



HAL
open science

The loss landscape of deep linear neural networks: a second-order analysis

El Mehdi Achour, François Malgouyres, Sébastien Gerchinovitz

► **To cite this version:**

El Mehdi Achour, François Malgouyres, Sébastien Gerchinovitz. The loss landscape of deep linear neural networks: a second-order analysis. 2022. hal-03299887v2

HAL Id: hal-03299887

<https://hal.science/hal-03299887v2>

Preprint submitted on 24 Feb 2022 (v2), last revised 19 Sep 2024 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The loss landscape of deep linear neural networks: a second-order analysis

El Mehdi Achour¹, François Malgouyres¹, and Sébastien Gerchinovitz^{2,1}

¹Institut de Mathématiques de Toulouse ; UMR5219 , Université de Toulouse ; CNRS , UPS IMT
F-31062 Toulouse Cedex 9, France

²IRT Saint Exupéry, 3 rue Tarfaya, 31400 Toulouse, France

February 24, 2022

Abstract

We study the optimization landscape of deep linear neural networks with the square loss. It is known that, under weak assumptions, there are no spurious local minima and no local maxima. However, the existence and diversity of non-strict saddle points, which can play a role in first-order algorithms' dynamics, have only been lightly studied. We go a step further with a full analysis of the optimization landscape at order 2. We characterize, among all critical points, which are global minimizers, strict saddle points, and non-strict saddle points. We enumerate all the associated critical values. The characterization is simple, involves conditions on the ranks of partial matrix products, and sheds some light on global convergence or implicit regularization that have been proved or observed when optimizing linear neural networks. In passing, we provide an explicit parameterization of the set of all global minimizers and exhibit large sets of strict and non-strict saddle points.

Keywords— Deep learning, landscape analysis, non-convex optimization, second-order geometry, strict saddle points, non-strict saddle points, global minimizers, implicit regularization

1 Introduction

Deep learning has been widely used recently due to its good empirical performances in image recognition, natural language processing, speech recognition, among other fields. However, there is still a gap between theory and practice. One of the aspects that are partially missing in the picture is why gradient-based algorithms can achieve low training error despite a highly non-convex function. Another partially open question is why they generalize well to unseen data despite many more parameters than the number of points in the training set, and how implicit regularization can help with that.

Various research directions have been explored to answer these questions, including double-descent and benign overfitting (e.g., [7, 23, 6]), gradient flow dynamics in Wasserstein space (e.g., [11, 33]), landscape analysis of the empirical risk (e.g., [12, 27, 22, 21]). In this paper, we follow the latter direction, by better characterizing the local structures around critical points of the empirical risk, in the case of deep linear neural networks with the square loss.

Before summarizing the related literature and our main contributions, we first recall definitions that will be key throughout the paper.

1.1 Reminder : minimizers, critical points of order 1 or 2, strict and non-strict saddle points

Let us recall the definitions of local structures of the landscape of the empirical risk, which are important from the statistical and optimization points of view.




Figure 1: Example of a landscape with a plateau (non-strict saddle point) and a global minimizer.




Figure 2: Example of a landscape with a strict saddle point at $(0,0)$.

For $\mathbf{w} \in \mathbb{R}^n$, denote by $\mathbf{w} \mapsto L(\mathbf{w})$ the function we want to minimize. Assume that $\mathbf{w} \mapsto L(\mathbf{w})$ is C^2 , and denote by ∇L and $\nabla^2 L$ its gradient and its Hessian.¹ We also write $A \succeq 0$ to say that a matrix $A \in \mathbb{R}^{n \times n}$ is positive semi-definite. Recall the following four definitions, which are nested:

- \mathbf{w}^* is a **global minimizer** if and only if $\forall \mathbf{w} \in \mathbb{R}^n, L(\mathbf{w}^*) \leq L(\mathbf{w})$.
- \mathbf{w}^* is a **local minimizer** if and only if there exists a neighbourhood $\mathcal{O} \subset \mathbb{R}^n$ of \mathbf{w}^* such that $\forall \mathbf{w} \in \mathcal{O}, L(\mathbf{w}^*) \leq L(\mathbf{w})$.
- \mathbf{w}^* is a **second-order critical point** if and only if $\nabla L(\mathbf{w}^*) = 0$ and $\nabla^2 L(\mathbf{w}^*) \succeq 0$. If, on the contrary, the Hessian has a negative eigenvalue, we say that the point has a negative curvature.
- \mathbf{w}^* is a **first-order critical point** if and only if $\nabla L(\mathbf{w}^*) = 0$.

We can also distinguish a specific type of first-order critical point: saddle points. As discussed below, they can be second-order critical points or not.

- \mathbf{w}^* is a **saddle point** if and only if it is a first-order critical point which is neither a local minimizer, nor a local maximizer.
 - A saddle point \mathbf{w}^* is **strict** if and only if it is not a second-order critical point (i.e., the Hessian $\nabla^2 L(\mathbf{w}^*)$ has a negative eigenvalue). Figure 2 gives an example.
 - A saddle point \mathbf{w}^* is **non-strict** if and only if it is a second-order critical point. In that case, the Hessian $\nabla^2 L(\mathbf{w}^*)$ is positive semi-definite and has at least one eigenvalue equal to zero. Typically, in the direction of the corresponding eigenvectors a higher-order term makes it a saddle point (e.g., $L(\mathbf{w}) = \sum_{i=1}^n w_i^3$ at $\mathbf{w}^* = 0$). Figure 1 gives an example.

1.2 On the importance of a landscape analysis at order 2

When the function we are trying to minimize is smooth, convex, and has a global minimizer, the gradient descent algorithm with a well-chosen learning rate converges to a first-order critical point, and this critical point is a global minimizer [36]. However, in general, finding a global optimum of a non-convex function is an NP-complete problem [35]; this is in particular the case for a simple 3-node neural network [8]. Despite that, when optimizing neural networks, the current practice is still to use gradient-based algorithms.

It has been known for decades that, even in the non-convex setting, under mild conditions gradient-based algorithms converge to a first-order critical point, in the sense that the iterates produced by the algorithm can reach an arbitrary small gradient after a finite (polynomial) number of iterations [36]. Adding smoothness conditions, recent works have shown that classical first-order algorithms escape strict saddle points in the long run [30, 29], and that

¹When the input parameter is not a vector, but, e.g., a sequence of matrices, the same definitions hold, where the gradient and the Hessian are computed with respect to the vectorized version of the input parameters.




Figure 3: The loss function during the iterative process, when initialized around a strict saddle point (in red) or a non-strict saddle point (in blue).




Figure 4: Histogram of escape epochs, when initialized around a strict (in red) or a non-strict saddle point (in blue). For clarity, the y -axis is endowed with two scales. The right axis corresponds to the blue curve and the left to the red one.

some of them can be stopped in polynomial time at a nearly second-order critical point² [24, 26, 13, 25]. However, nothing prevents these algorithms to converge to non-strict saddle points or to spend many epochs in their vicinity, which translates into a long plateau during training.

To see that this behavior actually occurs in practice, consider the simple experiment whose results are shown in Figures 3 and 4 (more details in Appendix G). For each run of this experiment, the parameters of a linear neural network of depth 5 are optimized to fit random input/output pairs. Discrepancy is measured with the square loss and we use the ADAM optimizer. Depending on the run, the algorithm is initialized in the vicinity either of a strict saddle point (in red) or of a non-strict saddle point (in blue). The distance between the random initial iterate and the saddle point is purposely not negligible: it is fixed to around 10% of the norm of the saddle point. Figure 3 shows the typical loss evolution for both cases. We can see that ADAM rapidly escapes from the strict saddle point but needs many epochs to escape the plateau in the vicinity of the non-strict saddle point. Figure 4 shows that this observation generalizes to most runs. We compare the empirical distributions of a random time (called *escape epoch*) defined as the epoch at which the loss has significantly decreased from its initial value. When initialized in the vicinity of non-strict saddle points, the algorithm suffers from an often large escape epoch and might be stopped there, without the possibility to distinguish this non-strict saddle point from a global minimum. Improving the analysis beyond local minimizers and characterizing strict and non-strict saddle points are therefore key to understand gradient descent dynamics and implicit regularization.

Beyond neural networks, the study of the loss landscape of specific non-convex optimization problems has revealed that they are tractable: phase retrieval [47], dictionary learning [46], tensor decomposition [17, 18, 16] and others [54, 34]. In fact, a landscape property which is shared by most of these problems is that every critical point is either a global minimizer or has a negative curvature. In other words, every second-order critical point is a global minimizer. For such problems, there are first-order algorithms which provably converge to global minimizers.

The general understanding of the landscape is not as good for neural networks. A regime which has been widely studied is the overparameterized regime (see [49] and [48] for a review), where it has been proved under some assumptions that for a wide non-linear fully connected neural network almost all local minima are global minima [38], or that there are no spurious valleys [37].

Many recent works have focused on linear neural networks, despite the fact that they are rarely used to solve real-world applications. They indeed compute a linear map between the input and output spaces. The motivation for these studies is that the empirical risk of linear networks is highly non-convex and shares similar properties to that of practical non-linear neural networks. A complete review of the scientific publications on the landscape for linear networks is given in Section 1.3 but we would like to emphasize at this point that many results are strongly related to

²More precisely, it is typically shown that, with high probability, such algorithms can be stopped in polynomial time and output a point with arbitrarily small gradient and nearly-positive semi-definite Hessian.

the properties of the landscape at order 2.

As shown by [43], linear networks exhibit nonlinear learning phenomena similar to those seen during the optimization of nonlinear networks, including long plateaus followed by rapid transitions to lower error solutions, as in Figure 3. Several works followed (e.g., [19, 20]) with formal proofs in several special cases. They express that for particular initializations the iterates trajectory spends some time on plateaus whose location defines an implicit regularization, if the algorithm is stopped early.

Other authors provide well-chosen initializations and/or architectures for which the convergence to a global minimizer can be established [1, 5, 15, 14].

The difference between implicit regularization and global convergence depends on whether the iterates trajectory considered in these articles avoids saddle points or not. To better understand convergence properties of first-order algorithms, to give some hints on their finite-time dynamics and better understand the implicit regularization, a detailed analysis of the empirical risk landscape at order 2 is needed.

Furthermore, although it has been proved under mild conditions that every local minimizer is a global minimizer and that there exists no local maximizer [27, 52], very little is known concerning saddle points and the landscape at order 2. For 1-hidden layer linear networks, a proved fact is that every saddle point is strict, therefore leading to the property that every second-order critical point is a global minimizer [56, 27]. Unfortunately, this is not the case for linear networks with two hidden layers or more. For such neural networks, it has only been noted by [27] that there exist non-strict saddle points (e.g., when all the weight matrices are equal to 0).

In this paper, we make the missing additional step and completely characterize the landscape of the empirical risk at order 2, for the square loss and deep linear networks. In particular, we derive a simple necessary and sufficient condition for a first-order critical point of the empirical risk to be either a global minimizer, a strict saddle point or a non-strict saddle point. In Sections 1.4 and 3.4 we detail our contributions and how they generalize earlier results and shed some light on global convergence or implicit regularization phenomena. We start below by describing related works in more details.

1.3 Related Works

The study of linear neural networks can be divided into two categories. The first line of research is about the geometric landscape of the empirical risk for linear neural networks, while the second line is about the trajectory of gradient descent dynamics in linear networks. Our work lies within the first category.

Geometric landscape for linear networks: This first started with [4]. They proved that for a 1-hidden layer linear network, under some conditions on the data matrices, and for the square loss, every local minimizer is a global minimizer. [27] later generalized and extended this result to deep linear neural networks under mild conditions and again proved that every local minimizer is a global minimizer (this part has been proved later by [31] with weaker assumptions on the data and simpler proofs). This author also proved that every other critical point is a saddle point, that for a 1-hidden layer linear network all saddle points are strict, while for deeper networks, there exist non-strict saddle points ([27] exhibits a space of non-strict saddle points where all but one weight matrix are equal to zero). [52] gave a condition for a critical point to be either a global minimizer or a saddle point. [55] removed all assumptions on the data and gave analytical forms for the critical points of the empirical risk. In the characterization, the weight matrices are defined recursively and can be found by solving equations; in particular they gave a characterization of global minimizers. [39] showed using assumptions only on the width of the layers that every local minimizer is a global minimizer. They prove that this assumption on the architecture is sharp in the sense that without it, and if we do not make assumptions on the data matrices as in previous works, then there exists a poor local minimizer. [56] used assumptions only on the input data matrix, to prove that for a 1-hidden layer linear network, every local minimizer is a global minimizer and every other critical point has a negative curvature. [28] proved for different general convex losses that, under assumptions on the architecture, all local minima are global. Finally, [50] and [32] used results from algebraic geometry to give other properties about critical points of linear networks.

Most of the previous works focus on local minimizers. None of these works provide simple necessary and sufficient conditions for a saddle point to be strict or not.³ In particular, in the case of more than two hidden

³By “simple”, we mean an easier-to-exploit condition than just looking at the smallest eigenvalue of the Hessian.

layers, only very specific examples of non-strict saddle points were described. Furthermore, global minimizers were characterized but not explicitly parameterized. See Section 3.4 for more details.

Gradient dynamics and implicit regularization for linear networks: In this line of research, authors study the dynamics of first-order algorithms for linear networks, which they sometimes combine with results about the loss landscape. [1] proved that gradient descent converges to a global minimum at a linear rate, under assumptions on the width of the layers, the initial iterate, and the loss at initialization. Other works also proved similar results with different assumptions [15, 5, 51]. However, as noted by [45], these works consider strong assumptions on the loss at initialization. Indeed, [45] gave a negative result on a deep linear network of width 1, by proving that for standard initializations, gradient descent can take exponential time to converge to the global minimizer. The author also provided empirical examples of the same phenomenon happening for larger widths. On the other hand, [14] proved that if the layers are wide enough, convergence to a global minimum can be achieved in polynomial time using a classical data-independent random Gaussian initialization (known as Xavier initialization). The required minimum width of the network depends on the norm of a global minimizer of the linear regression problem. As we will see in Section 3.4 this global convergence result can be re-interpreted in terms of the loss landscape at order 2.

On a similar line of research, [10] proved using assumptions on the architecture of the network and the data matrices that gradient flow almost surely converges to a global minimizer for a 1-hidden layer linear network. Later, [3] proved the same result under weaker assumptions on the data matrices. They also proved that, in deep linear networks, the gradient flow almost surely converges to global minimizers of the rank-constrained linear regression problem.

This is related to another consequence of the landscape properties: implicit regularization. [2] showed that, for matrix recovery, deep linear networks converge to low-rank solutions even when all the hidden layers are of size larger than or equal to the input and output sizes. [42] proved that, in deep matrix factorization, implicit regularization may not be explainable by norms, as all norms may go to infinity. They rather suggest seeing implicit regularization as a minimization of the rank. [44] and [19] proved with different assumptions on the data and a vanishing initialization that both gradient flow and discrete gradient dynamics sequentially learn solutions of a rank-constrained linear regression problem with a gradually increasing rank. Finally, [20] proved for a toy model that this incremental learning happens more often (with larger initialization), when the depth of the network increases. As we will see in Section 3.4, these results can be re-interpreted in the light of the landscape at order 2.

1.4 Summary of our contributions

Our contributions on the optimization landscape of deep linear networks can be summarized as follows.

- We characterize the square loss landscape of deep linear networks at order 2 (see Theorem 1 and Figure 6). That is, under some classical and weak assumptions on the data, we characterize, among all first-order critical points, which are global minimizers, strict saddle points, and non-strict saddle points. The characterization is simple and involves conditions on the ranks of partial matrix products. To the best of our knowledge, this is the first work that gives a simple necessary and sufficient condition for a saddle point to be strict or non-strict.
- Several results follow from the characterization: under the same assumptions,
 - we first immediately recover the fact that all saddle points are strict for one-hidden layer linear networks;
 - more importantly, for deeper networks, when proving that all cases considered in the characterization can indeed occur, we exhibit large sets of strict and non-strict saddle points (see Proposition 4 and its proof in Appendix B.8);
 - we show that the non-strict saddle points are associated with r_{max} plateau values of the empirical risk, where r_{max} is the size of the thinnest layer of the network (see Theorem 1). Typically these are values of the empirical risk that first-order algorithms can take for some time, as in Figure 3, and which might be confused with a global minimum.

- As a by-product of our analysis, we obtain explicit parameterizations of sets containing or included in the set of all first-order critical points (see Propositions 5 and 6). We also derive an explicit parameterization of the set of all global minimizers (see Proposition 7).

The above results are compared in details with previous works in Section 3.4. In particular, our second-order characterization sheds some light on two phenomena:

- **Implicit regularization:** we recover the fact that every non-strict saddle point corresponds to a global minimizer of the rank-constrained linear regression problem, as shown in [3, Proposition 35]. Our characterization additionally shows that only a fraction of the critical points corresponding to rank-constrained solutions are non-strict saddle points. The others are strict saddle points. Given the differences in the behavior of first-order algorithms in the vicinity of strict and non-strict saddle points as illustrated on Figures 3 and 4, our results open new research directions related to the very nature of implicit regularization and its stability.
- Our characterization can also be useful to understand recent global convergence results in terms of the loss landscape at order 2. In particular, we show how to re-interpret a proof of [14] to see that gradient descent with Xavier initialization on wide enough deep linear networks meets no non-strict saddle points on its trajectory.

1.5 Outline of the paper

The paper is organized as follows. We define the setting in Section 2 and state our results in Section 3. We prove our main result (Theorem 1) in Section 4. More precisely, we detail the proof structure and main arguments but defer all technical derivations to the appendix. We finally conclude our work in Section 5.

Most technical details can be found in the appendix, which is organized as follows. Section A contains additional notation and lemmas that will be useful in all subsequent sections. In Section B we provide proofs of propositions and lemmas related to first-order critical points, while Section C gathers the proofs for the parameterization of first-order critical points and global minimizers. Sections D, E, and F contain proofs corresponding to each subsection of Section 4. Finally, in Section G, we describe in more details the illustrative experiment underlying Figures 3 and 4.

2 Setting

In this section we formally define our setting (deep linear networks with square loss), set some notation, and describe our assumptions on the data.

Model and notation: We consider a fully-connected linear neural network of depth $H \geq 2$. The neural network consists of H layers and maps any input $x \in \mathbb{R}^{d_x}$ to an output $W_H \cdots W_1 x \in \mathbb{R}^{d_y}$, where $W_H \in \mathbb{R}^{d_y \times d_{H-1}}, \dots, W_h \in \mathbb{R}^{d_h \times d_{h-1}}, \dots, W_1 \in \mathbb{R}^{d_1 \times d_x}$, are the matrices associated with the H layers (d_h is the width of layer h). We set $d_H = d_y$ and $d_0 = d_x$. The input layer is of size d_x and the output layer is of size d_y . We also define the smallest width of the layers as $r_{max} = \min(d_H, \dots, d_0)$.⁴ We denote the parameters of the model by $\mathbf{W} = (W_H, \dots, W_1)$.

Let $(x_i, y_i)_{i=1..m}$ with $x_i \in \mathbb{R}^{d_x}$ and $y_i \in \mathbb{R}^{d_y}$, be the training set that we gather column-wise in matrices $X \in \mathbb{R}^{d_x \times m}$ and $Y \in \mathbb{R}^{d_y \times m}$. We consider the empirical risk L defined by:

$$L(\mathbf{W}) = \sum_{i=1}^m \|W_H W_{H-1} \cdots W_2 W_1 x_i - y_i\|_2^2 = \|W_H \cdots W_1 X - Y\|^2,$$

where $\|\cdot\|_2$ is the Euclidean norm and $\|\cdot\|$ denotes the Frobenius norm of a matrix.

We set:

$$\Sigma_{XX} = \sum_{i=1}^m x_i x_i^T = X X^T \in \mathbb{R}^{d_x \times d_x}, \quad \Sigma_{YY} = \sum_{i=1}^m y_i y_i^T = Y Y^T \in \mathbb{R}^{d_y \times d_y},$$

⁴The notation r_{max} comes from the fact that it is the maximum possible rank of the product $W_H \cdots W_1$.

$$\Sigma_{XY} = \sum_{i=1}^m x_i y_i^T = XY^T \in \mathbb{R}^{d_x \times d_y}, \quad \Sigma_{YX} = \sum_{i=1}^m y_i x_i^T = YX^T \in \mathbb{R}^{d_y \times d_x},$$

where, A^T denotes the transpose of A .

Assumption (\mathcal{H}). Throughout the article, we assume that $d_y \leq d_x \leq m$, that Σ_{XX} is invertible, and that Σ_{XY} is of full rank d_y . We define $\Sigma^{1/2} = \Sigma_{YX} \Sigma_{XX}^{-1} X \in \mathbb{R}^{d_y \times m}$ and $\Sigma = \Sigma^{1/2} (\Sigma^{1/2})^T = \Sigma_{YX} \Sigma_{XX}^{-1} \Sigma_{XY} \in \mathbb{R}^{d_y \times d_y}$. We assume that the singular values of $\Sigma^{1/2}$ are all distinct (i.e., that Σ has d_y distinct eigenvalues).

These assumptions are exactly the ones considered in [27]. Note that we do not make any assumption on the width of the hidden layers. As noted by [4], full rank matrices are dense, and deficient-rank matrices are of measure 0. In general, $m \geq d_x \geq d_y$, which is the classical learning regime, is essentially sufficient to have the other assumptions verified, due to the randomness of the data.

Let

$$\Sigma^{1/2} = U \Delta V^T \tag{1}$$

be a singular value decomposition of $\Sigma^{1/2}$, where $U \in \mathbb{R}^{d_y \times d_y}$ and $V \in \mathbb{R}^{m \times m}$ are orthogonal, and the diagonal elements of $\Delta \in \mathbb{R}^{d_y \times m}$ are in decreasing order.

Since $\Sigma = \Sigma^{1/2} (\Sigma^{1/2})^T$, Σ can be diagonalized as $\Sigma = U \Lambda U^T$ where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_{d_y})$, with $\lambda_1 > \dots > \lambda_{d_y} \geq 0$. Moreover, a consequence of Assumption (\mathcal{H}) is that Σ is positive definite (see Lemma 4); therefore, we have $\lambda_{d_y} > 0$.

Additional notation: We list below some notation and conventions that will use throughout the paper.

For all integers $a \leq b$, we denote by $\llbracket a, b \rrbracket$ the set of integers between a and b (including a and b). If $a > b$, $\llbracket a, b \rrbracket$ is the empty set (e.g. $\llbracket 1, 0 \rrbracket = \emptyset$).

If $\mathcal{S} = \emptyset$, then $\sum_{i \in \mathcal{S}} \lambda_i = 0$.

Given a matrix $A \in \mathbb{R}^{p \times q}$, $\text{col}(A)$, $\text{Ker}(A)$ and $\text{rk}(A)$, denote respectively the column space, the null space and the rank of A .

For a matrix $A \in \mathbb{R}^{p \times q}$, we write $A_i \in \mathbb{R}^p$ for the i -th column of A and $A_{\mathcal{J}} \in \mathbb{R}^{p \times |\mathcal{J}|}$ for the sub-matrix obtained by concatenating the column vectors A_i , for $i \in \mathcal{J}$. The identity matrix of size p will be denoted by I_p .

When we write $W_h \cdots W_{h'}$ for $h > h'$, the expression denotes the product of all W_j from $j = h$ to $j = h'$. For notational compactness, we allow two additional cases: when $h = h'$, the expression simply denotes W_h , and when $h' = h + 1$, it stands for the identity matrix $I_{d_h} \in \mathbb{R}^{d_h \times d_h}$.

Considering submatrices of compatible sizes, we define a block matrix by one of the three following ways:

- $[A, B]$ is the horizontal concatenation of the matrices A and B ;
- $\begin{bmatrix} G \\ H \end{bmatrix}$ is the vertical concatenation of G and H ;
- $\begin{bmatrix} C & D \\ E & F \end{bmatrix}$ is a 2×2 block matrix.

By convention, in block matrices, some blocks can have 0 lines or 0 columns; this means that such blocks do not exist. However if we have a product between two matrices that have 0 as the common size (the number of columns for the first matrix, of the lines for the second matrix), then their product equals a zero matrix, of the right size. More formally, if $A \in \mathbb{R}^{n \times 0}$ and $B \in \mathbb{R}^{0 \times p}$, then, by convention, $AB = 0_{n \times p}$. Note that the product of block matrices is compatible with this convention (e.g., $[A, B] \begin{bmatrix} C \\ D \end{bmatrix} = AC + BD$ is still true if $B \in \mathbb{R}^{n \times 0}$ and $D \in \mathbb{R}^{0 \times p}$).

Further notation that are used in the appendix can be found at the beginning of Appendix A.

3 Main results

In this section, we state the main results of this paper. We start with a necessary condition for being a first-order critical point of L (Proposition 1), to which we give a light reciprocal (Proposition 2). We then move to our main result (Theorem 1), which is a second-order classification of all first-order critical points. It distinguishes between global minimizers, strict saddle points and non-strict saddle points. Finally, the third result is a necessary parameterization for critical points (Proposition 5) and an explicit parameterization of all global minimizers (Proposition 7). These results are compared with previous works in Section 3.4. All the proofs can be found in Section 4 or in the appendix, where most technical derivations are deferred.

3.1 First-order critical points: preliminary results

In the next proposition we restate in our framework a necessary condition for being a first-order critical point, which was already present in [4] and most of the papers in this line of research. This proposition will serve later to distinguish between different types of critical points.

Proposition 1 (Global map and critical values). *Suppose Assumption \mathcal{H} in Section 2 holds true. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L and set $r = \text{rk}(W_H \cdots W_1) \in \llbracket 0, r_{max} \rrbracket$. There exists a unique subset $\mathcal{S} \subset \llbracket 1, d_y \rrbracket$ of size r such that:*

$$W_H \cdots W_1 = U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1},$$

where U was defined in (1). We say that the critical point \mathbf{W} is associated with \mathcal{S} . The associated critical value is

$$L(\mathbf{W}) = \text{tr}(\Sigma_{YY}) - \sum_{i \in \mathcal{S}} \lambda_i.$$

The proof can be found in Appendix B.2. The result is true even for $r = 0$, using the conventions from Section 2 (in this case, $\mathcal{S} = \emptyset$).

Note that $\Sigma_{YX} \Sigma_{XX}^{-1}$ corresponds to the solution of the classical linear regression problem. Therefore, we can see that for every critical point \mathbf{W} of L , the product $W_H \cdots W_1$ is the projection of this least-squares estimator onto a subspace generated by a subset of the eigenvectors of Σ . Note that $\text{tr}(\Sigma_{YY}) = \|Y\|^2$.

The following proposition is a light reciprocal to Proposition 1, by showing that all subsets \mathcal{S} and the corresponding critical values $\text{tr}(\Sigma_{YY}) - \sum_{i \in \mathcal{S}} \lambda_i$ are associated to an existing critical point. In particular the largest critical value is reached for $\mathcal{S} = \emptyset$ and the smallest critical value for $\mathcal{S} = \llbracket 1, r_{max} \rrbracket$.

Proposition 2. *Suppose Assumption \mathcal{H} in Section 2 holds true. For any $\mathcal{S} \subset \llbracket 1, d_y \rrbracket$ of size $r \in \llbracket 0, r_{max} \rrbracket$, there exists a first-order critical point \mathbf{W} associated with \mathcal{S} .*


The proof of Proposition 2 is deferred to Appendix B.6.

3.2 Second-order classification of the critical points of L

The main result of this section is Theorem 1 below, where we classify all first-order critical points into global minimizers, strict saddle points and non-strict saddle points. To state Theorem 1 we first need to introduce some definitions.

Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L . Below, we introduce the notions of *complementary block*, *tightened pivot* and *tightened critical point* that are key to the main results. Consider the sequence of H matrices W_H, \dots, W_2, W_1 and connect them by plugging Σ_{XY} between W_1 and W_H so as to form a cycle as on Figure 5. Note that the dimensions of these matrices allow us to consider any product of consecutive matrices on this cycle, e.g., $W_H W_{H-1} W_{H-2}$ or $W_2 W_1 \Sigma_{XY} W_H$ (the matrix Σ_{XY} between W_1 and W_H is key here). Such products of consecutive matrices in the cycle are what we call **“blocks”**. In the sequel, we call **“pivot”** any pair of indices $(i, j) \in \llbracket 1, H \rrbracket$, with $i > j$, and we consider blocks around a pivot (i, j) , as defined formally below.

First complementary block: $W_{j-1} \cdots W_1 \Sigma_{XY} W_H \cdots W_{i+1}$



Second complementary block: $W_{i-1} \cdots W_{j+1}$

Figure 5: Complementary blocks to the pivot (i, j) .

Definition 1 (Complementary blocks). *Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L . For any pivot $(i, j) \in \llbracket 1, H \rrbracket$, $(i > j)$, we define the two complementary blocks to (i, j) as:*

$$W_{j-1} \cdots W_1 \Sigma_{XY} W_H \cdots W_{i+1} \quad \text{and} \quad W_{i-1} \cdots W_{j+1}.$$

The general case is represented on Figure 5.

Note that, when $i = j + 1$, the second complementary block is $W_j W_{j+1}$, which using the convention in Section 2 is I_{d_j} . Similarly, if $i = H$ and $j = 1$, the first complementary block is Σ_{XY} . First we state a proposition about the ranks of the complementary blocks which is key to our analysis.

Proposition 3. *Suppose Assumption \mathcal{H} in Section 2 holds true. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L and $r = \text{rk}(W_H \cdots W_1)$. For any pivot (i, j) , the rank of each of the two complementary blocks is larger than or equal to r .*

The proof is in Appendix B.7. The boundary case when at least one of the two ranks is equal to r plays a special role in the loss landscape at order 2.

Definition 2 (Tightened pivot). *Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L and let $r = \text{rk}(W_H \cdots W_1)$.*

*We say that a pivot (i, j) is **tightened** if and only if at least one of the two complementary blocks to (i, j) is of rank r .*

Definition 3 (Tightened critical point). *Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L . We say that \mathbf{W} is **tightened** if and only if every pivot (i, j) is tightened.*

Note that a sufficient condition for a first-order critical point \mathbf{W} to be tightened is the existence of three weight matrices W_{h_1} , W_{h_2} and W_{h_3} of rank $r = \text{rk}(W_H \cdots W_1)$.

Note that when $H = 2$, there is no tightened critical point with $r < r_{max}$, because the pivot $(2, 1)$ is not tightened (both complementary blocks Σ_{XY} and I_{d_1} are of full rank, which is larger than or equal to $r_{max} = \min\{d_y, d_1, d_x\}$).

We can now state our main theorem, which characterizes the nature of any first-order critical point \mathbf{W} in terms of the associated index set \mathcal{S} and of tightening conditions. The corresponding classification, which is illustrated on Figure 6, complements the result of [27] stating that every critical point is a global minimizer or a saddle point. We recall that $r_{max} = \min(d_H, \dots, d_0)$ is the width of the thinnest layer, and that U corresponds to the eigenvectors of Σ (see (1)).




Figure 6: Second-order classification of the critical points of L .

Theorem 1 (Classification of the critical points of L). *Suppose Assumption \mathcal{H} in Section 2 holds true.*

Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L and set $r = \text{rk}(W_H \cdots W_1) \leq r_{max}$. Following Proposition 1, we consider the index set \mathcal{S} associated with \mathbf{W} .

- When $r = r_{max}$:
 - if $\mathcal{S} = \llbracket 1, r_{max} \rrbracket$, then \mathbf{W} is a global minimizer.
 - if $\mathcal{S} \neq \llbracket 1, r_{max} \rrbracket$, then \mathbf{W} is not a second-order critical point (\mathbf{W} is a strict saddle point).
- When $r < r_{max}$: \mathbf{W} is a saddle point.
 - if $\mathcal{S} \neq \llbracket 1, r \rrbracket$, then \mathbf{W} is not a second-order critical point (\mathbf{W} is a strict saddle point).
 - if $\mathcal{S} = \llbracket 1, r \rrbracket$: we have $W_H \cdots W_1 = U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \in \arg \min_{R \in \mathbb{R}^{d_y \times d_x}, \text{rk}(R) \leq r} \|RX - Y\|^2$.
 - * if \mathbf{W} is not tightened, then \mathbf{W} is not a second-order critical point (\mathbf{W} is a strict saddle point).
 - * if \mathbf{W} is tightened, then \mathbf{W} is a second-order critical point (\mathbf{W} is a non-strict saddle point).

The proof of Theorem 1 is given in Section 4, with most technical derivations deferred to the appendix. We now make two remarks. Note from the above that every non-strict saddle point corresponds to a global minimizer of the rank-constrained linear regression problem, as already shown by [3, Proposition 35]. In Remark 1 below we explain why our characterization sheds a new light on implicit regularization, and opens up research questions.

