

Nonlinear matrix recovery using optimization on the Grassmann manifold

Florentin Goyens^{*1,2}, Coralia Cartis^{†1,2}, and Armin Eftekhari^{2,3}

¹*Mathematical Institute, University of Oxford, Oxford, United Kingdom*

²*The Alan Turing Institute, London, United Kingdom*

³*Department of Mathematics and Mathematical Statistics, Umeå University, Umeå, Sweden*

December 8, 2022

Abstract

We investigate the problem of recovering a partially observed high-rank matrix whose columns obey a nonlinear structure such as a union of subspaces, an algebraic variety or grouped in clusters. The recovery problem is formulated as the rank minimization of a nonlinear feature map applied to the original matrix, which is then further approximated by a constrained non-convex optimization problem involving the Grassmann manifold. We propose two sets of algorithms, one arising from Riemannian optimization and the other as an alternating minimization scheme, both of which include first- and second-order variants. Both sets of algorithms have theoretical guarantees. In particular, for the alternating minimization, we establish global convergence and worst-case complexity bounds. Additionally, using the Kurdyka-Lojasiewicz property, we show that the alternating minimization converges to a unique limit point. We provide extensive numerical results for the recovery of union of subspaces and clustering under entry sampling and dense Gaussian sampling. Our methods are competitive with existing approaches and, in particular, high accuracy is achieved in the recovery using Riemannian second-order methods.

Keywords: nonlinear matrix recovery, nonconvex optimization, Riemannian optimization, second-order methods.

1 Introduction

In the matrix recovery problem, one tries to estimate a matrix $M \in \mathbb{R}^{n \times s}$ from partial information. The *low-rank matrix recovery problem* deals with instances where the matrix M is low-rank. This problem has received great attention in the literature, as applications abound in recommender systems and engineering (see [12] and the references therein for an overview). It was shown in [9] that solving a convex program allows to recover the original matrix M with very high probability, provided enough samples are available. However, solving this convex semi-definite problem for large instances is very costly in time and memory allocation. This has sparked the search for alternative nonconvex formulations of the problem [27, 40]. Riemannian optimization methods are used in some of the most efficient algorithm known to date for low-rank matrix completion. These methods solve optimization problems defined on smooth Riemannian manifolds, such as the manifold of fixed-rank matrices [39] or the Grassmann manifold [5].

*This author's work was supported by The Alan Turing Institute under the EPSRC Grant No. EP/N510129/1 and under the Turing Project Scheme.

†This author's work was supported by The Alan Turing Institute under the Turing Project Scheme.

All traditional approaches to matrix completion fail if the matrix M is high-rank. Our work is based on the recent discovery that an adaptation of traditional methods allows to recover specific classes of high-rank matrices [17, 32]. This problem is known as *nonlinear matrix recovery* (or high-rank matrix recovery). Recovering high-rank matrices requires one to make assumptions on the structure of M . Let m_1, \dots, m_s denote the columns of M . When the points $m_i \in \mathbb{R}^n$ belong to a low-dimensional subspace in \mathbb{R}^n , low-rank matrix recovery methods can be applied. Nonlinear matrix recovery attempts to recover M when the points m_i are related in a nonlinear way.

Classically, for some integer $m < ns$, the matrix M satisfies m linear equations of the type $\langle A_i, M \rangle = b_i$ for given matrices $A_i \in \mathbb{R}^{n \times s}$ and a given vector $b \in \mathbb{R}^m$, where we use the usual inner product $\langle A_i, M \rangle = \text{trace}(A_i^\top M)$. The matrices A_i are assumed to be randomly drawn from a known distribution. One defines the linear operator

$$\mathcal{A} : \mathbb{R}^{n \times s} \rightarrow \mathbb{R}^m \text{ where } \mathcal{A}(M)_i = \langle A_i, M \rangle \quad (1.1)$$

so as to have the compact notation $\mathcal{A}(M) = b$ for the measurements. When each matrix A_i has exactly one non-zero entry which is equal to 1, this is known as a matrix completion problem. The matrix M is then known on a subset Ω of the complete set of entries $\{1, \dots, n\} \times \{1, \dots, s\}$. Without loss of generality we assume $n \leq s$.

Problem description Nonlinear matrix recovery methods use features that map the columns of M to a space of higher dimension. The *feature map* is defined as

$$\varphi : \mathbb{R}^n \rightarrow \mathcal{F} : v \mapsto \varphi(v), \quad (1.2)$$

where \mathcal{F} is a Hilbert space. If \mathcal{F} is finite dimensional, we write $\mathcal{F} = \mathbb{R}^N$ where N is the dimension of the feature space, with $N \geq n$. We obtain the *feature matrix* $\Phi(M)$ by applying φ to each column of M ,

$$\Phi(M) = [\varphi(m_1) \ \dots \ \varphi(m_s)] \in \mathbb{R}^{N \times s}. \quad (1.3)$$

The map φ is chosen using a priori knowledge of the data so that the features of the data points $\varphi(m_i)$ for $i = 1, \dots, s$, all belong to the same subspace in \mathbb{R}^N . The nonlinear structure in M will cause a rank deficiency in the feature matrix $\Phi(M)$, even though M may be full-rank. This is illustrated in Figure 1.




Figure 1: The feature map φ is chosen to exploit the nonlinear structure.

