow (1).

Proof for (1) \Rightarrow (2). If \mathbf{U}_0 is a stationary point, then $\mathbf{0} = \nabla \mathcal{L}(\mathbf{U}_0) = \nabla f(\mathbf{W}_0) \mathbf{U}_0$. So

$$\nabla f(\mathbf{W}_0) \mathbf{W}_0 = (\nabla f(\mathbf{W}_0) \mathbf{U}_0) \mathbf{U}_0^\top = \mathbf{0}.$$

Proof for (3) \Rightarrow (1). If \mathbf{W}_0 is a critical point, then

$$0 = \langle g(\mathbf{W}_0), \nabla f(\mathbf{W}_0) \rangle = -2 \operatorname{Tr}(\nabla f(\mathbf{W}_0) \mathbf{W}_0 \nabla f(\mathbf{W}_0)) = -2 \|\nabla f(\mathbf{W}_0) \mathbf{U}_0\|_{\mathbb{F}}^2,$$

which implies $\nabla \mathcal{L}(\mathbf{U}_0) = \mathbf{0}$. □

Lemma C.2. For a stationary point $\mathbf{U}_0 \in \mathbb{R}^{d \times r}$ of $\mathcal{L}(\mathbf{U}) = \frac{1}{2}f(\mathbf{U}\mathbf{U}^\top)$ where $f(\cdot)$ is convex, $\mathbf{W}_0 := \mathbf{U}_0\mathbf{U}_0^\top$ attains the global minimum of $f(\cdot)$ in $\mathbb{S}_d^+ := \{\mathbf{W} : \mathbf{W} \succeq \mathbf{0}\}$ iff $\nabla f(\mathbf{W}_0) \succeq \mathbf{0}$.

Proof. Since $f(\mathbf{W})$ is a convex function and \mathbb{S}_d^+ is convex, we know that \mathbf{W}_0 is a global minimizer of $f(\mathbf{W})$ in \mathbb{S}_d^+ iff

$$\langle \nabla f(\mathbf{W}_0), \mathbf{W} - \mathbf{W}_0 \rangle \geq 0, \quad \forall \mathbf{W} \succeq \mathbf{0}. \quad (14)$$

Note that $\langle \nabla f(\mathbf{W}_0), \mathbf{W}_0 \rangle = \text{Tr}(\nabla f(\mathbf{W}_0)\mathbf{W}_0)$. By Lemma C.1, $\langle \nabla f(\mathbf{W}_0), \mathbf{W}_0 \rangle = 0$. Combining this with (14), we know that \mathbf{W}_0 is a global minimizer iff

$$\langle \nabla f(\mathbf{W}_0), \mathbf{W} \rangle \geq 0, \quad \forall \mathbf{W} \succeq \mathbf{0}. \quad (15)$$

It is easy to check that this condition is equivalent to $\nabla f(\mathbf{W}_0) \succeq \mathbf{0}$. \square

D PROOFS FOR COUNTER-EXAMPLE

Conjecture D.1 (Formal Statement, [Gunasekar et al. 2017](#)). Suppose $f : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$ is a quadratic function and $\min_{\mathbf{W} \succeq \mathbf{0}} f(\mathbf{W}) = 0$. Then for any $\mathbf{W}_{\text{init}} \succ \mathbf{0}$ if $\bar{\mathbf{W}}_1 = \lim_{\alpha \rightarrow 0} \lim_{t \rightarrow +\infty} \phi(\alpha \mathbf{W}_{\text{init}}, t)$ exists and $f(\bar{\mathbf{W}}_1) = 0$, then $\|\bar{\mathbf{W}}_1\|_* = \min_{\mathbf{W} \succeq \mathbf{0}} \|\mathbf{W}\|_*$ s.t. $f(\mathbf{W}) = 0$.

Proposition D.2 (Formal Statement for Example 5.9). For constant $R > 1$, let

$$\mathbf{M} = \begin{bmatrix} ? & ? & 1 & R \\ ? & ? & R & ? \\ 1 & R & ? & ? \\ R & ? & ? & ? \end{bmatrix}, \mathbf{M}_{\text{norm}} = \begin{bmatrix} R & 1 & 1 & R \\ 1 & R & R & 1 \\ 1 & R & R & 1 \\ R & 1 & 1 & R \end{bmatrix}, \text{ and } \mathbf{M}_{\text{rank}} = \begin{bmatrix} 1 & R & 1 & R \\ R & R^2 & R & R^2 \\ 1 & R & 1 & R \\ R & R^2 & R & R^2 \end{bmatrix}.$$

and

$$\mathcal{L}(\mathbf{U}) = \frac{1}{2}f(\mathbf{U}\mathbf{U}^\top), \quad f(\mathbf{W}) = \frac{1}{2} \sum_{(i,j) \in \Omega} (W_{ij} - M_{ij})^2$$

where $\Omega = \{(1, 3), (1, 4), (2, 3), (3, 1), (3, 2), (4, 1)\}$.

Then for any $\mathbf{W}_{\text{init}} \succeq \mathbf{0}$, s.t. $\mathbf{u}_1^\top \mathbf{W}_{\text{init}} \mathbf{u}_1 > 0$,

$$\lim_{\alpha \rightarrow 0} \lim_{t \rightarrow +\infty} \phi(\alpha \mathbf{W}_{\text{init}}, t) = \mathbf{M}_{\text{rank}}.$$

Moreover, we have

$$\|\mathbf{M}_{\text{rank}}\|_* = 2R^2 + 2 > 4R = \|\mathbf{M}_{\text{norm}}\|_* = \min_{\mathbf{W} \succeq \mathbf{0}, f(\mathbf{W})=0} \|\mathbf{W}\|_*.$$

Proof. We define $\mathbf{W}_{1,\epsilon}^G(t)$, $\mathbf{W}_1^G(t)$ in the same way as in Definition 5.1, Theorem 5.6.

$$\begin{aligned} \mathbf{W}_{1,\epsilon}^G(t) &:= \phi(\epsilon \mathbf{u}_1 \mathbf{u}_1^\top, t), \\ \mathbf{W}_1^G(t) &:= \lim_{\epsilon \rightarrow 0} \mathbf{W}_{1,\epsilon}^G\left(\frac{1}{2\mu_1} \log \frac{1}{\epsilon} + t\right). \end{aligned}$$

Below we will show

1. Assumption 5.7 and Assumption 5.5 are satisfied.
2. $\|\mathbf{W}_1^G(t)\|_F$ bounded for $t \geq 0$;
3. $\lim_{t \rightarrow +\infty} \mathbf{W}_1^G(t) = \mathbf{M}_{\text{rank}}$;
4. $\mathbf{M}_{\text{norm}} = \arg \min_{\mathbf{W} \succeq \mathbf{0}, f(\mathbf{W})=0} \|\mathbf{W}\|_*$.

Thus Since M_{rank} is a global minimizer of $f(\cdot)$, applying Theorem 5.8 finishes the proof.

Proof for Item 1. Let $M_0 := \nabla f(\mathbf{0})$, then

$$M_0 = \begin{bmatrix} 0 & 0 & 1 & R \\ 0 & 0 & R & 0 \\ 1 & R & 0 & 0 \\ R & 0 & 0 & 0 \end{bmatrix}.$$

Let $\mathbf{A} := [\frac{1}{R} \ R; 0 \ 0]$, then we have $\lambda_1(\mathbf{A}) = \frac{1+\sqrt{1+R^2}}{2}$, $\lambda_2(\mathbf{A}) = \frac{1-\sqrt{1+R^2}}{2}$, thus $\lambda_1(\mathbf{A}) > |\lambda_2(\mathbf{A})| > 0 > \lambda_2(\mathbf{A})$. As a result, $\lambda_1(\mathbf{A}) = \|\mathbf{A}\|_2$. Let $\mathbf{v}_1 \in \mathbb{R}^2$ be the top eigenvector of \mathbf{A} . We claim that $\mathbf{u}_1 = [\mathbf{v}_1^T]^T \in \mathbb{R}^4$ is the top eigenvector of $\nabla f(\mathbf{0})$. First by definition it is easy to check that $M_0 \mathbf{u}_1 = \lambda_1(\mathbf{A}) \mathbf{u}_1$. Further noticing that $M_0^2 = \begin{bmatrix} \mathbf{A}^2 & \mathbf{0} \\ \mathbf{0} & \mathbf{A}^2 \end{bmatrix}$, we know $\lambda_i^2(M_0) \in \{\lambda_1^2(\mathbf{A}), \lambda_2^2(\mathbf{A})\}$ for all eigenvalues $\lambda_i(M_0)$. That is, $\lambda_1(M_0) = \lambda_1(\mathbf{A})$, $\lambda_2(M_0) = -\lambda_2(\mathbf{A})$, $\lambda_3(M_0) = \lambda_2(\mathbf{A})$, and $\lambda_4(M_0) = -\lambda_1(\mathbf{A})$. Thus Assumption 5.5 is satisfied. Also note that f is quadratic, thus analytic, i.e., Assumption 5.7 is also satisfied.

Proof for Item 2. Let $(x_\epsilon(t), y_\epsilon(t)) \in \mathbb{R}^2$ be the gradient flow of $g(x, y) = \frac{1}{2}(x^2 - 1)^2 + (xy - R)^2$ starting from $(x_\epsilon(0), y_\epsilon(0)) = \sqrt{\epsilon} \mathbf{v}_1$.

$$\begin{aligned} \frac{dx(t)}{dt} &= (1 - x(t)^2)x(t) - 2y(t)(x(t)y(t) - R) \\ \frac{dy(t)}{dt} &= -2x(t)(x(t)y(t) - R) \end{aligned} \tag{16}$$

Let $\mathbf{W}_\epsilon(t)$ be the following matrix:

$$\mathbf{W}_\epsilon(t) := \begin{bmatrix} x_\epsilon(t) \\ y_\epsilon(t) \\ x_\epsilon(t) \\ y_\epsilon(t) \end{bmatrix} \begin{bmatrix} x_\epsilon(t) & y_\epsilon(t) & x_\epsilon(t) & y_\epsilon(t) \end{bmatrix}.$$

Then it is easy to verify that $\mathbf{W}_\epsilon(0) = \mathbf{W}_{1,\epsilon}^G(0)$ and $\mathbf{W}_\epsilon(t)$ satisfies (2). Thus by the existence and uniqueness theorem, we have $\mathbf{W}_\epsilon(t) = \mathbf{W}_{1,\epsilon}^G(t)$ for all t . Taking the limit $\epsilon \rightarrow 0$, we know that $\mathbf{W}_1^G(t)$ can also be written in the following form:

$$\mathbf{W}_1^G(t) = \begin{bmatrix} x(t) \\ y(t) \\ x(t) \\ y(t) \end{bmatrix} \begin{bmatrix} x(t) & y(t) & x(t) & y(t) \end{bmatrix},$$

and $(x_\epsilon(t), y_\epsilon(t)) \in \mathbb{R}^2$ is a gradient flow of $g(x, y) = \frac{1}{2}(x^2 - 1)^2 + (xy - R)^2$.

Since $g(x(t), y(t))$ is non-increasing overtime, and $\lim_{t \rightarrow -\infty} g(x(-t), y(-t)) = g(x(-\infty), y(-\infty)) = g(0, 0) = R^2 + 0.5$, we know $|x(t)y(t)| \leq 3R$ for all t . So whenever $y^2(t) - x^2(t) \geq 9R^2$, we have $x^2(t) \leq \frac{9R^2}{y^2(t)} \leq \frac{9R^2}{y^2(t) - x^2(t)} \leq 1$. In this case, $\frac{d(y^2(t) - x^2(t))}{dt} = 2x^2(t)(x^2(t) - 1) \leq 0$. Combining this with $y(-\infty)^2 - x(-\infty)^2 = 0 \leq 9R^2$, we have $y^2(t) - x^2(t) \leq 9R^2$ for all t , which also implies that $y(t)$ is bounded. Noticing that $9R^2 \geq g(x(t), y(t)) \geq (x^2(t) - 1)^2$, we know $x^2(t)$ is also bounded. Therefore, $\mathbf{W}_1^G(t)$ is bounded.

Proof for Item 3. Note that $(x(\infty), y(\infty))$ is a stationary point of $g(x, y)$. It is clear that $g(x, y)$ only has 3 stationary points — $(0, 0)$, $(1, R)$ and $(-1, -R)$. Thus $\overline{\mathbf{W}}_1$ can only be $\mathbf{0}$ or M_{rank} . However, since for all t , $f(\mathbf{W}_1^G(t)) < f(\mathbf{0})$, $\overline{\mathbf{W}}_1 = \lim_{t \rightarrow \infty} \mathbf{W}_1^G(t)$ cannot be $\mathbf{0}$. So $\overline{\mathbf{W}}_1$ must be M_{rank} .

Proof for Item 4. Let m_{ij} be (i, j) th element of M . Suppose $M \succeq \mathbf{0}$, we have

$$\begin{aligned} (\mathbf{e}_1 - \mathbf{e}_4)^T M (\mathbf{e}_1 - \mathbf{e}_4) \geq 0 &\implies m_{11} + m_{44} \geq m_{14} + m_{41} = 2R \\ (\mathbf{e}_2 - \mathbf{e}_3)^T M (\mathbf{e}_2 - \mathbf{e}_3) \geq 0 &\implies m_{22} + m_{33} \geq m_{23} + m_{32} = 2R \end{aligned}$$

Depth (L)	Simulation method
2	Constant LR, $\eta = 10^{-3}$ for 10^6 iterations
3	Adaptive LR, $\eta = 2 \times 10^{-5}$ and $\varepsilon = 10^{-4}$ for 10^6 iterations
4	Adaptive LR, $\eta = 3 \times 10^{-4}$ and $\varepsilon = 10^{-3}$ for 10^6 iterations

Table 1: Choice of hyperparameters for simulating gradient flow. For $L = 2$, gradient descent escapes saddles in $O(\log \frac{1}{\varepsilon})$ time, where ε is the distance between the initialization and the saddle.

Thus $4R = \min_{\mathbf{W} \succeq \mathbf{0}, f(\mathbf{W})=0} \|\mathbf{W}\|_*$, where the equality is only attained at $m_{ii} = R, i = 1, 2, 3, 4$. Otherwise, either $\begin{bmatrix} m_{11} & m_{14} \\ m_{41} & m_{44} \end{bmatrix}$ or $\begin{bmatrix} m_{22} & m_{23} \\ m_{32} & m_{33} \end{bmatrix}$ will have negative eigenvalues. Contradiction to that $\mathbf{M} \succeq \mathbf{0}$.

Below we will show the rest unknown off-diagonal entries must be 1. Let $\mathbf{V} = \begin{bmatrix} 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$, then

$$\mathbf{M} \succeq \mathbf{0} \implies \mathbf{V}\mathbf{M}\mathbf{V}^\top \succeq \mathbf{0} \implies \begin{bmatrix} 0 & m_{13} - m_{23} & m_{14} - m_{24} \\ m_{31} - m_{32} & R & R \\ m_{41} - m_{42} & R & R \end{bmatrix} \succeq \mathbf{0},$$

which implies $m_{13} = m_{23}, m_{14} = m_{24}$.

With the same argument for $\mathbf{V} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & -1 \end{bmatrix}$, we have $m_{13} = m_{14}, m_{23} = m_{24}$. Also note \mathbf{M} is symmetric and $m_{13} = 1$, thus $m_{ij} = m_{ji} = 1, \forall i = 1, 2, j = 3, 4$. Thus $\mathbf{M}_{\text{norm}} = \arg \min_{\mathbf{W} \succeq \mathbf{0}, f(\mathbf{W})=0} \|\mathbf{W}\|_*$, which is unique. \square

E EXPERIMENTS

E.1 GENERAL SETUP

The code is written in Julia (Bezanson et al., 2012) and PyTorch (Paszke et al., 2019).

The ground-truth matrix \mathbf{W}_* is low-rank by construction: we sample a random orthogonal matrix \mathbf{U} , a diagonal matrix \mathbf{S} with Frobenius norm $\|\mathbf{S}\|_F = 1$ and set $\mathbf{W}_* = \mathbf{U}\mathbf{S}\mathbf{U}^\top$. Each measurement \mathbf{X} in $\mathbf{X}_1, \dots, \mathbf{X}_m$ is generated by sampling two one-hot vectors \mathbf{u} and \mathbf{v} uniformly and setting $\mathbf{X} = \frac{1}{2}\mathbf{u}\mathbf{v}^\top + \frac{1}{2}\mathbf{v}\mathbf{u}^\top$.

In Figures 1, 2, 3 to 5 and 7, the ground truth matrix \mathbf{W}_* has shape 20×20 and rank 3, where $\|\mathbf{W}_*\|_F = 20, \lambda_1(\mathbf{W}_*) = 17.41, \lambda_2(\mathbf{W}_*) = 8.85, \lambda_3(\mathbf{W}_*) = 4.31$ and $\lambda_1(-\nabla f(\mathbf{0})) = 6.23, \lambda_2(-\nabla f(\mathbf{0})) = 5.41$. $p = 0.3$ is used for generating measurements, except $p = 0.25$ in Figure 3, i.e., each pair of entries of \mathbf{W}_{ij}^* and \mathbf{W}_{ji}^* is observed with probability p .

Gradient Descent. Let $\tilde{\varepsilon} > 0$ be the Frobenius norm of the target random initialization. For the depth-2 case, we sample 2 orthogonal matrices $\mathbf{V}_1, \mathbf{V}_2$ and a diagonal matrix \mathbf{D} with Frobenius norm $\tilde{\varepsilon}$, and we set $\mathbf{U} = \mathbf{V}_1\mathbf{D}^{1/2}\mathbf{V}_2^\top$; for the depth- L case with $L \geq 3$, we sample L orthogonal matrices $\mathbf{V}_1, \dots, \mathbf{V}_L$ and a diagonal matrix \mathbf{D} with Frobenius norm $\tilde{\varepsilon}$, and we set $\mathbf{U}_i := \mathbf{V}_i\mathbf{D}^{1/L}\mathbf{V}_{i+1}^\top$ ($\mathbf{V}_{L+1} = \mathbf{V}_1$). In this way, we can guarantee that the end-to-end matrix $\mathbf{W} = \mathbf{U}_1 \cdots \mathbf{U}_L$ is symmetric and the initialization is balanced for $L \geq 3$.

We discretize the time to simulate gradient flow. When $L > 2$, gradient flow stays around saddle points for most of the time, therefore we use full-batch GD with adaptive learning rate $\tilde{\eta}_t$, inspired by

RMSprop (Tieleman and Hinton, 2012), for faster convergence:

$$\begin{aligned} v_{t+1} &= \alpha v_t + (1 - \alpha) \|\nabla \mathcal{L}(\boldsymbol{\theta}_t)\|_2^2, \\ \tilde{\eta}_t &= \frac{\eta}{\sqrt{\frac{v_{t+1}}{1 - \alpha^{t+1}} + \epsilon}}, \\ \boldsymbol{\theta}_{t+1} &= \boldsymbol{\theta}_t - \tilde{\eta}_t \nabla \mathcal{L}(\boldsymbol{\theta}_t), \end{aligned}$$

where $\alpha = 0.99$, η is the (unadjusted) learning rate. The choices of hyperparameters are summarized in Table 1. The continuous time for $\boldsymbol{\theta}_t$ is measured as $\sum_{i=0}^{t-1} \tilde{\eta}_i$.

GLRL. In Figures 1, 2, 3 and 4, the GLRL’s trajectory is obtained by running Algorithm 1 with $\epsilon = 10^{-7}$ and $\eta = 10^{-3}$. The stopping criterion is that if the loop has been iterated for 10^7 times.

E.2 EXPERIMENTAL EQUIVALENCE BETWEEN GLRL AND GRADIENT DESCENT

Here we provide experimental evidence supporting our theoretical claims about the equivalence between GLRL and GF for both cases, $L = 2$ and $L \geq 3$.

In Figure 1, we show the distance from every point on GF (simulated by GD) from random initialization is close to the trajectory of GLRL. In Figure 2, we first run GLRL and obtain the critical points $\{\overline{\mathbf{W}}_r\}_{r=0}^3$ passed by GLRL. We also define the distance of a matrix \mathbf{W} to the critical points to be $\min_{0 \leq r \leq 3} \|\mathbf{W} - \overline{\mathbf{W}}_r\|_F$.

E.3 HOW WELL DOES GLRL WORK?

We compare GLRL with gradient descent (with not-so-small initialization), nuclear norm minimization and RIMP (Wang et al., 2014). We use CVXPY (Diamond and Boyd, 2016; Agrawal et al., 2018) for finding the nuclear norm solution. The results are shown in Figure 3. GLRL can fully recover the ground truth, while others have difficulty doing so.

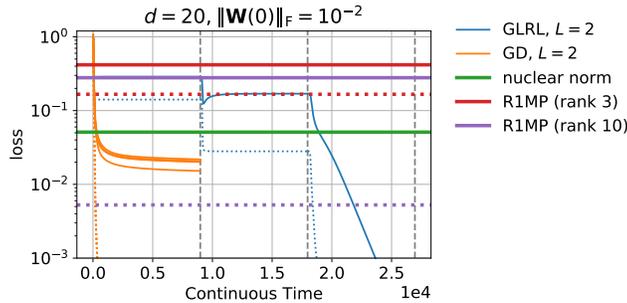


Figure 3: GD with small initialization outperforms RIMP and minimal nuclear norm solution on synthetic data with low-rank ground truth. Solid (dotted) curves correspond to test (training) loss. Here the loss $f(\mathbf{W}) := \frac{1}{d^2} \|\mathbf{W} - \mathbf{W}^*\|_F^2$ and $f(\mathbf{0}) = 1$. We run 10 random seeds for GD and plot them separately (most of them overlap).

E.4 HOW DOES INITIALIZATION AFFECT THE CONVERGENCE RATE TO THE RANK-1 GLRL TRAJECTORY?

We use the general setting in Appendix E.1. In these experiments, we use the constant learning rate 10^{-5} for 4×10^7 iterations. The reference matrix \mathbf{W}_{ref} is obtained by running the first stage of GLRL with $\|\mathbf{W}(0)\|_F = 10^{-48}$ and we pick one matrix in the trajectory with $\|\mathbf{W}_{\text{ref}}\|_F$ about 0.6.

For every $\epsilon = 10^i$, $i \in \{-1, -2, -3, -4, -5\}$, we run both gradient descent and the first phase of GLRL with $\|\mathbf{W}(0)\|_F = \epsilon$. For gradient descent, we use random initialization so $\|\mathbf{W}(0)\|_F$ is full rank w.p. 1. The distance of a trajectory to \mathbf{W}_{ref} is defined as $\min_{t \geq 0} \|\mathbf{W}(t) - \mathbf{W}_{\text{ref}}\|_F$. In practice,

as we discretized time to simulate gradient flow, we check every t during simulation to compute the distance. As a result, the estimation might be inaccurate when a trajectory is really close to \mathbf{W}_{ref} .

The result is shown at Figure 4. We observe that GLRL trajectories are closer to the reference matrix \mathbf{W}_{ref} by magnitudes. Thus the take home message here is that GLRL is in general a more computational efficient method to simulate the trajectory of GF (GD) with infinitesimal initialization, as one can start GLRL with a much larger initialization, while still maintaining high precision.

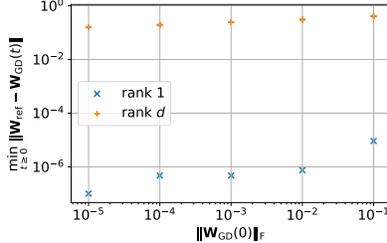


Figure 4: Using $\epsilon v_1 v_1^\top$ (denoted by “rank 1”) as initialization makes GD much closer to GLRL compared to using random initialization (denoted by “rank d ”), where v_1 is the top eigenvector of $-\nabla f(\mathbf{0})$. We take a fixed reference matrix on the trajectory of GLRL with constant norm and plot the distance of GD with each initialization to it respectively..

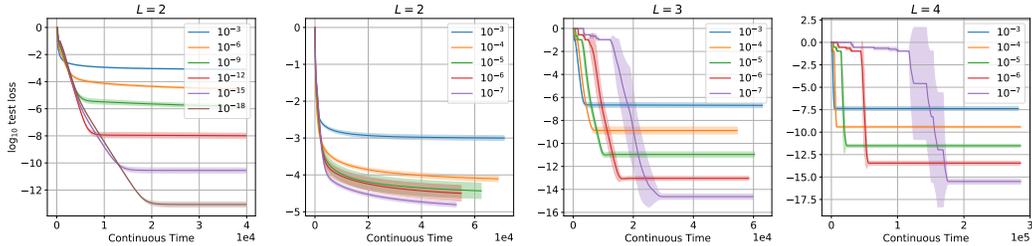


Figure 5: Deep matrix factorization encourages GF to find low rank solutions at a much practical initialization scale, e.g. 10^{-3} . Here the ground truth is rank-3. For each setting, we run 5 different random seeds. The solid curves are the mean and the shaded area indicates one standard deviation. We observe that performance of GD is quite robust to its initialization. Note that for $L > 2$, the shaded area with initialization scale 10^{-7} is large, as the sudden decrement of loss occurs at quite different continuous times for different random seeds in this case.

E.5 BENEFIT OF DEPTH: POLYNOMIAL VS EXPONENTIAL DEPENDENCE ON INITIALIZATION

To verify our theory in Section 6, we run gradient descent with different depth and initialization. The results are shown in Figure 5. We can see that as the initialization becomes smaller, the final solution gets closer to the ground truth. However, a depth-2 model requires exponentially small initialization, while deeper models require polynomial small initialization, though it takes much longer to converge.

F THE MARGINAL VALUE OF BEING DEEPER

Theorem F.1 shows that the end-to-end dynamics (17) converges point-wise while $L \rightarrow \infty$ if the product of learning rate and depth, ηL , is fixed as constant. Interestingly, (17) also allows us to simulate the dynamics of $\mathbf{W}(t)$ for all depths L while the computation time is independent of L . In Figure 6, we compare the effect of depth while fixing the initialization and ηL . We can see that deeper models converge faster. The difference between $L = 1, 2$, and 4 is large, while difference among $L \geq 16$ is marginal.

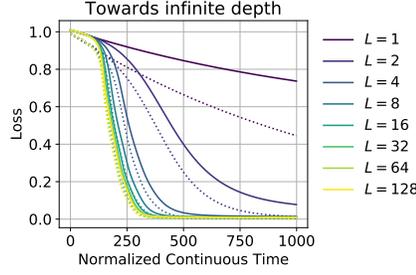


Figure 6: The marginal value of being deeper. The trajectory of GD converges when depth goes to infinity. Solid (dotted) curves correspond to test (train) loss. The x -axis stands for the normalized continuous time t (multiplied by L).

Theorem F.1. Suppose $\mathbf{W} = \tilde{\mathbf{U}}\tilde{\Sigma}\tilde{\mathbf{V}}^\top$ is the SVD decomposition of \mathbf{W} , where $\tilde{\Sigma} = \text{diag}(\sigma_1, \dots, \sigma_d)$. The dynamics of L -layer linear net is the following, \circ denotes the entry-wise multiplication:

$$\frac{d\mathbf{W}}{dt} = -L\tilde{\mathbf{U}} \left((\tilde{\mathbf{U}}^\top \nabla f(\mathbf{W}) \tilde{\mathbf{V}}) \circ \mathbf{K}^{(L)} \right) \tilde{\mathbf{V}}^\top, \quad (17)$$

where $K_{i,i}^{(L)} = \sigma_i^{2-2/L}$, $K_{i,j}^{(L)} = \frac{\sigma_i^2 - \sigma_j^2}{L\sigma_i^{2/L} - L\sigma_j^{2/L}}$ for $i \neq j$.