The next proposition shows the existence of both tightened and non-tightened critical points for $H \geq 3$ (there are no tightened critical points when $H = 2$ and $r < r_{max}$). Combining this result with Proposition 2 indicates that all conclusions of Theorem 1 can be observed.

Proposition 4. *Suppose Assumption \mathcal{H} in Section 2 holds true. For $H \geq 3$, for every $\mathcal{S} = \llbracket 1, r \rrbracket$ with $0 \leq r < r_{max}$, there exist both a tightened critical point and a non-tightened critical point associated with \mathcal{S} .*

The proof is postponed to Appendix B.8. It is constructive: we exhibit large sets of tightened and non-tightened critical points.

We can draw additional consequences from Theorem 1 and Propositions 2 and 4:

- For $H = 2$, for any $r < r_{max}$, there exist strict saddle points satisfying $W_H \cdots W_1 \in \arg \min_{R \in \mathbb{R}^{d_y \times d_x}, \text{rk}(R) \leq r} \|RX - Y\|^2$.
- For $H \geq 3$, for any $r < r_{max}$, there exist both strict and non-strict saddle points satisfying $W_H \cdots W_1 \in \arg \min_{R \in \mathbb{R}^{d_y \times d_x}, \text{rk}(R) \leq r} \|RX - Y\|^2$.
- In the special case $r = 0$, we have $\mathcal{S} = \emptyset$ and $\emptyset = \llbracket 1, r \rrbracket$ by convention (see Section 2), so that $\mathcal{S} = \llbracket 1, r \rrbracket$. In this case, Theorem 1 and Proposition 4 together imply that there exist both strict and non-strict saddle points \mathbf{W} such that $W_H \cdots W_1 = 0$ when $H \geq 3$.

Remark 1 (Implicit regularization). *As detailed in Sections 1.3 and 3.4, by analyzing the gradient dynamics in well-chosen settings, it has been established that the iterates trajectory passes in the vicinity of critical points \mathbf{W} such that $W_H \cdots W_1 = \arg \min_{R \in \mathbb{R}^{d_y \times d_x}, \text{rk}(R) \leq r} \|RX - Y\|^2$, for increasing $r \in \llbracket 0, r_{max} \rrbracket$. In such settings, the gradient dynamics sequentially finds the best linear regression predictor in*

$$\mathcal{D}_r = \{R \in \mathbb{R}^{d_y \times d_x}, \text{rk}(R) \leq r\},$$

for increasing r . The subset $\mathcal{D}_r \subset \mathbb{R}^{d_y \times d_x}$ is independent of X, Y and the network architecture, and plays the role of a regularization constraint in the predictor space.

From Theorem 1, we know that the critical points such that $W_H \cdots W_1 = \arg \min_{R \in \mathbb{R}^{d_y \times d_x}, \text{rk}(R) \leq r} \|RX - Y\|^2$ can be either strict saddle points or non-strict saddle points. From Proposition 4 we know that both cases exist. To the best of our knowledge, whether the saddle points approached by the iterates trajectory are strict or non-strict and the impact of this property on the implicit regularization phenomenon have not been studied. Though this study goes beyond the scope of this paper, let us sketch the main trends that we can anticipate from our results.

First, the experiments on Figures 3 and 4 suggest that, given a distance between the initial iterate and a saddle point, the number of iterations spent by a first-order algorithm in the vicinity of this saddle point is typically larger when the latter is non-strict. This suggests implicit regularization should be more easily observed around such points in practice. On the other hand, looking at the rank constraint in Definition 2 (which corresponds to the very last item of Theorem 1), we anticipate that there are much fewer non-strict saddle points than strict saddle points. Understanding where exactly implicit regularization typically occurs in realistic settings is a challenging question for future works.

3.3 Parameterization of first-order critical points and global minimizers

We now turn back to first-order critical points, and state all new related results. In our analysis, these results precede the proof of Theorem 1. The presentation has been reversed in Section 3 to highlight the main contribution of the article.

The next proposition provides an explicit parameterization of first-order critical points. Note that this is only a necessary condition.

Proposition 5. *Suppose Assumption \mathcal{H} in Section 2 holds true. at Assumption \mathcal{H} in Section 2 holds. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L associated with \mathcal{S} (cf Proposition 1), and let $Q = \llbracket 1, d_y \rrbracket \setminus \mathcal{S}$. Then, there exist invertible matrices $D_{H-1} \in \mathbb{R}^{d_{H-1} \times d_{H-1}}, \dots, D_1 \in \mathbb{R}^{d_1 \times d_1}$ and matrices $Z_H \in \mathbb{R}^{(d_y-r) \times (d_{H-1}-r)}, Z_1 \in \mathbb{R}^{(d_1-r) \times d_x}$ and $Z_h \in \mathbb{R}^{(d_h-r) \times (d_{h-1}-r)}$ for $h \in \llbracket 2, H-1 \rrbracket$ such that if we denote $\widetilde{W}_H = W_H D_{H-1}$, $\widetilde{W}_1 = D_1^{-1} W_1$ and $\widetilde{W}_h = D_h^{-1} W_h D_{h-1}$, for all $h \in \llbracket 2, H-1 \rrbracket$, then we have*

$$\widetilde{W}_H = [U_S, U_Q Z_H] \quad (2)$$

$$\widetilde{W}_1 = \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_1 \end{bmatrix} \quad (3)$$

$$\widetilde{W}_h = \begin{bmatrix} I_r & 0 \\ 0 & Z_h \end{bmatrix} \quad \forall h \in \llbracket 2, H-1 \rrbracket \quad (4)$$

$$\widetilde{W}_H \cdots \widetilde{W}_2 = [U_S, 0]. \quad (5)$$

The proposition is proved in Appendix C.1, and will be key to prove the last statement of Theorem 1. Next, we give a sufficient condition for any \mathbf{W} satisfying (2), (3) and (4), to be a first-order critical point of L .

Proposition 6. *Suppose Assumption \mathcal{H} in Section 2 holds true. Let $S \subset \llbracket 1, d_y \rrbracket$ of size $r \in \llbracket 0, r_{max} \rrbracket$ and $Q = \llbracket 1, d_y \rrbracket \setminus S$. Let $D_{H-1} \in \mathbb{R}^{d_{H-1} \times d_{H-1}}, \dots, D_1 \in \mathbb{R}^{d_1 \times d_1}$ be invertible matrices and let $Z_H \in \mathbb{R}^{(d_y-r) \times (d_{H-1}-r)}$, $Z_1 \in \mathbb{R}^{(d_1-r) \times d_x}$ and $Z_h \in \mathbb{R}^{(d_h-r) \times (d_{h-1}-r)}$ for $h \in \llbracket 2, H-1 \rrbracket$. Let the parameter of the network $\mathbf{W} = (W_H, \dots, W_1)$ be defined as follows:*

$$\begin{aligned} W_H &= [U_S, U_Q Z_H] D_{H-1}^{-1} \\ W_1 &= D_1 \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_1 \end{bmatrix} \\ W_h &= D_h \begin{bmatrix} I_r & 0 \\ 0 & Z_h \end{bmatrix} D_{h-1}^{-1} \quad \forall h \in \llbracket 2, H-1 \rrbracket. \end{aligned}$$

If $r = r_{max}$ or if there exist $h_1 \neq h_2$ such that $Z_{h_1} = 0$ and $Z_{h_2} = 0$, then, \mathbf{W} is a first-order critical point of L associated with S .

The proof of Proposition 6 is in Appendix B.5.

Note that, combining Propositions 5 and 6, we obtain an explicit parameterization of all critical points \mathbf{W} with a global map $W_H \cdots W_1$ of maximum rank r_{max} . In particular, it yields the next proposition, which provides an explicit parameterization of all the global minimizers of L .

Proposition 7 (Parameterization of all global minimizers). *Suppose Assumption \mathcal{H} in Section 2 holds true. Set $S_{max} = \llbracket 1, r_{max} \rrbracket$ and $Q_{max} = \llbracket 1, d_y \rrbracket \setminus S_{max} = \llbracket r_{max} + 1, d_y \rrbracket$. Then, $\mathbf{W} = (W_H, \dots, W_1)$ is a global minimizer of L if and only if there exist invertible matrices $D_{H-1} \in \mathbb{R}^{d_{H-1} \times d_{H-1}}, \dots, D_1 \in \mathbb{R}^{d_1 \times d_1}$, and matrices $Z_H \in \mathbb{R}^{(d_y-r_{max}) \times (d_{H-1}-r_{max})}$, $Z_h \in \mathbb{R}^{(d_h-r_{max}) \times (d_{h-1}-r_{max})}$ for $h \in \llbracket 2, H-1 \rrbracket$, and $Z_1 \in \mathbb{R}^{(d_1-r_{max}) \times d_x}$ such that:*

$$\begin{aligned} W_H &= [U_{S_{max}}, U_{Q_{max}} Z_H] D_{H-1}^{-1} \\ W_1 &= D_1 \begin{bmatrix} U_{S_{max}}^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_1 \end{bmatrix} \\ W_h &= D_h \begin{bmatrix} I_{r_{max}} & 0 \\ 0 & Z_h \end{bmatrix} D_{h-1}^{-1} \quad \forall h \in \llbracket 2, H-1 \rrbracket. \end{aligned}$$

The proof is in Appendix C.2. See in particular a remark in the same appendix on how to interpret the above formulas precisely (some blocks Z_h have 0 lines or columns).

3.4 Comparison with previous works

Next we further detail our contributions in light of earlier works.

Parameterization of global minimizers. To the best of our knowledge, Proposition 7 is the first explicit parameterization of the set of all global minimizers for deep linear networks and the square loss. For $H \geq 2$, it had been previously noted by [52] that a critical point \mathbf{W} is a global minimizer if and only if $\text{rk}(W_H \cdots W_1) = r_{max}$ and $\text{col}(W_H \cdots W_{d_{p+1}}) = \text{col}(U_{S_{max}})$, where $S_{max} = \llbracket 1, r_{max} \rrbracket$ and where p is any layer with the smallest width r_{max} . This is an implicit characterization.

Another previous work which characterized global minimizers is [55], but their characterization is not explicit: the weight matrices are defined recursively and should satisfy some equations, while in Proposition 7 the weight matrices are given explicitly. The same remark holds for their characterization of first-order critical points.

Saddle points. Among saddle points, we give a characterization of those that are strict and those that are not. Previously, for $H \geq 3$, it had been noted by [27] that $(0, \dots, 0)$ is a non-strict saddle point. This result also

follows from Theorem 1 since any critical point is tightened whenever at least 3 weight matrices are of rank $r = \text{rk}(W_H \cdots W_1)$ (which is the case for $(0, \dots, 0)$ with $r = 0$).

Also, Theorem 1 generalizes two results from [27] and [10] about sufficient conditions for strict saddle points. Indeed, [27] proved that, if \mathbf{W} is a saddle point such that $\text{rk}(W_{H-1} \cdots W_2) = r_{max}$, then \mathbf{W} is a strict saddle point. [10] proved under further assumptions on the data and the architecture that a sufficient condition for a saddle point to be strict is that $\text{rk}(W_{H-1} \cdots W_2) > r = \text{rk}(W_H \cdots W_1)$. Note that both results are special cases of Theorem 1, with the pivot $(H, 1)$. More precisely, assume that \mathbf{W} is a saddle point such that either $\text{rk}(W_{H-1} \cdots W_2) = r_{max} = r = \text{rk}(W_H \cdots W_1)$ or $\text{rk}(W_{H-1} \cdots W_2) > r = \text{rk}(W_H \cdots W_1)$ (which includes both conditions above). Then, if $\mathcal{S} \neq \llbracket 1, r \rrbracket$ (whether $r = r_{max}$ or not), by Theorem 1, \mathbf{W} is a strict saddle point without any condition on \mathbf{W} . But if $\mathcal{S} = \llbracket 1, r \rrbracket$ with $r < r_{max}$, our assumption above implies that the pivot $(H, 1)$, and therefore \mathbf{W} , is not tightened (recall that $\text{rk}(\Sigma_{XY}) = d_y \geq r_{max} > r$). In any case, \mathbf{W} is a strict saddle point.

Finally, Theorem 1 generalizes another result of [27] stating that all saddle points are strict for one-hidden layer linear networks. Indeed, let $H = 2$ and assume that we have a saddle point associated with $\mathcal{S} = \llbracket 1, r \rrbracket$ for $r < r_{max}$ (the only case where we can expect to see non-strict saddle points, by Theorem 1). Since $H = 2$, there is only one pivot which is $(2, 1)$; this pivot is not tightened because the complementary blocks are I_{d_1} and Σ_{XY} and both are of rank larger than or equal to r_{max} . Therefore, by Theorem 1, when $H = 2$ (and under Assumption \mathcal{H}), all saddle points are strict.

Convergence to global minimizer: an example where gradient descent meets no non-strict saddle points. Some recent works on deep linear networks proved under assumptions on the data, the initialization, or the minimum width of the network, that gradient descent or variants converge to a global minimum in polynomial time (e.g., [1, 5, 15, 14]). Since for general non-convex functions, gradient descent may get stuck at a non-strict saddle point, and since non-strict saddle points exist for any linear neural network of depth $H \geq 3$, it seemed impossible to deduce convergence to a global minimum using landscape results only. Instead, papers such as [14] chose to “directly analyze the trajectory generated by [...] gradient descent”.

It turns out that our characterization of strict saddle points can help re-interpret such global convergence results. Consider for instance the work of [14], who proved that with high probability gradient descent with Xavier initialization converges to a global minimum for any deep linear network which is wide enough. They analyse a network where all hidden layers have a width d_{hidden} at least proportional to the number H of layers and to other quantities depending on the data X, Y , the output dimension d_y , and the desired probability level. In their analysis, [14, Section 7] prove that with high probability, a condition $\mathcal{B}(t)$ holds at every iteration t . *Importantly, this condition implies that the point \mathbf{W} output by gradient descent at iteration t cannot be a non-strict saddle point.* Indeed, using our notation, the condition $\mathcal{B}(t)$ yields the lower-bound $\sigma_{\min}(W_H \cdots W_2) \geq \frac{3}{4} d_{\text{hidden}}^{(H-1)/2} > 0$, which in particular entails that the matrix product $W_H \cdots W_2$ is of full rank $\min\{d_{\text{hidden}}, d_y\} \geq r_{max}$. Let us check that if \mathbf{W} is a saddle point, then it is necessarily strict. By Theorem 1, either $r = \text{rk}(W_H \cdots W_1)$ is equal to r_{max} , in which case the saddle point \mathbf{W} is indeed strict, or $r < r_{max}$, in which case the pivot $(H, 1)$ is not tightened (since the two blocks Σ_{XY} and $W_{H-1} \cdots W_2$ are of rank at least r_{max}), so that the saddle point \mathbf{W} is strict, as previously claimed.

As a consequence, our characterization of strict saddle points in Theorem 1 helps re-interpret the analysis of [14, Section 7]: under Assumption \mathcal{H} , and for wide enough deep linear networks, gradient descent with Xavier initialization meets no non-strict saddle points on its trajectory.

Implicit regularization As previously explained, [29, 24, 26] and others proved that gradient-based algorithms can escape strict saddle points and be stopped at approximate second-order critical points after a number of iterations which is at most polynomial in the desired accuracy. The works [1, 5, 15, 14] mentioned above show that convergence to a global minimum can even be guaranteed under assumptions on the data, the initialization, or the minimum width of the network. Though global convergence outside these assumptions has been conjectured (e.g., [10] for gradient flow), when $H \geq 3$, we cannot yet rule out the possibility that non-stochastic gradient-based algorithms remain on a plateau around one of the r_{max} non-strict saddle points that we identified in Theorem 1 and Proposition 4. Since non-strict saddle points \mathbf{W} satisfy $W_H \cdots W_1 = \arg \min_{R \in \mathbb{R}^{d_y \times d_x}, \text{rk}(R) \leq r} \|RX - Y\|^2$, this can be seen as a form of long-term implicit regularization.

Furthermore, even under assumptions guaranteeing the convergence to a global minimum, gradient-based algo-

rithms can spend some time around non-strict saddle points before convergence (as in Figures 3 and 4), which would correspond to a form of short-term implicit regularization.

We now outline some relationships between our characterization and several earlier results on implicit regularization. [3] proved that gradient flow converges to global minimizers of the rank-constrained linear regression problem. In fact, what they proved can be understood as follows: "the only critical points the gradient flow almost surely converges to are the global minimizer or non-strict saddle points. The latter lead to global minimizers of the rank-constrained linear regression problem". In Theorem 1 and Proposition 4, we prove the existence and characterize such points, and in addition to non-strict saddle points we prove that some \mathbf{W} leading to the solution of the rank-constrained linear regression problem are strict saddle points. Doing so, we characterize and drastically reduce the implicit regularization set.

[2] found for small initializations and step size that for matrix recovery, deep matrix factorization favors solutions of low-rank. In the same context, [42] stated that, implicit regularization in deep linear networks should be seen as a minimization of rank rather than norms. Again, Theorem 1 and Proposition 4 specify the implicit regularization happening there.

The next two papers showcased the implicit regularization phenomenon for specific gradient flow or gradient descent trajectories. Our landscape analysis can help read their results from a new perspective. [19] proved that for $H = 2$, for a vanishing initialization and a small enough step size, the discrete gradient dynamics sequentially learns solutions of the rank-constrained linear regression problem with a gradually increasing rank. More precisely, the algorithm avoids all critical points associated with $\mathcal{S} \neq \llbracket 1, r \rrbracket$, but comes close to a critical point associated with $\mathcal{S} = \llbracket 1, r \rrbracket$, spends some time around it and decreases again. In the light of our work, we know that for $H = 2$ all saddle points are strict, but for $H \geq 3$, we know that there exist non-strict saddle points associated with $\mathcal{S} = \llbracket 1, r \rrbracket$. If we could extend [19] to $H \geq 3$, we expect that the gradient dynamics converges to a non-strict saddle point or spends much more time than for $H = 2$ around non-strict saddle points associated with $\mathcal{S} = \llbracket 1, r \rrbracket$. In both cases these facts would show that the implicit regularization outlined by [19] for $H = 2$ intensifies with depth. In fact, this is the result presented in [20] for a toy linear network, as they proved that, for $H = 2$, the algorithms need an exponentially vanishing initialization for this incremental learning to occur, while for $H \geq 3$, a polynomially vanishing initialization is enough. This indicates that this incremental learning arises more naturally in deep networks.

Perspectives. Even if gradient-based algorithms, and in particular SGD, could escape non-strict saddle points in infinite time with probability one in all cases, can we better formalize that the vicinity of these saddle points is finite-time stable? This way, early stopping would provably be a source of short-term implicit regularization and might explain good generalization guarantees as in [9, 53, 40, 41].

A related and interesting project would be to assess the dimensions and sizes of the implicit regularization sets for the different values of r . This should be strongly related to the expected number of epochs spent by a stochastic algorithm on the corresponding plateau.

Finally, another interesting future research direction is to determine the basin of attraction of the non-strict saddle points for non-stochastic gradient-based algorithms, for various initializations and various dimensions of the deep linear networks. Do they have zero Lebesgue measure in general, so that gradient dynamics would almost surely converge to a global minimizer, as conjectured by [10] for gradient flow? or do they have positive measure for some networks and some data, hence implying a long-term implicit regularization?

4 Proof of Theorem 1

The proof of Theorem 1 proceeds in several steps. In the end (see page 18), it will directly follow from Propositions 8, 9, 10 below and from Lemma 5 in Appendix A. In this section, we outline the overall proof structure and state the main intermediate results. We also provide proof sketches for these intermediate results, but defer many technical details to the appendix.

In our proofs, we will not compute the Hessian $\nabla^2 L(\mathbf{W})$ explicitly since this might be quite tedious. To show that a point \mathbf{W} is (or is not) a second-order critical point of L , we will instead Taylor-expand $L(\mathbf{W} + t\mathbf{W}')$ along any direction \mathbf{W}' and use the following lemma. Its proof follows directly from Taylor's theorem.

Lemma 1 (Characterization of first-order and second-order critical points). *Let $\mathbf{W} = (W_H, \dots, W_1)$. Assume that, for all $\mathbf{W}' = (W'_H, \dots, W'_1)$, the loss $L(\mathbf{W} + t\mathbf{W}')$ admits the following asymptotic expansion when $t \rightarrow 0$:*

$$L(\mathbf{W} + t\mathbf{W}') = L(\mathbf{W}) + c_1(\mathbf{W}, \mathbf{W}')t + c_2(\mathbf{W}, \mathbf{W}')t^2 + o(t^2). \quad (6)$$

Then:

- \mathbf{W} is a first-order critical point of L iff $c_1(\mathbf{W}, \mathbf{W}') = 0$ for all \mathbf{W}' .
- \mathbf{W} is a second-order critical point of L iff $c_1(\mathbf{W}, \mathbf{W}') = 0$ and $c_2(\mathbf{W}, \mathbf{W}') \geq 0$ for all \mathbf{W}' .
Therefore if for a first-order critical point \mathbf{W} , we can exhibit a direction \mathbf{W}' such that $c_2(\mathbf{W}, \mathbf{W}') < 0$, then \mathbf{W} is not a second-order critical point.

We divide the proof of Theorem 1 into three parts. Recall that from [27], we know that all first-order critical points are either global minimizers or saddle points (that is, there is no local extrema apart from global minimizers). We refine this classification.

4.1 Global minimizers and 'simple' strict saddle points

In this section, we start by identifying simple sufficient conditions on the support \mathcal{S} associated to a first-order critical point \mathbf{W} which guarantee that \mathbf{W} is either a global minimizer or a strict saddle point. More subtle strict saddle points and non-strict saddle points will be addressed in Sections 4.2 and 4.3.

Proposition 8. *Suppose Assumption \mathcal{H} in Section 2 holds true. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L associated with \mathcal{S} and set $r = \text{rk}(W_H \cdots W_1) \leq r_{max}$.*

- When $r = r_{max}$:
 - if $\mathcal{S} = \llbracket 1, r_{max} \rrbracket$, then \mathbf{W} is a global minimizer.
 - if $\mathcal{S} \neq \llbracket 1, r_{max} \rrbracket$, then \mathbf{W} is not a second-order critical point (\mathbf{W} is a strict saddle point).
- When $r < r_{max}$: \mathbf{W} is a saddle point.
 - if $\mathcal{S} \neq \llbracket 1, r \rrbracket$, then \mathbf{W} is not a second-order critical point (\mathbf{W} is a strict saddle point).

The proof is postponed to Appendix D. To prove that \mathbf{W} associated with $\mathcal{S} \neq \llbracket 1, r \rrbracket$, $r \leq r_{max}$ is not a second-order critical point, we explicitly exhibit a direction \mathbf{W}' such that the second-order coefficient $c_2(\mathbf{W}, \mathbf{W}')$ in the Taylor expansion of $L(\mathbf{W} + t\mathbf{W}')$ around $t = 0$, in (6), is negative. Using Lemma 1, we conclude that \mathbf{W} is not a second-order critical point.

Recall from Proposition 1 that the loss at any first-order critical point is given by $\text{tr}(\Sigma_{YY}) - \sum_{i \in \mathcal{S}} \lambda_i$. The spirit of the proof is that critical points associated with $\mathcal{S} \neq \llbracket 1, r \rrbracket$ capture a smaller singular value λ_j instead of a larger one λ_i with $i < j$. Thus, to see that the loss can be further decreased at order 2 (and is therefore not a second-order critical point by Lemma 1), a natural proof strategy is to perturb the singular vector corresponding to λ_j along the direction of the singular vector corresponding to λ_i . This part of the proof is an adaption of the proof of [4].

4.2 Strict saddle points associated with $\mathcal{S} = \llbracket 1, r \rrbracket$, $r < r_{max}$

We now address situations that to our knowledge, have never been addressed, in the literature. We prove the following.

Proposition 9. *Suppose Assumption \mathcal{H} in Section 2 holds true. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L associated with $\mathcal{S} = \llbracket 1, r \rrbracket$, with $0 \leq r < r_{max}$.
If \mathbf{W} is not tightened, then \mathbf{W} is not a second-order critical point (\mathbf{W} is a strict saddle point).*

We sketch the main arguments below. We will again construct a direction \mathbf{W}' such that the second-order coefficient $c_2(\mathbf{W}, \mathbf{W}')$ in the asymptotic expansion of $L(\mathbf{W} + t\mathbf{W}')$ around $t = 0$, in (6), is negative. More precisely, for a first-order critical point \mathbf{W} , for any $\beta \in \mathbb{R}$, we will consider a well-chosen \mathbf{W}'_β such that $c_2(\mathbf{W}, \mathbf{W}'_\beta) = a\beta^2 + c\beta$ for some constants a, c (possibly depending on \mathbf{W}) such that $a \geq 0$ and $c \neq 0$. Taking

$$\beta = \begin{cases} -c & \text{if } a = 0 \\ -\frac{c}{2a} & \text{if } a > 0 \end{cases} \quad (7)$$

we obtain

$$c_2(\mathbf{W}, \mathbf{W}'_\beta) = \begin{cases} -c^2 & \text{if } a = 0 \\ -\frac{c^2}{4a} & \text{if } a > 0 \end{cases}$$

and therefore

$$c_2(\mathbf{W}, \mathbf{W}'_\beta) < 0.$$

Using Lemma 1, we can conclude that \mathbf{W} is not a second-order critical point.

We now provide intuitions on how to choose \mathbf{W}' . Since \mathbf{W} is not tightened, there exists a pivot (i, j) , with $i > j$, which is not tightened. Depending on the values of i and j we will construct \mathbf{W}' differently. However, the strategy for constructing \mathbf{W}' is the same in all cases.

Recall again that from Proposition 1, at any first-order critical point \mathbf{W} , the value of the loss is given by $tr(\Sigma_{YY}) - \sum_{i \in \mathcal{S}} \lambda_i$. Contrary to the previous section, since $\mathcal{S} = \llbracket 1, r \rrbracket$ there is no immediate way to decrease the loss (at order 2) without increasing the rank of the product of the weight matrices. Indeed, we have $W_H \cdots W_1 = U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \in \arg \min_{\text{rk}(R) \leq r} \|RX - Y\|^2$.

Therefore, to be able to decrease the value of the loss, we need to perturb \mathbf{W} in a way that the product of the perturbed parameter weight matrices becomes of rank strictly larger than r . Also, to prove that \mathbf{W} is not a second-order critical point, we need to decrease the loss at order 2. This is possible when \mathbf{W} is not tightened. For the non-tightened pivot (i, j) , we choose a perturbation \mathbf{W}' with all $W'_h = 0$ except for W'_i and W'_j . Furthermore, our construction of W'_i and W'_j depends on whether i and/or j are on the boundary $\{1, H\}$. This is due to the fact that H and 1 play a special role in the product of the perturbed weights $(W_H + tW'_H) \cdots (W_1 + tW'_1)$. This is why we distinguish the four cases below:

- 1st case: $i \in \llbracket 2, H - 1 \rrbracket$ and $j = 1$. This case is treated in Appendix E.1.
- 2nd case: $i = H$ and $j = 1$. This case is treated in Appendix E.2.
- 3rd case: $i = H$ and $j \in \llbracket 2, H - 1 \rrbracket$. This case is treated in Appendix E.3.
- 4th case: $i, j \in \llbracket 2, H - 1 \rrbracket$ with $i > j$. This case is treated in Appendix E.4.

4.3 Non-strict saddle points

We now provide a sketch of the proof for the converse of Proposition 9, as stated in Proposition 10 below. All the proofs related to this section are deferred to Appendix F.

Proposition 10. *Suppose Assumption \mathcal{H} in Section 2 holds true. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L associated with $\mathcal{S} = \llbracket 1, r \rrbracket$, $0 \leq r < r_{max}$.*

If \mathbf{W} is tightened, then \mathbf{W} is a second-order critical point (\mathbf{W} is a non-strict saddle point).

To prove Proposition 10, we first state a proposition which indicates that multiplications by invertible matrices do not change the nature of the critical point.

Lemma 2. *For all $h \in \llbracket 1, H - 1 \rrbracket$, let $D_h \in \mathbb{R}^{d_h \times d_h}$ be an invertible matrix. We define $\widetilde{W}_H = W_H D_{H-1}$, $\widetilde{W}_1 = D_1^{-1} W_1$ and $\widetilde{W}_h = D_h^{-1} W_h D_{h-1}$, for all $h \in \llbracket 2, H - 1 \rrbracket$. Then*

- $\mathbf{W} = (W_H, \dots, W_1)$ is a first-order critical point of L if and only if $\widetilde{\mathbf{W}} = (\widetilde{W}_H, \dots, \widetilde{W}_1)$ is a first-order critical point of L .
- $\mathbf{W} = (W_H, \dots, W_1)$ is a second-order critical point of L if and only if $\widetilde{\mathbf{W}} = (\widetilde{W}_H, \dots, \widetilde{W}_1)$ is a second-order critical point of L .

The lemma is proved in Appendix B.4.

Proposition 10 is then obtained using Proposition 5 (note that when \mathbf{W} is tightened, $\widetilde{\mathbf{W}}$ is also tightened since the rank of a matrix does not change when multiplied by invertible matrices), by showing that $\widetilde{\mathbf{W}} = (\widetilde{W}_H, \dots, \widetilde{W}_1)$ as given by Proposition 5 is a second-order critical point of L and using Lemma 2 to conclude that \mathbf{W} is a second-order critical point. This is easier since $\widetilde{\mathbf{W}}$ has a simpler form.

More precisely, we have the following result, from which Proposition 10 follows (see Appendix F.2 for details).

Proposition 11. *Suppose Assumption \mathcal{H} in Section 2 holds true. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L associated with $\mathcal{S} = \llbracket 1, r \rrbracket$ with $0 \leq r < r_{max}$ such that there exist matrices $Z_H \in \mathbb{R}^{(d_y-r) \times (d_H-1-r)}$, $Z_1 \in \mathbb{R}^{(d_1-r) \times d_x}$ and $Z_h \in \mathbb{R}^{(d_h-r) \times (d_h-1-r)}$ for $h \in \llbracket 2, H-1 \rrbracket$ with*

$$W_H = [U_{\mathcal{S}}, U_Q Z_H] \quad (8)$$

$$W_1 = \begin{bmatrix} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_X^{-1} \\ Z_1 \end{bmatrix} \quad (9)$$

$$W_h = \begin{bmatrix} I_r & 0 \\ 0 & Z_h \end{bmatrix} \quad \forall h \in \llbracket 2, H-1 \rrbracket \quad (10)$$

$$W_H \cdots W_2 = [U_{\mathcal{S}}, 0], \quad (11)$$

where $Q = \llbracket 1, d_y \rrbracket \setminus \mathcal{S}$.

If \mathbf{W} is tightened, then \mathbf{W} is a second-order critical point of L .

Proposition 11 is proved in details in Section F.1. We provide a proof sketch below.

We denote, for t in the neighborhood of 0, and $h \in \llbracket 1, H \rrbracket$, $W_h(t) = W_h + tW'_h$ where $W'_h \in \mathbb{R}^{d_h \times d_{h-1}}$ is arbitrary. We define $\mathbf{W}(t) := (W_H(t), \dots, W_1(t))$ and $W(t) := W_H(t) \cdots W_1(t)$. As in the previous two sections, we use Lemma 1. However, this time, we show that the second-order coefficient $c_2(\mathbf{W}, \mathbf{W}')$ is non-negative for all directions \mathbf{W}' .

To compute the loss $\|W(t)X - Y\|^2$, we expand

$$\begin{aligned} W(t) &= W_H(t) \cdots W_1(t) \\ &= (W_H + tW'_H) \cdots (W_1 + tW'_1) \\ &= W_H \cdots W_1 + t \sum_{i=1}^H W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_1 \\ &\quad + t^2 \sum_{H \geq i > j \geq 1} W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_{j+1} W'_j W_{j-1} \cdots W_1 + o(t^2). \end{aligned}$$

Therefore,

$$\begin{aligned} L(\mathbf{W}(t)) &= \left\| W_H \cdots W_1 X - Y + t \sum_{i=1}^H W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_1 X \right. \\ &\quad \left. + t^2 \sum_{H \geq i > j \geq 1} W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_{j+1} W'_j W_{j-1} \cdots W_1 X + o(t^2) \right\|^2. \end{aligned}$$

We can now easily calculate the second-order coefficient $c_2(\mathbf{W}, \mathbf{W}')$ in the Taylor expansion of $L(\mathbf{W}(t))$ around $t = 0$ (in (6)).