If the features are infinite dimensional or that N is very large, the feature map should be represented using a kernel, which is known as the kernel trick. The set \mathcal{F} is then called a reproducing kernel Hilbert space. The kernel map represents the inner product between elements in the Hilbert space of features,

$$k : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R} : k(x, y) = \langle \varphi(x), \varphi(y) \rangle_{\mathcal{F}}. \quad (1.4)$$

This allows to define the *kernel matrix* of the data $K(M, M) \in \mathbb{R}^{s \times s}$, with $K_{ij}(M, M) = k(m_i, m_j)$. Throughout, we assume that $r = \text{rank}(\Phi(M))$ is known and smaller than $\min(N, s)$.

When $\mathcal{F} = \mathbb{R}^N$, we note that $\mathbf{K}(M, M) = \Phi(M)^\top \Phi(M)$ and therefore $\text{rank}(\mathbf{K}(M, M)) = \text{rank}(\Phi(M))$. We use the term *embedding* to denote a mapping to a higher-dimensional space, which may be performed using a kernel or a feature map. In [32], Ongie et al. use the monomial kernel for the completion of matrices whose columns belong to an algebraic variety (a set defined by a finite number of polynomial equations). This can notably be applied to a union of subspaces. In [17], Fan et al. use the monomial kernel and the Gaussian kernel on image inpainting problems. In Section 2, we detail why polynomial and Gaussian kernels may be used to model data which belongs to, respectively, an algebraic variety or a set of clusters.

Problem formulation For an appropriately chosen feature map, the nonlinear matrix recovery problem can be formulated as the rank minimization of the features matrix under the measurements constraint

$$\begin{cases} \min_{X \in \mathbb{R}^{n \times s}} & \text{rank}(\Phi(X)) \\ & \mathcal{A}(X) = b. \end{cases} \quad (1.5)$$

This seeks to find the matrix which fits the observations using a minimum number of independent features. As is the case for low-rank matrix recovery, minimizing the rank directly is NP-hard and should be avoided [9]. It is necessary to find a suitable approximation to the rank function.

Related work In essence, [17] and [32] apply different minimization algorithms to the Schatten p-norm of the features which is defined by

$$\left(\sum_{i=1}^{\min(N,s)} \sigma_i(\Phi(X))^p \right)^{1/p} \quad \text{for } 0 < p \leq 1. \quad (1.6)$$

When $p = 1$, the sum of the singular values is the nuclear norm. Both [17] and [32] use a kernel representation of the features, so that the features are never computed explicitly. In [16], the authors introduce the algorithm NLMC, which applies a quasi-Newton method to minimize the Schatten p-norm. The Schatten p-norm for $0 < p \leq 1$ is nonsmooth. This has the benefit of encouraging sparsity in the singular values, but it might prevent fast convergence near a minimizer. The algorithm VMC, introduced in [32], minimizes a smooth approximation of the Schatten p-norm. It uses a kernelized version of an iterative reweighted least-squares algorithm (IRLS). The IRLS method was originally proposed in [23] and [29] for low-rank matrix recovery and rank minimization. The IRLS framework has the advantage that it generalizes seamlessly to the kernel setting. In [20] a truncated version of the Schatten norm is proposed, where only the smallest singular values are minimized,

$$\left(\sum_{i=r+1}^{\min(N,s)} \sigma_i(\Phi(X))^p \right)^{1/p} \quad \text{for } 0 < p \leq 1, \quad (1.7)$$

where $r = \text{rank}(\Phi(M))$. They use the kernel trick and propose an algorithm which alternates between truncated singular value decompositions and a step of the Adam method with an additional tuning of the stepsizes.

In [19], the authors propose an extension to handle outliers in the data. This is achieved by introducing a sparse matrix in the model, which absorbs the outliers. In [31], the authors build a tensor representation of the data and apply known matrix completion techniques in the tensor space. Their algorithm, LADMC, is a simple and efficient approach for which they are able to show that the sampling requirements nearly match the information theoretic lower bounds for

recovery under a union of subspace model. This is remarkable as the sampling pattern in the tensor space is not random, and low-rank recovery results do not apply directly. Note that the approach in [31] is only applicable to matrix completion problems, not matrix sensing.

In [18], a new algorithm KFMC (kernelized factorization matrix completion) is proposed, which lends itself to online completion. In this setting, the columns of the matrix M are accessible as a stream and the matrix M is never stored in its entirety. They also develop a variant algorithm to deal with out of samples extensions. That is, how to complete a new column without recomputing the model. The offline formulation applies the kernel trick to

$$\begin{cases} \min_{X,D,Z} & \|\Phi(X) - \Phi(D)Z\|_F^2 + \alpha \|\Phi(D)\|_F^2 + \beta \|Z\|_F^2 \\ & X_{ij} = M_{ij}, (i,j) \in \Omega, \\ & D \in \mathbb{R}^{n \times r}, Z \in \mathbb{R}^{r \times s}, X \in \mathbb{R}^{n \times s}. \end{cases} \quad (1.8)$$

The variable $D \in \mathbb{R}^{n \times r}$ aims to find r points in \mathbb{R}^n so that their features will form a basis for $\Phi(M)$ in the Hilbert space. The last two terms in the objective are added as regularizers to improve the practical performances of the algorithm, as is often done in low-rank matrix completion [12]. In the method DMF, proposed in [16], the embedding is replaced by a deep-structure neural network who is trained to minimize the reconstruction error for the observable entries of M . This work showcases the applicability and performance of nonlinear matrix completion with numerous examples including image inpainting and collaborative filtering problems. For data drawn from multiple subspaces, [21] proposes a sparse factorization where each subspace is represented in a low rank decomposition. They solve this model with an algorithm in the spirit of PALM [3]. They are able to show sampling bounds to guarantee recovery.