Proof. We start from (11):

$$\begin{aligned} \frac{d\mathbf{W}}{dt} &= - \sum_{l=0}^{L-1} (\mathbf{W}\mathbf{W}^\top)^{\frac{l}{L}} \nabla f(\mathbf{W}) (\mathbf{W}^\top \mathbf{W})^{\frac{L-1-l}{L}} \\ &= - \sum_{l=0}^{L-1} \tilde{\mathbf{U}} \tilde{\Sigma}^{\frac{2l}{L}} \tilde{\mathbf{U}}^\top \nabla f(\mathbf{W}) \tilde{\mathbf{V}} \tilde{\Sigma}^{\frac{2(L-1-l)}{L}} \tilde{\mathbf{V}} \\ &= -L\tilde{\mathbf{U}} \left[L^{-1} \sum_{l=0}^{L-1} \tilde{\Sigma}^{\frac{2l}{L}} (\tilde{\mathbf{U}}^\top \nabla f(\mathbf{W}) \tilde{\mathbf{V}}) \tilde{\Sigma}^{\frac{2(L-1-l)}{L}} \right] \tilde{\mathbf{V}}. \end{aligned}$$

Note that $\tilde{\Sigma}$ is diagonal, so

$$\tilde{\Sigma}^{\frac{2l}{L}} (\tilde{\mathbf{U}}^\top \nabla f(\mathbf{W}) \tilde{\mathbf{V}}) \tilde{\Sigma}^{\frac{2(L-1-l)}{L}} = (\tilde{\mathbf{U}}^\top \nabla f(\mathbf{W}) \tilde{\mathbf{V}}) \circ \mathbf{H}^{(l)},$$

where $\mathbf{H}_{i,j}^{(l)} = \sigma_i^{\frac{2l}{L}} \sigma_j^{\frac{2(L-1-l)}{L}}$. Therefore,

$$\begin{aligned} L^{-1} \sum_{l=0}^{L-1} \tilde{\Sigma}^{\frac{2l}{L}} (\tilde{\mathbf{U}}^\top \nabla f(\mathbf{W}) \tilde{\mathbf{V}}) \tilde{\Sigma}^{\frac{2(L-1-l)}{L}} &= L^{-1} \sum_{l=0}^{L-1} (\tilde{\mathbf{U}}^\top \nabla f(\mathbf{W}) \tilde{\mathbf{V}}) \circ \mathbf{H}^{(l)} \\ &= (\tilde{\mathbf{U}}^\top \nabla f(\mathbf{W}) \tilde{\mathbf{V}}) \circ \mathbf{K}^{(L)}, \end{aligned}$$

where $\mathbf{K}^{(L)} = L^{-1} \sum_{l=0}^{L-1} \mathbf{H}^{(l)}$. Hence,

$$\frac{d\mathbf{W}}{dt} = -L\tilde{\mathbf{U}} \left[(\tilde{\mathbf{U}}^\top \nabla f(\mathbf{W}) \tilde{\mathbf{V}}) \circ \mathbf{K}^{(L)} \right] \tilde{\mathbf{V}}.$$

The entries of $\mathbf{K}^{(L)}$ can be directly calculated by

$$K_{i,j}^{(L)} = L^{-1} \sum_{l=0}^{L-1} \sigma_i^{\frac{2l}{L}} \sigma_j^{\frac{2(L-1-l)}{L}} = \begin{cases} \sigma_i^{2-2/L}, & i = j, \\ \frac{\sigma_i^2 - \sigma_j^2}{L\sigma_i^{2/L} - L\sigma_j^{2/L}}, & i \neq j. \end{cases}$$

□

Corollary F.2. As $L \rightarrow \infty$, $\mathbf{K}^{(L)}$ converges to \mathbf{K}^* , where $K_{i,i}^* = \sigma_i^2$, $K_{i,j}^* = \frac{\sigma_i^2 - \sigma_j^2}{\ln \sigma_i^2 - \ln \sigma_j^2}$ for $i \neq j$.

Experiment details. We follow the general setting in Appendix E.1. The ground truth \mathbf{W}^* is different but is generated in the same manner and has the same shape of 20×20 and $p = 0.3$ is used for observation generation. We directly apply (17), in which we compute $\tilde{\mathbf{V}}$ and $\tilde{\mathbf{U}}$ through SVD, to simulate the trajectory together with a constant learning rate of $\frac{10^{-3}}{L}$ for depth L . $\mathbf{W}(0)$ is sampled from $10^{-3} \times \mathcal{N}(0, \mathbf{I}_d)$.

G PROOFS FOR DYNAMICAL SYSTEM

In this section, we prove Theorem 5.3 in Section 5.1. In Appendix G.1, we show how to reduce Theorem 5.3 to the case where $\mathbf{J}(0)$ is exactly a diagonal matrix, then we prove this diagonal case in Appendix G.2. Finally, in Appendix G.3, we discuss how to extend it to the case where $\mathbf{J}(0)$ is non-diagonalizable.

G.1 REDUCTION TO THE DIAGONAL CASE

Theorem G.1. *If $\mathbf{J}(0) = \text{diag}(\tilde{\mu}_1, \dots, \tilde{\mu}_d)$ is diagonal, then the statement in Theorem 5.3 holds.*

Proof for Theorem 5.3. We show how to prove Theorem 5.3 based on Theorem G.1. Let $\frac{d\boldsymbol{\theta}}{dt} = \mathbf{g}(\boldsymbol{\theta})$ be the dynamical system in Theorem 5.3. Let $\mathbf{J}(0) = \tilde{\mathbf{V}}\tilde{\mathbf{D}}\tilde{\mathbf{V}}^{-1}$ be the eigendecomposition, where $\tilde{\mathbf{V}}$ is an invertible matrix and $\tilde{\mathbf{D}} = \text{diag}(\tilde{\mu}_1, \dots, \tilde{\mu}_d)$. Now we define the following new dynamics by changing the basis:

$$\hat{\boldsymbol{\theta}}(t) = \tilde{\mathbf{V}}^{-1}\boldsymbol{\theta}(t).$$

Then $\frac{d\hat{\boldsymbol{\theta}}(t)}{dt} = \hat{\mathbf{g}}(\hat{\boldsymbol{\theta}})$ for $\hat{\mathbf{g}}(\hat{\boldsymbol{\theta}}) := \tilde{\mathbf{V}}^{-1}\mathbf{g}(\tilde{\mathbf{V}}\hat{\boldsymbol{\theta}})$, and the associated Jacobian matrix is $\hat{\mathbf{J}}(\hat{\boldsymbol{\theta}}) := \tilde{\mathbf{V}}^{-1}\mathbf{J}(\tilde{\mathbf{V}}\hat{\boldsymbol{\theta}})\tilde{\mathbf{V}}$, and thus $\hat{\mathbf{J}}(0) = \text{diag}(\tilde{\mu}_1, \dots, \tilde{\mu}_d)$.

Now we apply Theorem G.1 to $\hat{\boldsymbol{\theta}}(t)$. Then $\hat{\mathbf{z}}_\alpha(t) := \tilde{\mathbf{V}}^{-1}\mathbf{z}_\alpha(t)$ converges to the limit $\hat{\mathbf{z}}(t) := \lim_{\alpha \rightarrow 0} \hat{\mathbf{z}}_\alpha(t)$. This shows that the limit $\mathbf{z}(t) = \tilde{\mathbf{V}}\hat{\mathbf{z}}(t)$ exists in Theorem 5.3. We can also verify that $\mathbf{z}(t)$ is a solution of (6).

Given $\boldsymbol{\delta}_\alpha$ converging to $\mathbf{0}$ with positive alignment with $\tilde{\mathbf{u}}_1$ as $\alpha \rightarrow 0$, we can define $\hat{\boldsymbol{\delta}}_\alpha := \tilde{\mathbf{V}}^{-1}\boldsymbol{\delta}_\alpha$, then $\hat{\boldsymbol{\delta}}_\alpha$ converges to $\mathbf{0}$ with positive alignment with \mathbf{e}_1 , where \mathbf{e}_1 is the first vector in the standard basis and is also the top eigenvector of $\hat{\mathbf{J}}(0)$. Therefore, for every $t \in (-\infty, +\infty)$, there is a constant $C > 0$ such that

$$\left\| \tilde{\mathbf{V}}^{-1}\phi\left(\boldsymbol{\delta}_\alpha, t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\langle \boldsymbol{\delta}_\alpha, \tilde{\mathbf{u}}_1 \rangle}\right) - \hat{\mathbf{z}}(t) \right\|_2 \leq C \cdot \|\hat{\boldsymbol{\delta}}_\alpha\|_2^{\frac{\tilde{\gamma}}{\tilde{\mu}_1 + \tilde{\gamma}}} \quad (18)$$

for every sufficiently small α . As $\tilde{\mathbf{V}}$ are invertible, this directly implies (7). \square

G.2 PROOF FOR THE DIAGONAL CASE

Now we only need to prove Theorem G.1. Let $\mathbf{e}_1, \dots, \mathbf{e}_d$ be the standard basis. Then $\tilde{\mathbf{u}}_1 = \tilde{\mathbf{v}}_1 = \mathbf{e}_1$ in this diagonal case. We only use \mathbf{e}_1 to stand for $\tilde{\mathbf{u}}_1$ and $\tilde{\mathbf{v}}_1$ in the rest of our analysis.

Let $R > 0$. Since $\mathbf{g}(\boldsymbol{\theta})$ is \mathcal{C}^2 -smooth, there exists $\beta > 0$ such that

$$\|\mathbf{J}(\boldsymbol{\theta}) - \mathbf{J}(\boldsymbol{\theta} + \mathbf{h})\|_2 \leq \beta \|\mathbf{h}\|_2 \quad (19)$$

for all $\|\boldsymbol{\theta}\|_2, \|\boldsymbol{\theta} + \mathbf{h}\|_2 \leq R$. Then the following can be proved by integration:

$$\mathbf{g}(\boldsymbol{\theta} + \mathbf{h}) - \mathbf{g}(\boldsymbol{\theta}) = \left(\int_0^1 \mathbf{J}(\boldsymbol{\theta} + \xi\mathbf{h}) d\xi \right) \mathbf{h}, \quad (20)$$

$$\|\mathbf{g}(\boldsymbol{\theta} + \mathbf{h}) - \mathbf{g}(\boldsymbol{\theta}) - \mathbf{J}(\boldsymbol{\theta})\mathbf{h}\|_2 \leq \beta \|\mathbf{h}\|_2^2. \quad (21)$$

By (21), we also have

$$\|\mathbf{g}(\boldsymbol{\theta}) - \mathbf{J}(0)\boldsymbol{\theta}\|_2 = \|\mathbf{g}(\boldsymbol{\theta}) - \mathbf{g}(0) - \mathbf{J}(0)\boldsymbol{\theta}\|_2 \leq \beta \|\boldsymbol{\theta}\|_2^2. \quad (22)$$

Let $\kappa := \beta/\tilde{\mu}_1$. We assume WLOG that $R \leq 1/\kappa$. Let $F(x) = \log x - \log(1 + \kappa x)$. It is easy to see that $F'(x) = \frac{1}{x+\kappa x^2}$ and $F(x)$ is an increasing function with range $(-\infty, \log(1/\kappa))$. We use $F^{-1}(y)$ to denote the inverse function of $F(x)$. Define $T_\alpha(r) := \frac{1}{\tilde{\mu}_1} (F(r) - F(\alpha)) = \frac{1}{\tilde{\mu}_1} \left(\log \frac{r}{\alpha} - \log \frac{1+\kappa r}{1+\kappa \alpha} \right)$.

Our proof only relies on the following properties of $\mathbf{J}(\mathbf{0})$ (besides that $\tilde{\mu}_1, \mathbf{e}_1$ are the top eigenvalue and eigenvector of $\mathbf{J}(\mathbf{0})$):

Lemma G.2. For $\mathbf{J}(\mathbf{0}) := \text{diag}(\tilde{\mu}_1, \dots, \tilde{\mu}_d)$, we have

1. For any $\mathbf{h} \in \mathbb{R}^d$, $\mathbf{h}^\top \mathbf{J}(\mathbf{0}) \mathbf{h} \leq \tilde{\mu}_1 \|\mathbf{h}\|_2^2$;
2. For any $t \geq 0$, $\|e^{t\mathbf{J}(\mathbf{0})} - e^{\tilde{\mu}_1 t} \mathbf{e}_1 \mathbf{e}_1^\top\|_2 = e^{\tilde{\mu}_2 t}$.

Proof. For Item 1, $\mathbf{h}^\top \mathbf{J}(\mathbf{0}) \mathbf{h} = \sum_{i=1}^d \tilde{\mu}_i h_i^2 \leq \tilde{\mu}_1 \|\mathbf{h}\|_2^2$. For Item 2, $\|e^{t\mathbf{J}(\mathbf{0})} - e^{\tilde{\mu}_1 t} \mathbf{e}_1 \mathbf{e}_1^\top\|_2 = \|\text{diag}(0, e^{\tilde{\mu}_2 t}, \dots, e^{\tilde{\mu}_d t})\|_2 = e^{\tilde{\mu}_2 t}$. \square

Lemma G.3. For $\boldsymbol{\theta}(t) = \phi(\boldsymbol{\theta}_0, t)$ with $\|\boldsymbol{\theta}_0\|_2 \leq \alpha$ and $t \leq T_\alpha(r)$,

$$\|\boldsymbol{\theta}(t)\|_2 \leq \frac{1 + \kappa r}{1 + \kappa \alpha} \alpha \cdot e^{\tilde{\mu}_1 t} \leq r.$$

Proof. By (22) and Lemma G.2, we have

$$\frac{1}{2} \frac{d\|\boldsymbol{\theta}(t)\|_2^2}{dt} = \langle \boldsymbol{\theta}(t), \mathbf{g}(\boldsymbol{\theta}(t)) \rangle \leq \langle \boldsymbol{\theta}(t), \mathbf{J}(\mathbf{0})\boldsymbol{\theta}(t) \rangle + \beta \|\boldsymbol{\theta}(t)\|_2^3 \leq \tilde{\mu}_1 \|\boldsymbol{\theta}(t)\|_2^2 + \beta \|\boldsymbol{\theta}(t)\|_2^3.$$

This implies $\frac{d\|\boldsymbol{\theta}(t)\|_2}{dt} \leq \tilde{\mu}_1 (\|\boldsymbol{\theta}(t)\|_2 + \kappa \|\boldsymbol{\theta}(t)\|_2^2)$. Since $F'(x) = \frac{1}{x+\kappa x^2}$, we further have

$$\frac{d}{dt} F(\|\boldsymbol{\theta}(t)\|_2) \leq \tilde{\mu}_1.$$

So $F(\|\boldsymbol{\theta}(t)\|_2) \leq F(\alpha) + \tilde{\mu}_1 t$. By definition of $T_\alpha(r)$, we then know that $\|\boldsymbol{\theta}(t)\|_2 \leq r$ for all $t \leq T_\alpha(r)$. So

$$\log \|\boldsymbol{\theta}(t)\|_2 \leq F(\|\boldsymbol{\theta}(t)\|_2) + \log(1 + \kappa r) \leq F(\alpha) + \tilde{\mu}_1 t + \log(1 + \kappa r).$$

Expanding $F(\alpha)$ proves the lemma. \square

Lemma G.4. For $\boldsymbol{\theta}(t) = \phi(\boldsymbol{\theta}_0, t)$ with $\|\boldsymbol{\theta}_0\|_2 \leq \alpha$ and $t \leq T_\alpha(r)$, we have

$$\boldsymbol{\theta}(t) = e^{t\mathbf{J}(\mathbf{0})} \boldsymbol{\theta}_0 + O(r^2).$$

Proof. Let $\hat{\boldsymbol{\theta}}(t) = e^{t\mathbf{J}(\mathbf{0})} \boldsymbol{\theta}_0$. Then we have

$$\begin{aligned} \frac{1}{2} \frac{d}{dt} \|\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t)\|_2^2 &\leq \left\langle \mathbf{g}(\boldsymbol{\theta}(t)) - \mathbf{J}(\mathbf{0})\hat{\boldsymbol{\theta}}(t), \boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t) \right\rangle \\ &= \left\langle \mathbf{g}(\boldsymbol{\theta}(t)) - \mathbf{J}(\mathbf{0})\boldsymbol{\theta}(t), \boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t) \right\rangle + (\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t))^\top \mathbf{J}(\mathbf{0})(\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t)) \\ &\leq \|\mathbf{g}(\boldsymbol{\theta}(t)) - \mathbf{J}(\mathbf{0})\boldsymbol{\theta}(t)\|_2 \cdot \|\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t)\|_2 + \tilde{\mu}_1 \|\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t)\|_2^2, \end{aligned}$$

where the last inequality is due to Lemma G.2. By (22) and Lemma G.3, we have

$$\|\mathbf{g}(\boldsymbol{\theta}(t)) - \mathbf{J}(\mathbf{0})\boldsymbol{\theta}(t)\|_2 \leq \beta \|\boldsymbol{\theta}(t)\|_2^2 \leq \beta \left(\frac{1 + \kappa r}{1 + \kappa \alpha} \alpha \right)^2 \cdot e^{2\tilde{\mu}_1 t}.$$

So we have $\frac{d}{dt} \|\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t)\|_2 \leq \beta \left(\frac{1 + \kappa r}{1 + \kappa \alpha} \alpha \right)^2 \cdot e^{2\tilde{\mu}_1 t} + \tilde{\mu}_1 \|\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t)\|_2$. By Grönwall's inequality,

$$\|\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t)\|_2 \leq \int_0^t \beta \left(\frac{1 + \kappa r}{1 + \kappa \alpha} \alpha \right)^2 \cdot e^{2\tilde{\mu}_1 \tau} e^{\tilde{\mu}_1(t-\tau)} d\tau.$$

Evaluating the integral gives

$$\|\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t)\|_2 \leq \beta \left(\frac{1 + \kappa r}{1 + \kappa \alpha} \alpha \right)^2 e^{\tilde{\mu}_1 t} \cdot \frac{e^{\tilde{\mu}_1 t} - 1}{\tilde{\mu}_1} \leq \kappa \left(\frac{1 + \kappa r}{1 + \kappa \alpha} \alpha \cdot e^{\tilde{\mu}_1 t} \right)^2 \leq \kappa r^2,$$

which proves the lemma. \square

Lemma G.5. Let $\boldsymbol{\theta}(t) = \phi(\boldsymbol{\theta}_0, t)$, $\hat{\boldsymbol{\theta}}(t) = \phi(\hat{\boldsymbol{\theta}}_0, t)$. If $\max\{\|\boldsymbol{\theta}_0\|_2, \|\hat{\boldsymbol{\theta}}_0\|_2\} \leq \alpha$, then for $t \leq T_\alpha(r)$,

$$\|\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t)\|_2 \leq e^{\tilde{\mu}_1 t + \kappa r} \|\boldsymbol{\theta}_0 - \hat{\boldsymbol{\theta}}_0\|_2.$$

Proof. For $t \leq T_\alpha(r)$, by (20),

$$\begin{aligned} \frac{1}{2} \frac{d}{dt} \|\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t)\|_2^2 &= \left\langle \mathbf{g}(\boldsymbol{\theta}(t)) - \mathbf{g}(\hat{\boldsymbol{\theta}}(t)), \boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t) \right\rangle \\ &= (\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t))^\top \left(\int_0^1 \mathbf{J}(\boldsymbol{\theta}_\xi(t)) d\xi \right) (\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t)), \end{aligned}$$

where $\boldsymbol{\theta}_\xi(t) := \xi\boldsymbol{\theta}(t) + (1 - \xi)\hat{\boldsymbol{\theta}}(t)$. By Lemma G.3, $\max\{\|\boldsymbol{\theta}(t)\|_2, \|\hat{\boldsymbol{\theta}}(t)\|_2\} \leq \frac{1+\kappa r}{1+\kappa\alpha} \alpha \cdot e^{\tilde{\mu}_1 t}$ for all $t \leq T_\alpha(r)$. So $\|\boldsymbol{\theta}_\xi(t)\|_2 \leq \frac{1+\kappa r}{1+\kappa\alpha} \alpha \cdot e^{\tilde{\mu}_1 t}$. Combining these with (19) and Lemma G.2, we have

$$\mathbf{h}^\top \mathbf{J}(\boldsymbol{\theta}_\xi(t)) \mathbf{h} = \mathbf{h}^\top \mathbf{J}(\mathbf{0}) \mathbf{h} + \mathbf{h}^\top (\mathbf{J}(\boldsymbol{\theta}_\xi(t)) - \mathbf{J}(\mathbf{0})) \mathbf{h} \leq \left(\tilde{\mu}_1 + \beta \cdot \frac{1 + \kappa r}{1 + \kappa \alpha} \alpha \cdot e^{\tilde{\mu}_1 t} \right) \|\mathbf{h}\|_2^2,$$

for all $\mathbf{h} \in \mathbb{R}^d$. Thus, $\frac{d}{dt} \|\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t)\|_2 \leq \left(\tilde{\mu}_1 + \beta \cdot \frac{1+\kappa r}{1+\kappa\alpha} \alpha \cdot e^{\tilde{\mu}_1 t} \right) \|\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t)\|_2$. This implies

$$\begin{aligned} \log \frac{\|\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t)\|_2}{\|\boldsymbol{\theta}(0) - \hat{\boldsymbol{\theta}}(0)\|_2} &\leq \int_0^t \left(\tilde{\mu}_1 + \beta \cdot \frac{1 + \kappa r}{1 + \kappa \alpha} \alpha \cdot e^{\tilde{\mu}_1 \tau} \right) d\tau \\ &\leq \tilde{\mu}_1 t + \kappa \cdot \frac{1 + \kappa r}{1 + \kappa \alpha} \alpha e^{\tilde{\mu}_1 t} \\ &\leq \tilde{\mu}_1 t + \kappa r. \end{aligned}$$

Therefore, $\|\boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t)\|_2 \leq e^{\tilde{\mu}_1 t + \kappa r} \|\boldsymbol{\theta}(0) - \hat{\boldsymbol{\theta}}(0)\|_2$. \square

Lemma G.6. For every $t \in (-\infty, +\infty)$, $\mathbf{z}(t)$ exists and $\mathbf{z}_\alpha(t)$ converges to $\mathbf{z}(t)$ in the following rate:

$$\|\mathbf{z}_\alpha(t) - \mathbf{z}(t)\|_2 = O(\alpha),$$

where O hides constants depending on $\mathbf{g}(\boldsymbol{\theta})$ and t .

Proof. We prove the lemma in the cases of $t \in (-\infty, F(R)/\tilde{\mu}_1]$ and $t > F(R)/\tilde{\mu}_1$ respectively.

Case 1. Fix $t \in (-\infty, F(R)/\tilde{\mu}_1]$. Let $\tilde{\alpha}$ be the unique number such that $\frac{\tilde{\alpha}}{1+\kappa\tilde{\alpha}} = \alpha$ (i.e., $F(\tilde{\alpha}) = \log \alpha$). Let α' be an arbitrary number less than α . Let $t_0 := \frac{1}{\tilde{\mu}_1} \log \frac{\alpha}{\alpha'}$. Then $t_0 = \frac{1}{\tilde{\mu}_1} (F(\tilde{\alpha}) - \log \alpha') \leq T_{\alpha'}(\tilde{\alpha})$. By Lemma G.4, we have

$$\|\phi(\alpha' \mathbf{e}_1, t_0) - \alpha \mathbf{e}_1\|_2 = \left\| \phi(\alpha' \mathbf{e}_1, t_0) - e^{t_0 \mathbf{J}(\mathbf{0})} \alpha' \mathbf{e}_1 \right\|_2 = O(\tilde{\alpha}^2).$$

Let $r := F^{-1}(\tilde{\mu}_1 t) \leq R$. Then $t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha} = T_{\tilde{\alpha}}(r)$ if $\tilde{\alpha} < r$.

By Lemma G.3, $\|\phi(\alpha' \mathbf{e}_1, t_0)\|_2 \leq \tilde{\alpha}$. Also, $\|\alpha \mathbf{e}_1\|_2 = \frac{\tilde{\alpha}}{1+\kappa\tilde{\alpha}} \leq \tilde{\alpha}$. By Lemma G.5,

$$\begin{aligned} \|\mathbf{z}_\alpha(t) - \mathbf{z}_{\alpha'}(t)\|_2 &= \left\| \phi\left(\alpha' \mathbf{e}_1, t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha'}\right) - \phi\left(\alpha \mathbf{e}_1, t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha}\right) \right\|_2 \\ &= \left\| \phi\left(\phi(\alpha' \mathbf{e}_1, t_0), t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha}\right) - \phi\left(\alpha \mathbf{e}_1, t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha}\right) \right\|_2 \\ &\leq O(\tilde{\alpha}^2 \cdot e^{\tilde{\mu}_1(t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha}) + \kappa r}) \\ &\leq O\left(\frac{\tilde{\alpha}^2}{\alpha}\right). \end{aligned}$$

For α small enough, we have $\tilde{\alpha} = O(\alpha)$, so for any $\alpha' \in (0, \alpha)$,

$$\|\mathbf{z}_\alpha(t) - \mathbf{z}_{\alpha'}(t)\|_2 = O(\alpha).$$

This implies that $\{\mathbf{z}_\alpha(t)\}$ satisfies Cauchy's criterion for every t , and thus the limit $\mathbf{z}(t)$ exists for $t \leq F(R)/\tilde{\mu}_1$. The convergence rate can be deduced by taking limits for $\alpha' \rightarrow 0$ on both sides.

Case 2. For $t = F(R)/\tilde{\mu}_1 + \tau$ with $\tau > 0$, $\phi(\boldsymbol{\theta}, \tau)$ is locally Lipschitz with respect to $\boldsymbol{\theta}$. So

$$\begin{aligned}\|\mathbf{z}_\alpha(t) - \mathbf{z}_{\alpha'}(t)\|_2 &= \|\phi(\mathbf{z}_\alpha(F(R)/\tilde{\mu}_1), \tau) - \phi(\mathbf{z}_{\alpha'}(F(R)/\tilde{\mu}_1), \tau)\|_2 \\ &= O(\|\mathbf{z}_\alpha(F(R)/\tilde{\mu}_1) - \mathbf{z}_{\alpha'}(F(R)/\tilde{\mu}_1)\|_2) \\ &= O(\alpha),\end{aligned}$$

which proves the lemma for $t > F(R)/\tilde{\mu}_1$. \square

Proof for Theorem G.1. The existence of $\mathbf{z}(t) := \lim_{\alpha \rightarrow 0} \mathbf{z}_\alpha(t) = \lim_{\alpha \rightarrow 0} \phi\left(\alpha \mathbf{e}_1, t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha}\right)$ has already been proved in Lemma G.6, where we show $\|\mathbf{z}_\alpha(t) - \mathbf{z}(t)\|_2 = O(\alpha)$.

By the continuity of $\phi(\cdot, t)$ for every $t \in \mathbb{R}$, we have

$$\mathbf{z}(t) = \lim_{\alpha \rightarrow 0} \phi\left(\alpha \tilde{\mathbf{v}}_1, t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha}\right) = \phi\left(\lim_{\alpha \rightarrow 0} \phi\left(\alpha \tilde{\mathbf{v}}_1, \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha}\right), t\right) = \phi(\mathbf{z}(0), t).$$

Now it is only left to prove (7). WLOG we can assume that $\|\boldsymbol{\delta}_\alpha\|_2$ is decreasing and $\frac{\alpha}{2} \leq \|\boldsymbol{\delta}_\alpha\|_2 \leq \alpha$ (otherwise we can do reparameterization). Then our goal becomes proving

$$\|\boldsymbol{\theta}_\alpha(t) - \mathbf{z}(t)\|_2 = O\left(\alpha^{\frac{\tilde{\gamma}}{\tilde{\mu}_1 + \tilde{\gamma}}}\right). \quad (23)$$

where $\boldsymbol{\theta}_\alpha(t) := \phi\left(\boldsymbol{\delta}_\alpha, t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\langle \boldsymbol{\delta}_\alpha, \mathbf{e}_1 \rangle}\right)$. We prove (23) in the cases of $t \in (-\infty, F(R)/\tilde{\mu}_1]$ and $t > F(R)/\tilde{\mu}_1$ respectively.