Recalling that $c_2(\mathbf{W}, \mathbf{W}')$ is such that $L(\mathbf{W}(t)) = L(\mathbf{W}) + c_2(\mathbf{W}, \mathbf{W}')t^2 + o(t^2)$ (since \mathbf{W} is a first-order critical point), we have

$$c_2(\mathbf{W}, \mathbf{W}') = \left\| \sum_{i=1}^H W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_1 X \right\|^2 + 2 \left\langle \sum_{H \geq i > j \geq 1} W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_{j+1} W'_j W_{j-1} \cdots W_1 X, W_H \cdots W_1 X - Y \right\rangle,$$

where $\langle A, B \rangle = \text{tr}(AB^T)$. In order to simplify the notation and equations, we define, for all $i \in \llbracket 1, H \rrbracket$,

$$T_i = W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_1 X, \quad (12)$$

and for all $i, j \in \llbracket 1, H \rrbracket$ with $i > j$:

$$T_{i,j} = \langle W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_{j+1} W'_j W_{j-1} \cdots W_1 X, W_H \cdots W_1 X - Y \rangle. \quad (13)$$

Then we set

$$FT = \left\| \sum_{i=1}^H T_i \right\|^2, \quad (14)$$

and

$$ST = 2 \sum_{H \geq i > j \geq 1} T_{i,j}. \quad (15)$$

The coefficient becomes

$$c_2(\mathbf{W}, \mathbf{W}') = \left\| \sum_{i=1}^H T_i \right\|^2 + 2 \sum_{H \geq i > j \geq 1} T_{i,j} = FT + ST.$$

Using the fact that \mathbf{W} is tightened, some weight products become simple (see Lemma 14) and we can simplify T_i and $T_{i,j}$ (see Lemmas 21 and 22 in Appendix F).

This allows us to establish that, for any \mathbf{W}' , there exist matrices A_2, A_3, A_4 and a non-negative scalar a_1 such that $FT = a_1 + \|A_2\|^2 + \|A_3\|^2 + \|A_4\|^2$ (see Appendix F.1.2) and $ST = -2 \langle A_3, A_4 \rangle$ (see Appendix F.1.3). Therefore

$$c_2(\mathbf{W}, \mathbf{W}') = FT + ST = a_1 + \|A_2\|^2 + \|A_3 - A_4\|^2 \geq 0,$$

and using Lemma 1 we conclude that \mathbf{W} is a second-order critical point.

We are now in a position to prove Theorem 1 as a direct corollary from the above results.

Proof of Theorem 1. The classification into global minimizers, strict saddle points, and non-strict saddle points follows directly from Propositions 8, 9, and 10 above. As for the fact that

$$W_H \cdots W_1 = U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \in \arg \min_{R \in \mathbb{R}^{d_y \times d_x}, \text{rk}(R) \leq r} \|RX - Y\|^2$$

when $S = \llbracket 1, r \rrbracket$, it follows from Proposition 1 above and from Lemma 5 in Appendix A. \square

5 Conclusion

We studied the optimization landscape of linear neural networks of arbitrary depth with the square loss. We first derived a necessary condition for being a first-order critical point by associating any of them with a set of eigenvectors of a data-dependent matrix. We then provided a complete characterization of the landscape at order 2 by distinguishing between global minimizers, strict saddle points, and non-strict saddle points. As a by-product of this analysis, we exhibited large sets of strict and non-strict saddle points and derived an explicit parameterization of all global minimizers. Our second-order characterization also sheds some light on the implicit regularization that may be induced by first-order algorithms, by proving that non-strict saddle points and some strict saddle points are among the global minimizers of the rank-constrained linear regression problem. It also helps re-interpret a recent convergence result, stating that gradient descent with Xavier initialization converges to a global minimum for any wide enough deep linear network.

Acknowledgements

Our work has benefited from the AI Interdisciplinary Institute ANITI. ANITI is funded by the French "Investing for the Future – PIA3" program under the Grant agreement n°ANR-19-PI3A-0004. The authors gratefully acknowledge the support of the DEEL project.⁵

A Notation and useful properties

In this section, we define some additional notation and terminology that will be used through all subsequent appendices. We also state simple linear algebra facts (Section A.2), together with some properties about the Moore-Penrose inverse (Section A.3). Since most of the proofs rely on linear algebra, we recommend the unfamiliar reader to check classical textbooks.

Additional notation: If a matrix A has already a subscript like W_H for example, we denote by $(W_H)_{.,i}$ the i -th column and by $(W_H)_{.,J}$ the sub-matrix obtained by concatenating the column vectors $(W_H)_{.,i}$, for all $i \in \mathcal{J}$. Also $(W_H)_{i,.}$ denotes the i -th row of W_H and $(W_H)_{\mathcal{I},.}$ the sub-matrix obtained by concatenating the line vectors $(W_H)_{i,.}$, for all $i \in \mathcal{I}$. More generally $(W_H)_{\mathcal{I},\mathcal{J}}$ denotes the matrix W_H restricted to the index set $\mathcal{I} \times \mathcal{J}$. For instance, $(W_H)_{1:r,r+1:d_{H-1}} \in \mathbb{R}^{r \times (d_{H-1}-r)}$ is the matrix formed from W_H by keeping the rows from 1 to r and the columns from $r + 1$ to d_{H-1} . The symbol $\delta_{i,j}$ denotes the Kronecker index which equal to 0 if $i \neq j$ and 1 if $i = j$.

Also, we define the partial gradients with respect to each weight matrix as follows.

A.1 Partial gradients

Definition 4 (gradient and partial gradients of L). *Since the input $\mathbf{W} = (W_H, \dots, W_1)$ of $L(\mathbf{W})$ is not a vector but a sequence of matrices, we define the gradient $\nabla L(\mathbf{W})$ of L at \mathbf{W} with a similar format :*

$$\nabla L(\mathbf{W}) = (\nabla_{W_H} L(\mathbf{W}), \dots, \nabla_{W_1} L(\mathbf{W})) ,$$

where each partial gradient $\nabla_{W_h} L(\mathbf{W}) \in \mathbb{R}^{d_h \times d_{h-1}}$ is the matrix whose entries are the partial derivatives $\frac{\partial L}{\partial (W_h)_{i,j}}$ for $i = 1, \dots, d_h$ and $j = 1, \dots, d_{h-1}$

The next lemma provides explicit formulas for the partial gradients of L . A proof can be found at the end of [52].

Lemma 3. *Let $h \in \llbracket 2, H - 1 \rrbracket$. The partial gradient of L with respect to W_h is:*

$$\nabla_{W_h} L(\mathbf{W}) = 2(W_H \cdots W_{h+1})^T (W_H \cdots W_1 \Sigma_{XX} - \Sigma_{YX}) (W_{h-1} \cdots W_1)^T .$$

⁵<https://www.deel.ai/>

We also have the partial gradient with respect to W_H :

$$\nabla_{W_H} L(\mathbf{W}) = 2(W_H \cdots W_1 \Sigma_{XX} - \Sigma_{YX})(W_{H-1} \cdots W_1)^T.$$

Finally, the partial gradient with respect to W_1 is:

$$\nabla_{W_1} L(\mathbf{W}) = 2(W_H \cdots W_2)^T (W_H \cdots W_1 \Sigma_{XX} - \Sigma_{YX}).$$

A.2 Simple linear algebra facts

Recall that $\Sigma^{1/2} = \Sigma_{YX} \Sigma_{XX}^{-1} X$ and $\Sigma = \Sigma^{1/2} (\Sigma^{1/2})^T = \Sigma_{YX} \Sigma_{XX}^{-1} \Sigma_{XY}$. Recall also from (1) that $\Sigma^{1/2} = U \Delta V^T$ is a Singular Value Decomposition, where $U \in \mathbb{R}^{d_y \times d_y}$ and $V \in \mathbb{R}^{m \times m}$ are orthogonal matrices.

Lemma 4. *Suppose Assumption \mathcal{H} in Section 2 holds true. Then Σ is invertible.*

Proof. Given the definition of $\Sigma^{1/2}$, it is a standard fact of linear algebra that $\text{rk}(\Sigma^{1/2}) = \text{rk}(\Sigma_{YX} \Sigma_{XX}^{-1} X) \leq \text{rk}(\Sigma_{YX})$. On the other hand, $\text{rk}(\Sigma^{1/2}) = \text{rk}(\Sigma_{YX} \Sigma_{XX}^{-1} X) \geq \text{rk}(\Sigma_{YX} \Sigma_{XX}^{-1} X X^T) = \text{rk}(\Sigma_{YX})$ since $\Sigma_{XX} = X X^T$. Therefore $\text{rk}(\Sigma^{1/2}) = \text{rk}(\Sigma_{YX}) = d_y$ by Assumption \mathcal{H} . Finally, using another fact of linear algebra we have $\text{rk}(\Sigma) = \text{rk}(\Sigma^{1/2} (\Sigma^{1/2})^T) = \text{rk}(\Sigma^{1/2})$, and therefore $\text{rk}(\Sigma) = d_y$. Hence, Σ is invertible. \square

The next lemma is about global minimizers of the rank-constrained linear regression problem.

Lemma 5. *Suppose Assumption \mathcal{H} in Section 2 holds true. Let $\mathcal{S} = \llbracket 1, r \rrbracket$. We have*

$$U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} \in \arg \min_{R \in \mathbb{R}^{d_y \times d_x}, \text{rk}(R) \leq r} \|RX - Y\|^2.$$

Proof. A proof of can be found in [52]. \square

We now present a lemma with elementary properties that we will use frequently and that are related to the orthogonality of U . The proof is straightforward.

Lemma 6. *We have the following properties related to the orthogonality of the matrix U :*

- We have $I_{d_y} = U U^T = U^T U$.
- For any $i, j \in \llbracket 1, d_y \rrbracket$, we have $U_i^T U_j = \delta_{i,j}$.
- For any $I, J \subset \llbracket 1, d_y \rrbracket$ such that $I \cap J = \emptyset$, we have $U_I^T U_J = 0_{|I| \times |J|}$.
- For any $I, J \subset \llbracket 1, d_y \rrbracket$ such that $I \cap J = \emptyset$ and $I \cup J = \llbracket 1, d_y \rrbracket$, we have $I_{d_y} = U_I U_I^T + U_J U_J^T$.
- For any $J \subset \llbracket 1, d_y \rrbracket$, we have $U_J^T U_J = I_{|J|}$ and $\text{rk}(U_J U_J^T) = |J|$.

Note that the same applies also to the other orthogonal matrix $V \in \mathbb{R}^{m \times m}$ appearing in the Singular Value Decomposition of $\Sigma^{1/2}$ (we only replace d_y by m).

Another useful lemma is the following:

Lemma 7. *For any $I, J \subset \llbracket 1, d_y \rrbracket$ such that $I \cap J = \emptyset$, we have*

$$U_I^T \Sigma U_J = 0_{|I| \times |J|}.$$

In particular, for any $\mathcal{S} \subset \llbracket 1, d_y \rrbracket$ and $Q = \llbracket 1, d_y \rrbracket \setminus \mathcal{S}$, we have $U_{\mathcal{S}}^T \Sigma U_Q = 0$.

Proof. We have, for any $k \in \llbracket 1, d_y \rrbracket$, $\Sigma U_k = \lambda_k U_k$. Hence for $j \neq k$ we have $U_j^T \Sigma U_k = \lambda_k U_j^T U_k = 0$ since U is orthogonal. Therefore, if we take two disjoint sets $J = \{j_1, \dots, j_p\}, K = \{k_1, \dots, k_n\} \subset \llbracket 1, d_y \rrbracket$, the coefficient in the position (l, m) of the matrix $U_J^T \Sigma U_K$ is equal to $U_{j_l} \Sigma U_{k_m}$ which is zero, since $j_l \neq k_m$. Therefore, $U_J^T \Sigma U_K = 0$. In particular, $U_{\mathcal{S}}^T \Sigma U_Q = 0$. \square

A.3 The Moore-Penrose inverse and its properties

The Moore-Penrose inverse is the most known and used generalized inverse⁶. It is defined as follows:

For $A \in \mathbb{R}^{m \times n}$, the pseudo-inverse of A is defined as the matrix $A^+ \in \mathbb{R}^{n \times m}$ which satisfies the 4 following criteria known as the Moore-Penrose conditions:

1. $AA^+A = A$.
2. $A^+AA^+ = A^+$.
3. $(AA^+)^T = AA^+$.
4. $(A^+A)^T = A^+A$.

A^+ exists for any matrix A and is unique. We also have the following properties:

- (i) $A^+ = (A^T A)^+ A^T$.
- (ii) $\text{rk}(A) = \text{rk}(A^+) = \text{rk}(AA^+) = \text{rk}(A^+A)$.
- (iii) If the linear system $Ax = b$ has any solutions, they are all given by

$$x = A^+b + (I - A^+A)w$$

for arbitrary vector w . This is equivalent to

$$x = A^+b + u$$

for arbitrary $u \in \text{Ker}(A)$.

- (iv) $P_A := AA^+$ is the orthogonal projection onto the range of A , and is therefore symmetric ($P_A^T = P_A$) (follows from 3) and idempotent ($P_A^2 = P_A$) (follows from 1).
- (v) $I_n - A^+A$ is the orthogonal projector onto the kernel of A .

B Propositions and lemmas for first-order critical points

In this section, we prove all lemmas about first-order critical points. We start by stating some preliminary results.

B.1 Preliminaries

The following lemma gives a necessary condition for \mathbf{W} to be a first-order critical point. It also provides the global map of the network, defined by $W_H \cdots W_1$. Finally, it states that the projection matrix P_K and Σ commute, where $K = W_H \cdots W_2$. This is key in the rest of the analysis.

Lemma 8. *Suppose Assumption \mathcal{H} in Section 2 holds true. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L . We define $K = W_H \cdots W_2$ and $W = W_H W_{H-1} \cdots W_1 = KW_1$. Then, we have*

$$W_1 = K^+ \Sigma_{YX} \Sigma_{XX}^{-1} + M,$$

where $M \in \mathbb{R}^{d_1 \times d_x}$ is such that $KM = 0$ and K^+ is the Moore-Penrose inverse of K (see Appendix A.3). As a consequence,

$$\begin{cases} W = P_K \Sigma_{YX} \Sigma_{XX}^{-1} \\ \text{rk}(W) = \text{rk}(P_K) = \text{rk}(K) \end{cases}$$

where we recall that $P_K = KK^+ \in \mathbb{R}^{d_y \times d_y}$ is the matrix of the orthogonal projection onto the range of K . Finally,

$$\Sigma P_K = P_K \Sigma.$$

⁶en.wikipedia.org/wiki/Moore-Penrose_inverse

Note that $\Sigma_{YX}\Sigma_{XX}^{-1}$ is the global minimizer of the problem with one layer (i.e the classical linear regression problem). Therefore, the global map $W_H \cdots W_1$ of any first-order critical point of L is equal to the global minimizer of the linear regression projected onto the column space of K .

Proof. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L . In particular, the partial gradients of L with respect to W_1 and W_H are equal to zero at \mathbf{W} . Using Lemma 3, this implies

$$\begin{cases} (W_H \cdots W_2)^T W_H \cdots W_1 \Sigma_{XX} = (W_H \cdots W_2)^T \Sigma_{YX} \\ W_H \cdots W_1 \Sigma_{XX} (W_{H-1} \cdots W_1)^T = \Sigma_{YX} (W_{H-1} \cdots W_1)^T . \end{cases}$$

We substitute in these equations $K = W_H W_{H-1} \cdots W_2$ and $W = W_H W_{H-1} \cdots W_1 = K W_1$. Using that Σ_{XX} is invertible, and multiplying the second equation on the right by W_H^T , we obtain that any critical point of L satisfies

$$\begin{cases} K^T K W_1 = K^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ W \Sigma_{XX} W^T = \Sigma_{YX} W^T . \end{cases} \quad (16)$$

The first equation implies $W_1 = (K^T K)^+ K^T \Sigma_{YX} \Sigma_{XX}^{-1} + M$, where $M \in \mathbb{R}^{d_1 \times d_x}$ is such that $K^T K M = 0$ (see Property (iii) in the reminder on Moore-Penrose inverse in Appendix A.3).

We have $(K^T K)^+ K^T = K^+$ (see Property (i) in Appendix A.3) and a standard fact of linear algebra is that $\text{Ker}(K^T K) = \text{Ker}(K)$.

Therefore, using these properties, we obtain $W_1 = K^+ \Sigma_{YX} \Sigma_{XX}^{-1} + M$, where $K M = 0$. This proves the first statement of the lemma. We then have,

$$W = K W_1 = K K^+ \Sigma_{YX} \Sigma_{XX}^{-1} + K M = P_K \Sigma_{YX} \Sigma_{XX}^{-1} . \quad (17)$$

where $P_K = K K^+$ is the orthogonal projection matrix onto the column space of K (see Appendix A.3). Using Assumption \mathcal{H} , we have that $\Sigma_{YX} \Sigma_{XX}^{-1}$ is of full row rank, hence

$$\text{rk}(W) = \text{rk}(P_K \Sigma_{YX} \Sigma_{XX}^{-1}) = \text{rk}(P_K) = \text{rk}(K) , \quad (18)$$

where the last equality comes from the property (ii) in Section A.3. Therefore, (16) and (18) prove the second statement of the lemma.

To prove that $\Sigma P_K = P_K \Sigma$, we remark that, using the second equation in (16), $\Sigma_{YX} W^T = W \Sigma_{XX} W^T$ and since $W \Sigma_{XX} W^T$ is symmetric and $(\Sigma_{YX})^T = \Sigma_{XY}$, we have

$$\Sigma_{YX} W^T = W \Sigma_{XY} .$$

Substituting the expression of W from (17), and since P_K and Σ_{XX}^{-1} are symmetric, we have

$$\Sigma_{YX} \Sigma_{XX}^{-1} \Sigma_{XY} P_K = P_K \Sigma_{YX} \Sigma_{XX}^{-1} \Sigma_{XY} .$$

Using the definition of Σ , this can be rewritten as

$$\Sigma P_K = P_K \Sigma ,$$

which concludes the proof. \square

Lemma 9. *Suppose Assumption \mathcal{H} in Section 2 holds true. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L . We set $K = W_H \cdots W_2$ and $r = \text{rk}(W_H \cdots W_1)$.*

There exists a unique subset $\mathcal{S} \subset \llbracket 1, d_y \rrbracket$ of size r such that:

$$P_K = U \mathcal{I}^{\mathcal{S}} U^T = U_{\mathcal{S}} U_{\mathcal{S}}^T ,$$

where $\mathcal{I}^{\mathcal{S}} \in \mathbb{R}^{d_y \times d_y}$ is the diagonal matrix such that, for all $i \in \llbracket 1, d_y \rrbracket$, $(\mathcal{I}^{\mathcal{S}})_{i,i} = 1$ if $i \in \mathcal{S}$ and 0 otherwise.

Proof. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L . Using Lemma 8, we have $\Sigma P_K = P_K \Sigma$. Substituting the diagonalization of Σ from Section 2, this becomes $U \Lambda U^T P_K = P_K U \Lambda U^T$. Since U is orthogonal, multiplying by U^T on the left and by U on the right we obtain $\Lambda U^T P_K U = U^T P_K U \Lambda$. Hence, $U^T P_K U$ commutes with a diagonal matrix whose diagonal elements are all distinct. Therefore, $\Gamma := U^T P_K U$ is diagonal, and $P_K = U \Gamma U^T$ is a diagonalization of P_K . From Lemma 8, we also have $r = \text{rk}(P_K)$. But, we know that $P_K = K K^+ \in \mathbb{R}^{d_y \times d_y}$ is the matrix of an orthogonal projection. Therefore, its eigenvalues are 1 with multiplicity r and 0 with multiplicity $d_y - r$.

Therefore, there exists an index set $\mathcal{S} \subset \llbracket 1, d_y \rrbracket$ of size r such that $\Gamma = \mathcal{I}^{\mathcal{S}}$ where $\mathcal{I}^{\mathcal{S}} \in \mathbb{R}^{d_y \times d_y}$ is the diagonal matrix such that, for all $i \in \llbracket 1, d_y \rrbracket$, $(\mathcal{I}^{\mathcal{S}})_{i,i} = 1$ if $i \in \mathcal{S}$ and 0 otherwise.

Therefore,

$$P_K = U \mathcal{I}^{\mathcal{S}} U^T = U \mathcal{I}^{\mathcal{S}} \mathcal{I}^{\mathcal{S}} U^T = U_{\mathcal{S}} U_{\mathcal{S}}^T.$$

If there exist \mathcal{S}' such that $\Gamma = \mathcal{I}^{\mathcal{S}'}$, we get $P_K = U \mathcal{I}^{\mathcal{S}} U^T = U \mathcal{I}^{\mathcal{S}'} U^T$ which implies $\mathcal{I}^{\mathcal{S}} = \mathcal{I}^{\mathcal{S}'}$, hence $\mathcal{S} = \mathcal{S}'$. Therefore, \mathcal{S} is unique. \square

B.2 Proof of Proposition 1

In this proof, we use Lemmas 8 and 9 stated and proved in the previous section.

Recall that $\lambda_1 > \dots > \lambda_{d_y}$ are the eigenvalues of $\Sigma = \Sigma_{YX} \Sigma_{XX}^{-1} \Sigma_{XY} \in \mathbb{R}^{d_y \times d_y}$.

Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L . We set $K = W_H \dots W_2$, $r = \text{rk}(W_H \dots W_1)$. Using Lemma 9, there exists a unique subset $\mathcal{S} \subset \llbracket 1, d_y \rrbracket$ of size r such that:

$$P_K = U_{\mathcal{S}} U_{\mathcal{S}}^T.$$

Therefore, using Lemma 8,

$$W_H \dots W_1 = P_K \Sigma_{YX} \Sigma_{XX}^{-1} = U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1}.$$

This proves the first statement of Proposition 1.

To prove the second statement, notice that we have

$$\begin{aligned} L(\mathbf{W}) &= \|WX - Y\|^2 \\ &= \|WX\|^2 - 2 \langle WX, Y \rangle + \|Y\|^2 \\ &= \text{tr}(W \Sigma_{XX} W^T) - 2 \text{tr}(W \Sigma_{XY}) + \text{tr}(\Sigma_{YY}) \\ &= \text{tr}(U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} \Sigma_{XX} \Sigma_{XX}^{-1} \Sigma_{XY} U_{\mathcal{S}} U_{\mathcal{S}}^T) - 2 \text{tr}(U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} \Sigma_{XY}) + \text{tr}(\Sigma_{YY}) \\ &= \text{tr}(U_{\mathcal{S}} U_{\mathcal{S}}^T U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma) - 2 \text{tr}(U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma) + \text{tr}(\Sigma_{YY}) \end{aligned}$$

Since $U_{\mathcal{S}}^T U_{\mathcal{S}} = I_r$ (see Lemma 6), using Lemma 9 and the fact that U diagonalizes Σ , this becomes

$$\begin{aligned} L(\mathbf{W}) &= \text{tr}(\Sigma_{YY}) - \text{tr}(U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma) \\ &= \text{tr}(\Sigma_{YY}) - \text{tr}(U \mathcal{I}^{\mathcal{S}} U^T U \Lambda U^T) \\ &= \text{tr}(\Sigma_{YY}) - \text{tr}(\mathcal{I}^{\mathcal{S}} U^T U \Lambda U^T U) \\ &= \text{tr}(\Sigma_{YY}) - \text{tr}(\mathcal{I}^{\mathcal{S}} \Lambda) \\ &= \text{tr}(\Sigma_{YY}) - \sum_{i \in \mathcal{S}} \lambda_i. \end{aligned}$$

This proves the second and last statement of Proposition 1.

B.3 Lemma 10

In this section we state and prove a lemma about first-order critical points which will be useful in various proofs. This lemma gives a simpler form for $K = W_H \cdots W_2$ and W_1 .

Lemma 10. *Suppose Assumption \mathcal{H} in Section 2 holds true. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L associated with \mathcal{S} . We set $r = \text{rk}(W_H \cdots W_1)$.*

Then there exists an invertible matrix $D \in \mathbb{R}^{d_1 \times d_1}$, a matrix $M \in \mathbb{R}^{d_1 \times d_x}$ satisfying $W_H \cdots W_2 M = 0$, such that:

$$K = W_H \cdots W_2 = \begin{bmatrix} U_{\mathcal{S}} & 0_{d_y \times (d_1 - r)} \end{bmatrix} D$$

and

$$W_1 = D^{-1} \begin{bmatrix} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ 0_{(d_1 - r) \times d_x} \end{bmatrix} + M.$$

Note that the result is still true when $r = 0$, provided that $U_{\emptyset} \in \mathbb{R}^{d_y \times 0}$.

To prove Lemma 10, we use Lemmas 8 and 9 stated and proved in the preliminaries of Appendix B.1. We will also need the following lemma

Lemma 11. *Let n be a positive integer and $\emptyset \neq \mathcal{S} \subset \llbracket 1, d_y \rrbracket$ such that $n \geq r := |\mathcal{S}|$. Let $A \in \mathbb{R}^{d_y \times n}$ such that $AA^+ = U_{\mathcal{S}} U_{\mathcal{S}}^T$. Then there exists an invertible matrix $D \in \mathbb{R}^{n \times n}$ such that*

$$A = [U_{\mathcal{S}} \quad 0_{d_y \times (n-r)}] D$$

and

$$A^+ = D^{-1} \begin{bmatrix} U_{\mathcal{S}}^T \\ 0_{(n-r) \times d_y} \end{bmatrix}.$$

Proof of Lemma 11. The matrix $I_n - A^+ A$ is the orthogonal projection onto $\text{Ker}(A)$ (see Appendix A.3), hence

$$\text{rk}(I_n - A^+ A) = \dim \text{Ker}(A) = n - \text{rk}(A)$$

But we have (see Property (ii) in Appendix A.3) $\text{rk}(A^+ A) = \text{rk}(A) = \text{rk}(AA^+)$ and, using Lemma 6, $\text{rk}(A^+ A) = \text{rk}(U_{\mathcal{S}} U_{\mathcal{S}}^T) = r$. Therefore, $\text{rk}(A) = r$ and

$$\text{rk}(I_n - A^+ A) = n - r.$$

Let $B \in \mathbb{R}^{n \times (n-r)}$ and $C \in \mathbb{R}^{(n-r) \times n}$ be such that $I_n - A^+ A = BC$ (such matrices can be obtained by considering the Singular Value Decomposition of $I_n - A^+ A$).

Denoting $D = \begin{bmatrix} U_{\mathcal{S}}^T A \\ C \end{bmatrix} \in \mathbb{R}^{n \times n}$, we have

$$[A^+ U_{\mathcal{S}}, B] D = [A^+ U_{\mathcal{S}}, B] \begin{bmatrix} U_{\mathcal{S}}^T A \\ C \end{bmatrix} = A^+ U_{\mathcal{S}} U_{\mathcal{S}}^T A + BC = A^+ AA^+ A + I_n - A^+ A.$$

Using Criteria 1 in Appendix A.3 we obtain

$$[A^+ U_{\mathcal{S}}, B] D = A^+ A + I_n - A^+ A = I_n.$$

Therefore, D is invertible and $D^{-1} = [A^+ U_{\mathcal{S}}, B]$. We have

$$[U_{\mathcal{S}}, 0_{d_y \times (n-r)}] D = [U_{\mathcal{S}}, 0_{d_y \times (n-r)}] \begin{bmatrix} U_{\mathcal{S}}^T A \\ C \end{bmatrix} = U_{\mathcal{S}} U_{\mathcal{S}}^T A = AA^+ A = A,$$

where the last equality follows from Criteria 1 in Appendix A.3. This proves the first equality of Lemma 11. Finally,

$$D^{-1} \begin{bmatrix} U_{\mathcal{S}}^T \\ 0_{(n-r) \times d_y} \end{bmatrix} = [A^+ U_{\mathcal{S}}, B] \begin{bmatrix} U_{\mathcal{S}}^T \\ 0_{(n-r) \times d_y} \end{bmatrix} = A^+ U_{\mathcal{S}} U_{\mathcal{S}}^T = A^+ AA^+ = A^+,$$

where the last equality follows again from Criteria 2 in Appendix A.3. This concludes the proof of Lemma 11. \square

Now we prove Lemma 10.

Proof of Lemma 10. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L associated with \mathcal{S} and $r = \text{rk}(W_H \cdots W_1)$.

Using Lemma 8, we have $r = \text{rk}(W_H \cdots W_2)$.

If $r = 0$, the conclusion of Lemma 10 is trivial because of the convention $U_\emptyset \in \mathbb{R}^{d_y \times 0}$.

When $r \geq 1$, using Lemma 8 and Proposition 1, we have $W_H \cdots W_1 = P_K \Sigma_{YX} \Sigma_{XX}^{-1} = U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1}$. Since Σ_{YX} is of full row rank this implies $P_K = K K^+ = U_S U_S^T$. Therefore, we can apply Lemma 11 with $n = d_1$ and $A = K$ to conclude that there exists an invertible matrix $D \in \mathbb{R}^{d_1 \times d_1}$ such that

$$K = [U_S, 0_{d_y \times (d_1 - r)}] D$$

which is the form of K in Lemma 10. Moreover, Lemma 11 also guarantees that

$$K^+ = D^{-1} \begin{bmatrix} U_S^T \\ 0_{(d_1 - r) \times d_y} \end{bmatrix}.$$

Using Lemma 8, we have $W_1 = K^+ \Sigma_{YX} \Sigma_{XX}^{-1} + M$ with $K M = 0$. Therefore, $W_1 = D^{-1} \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ 0_{(d_1 - r) \times d_x} \end{bmatrix} + M$, with $K M = 0$. This concludes the proof of Lemma 10. \square

B.4 Proof of Lemma 2

For any $h \in \llbracket 1, H - 1 \rrbracket$ let $D_h \in \mathbb{R}^{d_h \times d_h}$ be an invertible matrix. We define $\widetilde{\mathbf{W}} = (\widetilde{W}_H, \dots, \widetilde{W}_1)$ by $\widetilde{W}_H = W_H D_{H-1}$, $\widetilde{W}_1 = D_1^{-1} W_1$ and $\widetilde{W}_h = D_h^{-1} W_h D_{h-1}$ for all $h \in \llbracket 2, H - 1 \rrbracket$.

Assume that $\mathbf{W} = (W_H, \dots, W_1)$ is a first-order critical point. Then using Lemma 3 this is equivalent to

$$\begin{cases} \nabla_{W_h} L(\mathbf{W}) = 2(W_H \cdots W_{h+1})^T (W_H \cdots W_1 \Sigma_{XX} - \Sigma_{YX}) (W_{h-1} \cdots W_1)^T = 0 & \forall h \in \llbracket 2, H - 1 \rrbracket \\ \nabla_{W_H} L(\mathbf{W}) = 2(W_H \cdots W_1 \Sigma_{XX} - \Sigma_{YX}) (W_{H-1} \cdots W_1)^T = 0 \\ \nabla_{W_1} L(\mathbf{W}) = 2(W_H \cdots W_2)^T (W_H \cdots W_1 \Sigma_{XX} - \Sigma_{YX}) = 0. \end{cases} \quad (19)$$

Using the definition of $\widetilde{\mathbf{W}}$ above, we have

$$\begin{cases} W_H \cdots W_1 = \widetilde{W}_H \cdots \widetilde{W}_1 \\ W_H \cdots W_{h+1} = \widetilde{W}_H \cdots \widetilde{W}_{h+1} D_h^{-1} & \forall h \in \llbracket 1, H - 1 \rrbracket \\ W_{h-1} \cdots W_1 = D_{h-1} \widetilde{W}_{h-1} \cdots \widetilde{W}_1 & \forall h \in \llbracket 2, H \rrbracket. \end{cases}$$

Therefore (19) is equivalent to

$$\begin{cases} (D_h^{-1})^T (\widetilde{W}_H \cdots \widetilde{W}_{h+1})^T (\widetilde{W}_H \cdots \widetilde{W}_1 \Sigma_{XX} - \Sigma_{YX}) (\widetilde{W}_{h-1} \cdots \widetilde{W}_1)^T D_{h-1}^T = 0 & \forall h \in \llbracket 2, H - 1 \rrbracket \\ (\widetilde{W}_H \cdots \widetilde{W}_1 \Sigma_{XX} - \Sigma_{YX}) (\widetilde{W}_{H-1} \cdots \widetilde{W}_1)^T D_{H-1}^T = 0 \\ (D_1^{-1})^T (\widetilde{W}_H \cdots \widetilde{W}_2)^T (\widetilde{W}_H \cdots \widetilde{W}_1 \Sigma_{XX} - \Sigma_{YX}) = 0. \end{cases}$$

This is equivalent to

$$\begin{cases} \nabla_{W_h} L(\widetilde{\mathbf{W}}) = 2(\widetilde{W}_H \cdots \widetilde{W}_{h+1})^T (\widetilde{W}_H \cdots \widetilde{W}_1 \Sigma_{XX} - \Sigma_{YX}) (\widetilde{W}_{h-1} \cdots \widetilde{W}_1)^T = 0 & \forall h \in \llbracket 2, H - 1 \rrbracket \\ \nabla_{W_H} L(\widetilde{\mathbf{W}}) = 2(\widetilde{W}_H \cdots \widetilde{W}_1 \Sigma_{XX} - \Sigma_{YX}) (\widetilde{W}_{H-1} \cdots \widetilde{W}_1)^T = 0 \\ \nabla_{W_1} L(\widetilde{\mathbf{W}}) = 2(\widetilde{W}_H \cdots \widetilde{W}_2)^T (\widetilde{W}_H \cdots \widetilde{W}_1 \Sigma_{XX} - \Sigma_{YX}) = 0. \end{cases}$$

which is equivalent to $\nabla_{W_h} L(\widetilde{\mathbf{W}}) = 0$, for all $h \in \llbracket 1, H \rrbracket$. Therefore, \mathbf{W} is a first-order critical point if and only if $\widetilde{\mathbf{W}}$ is a first-order critical point. This proves the first part of the proposition.