Contribution and outline of the paper In Section 2 we describe the approach taken to recover high-rank matrices. It consists in using a feature map (or kernel) that exploits the nonlinear structure present in the matrix. This is applied to data which follows algebraic variety models or grouped in clusters. For these, we respectively use the monomial kernel and the Gaussian kernel. We demonstrate that the Gaussian kernel can be used to perform clustering with missing data, which expands the use cases of nonlinear matrix recovery.

In Section 3, we propose to use a new formulation for nonlinear matrix recovery. We use the feature map to write the recovery problem as a constrained nonconvex optimization problem on the Grassmann manifold. This extends the residual proposed in [15] in the context of low-rank matrix completion to the nonlinear case.

We propose to use Riemannian optimization methods to solve the recovery problem, which is new in the context of nonlinear matrix recovery. Riemannian optimization, as described in Section 4, provides a framework to design algorithms for problems with smooth constraints. This allows to seamlessly choose between standardized first- and second-order methods. The use of second-order methods allows to recover high-rank matrices up to high accuracy if desired.

Section 5 presents an alternating minimization algorithm to solve the recovery problem. First- and second-order variants of the alternating minimization are discussed. We prove global convergence of the algorithm to first-order stationary points in Section 6 and give a global complexity rate to achieve an arbitrary accuracy on the gradient norm from an arbitrary initial guess. In Section 7, we also show convergence of the sequence of iterates to a unique limit point using the Kurdyka-Lojasiewicz property. Our alternating minimization method is a similar approach to the method proposed in [20]. We provide extensive convergence analysis, which was not done in [20].

Section 8 summarizes the applications and algorithms covered in this paper with a framework to solve nonlinear matrix completion.

We conclude with an extensive set of numerical experiments that compare the performances of the optimization approaches and the quality of the solutions that can be obtained (Section 9). We discuss the influence of the complexity of the data and the role of model parameters on the recovery. Moreover, we showcase that our approach is very efficient at clustering data with missing information.

Notations Throughout the paper we use a notation consistent with [2] for the derivative of a function f defined on a Riemannian manifold. The unconstrained gradient of a function f is written ∇f , when the domain of f is extended to an embedding Euclidean space. Conversely, we use $\text{grad}f$ for the Riemannian gradient of f defined over a Riemannian manifold. For matrices $A, B \in \mathbb{R}^{n \times s}$, $\langle A, B \rangle = \text{trace}(A^\top B)$ is the canonical inner product, $\text{range}(A)$ is the column space of A , $\text{null}(A)$ is the null space of A . The identity matrix of size n is denoted by I_n and Id is the identity operator.

2 The feature map

As mentioned, our approach uses an embedding of the original matrix in a space of features, in the spirit of [17][32]. Through the case studies below (2.1, 2.2 and 2.3), we describe the embeddings that we use and some of the data structures to which they apply.

A1. *The feature map φ is chosen such that $\Phi(M)$ is low rank. In addition, $\Phi(X)$ should be high rank if X does not exhibit the same geometrical structure as M .*

The goal is to find an embedding that reveals the nonlinear relation between the points m_i , the columns of M . In [32] the authors use the polynomial features for data sets represented by algebraic varieties.

Case study 2.1 (Algebraic varieties [11]). *Let $\mathbb{R}[x]$ be the set of real valued polynomials over \mathbb{R}^n . A real (affine) algebraic variety is defined as the zero set of a system of polynomials $P \subset \mathbb{R}[x]$:*

$$V(P) = \{x \in \mathbb{R}^n : p(x) = 0 \text{ for all } p \in P\}. \quad (2.1)$$

We say that the matrix M follows an algebraic variety model if every column of M belongs to the same algebraic variety.

Let

$$N(n, d) = \binom{n+d}{n}, \quad (2.2)$$

which reads $n+d$ choose n , the number of monomials of degree d or less that can be formed with n variables. The monomial features φ_d for some degree d are defined as

$$\varphi_d: \mathbb{R}^n \rightarrow \mathbb{R}^{N(n,d)}: \varphi_d(x) = \begin{pmatrix} x^{\alpha^1} \\ x^{\alpha^2} \\ \vdots \\ x^{\alpha^{N(n,d)}} \end{pmatrix} \quad (2.3)$$

where, for $i = 1, 2, \dots, N(n, d)$, the exponent $\alpha^i = (\alpha_1^i, \alpha_2^i, \dots, \alpha_n^i)$ is a multi-index of non-negative integers; so that $x^{\alpha^i} := x_1^{\alpha_1^i} x_2^{\alpha_2^i} \dots x_n^{\alpha_n^i}$ and $\alpha_1^i + \alpha_2^i + \dots + \alpha_n^i \leq d$. The dimension of the feature space $N(n, d)$ increases exponentially in d . Therefore, a kernel implementation

