LEARNING DEEP MODELS: CRITICAL POINTS AND LOCAL OPENNESS

Anonymous authors
Paper under double-blind review

ABSTRACT

With the increasing interest in deeper understanding of the loss surface of many non-convex *deep models*, this paper presents a unifying framework to study the local/global equivalence of the optimization problem arising from training of such non-convex models. Using the *local openness* property of the underlying training models, we provide sufficient conditions under which any local optimum of the resulting optimization problem is global. Our result unifies and extends many of the existing results in the literature. For example, our theory shows that when the input data matrix X is full row rank, all *non-degenerate* local optima of the optimization problem for training linear deep model with squared loss error are global minima. Moreover, for two layer linear models, we show that all *degenerate* critical points are either global or second order saddles and the *non-degenerate* local optima are global. Unlike many existing results in the literature, our result assumes no assumption on the target data matrix Y. For non-linear deep models having certain pyramidal structure with invertible activation functions, we can show global/local equivalence with no assumption on the differentiability of the activation function. Our results are the direct consequence of our main theorem that provides necessary and sufficient conditions for the matrix multiplication mapping to be locally open in its range.

1 INTRODUCTION

Deep learning models have recently led to significant practical successes in various fields ranging from computer vision to natural language processing. Despite these significant empirical successes, the theoretical understanding of the behavior of these models is still very limited. While some works have tried to explain these successes through the lenses of *expressivity* of these models by showing their power in learning large class of mappings, other works find the root of the success in the *generalizability* of these models from learning perspective.

From optimization perspective, training deep models typically requires solving non-convex optimization problem due to the "deep" structure of the model. In fact, it has been shown by Blum & Rivest (1989) that training neural networks to global optimality is NP-complete in the worst case even for the simple case of three node networks. Despite this worst case barrier, the practical success of deep learning may suggest that most of the local optima of these models are close to the global optima. In particular, Choromanska et al. (2015) use spin glass theory and empirical experiments to show that the local optima of deep neural netrwork optimization problem are close to the global optima.

In an effort to better understand the landscape of training deep neural networks, Kawaguchi (2016); Lu & Kawaguchi (2017); Yun et al. (2017); Hardt & Ma (2016) studied the linear neural networks and provide sufficient conditions under which critical points (or local optima) of the optimization problem are globally optimal. For non-linear neural networks, various works have shown that when the number of parameters of the model is larger than the data dimension, then the local optima of the resulting optimization problems are easy and they can be found through local search procedures; see, e.g., Soltanolkotabi et al. (2017); Soudry & Carmon (2016); Nguyen & Hein (2017); Xie et al. (2017).

Despite the growing interest in studying the landscape of deep optimization problems, many of the results and mathematical analyses are problem specific and cannot be generalized to other problems

and network structures easily. As a first step toward reaching a unifying theory for these results, we propose the use of open mappings for characterizing the properties of the local optima of an optimization problem.

To study the landscape of shallow/deep models, let us start by the general optimization problem

$$\underset{w \in \mathcal{W}}{\text{minimize}} \ \ell(\mathcal{F}(w)) \tag{1}$$

for *training* of a learning model. Here $\ell(\cdot)$ is the loss function and $\mathcal{F}(\cdot)$ represents a statistical model with parameter w which needs to be learned by solving the above optimization problem. A simple example is the popular linear regression problem:

$$\underset{w}{\text{minimize}} \ \|Xw - y\|_2^2,$$

where y is a given constant response vector and X is a given constant feature matrix. In this example, the loss function is the ℓ_2 loss, i.e., $\ell(z) = \|z - y\|_2^2$, and the fitted model $\mathcal{F}$ is a linear model, i.e., $\mathcal{F}(w) = Xw$. While this linear regression problem is convex and easy, fitting many practical models, such as deep neural networks, requires solving non-trivial non-convex optimization problems. In addition to the training of deep neural networks, the well-studied matrix completion problem also lies in this category of non-convex problems. For this matrix completion problem, Park et al. (2016) shows that the non-convex matrix factorization formulation of the non-square matrix sensing problem has no spurious local optimum under restricted isometry property (RIP) conditions. Similar results were obtained for the symmetric matrix multiplication problem by Ge et al. (2016), and the non-convex factorized low-rank matrix recovery problem by Bhojanapalli et al. (2016). In this paper, we use the local openness of the mapping $\mathcal{F}$ to provide sufficient conditions under which every local optimum is in fact global.

To proceed, let us define our notations that will be used throughout the paper. We use the notation $A_{l,:}$, and $A_{:,l}$ to denote the lth row and column of matrix A respectively. Let $\|A\|$, $\mathcal{N}(A)$, $\mathcal{C}(A)$, r_A be the respective Frobenius norm, null space, column space, and rank of A. Given subspaces U and V, we say $U \perp V$ if U is orthogonal to V, and $U = V^{\perp}$ if U is the orthogonal complement of V. We say matrix $A \in \mathbb{R}^{d_1 \times d_0}$ is rank *deficient* if rank$(A) < \min(d_1, d_0)$, and *full rank* if rank$(A) = \min(d_1, d_0)$. We call a point $W = (W_h, \dots, W_1)$ *non-degenerate* if rank$(W) = \min_{0 \le i \le h} \ d_i$, and degenerate if rank$(W) < \ \min_{0 \le i \le h} \ d_i$. We also say a point $W = (W_h, \dots, W_1)$ is a *second order saddle point* if the hessian of the loss function at W has a negative eigenvalue.

Let us start by briefly highlighting two problems which will be used as a motivation for our analysis:

Feedforward Neural Networks: Consider the following multiple layer feedforward neural network optimization problem:

$$\underset{W}{\text{minimize}} \ \frac{1}{2}\|\mathcal{F}_h(W) \ - \ Y\|_F^2 \quad \text{with} \quad \mathcal{F}_1(W) \triangleq \sigma_1(W_1 X), \quad \mathcal{F}_k(W) \triangleq \sigma_k\big(W_k \mathcal{F}_{k-1}(W)\big),$$

for $k \in [2, h]$, where $\sigma_k(\cdot), k = 1, \dots, h$, are the activation functions for different layers, $W = \big(W_i\big)_{i=1}^h$, $W_i \in \mathbb{R}^{d_i \times d_{i-1}}$ are the weight matrices, $X \in \mathbb{R}^{d_0 \times n}$ is the input training data, and $Y \in \mathbb{R}^{d_h \times n}$ is the target training data, see Goodfellow & Courville (2016). To obtain the representation in (1), we need to set our loss function to the ℓ_2 loss, and set $\mathcal{F} = \mathcal{F}_h$.

A special instance of this optimization problem was studied in Nguyen & Hein (2017), which considers the non-linear neural network with pyramidal structure (i.e. $d_i \le d_{i-1} \ \forall \ i = 1, \dots, h$ and $d_0 \ge n$). (Nguyen & Hein, 2017, Theorem 3.8) shows that under some conditions, among which are the differentiability of the loss function $\ell(\cdot)$ and the activation function $\sigma(\cdot)$, if W is a critical point with W_i's being full row rank then it is a global minimum. In our paper, we show that by relaxing the differentiability condition on both $\ell(\cdot)$ and $\sigma(\cdot)$, we can still obtain a similar result under very minimal set of assumptions.

Another special case is the linear feedforward networks where the mapping $\sigma_k(\cdot)$ is the identity map in all layers:

$$\underset{W}{\text{minimize}} \ \frac{1}{2}\|W_h \cdots W_1 X \ - \ Y\|_F^2. \tag{2}$$

For this optimization problem Lu & Kawaguchi (2017) showed that every local optimum of the objective function is globally optimal under some assumptions. More precisely, by using perturbation

analysis, (Lu & Kawaguchi, 2017, Theorem 2.2) prove that when X and Y are full row rank, every local optimum in problem (2) is a local optimum in problem (3).

$$\begin{array}{ll} \underset{Z \in \mathbb{R}^{d_h \times n}}{\text{minimum}} & \frac{1}{2}||ZX - Y||_F^2 \\ \text{subject to} & \text{rank}(Z) \leq d_p \triangleq \min_{0 \leq i \leq h} d_i \end{array} \tag{3}$$

Moreover, they show that when X is full row rank, every local optimum of problem (3) is a global optimum. Thus, with the sufficient condition of X and Y being both full row rank, every local optimum of problem (2) is a global optimum. Yun et al. (2017) also show that the same result hold when XX^T and YX^T are both full rank. It is in fact not hard to see that one cannot relax the full rankness assumption of Y due to the following simple example:

$$X = I \quad W_3 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad W_2 = [\,0\,], \quad W_1 = [\ 1 \quad 0\ ], \quad Y = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$$

which is a local optimum of 3-layer deep linear model $\big($problem (2) with $h = 3\,\big)$ that is not global. However, if a given local optima is non-degenerate (which is a simple checkable condition), the full rankness of Y can be relaxed. In particular, we will show that if X is full row rank, then every non-degenerate critical point is either a global optimum or a saddle point, thus relaxing the full row rank assumption on Y.

Non-symmetric Matrix Completion and Matrix Factorization: Consider the following non-symmetric matrix completion optimization problem:

$$\underset{W_1 \in \mathbb{R}^{d_1 \times d_0}, W_2 \in \mathbb{R}^{d_2 \times d_1}}{\text{minimize}} \ \frac{1}{2}\|\Omega\left(W_2 W_1 - Y\right)\|_F^2 \tag{4}$$

where Ω is a linear mapping that represents the sensing process and $Y \in \mathbb{R}^{d_2 \times d_0}$ is a low rank target matrix. When Ω only selects a subset of entries, we get the famous Netflix prize problem, see Koren (2009). When Ω is the identity mapping, this problem becomes the low rank matrix estimation problem described in Srebro & Jaakkola (2003), which can be seen as a 2-layer linear neural network optimization problem

$$\underset{W_1 \in \mathbb{R}^{d_1 \times d_0}, W_2 \in \mathbb{R}^{d_2 \times d_1}}{\text{minimize}} \ \frac{1}{2}\|W_2 W_1 X - Y\|_F^2 \tag{5}$$

with $X = I$. In problem (5), the loss function is the ℓ_2 loss, and the mapping $\mathcal{F}$ is defined as $\mathcal{F}(W_1, W_2) = W_2 W_1$. In this paper, we show the following results for problem (5):

- Every degenerate critical point of (5) is either a global minimum or a second-order saddle point. This result can be generalized to general loss function $\ell(\cdot)$.
- We show that for problem (5), if X is full row rank, then every non-degenerate critical point is either a global minimum or a saddle point.

In addition to these results, we completely characterize the local openness of the matrix product map in its range. This result could be used in many other optimization problems for characterizing the local/global equivalence.