Now assume that $\mathbf{W} = (W_H, \dots, W_1)$ is a first-order critical point such that it is not a second-order critical point. Note that from the first part of the proof $\tilde{\mathbf{W}} = (\tilde{W}_H, \dots, \tilde{W}_1)$ is also a first-order critical point. Let us prove that $\tilde{\mathbf{W}}$ is not a second-order critical point. Using Lemma 1, since \mathbf{W} is not a second-order critical point, there exist $\mathbf{W}' = (W'_H, \dots, W'_1)$ such that, if we denote $\mathbf{W}(t) = \mathbf{W} + t\mathbf{W}'$, the second-order term of $L(\mathbf{W}(t))$ is strictly negative i.e $c_2(\mathbf{W}, \mathbf{W}') < 0$. We will prove that there exist $\tilde{\mathbf{W}}'$ such that $c_2(\tilde{\mathbf{W}}, \tilde{\mathbf{W}}') < 0$ and, using again Lemma 1, we conclude.

As already said, we set $W_h(t) = W_h + tW'_h$, for all $h \in \llbracket 1, H \rrbracket$. We denote

$$\begin{cases} \tilde{W}_H(t) = \tilde{W}_H + t\tilde{W}'_H = \tilde{W}_H + tW'_H D_{H-1} \\ \tilde{W}_1(t) = \tilde{W}_1 + t\tilde{W}'_1 = \tilde{W}_1 + tD_1^{-1}W'_1 \\ \tilde{W}_h(t) = \tilde{W}_h + t\tilde{W}'_h = \tilde{W}_h + tD_h^{-1}W'_h D_{h-1} \quad \forall h \in \llbracket 2, H-1 \rrbracket \\ \tilde{\mathbf{W}}' = (\tilde{W}'_H, \dots, \tilde{W}'_1) . \end{cases}$$

Hence, we have (where $\prod_{h=H-1}^2 A_h$ should read as $A_{H-1} \cdots A_2$)

$$\begin{aligned} \tilde{W}_H(t) \cdots \tilde{W}_1(t) &= (W_H D_{H-1} + tW'_H D_{H-1}) \left(\prod_{h=H-1}^2 (D_h^{-1}W_h D_{h-1} + tD_h^{-1}W'_h D_{h-1}) \right) (D_1^{-1}W_1 + tD_1^{-1}W'_1) \\ &= (W_H + tW'_H) \cdots (W_1 + tW'_1) \\ &= W_H(t) \cdots W_1(t) . \end{aligned}$$

Therefore, $L(\tilde{\mathbf{W}}(t)) = L(\mathbf{W}(t))$ and

$$c_2(\tilde{\mathbf{W}}, \tilde{\mathbf{W}}') = c_2(\mathbf{W}, \mathbf{W}') .$$

Since by hypothesis $c_2(\mathbf{W}, \mathbf{W}') < 0$, we conclude that $c_2(\tilde{\mathbf{W}}, \tilde{\mathbf{W}}') < 0$. Hence $(\tilde{W}_H, \dots, \tilde{W}_1)$ is not a second-order critical point.

We prove that if $\tilde{\mathbf{W}}$ is not a second-order critical point then \mathbf{W} is not a second-order critical point in the same way, by changing D_h with D_h^{-1} for all $h \in \llbracket 1, H \rrbracket$. This proves the second part of the proposition and concludes the proof.

B.5 Proof of Proposition 6

Let $\mathcal{S} \subset \llbracket 1, d_y \rrbracket$ of size $r \in \llbracket 0, r_{max} \rrbracket$ and $Q = \llbracket 1, d_y \rrbracket \setminus \mathcal{S}$. Let $Z_H \in \mathbb{R}^{(d_y-r) \times (d_{H-1}-r)}$, $Z_1 \in \mathbb{R}^{(d_1-r) \times d_x}$ and $Z_h \in \mathbb{R}^{(d_h-r) \times (d_{h-1}-r)}$ for $h \in \llbracket 2, H-1 \rrbracket$. Let the parameter of the network $\mathbf{W} = (W_H, \dots, W_1)$ be defined as follows:

$$\begin{cases} W_H = [U_{\mathcal{S}}, U_Q Z_H] \\ W_1 = \begin{bmatrix} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_1 \end{bmatrix} \\ W_h = \begin{bmatrix} I_r & 0 \\ 0 & Z_h \end{bmatrix} \quad \forall h \in \llbracket 2, H-1 \rrbracket . \end{cases} \quad (20)$$

Note that the above definition of \mathbf{W} does not involve the matrices $D_h \in \mathbb{R}^{d_h \times d_h}$. In fact, using Lemma 2, it suffices to prove that, when $r = r_{max}$ or there exist $h_1 \neq h_2$ such that $Z_{h_1} = 0$ and $Z_{h_2} = 0$, the \mathbf{W} defined above is a first-order critical point to conclude that Proposition 6 holds .

We have

$$\begin{aligned} W_H \cdots W_1 &= [U_{\mathcal{S}}, U_Q Z_H] \begin{bmatrix} I_r & 0 \\ 0 & Z_{H-1} \end{bmatrix} \cdots \begin{bmatrix} I_r & 0 \\ 0 & Z_2 \end{bmatrix} \begin{bmatrix} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_1 \end{bmatrix} \\ &= U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} + U_Q Z_H Z_{H-1} \cdots Z_2 Z_1 \end{aligned}$$

If there exists $h_1 \neq h_2$ such that $Z_{h_1} = 0$ and $Z_{h_2} = 0$, it immediately follows that $W_H \cdots W_1 = U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1}$. If $r = r_{max}$, then there exists $h \in \llbracket 0, H \rrbracket$ such that $r = d_h$.

- If $r = d_H = d_y$, then $U_Q \in \mathbb{R}^{d_y \times 0}$ and $Z_H \in \mathbb{R}^{0 \times (d_{H-1-r})}$, which, using conventions in Section 2, gives

$$U_Q Z_H = 0_{d_y \times (d_{H-1-r})}. \quad (21)$$

Therefore, $W_H \cdots W_1 = U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1}$.

- If $r = d_0 = d_x$, then, since $d_x \geq d_y$, we have $r = d_y$, which we have already treated in the previous item.
- If $r = d_h$ for some $h \in \llbracket 2, H-1 \rrbracket$, then $Z_{h+1} \in \mathbb{R}^{(d_{h+1-r}) \times 0}$ and $Z_h \in \mathbb{R}^{0 \times (d_{h-1-r})}$, which, using the conventions on Section 2, gives

$$Z_{h+1} Z_h = 0_{(d_{h+1-r}) \times (d_{h-1-r})}. \quad (22)$$

Therefore, $W_H \cdots W_1 = U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1}$.

- If $r = d_1$, then $Z_2 \in \mathbb{R}^{(d_2-r) \times 0}$ and $Z_1 \in \mathbb{R}^{0 \times d_x}$, which, using the conventions on Section 2, gives

$$Z_2 Z_1 = 0_{(d_2-r) \times d_x}. \quad (23)$$

Therefore, $W_H \cdots W_1 = U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1}$.

Note that these results still hold if there is more than one layer with the minimum width.

Therefore, in all cases, when $r = r_{max}$ or there exist $h_1 \neq h_2$ such that $Z_{h_1} = 0$ and $Z_{h_2} = 0$ we have,

$$W_H \cdots W_1 = U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1}. \quad (24)$$

Let us prove that the gradient of L at \mathbf{W} is equal to zero.

Recall that from Lemma 3 we have

$$\begin{aligned} \nabla_{W_h} L(\mathbf{W}) &= 2(W_H \cdots W_{h+1})^T (W_H \cdots W_1 \Sigma_{XX} - \Sigma_{YX}) (W_{h-1} \cdots W_1)^T \quad \forall h \in \llbracket 2, H-1 \rrbracket \\ \nabla_{W_H} L(\mathbf{W}) &= 2(W_H \cdots W_1 \Sigma_{XX} - \Sigma_{YX}) (W_{H-1} \cdots W_1)^T \\ \nabla_{W_1} L(\mathbf{W}) &= 2(W_H \cdots W_2)^T (W_H \cdots W_1 \Sigma_{XX} - \Sigma_{YX}). \end{aligned}$$

Using (24) and Lemma 6, we have

$$\begin{aligned} W_H \cdots W_1 \Sigma_{XX} - \Sigma_{YX} &= U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \Sigma_{XX} - \Sigma_{YX} \\ &= (U_S U_S^T - I_{d_y}) \Sigma_{YX} \\ &= -U_Q U_Q^T \Sigma_{YX}. \end{aligned}$$

Also, using (20), for all $h \in \llbracket 1, H-1 \rrbracket$,

$$\begin{aligned} W_H \cdots W_{h+1} &= [U_S, U_Q Z_H] \begin{bmatrix} I_r & 0 \\ 0 & Z_{H-1} \end{bmatrix} \cdots \begin{bmatrix} I_r & 0 \\ 0 & Z_{h+1} \end{bmatrix} \\ &= [U_S, U_Q Z_H Z_{H-1} \cdots Z_{h+1}] \end{aligned}$$

and, for all $h \in \llbracket 2, H \rrbracket$,

$$\begin{aligned} W_{h-1} \cdots W_1 &= \begin{bmatrix} I_r & 0 \\ 0 & Z_{h-1} \end{bmatrix} \cdots \begin{bmatrix} I_r & 0 \\ 0 & Z_2 \end{bmatrix} \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_1 \end{bmatrix} \\ &= \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_{h-1} \cdots Z_2 Z_1 \end{bmatrix}. \end{aligned}$$

We have, for all $h \in \llbracket 2, H-1 \rrbracket$,

$$\begin{aligned}
\frac{1}{2} (\nabla_{W_h} L(\mathbf{W}))^T &= (W_{h-1} \cdots W_1)(W_H \cdots W_1 \Sigma_{XX} - \Sigma_{YX})^T (W_H \cdots W_{h+1}) \\
&= - \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_{h-1} \cdots Z_2 Z_1 \end{bmatrix} (U_Q U_Q^T \Sigma_{YX})^T [U_S, U_Q Z_H Z_{H-1} \cdots Z_{h+1}] \\
&= - \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_{h-1} \cdots Z_2 Z_1 \end{bmatrix} \Sigma_{XY} U_Q U_Q^T [U_S, U_Q Z_H Z_{H-1} \cdots Z_{h+1}] \\
&= - \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \Sigma_{XY} U_Q \\ Z_{h-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q \end{bmatrix} [U_Q^T U_S, U_Q^T U_Q Z_H Z_{H-1} \cdots Z_{h+1}].
\end{aligned}$$

Using the definition of Σ , Lemma 6 and Lemma 7, we have

$$\begin{aligned}
\frac{1}{2} (\nabla_{W_h} L(\mathbf{W}))^T &= - \begin{bmatrix} U_S^T \Sigma U_Q \\ Z_{h-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q \end{bmatrix} [0_{(d_y-r) \times r}, Z_H Z_{H-1} \cdots Z_{h+1}] \\
&= - \begin{bmatrix} 0_{r \times (d_y-r)} \\ Z_{h-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q \end{bmatrix} [0_{(d_y-r) \times r}, Z_H Z_{H-1} \cdots Z_{h+1}] \\
&= - \begin{bmatrix} 0_{r \times r} & 0_{r \times (d_h-r)} \\ 0_{(d_{h-1}-r) \times r} & Z_{h-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q Z_H Z_{H-1} \cdots Z_{h+1} \end{bmatrix}.
\end{aligned}$$

Proceeding similarly, we obtain

$$\frac{1}{2} (\nabla_{W_H} L(\mathbf{W}))^T = - \begin{bmatrix} 0_{r \times d_y} \\ Z_{H-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q U_Q^T \end{bmatrix}$$

and

$$\frac{1}{2} (\nabla_{W_1} L(\mathbf{W}))^T = - [0_{d_x \times r}, \Sigma_{XY} U_Q Z_H Z_{H-1} \cdots Z_2].$$

If there exists $h_1 \neq h_2$ such that $Z_{h_1} = 0$ and $Z_{h_2} = 0$, we can easily see that the gradient is equal to zero, i.e., \mathbf{W} is a first-order critical point.

If $r = r_{max}$, then there exists $h' \in \llbracket 1, H \rrbracket$ such that $r = d_{h'}$. Using the same arguments as above that yielded (21), (22) and (23), we have,

- For $h = 1$,
 - if $r = d_1$, we have $Z_2 \in \mathbb{R}^{(d_2-r) \times 0}$ and therefore $\Sigma_{XY} U_Q Z_H Z_{H-1} \cdots Z_2 \in \mathbb{R}^{d_x \times 0}$.
 - if $r = d_H$, then $U_Q Z_H = 0_{d_y \times (d_{H-1}-r)}$ and therefore $\Sigma_{XY} U_Q Z_H Z_{H-1} \cdots Z_2 = 0_{d_x \times (d_1-r)}$.
 - if $r = d_{h'}$ for some $h' \in \llbracket 2, H-1 \rrbracket$, then $Z_{h'+1} Z_{h'} = 0_{(d_{h'+1}-r) \times (d_{h'-1}-r)}$ and therefore $\Sigma_{XY} U_Q Z_H Z_{H-1} \cdots Z_2 = 0_{d_x \times (d_1-r)}$.

Hence, in all cases, $\nabla_{W_1} L(\mathbf{W}) = 0$.

- For $h = H$,
 - if $r = d_H = d_y$, then $U_Q U_Q^T = 0_{d_y \times d_y}$ and therefore $Z_{H-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q U_Q^T = 0_{(d_{H-1}-r) \times d_y}$.
 - if $r = d_{H-1}$, then $Z_{H-1} \in \mathbb{R}^{0 \times (d_{H-2}-r)}$ and therefore $Z_{H-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q U_Q^T \in \mathbb{R}^{0 \times d_y}$.
 - if $r = d_{h'}$ for some $h' \in \llbracket 2, H-2 \rrbracket$, then $Z_{h'+1} Z_{h'} = 0_{(d_{h'+1}-r) \times (d_{h'-1}-r)}$ and therefore $Z_{H-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q U_Q^T = 0_{(d_{H-1}-r) \times d_y}$.
 - if $r = d_1$, then $Z_2 Z_1 = 0_{(d_2-r) \times d_x}$ and therefore $Z_{H-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q U_Q^T = 0_{(d_{H-1}-r) \times d_y}$.

Hence, in all cases, $\nabla_{W_H} L(\mathbf{W}) = 0$.

- For $h \in \llbracket 2, H-1 \rrbracket$,
 - if $r = d_{h-1}$, then $Z_{h-1} \in \mathbb{R}^{0 \times (d_{h-2-r})}$ and therefore $Z_{h-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q Z_H Z_{H-1} \cdots Z_{h+1} \in \mathbb{R}^{0 \times (d_h-r)}$.
 - if $r = d_h$, then $Z_{h+1} \in \mathbb{R}^{(d_{h+1-r}) \times 0}$ and therefore $Z_{h-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q Z_H Z_{H-1} \cdots Z_{h+1} \in \mathbb{R}^{(d_{h-1-r}) \times 0}$.
 - if $r = d_H$, then $U_Q Z_H = 0_{d_y \times (d_{H-1-r})}$ and therefore $Z_{h-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q Z_H Z_{H-1} \cdots Z_{h+1} = 0_{(d_{h-1-r}) \times (d_h-r)}$.
 - if $r = d_1$, then $Z_2 Z_1 = 0_{(d_2-r) \times d_x}$ and therefore $Z_{h-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q Z_H Z_{H-1} \cdots Z_{h+1} = 0_{(d_{h-1-r}) \times (d_h-r)}$.
 - if $r = d_{h'}$ for some $h' \in \llbracket 2, H-1 \rrbracket \setminus \{h, h-1\}$, then $Z_{h'+1} Z_{h'} = 0_{(d_{h'+1-r}) \times (d_{h'-1-r})}$ and therefore $Z_{h-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q Z_H Z_{H-1} \cdots Z_{h+1} = 0_{(d_{h-1-r}) \times (d_h-r)}$.

Hence, in all cases, $\nabla_{W_h} L(\mathbf{W}) = 0$.

Therefore, when $r = r_{max}$, \mathbf{W} is also a first-order critical point of L .

B.6 Proof of Proposition 2

Let $\mathcal{S} \subset \llbracket 1, d_y \rrbracket$ such that $|\mathcal{S}| = r \leq r_{max}$, and $Q = \llbracket 1, d_y \rrbracket \setminus \mathcal{S}$.

We define $\mathbf{W} = (W_H, \dots, W_1)$ by:

$$\begin{aligned} W_H &= [U_{\mathcal{S}}, 0_{d_y \times (d_{H-1-r})}] \\ W_h &= \begin{bmatrix} I_r & 0_{r \times (d_{h-1-r})} \\ 0_{(d_h-r) \times r} & 0_{(d_h-r) \times (d_{h-1-r})} \end{bmatrix} \quad \forall h \in \llbracket 2, H-1 \rrbracket \\ W_1 &= \begin{bmatrix} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ 0_{(d_1-r) \times d_x} \end{bmatrix}, \end{aligned}$$

By Proposition 6, \mathbf{W} is a first-order critical point of L . Moreover, we have $W_H \cdots W_1 = U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1}$. Therefore, \mathbf{W} is a first-order critical point associated with \mathcal{S} .

B.7 Proof of Proposition 3

Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point and $r = \text{rk}(W_H \cdots W_1)$, using Proposition 1 there exists a unique $\mathcal{S} \subset \llbracket 1, d_y \rrbracket$ of size r such that

$$W_H \cdots W_1 = U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1},$$

which implies

$$W_H \cdots W_1 \Sigma_{XY} = U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma.$$

Let $i, j \in \llbracket 1, H \rrbracket$ such that $i > j$. The complementary blocks are $W_{j-1} \cdots W_1 \Sigma_{XY} W_H \cdots W_{i+1}$ and $W_{i-1} \cdots W_{j+1}$. Using Lemma 4 and $U_{\mathcal{S}}^T U_{\mathcal{S}} = I_r$, we have, for the second complementary block,

$$\text{rk}(W_{i-1} \cdots W_{j+1}) \geq \text{rk}(W_H \cdots W_1 \Sigma_{XY}) = \text{rk}(U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma) \geq \text{rk}(U_{\mathcal{S}}^T (U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma) \Sigma^{-1} U_{\mathcal{S}}) = \text{rk}(I_r) = r.$$

For the first complementary block, using the same arguments, we have

$$\begin{aligned} \text{rk}(W_{j-1} \cdots W_1 \Sigma_{XY} W_H \cdots W_{i+1}) &\geq \text{rk}(W_H \cdots W_1 \Sigma_{XY} W_H \cdots W_1 \Sigma_{XY}) \\ &= \text{rk}(U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma) \\ &\geq \text{rk}(U_{\mathcal{S}}^T (U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma) \Sigma^{-1} U_{\mathcal{S}}) \\ &= \text{rk}(U_{\mathcal{S}}^T \Sigma U_{\mathcal{S}}). \end{aligned}$$

Recall that, from the diagonalization of Σ , we have $\Sigma U = U\Lambda$, hence, $\Sigma U_{\mathcal{S}} = U_{\mathcal{S}} \text{diag}((\lambda_s)_{s \in \mathcal{S}})$

$$\begin{aligned} \text{rk}(W_{j-1} \cdots W_1 \Sigma_{XY} W_H \cdots W_{i+1}) &\geq \text{rk}(U_{\mathcal{S}}^T U_{\mathcal{S}} \text{diag}((\lambda_s)_{s \in \mathcal{S}})) \\ &= \text{rk}(\text{diag}((\lambda_s)_{s \in \mathcal{S}})) \\ &= r. \end{aligned}$$

This concludes the proof.

B.8 Proof of Proposition 4

Let $H \geq 3$, $\mathcal{S} = \llbracket 1, r \rrbracket$ with $0 \leq r < r_{max}$. We define \mathbf{W} as follows:

$$\begin{cases} W_H = [U_{\mathcal{S}}, 0] \\ W_h = \begin{bmatrix} I_r & 0 \\ 0 & Z_h \end{bmatrix} \quad \text{for } h \in \llbracket 2, H-1 \rrbracket \\ W_1 = \begin{bmatrix} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ 0 \end{bmatrix}. \end{cases} \quad (25)$$

Using Proposition 6, \mathbf{W} is a first-order critical point associated with \mathcal{S} . Let us show that depending on the choice of $(Z_h)_{h=2..H-1}$, \mathbf{W} can be tightened or non-tightened.

Since $H \geq 3$, there exists $h \in \llbracket 2, H-1 \rrbracket$. If we choose Z_{H-1}, \dots, Z_2 such that $Z_{H-1} \cdots Z_2 \neq 0$ (e.g. when only the top left entry of each Z_h is nonzero, which is possible since $r < r_{max} = \min(d_H, \dots, d_0)$) then \mathbf{W} is non-tightened. Indeed, the pivot $(H, 1)$ is non-tightened because $\text{rk}(\Sigma_{XY}) = d_y > r$ and $\text{rk}(W_{H-1} \cdots W_2) = \text{rk}\left(\begin{bmatrix} I_r & 0 \\ 0 & Z_{H-1} \cdots Z_2 \end{bmatrix}\right) > r$.

If we choose Z_{H-1}, \dots, Z_2 such that $Z_{H-1} \cdots Z_2 = 0$ (e.g. $Z_2 = 0$), then \mathbf{W} is tightened. Indeed, the pivot $(H, 1)$ is tightened because $W_{H-1} \cdots W_2 = \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix}$ is of rank r , and by construction we have $\text{rk}(W_H) = \text{rk}(W_1) = r$. Hence, all the other pivots are tightened because at least one of their complementary blocks includes W_H or W_1 , and therefore, using Proposition 3, is of rank r . Therefore, \mathbf{W} is tightened.

C Parameterization of first-order critical points and global minimizers

In this section, we prove Propositions 5 and 7 that were stated in Section 3.3.

C.1 Proof of Proposition 5

Before proving Proposition 5, we introduce and prove two lemmas.

Lemma 12. *Let r be a nonnegative integer, and let n and p be two positive integers larger than or equal to r . Let $\mathcal{S} \subset \llbracket 1, d_y \rrbracket$ of size r and let $Q = \llbracket 1, d_y \rrbracket \setminus \mathcal{S}$. Let $A \in \mathbb{R}^{d_y \times n}$ and $B \in \mathbb{R}^{n \times p}$ be two matrices such that*

$$AB = [U_{\mathcal{S}}, 0].$$

Then, there exist an invertible matrix $D \in \mathbb{R}^{n \times n}$ and two matrices $N \in \mathbb{R}^{(d_y-r) \times (n-r)}$ and $B_{DR} \in \mathbb{R}^{(n-r) \times (p-r)}$ such that

$$AD = [U_{\mathcal{S}}, U_Q N] \quad (26)$$

$$D^{-1}B = \begin{bmatrix} I_r & 0 \\ 0 & B_{DR} \end{bmatrix}. \quad (27)$$

In the proof below, we can easily see that the result still holds for $r = 0$ and $r = \min(d_y, n, p)$ with the conventions adopted in Section 2.

Proof. Let n and p be non-negative integers such that $n, p \geq r$ and $A \in \mathbb{R}^{d_y \times n}$ and $B \in \mathbb{R}^{n \times p}$ such that

$$AB = [U_S, 0]. \quad (28)$$

Recall that for any matrix C with n columns we write $C = [C_1, C_2, \dots, C_n]$ where C_i represents the i -th column of C .

We have from (28),

$$A[B_1, B_2, \dots, B_r] = U_S. \quad (29)$$

Since the columns of U are linearly independent, we have

$$\text{rk}(A[B_1, B_2, \dots, B_r]) = \text{rk}(U_S) = r$$

and $\{B_1, \dots, B_r\}$ are necessarily also linearly independent. Using the incomplete basis theorem, we complement (B_1, \dots, B_r) to form a basis $(B_1, \dots, B_r, E_{r+1}, \dots, E_n)$. We set $E = [B_1, \dots, B_r, E_{r+1}, \dots, E_n] \in \mathbb{R}^{n \times n}$. By construction, the matrix E is invertible.

We now set $A' = AE$ and $B' = E^{-1}B$. In particular $A'B' = AB$. Also, note that

$$E \begin{bmatrix} I_r \\ 0 \end{bmatrix} = [B_1, \dots, B_r],$$

so that

$$E^{-1}[B_1, \dots, B_r] = \begin{bmatrix} I_r \\ 0 \end{bmatrix}.$$

Therefore, we can write

$$B' = E^{-1}B = \begin{bmatrix} I_r & B_{UR} \\ 0 & B_{DR} \end{bmatrix}, \quad (30)$$

with $B_{UR} \in \mathbb{R}^{r \times (p-r)}$ and $B_{DR} \in \mathbb{R}^{(n-r) \times (p-r)}$ such that

$$\begin{bmatrix} B_{UR} \\ B_{DR} \end{bmatrix} = E^{-1}[B_{r+1}, \dots, B_p].$$

We define $L \in \mathbb{R}^{r \times (n-r)}$ and $N \in \mathbb{R}^{(d_y-r) \times (n-r)}$ by $\begin{bmatrix} L \\ N \end{bmatrix} = [U_S, U_Q]^{-1}[AE_{r+1}, \dots, AE_n]$. We have

$$[AE_{r+1}, \dots, AE_n] = [U_S, U_Q] \begin{bmatrix} L \\ N \end{bmatrix} = U_S L + U_Q N. \quad (31)$$

We also define the invertible matrix $F = \begin{bmatrix} I_r & L \\ 0 & I_{n-r} \end{bmatrix} \in \mathbb{R}^{n \times n}$. Using (29) and (31) we have

$$\begin{aligned} A' &= AE \\ &= A[B_1, \dots, B_r, E_{r+1}, \dots, E_n] \\ &= [U_S, U_S L + U_Q N] \\ &= [U_S, U_Q N] \begin{bmatrix} I_r & L \\ 0 & I_{n-r} \end{bmatrix} \\ &= [U_S, U_Q N] F. \end{aligned}$$

Therefore, defining the invertible matrix $D = EF^{-1} \in \mathbb{R}^{n \times n}$, we finally have

$$AD = AEF^{-1} = [U_S, U_Q N]. \quad (32)$$

This proves (26).

We also have, using (30) and the definition of F

$$\begin{aligned} D^{-1}B &= FE^{-1}B \\ &= FB' \\ &= \begin{bmatrix} I_r & L \\ 0 & I_{n-r} \end{bmatrix} \begin{bmatrix} I_r & B_{UR} \\ 0 & B_{DR} \end{bmatrix} \\ &= \begin{bmatrix} I_r & B_{UR} + LB_{DR} \\ 0 & B_{DR} \end{bmatrix}. \end{aligned} \quad (33)$$

However, noticing that, since (28) holds,

$$(AD)(D^{-1}B) = AB = [U_S, 0],$$

and using (32) and (33) we obtain

$$[U_S, U_Q N] \begin{bmatrix} I_r & B_{UR} + LB_{DR} \\ 0 & B_{DR} \end{bmatrix} = [U_S, 0].$$

Therefore $U_S(B_{UR} + LB_{DR}) + U_Q NB_{DR} = 0$. Since $[U_S, U_Q]$ is invertible we get $B_{UR} + LB_{DR} = 0$ and $NB_{DR} = 0$.

Finally, (33) becomes

$$D^{-1}B = \begin{bmatrix} I_r & 0 \\ 0 & B_{DR} \end{bmatrix}.$$

This proves (27) and concludes the proof. \square

The second lemma states that if the product of two factors takes the format of (27), then up to the product by an invertible matrix, the two factors have the same format. In the proof of Proposition 5, we will use this property several times to establish (4).

Lemma 13. *Let r, q, n and p be positive integers such that $r \leq \min(q, n, p)$. Let $B \in \mathbb{R}^{q \times n}$, $C \in \mathbb{R}^{n \times p}$ and $P \in \mathbb{R}^{(q-r) \times (p-r)}$ such that*

$$BC = \begin{bmatrix} I_r & 0 \\ 0 & P \end{bmatrix}.$$

Then, there exist an invertible matrix $D \in \mathbb{R}^{n \times n}$ and two matrices $B_{DR} \in \mathbb{R}^{(q-r) \times (n-r)}$ and $C_{DR} \in \mathbb{R}^{(n-r) \times (p-r)}$ such that

$$BD = \begin{bmatrix} I_r & 0 \\ 0 & B_{DR} \end{bmatrix} \quad (34)$$

$$D^{-1}C = \begin{bmatrix} I_r & 0 \\ 0 & C_{DR} \end{bmatrix}. \quad (35)$$

In the proof below, we can easily see that the result still holds for $r = 0$ and $r = \min(q, n, p)$ with the conventions adopted in Section 2.

Proof. Let r, q, n and p be positive integers such that $r \leq \min(q, n, p)$. Let $B \in \mathbb{R}^{q \times n}$, $C \in \mathbb{R}^{n \times p}$ and $P \in \mathbb{R}^{(q-r) \times (p-r)}$ such that

$$BC = \begin{bmatrix} I_r & 0 \\ 0 & P \end{bmatrix}. \quad (36)$$

We have

$$B[C_1, C_2, \dots, C_r] = \begin{bmatrix} I_r \\ 0 \end{bmatrix}. \quad (37)$$

Since the columns of $\begin{bmatrix} I_r \\ 0 \end{bmatrix}$ are linearly independent,

$$\text{rk}(B[C_1, C_2, \dots, C_r]) = r$$

and the vectors C_1, \dots, C_r are necessarily also linearly independent. Using the incomplete basis theorem, we complement (C_1, \dots, C_r) to form a basis $(C_1, \dots, C_r, E_{r+1}, \dots, E_n)$. We denote $E = [C_1, \dots, C_r, E_{r+1}, \dots, E_n] \in \mathbb{R}^{n \times n}$. By construction, the matrix E is invertible.

We now set $B' = BE$ and $C' = E^{-1}C$. In particular

$$B'C' = BC. \quad (38)$$

Also notice that

$$E \begin{bmatrix} I_r \\ 0 \end{bmatrix} = [C_1, \dots, C_r],$$

so that

$$E^{-1}[C_1, \dots, C_r] = \begin{bmatrix} I_r \\ 0 \end{bmatrix}.$$

Therefore, we can write

$$C' = E^{-1}C = \begin{bmatrix} I_r & C_{UR} \\ 0 & C_{DR} \end{bmatrix}, \quad (39)$$

where $C_{UR} \in \mathbb{R}^{r \times (p-r)}$ and $C_{DR} \in \mathbb{R}^{(n-r) \times (p-r)}$ are such that $\begin{bmatrix} C_{UR} \\ C_{DR} \end{bmatrix} = E^{-1}[C_{r+1}, \dots, C_p]$.

Now notice that, using (37),

$$\begin{aligned} B' &= BE \\ &= B[C_1, \dots, C_r, E_{r+1}, \dots, E_n] \\ &= \begin{bmatrix} I_r & B_{UR} \\ 0 & B_{DR} \end{bmatrix}, \end{aligned} \quad (40)$$

where $B_{UR} \in \mathbb{R}^{r \times (n-r)}$ and $B_{DR} \in \mathbb{R}^{(q-r) \times (n-r)}$ are such that $\begin{bmatrix} B_{UR} \\ B_{DR} \end{bmatrix} = B[E_{r+1}, \dots, E_n]$.