is usually used in practice for moderate and large dimensions, or more precisely, whenever $s \leq N(n, d)$. The monomial kernel of degree d is defined for any $X, Y \in \mathbb{R}^{n \times s}$ as

$$K_d(X, Y) = (X^\top Y + c \mathbf{1}_{s \times s})^{\odot d}, \quad (2.4)$$

where the value $c \in \mathbb{R}$ is a parameter of the kernel, $\mathbf{1}_{s \times s}$ is a square matrix of size s full of ones and \odot is an entry-wise exponent. If the equations describing the variety are known to be homogeneous, one can set $c = 0$. Note that the monomial kernel in (2.4) is not exactly the kernel associated with the monomial features in (2.3). Instead, $K_d(X, X) = \tilde{\Phi}_d(X)^\top \tilde{\Phi}_d(X)$ for a map of monomials $\tilde{\Phi}_d$ that has non-unitary coefficients given by the multinomial theorem. For $x, y \in \mathbb{R}^n$, we have

$$k_d(x, y) = (x^\top y + c)^d = (x_1 y_1 + \dots + x_n y_n + c)^d \quad (2.5)$$

$$= \sum_{\alpha_1^i + \alpha_2^i + \dots + \alpha_{n+1}^i = d} \frac{d!}{\alpha_1^i! \alpha_2^i! \dots \alpha_{n+1}^i!} (x_1 y_1)^{\alpha_1^i} \dots (x_n y_n)^{\alpha_n^i} c^{\alpha_{n+1}^i} \quad (2.6)$$

$$= \sum_{\alpha_1^i + \alpha_2^i + \dots + \alpha_{n+1}^i = d} \left(\frac{\sqrt{d!} x_1^{\alpha_1^i} \dots x_n^{\alpha_n^i} \sqrt{c^{\alpha_{n+1}^i}}}{\sqrt{\alpha_1^i! \alpha_2^i! \dots \alpha_{n+1}^i!}} \right) \left(\frac{\sqrt{d!} y_1^{\alpha_1^i} \dots y_n^{\alpha_n^i} \sqrt{c^{\alpha_{n+1}^i}}}{\sqrt{\alpha_1^i! \alpha_2^i! \dots \alpha_{n+1}^i!}} \right). \quad (2.7)$$

It follows that $k_d(x, y) = \langle \tilde{\varphi}_d(x), \tilde{\varphi}_d(y) \rangle$ for a map $\tilde{\varphi}_d: \mathbb{R}^n \rightarrow \mathbb{R}^{N(n, d)}$ such that the entries of $\tilde{\varphi}_d(x)$ are of the form

$$\sqrt{d!} x_1^{\alpha_1^i} x_2^{\alpha_2^i} \dots x_n^{\alpha_n^i} \sqrt{c^{\alpha_{n+1}^i}} / \sqrt{\alpha_1^i! \alpha_2^i! \dots \alpha_{n+1}^i!} \quad (2.8)$$

for some natural numbers $\alpha_1^i + \alpha_2^i + \dots + \alpha_{n+1}^i = d$. The meaningful consequence is that the kernel K_d corresponds to features $\tilde{\varphi}_d$ which form a basis of the set of polynomials in n variables of degree at most d . Therefore, $\Phi_d(X)$ and $\tilde{\Phi}_d(X)$ have the same rank as $K_d(X, X)$ by virtue of $K_d(X, X) = \tilde{\Phi}_d(X)^\top \tilde{\Phi}_d(X)$.

In [32], the authors argue why using the monomial embedding is appropriate when the points m_i belong to an algebraic variety. Suppose the variety $V(P) \subset \mathbb{R}^n$ is defined by the set of polynomials $P = \{p_1, \dots, p_k\}$ where each p_i is at most of degree d . Then the columns of X belong to the variety $V(P)$ if and only if there exists $C \in \mathbb{R}^{N \times k}$ such that $\Phi_d(X)^\top C = 0$, where the columns of C define the coefficients of the polynomials p_i in the monomial basis. This implies that $\text{rank}(\Phi_d(X)) \leq \min(N - k, s)$. This justifies that $\Phi_d(X)$ is rank deficient when there are sufficiently many data points such that $s \geq N - k$, and X follows an algebraic variety model. The second case study below presents a union of subspaces as a particular type of algebraic variety.

Case study 2.2 (Union of subspaces). *Given two affine subspaces $\mathcal{S}_1, \mathcal{S}_2 \subset \mathbb{R}^n$ of dimension r_1 and r_2 respectively, we can write $\mathcal{S}_1 = \{x : q_i(x) = 0 \text{ for } i = 1, \dots, n - r_1\}$ and $\mathcal{S}_2 = \{x : p_j(x) = 0 \text{ for } j = 1, \dots, n - r_2\}$ where the q_i and p_j are affine functions. The union $\mathcal{S}_1 \cup \mathcal{S}_2$ can be expressed as the set where all possible products $q_i(x)p_j(x)$ vanish. Therefore, $\mathcal{S}_1 \cup \mathcal{S}_2$ is the solution of a system of $(n - r_1)(n - r_2)$ quadratic polynomial equations. Similarly, a union of k affine subspaces of dimensions r_1, \dots, r_k is a variety described by a system of $\prod_{i=1}^k (n - r_i)$ polynomial equations of degree k .*