2 Mathematical Framework

As discussed in the previous section, we are interested in solving

$$\underset{w \in \mathcal{W}}{\text{minimize}} \ \ell(\mathcal{F}(w)) \tag{6}$$

where $\mathcal{F} : \mathcal{W} \mapsto \mathcal{S}$ is a mapping and $\ell : \mathcal{S} \mapsto \mathbb{R}$ is a loss function. Here we assume the set $\mathcal{W}$ is closed and the mapping $\mathcal{F}$ is continuous. In non-convex scenarios, this optimization problem can only be solved up to "local optima" by local search procedures; see Lee et al. (2016) for an example. In this paper, *we study problems (6) and (7), and provide sufficient conditions under which any*

local optimum is in fact global. To proceed with our analysis, we define the auxiliary optimization problem

$$\underset{s \in \mathcal{S}}{\text{minimize}} \ \ell(s) \tag{7}$$

where $\mathcal{S}$ is the range of the mapping $\mathcal{F}$. Since problem (7) minimizes the function $\ell(\cdot)$ over the range of the mapping $\mathcal{F}$, the global optimal objective values for problems (6) and (7) are the same. Moreover, there is a connection between the global optima of the two optimization problem through the mapping $\mathcal{F}$. However, the connection between the local optima of the two optimization problem is not clear. This connection is in particular important when the local optima of (7) are "nice" (e.g. globally optimal or close to optimal). In what follows, we establish the connection between the local optima of the optimization problems (6) and (7) under some simple sufficient conditions. This connection is then used to study the relation between local and global optima of (6) and (7) for various deep learning models. Let us first define the following concepts:

Definition: A mapping $\mathcal{F} : \mathcal{W} \to \mathcal{S}$ is said to be open, if for every open set $U \in \mathcal{W}$, $\mathcal{F}(U)$ is (relatively) open in $\mathcal{S}$. The mapping $\mathcal{F}(w)$ is said to be locally open at w if for every $\epsilon > 0$, there exists $\delta > 0$ such that $\mathcal{B}_\delta\big(\mathcal{F}(w)\big) \subset \mathcal{F}\big(\mathcal{B}_\epsilon(w)\big)$, where $\mathcal{B}_\delta(w) \subseteq \mathcal{W}$ is a ball with radius δ centered at w, and $\mathcal{B}_\epsilon(\mathcal{F}(w)) \subseteq \mathcal{S}$ is a ball of radius ϵ centered at $\mathcal{F}(w)$. A useful property of (locally) open mappings is that the composition of two (locally) open maps is (locally) open.

We now state a simple intuitive result that allows us to establish a connection between the local optima of (6) and (7).

Observation 1. *Suppose $\mathcal{F}(\,\bar{w}\,)$ is locally open at $\bar{w}$. If $\bar{w}$ is a local minimum of problem (6), then $\bar{s} = \mathcal{F}(\bar{w})$ is a local minimum of problem (7).*

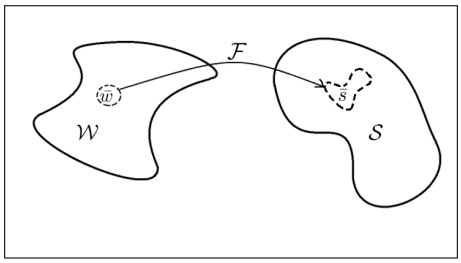

Figure 1: Sketch of the Proof of Observation 1.

Proof. Let $\bar{w}$ be a local minimum of problem (6). Then there exists $\epsilon > 0$ such that $\ell(\mathcal{F}(\bar{w})) \leq \ell(\mathcal{F}(w)),\ \forall w \in \mathcal{B}_\epsilon(\bar{w})$. By the definition of local openness,

$$\exists\, \delta > 0 \ \text{ such that } \mathcal{B}_\delta(\bar{s}) \ \subset \ \mathcal{F}\big(\mathcal{B}_\epsilon(\bar{w})\big) \Rightarrow \ell(\bar{s}) \leq \ell(s),\ \forall s \in \mathcal{B}_\delta(s)$$

which implies $\bar{s}$ is a local minimum of problem (7). □

Furthermore, if every local optimum of (7) is global, the above result implies that any local optimum of (6) is also global. This simple Lemma motivates us to study the local openness of some popular mappings, through which we can establish the local/global equivalence for various classes of optimization problems.

An example of a mapping that is widely used in many optimization problems, such as deep neural networks (2) and matrix completion (5), is the matrix multiplication mapping defined as

$$\begin{aligned} \mathcal{M} : \mathbb{R}^{m\times k} \times \mathbb{R}^{k\times n} \mapsto \mathcal{R}^u_{\mathcal{M}} \triangleq \{\, Z \in \mathbb{R}^{m\times n} \text{ with rank}(Z) \ \leq \ u \triangleq \min(m,n,k)\} \\ \text{with} \quad \mathcal{M}(X,Y) \triangleq XY. \end{aligned} \tag{8}$$

Although, the matrix multiplication mappings $\mathcal{M}(\,X,Y\,)$ appears naturally in deep models and is widely used as a non-convex factorization for rank constrained problems, see Wang et al. (2016); Bhojanapalli et al. (2016); Ge et al. (2016); Srebro & Jaakkola (2003); Sun (2015), to our knowledge,

the complete characterization of the openness of this mapping has not been studied in the optimization literature before. This motivated us to study the openness/local openness of the mapping $\mathcal{M}$ as one of our initial steps.

While the classical open mapping theorem in Rudin (1973) states that surjective continuous linear operators are open, this is not true in general for bilinear mappings such as matrix product. In fact, by providing a simple counterexample of a bilinear mapping that is not open, Horowitz (1975) shows that the linear case cannot be generally extended to multilinear maps. Several papers, see Balcerzak et al. (2013; 2005); Behrends (2011), investigate this bilinear mapping and provide a characterization of the points where this mapping is open. Moreover, Behrends (2017) studies the matrix multiplication mapping $\mathcal{M}$ which is a special example of bilinear mappings and also provides a complete characterization of the points where the mapping is locally open. However, in their study they consider the range of the mapping to be $\mathbb{R}^{m\times n}$, and not $\mathcal{R}^u_{\mathcal{M}}$; which due to the constraint of problem (7) that defines the feasible region to be the range of the mapping $\mathcal{F}$, does not allow us to establish the connection between local optima of problems (6) and (7). For that reason, we study the local openness of the mapping $\mathcal{M}$ in its range $\mathcal{R}^u_{\mathcal{M}}$. An intuitive definition of local openness of $\mathcal{M}(X,Y)$ at (X,Y) in $\mathcal{R}^u_{\mathcal{M}}$ is as follows. We say the multiplication mapping is locally open at (X,Y) if for any small perturbation $\tilde{Z} \in \mathcal{R}^u_{\mathcal{M}}$ of $Z = XY$, there exists a pair $(\tilde{X},\tilde{Y})$, small perturbations of (X,Y), such that $\tilde{Z} = \tilde{X}\tilde{Y}$.

Notice that when $k \geq \min(m,n)$, then $\mathcal{R}^{\min(m,n)}_{\mathcal{M}} = \mathbb{R}^{m\times n}$. However, in the case where $k < \min(m,n)$ the mapping is definitely not locally open in $\mathbb{R}^{m\times n}$, but can still be locally open in $\mathcal{R}^k_{\mathcal{M}}$. As a simple example, consider $X = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$ and $Y = \begin{bmatrix} 1 & 1 \end{bmatrix}$, then there does not exist $\tilde{X}$, $\tilde{Y}$ perturbations of X and Y respectively such that $\tilde{X}\tilde{Y} = \tilde{Z}$ when $\tilde{Z}$ is a full rank perturbation of $Z = XY$; however, for any rank 1 perturbation $\tilde{Z}$ of $Z = XY$, we can find a perturbed pair $(\tilde{X},\tilde{Y})$ such that $\tilde{Z} = \tilde{X}\tilde{Y}$. Motivated by Observation 1, we study in the next section the local openness of the mapping $\mathcal{M}$. We later use these results to analyze the behavior of local optima of deep neural networks.

3 Local Openness of Matrix Multiplication Mapping

Consider $X \in \mathbb{R}^{m\times k}$ and $Y \in \mathbb{R}^{k\times n}$ with $k \geq \min(m,n)$. Then the range of the mapping $\mathcal{M}$ is the entire space $\mathbb{R}^{m\times n}$. In this case, (Behrends, 2017, Theorem 2.5) provides a complete characterization of the pairs (X,Y) where the mapping is locally open. However, when $k \leq \min(m,n)$, i.e. the product of the matrix is rank deficient, the characterization of the set of points for which the mapping is locally open remains an unresolved problem. We settled this question in Theorem 3 in this section which provides a complete characterization of points (X,Y) for which the mapping $\mathcal{M}$ is locally open when $k < \min(m,n)$. Let us start by restating the main result in Behrends (2017):

Proposition 2. *[Behrends (2017)]: Assume $k \geq \min(m,n)$, then the the following statements are equivalent:*

1. $\begin{cases} \exists\, X_\epsilon \text{ such that } X_\epsilon Y = 0 \text{ and } X + X_\epsilon \text{ is full row rank.} \\ \qquad\qquad\qquad\qquad \text{or} \\ \exists\, Y_\epsilon \text{ such that } XY_\epsilon = 0 \text{ and } Y + Y_\epsilon \text{ is full column rank.} \end{cases}$
2. $\dim\big(\mathcal{N}(X) \cap \mathcal{C}(Y)\big) \leq k - m \quad \text{or} \quad n - (\text{rank}(Y) - \dim\big(\mathcal{N}(X) \cap \mathcal{C}(Y)\big) \leq k - \text{rank}(X).$
3. *$\mathcal{M}(X,Y)$ is locally open at (X,Y).*

The above proposition provides a checkable condition which completely characterizes the local openness of the mapping $\mathcal{M}$ at different points when the range of the mapping is the entire space. Now, let us state our result that characterizes the local openness of the mapping $\mathcal{M}$ in its range when $k < \min\{m,n\}$.

Theorem 3. *Let $X \in \mathbb{R}^{m\times k}$, $Y \in \mathbb{R}^{k\times n}$, and $k < \min(m,n)$. Then if rank$(X) \neq$ rank(Y), $\mathcal{M}(X,Y)$ is not locally open at (X,Y). Else if rank$(X) =$ rank(Y), then the following statements are equivalent:*

i) $\begin{cases} (\mathcal{A}_X): \exists\, X_\epsilon \text{ such that } X_\epsilon Y = 0 \text{ and } X + X_\epsilon \text{ is full column rank.} \\ \qquad\qquad\qquad\qquad \text{and} \\ (\mathcal{A}_Y): \exists\, Y_\epsilon \text{ such that } X Y_\epsilon = 0 \text{ and } Y + Y_\epsilon \text{ is full row rank.} \end{cases}$

ii) $dim\big(\mathcal{N}(X) \cap \mathcal{C}(Y)\big) = 0 \Leftrightarrow dim\big(\mathcal{N}(Y^T) \cap \mathcal{C}(X^T)\big) = 0.$

iii) $Z = \mathcal{M}(X, Y)$ *is locally open at* (X, Y) *in* $\mathcal{R}^u_{\mathcal{M}}$.

As previously mentioned, local openness can be described in terms of perturbation analysis. For example, $\mathcal{M}(X, Y)$ is locally open at (X, Y) if for a given $\epsilon > 0$, there exists $\delta > 0$ such that for any $\tilde{Z} = Z + R_\delta \in \mathcal{R}^u_{\mathcal{M}}$ with $||R_\delta|| \leq \delta$, there exists X_0, Y_0 with $||X_0|| \leq \epsilon$, $||Y_0|| \leq \epsilon$, such that $\tilde{Z} = (X + X_0)(Y + Y_0)$. As a perturbation bound on δ, we show that for any locally open pair (X, Y), given an ϵ, we need to choose δ to be of order ϵ. The details of our analysis can be found in the proof of Theorem 3 in Appendix $A.2$.