Plugging (40), (39) and (36) in the equality (38), we obtain

$$\begin{bmatrix} I_r & B_{UR} \\ 0 & B_{DR} \end{bmatrix} \begin{bmatrix} I_r & C_{UR} \\ 0 & C_{DR} \end{bmatrix} = \begin{bmatrix} I_r & 0 \\ 0 & P \end{bmatrix},$$

which yields

$$\begin{bmatrix} I_r & C_{UR} + B_{UR}C_{DR} \\ 0 & B_{DR}C_{DR} \end{bmatrix} = \begin{bmatrix} I_r & 0 \\ 0 & P \end{bmatrix}.$$

Therefore, $C_{UR} + B_{UR}C_{DR} = 0$ or, equivalently,

$$C_{UR} = -B_{UR}C_{DR}. \quad (41)$$

Define $F = \begin{bmatrix} I_r & -B_{UR} \\ 0 & I_{n-r} \end{bmatrix}$. The matrix F is invertible. Moreover, using (39) and (41) we have

$$\begin{aligned} C' &= \begin{bmatrix} I_r & -B_{UR}C_{DR} \\ 0 & C_{DR} \end{bmatrix} \\ &= \begin{bmatrix} I_r & -B_{UR} \\ 0 & I_{n-r} \end{bmatrix} \begin{bmatrix} I_r & 0 \\ 0 & C_{DR} \end{bmatrix} \\ &= F \begin{bmatrix} I_r & 0 \\ 0 & C_{DR} \end{bmatrix}. \end{aligned}$$

Therefore, if we define $D = EF$, D is invertible and

$$D^{-1}C = F^{-1}E^{-1}C = F^{-1}C' = \begin{bmatrix} I_r & 0 \\ 0 & C_{DR} \end{bmatrix}.$$

This proves (35).

In order to prove (34), we remark that, using (40) and the definition of F , we also have

$$\begin{aligned} BD &= BEF \\ &= B'F \\ &= \begin{bmatrix} I_r & B_{UR} \\ 0 & B_{DR} \end{bmatrix} \begin{bmatrix} I_r & -B_{UR} \\ 0 & I_{n-r} \end{bmatrix} \\ &= \begin{bmatrix} I_r & 0 \\ 0 & B_{DR} \end{bmatrix}. \end{aligned}$$

This proves (34) and concludes the proof. \square

Now we prove Proposition 5.

Proof of Proposition 5. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L . Then using Lemma 10 there exist $D \in \mathbb{R}^{d_1 \times d_1}$ invertible and a matrix $M \in \mathbb{R}^{d_1 \times d_x}$ which satisfies $W_H \cdots W_2 M = 0$ such that

$$W_H \cdots W_2 = [U_S, 0]D \quad (42)$$

$$W_1 = D^{-1} \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ 0 \end{bmatrix} + M. \quad (43)$$

Denoting $D_1 = D^{-1}$ and using (42), we have $W_H \cdots W_2 D_1 = [U_S, 0]$. Then applying Lemma 12 with $A = W_H$ and $B = W_{H-1} \cdots W_2 D_1$, there exist an invertible matrix $D_{H-1} \in \mathbb{R}^{d_{H-1} \times d_{H-1}}$, and matrices $Z_H \in \mathbb{R}^{(d_y-r) \times (d_{H-1}-r)}$ and $B_{DR} \in \mathbb{R}^{(d_{H-1}-r) \times (d_1-r)}$ such that

$$\begin{aligned} \widetilde{W}_H &:= W_H D_{H-1} = [U_S, U_Q Z_H] \\ D_{H-1}^{-1} W_{H-1} \cdots W_2 D_1 &= \begin{bmatrix} I_r & 0 \\ 0 & B_{DR} \end{bmatrix}. \end{aligned} \quad (44)$$

The first equality proves (2).

Then applying Lemma 13 to (44) with $B = D_{H-1}^{-1} W_{H-1}$ and $C = W_{H-2} \cdots W_2 D_1$ we get the existence of an invertible matrix $D_{H-2} \in \mathbb{R}^{d_{H-2} \times d_{H-2}}$, $C_{DR} \in \mathbb{R}^{(d_{H-2}-r) \times (d_1-r)}$ and $Z_{H-1} \in \mathbb{R}^{(d_{H-1}-r) \times (d_{H-2}-r)}$ such that

$$\widetilde{W}_{H-1} := D_{H-1}^{-1} W_{H-1} D_{H-2} = \begin{bmatrix} I_r & 0 \\ 0 & Z_{H-1} \end{bmatrix},$$

and

$$D_{H-2}^{-1} W_{H-2} \cdots W_2 D_1 = \begin{bmatrix} I_r & 0 \\ 0 & C_{DR} \end{bmatrix}.$$

Reiterating the process by using Lemma 13 multiple times with $B = D_h^{-1}W_h$ and $C = W_{h-1} \cdots W_2 D_1$ for h decreasing from $H - 2$ to 3, we can conclude that there exist invertible matrices $D_h \in \mathbb{R}^{d_h \times d_h}$ and matrices $Z_h \in \mathbb{R}^{(d_h-r) \times (d_{h-1}-r)}$, for $h \in \llbracket 2, H - 1 \rrbracket$, such that

$$\widetilde{W}_h := D_h^{-1}W_h D_{h-1} = \begin{bmatrix} I_r & 0 \\ 0 & Z_h \end{bmatrix} \quad \forall h \in \llbracket 2, H - 1 \rrbracket .$$

This entails (4).

We also have from (43) that $W_1 = D_1 \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ 0 \end{bmatrix} + M$ with $W_H \cdots W_2 M = 0$. Therefore,

$$D_1^{-1}W_1 = \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ 0 \end{bmatrix} + D_1^{-1}M .$$

Using (42), $D_1 = D^{-1}$ and $W_H \cdots W_2 M = 0$, we obtain

$$[U_S, 0] D_1^{-1}M = 0 .$$

Writing $D_1^{-1}M = \begin{bmatrix} Z_0 \\ Z_1 \end{bmatrix}$, where $Z_0 \in \mathbb{R}^{r \times d_x}$ and $Z_1 \in \mathbb{R}^{(d_1-r) \times d_x}$, we have

$$\begin{aligned} 0 &= [U_S, 0] D_1^{-1}M \\ &= [U_S, 0] \begin{bmatrix} Z_0 \\ Z_1 \end{bmatrix} \\ &= U_S Z_0 . \end{aligned}$$

Multiplying on the left by U_S^T we obtain

$$Z_0 = 0 .$$

Therefore $D_1^{-1}M = \begin{bmatrix} 0 \\ Z_1 \end{bmatrix}$, which yields

$$\widetilde{W}_1 := D_1^{-1}W_1 = \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_1 \end{bmatrix} .$$

This proves (3).

Finally we have

$$\begin{aligned} \widetilde{W}_H \cdots \widetilde{W}_2 &= (W_H D_{H-1}) (D_{H-1}^{-1} W_{H-1} D_{H-2}) \cdots (D_2^{-1} W_2 D_1) \\ &= W_H \cdots W_2 D_1 \\ &= [U_S, 0] , \end{aligned}$$

where the last equality is due to (42) and $D_1 = D^{-1}$. This entails (5) and concludes the proof. \square

C.2 Proof of Proposition 7

We first make a comment about notational subtleties to help understand the statement of Proposition 7, and then prove the proposition.

Recall that $r_{max} = \min(d_H, \dots, d_0)$, and $d_x = d_0 \geq d_y = d_H$ by assumption. Therefore, in the statement of Proposition 7, some blocks Z_h have 0 lines or 0 columns, and thus do not exist. For example, depending on the value of r_{max} , we have

$$\begin{cases} W_H = U_{S_{max}} D_{H-1}^{-1} & \text{if } r_{max} = d_{H-1} \\ W_1 = D_1 U_{S_{max}}^T \Sigma_{YX} \Sigma_{XX}^{-1} & \text{if } r_{max} = d_1 \end{cases}$$

and for $h \in \llbracket 2, H-1 \rrbracket$

$$W_h = \begin{cases} D_h \begin{bmatrix} I_{r_{max}} & 0 \end{bmatrix} D_{h-1}^{-1} & \text{if } r_{max} = d_h < d_{h-1} \\ D_h \begin{bmatrix} I_{r_{max}} \\ 0 \end{bmatrix} D_{h-1}^{-1} & \text{if } r_{max} = d_{h-1} < d_h \\ D_h I_{r_{max}} D_{h-1}^{-1} & \text{if } r_{max} = d_h = d_{h-1} \end{cases}$$

Also, if $r_{max} = d_y$, then $Q_{max} = \emptyset$, hence $U_{Q_{max}} \in \mathbb{R}^{d_y \times 0}$ and $Z_H \in \mathbb{R}^{0 \times (d_{H-1} - r_{max})}$. Then, using the convention in Section 2, $U_{Q_{max}} Z_H = 0_{d_y \times (d_{H-1} - r_{max})}$, so that $W_H = [U_{S_{max}}, 0_{d_y \times (d_{H-1} - r_{max})}] D_{H-1}^{-1} \in \mathbb{R}^{d_y \times d_{H-1}}$.

We are now ready to prove the proposition.

Proof of Proposition 7. Let $\mathcal{S}_{max} = \llbracket 1, r_{max} \rrbracket$. Let us first prove that \mathbf{W} is a global minimizer of L if and only if \mathbf{W} is a first-order critical point of L associated with \mathcal{S}_{max} . From Lemma 5, we have

$$U_{\mathcal{S}_{max}} U_{\mathcal{S}_{max}}^T \Sigma_{YX} \Sigma_{XX}^{-1} \in \underset{\substack{R \in \mathbb{R}^{d_y \times d_x} \\ \text{rk}(R) \leq r_{max}}}{\text{arg min}} \|RX - Y\|^2.$$

Let \mathbf{W} be a first-order critical point associated with \mathcal{S}_{max} (note that from Proposition 2, such \mathbf{W} exist). We have $W_H \cdots W_1 = U_{\mathcal{S}_{max}} U_{\mathcal{S}_{max}}^T \Sigma_{YX} \Sigma_{XX}^{-1}$, hence, for all $\mathbf{W}' = (W'_H, \dots, W'_1)$, since $\text{rk}(W'_H \cdots W'_1) \leq r_{max}$, we have

$$L(\mathbf{W}') \geq \min_{\substack{R \in \mathbb{R}^{d_y \times d_x} \\ \text{rk}(R) \leq r_{max}}} \|RX - Y\|^2 = \|W_H \cdots W_1 X - Y\|^2 = L(\mathbf{W}).$$

As a consequence, \mathbf{W} is a global minimizer of L .

Conversely, if \mathbf{W} is a global minimizer of L , then \mathbf{W} is a first-order critical point of L . From Proposition 1, there exist $\mathcal{S} \subset \llbracket 1, d_y \rrbracket$ of size $r \in \llbracket 0, r_{max} \rrbracket$ such that $W_H \cdots W_1 = U_{\mathcal{S}} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1}$, and we have $L(\mathbf{W}) = \text{tr}(\Sigma_{YY}) - \sum_{i \in \mathcal{S}} \lambda_i$. But we have from Assumption \mathcal{H} , $\lambda_1 > \dots > \lambda_{d_y}$, and, since Σ is invertible (see Lemma 4), then $\lambda_{d_y} > 0$. Therefore, using Proposition 2, \mathbf{W} is a global minimizer of L implies that $\mathcal{S} = \llbracket 1, r_{max} \rrbracket = \mathcal{S}_{max}$. Hence, \mathbf{W} is a global minimizer of L if and only if \mathbf{W} is a first-order critical point of L associated with \mathcal{S}_{max} .

Let us now prove Proposition 7.

Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point associated with $\mathcal{S}_{max} = \llbracket 1, r_{max} \rrbracket$. Using Proposition 5, there exist invertible matrices $D_{H-1} \in \mathbb{R}^{d_{H-1} \times d_{H-1}}, \dots, D_1 \in \mathbb{R}^{d_1 \times d_1}$, and matrices $Z_H \in \mathbb{R}^{(d_y - r_{max}) \times (d_{H-1} - r_{max})}$, $Z_h \in \mathbb{R}^{(d_h - r_{max}) \times (d_{h-1} - r_{max})}$ for $h \in \llbracket 2, H-1 \rrbracket$, and $Z_1 \in \mathbb{R}^{(d_1 - r_{max}) \times d_x}$ such that:

$$\begin{aligned} W_H &= [U_{\mathcal{S}_{max}}, U_{Q_{max}} Z_H] D_{H-1}^{-1} \\ W_1 &= D_1 \begin{bmatrix} U_{\mathcal{S}_{max}}^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_1 \end{bmatrix} \\ W_h &= D_h \begin{bmatrix} I_{r_{max}} & 0 \\ 0 & Z_h \end{bmatrix} D_{h-1}^{-1} \quad \forall h \in \llbracket 2, H-1 \rrbracket. \end{aligned}$$

Conversely, consider matrices D_h , for $h \in \llbracket 1, H-1 \rrbracket$ and Z_h , for $h \in \llbracket 1, H \rrbracket$ as in Proposition 7, and

$$\begin{aligned} W_H &= [U_{\mathcal{S}_{max}}, U_{Q_{max}} Z_H] D_{H-1}^{-1} \\ W_1 &= D_1 \begin{bmatrix} U_{\mathcal{S}_{max}}^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_1 \end{bmatrix} \\ W_h &= D_h \begin{bmatrix} I_{r_{max}} & 0 \\ 0 & Z_h \end{bmatrix} D_{h-1}^{-1} \quad \forall h \in \llbracket 2, H-1 \rrbracket. \end{aligned}$$

Since $|\mathcal{S}_{max}| = r_{max}$, using Proposition 6, we have that \mathbf{W} is a first-order critical point associated with \mathcal{S}_{max} . This concludes the proof. \square

D Global minimizers and simple strict saddle points (Proof of Proposition 8)

Recall that $r_{max} = \min(d_H, \dots, d_0)$.

Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L associated with \mathcal{S} of size $r = \text{rk}(W_H \cdots W_1) \leq r_{max}$.

Case 1: $\mathcal{S} = \llbracket 1, r_{max} \rrbracket = \mathcal{S}_{max}$. In this case, using Lemma 5,

$$W_H \cdots W_1 = U_{\mathcal{S}_{max}} U_{\mathcal{S}_{max}}^T \Sigma_{YX} \Sigma_{XX}^{-1} \in \arg \min_{\substack{R \in \mathbb{R}^{d_y \times d_x} \\ \text{rk}(R) \leq r_{max}}} \|RX - Y\|^2.$$

Moreover, for all $\mathbf{W}' = (W'_H, \dots, W'_1)$, since $\text{rk}(W'_H \cdots W'_1) \leq r_{max}$, we have

$$L(\mathbf{W}') \geq \min_{\substack{R \in \mathbb{R}^{d_y \times d_x} \\ \text{rk}(R) \leq r_{max}}} \|RX - Y\|^2 = \|W_H \cdots W_1 X - Y\|^2 = L(\mathbf{W}).$$

As a consequence, \mathbf{W} is a global minimizer of L .

Case 2: In order to prove the two remaining statements, we assume that $\mathcal{S} \neq \llbracket 1, r \rrbracket$ with $0 < r \leq r_{max}$, and show that \mathbf{W} is not a second-order critical point.

To do this we will find $\mathbf{W}' = (W'_H, \dots, W'_1)$ such that $c_2(\mathbf{W}, \mathbf{W}') < 0$ (see Lemma 1). More precisely, we find a linear trajectory of the form $W_h(t) = W_h + tW'_h$ such that the second-order coefficient of the asymptotic expansion of $L((W_h(t))_{h=1..H})$ around $t = 0$ is negative. This proves that \mathbf{W} is not a second-order critical point.

Since $\mathcal{S} \neq \llbracket 1, r \rrbracket$, and the eigenvalues $(\lambda_k)_{k \in \llbracket 1, d_y \rrbracket}$ are distinct and in decreasing order (see Section 2), there exist $j \in \mathcal{S}$ and $i \notin \mathcal{S}$ such that

$$\lambda_i > \lambda_j. \quad (45)$$

We denote by $\mathcal{S} = \{i_1, \dots, i_r\}$, hence there exists $g \in \llbracket 1, r \rrbracket$ such that $j = i_g$.

Note that,

$$U_{\mathcal{S}} = U \sum_{k=1}^r E_{i_k, k}$$

where $E_{l, k} \in \mathbb{R}^{d_y \times r}$ is the matrix whose entries are all 0 except the one in position (l, k) which is equal to 1.

Denote by U_t the matrix formed by replacing in $U_{\mathcal{S}}$ the column corresponding to u_j by $u_j + tu_i$. More precisely, set

$$U_t = U_{\mathcal{S}} + tU E_{i, g}.$$

Set $V = U E_{i, g} \in \mathbb{R}^{d_y \times r}$ and

$$V_t = \sum_{k=1}^r E_{i_k, k} + tE_{i, g} \in \mathbb{R}^{d_y \times r}. \quad (46)$$

Hence we have

$$U_t = U_{\mathcal{S}} + tV = UV_t. \quad (47)$$

Considering $D \in \mathbb{R}^{d_1 \times d_1}$ as provided by Lemma 10, we set

$$\begin{cases} W'_1 = D^{-1} \begin{bmatrix} V^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ 0_{(d_1-r) \times d_x} \end{bmatrix} \\ W'_h = 0 \quad \forall h \in \llbracket 2, H-1 \rrbracket \\ W'_H = V U_{\mathcal{S}}^T W_H. \end{cases}$$

and for all $h \in \llbracket 1, H \rrbracket$, $W_h(t) = W_h + tW'_h$. Note that

$$W_H(t) = W_H + tW'_H = (I_{d_y} + tVU_S^T)W_H ,$$

and therefore

$$K(t) := W_H(t) \cdots W_2(t) = (I_{d_y} + tVU_S^T)K ,$$

where $K = W_H \cdots W_2$. Using Lemma 10, there exists $M \in \mathbb{R}^{d_1 \times d_x}$ satisfying $KM = 0$ such that

$$W_1 = D^{-1} \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ 0_{(d_1-r) \times d_x} \end{bmatrix} + M .$$

Hence,

$$W_1(t) = D^{-1} \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ 0_{(d_1-r) \times d_x} \end{bmatrix} + M + tW'_1 = D^{-1} \begin{bmatrix} (U_S^T + tV^T) \Sigma_{YX} \Sigma_{XX}^{-1} \\ 0_{(d_1-r) \times d_x} \end{bmatrix} + M ,$$

where $M \in \mathbb{R}^{d_1 \times d_x}$ is such that $KM = 0$. Therefore

$$\begin{aligned} W_t &:= W_H(t) \cdots W_1(t) \\ &= K(t)W_1(t) \\ &= (I_{d_y} + tVU_S^T) \left(KD^{-1} \begin{bmatrix} (U_S^T + tV^T) \Sigma_{YX} \Sigma_{XX}^{-1} \\ 0_{(d_1-r) \times d_x} \end{bmatrix} + KM \right) . \end{aligned}$$

From Lemma 10, using that $KM = 0$ and $K = [U_S \quad 0_{d_y \times (d_1-r)}]D$, this becomes

$$\begin{aligned} W_t &= (I_{d_y} + tVU_S^T)[U_S \quad 0_{d_y \times (d_1-r)}]DD^{-1} \begin{bmatrix} (U_S^T + tV^T) \Sigma_{YX} \Sigma_{XX}^{-1} \\ 0_{(d_1-r) \times d_x} \end{bmatrix} \\ &= (I_{d_y} + tVU_S^T)U_S(U_S^T + tV^T) \Sigma_{YX} \Sigma_{XX}^{-1} . \end{aligned}$$

Using that $U_S^T U_S = I_r$ (see Lemma 6), we obtain

$$W_t = (U_S + tV)(U_S^T + tV^T) \Sigma_{YX} \Sigma_{XX}^{-1} = U_t U_t^T \Sigma_{YX} \Sigma_{XX}^{-1} . \quad (48)$$

Recall that our goal is to show that the asymptotic expansion of (49) around $t = 0$ has a negative second-order coefficient. We calculate

$$\begin{aligned} L((W_h(t))_{h=1..H}) &= \|W_t X - Y\|^2 \\ &= \text{tr}(W_t \Sigma_{XX} W_t^T) - 2 \text{tr}(W_t \Sigma_{XY}) + \text{tr}(\Sigma_{YY}) . \end{aligned} \quad (49)$$

Let us simplify $\text{tr}(W_t \Sigma_{XX} W_t^T)$ first. Using (48), we have

$$W_t \Sigma_{XX} W_t^T = U_t U_t^T \Sigma_{YX} \Sigma_{XX}^{-1} \Sigma_{XX} \Sigma_{XX}^{-1} \Sigma_{XY} U_t U_t^T = U_t U_t^T \Sigma U_t U_t^T .$$

Using (47), $U^T U = I_{d_y}$, $\Sigma = U \Lambda U^T$ and the cyclic property of the trace, we obtain

$$\text{tr}(W_t \Sigma_{XX} W_t^T) = \text{tr}(U V_t V_t^T U^T U \Lambda U^T U V_t V_t^T U^T) = \text{tr}(V_t V_t^T \Lambda V_t V_t^T) = \text{tr}\left((V_t V_t^T)^2 \Lambda\right) .$$

We define $(\bar{E}_{k,l})_{k=1..d_y, l=1..d_y}$ the canonical basis of $\mathbb{R}^{d_y \times d_y}$. More precisely, $\bar{E}_{k,l} \in \mathbb{R}^{d_y \times d_y}$ has all its entries equal to 0, except a 1 at position (k, l) . Note that for all $a, c \in \llbracket 1, d_y \rrbracket$ and $b, d \in \llbracket 1, r \rrbracket$

$$E_{a,b} E_{c,d}^T = \delta_{b,d} \bar{E}_{a,c} ,$$

where $\delta_{b,d}$ equals 1 if $b = d$ and 0 otherwise. Using the definition of V_t in (46) and $j = i_g$, for $g \in \llbracket 1, r \rrbracket$, we have

$$\begin{aligned}
V_t V_t^T &= \left(\sum_{k=1}^r E_{i_k, k} + t E_{i, g} \right) \left(\sum_{k'=1}^r E_{i_{k'}, k'}^T + t E_{i, g}^T \right) \\
&= \left(\sum_{k=1}^r \bar{E}_{i_k, i_k} \right) + t \bar{E}_{i_g, i} + t \bar{E}_{i, i_g} + t^2 \bar{E}_{i, i} \\
&= \left(\sum_{k \in \mathcal{S}} \bar{E}_{k, k} \right) + t \bar{E}_{j, i} + t \bar{E}_{i, j} + t^2 \bar{E}_{i, i}.
\end{aligned} \tag{50}$$

We also have for all $a, b, c, d \in \llbracket 1, d_y \rrbracket$

$$\bar{E}_{a, b} \bar{E}_{c, d} = \delta_{b, c} \bar{E}_{a, d}.$$

Recalling that $j \in \mathcal{S}$ and $i \notin \mathcal{S}$, we obtain

$$\begin{aligned}
(V_t V_t^T)^2 &= \left(\left(\sum_{k \in \mathcal{S}} \bar{E}_{k, k} \right) + t \bar{E}_{j, i} + t \bar{E}_{i, j} + t^2 \bar{E}_{i, i} \right) \left(\left(\sum_{k' \in \mathcal{S}} \bar{E}_{k', k'} \right) + t \bar{E}_{j, i} + t \bar{E}_{i, j} + t^2 \bar{E}_{i, i} \right) \\
&= \left(\left(\sum_{k \in \mathcal{S}} \bar{E}_{k, k} \right) + t \bar{E}_{j, i} + 0 + 0 \right) + (0 + 0 + t^2 \bar{E}_{j, j} + t^3 \bar{E}_{j, i}) \\
&\quad + (t \bar{E}_{i, j} + t^2 \bar{E}_{i, i} + 0 + 0) + (0 + 0 + t^3 \bar{E}_{i, j} + t^4 \bar{E}_{i, i}) \\
&= \left(\sum_{k \in \mathcal{S}} \bar{E}_{k, k} \right) + t^2 (1 + t^2) \bar{E}_{i, i} + t^2 \bar{E}_{j, j} + t (1 + t^2) \bar{E}_{i, j} + t (1 + t^2) \bar{E}_{j, i}.
\end{aligned}$$

Finally, since for all $a, b \in \llbracket 1, d_y \rrbracket$

$$\bar{E}_{a, b} \Lambda = \lambda_b \bar{E}_{a, b} \tag{51}$$

we have

$$\text{tr}(W_t \Sigma_{XX} W_t^T) = \text{tr}\left((V_t V_t^T)^2 \Lambda \right) = \sum_{k \in \mathcal{S}} \lambda_k + t^2 (1 + t^2) \lambda_i + t^2 \lambda_j. \tag{52}$$

Coming back to (49), we calculate the other term $\text{tr}(W_t \Sigma_{XY})$. Using (48), (47) and $\Sigma = U \Lambda U^T$, we obtain

$$\text{tr}(W_t \Sigma_{XY}) = \text{tr}(U_t U_t^T \Sigma) = \text{tr}(U V_t V_t^T U^T U \Lambda U^T) = \text{tr}(V_t V_t^T \Lambda).$$

Combining with (50) and (51), we get

$$\text{tr}(W_t \Sigma_{XY}) = \text{tr}(V_t V_t^T \Lambda) = \sum_{k \in \mathcal{S}} \lambda_k + t^2 \lambda_i. \tag{53}$$

Finally, substituting (52) and (53) in (49), we have

$$\begin{aligned}
L((W_h(t))_{h=1..H}) &= \text{tr}(\Sigma_{YY}) + \sum_{k \in \mathcal{S}} \lambda_k + t^2 (1 + t^2) \lambda_i + t^2 \lambda_j - 2 \sum_{k \in \mathcal{S}} \lambda_k - 2t^2 \lambda_i \\
&= \text{tr}(\Sigma_{YY}) - \sum_{k \in \mathcal{S}} \lambda_k + t^2 (\lambda_j - \lambda_i) + \lambda_i t^4.
\end{aligned}$$

Using Proposition 1 and recalling (45), we finally get as $t \rightarrow 0$,

$$L((W_h(t))_{h=1..H}) = L(\mathbf{W}) + ct^2 + o(t^2) \quad \text{with} \quad c = \lambda_j - \lambda_i < 0.$$

Therefore, we conclude from Lemma 1 that $\mathbf{W} = (W_H, \dots, W_1)$ is not a second-order critical point.

E Strict saddle points with $\mathcal{S} = \llbracket 1, r \rrbracket$, $r < r_{max}$ (Proof of Proposition 9)

We refer the reader to Section 4.2, which introduces the 4 cases proved below. Recall that $\mathcal{S} = \llbracket 1, r \rrbracket$ and we set $Q = \llbracket 1, d_y \rrbracket \setminus \mathcal{S} = \llbracket r+1, d_y \rrbracket$.

In this section, for each vector space \mathbb{R}^{d_h} , we will denote by e_m the m -th element of the canonical basis of \mathbb{R}^{d_h} . That is, the entries of $e_m \in \mathbb{R}^{d_h}$ are all equal to 0 except for the m -th coordinate which is equal to 1. The size of e_m will not be ambiguous, once in context, so we do not include it in the notation.

Remark about $r = 0$: Using the conventions of Section 2, in this case we have $\mathcal{S} = \emptyset$ and $Q = \llbracket 1, d_y \rrbracket$. Hence $U_{\mathcal{S}}$ is the matrix with no column, $U_Q = U$, and $U_{\mathcal{S}}U_{\mathcal{S}}^T = 0_{d_y \times d_y}$. For example, we still have $I_{d_y} = U_{\mathcal{S}}U_{\mathcal{S}}^T + U_QU_Q^T$. We can easily follow the proofs below with these conventions and see that the result still holds.

E.1 1st case: $i \in \llbracket 2, H-1 \rrbracket$ and $j = 1$

In this case, the two complementary blocks are $\Sigma_{XY}W_H \cdots W_{i+1}$ and $W_{i-1} \cdots W_2$. Recall that $\mathcal{S} = \llbracket 1, r \rrbracket$ and $r < r_{max} = \min(d_H, \dots, d_0)$. Note that $\text{rk}(\Sigma_{XY}W_H \cdots W_{i+1}) = \text{rk}(W_H \cdots W_{i+1})$ because Σ_{XY} is of full column rank (see Assumption \mathcal{H} , in Section 2).

Since the pivot (i, j) is not tightened, using Proposition 3, we have

$$\begin{cases} \text{rk}(W_H \cdots W_{i+1}) > r \\ \text{rk}(W_{i-1} \cdots W_2) > r. \end{cases} \quad (54)$$

Let us first show that there exists $k \in \llbracket r+1, d_y \rrbracket$ and $l \in \llbracket 1, d_i \rrbracket$ such that

$$U_k^T(W_H \cdots W_{i+1})_{.,l} \neq 0. \quad (55)$$

Indeed, assume by contradiction that for all $k \in \llbracket r+1, d_y \rrbracket$ and $l \in \llbracket 1, d_i \rrbracket$ we have

$$U_k^T(W_H \cdots W_{i+1})_{.,l} = 0.$$

Recalling that $Q = \llbracket 1, d_y \rrbracket \setminus \mathcal{S} = \llbracket r+1, d_y \rrbracket$, we obtain $U_Q^T W_H \cdots W_{i+1} = 0$. Using from Lemma 6 that $I_{d_y} = U_{\mathcal{S}}U_{\mathcal{S}}^T + U_QU_Q^T$, we have

$$\begin{aligned} W_H \cdots W_{i+1} &= (U_{\mathcal{S}}U_{\mathcal{S}}^T + U_QU_Q^T)W_H \cdots W_{i+1} \\ &= U_{\mathcal{S}}U_{\mathcal{S}}^T W_H \cdots W_{i+1}. \end{aligned}$$

Therefore,

$$\text{rk}(W_H \cdots W_{i+1}) = \text{rk}(U_{\mathcal{S}}U_{\mathcal{S}}^T W_H \cdots W_{i+1}).$$

The latter is impossible since $\text{rk}(U_{\mathcal{S}}U_{\mathcal{S}}^T W_H \cdots W_{i+1}) \leq |\mathcal{S}| = r$, which is not compatible with (54). Therefore (55) holds.