Case 1. Fix $t \in (-\infty, (F(R) + \log q)/\tilde{\mu}_1]$. Let $\tilde{\alpha}_1 = \alpha^{\frac{\tilde{\gamma}}{\tilde{\mu}_1 + \tilde{\gamma}}}$. Let $\alpha_1 := e^{F(\tilde{\alpha}_1)} = \frac{\tilde{\alpha}_1}{1 + \kappa \tilde{\alpha}_1}$. Let $t_0 := \frac{1}{\tilde{\mu}_1}(F(\tilde{\alpha}_1) - \log \alpha) \leq T_{\|\boldsymbol{\delta}_\alpha\|_2}(\tilde{\alpha}_1)$. At time t_0 , by Lemma G.2 we have

$$\left\|e^{t_0 \mathbf{J}(0)} - e^{\tilde{\mu}_1 t_0} \mathbf{e}_1 \mathbf{e}_1^\top\right\|_2 = e^{\tilde{\mu}_2 t_0} = e^{\frac{\tilde{\mu}_2}{\tilde{\mu}_1}(F(\tilde{\alpha}_1) - \log \alpha)} = \left(\frac{\alpha_1}{\alpha}\right)^{\frac{\tilde{\mu}_2}{\tilde{\mu}_1}}. \quad (24)$$

Let $q_\alpha := \langle \frac{\boldsymbol{\delta}_\alpha}{\alpha}, \mathbf{e}_1 \rangle$. By Definition 5.2, there exists $q > 0$ such that $q_\alpha \geq q$ for all sufficiently small α . Then we have

$$\begin{aligned}\|\phi(\boldsymbol{\delta}_\alpha, t_0) - \alpha_1 q_\alpha \mathbf{e}_1\|_2 &= \left\|\phi(\boldsymbol{\delta}_\alpha, t_0) - e^{t_0 \mathbf{J}(0)} \boldsymbol{\delta}_\alpha\right\|_2 + \left\|\left(e^{t_0 \mathbf{J}(0)} - e^{\tilde{\mu}_1 t_0} \mathbf{e}_1 \mathbf{e}_1^\top\right) \boldsymbol{\delta}_\alpha\right\|_2 \\ &= O(\tilde{\alpha}_1^2) + \left(\frac{\alpha_1}{\alpha}\right)^{\frac{\tilde{\mu}_2}{\tilde{\mu}_1}} \|\boldsymbol{\delta}_\alpha\|_2 \\ &= O(\tilde{\alpha}_1^2 + \alpha_1^{\tilde{\mu}_2/\tilde{\mu}_1} \alpha^{1 - \tilde{\mu}_2/\tilde{\mu}_1}) \\ &= O(\alpha_1^2).\end{aligned}$$

Let $r := F^{-1}(\tilde{\mu}_1 t + \log \frac{1}{q_\alpha}) \leq R$. Then $t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha_1 q_\alpha} = T_{\tilde{\alpha}}(r)$ if $\tilde{\alpha} < r$. By Lemma G.3, $\|\phi(\boldsymbol{\delta}_\alpha, t_0)\|_2 \leq \tilde{\alpha}_1$. Also, $\|\alpha_1 q_\alpha \mathbf{e}_1\|_2 \leq \alpha_1 = \frac{\tilde{\alpha}_1}{1 + \kappa \tilde{\alpha}_1} \leq \tilde{\alpha}_1$. By Lemma G.5,

$$\begin{aligned}\|\boldsymbol{\theta}_\alpha(t) - \mathbf{z}_{\alpha_1}(t)\|_2 &\leq \left\|\phi\left(\phi(\boldsymbol{\delta}_\alpha, t_0), t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha_1 q_\alpha}\right) - \phi\left(\alpha_1 q_\alpha \mathbf{e}_1, t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha_1 q_\alpha}\right)\right\|_2 \\ &= O\left(\alpha_1^2 \cdot e^{\tilde{\mu}_1\left(t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\alpha_1 q_\alpha}\right) + \kappa r}\right) \\ &= O(\alpha_1).\end{aligned}$$

Combining this with the convergence rate for $\mathbf{z}_{\alpha_1}(t)$, we have

$$\|\boldsymbol{\theta}_\alpha(t) - \mathbf{z}(t)\|_2 \leq \|\boldsymbol{\theta}_\alpha(t) - \mathbf{z}_{\alpha_1}(t)\|_2 + \|\mathbf{z}_{\alpha_1}(t) - \mathbf{z}(t)\|_2 = O(\alpha_1).$$

Case 2. For $t = (F(R) + \log q)/\tilde{\mu}_1 + \tau$ with $\tau > 0$, $\phi(\boldsymbol{\theta}, \tau)$ is locally Lipschitz with respect to $\boldsymbol{\theta}$. So

$$\begin{aligned}\|\boldsymbol{\theta}_\alpha(t) - \mathbf{z}(t)\|_2 &= \|\phi(\boldsymbol{\theta}_\alpha((F(R) + \log q)/\tilde{\mu}_1), \tau) - \phi(\mathbf{z}((F(R) + \log q)/\tilde{\mu}_1), \tau)\|_2 \\ &= O(\|\boldsymbol{\theta}_\alpha((F(R) + \log q)/\tilde{\mu}_1) - \mathbf{z}((F(R) + \log q)/\tilde{\mu}_1)\|_2) \\ &= O(\alpha_1),\end{aligned}$$

which proves (23) for $t > (F(R) + \log q)/\tilde{\mu}_1$. \square

there exists $\tilde{\mathbf{V}}_\delta$ such that $\mathbf{J}(\mathbf{0}) = \tilde{\mathbf{V}}_\delta \mathbf{J}_\delta \tilde{\mathbf{V}}_\delta^{-1}$, where $\mathbf{J}_\delta := \text{diag}(\mathbf{J}_{\delta[1]}, \dots, \mathbf{J}_{\delta[m]})$, $\mathbf{J}_{\delta[j]} := \mathbf{J}_{a,\delta}^{(r)}$ if $\mathbf{J}_{[j]} := \mathbf{J}_{a,1}^{(r)}$, or $\mathbf{J}_{\delta[j]} := \mathbf{J}_{a,b,\delta}^{(r)}$ if $\mathbf{J}_{[j]} := \mathbf{J}_{a,b,1}^{(r)}$. Since the top eigenvalue of $\mathbf{J}(\mathbf{0})$ is positive and unique, $\tilde{\mu}_1$ corresponds to only one block $[\tilde{\mu}_1] \in \mathbb{R}^{1 \times 1}$. WLOG we let $\mathbf{J}_1 = [\tilde{\mu}_1]$, and thus $\mathbf{J}_{\delta[1]} = [\tilde{\mu}_1]$.

We only need to select a parameter $\delta > 0$ and prove the theorem in the case of $\mathbf{J}(\mathbf{0}) = \mathbf{J}_\delta$ since we can change the basis in a similar way as we have done in Appendix G.1. By scrutinizing the proof for Theorem G.1, we can find that we only need to reprove Lemma G.2. However, Lemma G.2 may not be correct since $\mathbf{J}(\mathbf{0})$ is not diagonal anymore. Instead, we prove the following:

1. If $\delta \in (0, \tilde{\gamma})$, then $\mathbf{h}^\top \mathbf{J}_\delta \mathbf{h} \leq \tilde{\mu}_1 \|\mathbf{h}\|_2^2$ for all $\mathbf{h} \in \mathbb{R}^d$;
2. For any $\tilde{\mu}'_2 \in (\Re(\tilde{\mu}_2), \tilde{\mu}_1)$, if $\delta \in (0, \tilde{\mu}'_2 - \Re(\tilde{\mu}_2))$, then $\|e^{t\mathbf{J}_\delta} - e^{\tilde{\mu}_1 t} \mathbf{e}_1 \mathbf{e}_1^\top\|_2 \leq e^{\tilde{\mu}'_2 t}$ for all $t \geq 0$.

Proof for Item 1. Let \mathcal{K} be the set of pairs (k_1, k_2) such that $k_1 \neq k_2$ and the entry of \mathbf{J}_δ at the k_1 -th row and the k_2 -th column is non-zero. Then we have

$$\begin{aligned} \mathbf{h}^\top \mathbf{J}_\delta \mathbf{h} &= \mathbf{h}^\top \frac{\mathbf{J}_\delta + \mathbf{J}_\delta^\top}{2} \mathbf{h} = \sum_{k=1}^d \Re(\tilde{\mu}_k) h_k^2 + \sum_{(k_1, k_2) \in \mathcal{K}} h_{k_1} h_{k_2} \delta \\ &\leq \sum_{k=1}^d \Re(\tilde{\mu}_k) h_k^2 + \sum_{(k_1, k_2) \in \mathcal{K}} \frac{h_{k_1}^2 + h_{k_2}^2}{2} \delta. \end{aligned}$$

Note that $\Re(\tilde{\mu}_k) \leq \Re(\tilde{\mu}_2)$ for $k \geq 2$. Also note that there is no pair in \mathcal{K} has $k_1 = 1$ or $k_2 = 1$, and for every $k \geq 2$ there are at most two pairs in \mathcal{K} has $k_1 = k$ or $k_2 = k$. Combining all these together gives

$$\mathbf{h}^\top \mathbf{J}_\delta \mathbf{h} \leq \tilde{\mu}_1 h_1^2 + (\Re(\tilde{\mu}_2) + \delta) \sum_{k=2}^d h_k^2 \leq \tilde{\mu}_1 \|\mathbf{h}\|_2^2,$$

which proves Item 1.

Proof for Item 2. Since \mathbf{J}_δ is block diagonal, we only need to prove that $\|e^{t\mathbf{J}_{\delta[j]}}\|_2 \leq e^{\tilde{\mu}'_2 t}$ for every $j \geq 2$. If $\mathbf{J}_{\delta[j]} = \mathbf{J}_{a,\delta}^{(r)} = a\mathbf{I} + \delta\mathbf{N}$, where \mathbf{N} is the nilpotent matrix, then

$$e^{t\mathbf{J}_{\delta[j]}} = e^{at\mathbf{I} + \delta t\mathbf{N}} = e^{at\mathbf{I}} e^{\delta t\mathbf{N}} = e^{at} e^{\delta t\mathbf{N}},$$

where the second equality uses the fact that \mathbf{I} and \mathbf{N} are commutable. So we have

$$\|e^{t\mathbf{J}_{\delta[j]}}\|_2 \leq e^{at} \|e^{\delta t\mathbf{N}}\|_2 = e^{at} e^{\delta t \|\mathbf{N}\|_2} \leq e^{(a+\delta)t}.$$

If $\mathbf{J}_{\delta[j]} = \mathbf{J}_{a,\delta}^{(r)} = \mathbf{D} + \delta\mathbf{N}^2$, where $\mathbf{D} = \text{diag}(\mathbf{C}, \mathbf{C}, \dots, \mathbf{C})$ and \mathbf{N} is the nilpotent matrix, then

$$e^{t\mathbf{J}_{\delta[j]}} = e^{t\mathbf{D} + \delta t\mathbf{N}^2} = e^{t\mathbf{D}} e^{\delta t\mathbf{N}^2},$$

where the second equality uses the fact that \mathbf{D} and \mathbf{N}^2 are commutable. Note that $e^{t\mathbf{C}} = e^{at} \begin{bmatrix} \cos(bt) & -\sin(bt) \\ \sin(bt) & \cos(bt) \end{bmatrix}$, which implies $\|e^{t\mathbf{D}}\|_2 = \|e^{t\mathbf{C}}\|_2 = e^{at}$. So we have

$$\|e^{t\mathbf{J}_{\delta[j]}}\|_2 \leq \|e^{t\mathbf{D}}\|_2 \cdot \|e^{\delta t\mathbf{N}^2}\|_2 = e^{at} e^{\delta t \|\mathbf{N}^2\|_2} \leq e^{(a+\delta)t}.$$

Since $\delta \in (0, \tilde{\mu}'_2 - \Re(\tilde{\mu}_2))$, we know that $a + \delta < \tilde{\mu}'_2$, which completes the proof.

Proof for a fixed δ . Since Item 1 continues to hold for $\delta \in (0, \tilde{\gamma})$, Lemmas G.3 to G.6 also hold. This proves that $\mathbf{z}(t)$ exists and satisfies (6).

It remains to prove (25) for any $\epsilon > 0$. Let $\tilde{\gamma}' \in (0, \tilde{\gamma})$ be a number such that $\frac{\tilde{\gamma}'}{\tilde{\mu}_1 + \tilde{\gamma}'} \geq \frac{\tilde{\gamma}}{\tilde{\mu}_1 + \tilde{\gamma}} - \epsilon$. Fix $\tilde{\mu}'_2 = \tilde{\mu}_1 - \tilde{\gamma}'$, $\delta = \tilde{\mu}'_2 - \Re(\tilde{\mu}_2)$. By Item 2, we have $\|e^{t\mathbf{J}_\delta} - e^{\tilde{\mu}_1 t} \mathbf{e}_1 \mathbf{e}_1^\top\|_2 \leq e^{\tilde{\mu}'_2 t}$ for all $t \geq 0$. By scrutinizing the proof for Theorem G.1, we can find that the only place we use Item 2 in Lemma G.2

is in (24). For proving (25), we can repeat the proof while replacing all the occurrences of $\tilde{\mu}_2$ by $\tilde{\mu}'_2$. Then we know that for every $t \in \mathbb{R}$, there is a constant $C > 0$ such that

$$\left\| \tilde{\mathbf{V}}_\delta^{-1} \phi \left(\boldsymbol{\delta}_\alpha, t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\langle \boldsymbol{\delta}_\alpha, \tilde{\mathbf{u}}_1 \rangle} \right) - \tilde{\mathbf{V}}_\delta^{-1} \mathbf{z}(t) \right\|_2 \leq C \cdot \|\tilde{\mathbf{V}}_\delta^{-1} \boldsymbol{\delta}_\alpha\|_2^{\frac{\tilde{\gamma}'}{\tilde{\mu}_1 + \tilde{\gamma}'}} ,$$

for every sufficiently small α . By definition of $\tilde{\gamma}'$, $\frac{\tilde{\gamma}'}{\tilde{\mu}_1 + \tilde{\gamma}'} \geq \frac{\tilde{\gamma}}{\tilde{\mu}_1 + \tilde{\gamma}} - \epsilon$. Since $\boldsymbol{\delta}_\alpha \rightarrow \mathbf{0}$ as $\alpha \rightarrow 0$, we have $\|\tilde{\mathbf{V}}_\delta^{-1} \boldsymbol{\delta}_\alpha\|_2 < 1$ for sufficiently small α . Then we have

$$\begin{aligned} \left\| \phi \left(\boldsymbol{\delta}_\alpha, t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\langle \boldsymbol{\delta}_\alpha, \tilde{\mathbf{u}}_1 \rangle} \right) - \mathbf{z}(t) \right\|_2 &\leq \|\tilde{\mathbf{V}}_\delta\|_2 \cdot \left\| \tilde{\mathbf{V}}_\delta^{-1} \phi \left(\boldsymbol{\delta}_\alpha, t + \frac{1}{\tilde{\mu}_1} \log \frac{1}{\langle \boldsymbol{\delta}_\alpha, \tilde{\mathbf{u}}_1 \rangle} \right) - \tilde{\mathbf{V}}_\delta^{-1} \mathbf{z}(t) \right\|_2 \\ &\leq \|\tilde{\mathbf{V}}_\delta\|_2 \cdot C \cdot \|\tilde{\mathbf{V}}_\delta^{-1} \boldsymbol{\delta}_\alpha\|_2^{\frac{\tilde{\gamma}'}{\tilde{\mu}_1 + \tilde{\gamma}'}} \\ &\leq C \cdot \|\tilde{\mathbf{V}}_\delta\|_2 \cdot \|\tilde{\mathbf{V}}_\delta^{-1}\|_2^{\frac{\tilde{\gamma}'}{\tilde{\mu}_1 + \tilde{\gamma}'}} \cdot \|\boldsymbol{\delta}_\alpha\|_2^{\frac{\tilde{\gamma}}{\tilde{\mu}_1 + \tilde{\gamma}} - \epsilon}. \end{aligned}$$

Absorbing $\|\tilde{\mathbf{V}}_\delta\|_2 \cdot \|\tilde{\mathbf{V}}_\delta^{-1}\|_2^{\frac{\tilde{\gamma}'}{\tilde{\mu}_1 + \tilde{\gamma}'}}$ into C proves (25). \square

H EIGENVALUES OF JACOBIANS AND HESSIANS

In this section we analyze the eigenvalues of the Jacobian $\mathbf{J}(\mathbf{W})$ at critical points of (2).

For notation simplicity, we write $\text{sz}(\mathbf{A}) := \mathbf{A} + \mathbf{A}^\top$ to denote the symmetric matrix produced by adding up \mathbf{A} and its transpose, and write $\text{ac}\{\mathbf{A}, \mathbf{B}\} = \mathbf{AB} + \mathbf{BA}$ to denote the anticommutator of two matrices \mathbf{A}, \mathbf{B} . Then $\mathbf{g}(\mathbf{W})$ can be written as $\mathbf{g}(\mathbf{W}) := -\text{ac}\{\nabla f(\mathbf{W}), \mathbf{W}\}$.

Let $\mathbf{U}_0 \in \mathbb{R}^{d \times r}$ be a stationary point of the function $\mathcal{L} : \mathbb{R}^{d \times r} \rightarrow \mathbb{R}$, $\mathcal{L}(\mathbf{U}) = \frac{1}{2} f(\mathbf{U}\mathbf{U}^\top)$, i.e., $\nabla \mathcal{L}(\mathbf{U}_0) = \nabla f(\mathbf{U}_0 \mathbf{U}_0^\top) \mathbf{U}_0 = \mathbf{0}$. By Lemma C.1, this implies

$$\nabla f(\mathbf{W}_0) \mathbf{W}_0 = \mathbf{0} \quad (26)$$

for $\mathbf{W}_0 := \mathbf{U}_0 \mathbf{U}_0^\top$, and thus \mathbf{W}_0 is a critical point of (2).

For a real-valued or vector-valued function $F(\boldsymbol{\theta})$, we use $DF(\boldsymbol{\theta})[\boldsymbol{\delta}]$, $D^2F(\boldsymbol{\theta})[\boldsymbol{\delta}_1, \boldsymbol{\delta}_2]$ to denote the first- and second-order directional derivatives of $F(\cdot)$ at $\boldsymbol{\theta}$.

Let \mathcal{X} be a linear space, which can be $\mathbb{R}^{d \times d}$ or $\mathbb{R}^{d \times r}$. For a function $F : \mathcal{X} \rightarrow \mathcal{X}$, we use $DF(\boldsymbol{\theta})$ to denote the directional derivative of F at $\boldsymbol{\theta}$, represented by the linear operator

$$DF(\boldsymbol{\theta})[\boldsymbol{\Delta}] : \mathcal{X} \rightarrow \mathcal{X}, \boldsymbol{\Delta} \mapsto DF(\boldsymbol{\theta})[\boldsymbol{\Delta}] = \lim_{t \rightarrow 0} \frac{F(\boldsymbol{\theta} + t\boldsymbol{\Delta}) - F(\boldsymbol{\theta})}{t}.$$

We also write $DF(\boldsymbol{\theta})[\boldsymbol{\Delta}_1, \boldsymbol{\Delta}_2] := \langle DF(\boldsymbol{\theta})[\boldsymbol{\Delta}_1], \boldsymbol{\Delta}_2 \rangle$.

For a function $F : \mathcal{X} \rightarrow \mathbb{R}$, we use $D^2F(\boldsymbol{\theta}) = D(\nabla F(\boldsymbol{\theta}))$ to denote the second directional derivative of F at $\boldsymbol{\theta}$, i.e., $D^2F(\boldsymbol{\theta})[\boldsymbol{\Delta}] = D(\nabla F(\boldsymbol{\theta}))[\boldsymbol{\Delta}]$, $D^2F(\boldsymbol{\theta})[\boldsymbol{\Delta}_1, \boldsymbol{\Delta}_2] = D(\nabla F(\boldsymbol{\theta}))[\boldsymbol{\Delta}_1, \boldsymbol{\Delta}_2]$.

Define $\mathbf{J}(\mathbf{W}) := D\mathbf{g}(\mathbf{W})$. By simple calculus, we can compute the formula for $\mathbf{J}(\mathbf{W}_0)$:

$$\begin{aligned} \mathbf{J}(\mathbf{W}_0)[\boldsymbol{\Delta}] &= -\text{ac}\{\nabla f(\mathbf{W}_0), \boldsymbol{\Delta}\} - \text{ac}\{D^2f(\mathbf{W}_0)[\boldsymbol{\Delta}], \mathbf{W}_0\}, \\ \mathbf{J}(\mathbf{W}_0)[\boldsymbol{\Delta}_1, \boldsymbol{\Delta}_2] &= -\langle \nabla f(\mathbf{W}_0), \text{sz}(\boldsymbol{\Delta}_1 \boldsymbol{\Delta}_2^\top) \rangle - D^2f(\mathbf{W}_0)[\boldsymbol{\Delta}_1, \text{sz}(\mathbf{W}_0 \boldsymbol{\Delta}_2^\top)], \end{aligned}$$

where $\boldsymbol{\Delta}, \boldsymbol{\Delta}_1, \boldsymbol{\Delta}_2 \in \mathbb{R}^{d \times d}$.

We can also compute the formula for $D^2\mathcal{L}(\mathbf{U}_0)$:

$$\begin{aligned} D^2\mathcal{L}(\mathbf{U}_0)[\boldsymbol{\Delta}] &= \nabla f(\mathbf{W}_0) \boldsymbol{\Delta} + D^2f(\mathbf{W}_0)[\text{sz}(\boldsymbol{\Delta} \mathbf{U}_0^\top)] \mathbf{U}_0, \\ D^2\mathcal{L}(\mathbf{U}_0)[\boldsymbol{\Delta}_1, \boldsymbol{\Delta}_2] &= \frac{1}{2} (\langle \nabla f(\mathbf{W}_0), \text{sz}(\boldsymbol{\Delta}_1 \boldsymbol{\Delta}_2^\top) \rangle + D^2f(\mathbf{W}_0)[\text{sz}(\boldsymbol{\Delta}_1 \mathbf{U}_0^\top), \text{sz}(\boldsymbol{\Delta}_2 \mathbf{U}_0^\top)]), \end{aligned}$$

where $\boldsymbol{\Delta}, \boldsymbol{\Delta}_1, \boldsymbol{\Delta}_2 \in \mathbb{R}^{d \times r}$.

H.1 EIGENVALUES AT THE ORIGIN

The eigenvalues of $\mathbf{J}(\mathbf{0})$ is given in Lemma 5.4. Now we provide the proof.

Proof for Lemma 5.4. For $\mathbf{W}_0 = \mathbf{0}$, we have

$$\begin{aligned}\mathbf{J}(\mathbf{0})[\Delta] &= -\nabla f(\mathbf{0})\Delta - \Delta\nabla f(\mathbf{0}) \\ \mathbf{J}(\mathbf{0})[\Delta_1, \Delta_2] &= -\langle \nabla f(\mathbf{0}), \text{sz}(\Delta_1\Delta_2^\top) \rangle\end{aligned}$$

It is easy to see from the second equation that $\mathbf{J}(\mathbf{0})$ is symmetric.

Let $-\nabla f(\mathbf{0}) = \sum_{i=1}^d \mu_i \mathbf{u}_{1[i]} \mathbf{u}_{1[i]}^\top$ be the eigendecomposition of the symmetric matrix $-\nabla f(\mathbf{0})$. Then we have

$$\begin{aligned}\mathbf{J}(\mathbf{0})[\Delta] &= \sum_{i=1}^d \mu_i \left(\mathbf{u}_{1[i]} \mathbf{u}_{1[i]}^\top \Delta + \Delta \mathbf{u}_{1[i]} \mathbf{u}_{1[i]}^\top \right) \\ &= \sum_{i=1}^d \sum_{j=1}^d \mu_i \left(\mathbf{u}_{1[i]} \mathbf{u}_{1[i]}^\top \Delta \mathbf{u}_{1[j]} \mathbf{u}_{1[j]}^\top + \mathbf{u}_{1[j]} \mathbf{u}_{1[j]}^\top \Delta \mathbf{u}_{1[i]} \mathbf{u}_{1[i]}^\top \right) \\ &= \sum_{i=1}^d \sum_{j=1}^d (\mu_i + \mu_j) \mathbf{u}_{1[i]} \mathbf{u}_{1[i]}^\top \Delta \mathbf{u}_{1[j]} \mathbf{u}_{1[j]}^\top \\ &= \sum_{i=1}^d \sum_{j=1}^d (\mu_i + \mu_j) \left\langle \Delta, \mathbf{u}_{1[i]} \mathbf{u}_{1[j]}^\top \right\rangle \mathbf{u}_{1[i]} \mathbf{u}_{1[j]}^\top,\end{aligned}$$

which proves (8).

For $\Delta = \mathbf{u}_{1[i]} \mathbf{u}_{1[j]}^\top + \mathbf{u}_{1[j]} \mathbf{u}_{1[i]}^\top$, we have

$$\mathbf{J}(\mathbf{0})[\Delta] = (\mu_i + \mu_j) \mathbf{u}_{1[i]} \mathbf{u}_{1[j]}^\top + (\mu_i + \mu_j) \mathbf{u}_{1[j]} \mathbf{u}_{1[i]}^\top = (\mu_i + \mu_j) \Delta.$$

So $\mathbf{u}_{1[i]} \mathbf{u}_{1[j]}^\top + \mathbf{u}_{1[j]} \mathbf{u}_{1[i]}^\top$ is an eigenvector of $\mathbf{J}(\mathbf{0})$ associated with eigenvalue $\mu_i + \mu_j$. Note that $\{\mathbf{u}_{1[i]} \mathbf{u}_{1[j]}^\top + \mathbf{u}_{1[j]} \mathbf{u}_{1[i]}^\top : i, j \in [d]\}$ spans all the symmetric matrices, so these are all the eigenvectors in the space of symmetric matrices.

For every antisymmetric matrix Δ (i.e., $\Delta = -\Delta^\top$), we have

$$\mathbf{J}(\mathbf{0})[\Delta] = \mathbf{J}(\mathbf{0})[\Delta^\top] = \mathbf{J}(\mathbf{0})[-\Delta].$$

So $\mathbf{J}(\mathbf{0})[\Delta] = \mathbf{0}$ and every antisymmetric matrix is an eigenvector associated with eigenvalue 0.

Since every matrix can be expressed as the sum of a symmetric matrix and an antisymmetric matrix, we have found all the eigenvalues. \square

H.2 EIGENVALUES AT SECOND-ORDER STATIONARY POINTS

Now we study the eigenvalues of $\mathbf{J}(\mathbf{W}_0)$ when \mathbf{U}_0 is a second-order stationary point of $\mathcal{L}(\cdot)$, i.e., $\nabla \mathcal{L}(\mathbf{U}_0) = \mathbf{0}$, $D^2 \mathcal{L}(\mathbf{U}_0)[\Delta, \Delta] \geq \mathbf{0}$ for all $\Delta \in \mathbb{R}^{d \times r}$. We further assume that \mathbf{U}_0 is full-rank, i.e., $\text{rank}(\mathbf{U}_0) = r$. This condition is met if $\mathbf{W}_0 := \mathbf{U}_0 \mathbf{U}_0^\top$ is a local minimizer of $f(\cdot)$ in \mathbb{S}_d^+ but not a minimizer in \mathbb{S}_d^+ .