Remark 1 It follows from Theorem 3 that when X is full column rank, and Y is full row rank, the mapping $\mathcal{M}(X, Y)$ is locally open at (X, Y). This result was observed in other works; see, e.g., (Sun, 2015, Proposition 4.2). Also when $k < \min(m, n)$ if only one of the two matrices is full rank, then the mapping is not locally open. We have showed this result in the proof of Theorem 3, and below is a simple example:

Let

$$X_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \quad Y_1 = [\,0\,,\,0\,], \quad X_1 Y_1 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, \quad R_\delta = \begin{bmatrix} \delta & 0 \\ 0 & 0 \end{bmatrix},$$

then $X_1 Y_1 + R_\delta$ is feasible. On the other hand, for a perturbation $X_1^\epsilon = \begin{bmatrix} \epsilon_1 \\ \epsilon_2 \end{bmatrix}$ and $Y_1^\epsilon = [\,\epsilon_3\,,\,\epsilon_4\,]$, we have

$$(X_1 + X_1^\epsilon)(Y_1 + Y_1^\epsilon) = \begin{bmatrix} (1+\epsilon_1)\epsilon_3 & (1+\epsilon_1)\epsilon_4 \\ (1+\epsilon_2)\epsilon_3 & (1+\epsilon_2)\epsilon_4 \end{bmatrix}.$$

Hence, in order for this perturbation to be equal to $X_1 Y_1 + R_\delta$, we need ϵ_3 to be different from zero. However, when ϵ_3 is different from zero, for small enough ϵ_2, there does not exist such an X_1^ϵ, Y_1^ϵ, or equivalently, $\mathcal{M}(X_1, Y_1)$ is not locally open at (X_1, Y_1). Similarly, $X_2 = Y_1^T$ and $Y_2 = X_1^T$ constitutes an example of a rank deficient X and full rank Y for which $\mathcal{M}(X, Y)$ is not locally open.

In the next sections, we use our local openness result to characterize the cases where the local optima of various training optimization problem of the form (6) are globally optimal.

4 Non-linear Deep Neural Network with a Pyramidal Structure:

Consider the non-linear deep neural network optimization problem with a pyramidal structure

$$\underset{W}{\text{minimize}}\ \ell\big(\mathcal{F}_h(W)\big) \quad \text{with} \quad \mathcal{F}_1(W) \triangleq \boldsymbol{\sigma}_1(W_1 X); \quad \mathcal{F}_k(W) \triangleq \boldsymbol{\sigma}_k\big(W_k \mathcal{F}_{k-1}(W)\big), \tag{9}$$

for $k \in [2, h]$, where $\boldsymbol{\sigma}_k(\cdot)$ is the activation function applied component-wise, i.e., $\boldsymbol{\sigma}_k(B) = [\sigma_k(B_{ij})]_{i,j}$ with $\sigma_k : \mathbb{R} \mapsto \mathbb{R}$. Here $W = \big(W_i\big)_{i=1}^h$ where $W_i \in \mathbb{R}^{d_i \times d_{i-1}}$ is the weight matrix of layer i, and $X \in \mathbb{R}^{d_0 \times n}$ is the input training data. In this section, we consider the pyramidal network structure with $d_0 > n$ and $d_i \leq d_{i-1}$ for $1 \leq i \leq h$; see Nguyen & Hein (2017).

First notice that when X is full column rank, the image of the mapping $\mathcal{F}_h$ is in fact the entire space $\mathbb{R}^{d_h \times n}$ and hence every local optima of the auxiliary optimization problem (7) is global. We now show that when W_i's are all full row rank and $\sigma(\cdot)$ is invertible, the mapping $\mathcal{F}_h$ is locally open at W.

Lemma 4. *Assume the functions* $\sigma_k(\cdot) : \mathbb{R} \mapsto \mathbb{R}$ *are invertible. Then the mapping* $\mathcal{F}_h$ *defined in (9) is locally open at the point* $W = (W_1, \ldots, W_h)$ *if* W_i*'s are all full row rank.*

Before proving this result, we would like to remark that many of the popular activation functions such as logsitic, tangent hyperbolic, and leaky ReLu are invertible and satisfy the assumptions of this lemma.

Proof. Let us prove by induction. Since linear mappings are open, and since $\sigma_1(\cdot)$ is invertible; by using the composition property of open maps, we get that $\mathcal{F}_1$ is open.

Assume $\mathcal{F}_{k-1}\Big(\big(W_i\big)_{i=1}^{k-1}\Big)$ is locally open at $\big(W_i\big)_{i=1}^{k-1}$, then using Proposition 2, due to the full row rankness of W_k, the mapping $W_k\mathcal{F}_{k-1}\Big(\big(W_i\big)_{i=1}^{k-1}\Big)$ is locally open at $\big(W_k, (W_i)_{i=1}^{k-1}\big)$. Using the composition property of open maps and invertibility of $\sigma_k(\cdot)$, we get $\mathcal{F}_k\Big((W_i)_{i=1}^{k}\Big)$ is locally open at $\big(W_i\big)_{i=1}^{k}$. □

Thus, by Observation 1, if W is a local optimum of problem (9) with W_i's being full row rank, then $Z = \mathcal{F}_h(W)$ is a local optimum of the corresponding auxiliary problem:

$$\underset{Z\in\mathbb{R}^{d_h\times n}}{\text{minimize}}\ \ell(Z)$$

and is consequently a global optimum of problem (9) when the loss function $\ell(\cdot)$ is convex. Nguyen & Hein (2017) show that every critical point W of problem (9) with W_i's being full row rank is a global optimum when both $\sigma(\cdot)$ and $\ell(\cdot)$ are differentiable. Our result relaxes the differentiability assumption on both the activation and loss functions; however, we can only show all local optima are global.

5 Two-Layer Linear Neural Network

Consider the two layer linear neural network optimization problem

$$\underset{W}{\text{minimize}}\ \frac{1}{2}\|W_2W_1X\ -\ Y\ \|_F^2 \tag{10}$$

where $W_2 \in \mathbb{R}^{d_2\times d_1}$ and $W_1 \in \mathbb{R}^{d_1\times d_0}$ are weight matrices, $X \in \mathbb{R}^{d_0\times n}$ is the input data, and $Y \in \mathbb{R}^{d_2\times n}$ is the target training data. Using our transformation, the corresponding auxiliary optimization problem can be written as

$$\begin{array}{ll} \underset{Z}{\text{minimum}} & \frac{1}{2}||ZX - Y||_F^2 \\ \text{subject to} & \text{rank}(Z) \leq \min(d_2, d_1, d_0) \end{array} \tag{11}$$

(Kawaguchi, 2016, Theorem 2.3) shows that when XX^T and YX^T are full rank, $d_2 \leq d_0$, and when $YX^T(XX^T)^{-1}XY^T$ has d_2 distinct eigenvalues, every local optimum is global and all saddle points are second order saddles. While the local/global equivalence result holds for deeper networks, the property that all saddles are second order does not hold in that case. Another related result by (Yun et al., 2017, Theorem 2.2) shows that when XX^T, YX^T, and $YX^T(XX^T)^{-1}XY^T$ are full rank, every local optimum of a linear deep network is global. Moreover, they provide necessary and sufficient conditions for a critical point to be a global minimum. However, we notice that the full rankness assumption on YX^T was not used in showing the result for non-degenerate critical points and thus can be relaxed in that case. In this section, without any assumption on Y, we reconstruct the proof that shows the latter result for 2-layer networks using local openness, and then show a similar result for the degenerate case. The result for the degenerate case hold when replacing the square loss error by a general convex loss function as we will see in Colorollary 6. The proofs of the theorem and corollary stated below can be found in Appendix $A.1$

Theorem 5. *Every degenerate critical point of problem (10) is either a global minimum or a second order saddle. If X is full row rank, then every non-degenerate critical point of problem (10) is either a global minimum or a saddle point.*

Corollary 6. *Let the square loss error in (10) be replaced by a general convex loss function $\ell(\cdot)$. Then every degenerate critical point is either a global minimum or a second order saddle.*

Baldi & Hornik (1989) and Srebro & Jaakkola (2003) show the same result when both X and Y are full row rank. Theorem 5 generalizes their results by relaxing the assumptions on Y. Another implication of this theorem is the problem of the fully observed non-symmetric matrix completion. Ge et al. (2016) studied the symmetric matrix completion problem and showed that every local minimum is global (but not the non-symmetric case). By setting $X = I$, the theorem above extends the results in Ge et al. (2016) to the non-symmetric case. We summarize our results in the following chart:

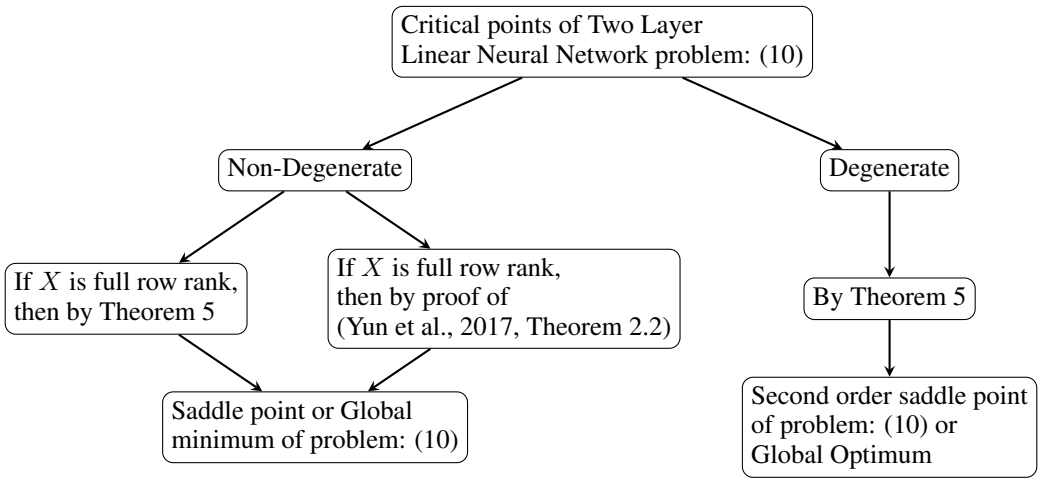

6 Multi-Layer Deep Linear Neural Network

Consider the training problem of multi-layer deep linear neural networks:

$$\underset{W}{\text{minimize}} \ \frac{1}{2}\|W_h \cdots W_1 X \ - \ Y \|_F^2. \tag{12}$$

Here $W = \big(W_i\big)_{i=1}^{h}$, $W_i \in \mathbb{R}^{d_i \times d_{i-1}}$ are the weight matrices, $X \in \mathbb{R}^{d_0 \times n}$ is the input training data, and $Y \in \mathbb{R}^{d_h \times n}$ is the target training data. Based on our general framework, the corresponding auxiliary optimization problem is given by

$$\begin{array}{ll} \underset{Z \in \mathbb{R}^{d_h \times n}}{\text{minimum}} & \frac{1}{2}||ZX \ - \ Y||_F^2 \\ \text{subject to} & \text{rank}(Z) \leq d_p \triangleq \min_{0 \leq i \leq h} \ d_i \end{array} \tag{13}$$

Lu & Kawaguchi (2017) attempted to prove that when X and Y are full row rank, every local minimum is global. However, the derivation does not constitute a formal proof of the desired result; this is because it uses Lemma 3.3 which states that given a full rank matrix $\bar{M}$ with singular value decomposition (SVD) $\bar{M} = \bar{U}\bar{\Sigma}\bar{V}^T$, for any perturbation M of $\bar{M}$, there exists U, Σ, and V perturbations of $\bar{U}$, $\bar{\Sigma}$, and $\bar{V}$ respectively, such that $U\Sigma V^T = M$ is an SVD in M. This statement is not true in general due to the follwowing counterexample provided by Stewart (1998):
Let

$$\bar{M} = \begin{bmatrix} 1 & 0 \\ 0 & 1+\epsilon \end{bmatrix} \quad \text{and} \quad M = \begin{bmatrix} 1 & \epsilon \\ \epsilon & 1 \end{bmatrix}.$$

Then the right singular vectors of $\bar{M}$ are given by

$$\bar{V} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},$$

while the right singular vectors of M are given by

$$V = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}.$$

However, this Lemma was only used to prove (Lu & Kawaguchi, 2017, Theorem 3.1) which, as we will show next, can be derived using Proposition 2 and Theorem 3. Before proceeding to the proof we define the following mapping:

$$\mathcal{M}_{i,j}(W_i, \dots, W_j) : \{W_i, \dots, W_j\} \to \mathcal{R}_{\mathcal{M}_{i,j}} \triangleq \{Z \in \mathbb{R}^{d_i \times d_{j-1}} \,|\, \text{rank}(Z) \leq \min_{j-1 \leq l \leq i} d_l \,\}$$
$$\text{with} \quad \mathcal{M}_{i,j}(W_i, \dots, W_j) = W_i \cdots W_j \quad \text{for } i > j$$

Now let $Z = W_h W_{h-1} \dots W_1 = \mathcal{M}(\mathcal{M}_{h,p+1}, \mathcal{M}_{p,1})$ be a rank d_p matrix, we state Theorem 3.1 of Lu & Kawaguchi (2017) using our notation.