Since \mathbf{W} is a first-order critical point, using Lemma 10, there exists an invertible matrix $D \in \mathbb{R}^{d_1 \times d_1}$ such that

$$W_H \cdots W_2 = [U_{\mathcal{S}}, 0_{d_y \times (d_1-r)}]D \quad (56)$$

and since \mathbf{W} is associated with \mathcal{S} , we have

$$W_H \cdots W_1 = U_{\mathcal{S}}U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1}. \quad (57)$$

Using (54) and D invertible, we have $\text{rk}(W_{i-1} \cdots W_2 D^{-1}) = \text{rk}(W_{i-1} \cdots W_2) > r$. Hence there exists $g \in \llbracket r+1, d_1 \rrbracket$ such that

$$(W_{i-1} \cdots W_2 D^{-1})_{.,g} \neq 0.$$

Therefore, there exists $a \in \mathbb{R}^{d_i-1}$ such that

$$a^T (W_{i-1} \cdots W_2 D^{-1})_{.,g} = 1. \quad (58)$$

Recall that k, l satisfy (55). We define $\mathbf{W}'_\beta = (W'_H, \dots, W'_1)^\beta$ by

$$\begin{cases} W'_i{}^\beta = \beta W'_i = \beta e_l a^T \in \mathbb{R}^{d_i \times d_{i-1}}, \text{ where } e_l \in \mathbb{R}^{d_i} \\ W'_1{}^\beta = W'_1 = D^{-1} e_g U_k^T \Sigma_{YX} \Sigma_{XX}^{-1} \in \mathbb{R}^{d_1 \times d_x}, \text{ where } e_g \in \mathbb{R}^{d_1} \\ W'_h{}^\beta = 0 \quad \forall h \in \llbracket 2, H \rrbracket \setminus \{i\} \end{cases}$$

We set $\mathbf{W}^\beta(t) = (W_H^\beta(t), \dots, W_1^\beta(t))$ such that $W_h^\beta(t) = W_h + tW'_h{}^\beta$ for $h \in \llbracket 1, H \rrbracket$. We have

$$\begin{aligned} W^\beta(t) &:= W_H^\beta(t) \cdots W_1^\beta(t) \\ &= W_H \cdots W_{i+1} (W_i + t\beta W'_i) W_{i-1} \cdots W_2 (W_1 + tW'_1) \\ &= W_H \cdots W_1 + t(\beta W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_1 + W_H \cdots W_2 W'_1) \\ &\quad + \beta t^2 W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_2 W'_1. \end{aligned}$$

Using (56) and (57), we obtain

$$\begin{aligned} W^\beta(t) &= U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} + t(\beta W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_1 + [U_S, 0] D D^{-1} e_g U_k^T \Sigma_{YX} \Sigma_{XX}^{-1}) \\ &\quad + \beta t^2 (W_H \cdots W_{i+1})_{.,l} a^T (W_{i-1} \cdots W_2 D^{-1})_{.,g} U_k^T \Sigma_{YX} \Sigma_{XX}^{-1}. \end{aligned}$$

Using (58) and $g \in \llbracket r+1, d_1 \rrbracket$, we have

$$W^\beta(t) = U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} + t\beta W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_1 + \beta t^2 (W_H \cdots W_{i+1})_{.,l} U_k^T \Sigma_{YX} \Sigma_{XX}^{-1}.$$

Denoting $N = W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_1$, we have

$$\begin{aligned} L(\mathbf{W}^\beta(t)) &= \|W^\beta(t)X - Y\|^2 \\ &= \|U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X - Y + t\beta N X + \beta t^2 (W_H \cdots W_{i+1})_{.,l} U_k^T \Sigma_{YX} \Sigma_{XX}^{-1} X\|^2. \end{aligned}$$

Expanding the square, the second-order term $c_2(\mathbf{W}, \mathbf{W}'_\beta)t^2$ has a coefficient equal to

$$\begin{aligned} c_2(\mathbf{W}, \mathbf{W}'_\beta) &= \beta^2 \|NX\|^2 + 2\beta \operatorname{tr}((W_H \cdots W_{i+1})_{.,l} U_k^T \Sigma_{YX} \Sigma_{XX}^{-1} X X^T \Sigma_{XX}^{-1} \Sigma_{XY} U_S U_S^T) \\ &\quad - 2\beta \operatorname{tr}((W_H \cdots W_{i+1})_{.,l} U_k^T \Sigma_{YX} \Sigma_{XX}^{-1} X Y^T) \\ &= \beta^2 \|NX\|^2 + 2\beta \operatorname{tr}((W_H \cdots W_{i+1})_{.,l} U_k^T \Sigma U_S U_S^T) - 2\beta \operatorname{tr}((W_H \cdots W_{i+1})_{.,l} U_k^T \Sigma) \\ &= \beta^2 \|NX\|^2 - 2\beta \lambda_k U_k^T (W_H \cdots W_{i+1})_{.,l}, \end{aligned}$$

where the last equality follows from Lemma 7 and $k \notin \mathcal{S}$, and $U^T \Sigma = \Lambda U^T$ and the cyclic property of the trace. Using Lemma 4 and (55), we have $\lambda_k U_k^T (W_H \cdots W_{i+1})_{.,l} \neq 0$, hence we can choose β according to (7), such that $c_2(\mathbf{W}, \mathbf{W}'_\beta) < 0$. Therefore, \mathbf{W} is not a second-order critical point.

E.2 2nd case: $i = H$ and $j = 1$

In this case, the two complementary blocks are Σ_{XY} and $W_{H-1} \cdots W_2$. We follow again the same lines as above. Since the pivot (i, j) is not tightened, using Proposition 3, we have

$$\operatorname{rk}(W_{H-1} \cdots W_2) > r. \quad (59)$$

Again, since \mathbf{W} is a first-order critical point, using Lemma 10, there exists an invertible matrix $D \in \mathbb{R}^{d_1 \times d_1}$ such that

$$W_H \cdots W_2 = [U_S, 0_{d_y \times (d_1-r)}] D \quad (60)$$

and since \mathbf{W} is associated with \mathcal{S} , we have

$$W_H \cdots W_1 = U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1}. \quad (61)$$

Using (59) and D invertible, we have $\text{rk}(W_{H-1} \cdots W_2 D^{-1}) = \text{rk}(W_{H-1} \cdots W_2) > r$. Hence there exists $g \in \llbracket r+1, d_1 \rrbracket$ such that

$$(W_{i-1} \cdots W_2 D^{-1})_{.,g} \neq 0.$$

Therefore, there exists $a \in \mathbb{R}^{d_{H-1}}$ such that

$$a^T (W_{H-1} \cdots W_2 D^{-1})_{.,g} = 1. \quad (62)$$

We define $\mathbf{W}'_\beta = (W'_H{}^\beta, \dots, W'_1{}^\beta)$ by

$$\begin{cases} W'_H{}^\beta = \beta W'_H = \beta U_{r+1} a^T \in \mathbb{R}^{d_y \times d_{H-1}} \\ W'_1{}^\beta = W'_1 = D^{-1} e_g U_{r+1}^T \Sigma_{YX} \Sigma_{XX}^{-1} \in \mathbb{R}^{d_1 \times d_x}, \text{ where } e_g \in \mathbb{R}^{d_1} \\ W'_h{}^\beta = 0 \quad \forall h \in \llbracket 2, H-1 \rrbracket. \end{cases}$$

We set $\mathbf{W}^\beta(t) = (W_H^\beta(t), \dots, W_1^\beta(t))$ such that $W_h^\beta(t) = W_h + t W'_h{}^\beta$, for all $h \in \llbracket 1, H \rrbracket$. We have

$$\begin{aligned} W^\beta(t) &:= W_H^\beta(t) \cdots W_1^\beta(t) \\ &= (W_H + t \beta W'_H) W_{H-1} \cdots W_2 (W_1 + t W'_1) \\ &= W_H \cdots W_1 + t(\beta W'_H W_{H-1} \cdots W_1 + W_H \cdots W_2 W'_1) + \beta t^2 W'_H W_{H-1} \cdots W_2 W'_1. \end{aligned}$$

Using (60) and (61), then (62) and $g \in \llbracket r+1, d_1 \rrbracket$, we obtain

$$\begin{aligned} W^\beta(t) &= U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} + t(\beta W'_H W_{H-1} \cdots W_1 + [U_S, 0] D D^{-1} e_g U_{r+1}^T \Sigma_{YX} \Sigma_{XX}^{-1}) \\ &\quad + \beta t^2 U_{r+1} a^T (W_{H-1} \cdots W_2 D^{-1})_{.,g} U_{r+1}^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ &= U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} + t \beta W'_H W_{H-1} \cdots W_1 + \beta t^2 U_{r+1} U_{r+1}^T \Sigma_{YX} \Sigma_{XX}^{-1}. \end{aligned}$$

Denoting by $N = W'_H W_{H-1} \cdots W_1$, we have

$$\begin{aligned} L(\mathbf{W}^\beta(t)) &= \|W^\beta(t)X - Y\|^2 \\ &= \|U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X - Y + t \beta N X + \beta t^2 U_{r+1} U_{r+1}^T \Sigma_{YX} \Sigma_{XX}^{-1} X\|^2. \end{aligned}$$

As previously, expanding the square, we can see that the second-order coefficient $c_2(\mathbf{W}, \mathbf{W}'_\beta)$ of the polynomial $L(\mathbf{W}^\beta(t))$ is given by

$$\begin{aligned} c_2(\mathbf{W}, \mathbf{W}'_\beta) &= \beta^2 \|N X\|^2 + 2\beta \text{tr}(U_{r+1} U_{r+1}^T \Sigma_{YX} \Sigma_{XX}^{-1} X X^T \Sigma_{XX}^{-1} \Sigma_{XY} U_S U_S^T) - 2\beta \text{tr}(U_{r+1} U_{r+1}^T \Sigma_{YX} \Sigma_{XX}^{-1} X Y^T) \\ &= \beta^2 \|N X\|^2 + 2\beta \text{tr}(U_{r+1} U_{r+1}^T \Sigma U_S U_S^T) - 2\beta \text{tr}(U_{r+1} U_{r+1}^T \Sigma). \end{aligned}$$

Using the cyclic property of the trace, $U_S^T U_{r+1} = 0$ (see Lemma 6), and $\Sigma U_{r+1} = \lambda_{r+1} U_{r+1}$, we obtain

$$\begin{aligned} c_2(\mathbf{W}, \mathbf{W}'_\beta) &= \beta^2 \|N X\|^2 - 2\beta \lambda_{r+1} U_{r+1}^T U_{r+1} \\ &= \beta^2 \|N X\|^2 - 2\beta \lambda_{r+1}. \end{aligned}$$

Using Lemma 4, we have $\lambda_{r+1} \neq 0$, hence we can choose β according to (7) such that $c_2(\mathbf{W}, \mathbf{W}'_\beta) < 0$. Therefore \mathbf{W} is not a second-order critical point.

E.3 3rd case: $i = H$ and $j \in \llbracket 2, H - 1 \rrbracket$

In this case, the two complementary blocks are $W_{j-1} \cdots W_1 \Sigma_{XY}$ and $W_{H-1} \cdots W_{j+1}$. We follow again the same lines as above. Since the pivot (i, j) is not tightened, using Proposition 3, we have

$$\begin{cases} \text{rk}(W_{H-1} \cdots W_{j+1}) > r \\ \text{rk}(W_{j-1} \cdots W_1 \Sigma_{XY}) > r. \end{cases} \quad (63)$$

Let us first show that there exist $k \in \llbracket r + 1, d_y \rrbracket$ and $l \in \llbracket 1, d_{j-1} \rrbracket$ such that

$$(W_{j-1} \cdots W_1)_{l, \Sigma_{XY}} U_k \neq 0. \quad (64)$$

Indeed, assume by contradiction that for all $k \in \llbracket r + 1, d_y \rrbracket$ and $l \in \llbracket 1, d_{j-1} \rrbracket$ we have

$$(W_{j-1} \cdots W_1)_{l, \Sigma_{XY}} U_k = 0.$$

Recalling that $Q = \llbracket 1, d_y \rrbracket \setminus \mathcal{S} = \llbracket r + 1, d_y \rrbracket$, we obtain $W_{j-1} \cdots W_1 \Sigma_{XY} U_Q = 0$, and using, from Lemma 6, that $I_{d_y} = U_S U_S^T + U_Q U_Q^T$, we have

$$\begin{aligned} W_{j-1} \cdots W_1 \Sigma_{XY} &= W_{j-1} \cdots W_1 \Sigma_{XY} (U_S U_S^T + U_Q U_Q^T) \\ &= W_{j-1} \cdots W_1 \Sigma_{XY} U_S U_S^T. \end{aligned}$$

Therefore,

$$\text{rk}(W_{j-1} \cdots W_1 \Sigma_{XY}) = \text{rk}(W_{j-1} \cdots W_1 \Sigma_{XY} U_S U_S^T).$$

The latter is impossible since $\text{rk}(W_{j-1} \cdots W_1 \Sigma_{XY} U_S U_S^T) \leq |\mathcal{S}| = r$ is not compatible with (63). Therefore (64) holds.

We know that $\text{rk}(W_H \cdots W_{j+1}) \geq \text{rk}(W_H \cdots W_1) = r$. Therefore, depending on the value of $\text{rk}(W_H \cdots W_{j+1})$, we distinguish two situations: either $\text{rk}(W_H \cdots W_{j+1}) > r$ or $\text{rk}(W_H \cdots W_{j+1}) = r$.

When $\text{rk}(W_H \cdots W_{j+1}) > r$, since Σ_{XY} is of full column rank, we have $\text{rk}(\Sigma_{XY} W_H \cdots W_{j+1}) = \text{rk}(W_H \cdots W_{j+1}) > r$. Also, using (63), we have $\text{rk}(W_{j-1} \cdots W_2) \geq \text{rk}(W_{j-1} \cdots W_1 \Sigma_{XY}) > r$. Hence, in this case, the pivot $(j, 1)$ is not tightened either. We have already proved in Section E.1 (beware that the pivot is denoted $(i, 1)$, not $(j, 1)$, in Section E.1) that, when such a pivot is not tightened, \mathbf{W} is not a second-order critical point. This concludes the proof in the case $\text{rk}(W_H \cdots W_{j+1}) > r$.

In the rest of the section we assume that $\text{rk}(W_H \cdots W_{j+1}) = r$.

Using (63), we have $\text{rk}(W_{H-1} \cdots W_{j+1}) > r = \text{rk}(W_H \cdots W_{j+1})$. Applying the rank-nullity theorem we obtain

$$\text{Ker}(W_{H-1} \cdots W_{j+1}) \subsetneq \text{Ker}(W_H \cdots W_{j+1}).$$

Therefore there exists $b \in \mathbb{R}^{d_j}$ such that

$$\begin{cases} b \in \text{Ker}(W_H \cdots W_{j+1}) \\ b \notin \text{Ker}(W_{H-1} \cdots W_{j+1}). \end{cases} \quad (65)$$

Hence, there also exists $a \in \mathbb{R}^{d_{H-1}}$ such that

$$a^T W_{H-1} \cdots W_{j+1} b = 1. \quad (66)$$

Recall that k, l satisfy (64). We define $\mathbf{W}'_\beta = (W_H'^\beta, \dots, W_1'^\beta)$ by

$$\begin{cases} W_H'^\beta = \beta W_H' = \beta U_k a^T \in \mathbb{R}^{d_y \times d_{H-1}} \\ W_j'^\beta = W_j' = b e_l^T \in \mathbb{R}^{d_j \times d_{j-1}}, \text{ where } e_l \in \mathbb{R}^{d_{j-1}} \\ W_h'^\beta = 0 \quad \forall h \in \llbracket 1, H \rrbracket \setminus \{i, j\} \end{cases}$$

We set $\mathbf{W}^\beta(t) = (W_H^\beta(t), \dots, W_1^\beta(t))$ such that $W_h^\beta(t) = W_h + tW_h'^\beta$ for $h \in \llbracket 1, H \rrbracket$. We have

$$\begin{aligned} W^\beta(t) &:= W_H^\beta(t) \cdots W_1^\beta(t) \\ &= (W_H + t\beta W_H') W_{H-1} \cdots W_{j+1} (W_j + tW_j') W_{j-1} \cdots W_1 \\ &= W_H \cdots W_1 + t(\beta W_H' W_{H-1} \cdots W_1 + W_H \cdots W_{j+1} W_j' W_{j-1} \cdots W_1) \\ &\quad + t^2 \beta W_H' \cdots W_{j+1} W_j' W_{j-1} \cdots W_1. \end{aligned}$$

Using Proposition 1 and the definition of \mathbf{W}_β' above, we obtain

$$\begin{aligned} W^\beta(t) &= U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} + t(\beta W_H' W_{H-1} \cdots W_1 + W_H \cdots W_{j+1} b e_t^T W_{j-1} \cdots W_1) \\ &\quad + \beta t^2 U_k a^T W_{H-1} \cdots W_{j+1} b (W_{j-1} \cdots W_1)_{l, \cdot}, \\ &= U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} + t\beta W_H' W_{H-1} \cdots W_1 + \beta t^2 U_k (W_{j-1} \cdots W_1)_{l, \cdot}, \end{aligned}$$

where the last equality follows from (65) and (66).

Denoting $N = W_H' W_{H-1} \cdots W_1$, we have

$$\begin{aligned} L(\mathbf{W}^\beta(t)) &= \|W^\beta(t)X - Y\|^2 \\ &= \|U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X - Y + t\beta N X + \beta t^2 U_k (W_{j-1} \cdots W_1)_{l, \cdot} X\|^2. \end{aligned}$$

Using the cyclic property of the trace, and, since $k \notin \mathcal{S}$, $U_S^T U_k = 0$, we get in this case a second-order coefficient equal to

$$\begin{aligned} c_2(\mathbf{W}, \mathbf{W}_\beta') &= \beta^2 \|NX\|^2 + 2\beta \operatorname{tr}(U_k (W_{j-1} \cdots W_1)_{l, \cdot} X X^T \Sigma_{XX}^{-1} \Sigma_{XY} U_S U_S^T) - 2\beta \operatorname{tr}(U_k (W_{j-1} \cdots W_1)_{l, \cdot} \Sigma_{XY}) \\ &= \beta^2 \|NX\|^2 - 2\beta (W_{j-1} \cdots W_1)_{l, \cdot} \Sigma_{XY} U_k. \end{aligned}$$

Since from (64), $(W_{j-1} \cdots W_1)_{l, \cdot} \Sigma_{XY} U_k \neq 0$, we can choose β according to (7), such that $c_2(\mathbf{W}, \mathbf{W}_\beta') < 0$. Therefore \mathbf{W} is not a second-order critical point.

E.4 4th case: $i, j \in \llbracket 2, H-1 \rrbracket$, with $i > j$

In this case, the two complementary blocks are $W_{j-1} \cdots W_1 \Sigma_{XY} W_H \cdots W_{i+1}$ and $W_{i-1} \cdots W_{j+1}$. We follow again the same lines as above. Since the pivot (i, j) is not tightened, using Proposition 3, we have

$$\begin{cases} \operatorname{rk}(W_{i-1} \cdots W_{j+1}) > r \\ \operatorname{rk}(W_{j-1} \cdots W_1 \Sigma_{XY} W_H \cdots W_{i+1}) > r. \end{cases} \quad (67)$$

Let us first show that there exist $k \in \llbracket 1, d_i \rrbracket$ and $l \in \llbracket 1, d_{j-1} \rrbracket$ such that

$$(W_{j-1} \cdots W_1)_{l, \cdot} \Sigma_{XY} U_Q U_Q^T (W_H \cdots W_{i+1})_{\cdot, k} \neq 0. \quad (68)$$

Indeed, assume by contradiction that, for all $k \in \llbracket 1, d_i \rrbracket$ and $l \in \llbracket 1, d_{j-1} \rrbracket$, we have

$$(W_{j-1} \cdots W_1)_{l, \cdot} \Sigma_{XY} U_Q U_Q^T (W_H \cdots W_{i+1})_{\cdot, k} = 0.$$

Then $W_{j-1} \cdots W_1 \Sigma_{XY} U_Q U_Q^T W_H \cdots W_{i+1} = 0$, and so, using $I_{d_y} = U_S U_S^T + U_Q U_Q^T$, we would have

$$\begin{aligned} W_{j-1} \cdots W_1 \Sigma_{XY} W_H \cdots W_{i+1} &= W_{j-1} \cdots W_1 \Sigma_{XY} I_{d_y} W_H \cdots W_{i+1} \\ &= W_{j-1} \cdots W_1 \Sigma_{XY} (U_S U_S^T + U_Q U_Q^T) W_H \cdots W_{i+1} \\ &= W_{j-1} \cdots W_1 \Sigma_{XY} U_S U_S^T W_H \cdots W_{i+1}. \end{aligned}$$

Therefore,

$$\text{rk}(W_{j-1} \cdots W_1 \Sigma_{XY} W_H \cdots W_{i+1}) = \text{rk}(W_{j-1} \cdots W_1 \Sigma_{XY} U_S U_S^T W_H \cdots W_{i+1}).$$

The latter is impossible since $\text{rk}(W_{j-1} \cdots W_1 \Sigma_{XY} U_S U_S^T W_H \cdots W_{i+1}) \leq |S| = r$ is not compatible with (67). Therefore (68) holds.

We know that $\text{rk}(W_H \cdots W_{j+1}) \geq \text{rk}(W_H \cdots W_1) = r$. Therefore, depending on the value of $\text{rk}(W_H \cdots W_{j+1})$, we distinguish two situations: either $\text{rk}(W_H \cdots W_{j+1}) > r$ or $\text{rk}(W_H \cdots W_{j+1}) = r$.

When $\text{rk}(W_H \cdots W_{j+1}) > r$, since Σ_{XY} is of full column rank, we have $\text{rk}(\Sigma_{XY} W_H \cdots W_{j+1}) = \text{rk}(W_H \cdots W_{j+1}) > r$. Also, using (67), we have $\text{rk}(W_{j-1} \cdots W_2) \geq \text{rk}(W_{j-1} \cdots W_1 \Sigma_{XY} W_H \cdots W_{i+1}) > r$. Hence, in this case, the pivot $(j, 1)$ is not tightened either. We have already proved in Section E.1 (beware that the pivot is denoted $(i, 1)$, not $(j, 1)$, in Section E.1) that, when such a pivot is not tightened, \mathbf{W} is not a second-order critical point. This concludes the proof when $\text{rk}(W_H \cdots W_{j+1}) > r$.

In the rest of the section we assume that $\text{rk}(W_H \cdots W_{j+1}) = r$.

Using (67), we have $\text{rk}(W_{i-1} \cdots W_{j+1}) > r = \text{rk}(W_H \cdots W_{j+1})$. Applying the rank-nullity theorem, we obtain

$$\text{Ker}(W_{i-1} \cdots W_{j+1}) \subsetneq \text{Ker}(W_H \cdots W_{j+1}).$$

Therefore there exists $b \in \mathbb{R}^{d_j}$ such that

$$\begin{cases} b \in \text{Ker}(W_H \cdots W_{j+1}) \\ b \notin \text{Ker}(W_{i-1} \cdots W_{j+1}). \end{cases} \quad (69)$$

Hence, there also exists $a \in \mathbb{R}^{d_{i-1}}$ such that

$$a^T W_{i-1} \cdots W_{j+1} b = 1. \quad (70)$$

Recall that k, l satisfy (68). We define $\mathbf{W}'_\beta = (W'_H{}^\beta, \dots, W'_1{}^\beta)$ by

$$\begin{cases} W'_i{}^\beta = \beta W'_i = \beta e_k a^T \in \mathbb{R}^{d_i \times d_{i-1}} \text{ where } e_k \in \mathbb{R}^{d_i} \\ W'_j{}^\beta = W'_j = b e_l^T \in \mathbb{R}^{d_j \times d_{j-1}} \text{ where } e_l \in \mathbb{R}^{d_{j-1}} \\ W'_h{}^\beta = 0 \quad \forall h \in \llbracket 1, H \rrbracket \setminus \{i, j\}. \end{cases}$$

We set $\mathbf{W}^\beta(t) = (W_H^\beta(t), \dots, W_1^\beta(t))$ with $W_h^\beta(t) = W_h + t W'_h{}^\beta$ for all $h \in \llbracket 1, H \rrbracket$. We have,

$$\begin{aligned} W^\beta(t) &:= W_H^\beta(t) \cdots W_1^\beta(t) \\ &= W_H \cdots W_{i+1} (W_i + t \beta W'_i) W_{i-1} \cdots W_{j+1} (W_j + t W'_j) W_{j-1} \cdots W_1 \\ &= W_H \cdots W_1 + t (\beta W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_1 + W_H \cdots W_{j+1} W'_j W_{j-1} \cdots W_1) \\ &\quad + \beta t^2 W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_{j+1} W'_j W_{j-1} \cdots W_1. \end{aligned}$$

Using Proposition 1 and the definition of \mathbf{W}'_β above, we obtain

$$\begin{aligned} W^\beta(t) &= U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} + t (\beta W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_1 + W_H \cdots W_{j+1} b e_l^T W_{j-1} \cdots W_1) \\ &\quad + \beta t^2 (W_H \cdots W_{i+1})_{.,k} a^T W_{i-1} \cdots W_{j+1} b (W_{j-1} \cdots W_1)_l, \\ &= U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} + t \beta W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_1 + \beta t^2 (W_H \cdots W_{i+1})_{.,k} (W_{j-1} \cdots W_1)_l, \end{aligned}$$

where the last equality follows from (69) and (70).

Denoting $N = W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_1$, we have

$$\begin{aligned} L(\mathbf{W}^\beta(t)) &= \|W^\beta(t) X - Y\|^2 \\ &= \|U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X - Y + t \beta N X + \beta t^2 (W_H \cdots W_{i+1})_{.,k} (W_{j-1} \cdots W_1)_l X\|^2. \end{aligned}$$

The second-order coefficient of $L(\mathbf{W}^\beta(t))$ is equal to

$$\begin{aligned} c_2(\mathbf{W}, \mathbf{W}'_\beta) &= \beta^2 \|NX\|^2 + 2\beta \operatorname{tr} \left((W_H \cdots W_{i+1})_{.,k} (W_{j-1} \cdots W_1)_{l,.} X X^T \Sigma_{XX}^{-1} \Sigma_{XY} U_S U_S^T \right) \\ &\quad - 2\beta \operatorname{tr} \left((W_H \cdots W_{i+1})_{.,k} (W_{j-1} \cdots W_1)_{l,.} \Sigma_{XY} \right) \\ &= \beta^2 \|NX\|^2 + 2\beta \operatorname{tr} \left((W_H \cdots W_{i+1})_{.,k} (W_{j-1} \cdots W_1)_{l,.} \Sigma_{XY} (U_S U_S^T - I_{d_y}) \right) . \end{aligned}$$

Using, from Lemma 6, that $U_S U_S^T - I_{d_y} = -U_Q U_Q^T$, and then the cyclic property of the trace, we obtain

$$\begin{aligned} c_2(\mathbf{W}, \mathbf{W}'_\beta) &= \beta^2 \|NX\|^2 - 2\beta \operatorname{tr} \left((W_H \cdots W_{i+1})_{.,k} (W_{j-1} \cdots W_1)_{l,.} \Sigma_{XY} U_Q U_Q^T \right) \\ &= \beta^2 \|NX\|^2 - 2\beta (W_{j-1} \cdots W_1)_{l,.} \Sigma_{XY} U_Q U_Q^T (W_H \cdots W_{i+1})_{.,k} . \end{aligned}$$

Since from (68), $(W_{j-1} \cdots W_1)_{l,.} \Sigma_{XY} U_Q U_Q^T (W_H \cdots W_{i+1})_{.,k} \neq 0$, we can choose β according to (7) such that $c_2(\mathbf{W}, \mathbf{W}'_\beta) < 0$. Therefore, \mathbf{W} is not a second-order critical point.

F Non-strict saddle points

In this section, we prove the results related to non-strict saddle points (see Section 4.3).

F.1 Proof of Proposition 11

To prove Proposition 11, we show that for any \mathbf{W}' , $c_2(\mathbf{W}, \mathbf{W}') \geq 0$, which is equivalent to say (see Lemma 1) that \mathbf{W} is a second-order critical point. We follow the proof strategy sketched in Section 4.3 after the statement of Proposition 11, and use the same notation introduced therein. Note that a first-order critical point can only be tightened if $H \geq 3$. Therefore, in all of this section we make the assumption $H \geq 3$. Recall that m is the number of examples in our sample, $\mathcal{S} = \llbracket 1, r \rrbracket$, with $r < r_{max}$. We set $Q = \llbracket r+1, d_y \rrbracket$. Recall also that

$$\Sigma^{1/2} = \Sigma_{YX} \Sigma_{XX}^{-1} X \in \mathbb{R}^{d_y \times m} .$$

and

$$\Sigma^{1/2} = U \Delta V^T$$

is a Singular Value Decomposition of $\Sigma^{1/2}$, where $\Delta \in \mathbb{R}^{d_y \times m}$ is such that $\Delta_{ii} = \sqrt{\lambda_i}$ for all $i \in \llbracket 1, d_y \rrbracket$, and $(\lambda_i)_{i=1..d_y}$ are the eigenvalues of Σ . We denote

$$\Delta^{(\mathcal{S})} = \operatorname{diag}(\sqrt{\lambda_1}, \dots, \sqrt{\lambda_r}) \in \mathbb{R}^{r \times r} \tag{71}$$

and

$$\Delta^{(Q)} = \operatorname{diag}(\sqrt{\lambda_{r+1}}, \dots, \sqrt{\lambda_{d_y}}) \in \mathbb{R}^{(d_y-r) \times (d_y-r)} . \tag{72}$$

Recall that, from Section 4.3, $c_2(\mathbf{W}, \mathbf{W}') = FT + ST$.

In what follows, we are going to present a key lemma, then various quick technical lemmas, then we simplify the expressions of FT and ST and conclude the proof of Proposition 11. Then, we prove all the lemmas of Appendix F.1. We present a lemma which uses that \mathbf{W} is tightened to simplify some products of weight matrices and lighten further calculations. This is a key lemma as it introduces indices p and q which will be used multiple times in the proof.

Lemma 14. *Suppose Assumption \mathcal{H} in Section 2 holds true. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L verifying the hypotheses of Proposition 11, and $r, \mathcal{S}, Q, (Z_h)_{h=1..H}$ as in Proposition 11. If \mathbf{W} is tightened, then, there exist $p \in \llbracket 3, H \rrbracket$ and $q \in \llbracket 1, \min(p-1, H-2) \rrbracket$ such that:*

$$\forall i \in \llbracket 1, p-1 \rrbracket, \quad W_H \cdots W_{i+1} = [U_{\mathcal{S}}, 0] \quad (73)$$

$$\forall i \in \llbracket p, H \rrbracket, \quad W_{i-1} \cdots W_2 = \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix} \quad (74)$$

$$\forall i \in \llbracket q+1, H \rrbracket, \quad Z_{i-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q = 0 \quad (75)$$

$$\forall i \in \llbracket 1, q \rrbracket, \quad W_{H-1} \cdots W_{i+1} = \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix}. \quad (76)$$

The proof of Lemma 14 is in Appendix F.1.5.

F.1.1 Useful technical lemmas

We now present technical lemmas which will be useful in Sections F.1.2, F.1.3 and F.1.4. In all of these Lemmas, we have $\mathcal{S} = \llbracket 1, r \rrbracket$ and $Q = \llbracket r+1, d_y \rrbracket$, and Assumption \mathcal{H} holds true.

Lemma 15. *We have*

$$\Sigma_{XY} U_Q = X V_Q \Delta^{(Q)}.$$

The proof of Lemma 15 is in Appendix F.1.6.

Lemma 16. *Let n be a positive integer. For any matrices $A \in \mathbb{R}^{d_y \times n}$ and $B \in \mathbb{R}^{r \times n}$ we have*

$$\|A + U_{\mathcal{S}} B\|^2 = \|U_{\mathcal{S}}^T A + B\|^2 + \|U_Q^T A\|^2.$$

The proof of Lemma 16 is in Appendix F.1.7.

Lemma 17. *Let n be any positive integer. For any matrices $A \in \mathbb{R}^{n \times r}$ and $B \in \mathbb{R}^{n \times (d_y - r)}$ we have:*

$$\langle A U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} X, B V_Q^T \rangle = 0.$$

The proof of Lemma 17 is in Appendix F.1.8.

Lemma 18. *Let n be any positive integer. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point of L verifying the hypotheses of Proposition 11, and $r, \mathcal{S}, Q, (Z_h)_{h=1..H}$ as in Proposition 11. If \mathbf{W} is tightened, then, for q as in Lemma 14, for any matrices $A \in \mathbb{R}^{n \times (d_q - r)}$ and $B \in \mathbb{R}^{n \times (d_y - r)}$, we have:*

$$\langle A Z_q \cdots Z_2 Z_1 X, B V_Q^T \rangle = 0.$$

The proof of Lemma 18 is in Appendix F.1.9.

Lemma 19. *For any matrix $A \in \mathbb{R}^{(d_y - r) \times r}$ we have*

$$\|A U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} X\|^2 = \sum_{a=1}^r \sum_{b=r+1}^{d_y} (\lambda_a - \lambda_b) (A_{b-r,a})^2 + \|\Delta^{(Q)} A\|^2.$$

The proof of Lemma 19 is in Appendix F.1.10.

Lemma 20. *Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point associated with \mathcal{S} . For any matrix $A \in \mathbb{R}^{d_y \times d_x}$, we have*

$$\langle A X, W_H \cdots W_1 X - Y \rangle = \langle A, -U_Q U_Q^T \Sigma_{YX} \rangle.$$

The proof of Lemma 20 is in Appendix F.1.11.

F.1.2 Simplifying FT

In this section and the next one, we simplify the expressions of FT and ST as defined in (14) and (15). In order to decompose $FT = a_1 + \|A_2\|^2 + \|A_3\|^2 + \|A_4\|^2$, with $a_1 \geq 0$, we first simplify the terms T_i , for $i \in \llbracket 1, H \rrbracket$, defined in (12). Let us first consider \mathbf{W} tightened satisfying the hypotheses of Proposition 11, and p and q defined as in Lemma 14. The simplification of T_i depends on the position of i with regard to 1, q , p and H . We define $J_1 = \llbracket p, H - 1 \rrbracket$, $J_2 = \llbracket q + 1, p - 1 \rrbracket$ and $J_3 = \llbracket 2, q \rrbracket$.

Note that, according to the convention in Section 2, these sets could be empty.

- if $p = H$, $J_1 = \emptyset$
- if $q = p - 1$, $J_2 = \emptyset$
- if $q = 1$, $J_3 = \emptyset$.

Note also that $\{1\}, J_3, J_2, J_1, \{H\}$ are disjoint and $\{1\} \cup J_3 \cup J_2 \cup J_1 \cup \{H\} = \llbracket 1, H \rrbracket$. Depending on the position of i , we need to distinguish four cases, in order to simplify T_i .