Proposition 2.1 (Rank of monomial features [32]). *If the columns of a matrix $X \in \mathbb{R}^{n \times s}$ belong to a union of p affine subspaces of dimension at most \tilde{r} , then for any degree $d \geq 1$, the*

matrix $\Phi_d(X) \in \mathbb{R}^{N(n,d) \times s}$ of monomial features, with $N(n,d)$ the dimension of the features space defined in equation (2.2), satisfies

$$\text{rank } \Phi_d(X) \leq p \binom{\tilde{r} + d}{d}. \quad (2.9)$$

In practice, choosing the degree d of the monomial kernel is a tricky task. In Section 9, we discuss the practical choice of this degree and how it impacts the rank of the feature matrix and the possibility to recover M . Previous works using the monomials kernel to recover high-rank matrices all restricted themselves to degrees two or three [17, 32]. Using a polynomial embedding of large degree would seem helpful to capture all the nonlinearity in some data sets. Unfortunately, increasing the degree will grow the dimensions of the optimization problem exponentially. Indeed, the dimension $N(n,d)$ blows up with d for even moderate values of n and the number of data points required is at least $N(n,d) - k$ where k is the number of polynomial equations that define the variety.

We now define the Gaussian kernel which will be used when the columns of the matrix M are grouped in several clusters.

Case study 2.3 (Clusters). For $X, Y \in \mathbb{R}^{n \times s}$, the entry (i, j) of the Gaussian kernel $K^G : \mathbb{R}^{n \times s} \times \mathbb{R}^{n \times s} \rightarrow \mathbb{R}^{s \times s}$ is defined as

$$K_{ij}^G(X, Y) = e^{-\frac{\|x_i - y_j\|_2^2}{2\sigma^2}}, \quad (2.10)$$

where $\sigma > 0$ is the width of the kernel. The Gaussian kernel acts as a proximity measure. For two columns of X , labelled x_i and x_j , we observe that x_i being close to x_j gives $K_{ij}^G(X, X) \approx 1$ and if x_i is far from x_j then $K_{ij}^G(X, X) \approx 0$. Therefore, the rank of the Gaussian kernel approximately coincides with the number of clusters in X . More precisely, one can show that the singular values whose index exceeds the number of clusters decay rapidly [34]. The value of σ should be chosen appropriately depending on the size of the clusters.

In Figure 2, we present a small data set of 100 data points divided in two clusters in \mathbb{R}^2 with the singular values of the Gaussian kernel (in log-scale). We see that the two largest singular values are much greater than the third one, and that the following singular values decrease at an approximately exponential rate. Therefore, the Gaussian kernel is near a low rank matrix for clustered data, which will allow us to complete such data sets from partial measurements. In [16], the Gaussian kernel was also used effectively on image inpainting and denoising problems.




Figure 2: Clustered data and the singular values of the Gaussian kernel in log-scale.

3 Nonlinear matrix recovery as an optimization problem

Noiseless measurements case Considering a noiseless measurements case, we would like to minimize the rank of the feature matrix, as in Equation (1.5). This is unfortunately known to be intractable, even in the case where M is low-rank [22]. We have to resort to approximations of this problem. The second difficulty is the nonlinearity of Φ .

As a nonconvex approximation to (1.5), we consider the formulation in [15] for low-rank problems, and extend it to the nonlinear case. Assuming that $\Phi(M)$ has rank r leads to the following formulation

$$\begin{cases} \min_{X, \mathcal{U}} f(X, \mathcal{U}) := \|\Phi(X) - P_{\mathcal{U}}\Phi(X)\|_F^2 \\ \mathcal{U} \in \text{Grass}(N, r) \\ \mathcal{A}(X) = b, \end{cases} \quad (3.1)$$

where $\text{Grass}(N, r)$ is the Grassmann manifold, the set of all subspaces of dimension r in \mathbb{R}^N , $P_{\mathcal{U}}$ is the orthogonal projection on the subspace \mathcal{U} and $\|\cdot\|_F$ denotes the Frobenius norm. The linear measurements, $\mathcal{A}(X) = b$, are defined in equation (1.1). Given $U \in \mathbb{R}^{N \times r}$ such that $\text{range}(U) = \mathcal{U}$ and $U^T U = I_r$, the projection is given by $P_{\mathcal{U}} = U U^T$. In (3.1), the objective function is expected to be nonconvex but smooth for practical choices of φ , such as case study 2.1, 2.3. If the variable \mathcal{U} is additionally constrained to be the range of the r leading singular vectors of $\Phi(X)$, the cost function becomes $\sum_{i=r+1}^{\min(N, s)} \sigma_i(\Phi(X))^2$. The advantage of the formulation in (3.1) is that it is straightforward to express it as a finite sum of s terms over all the data points. This allows to use stochastic sub-sampling algorithms that scale better to matrices with many columns (large s).

The advantage of using the Grassmann manifold, which is a quotient space, instead of the Stiefel manifold of orthogonal matrices $\text{St}(N, r) := \{U \in \mathbb{R}^{N \times r} : U^T U = I_r\}$ is that due to the invariance of the cost function with respect to the matrix that represents the subspace \mathcal{U} , local optimizers cannot possibly be isolated in a formulation over $\text{St}(N, r)$. Therefore, the fast local convergence rates of some second-order algorithms might not apply on $\text{St}(N, r)$, while they would apply on the quotient manifold.