Lemma 7. *If W is non-degenerate, then $\mathcal{M}_{h,1}(W) = W_h \cdots W_1$ is locally open at W.*

Proof. We construct a proof by induction on h to show the desired result. When $h = 2$, we either have $d_1 < \min(d_2, d_0)$ or $d_1 \geq \min(d_2, d_0)$. In the first case, since

$$d_1 = \text{rank}(\bar{W}_2 \bar{W}_1) \leq \text{rank}(\bar{W}_1) \leq d_1, \quad \text{and} \quad d_1 = \text{rank}(\bar{W}_2 \bar{W}_1) \leq \text{rank}(\bar{W}_2) \leq d_1,$$

then by Theorem 3, we get $\mathcal{M}_{2,1}$ is locally open at (W_2, W_1). In the second case, and since either

$$d_2 = \text{rank}(\bar{W}_2 \bar{W}_1) \leq \text{rank}(\bar{W}_2) \leq d_2, \quad \text{or} \quad d_0 = \text{rank}(\bar{W}_2 \bar{W}_1) \leq \text{rank}(\bar{W}_1) \leq d_0,$$

then by Proposition 2, $\mathcal{M}_{2,1}$ is locally open at (W_2, W_1).

Now assume the result holds for the product of h matrices $\mathcal{M}_{h,1}(W)$, we show it is true for $\mathcal{M}_{h+1,1}(W)$.

Since

$$d_p = \text{rank}(\bar{W}_h, \dots, \bar{W}_1) \leq \text{rank}(W_{p+1} W_p) \leq d_p$$

then using Proposition 2, we get $\mathcal{M}_{p+1,p}$ is locally open at (W_{p+1}, W_p). So we can replace $W_{p+1} W_p$ by a new matrix Z_p with rank d_p. Then by induction hypothesis, the product mapping $\mathcal{M}_{h+1,1} = W_{h+1} \cdots W_{p+2} Z_p W_{p-1} \cdots W_1$ is locally open at W. Since the composition of locally open maps is locally open, the result follows.

□

We now demonstrate our main results for this optimization problem which shows that under a set of necessary conditions, every critical point of problem (12) is either a saddle or a global minimum. Although the result for the non-degenerate case directly follows from (Yun et al., 2017, Theorem 2.2), we provide in Lemma 8 a more intuitive proof that uses local openness of $\mathcal{M}$.

Lemma 8. *Assume X is full row rank, then every non-degenerate critical point of (12) is either a saddle or a global minimum.*

Proof. Suppose $W = (W_h, \dots, W_1)$ is a non-degenerate local minimum. Then it follows by Lemma 7 that $\mathcal{M}_{h,1}$ is locally open at W. Then by Lemma 1, $Z = \mathcal{M}_h(W_h, \dots, W_1)$ is a local optimum of problem (13) which is in fact global by (Lu & Kawaguchi, 2017, Theorem 2.2). □

REFERENCES

M. Balcerzak, A. Wachowicz, and W. Wilczyński. Multiplying balls in the space of continuous functions on $[0, 1]$. *Studia Mathematica*, 170:203–209, 2005.

M. Balcerzak, A. Majchrzycki, and A. Wachowicz. Openness of multiplication in some function spaces. *Taiwanese J. Math*, 17:1115–1126, 2013.

P. Baldi and K. Hornik. Neural networks and principal component analysis: Learning from examples without local minima. *Neural networks*, 2(1):53–58, 1989.

E. Behrends. Products of n open subsets in the space of continuous functions on $[0, 1]$. *Studia Mathematica*, 204:73–95, 2011.

E. Behrends. Where is matrix multiplication locally open? *Linear Algebra and its Applications*, 517:167–176, 2017.

S. Bhojanapalli, B. Neyshabur, and N. Srebro. Global optimality of local search for low rank matrix recovery. In *Advances in Neural Information Processing Systems*, pp. 3873–3881, 2016.

A. Blum and R. L. Rivest. Training a 3-node neural network is np-complete. In *Advances in neural information processing systems*, pp. 494–501, 1989.

A. Choromanska, M. Henaff, M. Mathieu, G. B. Arous, and Y. LeCun. The loss surfaces of multilayer networks. In *Artificial Intelligence and Statistics*, pp. 192–204, 2015.

R. Ge, J. D. Lee, and T. Ma. Matrix completion has no spurious local minimum. In *Advances in Neural Information Processing Systems*, pp. 2973–2981, 2016.

I. Goodfellow and A. Courville. Deep learning. *Book in preparation for MIT Press, Cambridge*, 2016.

M. Hardt and T. Ma. Identity matters in deep learning. *arXiv preprint arXiv:1611.04231*, 2016.

C. Horowitz. An elementary counterexample to the open mapping principle for bilinear maps. *Proceedings of the American Mathematical Society*, 53(2):293–294, 1975.

K. Kawaguchi. Deep learning without poor local minima. In *Advances in Neural Information Processing Systems*, pp. 586–594, 2016.

Y. Koren. The bellkor solution to the netflix grand prize. *Netflix prize documentation*, 81:1–10, 2009.

J. D. Lee, M. Simchowitz, M. I. Jordan, and B. Recht. Gradient descent only converges to minimizers. In *Conference on Learning Theory*, pp. 1246–1257, 2016.

H. Lu and K. Kawaguchi. Depth creates no bad local minima. *arXiv preprint arXiv:1702.08580*, 2017.

Q. Nguyen and M. Hein. The loss surface of deep and wide neural networks. *arXiv preprint arXiv:1704.08045*, 2017.

D. Park, A. Kyrillidis, C. Caramanis, and S. Sanghavi. Non-square matrix sensing without spurious local minima via the burer-monteiro approach. *arXiv preprint arXiv:1609.03240*, 2016.

W. Rudin. Functional analysis, mcgraw-hill series in higher mathematics. 1973.

M. Soltanolkotabi, A. Javanmard, and J. D. Lee. Theoretical insights into the optimization landscape of over-parameterized shallow neural networks. *arXiv preprint arXiv:1707.04926*, 2017.

D. Soudry and Y. Carmon. No bad local minima: Data independent training error guarantees for multilayer neural networks. *arXiv preprint arXiv:1605.08361*, 2016.

N. Srebro and T. Jaakkola. Weighted low-rank approximations. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pp. 720–727, 2003.

G. W. Stewart. Perturbation theory for the singular value decomposition. Technical report, 1998.

R. Sun. *Matrix completion via nonconvex factorization: Algorithms and theory*. PhD thesis, University of Minnesota, 2015.

L. Wang, X. Zhang, and Q. Gu. A unified computational and statistical framework for nonconvex low-rank matrix estimation. *arXiv preprint arXiv:1610.05275*, 2016.

B. Xie, Y. Liang, and L. Song. Diverse neural network learns true target functions. In *Artificial Intelligence and Statistics*, pp. 1216–1224, 2017.

C. Yun, S. Sra, and A. Jadbabaie. Global optimality conditions for deep neural networks. *arXiv preprint arXiv:1707.02444*, 2017.

7 APPENDIX A.1

7.1 PROOFS OF THE THEOREM 5

Proof. The proof for the degenerate case is done by constructing a descent direction if the point is critical but not global. Let $(\bar{W}_2,\ \bar{W}_1)$ be a degenerate critical point, i.e. rank$(\bar{W}_2\bar{W}_1) < \min(d_2, d_1, d_0)$. Then, based on the dimensions of d_0, d_1, and d_2, we have one of the following cases:

$d_2 < d_1$ then $\exists\ b \neq 0$ such that $b \in \mathcal{N}\big(\bar{W}_2\big)$
$d_0 < d_1$ then $\exists\ b \neq 0$ such that $b \in \mathcal{N}\big(\bar{W}_1^T\big)$
$d_1 \leq d_2$, and $d_1 \leq d_0$ then either $\bar{W}_2$ is rank deficient and $\exists\ b \neq 0$ such that $b \in \mathcal{N}\big(\bar{W}_2\big)$ or $\bar{W}_1$ is rank deficient and $\exists\ b \neq 0$ such that $b \in \mathcal{N}\big(\bar{W}_1^T\big)$

So in all cases either $\mathcal{N}\big(\bar{W}_2\big) \neq 0$ or $\mathcal{N}\big(\bar{W}_1^T\big) \neq 0$. Also, let $\Delta = \bar{W}_2\bar{W}_1X - Y$. If $\Delta X^T = 0$, then by convexity of the square loss error function, the point $(\bar{W}_2, \bar{W}_1)$ is a global minimum of (11). Else, there exists (i, j) such that $\big\langle X_{i,:}, \Delta_{j,:}\big\rangle \neq 0$. We now use first and second order optimality conditions to construct a descent direction when the current critical point is not global.

First order optimality condition: By considering perturbation in the directions $A \in \mathbb{R}^{d_2\times d_1}$ and $B \in \mathbb{R}^{d_1\times d_0}$ for the optimization problem

$$\underset{t}{\text{minimize}}\ \frac{1}{2}\|(W_2 + tA)(W_1 + tB)\ -\ Y\ \|_F^2 \tag{14}$$

we obtain

$$\big\langle\ A\bar{W}_1X\ +\ \bar{W}_2BX\ ,\ \Delta\ \big\rangle = 0, \quad \forall A \in \mathbb{R}^{d_2\times d_1},\ B \in \mathbb{R}^{d_1\times d_0}$$

Second order optimality condition:

$$2\big\langle\ ABX,\ \Delta\ \big\rangle + ||A\bar{W}_1X + \bar{W}_2BX||_F^2\ \geq 0 \quad \forall A \in \mathbb{R}^{d_2\times d_1},\ B \in \mathbb{R}^{d_1\times d_0}$$

Suppose $(\bar{W}_1, \bar{W}_2)$ is a critical point and there exists $b \neq 0$, $b \in \mathcal{N}(\bar{W}_2)$. We define

$$B_{:,l} = \begin{cases} \alpha b & \text{if } l = i, \\ 0 & \text{otherwise} \end{cases} \qquad A_{l,:} = \begin{cases} e_s^T & \text{if } l = j, \\ 0 & \text{otherwise} \end{cases}$$

where α is a scalar constant, $s = \underset{k}{\text{argmax}}\,|b_k|$ and e_s is the sth unit vector. Then, using the second order optimality condition, for $c = ||A\bar{W}_1X||_F^2$, we get

$$\alpha\ \underbrace{b_s}_{\neq 0}\underbrace{\big\langle X_{i,:}, \Delta_{j,:}\big\rangle}_{\neq 0} + c \geq 0$$

Since this is true for every value of α, b_s should be zero which contradicts the assumption on the choice of b. Hence $\mathcal{N}(\bar{W}_2) = \emptyset$.