Lemma 21. *Suppose Assumption \mathcal{H} in Section 2 holds true. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point satisfying the hypotheses of Proposition 11, and $r, \mathcal{S}, Q, (Z_h)_{h=1..H}$ as in Proposition 11. Let $i \in \llbracket 1, H \rrbracket$. For any $\mathbf{W}' = (W'_H, \dots, W'_1)$, recall that, as defined in (12),*

$$T_i = W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_1 X.$$

If \mathbf{W} is tightened, then, for p and q as defined in Lemma 14 and J_1, J_2, J_3 as defined above, we have

- For $i = H$:

$$T_H = (W'_H)_{.,1:r} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} X \quad (77)$$

- For $i \in J_1$:

$$T_i = U_{\mathcal{S}}(W'_i)_{1:r,1:r} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} X + U_Q Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} X \quad (78)$$

- For $i \in J_2 \cup J_3$:

$$T_i = U_{\mathcal{S}}(W'_i)_{1:r,1:r} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} X + U_{\mathcal{S}}(W'_i)_{1:r,r+1:d_{i-1}} Z_{i-1} \cdots Z_2 Z_1 X \quad (79)$$

- For $i = 1$:

$$T_1 = U_{\mathcal{S}}(W'_1)_{1:r,.} X \quad (80)$$

The proof of Lemma 21 is in Appendix F.1.12.

We now simplify FT . Substituting the formulas of Lemma 21 in (14) we have

$$\begin{aligned} FT &= \left\| \sum_{i=1}^H T_i \right\|^2 \\ &= \left\| (W'_H)_{.,1:r} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} X \right. \\ &\quad + \sum_{i \in J_1} \left(U_{\mathcal{S}}(W'_i)_{1:r,1:r} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} X + U_Q Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} X \right) \\ &\quad \left. + \sum_{i \in J_2 \cup J_3} \left(U_{\mathcal{S}}(W'_i)_{1:r,1:r} U_{\mathcal{S}}^T \Sigma_{YX} \Sigma_{XX}^{-1} X + U_{\mathcal{S}}(W'_i)_{1:r,r+1:d_{i-1}} Z_{i-1} \cdots Z_2 Z_1 X \right) + U_{\mathcal{S}}(W'_1)_{1:r,.} X \right\|^2. \end{aligned}$$

FT can be identified with a term as $\|A + U_S B\|^2$ if we take

$$A = (W'_H)_{:,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + \sum_{i \in J_1} U_Q Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X .$$

and

$$B = \sum_{i \in J_1} (W'_i)_{1:r,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X \\ + \sum_{i \in J_2 \cup J_3} ((W'_i)_{1:r,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + (W'_i)_{1:r,r+1:d_{i-1}} Z_{i-1} \cdots Z_2 Z_1 X) + (W'_1)_{1:r, \cdot} X .$$

Applying Lemma 16, FT becomes:

$$FT = \|U_S^T A + B\|^2 + \|U_Q^T A\|^2 \\ = \left\| U_S^T (W'_H)_{:,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + \sum_{i \in J_1} U_S^T U_Q Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X \right. \\ \left. + \sum_{i \in J_1} (W'_i)_{1:r,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X \right. \\ \left. + \sum_{i \in J_2 \cup J_3} ((W'_i)_{1:r,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + (W'_i)_{1:r,r+1:d_{i-1}} Z_{i-1} \cdots Z_2 Z_1 X) + (W'_1)_{1:r, \cdot} X \right\|^2 \\ + \left\| U_Q^T (W'_H)_{:,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + \sum_{i \in J_1} U_Q^T U_Q Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X \right\|^2 .$$

Using Lemma 6, we have $U_S^T U_Q = 0$ and $U_Q^T U_Q = I_{d_y - r}$, hence we can write

$$FT = FT_1 + FT_2 ,$$

where

$$FT_1 = \left\| U_S^T (W'_H)_{:,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + \sum_{i \in J_1} (W'_i)_{1:r,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X \right. \\ \left. + \sum_{i \in J_2 \cup J_3} ((W'_i)_{1:r,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + (W'_i)_{1:r,r+1:d_{i-1}} Z_{i-1} \cdots Z_2 Z_1 X) + (W'_1)_{1:r, \cdot} X \right\|^2 ,$$

and

$$FT_2 = \left\| U_Q^T (W'_H)_{:,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + \sum_{i \in J_1} Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X \right\|^2 .$$

Let us first simplify FT_1 .

Recall that m is the number of examples in our sample, $V \in \mathbb{R}^{m \times m}$ is the orthogonal matrix defined in (1) and $Q = \llbracket r+1, d_y \rrbracket$. We set $S' = \mathcal{S} \cup \llbracket d_y+1, m \rrbracket = \llbracket 1, r \rrbracket \cup \llbracket d_y+1, m \rrbracket$ such that $S' \cup Q = \llbracket 1, m \rrbracket$.

Reordering the terms and, since V is orthogonal, using $I_m = VV^T = V_{S'} V_{S'}^T + V_Q V_Q^T$, we have

$$FT_1 = \left\| \left(U_S^T (W'_H)_{:,1:r} + \sum_{i \in J_1 \cup J_2 \cup J_3} (W'_i)_{1:r,1:r} \right) U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X \right\|^2$$

$$\begin{aligned}
& + \sum_{i \in J_2} (W'_i)_{1:r, r+1:d_{i-1}} Z_{i-1} \cdots Z_2 Z_1 X \\
& + \left(\sum_{i \in J_3} (W'_i)_{1:r, r+1:d_{i-1}} Z_{i-1} \cdots Z_2 Z_1 X + (W'_1)_{1:r, \cdot} X \right) (V_{S'} V_{S'}^T + V_Q V_Q^T) \Bigg\|^2.
\end{aligned}$$

Since for $i \in J_2$, we have $i - 1 \geq q$, we denote

$$N := \sum_{i \in J_2} (W'_i)_{1:r, r+1:d_{i-1}} Z_{i-1} \cdots Z_{q+1},$$

Recall that, using the convention in Section 2, for $i - 1 = q$, we have $Z_{i-1} \cdots Z_{q+1} = I_{d_q - r}$.

We also denote

$$\begin{aligned}
M & := \sum_{i \in J_3} (W'_i)_{1:r, r+1:d_{i-1}} Z_{i-1} \cdots Z_2 Z_1 X V_Q + (W'_1)_{1:r, \cdot} X V_Q, \\
J & := \sum_{i \in J_3} (W'_i)_{1:r, r+1:d_{i-1}} Z_{i-1} \cdots Z_2 Z_1 X V_{S'} + (W'_1)_{1:r, \cdot} X V_{S'}, \\
L & := U_S^T (W'_H)_{\cdot, 1:r} + \sum_{i \in J_1 \cup J_2 \cup J_3} (W'_i)_{1:r, 1:r}.
\end{aligned}$$

Therefore, we obtain

$$\begin{aligned}
FT_1 & = \|LU_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + NZ_q \cdots Z_2 Z_1 X + J V_{S'}^T + M V_Q^T\|^2 \\
& = \|LU_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + NZ_q \cdots Z_2 Z_1 X + J V_{S'}^T\|^2 + \|M V_Q^T\|^2 \\
& \quad + 2 \langle LU_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + NZ_q \cdots Z_2 Z_1 X + J V_{S'}^T, M V_Q^T \rangle.
\end{aligned}$$

Using Lemma 17 and Lemma 18 and $V_Q^T V_{S'} = 0$ (since V is orthogonal), the cross-product is equal to zero.

Noting also that since V is orthogonal $\|M V_Q^T\|^2 = \text{tr}(M V_Q^T V_Q M^T) = \text{tr}(M M^T) = \|M\|^2 = \|M^T\|^2$, we have

$$\begin{aligned}
FT_1 & = \|LU_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + NZ_q \cdots Z_2 Z_1 X + J V_{S'}^T\|^2 + \|M^T\|^2 \\
& = \|A_2\|^2 + \|A_4\|^2
\end{aligned}$$

where

$$\begin{aligned}
A_2 & := LU_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + NZ_q \cdots Z_2 Z_1 X + J V_{S'}^T \\
& = U_S^T (W'_H)_{\cdot, 1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + \sum_{i \in J_1} (W'_i)_{1:r, 1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X \\
& \quad + \sum_{i \in J_2} ((W'_i)_{1:r, 1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + (W'_i)_{1:r, r+1:d_{i-1}} Z_{i-1} \cdots Z_2 Z_1 X) \\
& \quad + \sum_{i \in J_3} ((W'_i)_{1:r, 1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + (W'_i)_{1:r, r+1:d_{i-1}} Z_{i-1} \cdots Z_2 Z_1 X V_{S'} V_{S'}^T) + (W'_1)_{1:r, \cdot} X V_{S'} V_{S'}^T \quad (81)
\end{aligned}$$

$$A_4 := M^T = \left(\sum_{i \in J_3} (W'_i)_{1:r, r+1:d_{i-1}} Z_{i-1} \cdots Z_2 Z_1 X V_Q + (W'_1)_{1:r, \cdot} X V_Q \right)^T. \quad (82)$$

Let us now simplify FT_2 .

We have $FT_2 = \|AU_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X\|^2$, with

$$A := U_Q^T (W'_H)_{\cdot, 1:r} + \sum_{i \in J_1} Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i, 1:r} \in \mathbb{R}^{(d_y - r) \times r}.$$

Hence, using Lemma 19, we have

$$\begin{aligned}
FT_2 &= \sum_{a=1}^r \sum_{b=r+1}^{d_y} (\lambda_a - \lambda_b) (A_{b-r,a})^2 + \|\Delta^{(Q)} A\|^2 \\
&= \sum_{a=1}^r \sum_{b=r+1}^{d_y} (\lambda_a - \lambda_b) \left(U_b^T (W'_H)_{.,a} + \sum_{i \in J_1} (Z_H)_{b-r,} Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,a} \right)^2 \\
&\quad + \left\| \Delta^{(Q)} \left(U_Q^T (W'_H)_{.,1:r} + \sum_{i \in J_1} Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r} \right) \right\|^2 \\
&= a_1 + \|A_3\|^2,
\end{aligned}$$

where

$$a_1 := \sum_{a=1}^r \sum_{b=r+1}^{d_y} (\lambda_a - \lambda_b) \left(U_b^T (W'_H)_{.,a} + \sum_{i \in J_1} (Z_H)_{b-r,} Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,a} \right)^2 \quad (83)$$

$$A_3 := \Delta^{(Q)} \left(U_Q^T (W'_H)_{.,1:r} + \sum_{i \in J_1} Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r} \right) \quad (84)$$

Finally,

$$\begin{aligned}
FT &= FT_1 + FT_2 \\
&= a_1 + \|A_2\|^2 + \|A_3\|^2 + \|A_4\|^2,
\end{aligned} \quad (85)$$

where a_1, A_2, A_3, A_4 are defined in (83), (81), (84), (82). Notice that, since $\lambda_1 > \cdots > \lambda_{d_y}$, we have

$$a_1 \geq 0. \quad (86)$$

F.1.3 Simplifying ST

In this section, we prove that $ST = -2 \langle A_3, A_4 \rangle$, where ST, A_3 and A_4 are defined in (15), (84) and (82). In order to do so, we first state a lemma that simplifies the terms $T_{i,j}$ defined in (13). We remind that the sets J_1, J_2 and J_3 are defined at the beginning of Section F.1.2.

Lemma 22. *Suppose Assumption \mathcal{H} in Section 2 holds true. Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point satisfying the hypotheses of Proposition 11, and $r, \mathcal{S}, Q, (Z_h)_{h=1..H}$ defined as in Proposition 11. Let $(i, j) \in \llbracket 1, H \rrbracket^2$, with $i > j$. For any $\mathbf{W}' = (W'_H, \dots, W'_1)$, recall that, as defined in (13),*

$$T_{i,j} = \langle W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_{j+1} W'_j W_{j-1} \cdots W_1 X, W_H \cdots W_1 X - Y \rangle.$$

If \mathbf{W} is tightened, then, for p and q as defined in Lemma 14 and J_1, J_2, J_3 as defined above, we have

- For $i = H$:

- For $j \in J_3$:

$$T_{H,j} = - \left\langle \Delta^{(Q)} U_Q^T (W'_H)_{.,1:r}, ((W'_j)_{1:r,r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 X V_Q)^T \right\rangle. \quad (87)$$

- For $j = 1$:

$$T_{H,1} = - \left\langle \Delta^{(Q)} U_Q^T (W'_H)_{.,1:r}, ((W'_1)_{1:r,} X V_Q)^T \right\rangle. \quad (88)$$

– For $j \in J_1 \cup J_2$:

$$T_{H,j} = 0. \quad (89)$$

• For $i \in J_1$:

– For $j \in J_3$:

$$T_{i,j} = - \left\langle \Delta^{(Q)} Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r}, ((W'_j)_{1:r,r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 X V_Q)^T \right\rangle. \quad (90)$$

– For $j = 1$:

$$T_{i,1} = - \left\langle \Delta^{(Q)} Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r}, ((W'_1)_{1:r}, X V_Q)^T \right\rangle. \quad (91)$$

– For $j \in J_1 \cup J_2$:

$$T_{i,j} = 0. \quad (92)$$

• For $i \in J_2 \cup J_3$, for all $j < i$, we have

$$T_{i,j} = 0. \quad (93)$$

The proof of Lemma 22 is in Appendix F.1.13.

Let us now prove that $ST = -2 \langle A_3, A_4 \rangle$. We remind that $\llbracket 1, H \rrbracket = \{H\} \cup J_1 \cup J_2 \cup J_3 \cup \{1\}$ and separate the sum appearing in (15) accordingly.

We then substitute the formulas of Lemma 22 in (15) and obtain

$$\begin{aligned} ST &= 2 \sum_{H \geq i > j \geq 1} T_{i,j} \\ &= 2 \left(\sum_{j \in J_1 \cup J_2} T_{H,j} + \sum_{j \in J_3} T_{H,j} + T_{H,1} + \sum_{i \in J_1} \sum_{\substack{j \in J_1 \cup J_2, \\ j < i}} T_{i,j} + \sum_{i \in J_1} \sum_{j \in J_3} T_{i,j} + \sum_{i \in J_1} T_{i,1} + \sum_{i \in J_2 \cup J_3} \sum_{j=1}^{i-1} T_{i,j} \right) \\ &= -2 \sum_{j \in J_3} \left\langle \Delta^{(Q)} U_Q^T (W'_H)_{.,1:r}, ((W'_j)_{1:r,r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 X V_Q)^T \right\rangle \\ &\quad - 2 \left\langle \Delta^{(Q)} U_Q^T (W'_H)_{.,1:r}, ((W'_1)_{1:r}, X V_Q)^T \right\rangle \\ &\quad - 2 \sum_{i \in J_1} \sum_{j \in J_3} \left\langle \Delta^{(Q)} Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r}, ((W'_j)_{1:r,r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 X V_Q)^T \right\rangle \\ &\quad - 2 \sum_{i \in J_1} \left\langle \Delta^{(Q)} Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r}, ((W'_1)_{1:r}, X V_Q)^T \right\rangle \\ &= -2 \left\langle \Delta^{(Q)} U_Q^T (W'_H)_{.,1:r}, \left(\sum_{j \in J_3} (W'_j)_{1:r,r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 X V_Q + (W'_1)_{1:r}, X V_Q \right)^T \right\rangle \\ &\quad - 2 \left\langle \Delta^{(Q)} \sum_{i \in J_1} Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r}, \left(\sum_{j \in J_3} (W'_j)_{1:r,r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 X V_Q + (W'_1)_{1:r}, X V_Q \right)^T \right\rangle \\ &= -2 \left\langle \Delta^{(Q)} \left(U_Q^T (W'_H)_{.,1:r} + \sum_{i \in J_1} Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r} \right), \right. \end{aligned}$$

$$\begin{aligned}
& \left(\sum_{j \in J_3} (W'_j)_{1:r, r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 X V_Q + (W_1)_{1:r, \cdot} X V_Q \right)^T \Bigg\rangle \\
& = -2 \langle A_3, A_4 \rangle ,
\end{aligned} \tag{94}$$

where we remind that A_3 and A_4 are defined in (84) and (82).

F.1.4 Concluding the proof of Proposition 11

Using the simplifications (85) and (94) above, for any \mathbf{W} satisfying the hypotheses of Proposition 11, if \mathbf{W} is tightened, then for any \mathbf{W}' ,

$$\begin{aligned}
c_2(\mathbf{W}, \mathbf{W}') &= FT + ST \\
&= a_1 + \|A_2\|^2 + \|A_3\|^2 + \|A_4\|^2 - 2 \langle A_3, A_4 \rangle \\
&= a_1 + \|A_2\|^2 + \|A_3 - A_4\|^2 .
\end{aligned}$$

Using (86), we find $c_2(\mathbf{W}, \mathbf{W}') \geq 0$.

Therefore, $\mathbf{W} = (W_H, \dots, W_1)$ is a second-order critical point.

F.1.5 Proof of Lemma 14

First note that, for $r = 0$, we can easily follow the same proof and see that the result still holds with the conventions adopted in Section 2.

Let us prove (73).

Consider the pivot $(i, j) = (2, 1)$. Its complementary blocks are $\Sigma_{XY} W_H \cdots W_3$ and I_{d_1} . Since \mathbf{W} is tightened and $\text{rk}(I_{d_1}) = d_1 \geq r_{max} > r$, we have $\text{rk}(\Sigma_{XY} W_H \cdots W_3) = r$. Since Σ_{XY} is full-column rank, we obtain $\text{rk}(W_H \cdots W_3) = r$.

Let $p \in \llbracket 3, H \rrbracket$ be the largest index such that

$$rk(W_H \cdots W_p) = r . \tag{95}$$

Using (8) and (10), we have $W_H \cdots W_p = [U_S, U_Q Z_H Z_{H-1} \cdots Z_p]$.

Since $rk(W_H \cdots W_p) = r$ and since the columns of $U_Q Z_H Z_{H-1} \cdots Z_p$ are in the vector space spanned by the columns of U_Q (which are orthogonal to the columns of U_S), (95) implies

$$Z_H Z_{H-1} \cdots Z_p = 0 .$$

Therefore,

$$W_H \cdots W_p = [U_S, 0] .$$

Using (10), for all $i \in \llbracket 1, p-1 \rrbracket$,

$$\begin{aligned}
W_H \cdots W_{i+1} &= (W_H \cdots W_p)(W_{p-1} \cdots W_{i+1}) \\
&= [U_S, 0] \begin{bmatrix} I_r & 0 \\ 0 & Z_{p-1} \cdots Z_{i+1} \end{bmatrix} \\
&= [U_S, 0] .
\end{aligned}$$

This proves (73).

Let us prove (74).

We consider the pivot $(p, 1)$. Its complementary blocks are $\Sigma_{XY} W_H \cdots W_{p+1}$ and $W_{p-1} \cdots W_2$. We have, by definition of p , $\text{rk}(W_H \cdots W_{p+1}) > r$. Therefore, since Σ_{XY} is full-column rank, we have $\text{rk}(\Sigma_{XY} W_H \cdots W_{p+1}) =$

$\text{rk}(W_H \cdots W_{p+1}) > r$. Note that this holds both for $p = H$ and for $p < H$. Hence, since \mathbf{W} is tightened, the second complementary block is of rank r , i.e.

$$\text{rk}(W_{p-1} \cdots W_2) = r .$$

Using (10), we also have $W_{p-1} \cdots W_2 = \begin{bmatrix} I_r & 0 \\ 0 & Z_{p-1} \cdots Z_2 \end{bmatrix}$.

Then, since $\text{rk}(W_{p-1} \cdots W_2) = r$, we have $Z_{p-1} \cdots Z_2 = 0$ and

$$W_{p-1} \cdots W_2 = \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix} .$$

Using (10) again, for all $i \in \llbracket p, H \rrbracket$,

$$\begin{aligned} W_{i-1} \cdots W_2 &= (W_{i-1} \cdots W_p)(W_{p-1} \cdots W_2) \\ &= \begin{bmatrix} I_r & 0 \\ 0 & Z_{i-1} \cdots Z_p \end{bmatrix} \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix} \\ &= \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix} . \end{aligned}$$

This proves (74).

Let us now prove (75).

Using Proposition 1, Lemma 4 and Lemma 6, we have

$$\text{rk}(W_{p-1} \cdots W_1 \Sigma_{XY}) \geq \text{rk}(W_H \cdots W_1 \Sigma_{XY}) = \text{rk}(U_S U_S^T \Sigma) \geq \text{rk}(U_S^T (U_S U_S^T \Sigma) \Sigma^{-1} U_S) = \text{rk}(I_r) = r .$$

Using (74) for $i = p$, we also have $\text{rk}(W_{p-1} \cdots W_1 \Sigma_{XY}) \leq \text{rk}(W_{p-1} \cdots W_2) = r$. Hence, $\text{rk}(W_{p-1} \cdots W_1 \Sigma_{XY}) = r$.

Notice that, considering the tightened pivot $(H, H - 1)$, since $\text{rk}(I_{d_{H-1}}) = d_{H-1} \geq r_{max} > r$, we obtain $\text{rk}(W_{H-2} \cdots W_1 \Sigma_{XY}) = r$.

We consider $q \in \llbracket 1, \min(p - 1, H - 2) \rrbracket$ the smallest index such that $\text{rk}(W_q \cdots W_1 \Sigma_{XY}) = r$.

Using (10) and (9), we have

$$\begin{aligned} W_q \cdots W_1 \Sigma_{XY} &= \begin{bmatrix} U_S^T \Sigma \\ Z_q \cdots Z_2 Z_1 \Sigma_{XY} \end{bmatrix} \\ &= \begin{bmatrix} \lambda_1 U_1^T \\ \vdots \\ \lambda_r U_r^T \\ Z_q \cdots Z_2 Z_1 \Sigma_{XY} \end{bmatrix} . \end{aligned}$$

Since $\text{rk}(W_q \cdots W_1 \Sigma_{XY}) = r$, every row of $Z_q \cdots Z_2 Z_1 \Sigma_{XY}$ lies in $\text{Vec}(U_1^T, \dots, U_r^T)$, hence we have

$$Z_q \cdots Z_2 Z_1 \Sigma_{XY} U_Q = 0 .$$

Finally, we conclude that, for all $i \in \llbracket q + 1, H \rrbracket$,

$$\begin{aligned} Z_{i-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q &= Z_{i-1} \cdots Z_{q+1} Z_q \cdots Z_2 Z_1 \Sigma_{XY} U_Q \\ &= Z_{i-1} \cdots Z_{q+1} 0 \\ &= 0 . \end{aligned}$$

This proves (75).

Let us now prove (76).

Consider the pivot (H, q) . Its complementary blocks are $W_{q-1} \cdots W_1 \Sigma_{XY}$ and $W_{H-1} \cdots W_{q+1}$. We have, by definition of q , $\text{rk}(W_{q-1} \cdots W_1 \Sigma_{XY}) > r$. Hence, since \mathbf{W} is tightened, the other complementary block is of rank r , i.e. $\text{rk}(W_{H-1} \cdots W_{q+1}) = r$. Using (10), we have

$$W_{H-1} \cdots W_{q+1} = \begin{bmatrix} I_r & 0 \\ 0 & Z_{H-1} \cdots Z_{q+1} \end{bmatrix}.$$

Therefore, $Z_{H-1} \cdots Z_{q+1} = 0$ and

$$W_{H-1} \cdots W_{q+1} = \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix}.$$

Finally, using (10), for all $i \in \llbracket 1, q \rrbracket$,

$$\begin{aligned} W_{H-1} \cdots W_{i+1} &= W_{H-1} \cdots W_{q+1} W_q \cdots W_{i+1} \\ &= \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} I_r & 0 \\ 0 & Z_q \cdots Z_{i+1} \end{bmatrix} \\ &= \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix}. \end{aligned}$$

This proves (76) and concludes the proof.

F.1.6 Proof of Lemma 15

Recall that $\Sigma^{1/2} = \Sigma_{YX} \Sigma_{XX}^{-1} X$. We have

$$\begin{aligned} \Sigma_{XY} &= XY^T \\ &= XX^T (XX^T)^{-1} XY^T \\ &= X(\Sigma^{1/2})^T. \end{aligned}$$

Using (1), we obtain

$$\Sigma_{XY} = XV \Delta^T U^T,$$

and, since U is orthogonal, we have

$$\Sigma_{XY} U = XV \Delta^T.$$

Restricting the equality to the columns in Q , we obtain

$$\Sigma_{XY} U_Q = XV_Q \Delta^{(Q)},$$

where $\Delta^{(Q)}$ is defined in (72). This concludes the proof.

F.1.7 Proof of Lemma 16

Let $A \in \mathbb{R}^{d_y \times n}$ and $B \in \mathbb{R}^{r \times n}$. We have

$$\begin{aligned} \|A + U_S B\|^2 &= \|A\|^2 + \|U_S B\|^2 + 2 \langle A, U_S B \rangle \\ &= \text{tr}(A^T A) + \text{tr}(B^T U_S^T U_S B) + 2 \langle U_S^T A, B \rangle. \end{aligned}$$

Using Lemma 6, this becomes

$$\begin{aligned} \|A + U_S B\|^2 &= \text{tr}(A^T (U_S U_S^T + U_Q U_Q^T) A) + \text{tr}(B^T B) + 2 \langle U_S^T A, B \rangle \\ &= \text{tr}(A^T U_Q U_Q^T A) + \text{tr}(A^T U_S U_S^T A) + \text{tr}(B^T B) + 2 \langle U_S^T A, B \rangle \\ &= \|U_Q^T A\|^2 + \|U_S^T A\|^2 + \|B\|^2 + 2 \langle U_S^T A, B \rangle \\ &= \|U_Q^T A\|^2 + \|U_S^T A + B\|^2. \end{aligned}$$

F.1.8 Proof of Lemma 17

Recall that $\Sigma^{1/2} = \Sigma_{YX} \Sigma_{XX}^{-1} X$ has a Singular Value Decomposition $\Sigma^{1/2} = U \Delta V^T$ (see (1)). Hence, we have $\Sigma^{1/2} V = U \Delta$ and therefore $\Sigma^{1/2} V_Q = U_Q \Delta^{(Q)}$, where $\Delta^{(Q)}$ is defined in (72). As a consequence,

$$\begin{aligned} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X V_Q &= U_S^T \Sigma^{1/2} V_Q \\ &= U_S^T U_Q \Delta^{(Q)} \\ &= 0, \end{aligned}$$

where the last equality follows from Lemma 6. Finally, we obtain for any $A \in \mathbb{R}^{n \times r}$, $B \in \mathbb{R}^{n \times (d_y - r)}$

$$\begin{aligned} \langle AU_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X, B V_Q^T \rangle &= \text{tr}(AU_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X V_Q B^T) \\ &= 0. \end{aligned}$$

F.1.9 Proof of Lemma 18

Using Lemma 15, we have $\Sigma_{XY} U_Q = X V_Q \Delta^{(Q)}$, then replacing this formula in (75) with $i = q + 1$, we have

$$Z_q \cdots Z_2 Z_1 X V_Q \Delta^{(Q)} = 0.$$

Since $\Delta^{(Q)}$ is diagonal and its diagonal elements are non-zero, it is invertible, hence

$$Z_q \cdots Z_2 Z_1 X V_Q = 0.$$

Finally, for any matrices $A \in \mathbb{R}^{n \times (d_q - r)}$ and $B \in \mathbb{R}^{n \times (d_y - r)}$, we have

$$\begin{aligned} \langle AZ_q \cdots Z_2 Z_1 X, B V_Q^T \rangle &= \text{tr}(AZ_q \cdots Z_2 Z_1 X V_Q B^T) \\ &= 0. \end{aligned}$$

F.1.10 Proof of Lemma 19

Recall that $\Delta^{(S)}$ is defined in (71) and $\Sigma = U \Lambda U^T$. Let $A \in \mathbb{R}^{(d_y - r) \times r}$, we have

$$\begin{aligned} \|AU_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X\|^2 &= \text{tr}(AU_S^T \Sigma U_S A^T) \\ &= \text{tr}(A \text{diag}(\lambda_1, \dots, \lambda_r) A^T) \\ &= \|A \Delta^{(S)}\|^2 \\ &= \sum_{a=1}^r \sum_{b=r+1}^{d_y} \lambda_a (A_{b-r,a})^2 \\ &= \sum_{a=1}^r \sum_{b=r+1}^{d_y} (\lambda_a - \lambda_b) (A_{b-r,a})^2 + \sum_{a=1}^r \sum_{b=r+1}^{d_y} \lambda_b (A_{b-r,a})^2 \\ &= \sum_{a=1}^r \sum_{b=r+1}^{d_y} (\lambda_a - \lambda_b) (A_{b-r,a})^2 + \|\Delta^{(Q)} A\|^2. \end{aligned}$$

F.1.11 Proof of Lemma 20

Let $\mathbf{W} = (W_H, \dots, W_1)$ be a first-order critical point associated with \mathcal{S} verifying the hypotheses of Proposition 11 and let $A \in \mathbb{R}^{d_y \times d_x}$. Using (11), (9), and Lemma 6, we have

$$\langle AX, W_H \cdots W_1 X - Y \rangle = \langle A, W_H \cdots W_1 X X^T - Y X^T \rangle$$

$$\begin{aligned}
&= \langle A, U_S U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X X^T - \Sigma_{YX} \rangle \\
&= \langle A, U_S U_S^T \Sigma_{YX} - \Sigma_{YX} \rangle \\
&= \langle A, -U_Q U_Q^T \Sigma_{YX} \rangle .
\end{aligned}$$

F.1.12 Proof of Lemma 21

Let $\mathbf{W} = (W_H, \dots, W_1)$ be a tightened first-order critical point satisfying the hypotheses of Proposition 11, and $r, \mathcal{S}, Q, (Z_h)_{h=1..H}$ defined as in Proposition 11. Since \mathbf{W} satisfies the hypotheses of Proposition 11, we are going to use all the equations (8), (9), (10) and (11) defined by these hypotheses and (73), (74), (75) and (76) of Lemma 14.

Let $\mathbf{W}' = (W'_H, \dots, W'_1)$ and $i \in \llbracket 1, H \rrbracket$. Recall that T_i is defined in (12) and $J_1 = \llbracket p, H-1 \rrbracket$, $J_2 = \llbracket q+1, p-1 \rrbracket$, $J_3 = \llbracket 2, q \rrbracket$, where p and q are defined as in Lemma 14.

Consider the case $i = H$.

Substituting (74) and (9) in (12), we have

$$\begin{aligned}
T_H &= W'_H (W_{H-1} \cdots W_2) W_1 X \\
&= W'_H \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_1 \end{bmatrix} X \\
&= W'_H \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ 0 \end{bmatrix} X \\
&= (W'_H)_{:,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X .
\end{aligned}$$

This proves (77).

Consider now the case $i \in J_1$.

Substituting (8), (10), (74) and (9), in (12), we have, for $i \in J_1$

$$\begin{aligned}
T_i &= W_H (W_{H-1} \cdots W_{i+1}) W'_i (W_{i-1} \cdots W_2) W_1 X \\
&= [U_S, U_Q Z_H] \begin{bmatrix} I_r & 0 \\ 0 & Z_{H-1} \end{bmatrix} \cdots \begin{bmatrix} I_r & 0 \\ 0 & Z_{i+1} \end{bmatrix} W'_i \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_1 \end{bmatrix} X \\
&= [U_S, U_Q Z_H Z_{H-1} \cdots Z_{i+1}] W'_i \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ 0 \end{bmatrix} X \\
&= [U_S, U_Q Z_H Z_{H-1} \cdots Z_{i+1}] (W'_i)_{:,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X \\
&= U_S (W'_i)_{1:r,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + U_Q Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X .
\end{aligned}$$

Note that the above calculations are still valid in the case $i = H-1$. In this case using the convention in Section 2,

$W_{H-1} \cdots W_{i+1} = I_{d_{H-1}}$ and $Z_{H-1} \cdots Z_{i+1} = I_{d_{H-1}-r}$.

This proves (78).

Consider now the case $i \in J_2 \cup J_3 = \llbracket 2, p-1 \rrbracket$.

Substituting (73), (10) and (9), in (12), we have, for $i \in J_2 \cup J_3$,

$$\begin{aligned}
T_i &= (W_H \cdots W_{i+1}) W'_i (W_{i-1} \cdots W_2) W_1 X \\
&= \begin{bmatrix} U_S, 0 \end{bmatrix} W'_i \begin{bmatrix} I_r & 0 \\ 0 & Z_{i-1} \end{bmatrix} \cdots \begin{bmatrix} I_r & 0 \\ 0 & Z_2 \end{bmatrix} \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_1 \end{bmatrix} X \\
&= U_S (W'_i)_{1:r, \cdot} \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_{i-1} \cdots Z_2 Z_1 \end{bmatrix} X \\
&= U_S (W'_i)_{1:r,1:r} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} X + U_S (W'_i)_{1:r,r+1:d_{i-1}} Z_{i-1} \cdots Z_2 Z_1 X .
\end{aligned}$$

Note that the above calculations are still valid in the case $i = 2$. In this case, using the conventions of Section 2, $W_{i-1} \cdots W_2 = I_{d_1}$ and $Z_{i-1} \cdots Z_2 = I_{d_2-r}$. This proves (79).