Consider U^\perp a basis of \mathcal{U}^\perp , the orthogonal complement of \mathcal{U} in \mathbb{R}^N . The variable $U^\perp \in \text{Grass}(N, N - r)$ has a nice interpretation since U^\perp spans $\text{null}(\Phi(X)^\top)$ when $f(X, \mathcal{U}) = 0$. In the case of algebraic varieties (case study 2.1), U^\perp gives the coefficients of the polynomials defining the variety in the basis given by Φ . Recovering the equations of the variety is of interest in some applications [8, 24].

Noisy measurements case When the measurements are known to be noisy, i.e. $\mathcal{A}(M) = b + \eta$ for some noise $\eta \in \mathbb{R}^m$, the measurement constraint can be lifted into the cost function as a penalty. This gives

$$\begin{cases} \min_{X, \mathcal{U}} f_\lambda(X, \mathcal{U}) := \|\Phi(X) - P_{\mathcal{U}}\Phi(X)\|_{\mathbb{F}}^2 + \lambda \|\mathcal{A}(X) - b\|_2^2 \\ \mathcal{U} \in \text{Grass}(N, r), \end{cases} \quad (3.2)$$

where the parameter $\lambda > 0$ has to be adjusted. This allows to satisfy the noisy measurements approximately.

3.1 Kernel representation of the features

When working with a kernel instead of a feature map, as shown in (1.4), we want to find a cost function equivalent to that of (3.1) which uses the kernel instead of the feature map. We find that they are related in the following way.

Proposition 3.1. *Given a feature map Φ and the associated kernel $K : X \times X \mapsto \Phi(X)^\top \Phi(X)$, for $\mathcal{W} \in \text{Grass}(s, r)$ we have*

$$\left\| \Phi(X)^\top - P_{\mathcal{W}}\Phi(X)^\top \right\|_{\mathbb{F}}^2 = \text{trace}\left(K(X, X) - P_{\mathcal{W}}K(X, X)\right). \quad (3.3)$$

Proof. We write $P_{\mathcal{W}^\perp} = I_{s \times s} - P_{\mathcal{W}}$ and find

$$\begin{aligned} \text{trace}\left(P_{\mathcal{W}^\perp}K(X, X)\right) &= \text{trace}\left(P_{\mathcal{W}^\perp}\Phi(X)^\top \Phi(X)\right) \\ &= \text{trace}\left(\Phi(X)P_{\mathcal{W}^\perp}\Phi(X)^\top\right) \\ &= \text{trace}\left(\Phi(X)(P_{\mathcal{W}^\perp})^\top P_{\mathcal{W}^\perp}\Phi(X)^\top\right) \\ &= \text{trace}\left((P_{\mathcal{W}^\perp}\Phi(X)^\top)^\top P_{\mathcal{W}^\perp}\Phi(X)^\top\right) \\ &= \left\| P_{\mathcal{W}^\perp}\Phi(X)^\top \right\|_{\mathbb{F}}^2. \end{aligned} \quad (3.4)$$

□

Using the kernel formula (3.3) corresponds to finding a subspace \mathcal{W} of dimension r in the row space of $\Phi(X)$. When a kernel is used for the embedding, the following optimization problem is solved,

$$\begin{cases} \min_{\mathcal{W}, X} f(X, \mathcal{W}) := \text{trace}\left(K(X, X) - P_{\mathcal{W}}K(X, X)\right) \\ \mathcal{W} \in \text{Grass}(s, r) \\ \mathcal{A}(X) = b. \end{cases} \quad (3.5)$$

Replacing the features Φ by the corresponding kernel becomes beneficial when the dimension N of the features is larger than the number of points s and when a convenient formula is available to compute the kernel and its derivatives. For example, in the case of clusters (case study 2.3), the features exist implicitly in an infinite dimensional space and we use the Gaussian kernel to represent them.

In the upcoming sections, we usually describe the algorithms and their properties using the notation of problem (3.1) with a feature map Φ and cost function $f(X, \mathcal{U})$. Unless specified otherwise, the developments also apply to problem (3.5) and the use of a kernel.

4 Riemannian optimization algorithms

Riemannian optimization methods provide an elegant and efficient way to solve optimization problems with smooth nonlinear constraints. The field of Riemannian optimization has rapidly developed over the past two decades. In particular, Riemannian optimization methods have proved very efficient in low-rank matrix completion [5, 39]. In this section, we investigate the use of Riemannian optimization methods to solve (3.1) or (3.5). For an overview of optimization algorithms on Riemannian manifolds, see [2, 4]. In order to formally express (3.1) as a Riemannian optimization problem, we define a notation for the affine subspace that represents the measurements on the matrix M ,

$$\mathcal{L}_{\mathcal{A},b} = \{X \in \mathbb{R}^{n \times s} : \mathcal{A}(X) = b\}. \quad (4.1)$$

We form the product manifold

$$\mathcal{M} = \mathcal{L}_{\mathcal{A},b} \times \text{Grass}(N, r), \quad (4.2)$$

so that Problem (3.1) can be viewed as the unconstrained minimisation of a smooth cost function defined on the manifold \mathcal{M} ,

$$\begin{cases} \min_{(X,\mathcal{U})} & \|\Phi(X) - P_{\mathcal{U}}\Phi(X)\|_{\text{F}}^2 \\ & (X,\mathcal{U}) \in \mathcal{M}. \end{cases} \quad (4.3)$$