Lemma H.1. *For $r \leq d$, if $\mathbf{U}_0 \in \mathbb{R}^{d \times r}$ is a second-order stationary point of $\mathcal{L}(\cdot)$, then either $\text{rank}(\mathbf{U}_0) = \text{rank}(\mathbf{W}_0) = r$, or \mathbf{W}_0 is a minimizer of $f(\cdot)$ in \mathbb{S}_d^+ , where $\mathbf{W}_0 = \mathbf{U}_0 \mathbf{U}_0^\top$.*

Proof. Assume to the contrary that \mathbf{U}_0 has $\text{rank} < r$ and \mathbf{W}_0 is a minimizer of $f(\cdot)$ in \mathbb{S}_d^+ . The former one implies that there exists a unit vector $\mathbf{q} \in \mathbb{R}^r$ such that $\mathbf{U}_0 \mathbf{q} = \mathbf{0}$, and the latter one implies that there exists $\mathbf{v} \in \mathbb{R}^d$ such that $\mathbf{v}^\top \nabla f(\mathbf{W}_0) \mathbf{v} < 0$ by Lemma C.2.

Let $\Delta = \mathbf{v}\mathbf{q}^\top$. Then we have

$$\begin{aligned} D^2\mathcal{L}(\mathbf{U}_0)[\Delta, \Delta] &= \langle \nabla f(\mathbf{W}_0), \mathbf{v}\mathbf{v}^\top \rangle + \frac{1}{2}D^2f(\mathbf{W}_0)[\text{sz}(\mathbf{v}(\mathbf{U}_0\mathbf{q})^\top), \text{sz}(\mathbf{v}(\mathbf{U}_0\mathbf{q})^\top)] \\ &= \langle \nabla f(\mathbf{W}_0), \mathbf{v}\mathbf{v}^\top \rangle + \frac{1}{2}D^2f(\mathbf{W}_0)[\mathbf{0}, \mathbf{0}] \\ &= \langle \nabla f(\mathbf{W}_0), \mathbf{v}\mathbf{v}^\top \rangle. \end{aligned}$$

So $D^2\mathcal{L}(\mathbf{U}_0)[\Delta, \Delta] < 0$, which leads to a contradiction. \square

By (26), the symmetric matrices $-\nabla f(\mathbf{W}_0)$ and \mathbf{W}_0 commute, so they can be simultaneously diagonalizable. Since (26) also implies that they have different column spans, we can have the following diagonalization:

$$-\nabla f(\mathbf{W}_0) = \sum_{i=1}^{d-r} \mu_i \mathbf{v}_i \mathbf{v}_i^\top, \quad \mathbf{W}_0 = \sum_{i=d-r+1}^d \mu_i \mathbf{v}_i \mathbf{v}_i^\top. \quad (27)$$

First we prove the following lemma on the eigenvalues and eigenvectors of the linear operator $-D^2\mathcal{L}(\mathbf{U}_0)$:

Lemma H.2. For every $\Delta \in \mathbb{R}^{d \times d}$, if

$$\mathbf{U}_0\Delta^\top + \Delta\mathbf{U}_0^\top = \mathbf{0} \quad (28)$$

then Δ is an eigenvector of the linear operator $-D^2\mathcal{L}(\mathbf{U}_0)[\cdot] : \mathbb{R}^{d \times r} \rightarrow \mathbb{R}^{d \times r}$ associated with eigenvalue 0. Moreover, the solutions of (28) spans a linear space of dimension $\frac{r(r-1)}{2}$.

Proof. Suppose $\mathbf{U}_0\Delta^\top + \Delta\mathbf{U}_0^\top = \mathbf{0}$. Then we have $\mathbf{U}_0\Delta^\top = -\Delta\mathbf{U}_0^\top$, and thus $\Delta^\top = -\mathbf{U}_0^+ \Delta \mathbf{U}_0^+$, where \mathbf{U}_0^+ is the pseudoinverse of the full-rank matrix \mathbf{U}_0 . This implies that there is a matrix $\mathbf{R} \in \mathbb{R}^{r \times r}$, such that $\Delta = \mathbf{U}_0 \mathbf{R}$. Then we have

$$\begin{aligned} -D^2\mathcal{L}(\mathbf{U}_0)[\Delta] &= -\nabla f(\mathbf{W}_0)\mathbf{U}_0\mathbf{R} - D^2f(\mathbf{W}_0)[\mathbf{U}_0\Delta^\top + \Delta\mathbf{U}_0^\top]\mathbf{U}_0 \\ &= -(\nabla f(\mathbf{W}_0)\mathbf{U}_0)\mathbf{R} - D^2f(\mathbf{W}_0)[\mathbf{0}]\mathbf{U}_0 \\ &= \mathbf{0}. \end{aligned}$$

Replacing Δ with $\mathbf{U}_0\mathbf{R}$ in (28) gives $\mathbf{U}_0(\mathbf{R} + \mathbf{R}^\top)\mathbf{U}_0^\top = \mathbf{0}$, which is equivalent to $\mathbf{R} = -\mathbf{R}^\top$ since \mathbf{U}_0 is full-rank. Since the dimension of $r \times r$ antisymmetric matrices is $\frac{r(r-1)}{2}$, the span spanned by the solutions of (28) also has dimension $\frac{r(r-1)}{2}$. \square

Definition H.3 (Eigendecomposition of $-D^2\mathcal{L}(\mathbf{U}_0)$). Let

$$-D^2\mathcal{L}(\mathbf{U}_0)[\Delta] = \sum_{p=1}^{rd} \xi_p \langle \mathbf{E}_p, \Delta \rangle \mathbf{E}_p$$

be the eigendecomposition of the symmetric linear operator $-D^2\mathcal{L}(\mathbf{U}_0)[\cdot] : \mathbb{R}^{d \times r} \rightarrow \mathbb{R}^{d \times r}$, where $\xi_1, \dots, \xi_{rd} \in \mathbb{R}$ are eigenvalues, $\mathbf{E}_1, \dots, \mathbf{E}_{rd} \in \mathbb{R}^{d \times r}$ are eigenvectors satisfying $\langle \mathbf{E}_p, \mathbf{E}_q \rangle = \delta_{pq}$. We enforce ξ_p to be 0 and \mathbf{E}_p to be a solution of (28) for every $rd - \frac{r(r-1)}{2} < p \leq rd$.

Lemma H.4. Let $\mathbf{A} \in \mathbb{R}^{D \times D}$ be a matrix. If $\{\hat{\mathbf{u}}_1, \dots, \hat{\mathbf{u}}_K\}$ is a set of linearly independent left eigenvectors associated with eigenvalues $\hat{\lambda}_1, \dots, \hat{\lambda}_K$ and $\{\tilde{\mathbf{v}}_1, \dots, \tilde{\mathbf{v}}_{D-K}\}$ is a set of linearly independent right eigenvectors associated with eigenvalues $\tilde{\lambda}_1, \dots, \tilde{\lambda}_{D-K}$, and $\langle \hat{\mathbf{u}}_i, \tilde{\mathbf{v}}_j \rangle = 0$ for all $1 \leq i \leq K, 1 \leq j \leq D - K$, then $\hat{\lambda}_1, \dots, \hat{\lambda}_K, \tilde{\lambda}_1, \dots, \tilde{\lambda}_{D-K}$ are all the eigenvalues of \mathbf{A} .

Proof. Let $\hat{\mathbf{U}} := (\hat{\mathbf{u}}_1, \dots, \hat{\mathbf{u}}_K)^\top \in \mathbb{R}^{K \times D}$ and $\tilde{\mathbf{V}} := (\tilde{\mathbf{v}}_1, \dots, \tilde{\mathbf{v}}_{D-K}) \in \mathbb{R}^{D \times (D-K)}$. Then both $\hat{\mathbf{U}}$ and $\tilde{\mathbf{V}}$ are full-rank. Let $\hat{\mathbf{U}}^+ = \hat{\mathbf{U}}^\top (\hat{\mathbf{U}}\hat{\mathbf{U}}^\top)^{-1}$, $\tilde{\mathbf{V}}^+ = (\tilde{\mathbf{V}}^\top \tilde{\mathbf{V}})^{-1} \tilde{\mathbf{V}}^\top$ be the pseudoinverses of $\hat{\mathbf{U}}$ and $\tilde{\mathbf{V}}$.

Now we define

$$\mathbf{P} := \begin{bmatrix} \hat{\mathbf{U}} \\ \tilde{\mathbf{V}}^+ \end{bmatrix}, \quad \mathbf{Q} := [\hat{\mathbf{U}}^+ \quad \tilde{\mathbf{V}}].$$

Then we have

$$\mathbf{P}\mathbf{Q} = \begin{bmatrix} \hat{\mathbf{U}}\hat{\mathbf{U}}^+ & \hat{\mathbf{U}}\tilde{\mathbf{V}} \\ \tilde{\mathbf{V}}^+\hat{\mathbf{U}}^+ & \tilde{\mathbf{V}}^+\tilde{\mathbf{V}} \end{bmatrix}.$$

Note that $\hat{\mathbf{U}}\hat{\mathbf{U}}^+ = \mathbf{I}_K$, $\hat{\mathbf{U}}\tilde{\mathbf{V}} = \mathbf{0}$, $\tilde{\mathbf{V}}^+\hat{\mathbf{U}}^+ = (\tilde{\mathbf{V}}^\top \tilde{\mathbf{V}})^{-1}(\hat{\mathbf{U}}\tilde{\mathbf{V}})^\top (\hat{\mathbf{U}}\hat{\mathbf{U}}^\top)^{-1} = \mathbf{0}$, $\tilde{\mathbf{V}}^+\tilde{\mathbf{V}} = \mathbf{I}_{D-K}$. So $\mathbf{P}\mathbf{Q} = \mathbf{I}_D$, or equivalently $\mathbf{Q} = \mathbf{P}^{-1}$. Then we have

$$\mathbf{P}^{-1}\mathbf{A}\mathbf{P} = \begin{bmatrix} \text{diag}(\hat{\lambda}_1, \dots, \hat{\lambda}_K) & * \\ \mathbf{0} & \text{diag}(\tilde{\lambda}_1, \dots, \tilde{\lambda}_{D-K}) \end{bmatrix},$$

where $*$ can be any $K \times (D-K)$ matrix. Since $\mathbf{P}^{-1}\mathbf{A}\mathbf{P}$ is upper-triangular, we know that $\mathbf{P}^{-1}\mathbf{A}\mathbf{P}$ has eigenvalues $\hat{\lambda}_1, \dots, \hat{\lambda}_K, \tilde{\lambda}_1, \dots, \tilde{\lambda}_{D-K}$, and so does \mathbf{A} . \square

Theorem H.5. *The eigenvalues of $\mathbf{J}(\mathbf{W}_0)$ can be fully classified into the following 3 types:*

1. $\mu_i + \mu_j$ is an eigenvalue for every $1 \leq i \leq j \leq d - r$, and $\hat{\mathbf{U}}_{ij} := \mathbf{v}_i \mathbf{v}_j^\top + \mathbf{v}_j \mathbf{v}_i^\top$ is an associated left eigenvector.
2. ξ_p is an eigenvalue for every $1 \leq p \leq rd - \frac{r(r-1)}{2}$, and $\tilde{\mathbf{V}}_p := \mathbf{E}_p \mathbf{U}_0^\top + \mathbf{U}_0 \mathbf{E}_p^\top$ is an associated right eigenvector.
3. 0 is an eigenvalue, and any antisymmetric matrix is an associated right eigenvector, which spans a linear space of dimension $\frac{d(d-1)}{2}$.

Proof of Theorem H.5. We first prove each item respectively, and then prove that these are all the eigenvalues of $\mathbf{J}(\mathbf{W}_0)$.

Proof for Item 1. For $\hat{\mathbf{U}}_{ij} = \mathbf{v}_i \mathbf{v}_j^\top + \mathbf{v}_j \mathbf{v}_i^\top$, it is easy to check:

$$\begin{aligned} \text{ac}\{-\nabla f(\mathbf{W}_0), \hat{\mathbf{U}}_{ij}\} &= (\lambda_i + \lambda_j) \hat{\mathbf{U}}_{ij} \\ \hat{\mathbf{U}}_{ij} \mathbf{W}_0 &= \mathbf{0} \\ \mathbf{W}_0 \hat{\mathbf{U}}_{ij} &= \mathbf{0} \end{aligned}$$

So we have

$$\mathbf{J}(\mathbf{W}_0)[\Delta, \hat{\mathbf{U}}_{ij}] = (\lambda_i + \lambda_j) \langle \Delta, \hat{\mathbf{U}}_{ij} \rangle - D^2 f(\mathbf{W}_0)[\Delta, \mathbf{0}] = (\lambda_i + \lambda_j) \langle \Delta, \hat{\mathbf{U}}_{ij} \rangle,$$

which shows that $\hat{\mathbf{U}}_{ij}$ is a left eigenvector associated with eigenvalue $\lambda_i + \lambda_j$.

Proof for Item 2. By definition of eigenvector, we have $-D^2 \mathcal{L}(\mathbf{U}_0)[\mathbf{E}_p] = \xi_p \mathbf{E}_p$, so

$$\xi_p \mathbf{E}_p = -\nabla f(\mathbf{W}_0) \mathbf{E}_p - D^2 f(\mathbf{W}_0)[\mathbf{U}_0 \mathbf{E}_p^\top + \mathbf{E}_p \mathbf{U}_0^\top] \mathbf{U}_0.$$

Right-multiplying both sides by \mathbf{U}_0^\top , we get

$$\begin{aligned} \xi_p \mathbf{E}_p \mathbf{U}_0^\top &= -\nabla f(\mathbf{W}_0) \mathbf{E}_p \mathbf{U}_0^\top - D^2 f(\mathbf{W}_0)[\tilde{\mathbf{V}}_p] \mathbf{W}_0 \\ &= -\nabla f(\mathbf{W}_0)(\mathbf{E}_p \mathbf{U}_0^\top + \mathbf{U}_0 \mathbf{E}_p^\top) - D^2 f(\mathbf{W}_0)[\tilde{\mathbf{V}}_p] \mathbf{W}_0 \\ &= -\nabla f(\mathbf{W}_0) \tilde{\mathbf{V}}_p - D^2 f(\mathbf{W}_0)[\tilde{\mathbf{V}}_p] \mathbf{W}_0, \end{aligned}$$

where the second equality uses the fact that $\nabla f(\mathbf{W}_0) \mathbf{U}_0 = \mathbf{0}$ since \mathbf{U}_0 is a critical point. Taking both sides into $\text{sz}(\cdot)$ gives

$$\begin{aligned} \xi_p \tilde{\mathbf{V}}_p &= -\text{sz}(\nabla f(\mathbf{W}_0) \tilde{\mathbf{V}}_p) - \text{sz}(D^2 f(\mathbf{W}_0)[\tilde{\mathbf{V}}_p] \mathbf{W}_0) \\ &= \mathbf{J}(\mathbf{W}_0)[\tilde{\mathbf{V}}_p], \end{aligned}$$

which proves that $\tilde{\mathbf{V}}_p$ is a right eigenvector associated with eigenvalue ξ_p .

Proof for Item 3. Since $\nabla f(\mathbf{W})$ is symmetric, $\mathbf{g}(\mathbf{W})$ is also symmetric. For any $\Delta = -\Delta^\top$,

$$\mathbf{J}(\mathbf{W}_0)[\Delta] = \mathbf{J}(\mathbf{W}_0)[\Delta^\top] = \mathbf{J}(\mathbf{W}_0)[- \Delta].$$

So $\mathbf{J}(\mathbf{W}_0)[\Delta] = \mathbf{0}$ and Δ is an eigenvector associated with eigenvalue 0.

No other eigenvalues. Let \mathbb{S}_d be the space of symmetric matrices and \mathbb{A}_d be the space of anti-symmetric matrices. It is easy to see that \mathbb{S}_d and \mathbb{A}_d are orthogonal to each other, and \mathbb{S}_d and \mathbb{A}_d are invariant subspaces of $\mathbf{J}(\mathbf{W}_0)[\Delta]$. Let $\mathbf{h} : \mathbb{S}_d \rightarrow \mathbb{S}_d$, $\Delta \mapsto \mathbf{J}(\mathbf{W}_0)[\Delta]$ be the linear operator $\mathbf{J}(\mathbf{W}_0)[\Delta]$ restricted on symmetric matrices. We only need to prove that \mathbf{h} is diagonalizable.

It is easy to see that $\{\hat{\mathbf{U}}_{ij}\}$ are linearly independent to each other and thus spans a subspace of \mathbb{S}_d with dimension $\frac{(d-r)(d-r+1)}{2}$. We can also prove that $\{\tilde{\mathbf{V}}_p\}$ spans a subspace of \mathbb{S}_d with dimension $rd - \frac{r(r-1)}{2}$ by contradiction. Assume to the contrary that there exists scalars α_p for $1 \leq p \leq rd - r(r-1)/2$, not all zero, such that $\sum_{p=1}^{rd-r(r-1)/2} \alpha_p \tilde{\mathbf{V}}_p = \mathbf{0}$. Then $\sum_{p=1}^{rd-r(r-1)/2} \alpha_p \mathbf{E}_p$ is a solution of (28). However, this suggests that $\sum_{p=1}^{rd-r(r-1)/2} \alpha_p \mathbf{E}_p$ lies in the span of $\{\mathbf{E}_p\}_{rd-r(r-1)/2 < p \leq rd}$, which contradicts to the linear independence of $\{\mathbf{E}_p\}_{1 \leq p \leq rd}$.

Note that

$$\frac{(d-r)(d-r+1)}{2} + \left(rd - \frac{r(r-1)}{2} \right) = \frac{d(d+1)}{2} = \dim(\mathbb{S}_d).$$

Also note that $\langle \hat{\mathbf{U}}_{ij}, \tilde{\mathbf{V}}_p \rangle = 2\mathbf{v}_i^\top \mathbf{E}_p \mathbf{U}_0^\top \mathbf{v}_j + 2\mathbf{v}_j^\top \mathbf{E}_p \mathbf{U}_0^\top \mathbf{v}_i = 0$. By Lemma H.4, Items 1 and 2 give all the eigenvalues of \mathbf{h} , and thus Items 1, 2, 3 give all the eigenvalues of $\mathbf{J}(\mathbf{W}_0)$. \square

I PROOFS FOR THE DEPTH-2 CASE

I.1 PROOF FOR THEOREM 5.6

Proof for Theorem 5.6. Since $\mathbf{W}(t)$ is always symmetric, it suffices to study the dynamics of the lower triangle of $\mathbf{W}(t)$. For any symmetric matrix $\mathbf{W} \in \mathbb{S}_d$, let $\text{vec}_{\text{LT}}(\mathbf{W}) \in \mathbb{R}^{\frac{d(d+1)}{2}}$ be the vector consisting of the $\frac{d(d+1)}{2}$ entries of \mathbf{W} in the lower triangle, permuted according to some fixed order.

Let $\mathbf{g}(\mathbf{W})$ be the function defined in (2), which always maps symmetric matrices to symmetric matrices. Let $\tilde{\mathbf{g}} : \mathbb{R}^{\frac{d(d+1)}{2}} \rightarrow \mathbb{R}^{\frac{d(d+1)}{2}}$ be the function such that $\tilde{\mathbf{g}}(\text{vec}_{\text{LT}}(\mathbf{W})) = \text{vec}_{\text{LT}}(\mathbf{g}(\mathbf{W}))$ for any $\mathbf{W} \in \mathbb{S}_d$. For $\mathbf{W}(t)$ evolving with (2), we view $\text{vec}_{\text{LT}}(\mathbf{W}(t))$ as a dynamical system.

$$\frac{d}{dt} \text{vec}_{\text{LT}}(\mathbf{W}(t)) = \tilde{\mathbf{g}}(\text{vec}_{\text{LT}}(\mathbf{W}(t))).$$

By Lemma 5.4, the spaces of symmetric matrices \mathbb{S}_d and antisymmetric matrices \mathbb{A}_d are invariant subspaces of $\mathbf{J}(\mathbf{0})$, and $\left\{ (\mu_i + \mu_j, \mathbf{u}_{1[i]} \mathbf{u}_{1[j]}^\top + \mathbf{u}_{1[j]} \mathbf{u}_{1[i]}^\top) \right\}_{1 \leq i < j \leq d}$ is the set of all the eigenvalues and eigenvectors in the invariant subspace \mathbb{S}_d . Thus, $\tilde{\mu}_1 := 2\mu_1$ and $\tilde{\mu}_2 := \mu_1 + \mu_2$ are the largest and second largest eigenvalues of the Jacobian of $\tilde{\mathbf{g}}(\cdot)$ at $\text{vec}_{\text{LT}}(\mathbf{W}) = \mathbf{0}$, and $\tilde{\mathbf{u}}_1 = \tilde{\mathbf{v}}_1 = \mathbf{u}_1 \mathbf{u}_1^\top$ are the corresponding left and right eigenvectors of the top eigenvalue. Then it is easy to translate Theorem 5.3 to Theorem 5.6. \square

I.2 PROOF FOR THEOREM 5.8

The proof for Theorem 5.8 relies on the following Lemma on the gradient flow around a local minimizer:

Lemma I.1. *If $\bar{\theta}$ is a local minimizer of $\mathcal{L}(\theta)$ and for all $\|\theta - \bar{\theta}\|_2 \leq r$, θ satisfies Łojasiewicz inequality:*

$$\|\nabla \mathcal{L}(\theta)\|_2 \geq c (\mathcal{L}(\theta) - \mathcal{L}(\bar{\theta}))^\mu$$

for some $\mu \in [1/2, 1)$, then the gradient flow $\theta(t) = \phi(\theta_0, t)$ converges to a point θ_∞ near $\bar{\theta}$ if θ_0 is close enough to $\bar{\theta}$, and the distance can be bounded by $\|\theta_\infty - \bar{\theta}\|_2 = O(\|\theta_0 - \bar{\theta}\|_2^{2(1-\mu)})$.

Proof. For every $t \geq 0$, if $\|\boldsymbol{\theta}(t) - \bar{\boldsymbol{\theta}}\|_2 \leq r$,

$$\begin{aligned} \frac{d}{dt} (\mathcal{L}(\boldsymbol{\theta}(t)) - \mathcal{L}(\bar{\boldsymbol{\theta}}))^{1-\mu} &= (1-\mu) (\mathcal{L}(\boldsymbol{\theta}(t)) - \mathcal{L}(\bar{\boldsymbol{\theta}}))^{-\mu} \cdot \left\langle \nabla \mathcal{L}, \frac{d\boldsymbol{\theta}}{dt} \right\rangle \\ &= -(1-\mu) (\mathcal{L}(\boldsymbol{\theta}(t)) - \mathcal{L}(\bar{\boldsymbol{\theta}}))^{-\mu} \cdot \|\nabla \mathcal{L}\|_2 \cdot \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|_2 \\ &\leq -(1-\mu)c \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|_2. \end{aligned}$$

Therefore, $\|\boldsymbol{\theta}(t) - \boldsymbol{\theta}_0\|_2 \leq \int_0^t \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|_2 dt \leq \frac{1}{(1-\mu)c} \mathcal{L}(\boldsymbol{\theta}_0)^{1-\mu} = O(\|\boldsymbol{\theta}_0 - \bar{\boldsymbol{\theta}}\|_2^{2(1-\mu)})$. If we choose $\|\boldsymbol{\theta}(t) - \bar{\boldsymbol{\theta}}\|_2$ small enough, then $\|\boldsymbol{\theta}(t) - \bar{\boldsymbol{\theta}}\|_2 \leq \|\boldsymbol{\theta}(t) - \boldsymbol{\theta}_0\|_2 + \|\boldsymbol{\theta}_0 - \bar{\boldsymbol{\theta}}\|_2 = O(\|\boldsymbol{\theta}_0 - \bar{\boldsymbol{\theta}}\|_2^{2(1-\mu)}) < r$, and thus $\int_0^{+\infty} \left\| \frac{d\boldsymbol{\theta}}{dt} \right\|_2 dt$ is convergent and finite. This implies that $\boldsymbol{\theta}_\infty := \lim_{t \rightarrow +\infty} \boldsymbol{\theta}(t)$ exists and $\|\boldsymbol{\theta}_\infty - \bar{\boldsymbol{\theta}}\|_2 = O(\|\boldsymbol{\theta}_0 - \bar{\boldsymbol{\theta}}\|_2^{2(1-\mu)})$. \square

Proof for Theorem 5.8. Since $\mathbf{W}_1^G(t) \in \mathbb{S}_{d, \leq 1}^+$ satisfies (2), there exists $\mathbf{u}(t) \in \mathbb{R}^d$ such that $\mathbf{u}(t)\mathbf{u}(t)^\top = \mathbf{W}_1^G(t)$ and $\mathbf{u}(t)$ satisfies (1), i.e., $\frac{d\mathbf{u}}{dt} = -\nabla \mathcal{L}(\mathbf{u})$, where $\mathcal{L} : \mathbb{R}^d \rightarrow \mathbb{R}, \mathbf{u} \mapsto \frac{1}{2}f(\mathbf{u}\mathbf{u}^\top)$. If $\mathbf{W}_1^G(t)$ does not diverge to infinity, then so does $\mathbf{u}(t)$. This implies that there is a limit point $\bar{\mathbf{u}}$ of the set $\{\mathbf{u}(t) : t \geq 0\}$.

Let $\mathcal{U} := \{\mathbf{u} : \mathcal{L}(\mathbf{u}) \geq \mathcal{L}(\bar{\mathbf{u}})\}$. Since $\mathcal{L}(\mathbf{u}(t))$ is non-increasing, we have $\mathbf{u}(t) \in \mathcal{U}$ for all t . Note that $\bar{\mathbf{u}}$ is a local minimizer of $\mathcal{L}(\cdot)$ in \mathcal{U} . By analyticity of $f(\cdot)$, Łojasiewicz inequality holds for $\mathcal{L}(\cdot)$ around $\bar{\mathbf{u}}$ (Łojasiewicz, 1965). Applying Lemma I.1 for \mathcal{L} restricted on \mathcal{U} , we know that if $\mathbf{u}(t_0)$ is sufficiently close to $\bar{\mathbf{u}}$, the remaining length of the trajectory of $\mathbf{u}(t)$ ($t \geq t_0$) is finite and thus $\lim_{t \rightarrow +\infty} \mathbf{u}(t)$ exists. As $\bar{\mathbf{u}}$ is a limit point, this limit can only be $\bar{\mathbf{u}}$. Therefore, $\bar{\mathbf{W}}_1 := \lim_{t \rightarrow +\infty} \mathbf{W}_1^G(t) = \bar{\mathbf{u}}\bar{\mathbf{u}}^\top$ exists.

If $\bar{\mathbf{W}}_1$ is a minimizer of $f(\cdot)$, $\bar{\mathbf{U}} = (\bar{\mathbf{u}}, \mathbf{0}, \dots, \mathbf{0}) \in \mathbb{R}^{d \times d}$ is also a minimizer of $\mathcal{L} : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}, \mathbf{U} \mapsto \frac{1}{2}f(\mathbf{U}\mathbf{U}^\top)$. By analyticity of $f(\cdot)$, Łojasiewicz inequality holds for $\mathcal{L}(\cdot)$ around $\bar{\mathbf{U}}$. For every $\epsilon > 0$, we can always find a time t_ϵ such that $\|\mathbf{u}(t_\epsilon) - \bar{\mathbf{u}}\|_2 \leq \epsilon/2$. On the other hand, by Theorem 5.6, there exists a number α_ϵ such that for every $\alpha < \alpha_\epsilon$,

$$\|\phi(\mathbf{W}_\alpha, T(\mathbf{W}_\alpha) + t_\epsilon) - \mathbf{W}_1^G(t_\epsilon)\|_2 \leq \epsilon/2, \quad \text{where } T(\mathbf{W}) := \frac{1}{2\mu_1} \log \frac{1}{\langle \mathbf{W}, \mathbf{u}_1 \mathbf{u}_1^\top \rangle}.$$

Combining these together we have $\|\phi(\mathbf{W}_\alpha, T(\mathbf{W}_\alpha) + t_\epsilon) - \bar{\mathbf{W}}_1\|_2 \leq \epsilon$.