Consider finally the case $i = 1$.
Substituting (73) in (12), we have

$$\begin{aligned} T_1 &= (W_H \cdots W_2) W_1' X \\ &= [U_S, 0] W_1' X \\ &= U_S (W_1')_{1:r, \cdot} X. \end{aligned}$$

This proves (80).

Note that, using the conventions of Section 2, the proof still holds for $r = 0$. In this case, $T_i = 0, \forall i$. This concludes the proof.

F.1.13 Proof of Lemma 22

Let $\mathbf{W} = (W_H, \dots, W_1)$ be a tightened first-order critical point satisfying the hypotheses of Proposition 11, and $r, \mathcal{S}, Q, (Z_h)_{h=1..H}$ defined as in Proposition 11. Since \mathbf{W} satisfies the hypotheses of Proposition 11, we are going to use all the equations (8), (9), (10) and (11) defined by these hypotheses and (73), (74), (75) and (76) of Lemma 14. Let $\mathbf{W}' = (W_H', \dots, W_1')$ and $(i, j) \in \llbracket 1, H \rrbracket^2$, with $i > j$. Recall that $T_{i,j}$ is defined in (13) and $J_1 = \llbracket p, H-1 \rrbracket$, $J_2 = \llbracket q+1, p-1 \rrbracket$, $J_3 = \llbracket 2, q \rrbracket$, where p and q are defined as in Lemma 14.

Consider the case $i \in \{H\} \cup J_1$ and $j \in J_1 \cup J_2$ with $i > j$.
Applying Lemma 20 to (13) and using (10) and (9), we obtain

$$\begin{aligned} T_{i,j} &= \langle W_H \cdots W_{i+1} W_i' W_{i-1} \cdots W_{j+1} W_j' W_{j-1} \cdots W_1 X, W_H \cdots W_1 X - Y \rangle \\ &= \langle W_H \cdots W_{i+1} W_i' W_{i-1} \cdots W_{j+1} W_j' W_{j-1} \cdots W_1, -U_Q U_Q^T \Sigma_{YX} \rangle \\ &= -\text{tr} (W_H \cdots W_{i+1} W_i' W_{i-1} \cdots W_{j+1} W_j' W_{j-1} \cdots W_1 \Sigma_{XY} U_Q U_Q^T) \\ &= -\text{tr} \left((W_H \cdots W_{i+1} W_i' W_{i-1} \cdots W_{j+1} W_j' \begin{bmatrix} U_S^T \Sigma U_Q \\ Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q \end{bmatrix} U_Q^T) \right). \end{aligned}$$

Using Lemma 7 and since $j \geq q+1$, using (75), we obtain

$$T_{i,j} = 0.$$

This proves (89) and (92).

Consider now the case $i = H$ and $j \in J_3$.
Applying Lemma (20) to (13) and using (76), (10) and (9), we obtain

$$\begin{aligned} T_{H,j} &= \langle W_H' W_{H-1} \cdots W_{j+1} W_j' W_{j-1} \cdots W_1 X, W_H \cdots W_1 X - Y \rangle \\ &= \langle W_H' W_{H-1} \cdots W_{j+1} W_j' W_{j-1} \cdots W_1, -U_Q U_Q^T \Sigma_{YX} \rangle \\ &= \left\langle W_H' \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix} W_j' \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_{j-1} \cdots Z_2 Z_1 \end{bmatrix}, U_Q U_Q^T \Sigma_{YX} \right\rangle \\ &= -\text{tr} \left(W_H' \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix} W_j' \begin{bmatrix} U_S^T \Sigma U_Q U_Q^T \\ Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q U_Q^T \end{bmatrix} \right). \end{aligned}$$

Using Lemma 15, Lemma 7 and the cyclic property of the trace, we have

$$T_{H,j} = -\text{tr} \left((W_H')_{\cdot, 1:r} (W_j')_{1:r, \cdot} \begin{bmatrix} 0 \\ Z_{j-1} \cdots Z_2 Z_1 X V_Q \Delta^{(Q)} U_Q^T \end{bmatrix} \right)$$

$$\begin{aligned}
&= -\text{tr} \left((W'_H)_{.,1:r} (W'_j)_{1:r,r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 X V_Q \Delta^{(Q)} U_Q^T \right) \\
&= -\text{tr} \left(\Delta^{(Q)} U_Q^T (W'_H)_{.,1:r} (W'_j)_{1:r,r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 X V_Q \right) \\
&= -\left\langle \Delta^{(Q)} U_Q^T (W'_H)_{.,1:r}, ((W'_j)_{1:r,r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 X V_Q)^T \right\rangle.
\end{aligned}$$

This proves (87).

Consider now the case $i = H$ and $j = 1$.

Applying Lemma 20 to (13) and using (76) and Lemma 15, we obtain

$$\begin{aligned}
T_{H,1} &= \langle W'_H W_{H-1} \cdots W_2 W'_1 X, W_H \cdots W_1 X - Y \rangle \\
&= \langle W'_H W_{H-1} \cdots W_2 W'_1, -U_Q U_Q^T \Sigma_{YX} \rangle \\
&= -\left\langle W'_H \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix} W'_1, U_Q (X V_Q \Delta^{(Q)})^T \right\rangle \\
&= -\left\langle (W'_H)_{.,1:r} (W'_1)_{1:r,}, U_Q \Delta^{(Q)} V_Q^T X^T \right\rangle \\
&= -\left\langle \Delta^{(Q)} U_Q^T (W'_H)_{.,1:r}, V_Q^T X^T ((W'_1)_{1:r,})^T \right\rangle \\
&= -\left\langle \Delta^{(Q)} U_Q^T (W'_H)_{.,1:r}, ((W'_1)_{1:r,} X V_Q)^T \right\rangle.
\end{aligned}$$

This proves (88).

Consider now the case $i \in J_1$ and $j \in J_3$.

Applying Lemma 20 to (13) and using (8), (10) and (9), we obtain

$$\begin{aligned}
T_{i,j} &= \langle W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_{j+1} W'_j W_{j-1} \cdots W_1 X, W_H \cdots W_1 X - Y \rangle \\
&= \langle W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_{j+1} W'_j W_{j-1} \cdots W_1, -U_Q U_Q^T \Sigma_{YX} \rangle \\
&= -\left\langle [U_S, U_Q Z_H Z_{H-1} \cdots Z_{i+1}] W'_i W_{i-1} \cdots W_{j+1} W'_j \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_{j-1} \cdots Z_2 Z_1 \end{bmatrix}, U_Q U_Q^T \Sigma_{XY} \right\rangle \\
&= -\text{tr} \left([U_S, U_Q Z_H Z_{H-1} \cdots Z_{i+1}] W'_i W_{i-1} \cdots W_{j+1} W'_j \begin{bmatrix} U_S^T \Sigma \\ Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} \end{bmatrix} U_Q U_Q^T \right) \\
&= -\text{tr} \left([U_Q^T U_S, U_Q^T U_Q Z_H Z_{H-1} \cdots Z_{i+1}] W'_i W_{i-1} \cdots W_{j+1} W'_j \begin{bmatrix} U_S^T \Sigma U_Q \\ Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q \end{bmatrix} \right).
\end{aligned}$$

Using Lemma 6 and Lemma 7, we have

$$\begin{aligned}
T_{i,j} &= -\text{tr} \left([0, Z_H Z_{H-1} \cdots Z_{i+1}] W'_i W_{i-1} \cdots W_{j+1} W'_j \begin{bmatrix} 0 \\ Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q \end{bmatrix} \right) \\
&= -\text{tr} (Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,} W_{i-1} \cdots W_{j+1} (W'_j)_{.,r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q). \quad (96)
\end{aligned}$$

Here, since \mathbf{W} is tightened, taking the tightened pivot (i, j) we have two possible cases: either $\text{rk}(W_{i-1} \cdots W_{j+1}) = r$ or $\text{rk}(W_{j-1} \cdots W_1 \Sigma_{XY} W_H \cdots W_{i+1}) = r$. We treat the two cases separately.

In the first case, using (10) we have

$$\begin{aligned}
W_{i-1} \cdots W_{j+1} &= \begin{bmatrix} I_r & 0 \\ 0 & Z_{i-1} \end{bmatrix} \cdots \begin{bmatrix} I_r & 0 \\ 0 & Z_{j+1} \end{bmatrix} \\
&= \begin{bmatrix} I_r & 0 \\ 0 & Z_{i-1} \cdots Z_{j+1} \end{bmatrix}.
\end{aligned}$$

Hence, $\text{rk}(W_{i-1} \cdots W_{j+1}) = r$ implies $Z_{i-1} \cdots Z_{j+1} = 0$ and we conclude that

$$W_{i-1} \cdots W_{j+1} = \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix}.$$

Then, using this last equality, (96) becomes

$$\begin{aligned} T_{i,j} &= -\text{tr} \left(Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i, \cdot} \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix} (W'_j)_{\cdot, r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q \right) \\ &= -\text{tr} \left(Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i, 1:r} (W'_j)_{1:r, r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q \right). \end{aligned} \quad (97)$$

In the second case, we have $\text{rk}(W_{j-1} \cdots W_1 \Sigma_{XY} W_H \cdots W_{i+1}) = r$. Let us prove that (97) also holds in this case. Using (10), (9), (8), Lemma 7 and $\mathcal{S} = \llbracket 1, r \rrbracket$, we have

$$\begin{aligned} &W_{j-1} \cdots W_1 \Sigma_{XY} W_H \cdots W_{i+1} \\ &= \begin{bmatrix} U_{\mathcal{S}}^T \Sigma \\ Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} \end{bmatrix} [U_{\mathcal{S}}, U_Q Z_H Z_{H-1} \cdots Z_{i+1}] \\ &= \begin{bmatrix} U_{\mathcal{S}}^T \Sigma U_{\mathcal{S}} & U_{\mathcal{S}}^T \Sigma U_Q Z_H Z_{H-1} \cdots Z_{i+1} \\ Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_{\mathcal{S}} & Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q Z_H Z_{H-1} \cdots Z_{i+1} \end{bmatrix} \\ &= \begin{bmatrix} \text{diag}(\lambda_1, \dots, \lambda_r) & 0 \\ Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_{\mathcal{S}} & Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q Z_H Z_{H-1} \cdots Z_{i+1} \end{bmatrix}. \end{aligned}$$

Therefore since $\text{rk}(W_{j-1} \cdots W_1 \Sigma_{XY} W_H \cdots W_{i+1}) = r$ and for all $i \in \llbracket 1, r \rrbracket$, $\lambda_i \neq 0$, we must have

$$Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q Z_H Z_{H-1} \cdots Z_{i+1} = 0. \quad (98)$$

Using the above equation, and the cyclic property of the trace, (96) becomes

$$\begin{aligned} T_{i,j} &= -\text{tr} \left(Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i, \cdot} W_{i-1} \cdots W_{j+1} (W'_j)_{\cdot, r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q \right) \\ &= -\text{tr} \left(Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i, \cdot} W_{i-1} \cdots W_{j+1} (W'_j)_{\cdot, r+1:d_{j-1}} \right) \\ &= 0. \end{aligned}$$

We can use (98) again to write the equation $T_{i,j} = 0$ in the format of equation (97). Indeed, we have

$$\begin{aligned} &-\text{tr} \left(Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i, 1:r} (W'_j)_{1:r, r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q \right) \\ &= -\text{tr} \left(Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i, 1:r} (W'_j)_{1:r, r+1:d_{j-1}} \right) \\ &= 0 \\ &= T_{i,j}. \end{aligned}$$

Therefore, in both cases we have

$$T_{i,j} = -\text{tr} \left(Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i, 1:r} (W'_j)_{1:r, r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 \Sigma_{XY} U_Q \right).$$

Using Lemma 15, it becomes

$$\begin{aligned} T_{i,j} &= -\text{tr} \left(Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i, 1:r} (W'_j)_{1:r, r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 X V_Q \Delta^{(Q)} \right) \\ &= -\text{tr} \left(\Delta^{(Q)} Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i, 1:r} (W'_j)_{1:r, r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 X V_Q \right) \\ &= -\left\langle \Delta^{(Q)} Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i, 1:r}, \left((W'_j)_{1:r, r+1:d_{j-1}} Z_{j-1} \cdots Z_2 Z_1 X V_Q \right)^T \right\rangle. \end{aligned}$$

This proves (90).

Consider now the case $i \in J_1$ and $j = 1$.
Using Lemma 20 to simplify (13), we have

$$\begin{aligned} T_{i,1} &= \langle W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_2 W'_1 X, W_H \cdots W_1 X - Y \rangle \\ &= \langle W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_2 W'_1, -U_Q U_Q^T \Sigma_{YX} \rangle. \end{aligned}$$

Using Lemma 15 and substituting (8), (10), and since $i \geq p$, using (74), this becomes

$$\begin{aligned} T_{i,1} &= - \left\langle [U_S, U_Q Z_H Z_{H-1} \cdots Z_{i+1}] W'_i \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix} W'_1, U_Q (XV_Q \Delta^{(Q)})^T \right\rangle \\ &= - \left\langle \Delta^{(Q)} [U_Q^T U_S, U_Q^T U_Q Z_H Z_{H-1} \cdots Z_{i+1}] (W'_i)_{:,1:r} (W'_1)_{1:r,}, (XV_Q)^T \right\rangle. \end{aligned}$$

Using Lemma 6, it becomes

$$\begin{aligned} T_{i,1} &= - \left\langle \Delta^{(Q)} [0, Z_H Z_{H-1} \cdots Z_{i+1}] (W'_i)_{:,1:r} (W'_1)_{1:r,}, (XV_Q)^T \right\rangle \\ &= - \left\langle \Delta^{(Q)} Z_H Z_{H-1} \cdots Z_{i+1} (W'_i)_{r+1:d_i,1:r}, ((W'_1)_{1:r,} XV_Q)^T \right\rangle. \end{aligned}$$

This proves (91).

Consider now the case $i \in J_2 \cup J_3 = \llbracket 2, p-1 \rrbracket$ and $j < i$.
Applying Lemma 20 to (13) and, since $i < p$, using (73), we obtain

$$\begin{aligned} T_{i,j} &= \langle W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_{j+1} W'_j W_{j-1} \cdots W_1 X, W_H \cdots W_1 X - Y \rangle \\ &= \langle W_H \cdots W_{i+1} W'_i W_{i-1} \cdots W_{j+1} W'_j W_{j-1} \cdots W_1, -U_Q U_Q^T \Sigma_{YX} \rangle \\ &= -\text{tr}([U_S, 0] W'_i W_{i-1} \cdots W_{j+1} W'_j W_{j-1} \cdots W_1 \Sigma_{XY} U_Q U_Q^T) \end{aligned}$$

The cyclic property of the trace and Lemma 6 lead to

$$\begin{aligned} T_{i,j} &= -\text{tr}([U_Q^T U_S, 0] W'_i W_{i-1} \cdots W_{j+1} W'_j W_{j-1} \cdots W_1 \Sigma_{XY} U_Q) \\ &= 0. \end{aligned}$$

This proves (93) and concludes the proof.

Note that, with the convention of Section 2, the proof still holds for $r = 0$. In this case, $T_{i,j} = 0, \forall i > j$.

F.2 Proof of Proposition 10

Let $\mathbf{W} = (W_H, \dots, W_1)$ be a tightened first-order critical point associated with $\mathcal{S} = \llbracket 1, r \rrbracket$ with $r < r_{max}$. Then, using Proposition 5 there exist invertible matrices $D_{H-1} \in \mathbb{R}^{d_{H-1} \times d_{H-1}}, \dots, D_1 \in \mathbb{R}^{d_1 \times d_1}$ and matrices $Z_H \in \mathbb{R}^{(d_y-r) \times (d_{H-1}-r)}, Z_1 \in \mathbb{R}^{(d_1-r) \times d_x}$ and $Z_h \in \mathbb{R}^{(d_h-r) \times (d_{h-1}-r)}$ for $h \in \llbracket 2, H-1 \rrbracket$ such that if we denote $\widetilde{W}_H = W_H D_{H-1}, \widetilde{W}_1 = D_1^{-1} W_1$ and $\widetilde{W}_h = D_h^{-1} W_h D_{h-1}$ for all $h \in \llbracket 2, H-1 \rrbracket$, and $\widetilde{\mathbf{W}} = (\widetilde{W}_H, \dots, \widetilde{W}_1)$, then

$$\begin{aligned} \widetilde{W}_H &= [U_S, U_Q Z_H] \\ \widetilde{W}_1 &= \begin{bmatrix} U_S^T \Sigma_{YX} \Sigma_{XX}^{-1} \\ Z_1 \end{bmatrix} \\ \widetilde{W}_h &= \begin{bmatrix} I_r & 0 \\ 0 & Z_h \end{bmatrix} \quad \forall h \in \llbracket 2, H-1 \rrbracket \\ \widetilde{W}_H \cdots \widetilde{W}_2 &= [U_S, 0]. \end{aligned}$$

Then, due to Lemma 2, and since \mathbf{W} is a first-order critical point, we have that $\widetilde{\mathbf{W}}$ is a first-order critical point. We also have $\widetilde{W}_H \cdots \widetilde{W}_1 = W_H \cdots W_1$. Hence, according to Proposition 1 $\widetilde{\mathbf{W}}$ is also associated with \mathcal{S} .

Since \mathbf{W} is tightened and multiplication by invertible matrices does not change the rank, $\tilde{\mathbf{W}}$ is also tightened. Hence, $\tilde{\mathbf{W}}$ satisfies the hypotheses of Proposition 11 and therefore is a second-order critical point. Finally, using Lemma 2, we conclude that \mathbf{W} is a second-order critical point. Since $r < r_{max}$ and Σ is invertible (Lemma 4), using Proposition 1, we have

$$L(\mathbf{W}) = \text{tr}(\Sigma_{YY}) - \sum_{i=1}^r \lambda_i > \text{tr}(\Sigma_{YY}) - \sum_{i=1}^{r_{max}} \lambda_i .$$

Therefore, \mathbf{W} is not a global minimizer, hence \mathbf{W} is a non-strict saddle point.

G A simple illustrative experiment

Next we provide more details on the experiment whose results were plotted in Figures 3 and 4. The goal is to illustrate the behavior of the ADAM optimizer in the vicinity of strict or non-strict saddle points.

Experimental setting. We optimize a linear neural network starting in the vicinity either of a strict saddle point (10000 runs in total) or of a non-strict saddle point (10000 runs in total). For each run, the setting is the following:

- Network architecture: $d_x = 10$, $d_y = 4$, $H = 5$ and $d_4 = d_3 = d_2 = d_1 = 10$.
- Data construction: $m = 100$ i.i.d. data points $(x_1, y_1), \dots, (x_m, y_m) \in \mathbb{R}^{d_x} \times \mathbb{R}^{d_y}$ such that, for all $i = 1, \dots, m$, the points x_i and y_i are drawn independently at random from the Gaussian distributions $\mathcal{N}(0, I_{d_x})$ and $\mathcal{N}(0, I_{d_y})$ respectively.
- Initial iterate: we define it as $(W_1, \dots, W_H) = \mathbf{W}^{cp} + (V_1, \dots, V_H)$, for a critical point \mathbf{W}^{cp} (defined later) and a random perturbation (V_1, \dots, V_H) whose components $(V_h)_{i,j}$ are drawn independently from the distributions $\mathcal{N}(0, \sigma_h^2)$, with $\sigma_h = 0.1 \frac{\|\mathbf{W}_h^{cp}\|_F}{\sqrt{d_{h-1} d_h}}$. The critical point \mathbf{W}^{cp} is defined as in (25) in Appendix B.8, for $r = 2$ ($\mathcal{S} = \{1, 2\}$) and

$$Z_h = \begin{cases} I_{d_h-2} & \text{for all } h \in \llbracket 2, 4 \rrbracket, \text{ for runs starting at a strict saddle point;} \\ 0_{(d_h-2) \times (d_h-2)} & \text{for all } h \in \llbracket 2, 4 \rrbracket, \text{ for runs starting at a non-strict saddle point.} \end{cases}$$

Since $d_4 = d_3 = d_2 = d_1$, note that the sizes of the above matrices Z_h are consistent with (25). As explained in Appendix B.8, when $Z_h = I_{d_h-2}$ for all $h \in \llbracket 2, 4 \rrbracket$, the critical point \mathbf{W}^{cp} is non-tightened and therefore Theorem 1 guarantees that it is a strict saddle point. Similarly, when $Z_h = 0_{(d_h-2) \times (d_h-2)}$ for all $h \in \llbracket 2, 4 \rrbracket$, the critical point \mathbf{W}^{cp} is tightened and Theorem 1 guarantees that it is a non-strict saddle point.

- Optimizer: we use the ADAM optimizer of the Keras library, with the default parameters.

Observations. Figure 3 in Section 1.2 shows the evolution of the loss along the optimization process for two representative runs (initialization near a strict or a non-strict saddle point). We can see that, when initialized in the vicinity of the strict saddle point, ADAM rapidly decreases below the initial value $L(\mathbf{W}^{cp})$. On the contrary, ADAM needs many epochs to exit the plateau at the critical value of the non-strict saddle point.

In order to assess the importance of this phenomenon, we repeated the above experiment 10000 times for both strict saddle points and non-strict saddle points. For each run, we define and compute the *escape epoch* as the first epoch such that $L(\mathbf{W}) < L(\mathbf{W}^{cp}) - \frac{\lambda_3}{2}$ (the average of the critical values associated with $\mathcal{S} = \{1, 2\}$ and $\mathcal{S}' = \{1, 2, 3\}$). On Figure 4 (Section 1.2) the histograms of the escape epoch are displayed separately for runs corresponding to strict saddle points (in red) or non-strict saddle points (in blue). We can see that, while ADAM quickly escapes from the vicinity of the strict saddle points, it takes many more epochs to escape from the vicinity of the non-strict saddle points. In the last case, the plateau can easily be confused with a global minimum.

References

- [1] Sanjeev Arora, Nadav Cohen, Noah Golowich, and Wei Hu, *A convergence analysis of gradient descent for deep linear neural networks*, International Conference on Learning Representations, 2019.
- [2] Sanjeev Arora, Nadav Cohen, Wei Hu, and Yuping Luo, *Implicit regularization in deep matrix factorization*, Advances in Neural Information Processing Systems, 2019, pp. 7413–7424.
- [3] Bubacarr Bah, Holger Rauhut, Ulrich Terstiege, and Michael Westdickenberg, *Learning deep linear neural networks: Riemannian gradient flows and convergence to global minimizers*, Information and Inference: A Journal of the IMA (2021), 10.1093/imaiai/iaaa039.
- [4] P. Baldi and K. Hornik, *Neural networks and principal component analysis: Learning from examples without local minima*, Neural Netw. **2** (1989), no. 1, 53–58.
- [5] Peter Bartlett, Dave Helmbold, and Philip Long, *Gradient descent with identity initialization efficiently learns positive definite linear transformations by deep residual networks*, International conference on machine learning, 2018, pp. 521–530.
- [6] Peter L. Bartlett, Philip M. Long, Gábor Lugosi, and Alexander Tsigler, *Benign overfitting in linear regression*, Proceedings of the National Academy of Sciences **117** (2020), no. 48, 30063–30070.
- [7] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal, *Reconciling modern machine-learning practice and the classical bias–variance trade-off*, Proceedings of the National Academy of Sciences **116** (2019), no. 32, 15849–15854.
- [8] Avrim Blum and Ronald L Rivest, *Training a 3-node neural network is np-complete*, Advances in neural information processing systems, 1989, pp. 494–501.
- [9] Peter Bühlmann and Bin Yu, *Boosting with the l2 loss*, Journal of the American Statistical Association **98** (2003), no. 462, 324–339.
- [10] Yacine Chitour, Zhenyu Liao, and Romain Couillet, *A geometric approach of gradient descent algorithms in neural networks*, arXiv preprint arXiv:1811.03568 (2018).
- [11] Lénaïc Chizat and Francis Bach, *On the global convergence of gradient descent for over-parameterized models using optimal transport*, Proceedings of the 32nd International Conference on Neural Information Processing Systems, 2018, pp. 3040–3050.
- [12] Anna Choromanska, Mikael Henaff, Michael Mathieu, Gerard Ben Arous, and Yann LeCun, *The Loss Surfaces of Multilayer Networks*, Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics, vol. 38, 2015, pp. 192–204.
- [13] Hadi Daneshmand, Jonas Kohler, Aurelien Lucchi, and Thomas Hofmann, *Escaping saddles with stochastic gradients*, International Conference on Machine Learning, 2018, pp. 1155–1164.
- [14] Simon Du and Wei Hu, *Width provably matters in optimization for deep linear neural networks*, International Conference on Machine Learning, 2019, pp. 1655–1664.
- [15] Armin Eftekhari, *Training linear neural networks: Non-local convergence and complexity results*, Proceedings of the 37th International Conference on Machine Learning, vol. 119, 2020, pp. 2836–2847.
- [16] Abraham Frandsen and Rong Ge, *Optimization landscape of tucker decomposition*, Mathematical Programming (2020), 1–26.
- [17] Rong Ge, Furong Huang, Chi Jin, and Yang Yuan, *Escaping from saddle points—online stochastic gradient for tensor decomposition*, Conference on Learning Theory, 2015, pp. 797–842.

- [18] Rong Ge and Tengyu Ma, *On the optimization landscape of tensor decompositions*, *Mathematical Programming* (2020), 1–47.
- [19] Gauthier Gidel, Francis Bach, and Simon Lacoste-Julien, *Implicit regularization of discrete gradient dynamics in linear neural networks*, *Advances in Neural Information Processing Systems*, 2019, pp. 3202–3211.
- [20] Daniel Gissin, Shai Shalev-Shwartz, and Amit Daniely, *The implicit bias of depth: How incremental learning drives generalization*, *International Conference on Learning Representations*, 2019.
- [21] Benjamin D Haeffele and René Vidal, *Global optimality in neural network training*, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 7331–7339.
- [22] Moritz Hardt and Tengyu Ma, *Identity matters in deep learning*, *5th International Conference on Learning Representations*, 2017.
- [23] Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J. Tibshirani, *Surprises in high-dimensional ridgeless least squares interpolation*, 2020, arXiv:1903.08560.
- [24] Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M Kakade, and Michael I Jordan, *How to escape saddle points efficiently*, *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 2017, pp. 1724–1732.
- [25] Chi Jin, Praneeth Netrapalli, Rong Ge, Sham M. Kakade, and Michael I. Jordan, *On nonconvex optimization for machine learning: Gradients, stochasticity, and saddle points*, *J. ACM* **68** (2021), no. 2.
- [26] Chi Jin, Praneeth Netrapalli, and Michael I Jordan, *Accelerated gradient descent escapes saddle points faster than gradient descent*, *Conference On Learning Theory*, 2018, pp. 1042–1085.
- [27] Kenji Kawaguchi, *Deep learning without poor local minima*, *Advances in Neural Information Processing Systems* **29** (2016), 586–594.
- [28] Thomas Laurent and James von Brecht, *Deep linear networks with arbitrary loss: All local minima are global*, *ICML*, 2018, pp. 2908–2913.
- [29] Jason D Lee, Ioannis Panageas, Georgios Piliouras, Max Simchowitz, Michael I Jordan, and Benjamin Recht, *First-order methods almost always avoid strict saddle points*, *Mathematical programming* **176** (2019), no. 1-2, 311–337.
- [30] Jason D. Lee, Max Simchowitz, Michael I. Jordan, and Benjamin Recht, *Gradient descent only converges to minimizers*, *Proceedings of the 29th Conference on Learning Theory*, vol. 49, 2016, pp. 1246–1257.
- [31] Haihao Lu and Kenji Kawaguchi, *Depth creates no bad local minima*, *arXiv preprint arXiv:1702.08580* (2017).
- [32] Dhagash Mehta, Tianran Chen, Tingting Tang, and Jonathan Hauenstein, *The loss surface of deep linear networks viewed through the algebraic geometry lens*, *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021), 10.1109/TPAMI.2021.3071289.
- [33] Song Mei, Andrea Montanari, and Phan-Minh Nguyen, *A mean field view of the landscape of two-layer neural networks*, *Proceedings of the National Academy of Sciences* **115** (2018), no. 33, E7665–E7671.
- [34] Igor Molybog, Somayeh Sojoudi, and Javad Lavaei, *Role of sparsity and structure in the optimization landscape of non-convex matrix sensing*, *Mathematical Programming* (2020), 1–37.
- [35] Katta G Murty and Santosh N Kabadi, *Some np-complete problems in quadratic and nonlinear programming*, *Mathematical Programming* **39** (1987), no. 2, 117–129.
- [36] Yurii Nesterov, *Introductory lectures on convex programming volume i: Basic course*, *Lecture notes* **3** (1998), no. 4, 5.

- [37] Quynh Nguyen, *On connected sublevel sets in deep learning*, International Conference on Machine Learning, 2019, pp. 4790–4799.
- [38] Quynh Nguyen and Matthias Hein, *The loss surface of deep and wide neural networks*, International conference on machine learning, 2017, pp. 2603–2612.
- [39] Maher Nouiehed and Meisam Razaviyayn, *Learning deep models: Critical points and local openness*, 6th International Conference on Learning Representations, 2018.
- [40] Lutz Prechelt, *Early stopping-but when?*, Neural Networks: Tricks of the trade, Springer, 1998, pp. 55–69.
- [41] G. Raskutti, M. J. Wainwright, and B. Yu, *Early stopping for non-parametric regression: An optimal data-dependent stopping rule*, 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton), 2011, pp. 1318–1325.
- [42] Noam Razin and Nadav Cohen, *Implicit regularization in deep learning may not be explainable by norms*, Advances in Neural Information Processing Systems, vol. 33, 2020, pp. 21174–21187.
- [43] Andrew M. Saxe, James L. McClelland, and Surya Ganguli, *Exact solutions to the nonlinear dynamics of learning in deep linear neural networks*, 2nd International Conference on Learning Representations, 2014.
- [44] ———, *A mathematical theory of semantic development in deep neural networks*, Proceedings of the National Academy of Sciences **116** (2019), no. 23, 11537–11546.
- [45] Ohad Shamir, *Exponential convergence time of gradient descent for one-dimensional deep linear neural networks*, Conference on Learning Theory, 2019, pp. 2691–2713.
- [46] Ju Sun, Qing Qu, and John Wright, *Complete dictionary recovery over the sphere i: Overview and the geometric picture*, IEEE Transactions on Information Theory **63** (2016), no. 2, 853–884.
- [47] ———, *A geometric analysis of phase retrieval*, Foundations of Computational Mathematics **18** (2018), no. 5, 1131–1198.
- [48] Ruoyu Sun, *Optimization for deep learning: theory and algorithms*, arXiv preprint arXiv:1912.08957 (2019).
- [49] Ruoyu Sun, Dawei Li, Shiyu Liang, Tian Ding, and Rayadurgam Srikant, *The global landscape of neural networks: An overview*, IEEE Signal Processing Magazine **37** (2020), no. 5, 95–108.
- [50] Matthew Trager, Kathlén Kohn, and Joan Bruna, *Pure and spurious critical points: a geometric study of linear networks*, International Conference on Learning Representations, 2020.
- [51] Lei Wu, Qingcan Wang, and Chao Ma, *Global convergence of gradient descent for deep linear residual networks*, Advances in Neural Information Processing Systems, 2019, pp. 13389–13398.
- [52] Chulhee Yun, Suvrit Sra, and Ali Jadbabaie, *Global optimality conditions for deep neural networks*, International Conference on Learning Representations, 2018.
- [53] Tong Zhang and Bin Yu, *Boosting with early stopping: Convergence and consistency*, Ann. Statist. **33** (2005), no. 4, 1538–1579.
- [54] Yuqian Zhang, Qing Qu, and John Wright, *From symmetry to geometry: Tractable nonconvex problems*, arXiv preprint arXiv:2007.06753 (2020).
- [55] Yi Zhou and Yingbin Liang, *Critical points of linear neural networks: Analytical forms and landscape properties*, International Conference on Learning Representations, 2018.
- [56] Zhihui Zhu, Daniel Soudry, Yonina C Eldar, and Michael B Wakin, *The global optimization geometry of shallow linear neural networks*, Journal of Mathematical Imaging and Vision **62** (2020), no. 3, 279–292.