We introduce the notation $z = (X,\mathcal{U}) \in \mathcal{M}$ to denote both variables that appear in the optimization problem. Riemannian optimization algorithms are feasible methods that iteratively exploit the local geometry of the feasible set. Analogously to unconstrained optimization in Euclidean spaces, each iteration of a Riemannian optimization algorithm uses derivatives to build a model that locally approximates the cost function. This model is then fully or approximately minimized. Most commonly, the model uses the gradient and possibly the Hessian or an approximation of it. This yields respectively a first- or second-order method. Riemannian optimization follows these principles, only the model is defined on a local linearisation of the manifold, namely, the tangent space. The Riemannian gradient, written $\text{grad}f$ is a vector belonging to the tangent space of the manifold. The Riemannian Hessian, written $\text{Hess}f$ is a symmetric operator on that tangent space. In Appendix A, we show how to compute the Euclidean gradient and Hessian for the cost function of (3.5) in the case of the kernels presented in Section 2 (monomial kernel, Gaussian kernel). Then, to find their Riemannian counterparts, ∇f and $\nabla^2 f$ are projected onto the tangent space of \mathcal{M} using the tools defined later on in this section. When exploring the tangent space, it is necessary to have a tool that allows to travel on the manifold in a direction prescribed by a tangent vector. This operation is called a retraction [2, Def. 4.1.1].

Definition 4.1 (Retraction). *A retraction on a manifold \mathcal{M} is a smooth mapping Retr from the tangent bundle $\text{T}\mathcal{M}$ to \mathcal{M} with the following properties. Let $\text{Retr}_z: \text{T}_z\mathcal{M} \rightarrow \mathcal{M}$ denote the restriction of Retr to $\text{T}_z\mathcal{M}$.*

- (i) $\text{Retr}_z(0_z) = z$, where 0_z is the zero vector in $\text{T}_z\mathcal{M}$;
- (ii) the differential of Retr_z at 0_z , $\text{DRetr}_z(0_z)$, is the identity map.

The retraction curves $t \mapsto \text{Retr}_z(t\eta)$ agree up to first order with geodesics passing through z with velocity η , around $t = 0$.

Let us detail the tools necessary to use Riemannian optimization methods on the two manifolds that compose our search space \mathcal{M} . Note that the Cartesian product of two Riemannian manifolds is a Riemannian manifold. The geometry of $\mathcal{L}_{\mathcal{A},b}$ is rather trivial because the manifold is affine. It must nonetheless be implemented so we will describe how to handle this constraint in a Riemannian way. More generally, this gives a straightforward way to deal with affine equality constraints.

Measurement subspace At any point $X \in \mathcal{L}_{\mathcal{A},b}$, the tangent space is the null space of \mathcal{A} ,

$$\mathrm{T}_X \mathcal{L}_{\mathcal{A},b} = \mathrm{null}(\mathcal{A}) = \{\Delta \in \mathbb{R}^{n \times s} : \mathcal{A}(\Delta) = 0\}. \quad (4.4)$$

Since the tangent space does not depend on X , we write $\mathrm{T}\mathcal{L}_{\mathcal{A},b}$. This tangent space inherits an inner product from the embedding space $\mathbb{R}^{n \times s}$,

$$\langle \Delta_1, \Delta_2 \rangle = \mathrm{trace}(\Delta_1^\top \Delta_2) \quad \text{for all } \Delta_1, \Delta_2 \in \mathrm{T}\mathcal{L}_{\mathcal{A},b}. \quad (4.5)$$

The Riemannian gradient is the orthogonal projection of the Euclidean gradient onto the tangent space $\mathrm{T}\mathcal{L}_{\mathcal{A},b}$. From the fundamental theorem of algebra, $\mathrm{null}(\mathcal{A}) = \mathrm{range}(\mathcal{A}^\top)^\perp$. Therefore we can express $\mathrm{P}_{\mathrm{null}(\mathcal{A})} = \mathrm{Id} - \mathrm{P}_{\mathrm{range}(\mathcal{A}^\top)}$.

The application \mathcal{A} is represented by a flat matrix $A \in \mathbb{R}^{m \times ns}$ such that $\mathcal{A}(X) = AX(\cdot) \in \mathbb{R}^m$, where $X(\cdot)$ is a vector of length ns made of the columns of X taken from left to right and stacked on top of each other. Visually this gives,

$$\mathcal{A}(X) = \begin{pmatrix} \langle A_1, X \rangle \\ \langle A_2, X \rangle \\ \vdots \\ \langle A_m, X \rangle \end{pmatrix} = \begin{pmatrix} A_1(\cdot)^\top \\ A_2(\cdot)^\top \\ \vdots \\ A_m(\cdot)^\top \end{pmatrix} X(\cdot) =: AX(\cdot). \quad (4.6)$$

The tall matrix A^\top represents the linear application \mathcal{A}^\top . Viewing A^\top as the matrix of an overdetermined system of linear equations convinces us that $\mathrm{P}_{\mathrm{range}(\mathcal{A}^\top)} = A^\top(AA^\top)^{-1}A$. Equivalently, if $Q \in \mathbb{R}^{m \times m}$ is an orthogonal basis for $\mathrm{range}(\mathcal{A}^\top)$ (which can be obtained by a reduced QR factorization of A^\top), we can apply $\mathrm{P}_{\mathrm{range}(\mathcal{A}^\top)} = QQ^\top$. The projection is given by

$$\mathrm{P}_{\mathrm{null}(\mathcal{A})} = \mathrm{I}_{ns} - A^\top(AA^\top)^{-1}A = \mathrm{I}_{ns} - QQ^\top. \quad (4.7)$$