It is easy to construct a factorization $\phi(\mathbf{W}_\alpha, T(\mathbf{W}_\alpha) + t_\epsilon) := \mathbf{U}_{\alpha, \epsilon} \mathbf{U}_{\alpha, \epsilon}^\top$ such that $\|\mathbf{U}_{\alpha, \epsilon} - \bar{\mathbf{U}}\|_2 = O(\epsilon)$, e.g., we can find an arbitrary factorization and then right-multiply an orthogonal matrix so that the row vector with the largest norm aligns with the direction of $\bar{\mathbf{u}}$. Applying Lemma I.1, we know that gradient flow starting with $\mathbf{U}_{\alpha, \epsilon}$ converges to a point that is only $O(\epsilon^{2(1-\mu)})$ far from $\bar{\mathbf{u}}$. So we have

$$\left\| \lim_{t \rightarrow +\infty} \phi(\mathbf{W}_\alpha, T(\mathbf{W}_\alpha) + t) - \bar{\mathbf{W}}_1 \right\|_2 = O(\epsilon^{2(1-\mu)}).$$

Taking $\epsilon \rightarrow 0$ complete the proof. \square

I.3 PROOF FOR THEOREM 5.11

Theorem I.2. Let $\bar{\mathbf{W}}$ be a critical point of (2) satisfying that $\bar{\mathbf{W}}$ is a local minimizer of $f(\cdot)$ in $\mathbb{S}_{d, \leq r}^+$ for some $r \geq 1$ but not a minimizer in \mathbb{S}_d^+ . Let $-\nabla f(\bar{\mathbf{W}}) = \sum_{i=1}^d \mu_i \mathbf{v}_i \mathbf{v}_i^\top$ be the eigendecomposition of $-\nabla f(\bar{\mathbf{W}})$. If $\mu_1 > \mu_2$, the following limit exists and is a solution of (2).

$$\mathbf{W}^G(t) := \lim_{\epsilon \rightarrow 0} \phi \left(\bar{\mathbf{W}} + \epsilon \mathbf{v}_1 \mathbf{v}_1^\top, \frac{1}{2\mu_1} \log \frac{1}{\epsilon} + t \right).$$

For $\{\mathbf{W}_\alpha\} \subseteq \mathbb{S}_d^+$, if there exists time $T_\alpha \in \mathbb{R}$ for every α so that $\phi(\mathbf{W}_\alpha, T_\alpha)$ converges to $\bar{\mathbf{W}}$ with positive alignment with the top principal component $\mathbf{v}_1 \mathbf{v}_1^\top$ as $\alpha \rightarrow 0$, then $\forall t \in \mathbb{R}$,

$$\lim_{\alpha \rightarrow 0} \phi \left(\mathbf{W}_\alpha, T_\alpha + \frac{1}{2\mu_1} \log \frac{1}{\langle \phi(\mathbf{W}_\alpha, T_\alpha), \mathbf{v}_1 \mathbf{v}_1^\top \rangle} + t \right) = \mathbf{W}^G(t).$$

Moreover, there exists a constant $C > 0$ such that

$$\left\| \phi \left(\mathbf{W}_\alpha, T_\alpha + \frac{1}{2\mu_1} \log \frac{1}{\langle \phi(\mathbf{W}_\alpha, T_\alpha), \mathbf{v}_1 \mathbf{v}_1^\top \rangle} + t \right) - \mathbf{W}^G(t) \right\|_{\mathbb{F}} \leq C \|\phi(\mathbf{W}_\alpha, T_\alpha)\|_{\mathbb{F}}^{\frac{\tilde{\gamma}}{2\mu_1 + \tilde{\gamma}}}$$

for every sufficiently small α , where $\tilde{\gamma} := 2\mu_1 - \max\{\mu_1 + \mu_2, 0\}$.

Proof. Following Appendix I.1, we view $\text{vec}_{\text{LT}}(\mathbf{W}(t))$ as a dynamical system.

$$\frac{d}{dt} \text{vec}_{\text{LT}}(\mathbf{W}(t)) = \tilde{\mathbf{g}}(\text{vec}_{\text{LT}}(\mathbf{W}(t))).$$

Let $\overline{\mathbf{W}} = \overline{\mathbf{U}}\overline{\mathbf{U}}^\top$ be a factorization of $\overline{\mathbf{W}}$, where $\overline{\mathbf{U}} \in \mathbb{R}^{d \times r}$. Since $\overline{\mathbf{W}}$ is a local minimizer of $f(\cdot)$ in $\mathbb{S}_{d, \leq r}^+$, $\overline{\mathbf{U}}$ is also a local minimizer of $\mathcal{L} : \mathbb{R}^{d \times r} \rightarrow \mathbb{R}, \mathbf{U} \mapsto \frac{1}{2}f(\mathbf{U}\mathbf{U}^\top)$. Since $\overline{\mathbf{W}}$ is not a minimizer of $f(\cdot)$ in \mathbb{S}_d^+ , by Lemma H.1, $\overline{\mathbf{U}}$ is full-rank. By Theorem H.5, $\mathbf{J}(\overline{\mathbf{W}})$ has eigenvalues $\mu_i + \mu_j, \xi_p, 0$. By a similar argument as in Appendix I.1, the Jacobian of $\tilde{\mathbf{g}}$ at $\text{vec}_{\text{LT}}(\overline{\mathbf{W}})$ has eigenvalues $\mu_i + \mu_j, \xi_p$.

Since $\overline{\mathbf{U}}$ is a local minimizer, $\xi_p \leq 0$ for all p . If $\mu_1 > \mu_2$, then $2\mu_1$ is the unique largest eigenvalue, and Theorem H.5 shows that $\text{vec}_{\text{LT}}(\mathbf{v}_1 \mathbf{v}_1^\top)$ is a left eigenvector associated with $2\mu_1$. The eigenvalue gap $\tilde{\gamma} := 2\mu_1 - \max\{\mu_1 + \mu_2, \max\{\xi_p : 1 \leq p \leq rd - \frac{r(r-1)}{2}\}\} \geq 2\mu_1 - \max\{\mu_1 + \mu_2, 0\}$.

Also note that $\langle \phi(\mathbf{W}_\alpha, T_\alpha) - \overline{\mathbf{W}}, \mathbf{v}_1 \mathbf{v}_1^\top \rangle = \langle \phi(\mathbf{W}_\alpha, T_\alpha), \mathbf{v}_1 \mathbf{v}_1^\top \rangle$ because $\langle \overline{\mathbf{W}}, \mathbf{v}_1 \mathbf{v}_1^\top \rangle = 0$ by (27). If $\phi(\mathbf{W}_\alpha, T_\alpha)$ converges to $\overline{\mathbf{W}}$ as $\alpha \rightarrow 0$, then it has positive alignment with $\mathbf{v}_1 \mathbf{v}_1^\top$ iff $\liminf_{\alpha \rightarrow 0} \frac{\langle \phi(\mathbf{W}_\alpha, T_\alpha), \mathbf{v}_1 \mathbf{v}_1^\top \rangle}{\|\phi(\mathbf{W}_\alpha, T_\alpha) - \overline{\mathbf{W}}\|_{\mathbb{F}}} > 0$. Then it is easy to translate Theorem 5.3 to Theorem I.2. \square

I.4 GRADIENT FLOW ONLY FINDS MINIMIZERS (PROOF FOR THEOREM 5.10)

The proof for Theorem 5.10 is based on the following two theorems from the literature.

Theorem I.3 (Theorem 3.1 in Du and Lee 2018). *Let $f : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$ be a \mathcal{C}^2 convex function. Then $\mathcal{L} : \mathbb{R}^{d \times k} \rightarrow \mathbb{R}, \mathcal{L}(\mathbf{U}) = f(\mathbf{U}\mathbf{U}^\top), k \geq d$ satisfies that (1). Every local minimizer of \mathcal{L} is also a global minimizer; (2). All saddles are strict. Here saddles denote those stationary points whose hessian are not positive semi-definite (thus including local maximizers).* ²

Theorem I.4 (Theorem 2 in Lee et al. 2017). *Let \mathbf{g} be a \mathcal{C}^1 mapping from $\mathcal{X} \rightarrow \mathcal{X}$ and $\det(D\mathbf{g}(x)) \neq 0$ for all $x \in \mathcal{X}$. Then the set of initial points that converge to an unstable fixed point has measure zero, $\mu(\{x_0 : \lim_{k \rightarrow \infty} \mathbf{g}^k(x_0) \in \mathcal{A}_\mathbf{g}^*\}) = 0$, where $\mathcal{A}_\mathbf{g}^* = \{x : \mathbf{g}(x) = x, \max_i |\lambda_i(D\mathbf{g}(x))| > 1\}$.*

Theorem I.5 (GF only finds minimizers, a continuous analog of Theorem I.4). *Let $\mathbf{f} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be a \mathcal{C}^1 -smooth function, and $\phi : \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^d$ be the solution of the following differential equation,*

$$\frac{d\phi(\mathbf{x}, t)}{dt} = \mathbf{f}(\phi(\mathbf{x}, t)), \quad \phi(\mathbf{x}, 0) = \mathbf{x}, \quad \forall \mathbf{x} \in \mathbb{R}^d, t \in \mathbb{R}.$$

Then the set of initial points that converge to a unstable critical point has measure zero, $\mu\left(\left\{x_0 : \lim_{t \rightarrow \infty} \phi(x_0, t) \in \mathcal{U}_\mathbf{f}^\right\}\right) = 0$, where $\mathcal{U}_\mathbf{f}^* = \{x : \mathbf{f}(x) = \mathbf{0}, \lambda_1(D\mathbf{f}(x)) > 0\}$ and $D\mathbf{f}$ is the Jacobian matrix of \mathbf{f} .*

Proof of Theorem I.5. By Theorem 1 in Section 2.3, Perko (2013), we know $\phi(\cdot, \cdot)$ is \mathcal{C}^1 -smooth for both x, t . We let $\mathbf{g}(x) = \phi(x, 1)$, then we know $\mathbf{g}^{-1}(x) = \phi(x, -1)$ and both $\mathbf{g}, \mathbf{g}^{-1}$ are \mathcal{C}^1 -smooth. Note that $D\mathbf{g}^{-1}(x)$ is the inverse matrix of $D\mathbf{g}(x)$. So both of the two matrices are invertible. Thus we can apply Theorem I.4 and we know $\mu(\{x_0 : \lim_{k \rightarrow \infty} \mathbf{g}^k(x_0) \in \mathcal{A}_\mathbf{g}^*\}) = 0$.

Note that if $\lim_{t \rightarrow \infty} \phi(x, t)$ exists, then $\lim_{k \rightarrow \infty} \mathbf{g}^k(x) = \lim_{t \rightarrow \infty} \phi(x, t)$. It remains to show that $\mathcal{U}_\mathbf{f}^* \subseteq \mathcal{A}_\mathbf{g}^*$. For $\mathbf{f}(x_0) = \mathbf{0}$, we have $\phi(x_0, t) = x_0$ and thus $\mathbf{g}(x_0) = x_0$. Now it suffices to prove

²Though the original theorem is proven for convex functions of form $\sum_{i=1}^n \ell(\mathbf{x}_i \mathbf{U} \mathbf{U}^\top \mathbf{x}_i^\top, y_i)$, where $\ell(\cdot, \cdot)$ is \mathcal{C}^2 convex for its first variable. By scrutinizing their proof, we can see the assumption can be relaxed to f is \mathcal{C}^2 convex.

that $\lambda_1(D\mathbf{g}(\mathbf{x}_0)) > 1$. For every $t \in [0, 1]$, by Corollary of Theorem 1 in Section 2.3, [Perko \(2013\)](#), we have $\frac{\partial}{\partial t} D\phi(\mathbf{x}, t) = D\mathbf{f}(\phi(\mathbf{x}, t))D\phi(\mathbf{x}, t)$, $\forall \mathbf{x}, t$. Thus,

$$\frac{\partial}{\partial t} D\phi(x_0, t) = D\mathbf{f}(\phi(x_0, t))D\phi(x_0, t) = D\mathbf{f}(x_0)D\phi(x_0, t).$$

Solving this ODE gives $D\mathbf{g}(\mathbf{x}_0) = D\phi(\mathbf{x}, 1) = e^{D\mathbf{f}(x_0)}D\phi(\mathbf{x}, 0) = e^{D\mathbf{f}(x_0)}$, where the last equality is due to $D\phi(\mathbf{x}, 0) \equiv \mathbf{I}$, $\forall \mathbf{x}$. Combining this with $\lambda_1(D\mathbf{f}(x_0)) > 0$, we have $\lambda_1(D\mathbf{g}(x_0)) > 1$.

Thus we have $\mathcal{U}_{\mathbf{f}}^* := \{\mathbf{x}_0 : \mathbf{f}(\mathbf{x}_0) = \mathbf{0}, \lambda_1(D\mathbf{f}(\mathbf{x}_0)) > 0\} \subseteq \mathcal{A}_{\mathbf{g}}^*$, which implies that $\{\mathbf{x}_0 : \lim_{t \rightarrow \infty} \phi(\mathbf{x}_0, t) \in \mathcal{U}^*\} \subseteq \{\mathbf{x}_0 : \lim_{k \rightarrow \infty} \mathbf{g}^k(\mathbf{x}_0) \in \mathcal{A}_{\mathbf{g}}^*\}$ \square

Theorem 5.10. *Let $f : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$ be a convex \mathcal{C}^2 -smooth function. (1). All stationary points of $\mathcal{L} : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$, $\mathcal{L}(\mathbf{U}) = \frac{1}{2}f(\mathbf{U}\mathbf{U}^\top)$ are either strict saddles or global minimizers; (2). For any random initialization, GF (1) converges to strict saddles of $\mathcal{L}(\mathbf{U})$ with probability 0.*

Proof of Theorem 5.10. For (1), by Theorem I.3, we immediately know all the stationary points of $\mathcal{L}(\cdot)$ are either global minimizers or strict saddles. (2) is just a direct consequence of Theorem I.5 by setting \mathbf{f} in the above proof to $-\nabla \mathcal{L}$. \square

J EQUIVALENCE BETWEEN GF AND GLRL

In this section we elaborate on the theoretical evidence that GF and GLRL are equivalent generically, including the case where GLRL does not end in the first phase. The word ‘‘generically’’ used when we want to assume one of the following regularity conditions:

1. We want to assume that GF converges to a local minimizer (i.e., GF does not get stuck on saddle points);
2. We want to assume that the top eigenvalue $\lambda_1(-\nabla f(\mathbf{W}))$ is unique for a critical point \mathbf{W} of (2) that is not a minimizer of $f(\cdot)$ in \mathbb{S}_d^+ ;
3. We want to assume that a convergent sequence of PSD matrices $\mathbf{W}_\alpha \rightarrow \overline{\mathbf{W}}$ has positive alignment with $\mathbf{v}\mathbf{v}^\top$ for some fixed vector \mathbf{v} with $\langle \overline{\mathbf{W}}, \mathbf{v}\mathbf{v}^\top \rangle = \mathbf{0}$, i.e., for a convergent sequence of PSD matrices $\mathbf{W}_\alpha \rightarrow \overline{\mathbf{W}}$, it holds for sure that $\liminf_{\alpha \rightarrow 0} \left\langle \frac{\mathbf{W}_\alpha - \overline{\mathbf{W}}}{\|\mathbf{W}_\alpha - \overline{\mathbf{W}}\|_{\mathbb{F}}}, \mathbf{v}\mathbf{v}^\top \right\rangle = \liminf_{\alpha \rightarrow 0} \frac{\langle \mathbf{W}_\alpha, \mathbf{v}\mathbf{v}^\top \rangle}{\|\mathbf{W}_\alpha - \overline{\mathbf{W}}\|_{\mathbb{F}}} \geq 0$, and we further assume that the inequality is strict generically.

Theorem I.2 uncovers how GF with infinitesimal initialization generically behaves. Let $\overline{\mathbf{W}}_0 := \mathbf{0}$. For every $r \geq 1$, if $\overline{\mathbf{W}}_{r-1}$ is a local minimizer in $\mathbb{S}_{d, \leq r-1}^+$ but not a minimizer in \mathbb{S}_d^+ , then $\lambda_1(-\nabla f(\overline{\mathbf{W}}_{r-1})) > 0$ by Lemma C.2. Generically, the top eigenvalue $\lambda_1(-\nabla f(\overline{\mathbf{W}}_{r-1}))$ should be unique, i.e., $\lambda_1(-\nabla f(\overline{\mathbf{W}}_{r-1})) > \lambda_2(-\nabla f(\overline{\mathbf{W}}_{r-1}))$. This enables us to apply Theorem I.2 and deduce that the limiting trajectory

$$\mathbf{W}_r^{\text{G}}(t) := \lim_{\epsilon \rightarrow 0} \phi \left(\overline{\mathbf{W}}_{r-1} + \epsilon \mathbf{u}_r \mathbf{u}_r^\top, \frac{1}{2\lambda_1(-\nabla f(\overline{\mathbf{W}}_{r-1}))} \log \frac{1}{\epsilon} + t \right)$$

exists, where \mathbf{u}_r is the top eigenvector of $-\nabla f(\overline{\mathbf{W}}_{r-1})$. This $\mathbf{W}_r^{\text{G}}(\cdot)$ is exactly the trajectory of GLRL in phase r as $\epsilon \rightarrow 0$.

Note that $\mathbf{W}_r^{\text{G}}(\cdot)$ corresponds to a trajectory of GF minimizing $\mathcal{L}(\cdot)$ in $\mathbb{R}^{d \times r}$, which should generically converge to a local minimizer of $\mathcal{L}(\cdot)$ in $\mathbb{R}^{d \times r}$. This means the limit $\overline{\mathbf{W}}_r := \lim_{t \rightarrow +\infty} \mathbf{W}_r^{\text{G}}(t)$ should generically be a local minimizer of $f(\cdot)$ in $\mathbb{S}_{d, \leq r}^+$. If $\overline{\mathbf{W}}_r$ is further a minimizer in \mathbb{S}_d^+ , then $\lambda_1(-\nabla f(\overline{\mathbf{W}}_r)) \leq 0$ and GLRL exits with $\overline{\mathbf{W}}_r$; otherwise GLRL enters phase $r + 1$.

If GF aligns well with GLRL in the beginning of phase r (defined below), then by Theorem I.2, as $\alpha \rightarrow 0$, the minimum distance from GF to $\mathbf{W}_r^{\text{G}}(t)$ converges to 0 for every $t \in \mathbb{R}$. Therefore, GF can get arbitrarily close to the r -th critical point $\overline{\mathbf{W}}_r$ of GLRL, i.e., there exists a suitable choice $T_\alpha^{(r)}$ so that $\lim_{\alpha \rightarrow 0} \phi(\mathbf{W}_\alpha, T_\alpha^{(r)}) = \overline{\mathbf{W}}_r$. Note that $\langle \overline{\mathbf{W}}_r, \mathbf{u}_r \mathbf{u}_r^\top \rangle = 0$ by (27) and thus

$\liminf_{\alpha \rightarrow 0} \left\langle \frac{\phi(\mathbf{W}_\alpha, T_\alpha^{(r)}) - \overline{\mathbf{W}}_r}{\|\phi(\mathbf{W}_\alpha, T_\alpha^{(r)}) - \overline{\mathbf{W}}_r\|_F}, \mathbf{u}_r \mathbf{u}_r^\top \right\rangle = \liminf_{\alpha \rightarrow 0} \frac{\langle \phi(\mathbf{W}_\alpha, T_\alpha^{(r)}), \mathbf{u}_r \mathbf{u}_r^\top \rangle}{\|\phi(\mathbf{W}_\alpha, T_\alpha^{(r)}) - \overline{\mathbf{W}}_r\|_F} \geq 0$. Generically, there should exist a suitable choice of $T_\alpha^{(r)}$ so that $\phi(\mathbf{W}_\alpha, T_\alpha^{(r)})$ not only converges to $\overline{\mathbf{W}}_r$ but also has positive alignment with $\mathbf{u}_r \mathbf{u}_r^\top$, that is, GF should generically align well with GLRL in the beginning of phase $r + 1$.

Definition J.1. We say that GF aligns well with GLRL in the beginning of phase r if there exists $T_\alpha^{(r)}$ for every $\alpha > 0$ such that $\phi(\mathbf{W}_\alpha, T_\alpha^{(r)})$ converges to $\overline{\mathbf{W}}_{r-1}$ with positive alignment with $\mathbf{u}_r \mathbf{u}_r^\top$ as $\alpha \rightarrow 0$.

If the initialization satisfies that \mathbf{W}_α converges to $\mathbf{0}$ with positive alignment with $\mathbf{u}_1 \mathbf{u}_1^\top$ as $\alpha \rightarrow 0$, then GF aligns well with GLRL in the beginning of phase 1, which can be seen by taking $T_\alpha^{(1)} = 0$. Now assume that GF aligns well with GLRL in the beginning of phase $r - 1$, then the above argument shows that GF should generically align well with GLRL in the beginning of phase r , if GLRL does not exit in phase $r - 1$. In the other case, we can use a similar argument as in Theorem 5.8 to show that GF converges to a solution near the minimizer $\overline{\mathbf{W}}_r$ of $f(\cdot)$ as $t \rightarrow \infty$, and the distance between the solution and $\overline{\mathbf{W}}_r$ converges to 0 as $\alpha \rightarrow 0$. By this induction we prove that GF with infinitesimal initialization is equivalent to GLRL generically.

K PROOFS FOR DEEP MATRIX FACTORIZATION

K.1 PRELIMINARY LEMMAS

Lemma K.1. *If $\mathbf{W}(0) \succeq \mathbf{0}$, then $\mathbf{W}(t) \succeq \mathbf{0}$ and $\text{rank}(\mathbf{W}(t)) = \text{rank}(\mathbf{W}(0))$ for all t .*

Proof. Note that we can always find a set of balanced $\mathbf{U}_i(t)$, such that $\mathbf{U}_1(t) \dots \mathbf{U}_L(t) = \mathbf{W}(t)$, $d_2 = d_3 = \dots = d_L = \text{rank}(\mathbf{W}(t))$ and write the dynamics of $\mathbf{W}(t)$ in the space of $\{\mathbf{U}_i\}_{i=1}^L$. Thus it is clear that for all t' , $\text{rank}(\mathbf{W}(t')) \leq \text{rank}(\mathbf{W}(t))$. We can apply the same argument for t' and we know $\text{rank}(\mathbf{W}(t)) \leq \text{rank}(\mathbf{W}(t'))$. Thus $\text{rank}(\mathbf{W}(t))$ is constant over time, and we denote it by k . Since eigenvalues are continuous matrix functions, and $\forall t, \lambda_i(\mathbf{W}(t)), i \in [k] \neq 0$. Thus they cannot change their signs and it must hold that $\mathbf{W}(t) \succeq \mathbf{0}$. \square

Lemma K.2. *$\forall a, b, P \in \mathbb{R}$, if $a > b \geq 0, P \geq 1$, then $\frac{a^P - b^P}{a - b} \leq P a^{P-1}$.*

Proof. Let $f(x) = P(1 - x) - (1 - x^P)$. Since $f'(x) = -P + P x^{P-1} < 0$ for all $x \in [0, 1)$, $f(x) \geq f(0) = 0$. Then substituting x by $\frac{b}{a}$ completes the proof. \square

Recall we use $D\mathbf{F}(\mathbf{N})[\mathbf{M}]$ to denote the directional derivative along \mathbf{M} of \mathbf{F} at \mathbf{N} .

Lemma K.3. *Let $\mathbf{F} : \mathbb{S}_d^+ \rightarrow \mathbb{S}_d^+, \mathbf{M} \mapsto \mathbf{M}^P$, where $P \geq 1$ and $P \in \mathbb{Q}$. Then $\forall \mathbf{M}, \mathbf{N} \succeq \mathbf{0}$,*

$$\|D\mathbf{F}(\mathbf{N})[\mathbf{M}]\|_F \leq P \|\mathbf{N}\|_2^{P-1} \|\mathbf{M}\|_F,$$

where $D\mathbf{F}(\mathbf{N})[\mathbf{M}] := \lim_{t \rightarrow 0} \frac{\mathbf{F}(\mathbf{N} + t\mathbf{M}) - \mathbf{F}(\mathbf{N})}{t}$ is the directional derivative of \mathbf{F} along \mathbf{M} .

Proof. Let $\mathbf{N} = \mathbf{U}\Sigma\mathbf{U}^\top$, where $\mathbf{U}\mathbf{U}^\top = \mathbf{I}$ and $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_d)$. Note that $\mathbf{F}(\mathbf{U}\mathbf{M}\mathbf{U}^\top) = \mathbf{U}\mathbf{F}(\mathbf{M})\mathbf{U}^\top$ for any $\mathbf{M} \in \mathbb{S}_d^+$. Then we have

$$\begin{aligned} \|D\mathbf{F}(\mathbf{N})[\mathbf{M}]\|_F &= \lim_{t \rightarrow 0} \frac{\|\mathbf{F}(\mathbf{N} + t\mathbf{M}) - \mathbf{F}(\mathbf{N})\|_F}{t} \\ &= \lim_{t \rightarrow 0} \frac{\|\mathbf{F}(\Sigma + t\mathbf{U}^\top \mathbf{M} \mathbf{U}) - \mathbf{F}(\Sigma)\|_F}{t} \\ &= \|D\mathbf{F}(\Sigma)[\mathbf{U}^\top \mathbf{M} \mathbf{U}]\|_F. \end{aligned}$$

Therefore, it suffices to prove the lemma for the case where \mathbf{N} is diagonal, i.e., $\mathbf{N} = \Sigma$.