Similarly, suppose $(\bar{W}_1, \bar{W}_2)$ is a critical point and exists $a^T \neq 0$, $a^T \in \mathcal{N}(\bar{W}_1)$. Let

$$A_{l,:} = \begin{cases} \alpha a^T & \text{if } l = j, \\ 0 & \text{otherwise} \end{cases} \qquad B_{:,l} = \begin{cases} e_s & \text{if } l = i, \\ 0 & \text{otherwise} \end{cases}$$

where α is a constant, $s = \underset{i}{\text{argmax}}\,|a_i|$ and e_s is the sth unit vector. Then, for $c = ||\bar{W}_2BX||_F^2$, we get

$$\alpha\ \underbrace{a_s}_{\neq 0}\big\langle X_{i,:}, \Delta_{j,:}\big\rangle + c \geq 0$$

Using the same argument, we can show that $(\bar{W}_2, \bar{W}_1)$ is a second order saddle point of problem (10).

We now show the result for the degenerate case. Let $(\bar{W}_2,\ \bar{W}_1)$ be a non-degenerate local optimum, i.e. rank$(\bar{W}_2\bar{W}_1) = \min(d_2, d_1, d_0)$. Then it follows by Lemma 7 that the matrix multiplication $\mathcal{M}(\bar{W}_2\bar{W}_1)$ is locally open at $(\bar{W}_2, \bar{W}_1)$. Then by Lemma 1, $Z = \bar{W}_2\bar{W}_1$ is a local optimum of problem (13) which is in fact global by (Lu & Kawaguchi, 2017, Theorem 2.2).

□

7.2 PROOF OF COROLLARY 6

Proof. First order optimality condition:

$$\langle A\bar{W}_1X + \bar{W}_2BX, \nabla\ell(\bar{W}_2\bar{W}_1X - Y)\rangle = 0 \ \forall A \in \mathbb{R}^{d_2\times d_1},\ B \in \mathbb{R}^{d_1\times d_0}$$

Second order optimality condition:

$$2\langle\, ABX,\ \nabla\ell(\bar{W}_2\bar{W}_1X - Y)\,\rangle + h\big(A\bar{W}_1X,\ \bar{W}_2BX, \bar{W}_2\bar{W}_1X\big) \geq 0 \ \ \forall A \in \mathbb{R}^{d_2\times d_1},\ B \in \mathbb{R}^{d_1\times d_0}$$

where $h(\cdot)$ is a function that has a tensor representation. But we only need to know that it is a function of $A\bar{W}_1X$, $\bar{W}_2BX$, and $\bar{W}_2\bar{W}_1X$.

If $\nabla\ell(\bar{W}_2\bar{W}_1X - Y)X^T$, then by convexity of $\ell(\cdot)$, $(\bar{W}_2, \bar{W}_1)$ is a global minimum. Otherwise, there exists (i,j) such that $\big\langle X_{i,:}, \big(\nabla\ell(\bar{W}_2\bar{W}_1X - Y)\big)_{j,:}\big\rangle \neq 0$. Using the same former argument in proof of Theorem 5, we choose A and B such that $h(A\bar{W}_1X,\ \bar{W}_2BX,\ \bar{W}_2\bar{W}_1X)$ is some constant that does not depend on α, and $\big\langle\, ABX, \nabla\ell(\bar{W}_2\bar{W}_1X - Y)\,\big\rangle = \alpha \underbrace{\big\langle X_{i,:}, \big(\nabla\ell(\bar{W}_2\bar{W}_1X - Y)\big)_{j,:}\big\rangle \neq 0}_{\neq 0}$.

Then by proper choice of α we show that the point $(\bar{W}_2, \bar{W}_1)$ is a second order saddle point.

□

8 APPENDIX A.2

8.1 PROOF OF THEOREM 3

In this section, we prove Theorem 3 which provides a complete characterization of points $(X \in \mathbb{R}^{m\times k}, Y \in \mathbb{R}^{k\times n})$ where the matrix multiplication mapping $\mathcal{M}(X,Y)$ is locally open in its range $\mathcal{R}^u_{\mathcal{M}} = \{Z \in \mathbb{R}^{m\times n} \text{ with rank}(Z) \leq u \triangleq \min(m,k,n)\}$. Our result shows that when $u = k < \min(m,n)$, $\mathcal{M}$ is locally open at a point (X,Y) if and only if the following two conditions both hold:

$$\begin{cases} (\mathcal{A}_X) : \exists\, X_\epsilon \text{ such that } X_\epsilon Y = 0 \text{ and } X + X_\epsilon \text{ is full column rank.} \\ \qquad\qquad\qquad\qquad \text{and} \\ (\mathcal{A}_Y) : \exists\, Y_\epsilon \text{ such that } XY_\epsilon = 0 \text{ and } Y + Y_\epsilon \text{ is full row rank.} \end{cases}$$

To do so, we first show that the local openness of $\mathcal{M}(X,Y)$ at (X,Y) is equivalent to the local openness of $\mathcal{M}(U^TX, YV)$ at (U^TX, YV) where $U \in \mathbb{R}^{m\times m}$ and $V \in \mathbb{R}^{n\times n}$ are left and right singular vectors of the product XY respectively. This allows us to focus our study on the local openness of the mapping $\mathcal{M}$ to matrix pairs whose product is a diagonal matrix. We then show in Lemma 10 that when $u = k$ and $r_X = r_Y$, $(\mathcal{A}_X)$ holds if and only if $(\mathcal{A}_Y)$ holds. Finally, we show in Proposition 13, that these conditions hold if and only if the mapping $\mathcal{M}(X,Y)$ is locally open at (X,Y).

Lemma 9. *Consider $X \in \mathbb{R}^{m\times k}$, $Y \in \mathbb{R}^{k\times n}$, and let $XY = U\Sigma V^T$ with $U \in \mathbb{R}^{m\times m}$, $V \in \mathbb{R}^{n\times n}$, and $\Sigma \in \mathbb{R}^{m\times n}$ be an SVD decomposition of the matrix product XY. Then,*

$$\mathcal{M}(X,Y) \textit{ is locally open at } (X,Y) \Leftrightarrow \mathcal{M}(U^TX, YV) \textit{ is locally open at } (U^TX, YV)$$

Proof. We first show the direction " $\Rightarrow$ ". Suppose $\mathcal{M}(X,Y)$ is locally open at (X,Y), then by definition of local openness, for any given $\epsilon > 0$, there exists $\delta > 0$ such that

$$\mathbb{B}_\delta(XY) \cap \mathcal{R}^u_{\mathcal{M}} \subseteq \mathcal{M}\big(\mathbb{B}_\epsilon(X), \mathbb{B}_\epsilon(Y)\big) = \{(X + X_\epsilon)(Y + Y_\epsilon) \,|\, \|X_\epsilon\| \leq \epsilon,\ \|Y_\epsilon\| \leq \epsilon\} \tag{15}$$

We now show that

$$\mathbb{B}_\delta(U^TXYV) \cap \mathcal{R}^u_{\mathcal{M}} \subseteq \mathcal{M}\big(\mathbb{B}_\epsilon(U^TX), \mathbb{B}_\epsilon(YV)\big).$$

Consider $\tilde{\Sigma} \in \mathbb{B}_\delta(U^TXYV) \cap \mathcal{R}^u_{\mathcal{M}}$, i.e. $\tilde{\Sigma} = U^TXYV + R_\delta$ with rank($\tilde{\Sigma}$)$\leq u$ and $\|R_\delta\| \leq \delta$. Since U and V are orthonormal left and right singular vectors, we get $U\tilde{\Sigma}V^T = XY + UR_\delta V^T$ with rank($U\tilde{\Sigma}V^T$)=rank($\tilde{\Sigma}$)$\leq u$ and $\|UR_\delta V^T\| = \|R_\delta\| \leq \delta$. According to (15), this translates to the following:

$$U\tilde{\Sigma}V^T \in \mathbb{B}_\delta(XY) \cap \mathcal{R}^u_{\mathcal{M}} \subseteq \{(X + X_\epsilon)(Y + Y_\epsilon) \,|\, \|X_\epsilon\| \leq \epsilon,\ \|Y_\epsilon\| \leq \epsilon\}$$

which implies,

$$\begin{aligned}\tilde{\Sigma} &\in \{(U^T X + U^T X_\epsilon)(YV + Y_\epsilon V) \mid \|U^T X_\epsilon\| \leq \epsilon,\ \|Y_\epsilon V\| \leq \epsilon\} \\ &= \{(U^T X + U^T X_\epsilon)(YV + Y_\epsilon V) \mid \|X_\epsilon\| \leq \epsilon,\ \|Y_\epsilon\| \leq \epsilon\}\end{aligned}$$

Since $\tilde{\Sigma}$ was arbitrarily chosen, we get $\mathbb{B}_\delta(U^T XYV) \cap \mathcal{R}^u_{\mathcal{M}} \subseteq \mathcal{M}\big(\mathbb{B}_\epsilon(U^T X), \mathbb{B}_\epsilon(YV)\big)$.

Proving the converse direction " $\Leftarrow$ " is similar. Suppose $\mathcal{M}(U^T X, YV)$ is locally open at $(U^T X, YV)$, then by definition of local openness, for any given $\epsilon > 0$, there exists $\delta > 0$ such that

$$\mathbb{B}_\delta(U^T XYV) \cap \mathcal{R}^u_{\mathcal{M}} \subseteq \mathcal{M}\big(\mathbb{B}_\epsilon(U^T X), \mathbb{B}_\epsilon(YV)\big) = \{(U^T X + X_\epsilon)(YV + Y_\epsilon) \mid \|X_\epsilon\| \leq \epsilon,\ \|Y_\epsilon\| \leq \epsilon\} \tag{16}$$

We now show that

$$\mathbb{B}_\delta(XY) \cap \mathcal{R}^u_{\mathcal{M}} \subseteq \mathcal{M}\big(\mathbb{B}_\epsilon(X), \mathbb{B}_\epsilon(Y)\big).$$

Consider $\tilde{Z} \in \mathbb{B}_\delta(XY) \cap \mathcal{R}^u_{\mathcal{M}}$, i.e. $\tilde{Z} = XY + R_\delta$ with rank($\tilde{Z}$)$\leq u$ and $\|R_\delta\| \leq \delta$. Since U and V are orthonormal left and right singular vectors, we get $U^T \tilde{Z} V = U^T XYV + U^T R_\delta V$ with rank($U^T \tilde{Z} V$)=rank($\tilde{Z}$)$\leq u$ and $\|U^T R_\delta V\| = \|R_\delta\| \leq \delta$. According to (16), this translates to the following:

$$U^T \tilde{Z} V \in \mathbb{B}_\delta(U^T XYV) \cap \mathcal{R}^u_{\mathcal{M}} \subseteq \mathcal{M}\big(\mathbb{B}_\epsilon(U^T X), \mathbb{B}_\epsilon(YV)\big)$$

which implies,

$$\begin{aligned}\tilde{Z} &\in \{(X + UX_\epsilon)(Y + Y_\epsilon V^T) \mid \|UX_\epsilon\| \leq \epsilon,\ \|Y_\epsilon V^T\| \leq \epsilon\} \\ &= \{(X + UX_\epsilon)(Y + Y_\epsilon V^T) \mid \|X_\epsilon\| \leq \epsilon,\ \|Y_\epsilon\| \leq \epsilon\}\end{aligned}$$

Since $\tilde{Z}$ was arbitrarily chosen, we get $\mathbb{B}_\delta(XY) \cap \mathcal{R}^u_{\mathcal{M}} \subseteq \mathcal{M}\big(\mathbb{B}_\epsilon(X), \mathbb{B}_\epsilon(Y)\big)$, which completes the proof.