Assume $P = \frac{q}{p}$, where $p, q \in \mathbb{N}$ and $q \geq p > 0$. Define $\mathbf{G}(\mathbf{N}) = \mathbf{N}^{\frac{1}{p}}$. Then $\mathbf{G}(\boldsymbol{\Sigma})^p = \boldsymbol{\Sigma}$. Taking directional derivative on both sides along direction \mathbf{M} , we have

$$\sum_{i=1}^p \mathbf{G}(\boldsymbol{\Sigma})^{i-1} D\mathbf{G}(\boldsymbol{\Sigma})[\mathbf{M}]\mathbf{G}(\boldsymbol{\Sigma})^{p-1} = \mathbf{M},$$

So we have

$$[D\mathbf{G}(\boldsymbol{\Sigma})[\mathbf{M}]]_{ij} = \frac{m_{ij}}{\sum_{k=1}^p \sigma_i^{\frac{k-1}{p}} \sigma_j^{\frac{p-k}{p}}}.$$

Let $\mathbf{H}(\mathbf{G}) = \mathbf{G}^q$. With the same argument, we know

$$[D\mathbf{H}(\mathbf{G}(\boldsymbol{\Sigma}))[\mathbf{M}]]_{ij} = m_{ij} \sum_{k=1}^q \sigma_i^{\frac{k-1}{q}} \sigma_j^{\frac{q-k}{q}}.$$

Note that $\mathbf{H}(\mathbf{G}(\boldsymbol{\Sigma})) = \mathbf{F}(\boldsymbol{\Sigma})$. By chain rule, we have

$$D\mathbf{F}(\boldsymbol{\Sigma})[\mathbf{M}] = D\mathbf{H}(\mathbf{G}(\boldsymbol{\Sigma})) [D\mathbf{G}(\boldsymbol{\Sigma})[\mathbf{M}]].$$

That is,

$$[D\mathbf{F}(\boldsymbol{\Sigma})[\mathbf{M}]]_{ij} = m_{ij} \frac{\sum_{k=1}^q \sigma_i^{\frac{k-1}{q}} \sigma_j^{\frac{q-k}{q}}}{\sum_{k=1}^p \sigma_i^{\frac{k-1}{p}} \sigma_j^{\frac{p-k}{p}}}.$$

When $\sigma_i = \sigma_j$, clearly $[D\mathbf{F}(\boldsymbol{\Sigma})[\mathbf{M}]]_{ij} = m_{ij} \cdot \frac{q}{p} \cdot \sigma_i^{\frac{q-p}{p}} = P m_{ij} \sigma_i^{P-1}$. Otherwise, we assume WLOG that $\sigma_i > \sigma_j$, we multiply $\sigma_i - \sigma_j$ to both numerator and denominator and we have

$$|[D\mathbf{F}(\boldsymbol{\Sigma})[\mathbf{M}]]_{ij}| = |m_{ij}| \frac{\sigma_i^P - \sigma_j^P}{\sigma_i - \sigma_j} \leq |m_{ij}| P \sigma_i^{P-1} \leq |m_{ij}| P \|\boldsymbol{\Sigma}\|_2^{P-1}.$$

where the first inequality is by Lemma K.2. Thus we conclude the proof. \square

Lemma K.4. For any $\mathbf{A}, \mathbf{B} \succeq \mathbf{0}$ and $P \in \mathbb{R}, P \geq 1$,

$$\|\mathbf{A}^P - \mathbf{B}^P\|_{\mathbb{F}} \leq P \|\mathbf{A} - \mathbf{B}\|_{\mathbb{F}} \max \left\{ \|\mathbf{A}\|_2^{P-1}, \|\mathbf{B}\|_2^{P-1} \right\}.$$

Proof. Since both sides are continuous in P and \mathbb{Q} is dense in \mathbb{R} , it suffices to prove the lemma for $P \in \mathbb{Q}$. Let $\rho := \max \{ \|\mathbf{A}\|_2, \|\mathbf{B}\|_2 \}$ and $\mathbf{F}(\mathbf{M}) = \mathbf{M}^P$. Define $\mathbf{N} : [0, 1] \rightarrow \mathbb{S}_d^+$, $\mathbf{N}(t) = (1-t)\mathbf{A} + t\mathbf{B}$, we have

1. $\|\mathbf{N}(t)\|_2 \leq \rho$, since $\|\cdot\|_2$ is convex.
2. $\|D\mathbf{F}(\mathbf{N}(t))[\mathbf{B} - \mathbf{A}]\|_{\mathbb{F}} \leq P \|\mathbf{N}(t)\|_2^{P-1} \|\mathbf{B} - \mathbf{A}\|_{\mathbb{F}}$ by Lemma K.3.

Therefore,

$$\begin{aligned} \|\mathbf{F}(\mathbf{N}(1)) - \mathbf{F}(\mathbf{N}(0))\|_{\mathbb{F}} &\leq \int_0^1 \left\| \frac{d\mathbf{F}(\mathbf{N}(t))}{dt} \right\|_{\mathbb{F}} dt \\ &= \int_{t=0}^1 \|D\mathbf{F}(\mathbf{N}(t))[\mathbf{B} - \mathbf{A}]\|_{\mathbb{F}} dt \\ &\leq P \|\mathbf{A} - \mathbf{B}\|_{\mathbb{F}} \rho^{P-1}, \end{aligned}$$

which completes the proof. \square

For a locally Lipschitz function $f(\cdot)$, the Clarke subdifferential (Clarke, 1975; 1990; Clarke et al., 2008) of f at any point \mathbf{x} is the following convex set

$$\frac{\partial^\circ f(\mathbf{x})}{\partial \mathbf{x}} := \text{co} \left\{ \lim_{k \rightarrow \infty} \nabla f(\mathbf{x}_k) : \mathbf{x}_k \rightarrow \mathbf{x}, f \text{ is differentiable at } \mathbf{x}_k \right\},$$

where co denotes the convex hull.

Clarke subdifferential generalize the standard notion of gradients in the sense that, when f is smooth, $\frac{\partial^\circ f(\mathbf{x})}{\partial \mathbf{x}} = \{\nabla f(\mathbf{x})\}$. Clarke subdifferential satisfies the chain rule:

Theorem K.5 (Theorem 2.3.10, [Clarke 1990](#)). *Let $F : \mathbb{R}^k \rightarrow \mathbb{R}^d$ be a differentiable function and $g : \mathbb{R}^d \rightarrow \mathbb{R}$ Lipschitz around $F(x)$. Then $f = g \circ F$ is Lipschitz around x and one has*

$$\frac{\partial^\circ f(x)}{\partial x} \subseteq \frac{\partial^\circ g(F(x))}{\partial F} \circ \frac{dF(x)}{dx}.$$

Let $\lambda_m : \mathbb{S}_d \rightarrow \mathbb{R}$, $M \mapsto \lambda_m(M)$ be the m -th largest eigenvalue of a symmetric matrix M . The following theorem gives the Clarke's subdifferentials of the eigenvalue:

Theorem K.6 (Theorem 5.3, [Hiriart-Urruty and Lewis 1999](#)). *The Clarke subdifferential of the eigenvalue function λ_m is given below, where co denotes the convex hull:*

$$\frac{\partial^\circ \lambda_m(M)}{\partial M} = \text{co}\{\mathbf{v}\mathbf{v}^\top : M\mathbf{v} = \lambda_m(M)\mathbf{v}, \|\mathbf{v}\|_2 = 1\}.$$

K.2 PROOF OF LEMMA 6.1

The equation to be proved is:

$$\frac{dM}{dt} = -\nabla f(M^{L/2})M^{L/2} - M^{L/2}\nabla f(M^{L/2}). \quad (29)$$

Since $W(t) \succeq 0$ by Lemma K.1, (11) can be rewritten as the following:

$$\frac{dW}{dt} = -\sum_{i=0}^{L-1} W^{\frac{2i}{L}} \nabla f(W) W^{2-\frac{2i+2}{L}}. \quad (30)$$

Proof for Lemma 6.1. Suppose $W(t)$ is a symmetric solution of (11). By Lemma K.1, we know $W(t)$ also satisfies (30). Now we let $R(t)$ be the solution of the following ODE with $R(0) := (W(0))^{\frac{1}{L}}$. Note we don't define $R(t)$ by $(W(t))^{\frac{1}{L}}$.

$$\frac{dR}{dt} = -\sum_{i=0}^{L-1} (-1)^i R^i \nabla f(R^L) R^{L-1-i}. \quad (31)$$

The calculation below shows that $R^L(t)$ also satisfies (30).

$$\begin{aligned} \frac{dR^L}{dt} &= \sum_{j=0}^{L-1} R^j \frac{dR}{dt} R^{L-1-j} = \sum_{j=0}^{L-1} \sum_{i=0}^{L-1} (-1)^i R^{i+j} \nabla f(R^L) R^{2L-2-i-j} \\ &= \sum_{i=0}^{2L-2} \left(\sum_{j=0}^i (-1)^j \right) R^i \nabla f(R^L) R^{2L-2-i} \\ &= \sum_{i=0}^{L-1} (R^L)^{\frac{2i}{L}} \nabla f(R^L) (R^L)^{2-\frac{2+2i}{L}}. \end{aligned}$$

Since $R^L(0) = W(0)$, by existence and uniqueness theorem, $R^L(t) = W(t)$, $\forall t \in \mathbb{R}$. So

$$\frac{dM}{dt} = R \frac{dR}{dt} + \frac{dR}{dt} R = -\nabla f(M^{L/2})M^{L/2} - M^{L/2}\nabla f(M^{L/2}),$$

which completes the proof. \square

K.3 PROOF FOR THEOREM 6.2

Now we turn to prove Theorem 6.2. Let $P = L/2$. Then (29) can be rewritten as

$$\frac{dM}{dt} = -(\nabla f(M^P)M^P + M^P\nabla f(M^P)). \quad (32)$$

The following lemma about the growth rate of $\lambda_k(M)$ is used later in the proof.

Lemma K.7. Suppose $\mathbf{M}(t)$ satisfies (32), we have for any $T' > T$, and $k \in [d]$,

$$\lambda_k(\mathbf{M}(T')) - \lambda_k(\mathbf{M}(T)) \leq \int_T^{T'} 2\lambda_k(\mathbf{M}(t))^P \|\nabla f(\mathbf{M}^P(t))\|_2 dt. \quad (33)$$

and

$$\frac{1}{P-1} (\lambda_k^{1-P}(\mathbf{M}(T)) - \lambda_k^{1-P}(\mathbf{M}(T'))) \leq \int_T^{T'} 2\|\nabla f(\mathbf{M}^P(t))\|_2 dt. \quad (34)$$

Proof. Since $\lambda_k(\mathbf{M}(t))$ is locally Lipschitz in t , by Rademacher's theorem, we know $\lambda_k(\mathbf{M}(t))$ is differentiable almost everywhere, and the following holds

$$\lambda_k(\mathbf{M}(T')) - \lambda_k(\mathbf{M}(T)) = \int_T^{T'} \frac{d\lambda_k(\mathbf{M}(t))}{dt} dt.$$

When $\frac{d\lambda_k(\mathbf{M}(t))}{dt}$ exists, we have

$$\begin{aligned} \frac{d\lambda_k(\mathbf{M}(t))}{dt} &\in \left\{ \left\langle \mathbf{G}, \frac{d\mathbf{M}(t)}{dt} \right\rangle : \mathbf{G} \in \frac{\partial^\circ \lambda_k(\mathbf{M})}{\partial \mathbf{M}} \right\} \\ &= \left\{ 2\lambda_k(\mathbf{M}^P(t)) \langle \mathbf{G}, -\nabla f(\mathbf{M}^P(t)) \rangle : \mathbf{G} \in \frac{\partial^\circ \lambda_k(\mathbf{M})}{\partial \mathbf{M}} \right\} \end{aligned}$$

Note that $\|\mathbf{G}\|_F \leq \|\mathbf{G}\|_* = 1$. So $|\langle \mathbf{G}, -\nabla f(\mathbf{M}^P(t)) \rangle| \leq \|\nabla f(\mathbf{M}^P(t))\|_2$. We can prove (34) with a similar argument. \square

To prove Theorem 6.2, it suffices to consider the case that $\mathbf{M}(0) = \hat{\alpha} \mathbf{I}$ where $\hat{\alpha} := \alpha^{1/P}$. WLOG we can assume $-\nabla f(\mathbf{0}) = \text{diag}(\mu_1, \dots, \mu_d)$ by choosing a suitable standard basis. By assumption in Theorem 6.2, we have $\mu_1 > \max\{\mu_2, 0\}$ and $\mu_1 = \|\nabla f(\mathbf{0})\|_2$. We use $\phi_m(\mathbf{M}_0, t)$ to denote the solution of $\mathbf{M}(t)$ when $\mathbf{M}(0) = \mathbf{M}_0$.

Let $R > 0$. Since $f(\cdot)$ is \mathcal{C}^3 -smooth, there exists $\beta > 0$ such that

$$\|\nabla f(\mathbf{W}_1) - \nabla f(\mathbf{W}_2)\|_F \leq \beta \|\mathbf{W}_1 - \mathbf{W}_2\|_2$$

for all $\mathbf{W}_1, \mathbf{W}_2$ with $\|\mathbf{W}_1\|_2, \|\mathbf{W}_2\|_2 \leq R$.

Let $\kappa = \beta/\mu_1$. We assume WLOG that $R \leq \frac{1}{\kappa(P-1)}$. Let $F_{\hat{\alpha}}(x) := \int_{x^{-(P-1)}}^{\hat{\alpha}^{-(P-1)}} \frac{dz}{1 + \kappa z^{-P/(P-1)}}$. Then $F'_{\hat{\alpha}}(x) = \frac{(P-1)x^{-P}}{1 + \kappa x^P} = \frac{P-1}{(1 + \kappa x^P)x^P}$. We will use this function to bound norm growth. Let $g_{\hat{\alpha}, c}(t) = \frac{1}{\hat{\alpha}^{-(P-1) - \kappa(P-1)c - 2\mu_1(P-1)t}}$. Define $T_{\hat{\alpha}}(r) = \frac{\hat{\alpha}^{-(P-1) - \kappa(P-1)r - r^{-(P-1)}}}{2\mu_1(P-1)}$. It is easy to verify that $g_{\hat{\alpha}, r}(T_{\hat{\alpha}}(r)) = r^{P-1}$.

Lemma K.8. For any $x \in [\hat{\alpha}, R]$ we have

$$\left(\hat{\alpha}^{-(P-1)} - x^{-(P-1)} \right) - F_{\hat{\alpha}}(x) \in [0, \kappa(P-1)x].$$

Proof. On the one hand, we have

$$\hat{\alpha}^{-(P-1)} - x^{-(P-1)} - F_{\hat{\alpha}}(x) = \int_{x^{-(P-1)}}^{\hat{\alpha}^{-(P-1)}} \left(1 - \frac{1}{1 + \kappa z^{-P/(P-1)}} \right) dz \geq 0.$$

On the other hand,

$$\begin{aligned} \hat{\alpha}^{-(P-1)} - x^{-(P-1)} - F_{\hat{\alpha}}(x) &= \int_{x^{-(P-1)}}^{\hat{\alpha}^{-(P-1)}} \frac{\kappa}{z^{P/(P-1)} + \kappa} dz \leq \kappa \int_{x^{-(P-1)}}^{\hat{\alpha}^{-(P-1)}} \frac{1}{z^{P/(P-1)}} dz \\ &= \kappa(P-1) \cdot \frac{-1}{z^{1/(P-1)}} \Big|_{x^{-(P-1)}}^{\hat{\alpha}^{-(P-1)}} \\ &\leq \kappa(P-1)x, \end{aligned}$$

which completes the proof. \square

Lemma K.9. Let \mathbf{M}_0 be a PSD matrix with $\|\mathbf{M}_0\|_2 \leq 1$. For $\mathbf{M}(t) := \phi_m(\hat{\alpha}\mathbf{M}_0, t)$ and $t \leq T_{\hat{\alpha}}(c)$,

$$\|\mathbf{M}(t)\|_2 = \lambda_1(\mathbf{M}(t)) \leq g_{\hat{\alpha},c}(t)^{\frac{1}{P-1}}.$$

Proof. Since $\|\nabla f(\mathbf{M}^P)\|_2 \leq \|\nabla f(\mathbf{0})\|_2 + \beta\|\mathbf{M}\|_2^P \leq \mu_1 + \beta(\lambda_1(\mathbf{M}))^P$, by Lemma K.7, we have

$$\begin{aligned} \lambda_1(\mathbf{M}(t)) &\leq \lambda_1(\mathbf{M}(0)) + \int_0^t 2(\mu_1 + \beta(\lambda_1(\mathbf{M}(\tau)))^P)(\lambda_1(\mathbf{M}(\tau)))^P d\tau \\ &= \hat{\alpha} + 2\mu_1(P-1) \int_0^t \frac{d\tau}{F'_{\hat{\alpha}}(\lambda_1(\mathbf{M}(\tau)))} \end{aligned}$$

So

$$F_{\hat{\alpha}}(\lambda_1(\mathbf{M}(t))) \leq 2\mu_1(P-1)t.$$

If $\|\mathbf{M}(t)\|_2 < \hat{\alpha}$, then $\|\mathbf{M}(t)\|_2 \leq g_{\hat{\alpha},c}(t)^{\frac{1}{P-1}}$. If $\|\mathbf{M}(t)\|_2 \geq \hat{\alpha}$, then by Lemma K.8,

$$F_{\hat{\alpha}}(\|\mathbf{M}(t)\|_2) \leq 2\mu_1(P-1)T_{\hat{\alpha}}(c) = \hat{\alpha}^{-(P-1)} - \kappa(P-1)c - c^{-(P-1)} \leq F_{\hat{\alpha}}(c),$$

so $\|\mathbf{M}(t)\|_2 \leq c$ for all $t \leq T_{\hat{\alpha}}(c)$. Applying Lemma K.8 again, we have

$$\hat{\alpha}^{-(P-1)} - \|\mathbf{M}(t)\|_2^{-(P-1)} \leq F(\|\mathbf{M}(t)\|_2) + \kappa(P-1)c \leq 2\mu_1(P-1)t + \kappa(P-1)c,$$

which implies $\|\mathbf{M}(t)\|_2 \leq g_{\hat{\alpha},c}(t)^{\frac{1}{P-1}}$ by definition. \square

Consider the following ODE:

$$\frac{d\widehat{\mathbf{M}}}{dt} = -\left(\nabla f(\mathbf{0})\widehat{\mathbf{M}}^P + \widehat{\mathbf{M}}^P\nabla f(\mathbf{0})\right).$$

We use $\hat{\phi}_m(\widehat{\mathbf{M}}_0, t)$ to denote the solution of $\widehat{\mathbf{M}}(t)$ when $\widehat{\mathbf{M}}(0) = \widehat{\mathbf{M}}_0$. For diagonal matrix $\widehat{\mathbf{M}}_0$, $\widehat{\mathbf{M}}(t)$ is also diagonal for any t , and it is easy to show that

$$\mathbf{e}_i^\top \widehat{\mathbf{M}}(t) \mathbf{e}_i = \begin{cases} \left(\frac{1}{(\hat{\alpha} \mathbf{e}_i^\top \widehat{\mathbf{M}}_0 \mathbf{e}_i)^{-(P-1)} - 2\mu_i(P-1)t}\right)^{\frac{1}{P-1}} & \mathbf{e}_i^\top \widehat{\mathbf{M}}_0 \mathbf{e}_i \neq 0, \\ 0 & \mathbf{e}_i^\top \widehat{\mathbf{M}}_0 \mathbf{e}_i = 0. \end{cases} \quad (35)$$

Remark K.10. Unlike depth-2 case, the closed form solution, $\widehat{\mathbf{M}}(t)$ is only tractable for diagonal initialization, i.e., (35) (note that the identity matrix is diagonal). And this is the main barrier for extending our two-phase analysis to the case of general initialization when $L \geq 3$. In Appendix L, we give a more detailed discussion on this barrier.

The following lemma shows that the trajectory of $\mathbf{M}(t)$ is close to $\widehat{\mathbf{M}}(t)$.

Lemma K.11. Let \mathbf{M}_0 be a diagonal PSD matrix with $\|\mathbf{M}_0\|_2 \leq 1$. For $\mathbf{M}(t) := \phi_m(\hat{\alpha}\mathbf{M}_0, t)$ and $\widehat{\mathbf{M}}(t) := \hat{\phi}_m(\hat{\alpha}\mathbf{M}_0, t)$, we have

$$\|\mathbf{M}(T_{\hat{\alpha}}(r)) - \widehat{\mathbf{M}}(T_{\hat{\alpha}}(r))\|_{\text{F}} = O(r^{P+1}).$$

Proof. We bound the difference $\mathbf{D} := \mathbf{M} - \widehat{\mathbf{M}}$ between \mathbf{M} and $\widehat{\mathbf{M}}$.

$$\begin{aligned} \left\| \frac{d\mathbf{D}}{dt} \right\|_{\text{F}} &= 2 \left\| \nabla f(\mathbf{0}) \left(\mathbf{M}^P - \widehat{\mathbf{M}}^P \right) + \left(\nabla f(\mathbf{M}^P) - \nabla f(\mathbf{0}) \right) \mathbf{M}^P \right\|_{\text{F}} \\ &\leq 2 \left(\|\nabla f(\mathbf{0})\|_2 \|\mathbf{M}^P - \widehat{\mathbf{M}}^P\|_{\text{F}} + \|\nabla f(\mathbf{M}^P) - \nabla f(\mathbf{0})\|_{\text{F}} \|\mathbf{M}^P\|_2 \right) \\ &\leq 2 \left(\mu_1 P \max\{\|\mathbf{M}\|_2^{P-1}, \|\widehat{\mathbf{M}}\|_2^{P-1}\} \|\mathbf{D}\|_{\text{F}} + \beta \|\mathbf{M}\|_2^{2P} \right), \end{aligned}$$

where the last step is by Lemma K.4. This implies that

$$\|\mathbf{D}(t)\|_{\text{F}} \leq \int_{\tau=0}^t \left\| \frac{d\mathbf{D}(\tau)}{d\tau} \right\|_{\text{F}} d\tau \leq \int_0^t 2 \left(\mu_1 P g_{\hat{\alpha},r}(\tau) \|\mathbf{D}(\tau)\|_{\text{F}} + \beta g_{\hat{\alpha},r}(\tau)^{\frac{2P}{P-1}} \right) d\tau.$$

So

$$\begin{aligned}
\|\mathbf{D}(T_{\hat{\alpha}}(r))\|_{\mathbb{F}} &\leq \int_0^{T_{\hat{\alpha}}(r)} 2\beta g_{\hat{\alpha},r}(t)^{\frac{2P}{P-1}} \exp\left(2\mu_1 P \int_t^{T_{\hat{\alpha}}(r)} g_{\hat{\alpha},r}(\tau) d\tau\right) dt \\
&= \int_0^{T_{\hat{\alpha}}(r)} 2\beta g_{\hat{\alpha},r}(t)^{\frac{2P}{P-1}} \exp\left(\frac{P}{P-1} \ln \frac{g_{\hat{\alpha},r}(T_{\hat{\alpha}}(r))}{g_{\hat{\alpha},r}(t)}\right) dt \\
&= \int_0^{T_{\hat{\alpha}}(r)} 2\beta g_{\hat{\alpha},r}(t)^{\frac{P}{P-1}} g_{\hat{\alpha},r}(T_{\hat{\alpha}}(r))^{\frac{P}{P-1}} dt \\
&= 2\beta \cdot \frac{1}{2\mu_1} g_{\hat{\alpha},r}(T_{\hat{\alpha}}(r))^{\frac{1}{P-1}} \cdot g_{\hat{\alpha},r}(T_{\hat{\alpha}}(r))^{\frac{P}{P-1}} \\
&= \kappa g_{\hat{\alpha},r}(T_{\hat{\alpha}}(r))^{\frac{P+1}{P-1}} \\
&= \kappa r^{P+1}.
\end{aligned}$$

which proves the bound. \square

Lemma K.12. Let $\mathbf{M}(t) = \phi_m(\hat{\alpha}\mathbf{M}_0, t)$, $\widetilde{\mathbf{M}}(t) = \phi_m(\hat{\alpha}\widetilde{\mathbf{M}}_0, t)$. If $\max\{\|\mathbf{M}_0\|_2, \|\widetilde{\mathbf{M}}_0\|_2\} \leq 1$. For $t \leq T_{\hat{\alpha}}(r)$, we have

$$\|\mathbf{M}(t) - \widetilde{\mathbf{M}}(t)\|_{\mathbb{F}} \leq \left(\frac{r}{\hat{\alpha}}\right)^P e^{2\kappa r^P} \|\mathbf{M}(0) - \widetilde{\mathbf{M}}(0)\|_{\mathbb{F}}.$$

Proof. Define $\mathbf{D}(t) = \mathbf{M}(t) - \widetilde{\mathbf{M}}(t)$. Then we have

$$\begin{aligned}
\left\|\frac{d\mathbf{D}}{dt}\right\|_{\mathbb{F}} &= 2 \left\| \left(\nabla f(\mathbf{M}^P) (\mathbf{M}^P - \widetilde{\mathbf{M}}^P) + (\nabla f(\mathbf{M}^P) - \nabla f(\widetilde{\mathbf{M}}^P)) \widetilde{\mathbf{M}}^P \right) \right\|_{\mathbb{F}} \\
&\leq 2 \left(\|\nabla f(\mathbf{M}^P)\|_2 \|\mathbf{M}^P - \widetilde{\mathbf{M}}^P\|_{\mathbb{F}} + \beta \|\mathbf{M}^P - \widetilde{\mathbf{M}}^P\|_{\mathbb{F}} \|\widetilde{\mathbf{M}}^P\|_2 \right) \\
&\leq 2 \left(\mu_1 + \beta \|\widetilde{\mathbf{M}}\|_2^P + \beta \|\mathbf{M}\|_2^P \right) P \max\{\|\mathbf{M}\|_2^{P-1}, \|\widetilde{\mathbf{M}}\|_2^{P-1}\} \|\mathbf{D}\|_{\mathbb{F}},
\end{aligned}$$

where the last step is by Lemma K.4. So

$$\begin{aligned}
\|\mathbf{D}(T_{\hat{\alpha}}(r))\|_{\mathbb{F}} &\leq \|\mathbf{D}(0)\|_{\mathbb{F}} \cdot \exp\left(2P\mu_1 \int_0^{T_{\hat{\alpha}}(r)} \left(1 + 2\kappa g_{\hat{\alpha},r}(t)^{\frac{P}{P-1}}\right) g_{\hat{\alpha},r}(t) dt\right) \\
&\leq \|\mathbf{D}(0)\|_{\mathbb{F}} \cdot \exp\left(\frac{P}{P-1} \ln \frac{g_{\hat{\alpha},r}(T_{\hat{\alpha}}(r))}{g_{\hat{\alpha},r}(0)} + 2\kappa g_{\hat{\alpha},r}(T_{\hat{\alpha}}(r))^{\frac{P}{P-1}}\right) \\
&\leq \|\mathbf{D}(0)\|_{\mathbb{F}} \left(\frac{r}{\hat{\alpha}}\right)^P e^{2\kappa r^P},
\end{aligned}$$

which proves the bound. \square

Let $\mathbf{M}_{\hat{\alpha}}^G(t) := \phi_m\left(\alpha \mathbf{e}_1 \mathbf{e}_1^\top, \frac{\hat{\alpha}^{-(P-1)}}{2\mu_1(P-1)} + t\right)$. Let $\overline{\mathbf{M}}(t) := \lim_{\alpha \rightarrow 0} \mathbf{M}_{\hat{\alpha}}^G(t)$.

Lemma K.13. For every $t \in (-\infty, +\infty)$, $\overline{\mathbf{M}}(t)$ exists and $\mathbf{M}_{\hat{\alpha}}^G(t)$ converges to $\overline{\mathbf{M}}(t)$ in the following rate:

$$\|\mathbf{M}_{\hat{\alpha}}^G(t) - \overline{\mathbf{M}}(t)\|_{\mathbb{F}} = O(\hat{\alpha}).$$

Proof. Let c be a sufficiently small constant. Let $\bar{T} := \frac{-\kappa(P-1)c - c^{-(P-1)}}{2\mu_1(P-1)}$. We prove this lemma in the cases of $t \in (-\infty, \bar{T}]$ and $t > \bar{T}$ respectively.

Case 1. Fix $t \in (-\infty, \bar{T}]$. Then $\frac{\hat{\alpha}^{-(P-1)}}{2\mu_1(P-1)} + t \leq T_{\hat{\alpha}}(c)$. Let $\tilde{\alpha}$ be the unique number such that $\kappa(P-1)\tilde{\alpha} + \tilde{\alpha}^{-(P-1)} = \hat{\alpha}^{-(P-1)}$. Let $\hat{\alpha}' < \hat{\alpha}$ be an arbitrarily small number. Let $t_0 := T_{\hat{\alpha}'}(\tilde{\alpha}) = \frac{(\hat{\alpha}')^{-(P-1)} - \tilde{\alpha}^{-(P-1)}}{2\mu_1(P-1)}$. By Lemma K.11 and (35), we have

$$\|\phi_m(\hat{\alpha}' \mathbf{e}_1 \mathbf{e}_1^\top, t_0) - \hat{\alpha} \mathbf{e}_1 \mathbf{e}_1^\top\|_{\mathbb{F}} \leq \left\| \phi_m(\hat{\alpha}' \mathbf{e}_1 \mathbf{e}_1^\top, t_0) - \hat{\phi}_m(\hat{\alpha}' \mathbf{e}_1 \mathbf{e}_1^\top, t_0) \right\|_{\mathbb{F}} \leq O(\tilde{\alpha}^{P+1}).$$

By Lemma K.9, $\|\phi_m(\hat{\alpha}'\mathbf{e}_1\mathbf{e}_1^\top, t_0)\|_2 \leq \tilde{\alpha}$. Then by Lemma K.12, we have

$$\|\phi_m(\hat{\alpha}'\mathbf{e}_1\mathbf{e}_1^\top, t_0 + t) - \phi_m(\hat{\alpha}\mathbf{e}_1\mathbf{e}_1^\top, t)\|_{\mathbb{F}} \leq \left(\frac{c}{\tilde{\alpha}}\right)^P e^{2\kappa c^P} \cdot O(\tilde{\alpha}^{P+1}) = O(\tilde{\alpha}) = O(\hat{\alpha}).$$

This implies that $\{M_{\hat{\alpha}}^G(t)\}$ satisfies Cauchy's criterion for every t , and thus the limit $\bar{M}(t)$ exists for $t \leq \bar{T}$. The convergence rate can be deduced by taking limits for $\hat{\alpha}' \rightarrow 0$ on both sides.