□

We next derive some facts that we will later need to prove the main result in Theorem 3.

Lemma 10. *Consider $X \in \mathbb{R}^{m\times k}$, $Y \in \mathbb{R}^{k\times n}$, $k < \min(m,n)$ with rank(X) = rank$(Y) = r$, then*

i) $(\mathcal{A}_Y)$*: $\exists\ Y_\epsilon$ such that $XY_\epsilon = 0$ and $Y + Y_\epsilon$ is full row rank $\Leftrightarrow$ dim$\big(\mathcal{N}(X) \cap \mathcal{C}(Y)\big) = 0$.*

ii) $(\mathcal{A}_X)$*: $\exists\ X_\epsilon$ such that $X_\epsilon Y = 0$ and $X + X_\epsilon$ is full column rank $\Leftrightarrow$ dim$\big(\mathcal{N}(Y^T) \cap \mathcal{C}(X^T)\big) = 0$.*

iii) dim$\big(\mathcal{N}(X) \cap \mathcal{C}(Y)\big) = 0 \Leftrightarrow$ dim$\big(\mathcal{N}(Y^T) \cap \mathcal{C}(X^T)\big) = 0$.

Proof. $i)$ We first show the direction " $\Rightarrow$ ". Consider $X \in \mathbb{R}^{m\times k}, Y \in \mathbb{R}^{k\times n}$ both rank r matrices. Suppose $(\mathcal{A}_Y)$ holds, then

$$\mathcal{C}(Y_\epsilon) \subseteq \mathcal{N}(X) \Rightarrow \text{rank}(Y_\epsilon) \leq \dim\big(\mathcal{N}(X)\big) = k - r. \tag{17}$$

Also,

$$k = \text{rank}(Y + Y_\epsilon) \leq \text{rank}(Y) + \text{rank}(Y_\epsilon) = r + \text{rank}(Y_\epsilon) \tag{18}$$

From inequalities (17) and (18), we get

$$k - r \ \leq\ \text{rank}(Y_\epsilon) \ \leq\ k - r \ \Rightarrow\ \text{rank}(Y_\epsilon) \ =\ k - r$$

Note that $\dim\big(\mathcal{C}(Y_\epsilon)\big) = \dim\big(\mathcal{N}(X)\big)$ and $\mathcal{C}(Y_\epsilon) \subseteq \mathcal{N}(X)$, which implies that $\mathcal{C}(Y_\epsilon) = \mathcal{N}(X)$.

Then since $\text{rank}(Y + Y_\epsilon) = \text{rank}(Y) + \text{rank}(Y_\epsilon)$, we get

$$\emptyset = \mathcal{C}(Y_\epsilon) \cap \mathcal{C}(Y) = \mathcal{N}(X) \cap \mathcal{C}(Y) \ \Rightarrow\ \dim\big(\mathcal{N}(X) \cap \mathcal{C}(Y)\big) = 0.$$

We now show the other direction " $\Leftarrow$ ". Without loss of generality, let $Y = \left[Y_0^{k\times r} A^{r\times n-r} \;\; Y_0^{k\times r}\right]$ where columns of Y_0 are linearly independent and let $Y_\epsilon = \epsilon\left[y^1, \dots, y^{k-r}, 0, \dots, 0\right] \in \mathbb{R}^{k\times n}$ be a rank $k-r$ matrix where y^i are unit basis of $\mathcal{N}(\,X\,)$ which yields $\mathcal{C}(\,Y_\epsilon\,) = \mathcal{N}(\,X\,)$. Then since $\dim\big(\mathcal{N}(X) \cap \mathcal{C}(Y)\big) = 0$, $\text{rank}(Y + Y_\epsilon) = k$ for some value of ϵ. This completes the proof.

$ii)$ Note that by setting $X = Y^T$ and $Y = X^T$, the same proof can be used to show (ii).

$iii)$ Notice that

$$\begin{aligned}\dim\Big(\text{span}\big(\mathcal{N}(\,X\,) \cup \mathcal{C}(\,Y\,)\big)\Big) &= \dim\big(\mathcal{N}(\,X\,)\big) + \dim\big(\mathcal{C}(\,Y\,)\big) - \dim\big(\mathcal{N}(\,X\,) \cap \mathcal{C}(\,Y\,)\big)\\ &= k - r + r - \dim\big(\mathcal{N}(\,X\,) \cap \mathcal{C}(\,Y\,)\big)\\ &= k - \dim\big(\mathcal{N}(\,X\,) \cap \mathcal{C}(\,Y\,)\big)\end{aligned}$$

Thus,

$$\begin{aligned}\dim\big(\mathcal{N}(\,X\,) \cap \mathcal{C}(\,Y\,)\big) \neq 0 &\Leftrightarrow \dim\Big(\text{span}\big(\mathcal{N}(\,X\,) \cup \mathcal{C}(\,Y\,)\big)\Big) < k\\ &\Leftrightarrow \exists\, a \text{ such that } a \perp \mathcal{C}(\,Y\,), \text{ and } a \perp \mathcal{N}(\,X\,)\\ &\Leftrightarrow \exists\, a \text{ such that } a \in \mathcal{N}(\,Y^T\,), \text{ and } a \in \mathcal{C}(\,X^T\,)\\ &\Leftrightarrow \dim\big(\mathcal{N}(\,Y^T\big) \cap \mathcal{C}(\,X^T\,)\big) \neq 0.\end{aligned}$$

□

Interestingly, Lemma 10 shows that when $u = k < \min(m, n)$ and X and Y have the same rank, $(\mathcal{A}_X)$ and $(\mathcal{A}_Y)$ are equivalent. The next result shows that if these conditions hold, then r_X is equal to r_Y, and the last $n - r$ rows of $U^T X$ and last $n - r$ columns of YV, where $U \in \mathbb{R}^{m\times m}$ and $V \in \mathbb{R}^{n\times n}$ are the left and right singular vectors of the product XY respectively, are all zeros.

Lemma 11. *Consider $X \in \mathbb{R}^{m\times k}$, $Y \in \mathbb{R}^{k\times n}$ with $k < \min(m, n)$, and let $XY = U\Sigma V^T$ with $U \in \mathbb{R}^{m\times m}$, and $V \in \mathbb{R}^{n\times n}$, and $\Sigma \in \mathbb{R}^{m\times n}$ be an SVD decomposition of the product XY. If*

$$\begin{cases}(\mathcal{A}_X) : \exists\, X_\epsilon \textit{ such that } X_\epsilon Y = 0 \textit{ and } X + X_\epsilon \textit{ is full column rank.}\\ \qquad\qquad\qquad\qquad \textit{and}\\ (\mathcal{A}_Y) : \exists\, Y_\epsilon \textit{ such that } XY_\epsilon = 0 \textit{ and } Y + Y_\epsilon \textit{ is full row rank.}\end{cases}$$

then

$$r_X = r_Y \triangleq r, \qquad \big(YV\big)_{:,r+1:n} = 0, \quad \textit{and} \quad \big(U^T X\big)_{r+1:n,:} = 0$$

Proof. Suppose that $(\mathcal{A}_Y)$ holds, then

$$\mathcal{C}(\,Y_\epsilon\,) \subseteq \mathcal{N}(\,X\,) \Rightarrow \text{rank}(Y_\epsilon) \leq \dim\big(\,\mathcal{N}(\,X\,)\,\big) = k - r_X. \tag{19}$$

Also,

$$k = \text{rank}(Y + Y_\epsilon) \leq \text{rank}(Y) + \text{rank}(Y_\epsilon) = r_Y + \text{rank}(Y_\epsilon) \tag{20}$$

From inequalities (19) and (20), we get

$$k - r_Y \leq \text{rank}(Y_\epsilon) \leq k - r_X \Rightarrow r_Y \geq r_X \tag{21}$$

Suppose that $(\mathcal{A}_X)$ holds, then

$$\mathcal{C}(\,X_\epsilon^T\,) \subseteq \mathcal{N}(\,Y^T\,) \Rightarrow \text{rank}(X_\epsilon) \leq \dim\big(\,\mathcal{N}(\,Y^T\,)\big) = k - r_Y. \tag{22}$$

Also,

$$k = \text{rank}(X + X_\epsilon) \leq \text{rank}(X) + \text{rank}(X_\epsilon) = r_X + \text{rank}(X_\epsilon) \tag{23}$$

From inequalities (22) and (23), we get

$$k - r_X \leq \text{rank}(X_\epsilon) \leq k - r_Y \Rightarrow r_X \geq r_Y \tag{24}$$

From inequalities (21) and (24), we get $r_X = r_Y$, which by Lemma 10 implies $\dim\big(\mathcal{N}(X) \cap \mathcal{C}(Y)\big) = 0$. Let $r = r_X = r_Y$ be the rank of the matrices X and Y, then the rank of the product matrix XY is less then or equal to r. Then it directly follows from the SVD decomposition of the matrix XY, that $U^T X\big(YV\big)_{:,r+1:n} = \Sigma_{:,r+1:n} = 0$, or equivalently $XYV_{:,r+1:n} = 0$. On the other hand, since $\mathcal{C}\big(YV_{:,r+1:n}\big) \subset \mathcal{C}\big(Y\big)$ and $\mathcal{N}\big(X\big) \cap \mathcal{C}\big(Y\big) = \emptyset$, we conclude that $\big(YV\big)_{:,r+1:n} = 0$. Similarly, one can show that $\big(U^T X\big)_{r+1:n,:} = 0$. □

We now state and prove the following lemma which will be used latter in the proof of Proposition 13.

Lemma 12. *Consider a matrix $V \in \mathbb{R}^{m\times n}$ with rank$(V) = r$. Then there exists a basis $\mathcal{B} = \{i_1, \dots, i_r\}$ and a matrix $A \in \mathbb{R}^{r\times(m-r)}$ with bounded norm $\|A\| \le m2^m$, such that*

$$V_{\mathcal{B}^c} = AV_{\mathcal{B}}$$

where $V_{\mathcal{B}}$ is a matrix with rows $\{V_{i,:}\}_{i\in\mathcal{B}}$ and $V_{\mathcal{B}^c}$ is a matrix with rows $\{V_{i,:}\}_{i\in\mathcal{B}^c}$

Proof. To ease the notation, we denote the ith row of V by v_i. We use induction on m to show that there exists a basis $\mathcal{B} = \{i_1, \dots, i_r\}$ such that $\forall\ j \in \mathcal{B}^c$,

$$v_j = \sum_{i\in\mathcal{B}} a_{j,i} v_i \qquad \text{with } |a_{j,i}| \le 2^m \quad \forall\ i \in \mathcal{B}.$$

Base step $m = r + 1$:

Without loss generality assume $\mathcal{B} = \{1, \dots, r\}$. Since the case of $v_{r+1} = 0$ trivially holds, we consider $v_{r+1} \neq 0$. By the property of basis, there exists $a_{r+1} \neq 0$ such that $v_{r+1} = \sum_{i=1}^{r} a_{r+1,i} v_i$.

Let $i^* = \underset{i\in\mathcal{B}}{\operatorname{argmin}} |a_{r+1,i}|$. We only need to consider the case when $|a_{r+1,i^*}| \ge 2^{r+1}$. In that case,

$$\begin{aligned} v_{i^*} &= \underbrace{\frac{1}{a_{r+1,i^*}}}_{\bar{a}_{r+1,r+1}} v_{r+1} - \sum_{i=1;\, i\neq i^*}^{r} \underbrace{\frac{a_{r+1,i}}{a_{r+1,i^*}}}_{\bar{a}_{r+1,i}} v_i \\ &= \sum_{i\in\mathcal{B}^*} \bar{a}_{r+1,i} v_i \qquad \text{where } \ \mathcal{B}^* = \mathcal{B} - \{i^*\} + \{r+1\} \end{aligned}$$

Since $\bar{a}_{r+1,i} \le 1 < 2^{r+1}$, the induction step holds.

Assume it is true for $m > r$, we show the result for $m + 1$. Without loss of generality assume $\mathcal{B} = \{1, \dots, r\}$. By induction hypothesis,

$$v_j = \sum_{i=1}^{r} a_{j,i}\, v_i \qquad \text{with } \ |a_{j,i}| \le 2^m \quad \forall\, j = \{r+1, \dots, m\};\ \ i \in \mathcal{B}$$

Since the case of $v_{m+1} = 0$ trivially holds, we consider $v_{m+1} \neq 0$. By the property of basis, there exists $a_{m+1} \neq 0$ such that $v_{m+1} = \sum_{i=1}^{r} a_{m+1,i}\, v_i$. Let $i^* = \underset{i\in\mathcal{B}}{\operatorname{argmin}} |a_{m+1,i}|$. We only need to consider the case when $|a_{m+1,i^*}| \ge 2^{m+1}$. In that case,

$$\begin{aligned} v_{i^*} &= \underbrace{\frac{1}{a_{m+1,i^*}}}_{\bar{a}_{m+1,m+1}} v_{m+1} - \sum_{i=1;\, i\neq i^*}^{r} \underbrace{\frac{a_{m+1,i}}{a_{m+1,i^*}}}_{\bar{a}_{m+1,i}} v_i \\ &= \sum_{i\in\mathcal{B}^*} \bar{a}_{m+1,i} v_i \end{aligned}$$

where $\mathcal{B}^* = \mathcal{B} - \{i^*\} + \{m+1\}$ and $|\bar{a}_{m+1,i}| \leq 1 \; \forall \; i \in \mathcal{B}^*$. For all $j \in \{r+1, \ldots, m\}$

$$\begin{aligned} v_j &= \sum_{i=1;\, i\neq i^*}^{r} a_{j,i}\, v_i + a_{j,i^*} v_{i^*} \\ &= \sum_{i=1;\, i\neq i^*}^{r} a_{j,i}\, v_i + \frac{a_{j,i^*}}{a_{m+1,i^*}} v_{m+1} - \sum_{i=1;\, i\neq i^*}^{r} \frac{a_{m+1,i}\, a_{j,i^*}}{a_{m+1,i^*}} v_i \\ &= \sum_{i=1;\, i\neq i^*}^{r} \underbrace{\big(a_{j,i} - \frac{a_{j,i^*}\, a_{m+1,i}}{a_{m+1,i^*}} \big)}_{\bar{a}_{j,i}} v_i + \underbrace{\frac{a_{j,i^*}}{a_{m+1,i^*}}}_{\bar{a}_{j,m+1}} v_{m+1} \\ &= \sum_{i\in\mathcal{B}^*} \bar{a}_{j,i} v_i \end{aligned}$$

It remains to show that $|\bar{a}_{j,i}| \leq 2^{m+1}$ for all $i \in \mathcal{B}$, $j \in \mathcal{B}^c$. Lets first consider $i \in \mathcal{B}^* - \{m+1\}$ and $j \in \mathcal{B}^c$.

$$\begin{aligned} |\bar{a}_{j,i}| &\leq |a_{j,i}| + \left| a_{j,i^*} \frac{a_{m+1,i}}{a_{m+1,i^*}} \right| \qquad \text{by induction hypothesis} \\ &\leq 2^m + 2^m \left| \frac{a_{m+1,i}}{a_{m+1,i^*}} \right| \\ &\leq 2^{m+1} \end{aligned}$$

For $i = m+1$, $|\bar{a}_{j,m+1}| = \left| \frac{a_{j,i^*}}{a_{m+1,i^*}} \right| \leq 2^m$. This completes the induction proof.

By setting the rows of A to $\big(\bar{a}_j\big)_{j\in\mathcal{B}}$, it directly follows that $V_{\mathcal{B}^c} = AV_{\mathcal{B}}$ with $\|A\| \leq \sqrt{|\mathcal{B}||\mathcal{B}^c| 2^m} \leq \sqrt{m^2 2^m} = m 2^m$

□

Note that the bound we got on A is not strict, but we only need to use the fact that it is bounded. We are now ready to state and prove Proposition 13 which is the main block in Theorem 3.

Proposition 13. *Consider $X \in \mathbb{R}^{m\times k}$, $Y \in \mathbb{R}^{k\times n}$, $k < \min(m,n)$ then*

$$\mathcal{M}(X,Y) \text{ is locally open at } (X,Y) \Leftrightarrow \begin{cases} (\mathcal{A}_X) : \exists\, X_\epsilon \text{ such that } X_\epsilon Y = 0 \text{ and } X + X_\epsilon \text{ is full column rank.} \\ \qquad\qquad \text{and} \\ (\mathcal{A}_Y) : \exists\, Y_\epsilon \text{ such that } X Y_\epsilon = 0 \text{ and } Y + Y_\epsilon \text{ is full row rank.} \end{cases}$$

Proof. Let $U\Sigma V^T$ be the singular value decomposition of Z with $U \in \mathbb{R}^{m\times m}$, $V \in \mathbb{R}^{n\times n}$, and $\Sigma \in \mathbb{R}^{m\times n}$, and define $\bar{X} = U^T X$, $\bar{Y} = YV$. Then using Lemma 9, $Z = \mathcal{M}(X,Y)$ is locally open at (X,Y) if and only if $\Sigma = \mathcal{M}(\bar{X},\bar{Y})$ is locally open at $(\bar{X},\bar{Y})$. Notice that

$$XY_\epsilon = 0 \Leftrightarrow \bar{X}\bar{Y}_\epsilon = U^T X Y_\epsilon V = 0, \quad \text{and} \quad Y + Y_\epsilon \text{ is full row rank} \Leftrightarrow \bar{Y} + \bar{Y}_\epsilon \text{ is full row rank}$$

A similar equivalence relation holds for the condition on X_ϵ and $\bar{X}_\epsilon$.

We now prove the direction " $\Rightarrow$ ". While the result obviously holds when $r_{\bar{X}} = r_{\bar{Y}} = k$ by choosing $\bar{X}_\epsilon = \bar{Y}_\epsilon = 0$, we show it is true when at least one of the matrices is rank deficient.

$\bar{X}$ is full column rank and $\bar{Y}$ is rank deficient:
Suppose $\mathcal{M}(\bar{X},\bar{Y})$ is locally open at $(\bar{X},\bar{Y})$, it follows from the definition of openness that $\mathcal{M}(\bar{X},\bar{Y}_1)$ is locally open at $(\bar{X},\bar{Y}_1)$ where $\bar{Y}_1 = \bar{Y}_{:,1:k}$ is the first k columns of $\bar{Y}$.
Since the range of the mapping $\mathcal{M}$ at $(\bar{X},\bar{Y}_1)$ is $\mathcal{R}^k_{\mathcal{M}} = \mathbb{R}^{m\times k}$, using Proposition 2, we get that either

$$\begin{cases} \exists\, \bar{X}_\epsilon \text{ such that } \bar{X}_\epsilon \bar{Y}_1 = 0 \text{ and } \bar{X} + \bar{X}_\epsilon \text{ is full row rank.} \\ \qquad\qquad \text{or} \\ \exists\, \bar{Y}^1_\epsilon \text{ such that } \bar{X}\bar{Y}^1_\epsilon = 0 \text{ and } \bar{Y}_1 + \bar{Y}^1_\epsilon \text{ is full column rank.} \end{cases}$$

But since $\bar{X} \in \mathbb{R}^{m\times k}$ and $m > k$, it is impossible to have $\bar{X}_\epsilon$ such that $\bar{X} + \bar{X}_\epsilon$ full row rank. Moreover, since $\bar{X}$ is full rank, for any $\bar{Y}_\epsilon^1$ satisfying $\bar{X}\bar{Y}_\epsilon^1 = 0$ we have $\bar{Y}_\epsilon^1 = 0 \Rightarrow \bar{Y}_1 + \bar{Y}_\epsilon^1$ is not full column rank. This contradicts the conditions of local optimality of $\mathcal{M}(\bar{X}, \bar{Y}_1)$.

$\bar{X}$ is rank deficient and $\bar{Y}_1$ is full row rank:
The proof of this case if similar to the one above, just substitute $\bar{X} = \bar{Y}^T$ and $\bar{Y} = \bar{X}^T$.

$\bar{X}$ and $\bar{Y}$ are both rank deficient:
Suppose $\mathcal{M}(\bar{X}, \bar{Y})$ is locally open at $(\bar{X}, \bar{Y})$, then it directly follows that $\mathcal{M}(\bar{X}, \bar{Y}_1)$ is locally open at $(\bar{X}, \bar{Y}_1)$. By Proposition 2, and since there does not exist $\bar{X}_\epsilon$ such that $\bar{X} + \bar{X}_\epsilon$ is full row rank, there exists $\bar{Y}_\epsilon^1$ such that $\bar{X}\bar{Y}_\epsilon^1 = 0$ and $\bar{Y}_1 + \bar{Y}_\epsilon^1$ is full rank.

By letting $\bar{Y}_\epsilon = \left[\begin{array}{c|c} \bar{Y}_\epsilon^1 & 0 \end{array}\right]$, we satisfy the desired condition.

Similarly, since $\mathcal{M}(\bar{Y}^T, \bar{X}^T)$ is locally open at $(\bar{Y}^T, \bar{X}^T)$, the same proof can be used to show that there exists $\bar{X}_\epsilon$ such that $\bar{X} + \bar{X}_\epsilon$ is full column rank, and $\bar{X}_\epsilon\bar{Y} = 0$.

Note that the matrix multiplication is not locally open at $(\bar{X}, \bar{Y})$ when only one of the matrices is rank deficient.