Case 2. For $t = \bar{T} + \tau$ with $\tau > 0$, $\phi_m(\mathbf{M}, \tau)$ is locally Lipschitz with respect to \mathbf{M} . So

$$\begin{aligned} \|M_{\hat{\alpha}}^G(t) - M_{\hat{\alpha}'}^G(t)\|_{\mathbb{F}} &= \|\phi_m(M_{\hat{\alpha}}^G(\bar{T}), \tau) - \phi_m(M_{\hat{\alpha}'}^G(\bar{T}), \tau)\|_{\mathbb{F}} \\ &= O(\|M_{\hat{\alpha}}^G(\bar{T}) - M_{\hat{\alpha}'}^G(\bar{T})\|_{\mathbb{F}}) \\ &= O(\hat{\alpha}), \end{aligned}$$

which proves the lemma for $t > \bar{T}$. \square

Theorem K.14. For every $t \in (-\infty, +\infty)$, as $\alpha \rightarrow 0$, we have:

$$\left\| \phi_m \left(\hat{\alpha}\mathbf{I}, \frac{\hat{\alpha}^{-(P-1)}}{2\mu_1(P-1)} + t \right) - \bar{M}(t) \right\|_{\mathbb{F}} = O(\hat{\alpha}^{\frac{1}{P+1}}), \quad (36)$$

and for any $2 \leq k \leq d$,

$$\lambda_k \left(\phi_m \left(\hat{\alpha}\mathbf{I}, \frac{\hat{\alpha}^{-(P-1)}}{2\mu_1(P-1)} + t \right) \right) = O(\hat{\alpha}). \quad (37)$$

Proof. Let $M_{\hat{\alpha}}(t) := \phi_m \left(\hat{\alpha}\mathbf{I}, \frac{\hat{\alpha}^{-(P-1)}}{2\mu_1(P-1)} + t \right)$. Again we let c be a sufficiently small constant and $\bar{T} := \frac{-\kappa(P-1)c - c^{-(P-1)}}{2\mu_1(P-1)}$. We prove in the cases of $t \in (-\infty, \bar{T}]$ and $t > \bar{T}$ respectively.

Case 1. Fix $t \in (-\infty, \bar{T}]$. Let $\hat{\alpha}_1 := \hat{\alpha}^{\frac{1}{P+1}}$. Let $\tilde{\alpha}_1$ be the unique number such that $\kappa(P-1)\tilde{\alpha}_1 + \tilde{\alpha}_1^{-(P-1)} = \hat{\alpha}_1^{-(P-1)}$. Let $t_0 := T_{\hat{\alpha}}(\tilde{\alpha}_1) = \frac{\hat{\alpha}^{-(P-1)} - \tilde{\alpha}_1^{-(P-1)}}{2\mu_1(P-1)}$. Then

$$\begin{aligned} \|\phi_m(\hat{\alpha}\mathbf{I}, t_0) - \hat{\alpha}_1\mathbf{e}_1\mathbf{e}_1^\top\|_{\mathbb{F}} &\leq \left\| \phi_m(\hat{\alpha}\mathbf{I}, t_0) - \hat{\phi}_m(\hat{\alpha}\mathbf{I}, t_0) \right\|_{\mathbb{F}} + \left\| \hat{\phi}_m(\hat{\alpha}\mathbf{I}, t_0) - \hat{\alpha}_1\mathbf{e}_1\mathbf{e}_1^\top \right\|_{\mathbb{F}} \\ &= O(\tilde{\alpha}_1^{P+1} + \hat{\alpha}) \\ &= O(\hat{\alpha}). \end{aligned}$$

By Lemma K.9, $\|\phi_m(\hat{\alpha}'\mathbf{I}, t_0)\|_2 \leq \tilde{\alpha}_1$. Then by Lemma K.12, we have

$$\begin{aligned} \|M_{\hat{\alpha}}(t) - M_{\hat{\alpha}_1}^G(t)\|_{\mathbb{F}} &= \|\phi_m(\hat{\alpha}\mathbf{I}, t_0 + t) - \phi_m(\hat{\alpha}_1\mathbf{e}_1\mathbf{e}_1^\top, t)\|_{\mathbb{F}} \\ &\leq \left(\frac{c}{\tilde{\alpha}_1}\right)^P e^{2\kappa c^P} \cdot O(\hat{\alpha}) = O(\hat{\alpha}^{\frac{1}{P+1}}). \end{aligned}$$

Combining this with the convergence rate for $M_{\hat{\alpha}_1}^G(t)$ proves the bound (36).

For (37), by Lemma K.7, we have

$$\begin{aligned} \lambda_k^{1-P}(M_{\hat{\alpha}}(\bar{T})) - \lambda_k^{1-P}(M_{\hat{\alpha}}(t_0)) &\leq \int_{t_0}^{\bar{T}} 2(P-1) \|\nabla f(M_{\hat{\alpha}}^P(t))\|_2 dt \\ &\leq \int_{t_0}^{\bar{T}} 2(P-1)(\mu_1 + \beta \|M_{\hat{\alpha}}(t)\|_2^P) dt \\ &\leq -2(P-1) \left(\mu_1(t - T_1) + \frac{\kappa}{2} \cdot g_{\hat{\alpha}, c}(t)^{\frac{1}{P-1}} \right). \end{aligned} \quad (38)$$

By Lemma K.11, $\lambda_1(M_{\hat{\alpha}}(\bar{T})) = \|M_{\hat{\alpha}}(\bar{T})\|_2 = c + O(c^{P+1})$. For $k \geq 2$,

$$\begin{aligned} \lambda_k(M_{\hat{\alpha}}(\bar{T}))^{-(P-1)} &\geq \Omega(\hat{\alpha}^{-(P-1)}) - 2(P-1) \left(\mu_1(\bar{T} - T_1) + \frac{\kappa}{2} \cdot c \right) \\ &\geq \Omega(\hat{\alpha}^{-(P-1)}) - \frac{\hat{\alpha}^{-\frac{P-1}{P+1}} - c^{-(P-1)}}{2\mu_1(P-1)} - O(c) \\ &\geq \Omega(\hat{\alpha}^{-(P-1)}). \end{aligned}$$

Thus $\lambda_k(\mathbf{M}_{\hat{\alpha}}(\bar{T})) \leq O(\hat{\alpha})$.

Case 2. For $t = \bar{T} + \tau$ with $\tau > 0$, $\phi_m(\mathbf{M}, \tau)$ is locally Lipschitz with respect to \mathbf{M} . So

$$\begin{aligned} \|\mathbf{M}_{\hat{\alpha}}(t) - \mathbf{M}_{\hat{\alpha}_1}^G(t)\|_F &= \|\phi_m(\mathbf{M}_{\hat{\alpha}}(\bar{T}), \tau) - \phi_m(\mathbf{M}_{\hat{\alpha}_1}^G(\bar{T}), \tau)\|_F \\ &= O(\|\mathbf{M}_{\hat{\alpha}}(\bar{T}) - \mathbf{M}_{\hat{\alpha}_1}^G(\bar{T})\|_F) = O(\hat{\alpha}^{\frac{1}{P+1}}), \end{aligned}$$

which proves the bound (36).

For (37), again by Lemma K.7, we have

$$\begin{aligned} &\lambda_k^{1-P}(\mathbf{M}_{\hat{\alpha}}(\bar{T})) - \lambda_k^{1-P}(\mathbf{M}_{\hat{\alpha}}(\bar{T} + \tau)) \\ &\leq \int_{\bar{T}}^{\bar{T}+\tau} 2(P-1) \|\nabla f(\mathbf{M}_{\hat{\alpha}}^P(t))\|_2 dt \\ &\leq \int_{\bar{T}}^{\bar{T}+\tau} 2(P-1) (\beta \|\mathbf{M}_{\hat{\alpha}}^P(t) - (\mathbf{M}^G)^P(t)\|_2 + \|\nabla f((\mathbf{M}^G)^P(t))\|_2) dt \\ &\leq \int_{\bar{T}}^{\bar{T}+\tau} 2(P-1)(O(\hat{\alpha}^{\frac{1}{P+1}}) + \beta \|\mathbf{M}^G(t)\|_2^P) dt \\ &\leq O(1). \end{aligned}$$

Thus $\lambda_k^{1-P}(\mathbf{M}_{\hat{\alpha}}(\bar{T} + \tau)) = \Omega(\hat{\alpha}^{-(P-1)})$, that is, $\lambda_k(\mathbf{M}_{\hat{\alpha}}(\bar{T} + \tau)) = O(\hat{\alpha})$, $\forall k \geq 2$. \square

Proof of Theorem 6.2. Note that $(\overline{\mathbf{M}}(t))^P = \overline{\mathbf{W}}(t)$ and

$$\left(\phi_m \left(\hat{\alpha} \mathbf{I}, \frac{\hat{\alpha}^{-(P-1)}}{2\mu_1(P-1)} + t \right) \right)^P = \phi \left(\alpha \mathbf{I}, \frac{\hat{\alpha}^{-(P-1)}}{2\mu_1(P-1)} + t \right).$$

By Theorem K.14, We have

$$\begin{aligned} &\left\| \phi \left(\alpha \mathbf{I}, \frac{\alpha^{-(1-1/P)}}{2\mu_1(P-1)} + t \right) - \overline{\mathbf{W}}(t) \right\|_F \\ &\leq \left\| \left(\phi_m \left(\hat{\alpha} \mathbf{I}, \frac{\hat{\alpha}^{-(P-1)}}{2\mu_1(P-1)} + t \right) \right)^P - (\overline{\mathbf{M}}(t))^P \right\|_F \\ &\leq P \left\| \phi_m \left(\hat{\alpha} \mathbf{I}, \frac{\hat{\alpha}^{-(P-1)}}{2\mu_1(P-1)} + t \right) - \overline{\mathbf{M}}(t) \right\|_F \max \left(\left\| \phi_m \left(\hat{\alpha} \mathbf{I}, \frac{\hat{\alpha}^{-(P-1)}}{2\mu_1(P-1)} + t \right) \right\|_2, \|\overline{\mathbf{M}}(t)\|_2 \right)^{P-1} \\ &= O(\hat{\alpha}^{\frac{1}{P+1}})O(1) = O(\alpha^{\frac{1}{P(P+1)}}), \end{aligned}$$

and for $2 \leq k \leq d$,

$$\lambda_k \left(\phi \left(\alpha \mathbf{I}, \frac{\alpha^{-(1-1/P)}}{2\mu_1(P-1)} + t \right) \right) = \lambda_k \left(\phi_m \left(\hat{\alpha} \mathbf{I}, \frac{\hat{\alpha}^{-(P-1)}}{2\mu_1(P-1)} + t \right) \right) = O(\hat{\alpha}^P) = O(\alpha).$$

\square

L ESCAPING DIRECTION FOR DEEP MATRIX FACTORIZATION

For deep matrix factorization, recall that we only prove that GF with infinitesimal identity initialization escapes in the direction of the top eigenvector. The main burden for us to generalize this proof to general initialization is that we don't know how to analyze the early phase dynamics of (29), i.e., the analytical solution of (39) is difficult to compute, when $L \geq 3$. Intuitively, the direction that the infinitesimal initialization escapes $\mathbf{0}$ is exactly $\overline{\mathbf{M}} := \lim_{t \rightarrow \infty} \frac{\mathbf{M}(t)}{\|\mathbf{M}(t)\|_F}$, where $\mathbf{M}(t)$ is the solution of (39). Showing $\overline{\mathbf{M}} = \mathbf{v}_1 \mathbf{v}_1^\top$ is a critical step in our analysis towards convergence to GLRL.

$$\frac{d\mathbf{M}}{dt} = -\nabla f(\mathbf{0})\mathbf{M}^{L/2} - \mathbf{M}^{L/2}\nabla f(\mathbf{0}). \quad (39)$$

However, unlike the depth-2 case, $\overline{\mathbf{M}}$ can be different from $\mathbf{v}_1\mathbf{v}_1^\top$ even if $\mathbf{v}_1^\top\mathbf{M}(0)\mathbf{v}_1 > 0$. We here give an example for diagonal $\mathbf{M}(0)$ and $\nabla f(\mathbf{0})$ at Appendix L.2. Nevertheless, we still conjecture that except for a zero measure set of $\mathbf{M}(0)$, $\overline{\mathbf{M}} = \mathbf{v}_1\mathbf{v}_1^\top$, based on the following theoretical and experimental evidences:

- If $\mathbf{v}_1^\top\mathbf{M}(0)\mathbf{v}_1 > 0$ and $\text{rank}(\mathbf{M}(0)) = 1$, we prove that $\overline{\mathbf{M}} = \mathbf{v}_1\mathbf{v}_1^\top$. (See Theorem L.1)
- For the counter-example, we show experimentally, even with perturbation of only magnitude 10^{-5} , $\overline{\mathbf{M}} \neq \mathbf{v}_1\mathbf{v}_1^\top$. The results are shown at Figure 7. The y -axis indicates $\langle \mathbf{v}_1, \mathbf{u}_1(t) \rangle$ where $\mathbf{u}_1(t)$ is the top eigenvector of $\mathbf{M}(t)$. As $\|\mathbf{W}(t)\|_F$ becomes larger, $\mathbf{u}_1(t)$ aligns better with \mathbf{v}_1 , which means the noise helps \mathbf{M} escaping from \mathbf{v}_1 . The larger the noise is, the faster $\mathbf{u}_1(t)$ converges to \mathbf{v}_1 .

L.1 RANK-ONE CASE

Theorem L.1 (rank-1 initialization escapes along the top eigenvector). *When $\text{rank}(\mathbf{M}(0)) = 1$, $\lim_{t \rightarrow \infty} \frac{\mathbf{M}(t)}{\|\mathbf{M}(t)\|_F} = \mathbf{v}_1\mathbf{v}_1^\top$, if $\mathbf{v}_1^\top\mathbf{M}(0)\mathbf{v}_1 > 0$.*

Proof. Let $\mathbf{u}(0)$ be the vector such that $\mathbf{M}(0) = \mathbf{u}(0)\mathbf{u}(0)^\top$ and $\mathbf{u}(t) \in \mathbb{R}^d$ be the solution of

$$\frac{d\mathbf{u}(t)}{dt} = \|\mathbf{u}(t)\|_2^{L-2} \nabla f(\mathbf{0})\mathbf{u}(t).$$

It is easy to check that $\mathbf{M}(t) = \mathbf{u}(t)\mathbf{u}(t)^\top$ is the solution of (39), because

$$\begin{aligned} \frac{d\mathbf{M}}{dt} &= \frac{d\mathbf{u}}{dt}\mathbf{u}^\top + \mathbf{u}\frac{d\mathbf{u}^\top}{dt} = -\nabla f(\mathbf{0})\mathbf{M}(t)\|\mathbf{u}(t)\|_2^{L-2} - \mathbf{M}(t)\nabla f(\mathbf{0})\|\mathbf{u}(t)\|_2^{L-2} \\ &= -\nabla f(\mathbf{0})\mathbf{M}^{L/2} - \mathbf{M}^{L/2}\nabla f(\mathbf{0}). \end{aligned}$$

Let $\tau(t) = \int_0^t \|\mathbf{u}(s)\|_2^{L-2} ds$. Then

$$\frac{d\mathbf{u}}{d\tau} = \frac{d\mathbf{u}}{dt} \frac{dt}{d\tau} = -\frac{1}{\frac{d\tau}{dt}} \|\mathbf{u}\|_2^{L-2} \nabla f(\mathbf{0})\mathbf{u} = -\nabla f(\mathbf{0})\mathbf{u}.$$

That is, under time rescaling $t \rightarrow \tau(t)$, the trajectory of $\mathbf{u}(t)$ still follows the power iteration, regardless of the depth L . \square

L.2 COUNTER-EXAMPLE FOR ESCAPING DIRECTION

Let $\nabla f(\mathbf{0}) = \text{diag}(2, 0.9, 0.8, \dots, 0.1) \in \mathbb{R}^{10 \times 10}$ be diagonal. Let $\mathbf{W}(0)$ be also diagonal and $\mathbf{W}(0)_{i,i} \sim \text{Unif}[0.9, 1.1] \cdot \alpha$ for $i \in [10] \setminus \{2\}$, $\mathbf{W}(0)_{2,2} = 16\alpha$, where $\alpha = 10^{-16}$ is a small constant. Let the depth be 4.

Lemma L.2. *With $\nabla f(\mathbf{0})$ and $\mathbf{W}(0)$ constructed above, $\mathbf{v}_1\mathbf{M}(0)\mathbf{v}_1^\top > 0$ and $\overline{\mathbf{M}} \neq \mathbf{v}_1\mathbf{v}_1^\top$.*

Proof. It is easy to check that $\mathbf{v}_1 = \mathbf{e}_1$, so $\mathbf{v}_1\mathbf{M}(0)\mathbf{v}_1^\top > 0$. Now we prove that $\overline{\mathbf{M}}(\infty) \neq \mathbf{v}_1\mathbf{v}_1^\top$.

As both $\mathbf{W}(0)$ and $\nabla f(\mathbf{0})$ are diagonal, $\mathbf{W}(t)$ is always diagonal and has dynamics

$$\frac{dM(t)_{i,i}}{dt} = -2\nabla f(\mathbf{0})_{i,i}M(t)_{i,i}^2, \quad \forall i \in [10],$$

therefore we have closed form of $\mathbf{M}(t)$:

$$M(t)_{i,i}^{-1} = M(0)_{i,i}^{-1} - 2\nabla f(\mathbf{0})_{i,i}t, \quad \forall i \in [10].$$

For $i \in [10]$, the time for $M(t)_{i,i}$ going to infinity is $(2M(0)_{i,i}\nabla f(\mathbf{0})_{i,i})^{-1}$. By simple calculation, $M(t)_{2,2}$ goes to infinity the fastest, thus $\overline{\mathbf{M}} = \mathbf{e}_2\mathbf{e}_2^\top \neq \mathbf{v}_1\mathbf{v}_1^\top$. \square

We remark that the scales of $\mathbf{W}(0)$ and $\nabla f(\mathbf{0})$ do not matter as in gradient flow, as scaling $\nabla f(\mathbf{0})$ is equivalent to scaling time (by Lemma L.3 below). And for this reason, the x -axis is the chosen as $\frac{\|\mathbf{W}(t)\|_F}{\|\mathbf{W}(0)\|_F}$, the relative growth rate.

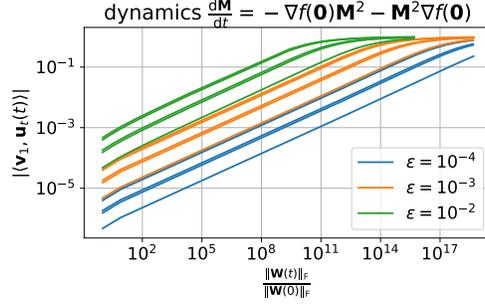


Figure 7: Dynamics of $\frac{dM}{dt} = -\nabla f(\mathbf{0})M^{L/2} - M^{L/2}\nabla f(\mathbf{0})$ plotted, where $L = 4$, $\mathbf{u}_1(t)$ is the top eigenvector of $\mathbf{W}(t)$ and ϵ is the relative magnitude of noise. The initialization we use in this experiment is $\mathbf{W}_{\text{noise}}(0) = \mathbf{W}(0) + \frac{\alpha\epsilon}{2}(\mathbf{Z} + \mathbf{Z}^\top)$, where $\mathbf{W}(0)$ is what we construct at Appendix L.2, and \mathbf{Z} is a matrix where entries are i.i.d. samples from the standard Gaussian distribution $\mathcal{N}(0, 1)$. We run 5 fixed random seeds (the noise matrix) for each ϵ . The trajectory of \mathbf{W} is calculated by simulating gradient flow on M with small timestep and RMSprop (Tieleman and Hinton, 2012) for faster convergence.

Lemma L.3. Suppose $\mathbf{g} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a P -homogeneous function, that is, $\mathbf{g}(\alpha\boldsymbol{\theta}) = \lambda^P \mathbf{g}(\boldsymbol{\theta})$ for any $\alpha > 0$, and $\frac{d\boldsymbol{\theta}'(t)}{dt} = \mathbf{g}(\boldsymbol{\theta}'(t))$. Then $\alpha\boldsymbol{\theta}'(\alpha^{P-1}t)$ is the solution of

$$\frac{d\boldsymbol{\theta}(t)}{dt} = \mathbf{g}(\boldsymbol{\theta}(t)), \quad \boldsymbol{\theta}(0) = \alpha\boldsymbol{\theta}'(0). \quad (40)$$

Proof. Simply plug in $\boldsymbol{\theta}(t) = \alpha\boldsymbol{\theta}'(\alpha^{P-1}t)$, then we have

$$\frac{d\boldsymbol{\theta}(t)}{dt} = \frac{d\alpha\boldsymbol{\theta}'(\alpha^{P-1}t)}{dt} = \alpha^P \frac{d\boldsymbol{\theta}'(\alpha^{P-1}t)}{d(\alpha^{P-1}t)} = \alpha^P \mathbf{g}(\boldsymbol{\theta}'(\alpha^{P-1}t)) = \mathbf{g}(\alpha\boldsymbol{\theta}'(\alpha^{P-1}t)) = \mathbf{g}(\boldsymbol{\theta}(t)).$$

□

M PROOF OF LINEAR CONVERGENCE TO MINIMIZER

In this section, we will present the theorems that guarantee the linear convergence to a minimizer \mathbf{W}_0 of $f(\cdot)$ if the dynamics (41) is initialized sufficiently close to \mathbf{W}_0 , i.e., $\|\mathbf{W}(0) - \mathbf{W}_0\|_F$ is sufficiently small. In Appendix M.3, we will apply this result to prove Theorem 6.4.

$$\frac{d\mathbf{W}}{dt} = -\sum_{i=0}^{L-1} \mathbf{W}^{\frac{2i}{L}} \nabla f(\mathbf{W}) \mathbf{W}^{2-\frac{2i+2}{L}} =: \mathbf{g}(\mathbf{W}). \quad (41)$$

Throughout this section, we assume $\text{rank}(\mathbf{W}_0) = k$ and use $m := \lambda_k(\mathbf{W}_0)$ to denote the k -th smallest non-zero eigenvalue of \mathbf{W}_0 . The tangent space of manifold of rank- k symmetric matrices at \mathbf{W}_0 is $\mathcal{T} = \{\mathbf{V}\mathbf{W}_0^\top + \mathbf{W}_0\mathbf{V}^\top : \mathbf{V} \in \mathbb{R}^{d \times k}\}$. It can be shown that $\dim(\mathcal{T}) = k(d-k) + \frac{k(k+1)}{2} = \frac{k(2d-k+1)}{2}$.

Let $\mathbf{J}(\mathbf{W})$ be the Jacobian of $\mathbf{g}(\mathbf{W})$ in (41). For depth-2 case, we have shown that \mathcal{T} is an invariant subspace of $\mathbf{J}(\mathbf{W}_0)$ in Theorem H.5, property 2. This can be generalize to the deep case where $L \geq 3$. Therefore, we can use $\mathbf{J}(\mathbf{W}_0)|_{\mathcal{T}} : \mathcal{T} \rightarrow \mathcal{T}$ to denote the linear operator $\mathbf{J}(\mathbf{W}_0)$ restricted on \mathcal{T} . We also define $\Pi_1^{d^2}(\mathbf{W})$ as the projection of $\mathbf{W} \in \mathbb{R}^{d \times d}$ on \mathcal{T} , and $\Pi_2^{d^2}(\mathbf{W}) := \mathbf{W} - \Pi_1^{d^2}(\mathbf{W})$.

Towards showing the main convergence result in the section, we make the following assumption.

Assumption M.1. Suppose $\mathbf{J}(\mathbf{W}_0)|_{\mathcal{T}}$ diagonalizable and all eigenvalues are negative real numbers.

\mathbf{W}_0 is a minimizer, so it is clear that $\mathbf{J}(\mathbf{W}_0)|_{\mathcal{T}}$ has no eigenvalues with positive real parts (otherwise there is a descending direction of $f(\cdot)$ from \mathbf{W}_0 , since the loss $f(\cdot)$ strictly decreases along

the trajectory of (41)). If further Assumption M.1 holds, then we know $\mathbf{J}(\mathbf{W}_0)|_{\mathcal{T}}: \mathcal{T} \rightarrow \mathcal{T}$ can be diagonalized as $\mathbf{J}(\mathbf{W}_0)|_{\mathcal{T}}[\cdot] = \mathcal{V}(\Sigma\mathcal{V}^{-1}(\cdot))$, where $\Sigma_i = \text{diag}(-\mu_1, \dots, -\mu_{\dim(\mathcal{T})})$, $\mathcal{V}: \mathbb{R}^{\dim(\mathcal{T})} \rightarrow \mathcal{T}$, $\mathcal{V}(\mathbf{x}) = \sum_{i=1}^{\dim(\mathcal{T})} x_i \mathbf{V}_i$, and \mathbf{V}_i is the eigenvector associated with eigenvalue $-\mu_i$.

As shown in Theorem M.3 below, this assumption implies that if $\mathbf{W}(0)$ is rank- k and is sufficiently close to \mathbf{W}_0 , then $\|\mathbf{W}(t) - \mathbf{W}_0\|_F \leq Ce^{-\mu_1 t}$ for some constant C . For depth-2 case, the above assumption is equivalent to that $\mathcal{L}(\mathcal{U}_0)$ is ‘‘strongly convex’’ at \mathcal{U}_0 , except those 0 eigenvalues due to symmetry, by property 2 of Theorem H.5). For the case where $L \geq 3$, because this dynamics is not gradient flow, in general it does not correspond to a loss function and strongly convexity does not make any sense. Nevertheless, in experiments we do observe linear convergence to \mathbf{W}_0 , so this assumption is reasonable.

M.1 RANK- k INITIALIZATION

For convenience, we define for all $\mathbf{W} \in \mathbb{S}_d$,

$$\|\mathbf{W}\|_{\mathcal{V}} := \left\| \mathcal{V}^{-1} \left(\Pi_1^{d^2}(\mathbf{W}) \right) \right\|_F, \quad \|\mathbf{W}\|_{F,1} := \left\| \Pi_1^{d^2}(\mathbf{W}) \right\|_F, \quad \|\mathbf{W}\|_{F,2} := \left\| \Pi_2^{d^2}(\mathbf{W}) \right\|_F.$$

The reason for such definition of norms, as we will see later, is that the norm (or the difference) in the tangent space of the manifold of symmetric rank- r matrices, $\|\mathbf{W} - \mathbf{W}'\|_{F,1}$, dominates that in the orthogonal complement of the tangent space, $\|\mathbf{W} - \mathbf{W}'\|_{F,2}$, when both \mathbf{W}, \mathbf{W}' get very close to the \mathbf{W}_0 (see a more rigorous statement in Lemma M.2). WLOG, we can assume

$$\frac{\|\cdot\|_{F,1}}{K} \leq \|\cdot\|_{\mathcal{V}} \leq \|\cdot\|_{F,1},$$

for some constant K , which may depend on f and \mathbf{W}_0 . This also implies that $\|\cdot\|_{\mathcal{V}} \leq \|\cdot\|_F$. Below we also assume for sufficiently small R , and any \mathbf{W} such that $\|\mathbf{W} - \mathbf{W}_0\|_F \leq R$, we have $\|\nabla f(\mathbf{W})\|_2 \leq \rho$ and $\|\mathbf{J}(\mathbf{W})[\Delta]\|_F \leq \beta \|\Delta\|_F$ for any Δ . In the proof below, we assume such properties hold as long as we can show the boundedness of $\mathbf{W}(t) - \mathbf{W}_0$.