We now prove the converse direction " $\Leftarrow$ ". Suppose $(\mathcal{A}_\mathcal{X})$ and $(\mathcal{A}_\mathcal{Y})$ hold,

$$\begin{cases} \exists\, X_\epsilon \text{ such that } X_\epsilon Y = 0 \text{ and } X + X_\epsilon \text{ is full column rank.} \\ \qquad\qquad\qquad\qquad \text{and} \\ \exists\, Y_\epsilon \text{ such that } XY_\epsilon = 0 \text{ and } Y + Y_\epsilon \text{ is full row rank.} \end{cases}$$

then by Lemma 11, we get that $r_{\bar{X}} = r_{\bar{Y}} \triangleq r$, and the last $n - r$ columns of $\bar{Y}$ are all zeros. We need to show that for any given $\epsilon > 0$, there exists $\delta > 0$, such that

$$\mathbb{B}_\delta\big(\bar{X}\bar{Y}\big) \cap \mathcal{R}^u_\mathcal{M} \subseteq \mathcal{M}\Big(\mathbb{B}_\epsilon\big(\bar{X}\big), \mathbb{B}_\epsilon\big(\bar{Y}\big)\Big)$$

Let $\Sigma = \bar{X}\bar{Y} = \left[\begin{array}{cc} \Sigma_{:,1:r} & 0 \end{array}\right]$ be a rank r matrix. We show that for any $\tilde{\Sigma} \in \mathbb{B}_\delta\big(\bar{X}\bar{Y}\big) \cap \mathcal{R}^u_\mathcal{M}$, which is a permutation of Σ, $\tilde{\Sigma} \in \mathbb{B}_\delta\big(\bar{X}\bar{Y}\big) \cap \mathcal{R}^u_\mathcal{M}$. Without loss of generality, and by permuting the columns of $\tilde{\Sigma}$ if necessary, $\tilde{\Sigma}$ can be expressed as

$$\tilde{\Sigma} = \left[\begin{array}{c|c|c} \underbrace{\Sigma_{:,1:r} + R^1_\delta}_{m\times r} & \underbrace{R^2_\delta}_{m\times(k-r)} & \underbrace{(\Sigma_{:,1:r} + R^1_\delta)A_1 + R^2_\delta A_2}_{m\times(n-k)} \end{array}\right]$$

or its transpose.

Where $A_1 \in \mathbb{R}^{r\times(n-k)}$, $A_2 \in \mathbb{R}^{(k-r)\times(n-k)}$ assure that the rank condition is satisfied, i.e. rank($\tilde{\Sigma}$)$\leq k$. Moreover, $\|\tilde{\Sigma} - \Sigma\| \leq \delta$ implies that the perturbation matrix

$$R_\delta = \left[\begin{array}{c|c|c} R^1_\delta & R^2_\delta & (\Sigma_{:,1:r} + R^1_\delta)A_1 + R^2_\delta A_2 \end{array}\right]$$

has norm less than or equal δ, i.e. $\|R_\delta\| \leq \delta$.

Given that there exist $\bar{Y}_\epsilon$ such that $\bar{X}\bar{Y}_\epsilon = 0$ and $\bar{Y} + \bar{Y}_\epsilon$ is full row rank, and knowing that the last $n - r$ columns of $\bar{Y}$ are all zeros, there exist a perturbation matrix

$$\bar{Y}^1_\epsilon \triangleq \delta \left[\begin{array}{cc} \overbrace{0}^{k\times r} & \overbrace{y^1_\epsilon \ \ldots, y^{k-r}_\epsilon}^{k\times(k-r)} \\ & \end{array}\right]$$

where y_ϵ^i are unit basis of $\bar{Y}_\epsilon$ such that $\bar{Y}_1 + \bar{Y}_\epsilon^1$ is a full rank $k \times k$ matrix, and $\bar{X}\bar{Y}_\epsilon^1 = 0$. Note that there exist $k - r$ unit basis vectors for $\bar{Y}_\epsilon$ since the rank$(\bar{Y}_\epsilon) \geq$ rank$(\bar{Y} + \bar{Y}_\epsilon) -$ rank$(\bar{Y}) = k - r$. Let

$$\bar{X}_0 = \left[\begin{array}{c|c} R_\delta^1 & R_\delta^2 \end{array}\right] (\bar{Y}_1 + \bar{Y}_\epsilon^1)^{-1} \qquad \bar{Y}_0 = \left[\begin{array}{c|c} \bar{Y}_\epsilon^1 & \left(\bar{Y}_1\right)_{:,1:r} A_1 + \bar{Y}_\epsilon^1 A_2 \end{array}\right]$$

Then

$$(\bar{X} + \bar{X}_0)(\bar{Y} + \bar{Y}_0) = \bar{X}\bar{Y} + \left[\begin{array}{c|c|c} R_\delta^1 & R_\delta^2 & (\bar{X} + \bar{X}_0)\big((\bar{Y}_1)_{:,1:r} A_1 + \bar{Y}_\epsilon^1 A_2\big) \end{array}\right] = \bar{X}\bar{Y} + R_\delta = \tilde{\Sigma} \tag{25}$$

To complete the proof, it remains to show that $\|\bar{X}_0\| \leq \epsilon$ and $\|\bar{Y}_0\| \leq \epsilon$ which implies $\tilde{\Sigma} \in \mathcal{M}\Big(\mathbb{B}_\epsilon\big(\bar{X}\big), \mathbb{B}_\epsilon\big(\bar{Y}\big)\Big)$.

Perturbation Bound

Let $\tilde{r} \geq r$ be the rank of $\tilde{\Sigma}$, then by lemma 12 and by possibly permuting the columns of $\tilde{\Sigma}$, it can be expressed as

$$\tilde{\Sigma} = \left[\begin{array}{cc} \tilde{\Sigma}_1 & \tilde{\Sigma}_1 \bar{A} \end{array}\right]$$

where $\tilde{\Sigma}_1 \in \mathbb{R}^{m \times \tilde{r}}$ is full column rank, and $\bar{A}$ has a bounded norm ($\|\bar{A}\| \leq n2^n$). Notice that given $\bar{X}_0$ and $\bar{Y}_0$ that satisfy (25), permuting the columns of $\tilde{\Sigma}$ correspond to permuting the columns of $(\bar{Y} + \bar{Y}_0)$. If we can show that the first r columns are not among the permuted ones, then using the fact that $\bar{Y}$ has only its first r columns non-zero, it follows that this permutation of the columns of $\tilde{\Sigma}$ corresponds to the same permutation of the columns of $\bar{Y} + \bar{Y}_0$ which in turn corresponds to the same permutation of the columns of $\bar{Y}_0$. Moreover, if the first r columns are not among the permuted ones, then without loss of generality we can express

$$\tilde{\Sigma} = \left[\begin{array}{c|c|c} \underbrace{\Sigma_{:,1:r} + R_\delta^1}_{m \times r} & \underbrace{R_\delta^2}_{m \times (k-r)} & \underbrace{(\Sigma_{:,1:k} + R_\delta^1)\bar{A}_1 + R_\delta^2 \bar{A}_2}_{m \times (n-k)} \end{array}\right]$$

where $\left[\begin{array}{c} \bar{A}_1 \\ \bar{A}_2 \end{array}\right] = \bar{A}$ which has a bounded norm.

We now show that the first r columns of $\tilde{\Sigma}$ before permutation $\Sigma_{:,1:r} + \left(R_\delta^1\right)_{:,1:r} \subseteq \tilde{\Sigma}_1$. Assume the contrary, then there exists at least one column call it $\Sigma_{:,j} + \left(R_\delta^1\right)_{j,:}$ that is not in $\tilde{\Sigma}_1$, which implies $\Sigma_{:,j} + \left(R_\delta^1\right)_{j,:} \subseteq \tilde{\Sigma}_1 \bar{A}$. Without loss of generality let $\Sigma_{:,j} + \left(R_\delta^1\right)_{j,:} = \tilde{\Sigma}_1 \bar{A}_{:,1}$. It follows that

$$\Sigma_{j,j} + \left(R_\delta^1\right)_{j,j} = (\tilde{\Sigma}_1)_{j,:} \bar{A}_{:,1}$$

But since $\Sigma_{j,j} + \left(R_\delta^1\right)_{j,j}$ is a non-zero perturbed singular value, and since elements of $(\tilde{\Sigma}_1)_{j,:}$ are all of order δ, then by choosing δ sufficiently small, we get $|\bar{A}| > n2^n$ which contradicts the statement of lemma 12.

We now get a bound on $\|Y_0\|$. Since the norm of $\bar{A}$ is bounded, then the norm of $\bar{A}_2$ is bounded by some constant, denote it by K.

$$\begin{aligned} \delta &\geq \|(\Sigma_{:,1:r} + R_\delta^1)\bar{A}_1 + R_\delta^2 \bar{A}_2\| \\ &\geq \|(\Sigma_{:,1:r} + R_\delta^1)\bar{A}_1\| - \|R_\delta^2 A_2\| \\ &\geq \|(\Sigma_{:,1:r} + R_\delta^1)\bar{A}_1\| - K\delta \\ &\geq \sigma_{\min}\|\bar{A}_1\| - K\delta \end{aligned}$$

where $\sigma_{\min}$ is the minimum singular value of the full column rank matrix $\Sigma_{:,1:r} + R_\delta^1$ which is bounded away from zero. Then it follows that

$$\|\bar{A}_1\| \leq \frac{(1+K)\delta}{\sigma_{\min}}$$

Then,

$$\|Y_0\|^2 \leq \|\bar{Y}_\epsilon^1\|^2 + \|\bar{Y}_1\|^2 \, \|\bar{A}_1\|^2 + \|\bar{Y}_\epsilon^1\|^2 \, \|\bar{A}_2\|^2 \qquad \text{(by triangular inequality and Cauchy Shwarz)}$$
$$\leq \delta^2 \Big[1 + \left(\frac{1+K}{\sigma_{min}} \right)^2 + K^2 \Big]$$

For a given $\epsilon > 0$, choose $\delta \leq \dfrac{\epsilon}{1 + \max\left(\|(\bar{Y}_1 + \bar{Y}_\epsilon^1)^{-1}\|, \sqrt{1 + K^2 + \left(\frac{1+K}{\sigma_{min}} \right)^2} \right.}$, then we get $\|Y_0\| < \epsilon$ and

$$||\bar{X}_0|| \leq ||R_\delta|| \, ||(\bar{Y}_1 + \bar{Y}_\epsilon^1)^{-1}|| \leq \delta ||(Y + Y_\epsilon)^{-1}|| \leq \frac{\epsilon(||(Y+Y_\epsilon)^{-1}||}{||(Y+Y_\epsilon)^{-1}|| + 1} \leq \epsilon$$

Note that δ is of order ϵ. This completes the proof.

□

We now use Proposition 13, Lemma 10 and Lemma 11 to complete the proof of Theorem 3

Proof. Consider $X \in \mathbb{R}^{m \times k}$, and $Y \in \mathbb{R}^{k \times n}$ with $u = k \leq (m.n)$. If $r_X = r_Y$, it follows from Lemma 10 that $(\mathcal{A}_X)$ and $(\mathcal{A}_Y)$ are equivalent. If the two conditions hold, then by Proposition 13, the mapping $\mathcal{M}$ is locally open at (X, Y).

But if $r_X \neq r_Y$, then, by Lemma 11, $(\mathcal{A}_X)$ and $(\mathcal{A}_Y)$ cannot both hold which implies, by Theorem 3, that $\mathcal{M}$ is not locally open at (X, Y).This completes the proof. □