Lemma M.2. *Let $\max\{\|\mathbf{W} - \mathbf{W}_0\|_{F,1}, \|\mathbf{W}' - \mathbf{W}_0\|_{F,1}\} = r$, when $r \leq \frac{m}{2}$, we have*

$$\|\mathbf{W} - \mathbf{W}'\|_{F,2} \leq \frac{5r}{m} \|\mathbf{W} - \mathbf{W}'\|_{F,1}.$$

As a special case, we have

$$\|\mathbf{W} - \mathbf{W}_0\|_{F,2} \leq \frac{5 \|\mathbf{W} - \mathbf{W}'\|_{F,1}^2}{m}.$$

Proof. WLOG we can assume \mathbf{W}_0 is only non-zero in the first k dimension, i.e., $[\mathbf{W}_0]_{ij} = 0$, for all $i \geq k+1, j \geq k+1$. We further denote \mathbf{W} and \mathbf{W}' by

$$\mathbf{W} = \begin{bmatrix} \mathbf{A} & \mathbf{B}^\top \\ \mathbf{B} & \mathbf{C} \end{bmatrix} \quad \text{and} \quad \mathbf{W}' = \begin{bmatrix} \mathbf{A}' & \mathbf{B}'^\top \\ \mathbf{B}' & \mathbf{C}' \end{bmatrix},$$

where $\mathbf{A}, \mathbf{A}' \in \mathbb{R}^{k \times k}$, $\mathbf{B}, \mathbf{B}' \in \mathbb{R}^{(d-k) \times k}$, $\mathbf{C}, \mathbf{C}' \in \mathbb{R}^{(d-k) \times (d-k)}$. By definition, we have $\|\mathbf{A} - \mathbf{A}'\|_F, \|\mathbf{B} - \mathbf{B}'\|_F \leq \|\mathbf{W} - \mathbf{W}'\|_{F,1}$ and $\|\mathbf{W} - \mathbf{W}'\|_{F,2} = \|\mathbf{C} - \mathbf{C}'\|_F$. Moreover, we have $\lambda_{\min}(\mathbf{A}) \geq m - \|\mathbf{A} - \mathbf{W}_0\|_F \geq m - \|\mathbf{W} - \mathbf{W}_0\|_{F,1} \geq \frac{m}{2}$.

Since \mathbf{W}, \mathbf{W}' is rank- k , we have $\mathbf{C} = \mathbf{B}\mathbf{A}^{-1}\mathbf{B}^\top, \mathbf{C}' = \mathbf{B}'\mathbf{A}'^{-1}\mathbf{B}'^\top$. Thus

$$\begin{aligned} & \|\mathbf{W} - \mathbf{W}'\|_{F,2} \\ &= \|\mathbf{C} - \mathbf{C}'\|_F \\ &= \left\| \mathbf{B}\mathbf{A}^{-1}\mathbf{B}^\top - \mathbf{B}'\mathbf{A}'^{-1}\mathbf{B}'^\top \right\|_F \\ &\leq \|\mathbf{B} - \mathbf{B}'\|_F \|\mathbf{A}^{-1}\mathbf{B}^\top\|_F + \|\mathbf{B}\mathbf{A}^{-1}\|_F \|\mathbf{A}' - \mathbf{A}\|_F \|\mathbf{A}'^{-1}\mathbf{B}'^\top\|_F + \|\mathbf{B}'\mathbf{A}'^{-1}\|_F \|\mathbf{B}^\top - \mathbf{B}'^\top\|_F \\ &\leq \|\mathbf{W} - \mathbf{W}'\|_{F,1} \frac{2r}{m} + \|\mathbf{W} - \mathbf{W}'\|_{F,1} \left(\frac{2r}{m} \right)^2 + \|\mathbf{W} - \mathbf{W}'\|_{F,1} \frac{2r}{m} \\ &\leq \|\mathbf{W} - \mathbf{W}'\|_{F,1} \frac{5r}{m}. \end{aligned}$$

□

Theorem M.3 (Linear convergence of rank- k matrices). *Suppose that $\text{rank}(\mathbf{W}(0)) = \text{rank}(\mathbf{W}_0) = k$ and*

$$\|\mathbf{W}(0) - \mathbf{W}_0\|_{\mathcal{V}} \leq R := \max \left\{ \frac{m}{2K}, \frac{\mu_1}{K^2(29\beta + 10\rho/m)} \right\},$$

we have $\|\mathbf{W}(t) - \mathbf{W}_0\|_{\mathcal{V}} \leq Ce^{-\mu_1 t} \|\mathbf{W}(0) - \mathbf{W}_0\|_{\mathcal{V}}$ for some constant C depending on \mathbf{W}_0 , where $\mathbf{W}(t)$ satisfies (41).

Proof. For convenience, we define $\mathbf{W}_1(t) := \Pi_1^{d^2}(\mathbf{W}(t) - \mathbf{W}_0)$, $\mathbf{W}_2(t) := \Pi_2^{d^2}(\mathbf{W}(t) - \mathbf{W}_0) = \Pi_2^{d^2}(\mathbf{W}(t))$. We also use $\langle \cdot, \cdot \rangle_{\mathcal{V}^{-1}} = \langle \mathcal{V}^{-1}(\cdot), \mathcal{V}^{-1}(\cdot) \rangle$ for short.

$$\begin{aligned} \frac{d \|\mathbf{W}_1(t)\|_{\mathcal{V}}^2}{dt} &= \frac{d \|\mathbf{W}(t) - \mathbf{W}_0\|_{\mathcal{V}}^2}{dt} \\ &= 2 \left\langle \Pi_1^{d^2} \left(\frac{d\mathbf{W}(t)}{dt} \right), \Pi_1^{d^2}(\mathbf{W}(t) - \mathbf{W}_0) \right\rangle_{\mathcal{V}^{-1}} \\ &= 2 \left\langle \Pi_1^{d^2}(g(\mathbf{W}(t))), \mathbf{W}_1(t) \right\rangle_{\mathcal{V}^{-1}} \\ &\leq 2 \left\langle \Pi_1^{d^2}(\mathbf{J}(\mathbf{W}_0)[\mathbf{W}(t) - \mathbf{W}_0]), \mathbf{W}_1(t) \right\rangle_{\mathcal{V}^{-1}} \\ &\quad + 2 \|g(\mathbf{W}(t) - \mathbf{W}_0) - \mathbf{J}(\mathbf{W}_0)[\mathbf{W}(t) - \mathbf{W}_0]\|_{\mathcal{V}} \|\mathbf{W}(t) - \mathbf{W}_0\|_{\mathcal{V}} \\ &= 2 \left\langle \Pi_1^{d^2}(\mathbf{J}(\mathbf{W}_0)[\mathbf{W}_1(t) + \mathbf{W}_2(t)]), \mathbf{W}_1(t) \right\rangle_{\mathcal{V}^{-1}} \\ &\quad + 2 \|g(\mathbf{W}(t) - \mathbf{W}_0) - \mathbf{J}(\mathbf{W}_0)[\mathbf{W}(t) - \mathbf{W}_0]\|_{\mathcal{V}} \|\mathbf{W}_1(t)\|_{\mathcal{V}} \\ &= 2 \left\langle \Pi_1^{d^2}(\mathbf{J}(\mathbf{W}_0)[\mathbf{W}_1(t)]), \mathbf{W}_1(t) \right\rangle_{\mathcal{V}^{-1}} + 2 \|\mathbf{J}(\mathbf{W}_0)[\mathbf{W}_2(t)]\|_{\mathcal{V}} \|\mathbf{W}_1(t)\|_{\mathcal{V}} \\ &\quad + 2 \|g(\mathbf{W}(t) - \mathbf{W}_0) - \mathbf{J}(\mathbf{W}_0)[\mathbf{W}(t) - \mathbf{W}_0]\|_{\mathcal{V}} \|\mathbf{W}_1(t)\|_{\mathcal{V}}. \end{aligned}$$

For the first term $\left\langle \Pi_1^{d^2}(\mathbf{J}(\mathbf{W}_0)[\mathbf{W}_1(t)]), \mathbf{W}_1(t) \right\rangle_{\mathcal{V}^{-1}}$, we know $\mathbf{W}_1(t) \in \mathcal{T}$, and \mathcal{T} is an invariant space of $\mathbf{J}(\mathbf{W}_0)$. Recall $\mathbf{J}(\mathbf{W}_0)|_{\mathcal{T}}[\cdot] = \mathcal{V}(\Sigma\mathcal{V}^{-1}(\cdot))$, we have

$$2 \left\langle \Pi_1^{d^2}(\mathbf{J}(\mathbf{W}_0)[\mathbf{W}_1(t)]), \mathbf{W}_1(t) \right\rangle_{\mathcal{V}^{-1}} = 2 \langle \Sigma\mathcal{V}^{-1}(\mathbf{W}_1(t)), \mathcal{V}^{-1}(\mathbf{W}_1(t)) \rangle \leq -2\mu_1 \|\mathbf{W}_1(t)\|_{\mathbb{F},1}.$$

For the second term $2\beta \|\mathbf{J}(\mathbf{W}_0)[\mathbf{W}_2(t)]\|_{\mathcal{V}} \|\mathbf{W}_1(t)\|_{\mathcal{V}}$, we have

$$2 \|\mathbf{J}(\mathbf{W}_0)[\mathbf{W}_2(t)]\|_{\mathcal{V}} \leq 2 \|\mathbf{J}(\mathbf{W}_0)[\mathbf{W}_2(t)]\|_{\mathbb{F}} \leq 2 \|\mathbf{J}(\mathbf{W}_0)\|_2 \|\mathbf{W}_2(t)\|_{\mathbb{F}} = 2\rho \|\mathbf{W}_2(t)\|_{\mathbb{F}}.$$

For the third term $2 \|g(\mathbf{W}(t) - \mathbf{W}_0) - \mathbf{J}(\mathbf{W}_0)[\mathbf{W}(t) - \mathbf{W}_0]\|_{\mathcal{V}} \|\mathbf{W}_1(t)\|_{\mathcal{V}}$, we have

$$\begin{aligned} 2 \|g(\mathbf{W}(t) - \mathbf{W}_0) - \mathbf{J}(\mathbf{W}_0)[\mathbf{W}(t) - \mathbf{W}_0]\|_{\mathcal{V}} &\leq 2\beta \|\mathbf{W}(t) - \mathbf{W}_0\|_{\mathbb{F}}^2 \\ &\leq 4\beta(\|\mathbf{W}_1(t)\|_{\mathbb{F}}^2 + \|\mathbf{W}_2(t)\|_{\mathbb{F}}^2) \\ &\leq 4\beta(K^2 \|\mathbf{W}_1(t)\|_{\mathcal{V}}^2 + \|\mathbf{W}_2(t)\|_{\mathbb{F}}^2). \end{aligned}$$

Thus we have shown the following. Note so far we have not used the assumption that \mathbf{W} is rank- k .

$$\frac{d \|\mathbf{W}_1(t)\|_{\mathcal{V}}^2}{dt} \leq -2\mu_1 \|\mathbf{W}_1(t)\|_{\mathcal{V}}^2 + 2 \|\mathbf{W}_1(t)\|_{\mathcal{V}} \left(\rho \|\mathbf{W}_2(t)\|_{\mathbb{F}} + 2\beta K^2 \|\mathbf{W}_1(t)\|_{\mathcal{V}}^2 + 2\beta \|\mathbf{W}_2(t)\|_{\mathbb{F}}^2 \right),$$

that is,

$$\frac{d \log \|\mathbf{W}_1(t)\|_{\mathcal{V}}^2}{dt} \leq -2\mu_1 + 4\beta K^2 \|\mathbf{W}_1(t)\|_{\mathcal{V}} + \frac{4\beta \|\mathbf{W}_2(t)\|_{\mathbb{F}}^2 + 2\rho \|\mathbf{W}_2(t)\|_{\mathbb{F}}}{\|\mathbf{W}_1(t)\|_{\mathcal{V}}}. \quad (42)$$

Let $T := \sup\{t \geq 0 : \|\mathbf{W}_1(t)\|_{\mathcal{V}} \leq \frac{m}{2K}\}$. Setting $\mathbf{W}' = \mathbf{W}_0$ in Lemma M.2, we have for $t < T$, $r = \|\mathbf{W}(t) - \mathbf{W}_0\|_{\mathbb{F},1} \leq \|\mathbf{W}(t) - \mathbf{W}_0\|_{\mathbb{F}} \leq K \|\mathbf{W}(t) - \mathbf{W}_0\|_{\mathcal{V}} \leq \frac{m}{2}$. Thus,

$$\|\mathbf{W}_2(t)\|_{\mathbb{F}} = \|\mathbf{W}_2(t)\|_{\mathbb{F},2} \leq \frac{5 \|\mathbf{W}(t) - \mathbf{W}_0\|_{\mathbb{F},1}^2}{m} \leq \frac{5K^2 \|\mathbf{W}(t) - \mathbf{W}_0\|_{\mathcal{V}}^2}{m} \leq \frac{5}{4}m.$$

Thus, from (42) we can derive that

$$\frac{d \log \|\mathbf{W}_1(t)\|_{\mathcal{V}}^2}{dt} \leq -2\mu_1 + K^2(29\beta + 10\rho/m) \|\mathbf{W}_1(t)\|_{\mathcal{V}} \leq -\mu_1. \quad (43)$$

Since $\mu_1 < 0$, $\|\mathbf{W}_1(t)\|_{\mathcal{V}}$ decreases for $[0, T)$. Thus T must be ∞ , otherwise $\|\mathbf{W}_1(T)\|_{\mathcal{V}} = \lim_{t \rightarrow T^-} \|\mathbf{W}_1(t)\|_{\mathcal{V}} < R_1$. Contradiction.

Therefore, for any $t \in [0, \infty)$, we have $\|\mathbf{W}_1(t)\|_{\mathcal{V}} \leq \|\mathbf{W}_1(0)\|_{\mathcal{V}} e^{-\frac{\mu_1}{2}t}$. That is,

$$\int_0^{\infty} \|\mathbf{W}_1(t)\|_{\mathcal{V}} dt \leq \frac{2}{\mu_1} \|\mathbf{W}_1(0)\|_{\mathcal{V}} \leq \frac{2R}{\mu_1}.$$

Thus from (43), we have

$$\begin{aligned} \|\mathbf{W}(t)\|_{\mathcal{V}} &= \|\mathbf{W}_1(t)\|_{\mathcal{V}} \leq \|\mathbf{W}_1(0)\|_{\mathcal{V}} \exp\left(-\mu_1 t + \frac{K^2}{2}(29\beta + 10\rho/m) \int_0^{\infty} \|\mathbf{W}_1(t)\|_{\mathcal{V}} dt\right) \\ &\leq \|\mathbf{W}_1(0)\|_{\mathcal{V}} \exp\left(-\mu_1 t + \frac{K^2 R}{\mu_1}(29\beta + 10\rho/m)\right) \\ &=: C \|\mathbf{W}(0)\|_{\mathcal{V}} e^{-\mu_1 t}, \end{aligned}$$

which completes the proof. \square

M.2 ALMOST RANK- k INITIALIZATION

We use $\mathbf{M}(t)$ to denote the top- k components of $\mathbf{W}(t)$ in SVD, and $\mathbf{N}(t)$ to denote the rest part, i.e., $\mathbf{W}(t) - \mathbf{M}(t)$. One can think $\mathbf{M}(t)$ as the main part and $\mathbf{N}(t)$ as the negligible part.

Below we show that for deep overparametrized matrix factorization, where $\mathbf{W}(t)$ satisfies (41), if the trajectory is initialized at some $\mathbf{W}(0)$ in a small neighborhood of the k -th critical point \mathbf{W}_0 of deep GLRL, and $\mathbf{W}(0)$ is approximately rank- k , in the sense that $\mathbf{N}(0)$ is very small, then $\inf_{t \geq 0} \|\mathbf{W}(t) - \mathbf{W}_0\|_{\mathcal{V}}$ is roughly at the same magnitude of $\mathbf{N}(0)$.

Theorem M.4 (Linear convergence of almost rank- k matrices, deep case). *Suppose \mathbf{W}_0 is a critical point of rank k and \mathbf{W}_0 satisfies Assumption M.1, there exists constants C_0 and r , such that if $C_0 \|\mathbf{N}(0)\|_{\mathbb{F}} \leq \|\mathbf{W}_1(0)\|_{\mathcal{V}} \leq r$, then there exists a time T and constants C, C' , such that*

- (1). $\|\mathbf{W}(t) - \mathbf{W}_0\|_{\mathcal{V}} \leq C e^{-\mu_1 t/2} \|\mathbf{W}(0) - \mathbf{W}_0\|_{\mathcal{V}}$, for $t \leq T$.
- (2). $\|\mathbf{W}(T) - \mathbf{W}_0\|_{\mathbb{F}} \leq C' \|\mathbf{N}(0)\|_{\mathbb{F}}$.

Proof. When $\|\mathbf{W}(t) - \mathbf{W}_0\|_{\mathbb{F}} \leq \frac{\lambda_{\min}(\mathbf{W}_0)}{4}$, $\|\mathbf{N}(t)\|_{\mathbb{F}} \leq \frac{\lambda_{\min}(\mathbf{W}_0)}{4}$, thus we have

$$\|\mathbf{M}(t) - \mathbf{W}_0\|_{\mathbb{F},1} \leq \|\mathbf{W}(t) - \mathbf{W}_0\|_{\mathbb{F},1} + \|\mathbf{N}(t)\|_{\mathbb{F},1} \leq \frac{\lambda_{\min}(\mathbf{W}_0)}{2},$$

thus by Lemma M.2, we have

$$\begin{aligned} \|\mathbf{W}_2(t)\|_{\mathbb{F},2} &\leq \|\mathbf{M}(t) - \mathbf{W}_0\|_{\mathbb{F},2} + \|\mathbf{N}(t)\|_{\mathbb{F},2} \\ &\leq \frac{5 \|\mathbf{M}(t) - \mathbf{W}_0\|_{\mathbb{F},1}^2}{\lambda_{\min}(\mathbf{W}_0)} + \|\mathbf{N}(t)\|_{\mathbb{F},2} \\ &\leq \frac{10 \|\mathbf{W}_1(t)\|_{\mathbb{F},1}^2 + 10 \|\mathbf{N}(t)\|_{\mathbb{F}}^2}{\lambda_{\min}(\mathbf{W}_0)} + \|\mathbf{N}(t)\|_{\mathbb{F},2} \\ &\leq \frac{10K^2 \|\mathbf{W}_1(t)\|_{\mathcal{V}}^2 + 10 \|\mathbf{N}(t)\|_{\mathbb{F}}^2}{\lambda_{\min}(\mathbf{W}_0)} + \|\mathbf{N}(t)\|_{\mathbb{F},2}. \end{aligned}$$

Thus we can pick constant C_0 large enough and r small enough, such that for any $t \geq 0$, if $C_0 \|\mathbf{N}(t)\|_{\mathbb{F}} \leq \|\mathbf{W}_1(t)\|_{\mathcal{V}} \leq r$, then it holds that:

- The “small terms” in the RHS of (42) satisfies that

$$4\beta K^2 \|\mathbf{W}_1(t)\|_{\mathcal{V}} + \frac{4\beta \|\mathbf{W}_2(t)\|_{\mathbb{F}}^2 + 2\rho \|\mathbf{W}_2(t)\|_{\mathbb{F}}}{\|\mathbf{W}_1(t)\|_{\mathcal{V}}} \leq C_1 \|\mathbf{W}_1(t)\|_{\mathcal{V}} + C_2 \|\mathbf{N}(t)\|_{\mathbb{F}} \leq \mu_1$$

for some C_1 and C_2 independent of t .

- The spectral norm $\frac{1}{2} \|\nabla f(\mathbf{W}(t))\|_2 \leq \|\nabla f(\mathbf{W}_0)\|_2 =: \rho$ for all $t \geq 0$.
- $\forall x < r, \frac{\kappa_L x^{\frac{2}{L}-1}}{(L-2)\rho} > \frac{2}{\mu_1} \ln \frac{2r}{C_0 x}$, where $\kappa_L = 1 - 0.5^{\frac{L-2}{L}}$.

Note these conditions can always be satisfied by some C_0 and r because we can first find 3 groups (C_0, r) to satisfy each individual condition, and then take the maximal C_0 and minimal r , it's easy to check these conditions are still verified. And we let $T_{C_0, r}$ be the earliest time that such condition, i.e., $C_0 \|\mathbf{N}(t)\|_{\mathbb{F}} \leq \|\mathbf{W}_1(t)\|_{\mathcal{V}} \leq r$ fails. Thus by (42), for $t \in [0, T_{C_0, r})$, we have $\|\mathbf{W}(t)\|_{\mathcal{V}} = \|\mathbf{W}_1(t)\|_{\mathcal{V}} \leq \|\mathbf{W}_1(0)\|_{\mathcal{V}} e^{-\frac{\mu_1 t}{2}} = \|\mathbf{W}(0)\|_{\mathcal{V}} e^{-\frac{\mu_1 t}{2}}$. Thus (1) holds for any T smaller than $T_{C_0, r}$. If $T_{C_0, r} = \infty$, then clearly we can pick a sufficiently large T , such that (2) holds. Therefore, below it suffices to consider the case where $T_{C_0, r}$ is finite. And we know the condition that fails must be $C_0 \|\mathbf{N}(t)\|_{\mathbb{F}} \leq \|\mathbf{W}_1(t)\|_{\mathcal{V}}$, i.e. $C_0 \|\mathbf{N}(T_{C_0, r})\|_{\mathbb{F}} = \|\mathbf{W}_1(T_{C_0, r})\|_{\mathcal{V}}$.

By (34) in Lemma K.7, we have

$$\left| \|\mathbf{N}(0)\|_2^{\frac{2}{L}-1} - \|\mathbf{N}(t)\|_2^{\frac{2}{L}-1} \right| \leq (L-2)\rho t.$$

Define $T' := \frac{\kappa_L \|\mathbf{N}(0)\|_2^{\frac{2}{L}-1}}{(L-2)\rho}$, we know for any $t < T'$, we have $\left| \|\mathbf{N}(0)\|_2^{\frac{2}{L}-1} - \|\mathbf{N}(t)\|_2^{\frac{2}{L}-1} \right| \leq \kappa_L \|\mathbf{N}(t)\|_2^{\frac{2}{L}-1}$. That is,

$$\frac{\|\mathbf{N}(t)\|_2^{\frac{2}{L}-1}}{\|\mathbf{N}(0)\|_2^{\frac{2}{L}-1}} \in \left[1 - \kappa_L, \frac{1}{1 - \kappa_L} \right] = \left[0.5^{\frac{L-2}{L}}, 0.5^{-\frac{L-2}{L}} \right] \implies \frac{\|\mathbf{N}(t)\|_2}{\|\mathbf{N}(0)\|_2} \in [1/2, 2].$$

Now we claim it must hold that $T' \geq T_{C_0, r}$. Otherwise, we have

$$\frac{C_0}{2} \|\mathbf{N}(0)\|_2 \leq C_0 \|\mathbf{N}(T')\|_{\mathbb{F}} \leq \|\mathbf{W}_1(T')\|_{\mathcal{V}} \leq e^{-\mu_1 T'/2} \|\mathbf{W}_1(0)\|_{\mathcal{V}} \leq e^{-\mu_1 T'/2} r.$$

Therefore, $\frac{\kappa_L \|\mathbf{N}(0)\|_2^{\frac{2}{L}-1}}{(L-2)\rho} = T' \leq \frac{2}{\mu_1} \ln \frac{2r}{C_0 \|\mathbf{N}(0)\|_2}$, which contradicts to the definition of C_0 and r .

As a result, we have

$$\begin{aligned} 2C_0 \sqrt{d} \|\mathbf{N}(0)\|_2 &\geq 2C_0 \|\mathbf{N}(0)\|_{\mathbb{F}} \geq C_0 \|\mathbf{N}(T_{C_0, r})\|_{\mathbb{F}} = \|\mathbf{W}_1(T_{C_0, r})\|_{\mathcal{V}} \\ &\geq \|\mathbf{W}_1(0)\|_{\mathcal{V}} e^{-\mu_1 T_{C_0, r}/2}, \end{aligned}$$

and therefore,

$$T_{C_0, r} \leq \frac{2}{\mu_1} \ln \frac{\|\mathbf{W}_1(0)\|_{\mathcal{V}}}{2\sqrt{d}C_0 \|\mathbf{N}(0)\|_{\mathbb{F}}}.$$

Thus by Lemma M.2, we know

$$\begin{aligned} \|\mathbf{W}(T_{C_0, r}) - \mathbf{W}_0\|_{\mathbb{F}} &\leq \|\mathbf{W}(T_{C_0, r}) - \mathbf{W}_0\|_{\mathbb{F}, 1} + \|\mathbf{W}(T_{C_0, r}) - \mathbf{W}_0\|_{\mathbb{F}, 2} \\ &\leq K \|\mathbf{W}_1(T_{C_0, r})\|_{\mathcal{V}} + \|\mathbf{M}(T_{C_0, r}) - \mathbf{W}_0\|_{\mathbb{F}, 2} + \|\mathbf{N}(T_{C_0, r})\|_{\mathbb{F}, 2} \\ &\leq O(\|\mathbf{N}(0)\|_{\mathbb{F}}) + O(\|\mathbf{N}(0)\|_{\mathbb{F}}^2) + O(\|\mathbf{N}(0)\|_{\mathbb{F}}) \\ &= O(\|\mathbf{N}(0)\|_{\mathbb{F}}). \end{aligned}$$

□

M.3 PROOF FOR THEOREM 6.4

Proof for Theorem 6.4. Let C_0, r be the constants predicted by Theorem M.4 w.r.t. to $\overline{\mathbf{W}}(\infty)$. We claim that we can pick large enough constant T , and α_0 sufficiently small, such that for all $\alpha \leq \alpha_0$, the initial condition in Theorem M.4 holds, i.e. $C_0 \|\mathbf{N}(0)\|_{\mathbb{F}} \leq \|\mathbf{W}_1(0)\|_{\mathcal{V}} \leq r$, where $\mathbf{W}(0) := \phi\left(\alpha \mathbf{I}, \frac{\alpha^{-(P-1)}}{2\mu_1^{1-(P-1)}} + T\right)$.

This is because we can first ensure $\|\overline{\mathbf{W}}(T) - \overline{\mathbf{W}}(\infty)\|_2$ is sufficiently small, i.e., smaller than $\frac{\epsilon}{2}$. By Theorem 6.2, we know when $\alpha \rightarrow 0$, $\|\overline{\mathbf{W}}(T) - \mathbf{W}(0)\|_{\mathcal{V}} \leq K \|\overline{\mathbf{W}}(T) - \mathbf{W}(0)\|_{\mathbb{F}} = o(1)$ and $\|\mathbf{N}(0)\|_{\mathbb{F}} = O(\alpha)$.

By Theorem M.4, we know there is a time T (either $T_{C_0, r}$ or some sufficiently large number when $T_{C_0, r} = \infty$), such that $\|\mathbf{W}(T) - \mathbf{W}_0\|_{\mathbb{F}} = O(\|\mathbf{N}(0)\|_{\mathbb{F}}) = O(\alpha)$. \square