Therefore,

$$\mathrm{P}_{\mathrm{T}\mathcal{L}_{\mathcal{A},b}} : \mathbb{R}^{n \times s} \rightarrow \mathrm{T}\mathcal{L}_{\mathcal{A},b} : \mathrm{P}_{\mathrm{T}\mathcal{L}_{\mathcal{A},b}}(\Delta) = (\mathrm{I}_{ns} - QQ^\top)\Delta \quad (4.8)$$

where $Q \in \mathbb{R}^{m \times m}$ is an orthogonal basis for $\mathrm{range}(\mathcal{A}^\top)$. In the case of matrix completion, the operator \mathcal{A} selects the known entries of M . The description of the feasible subspace is simplified. We write $\mathcal{L}_{\Omega,b} = \{X : X_{ij} = M_{ij}, ij \in \Omega\}$ to make explicit that the measurements correspond to matrix completion. The tangent space is $\mathrm{TL}_{\Omega,b} = \{\Delta : \Delta_{ij} = 0 \text{ for } ij \in \Omega\}$ and the projection onto $\mathrm{TL}_{\Omega,b}$ simply amounts to setting the entries outside of Ω to zero,

$$\mathrm{P}_{\mathrm{TL}_{\Omega,b}}(\Delta) = \begin{cases} 0 & \text{for } ij \in \Omega \\ \Delta_{ij} & \text{for } ij \notin \Omega. \end{cases}$$

The natural retraction on $\mathcal{L}_{\mathcal{A},b}$ for $X \in \mathcal{L}_{\mathcal{A},b}$ and $\Delta \in \mathrm{T}\mathcal{L}_{\mathcal{A},b}$ is given by

$$\mathrm{Retr}_X : \mathrm{T}\mathcal{L}_{\mathcal{A},b} \rightarrow \mathcal{L}_{\mathcal{A},b} : \mathrm{Retr}_X(\Delta) = X + \Delta, \quad (4.9)$$

because the manifold is flat. These tools are also needed for the Grassmann manifold and we follow the representation given in [5].

Grassmann manifold The Grassmann manifold, written $\text{Grass}(N, r)$, is the set of all linear subspaces of dimension r in \mathbb{R}^N . A point $\mathcal{U} \in \text{Grass}(N, r)$ is represented by a full-rank matrix $U \in \mathbb{R}^{N \times r}$ such that $\text{range}(U) = \mathcal{U}$. For any orthogonal matrix $Y \in \mathbb{R}^{r \times r}$, the matrix UY is also a valid representation of \mathcal{U} , since $\text{range}(UY) = \text{range}(U)$. The set of matrices with orthonormal columns is defined as the Stiefel manifold, $\text{St}(N, r) = \{U \in \mathbb{R}^{N \times r} : U^\top U = \mathbf{I}_r\}$, and the orthogonal group is defined as $O(r) = \{Y \in \mathbb{R}^{r \times r} : Y^\top Y = \mathbf{I}_r\}$. The orthogonal group induces an equivalence relation on the Stiefel manifold, where any two matrices are equivalent if they have the same column space. In this regard, the Grassmann is a quotient space

$$\text{Grass}(N, r) = \text{St}(N, r)/O(r), \quad (4.10)$$

and each equivalence class consists of all matrices with the same span. The tangent space to the Stiefel manifold at $U \in \text{St}(N, r)$ has the form

$$\text{T}_U \text{St}(N, r) = \left\{UZ + U^\perp B : Z \in \text{Skew}(r), B \in \mathbb{R}^{(N-r) \times r}\right\}, \quad (4.11)$$

where $\text{Skew}(r)$ is the set of skew-symmetric matrices of size r [4, Section 7.3]. The equivalence class of $U \in \text{St}(N, r)$ —seen as a submanifold of $\text{St}(N, r)$ —has a tangent space at U , which is called the vertical space $\text{V}_U \text{St}(N, r) = \{UZ : Z \in \text{Skew}(r)\} \subseteq \text{T}_U \text{St}(N, r)$. The orthogonal complement of $\text{V}_U \text{St}(N, r)$ in $\text{T}_U \text{St}(N, r)$ is called the horizontal space and is given by

$$\text{H}_U \text{St}(N, r) = \left\{U^\perp B : B \in \mathbb{R}^{(N-r) \times r}\right\} = \text{range}(U^\perp) \subseteq \text{T}_U \text{St}(N, r). \quad (4.12)$$

As is common in differential geometry, the tangent space to $\text{Grass}(N, r)$ at \mathcal{U} is represented by the horizontal space $\text{H}_U \text{St}(N, r)$, that is, $\text{H}_U \text{St}(N, r) \simeq \text{T}_\mathcal{U} \text{Grass}(N, r)$. Any tangent vector $H_\mathcal{U} \in \text{T}_\mathcal{U} \text{Grass}(N, r)$ is represented by an horizontal vector $H_U \in \text{H}_U \text{St}(N, r)$ called the horizontal lift of $H_\mathcal{U}$ at U . A thorough treatment of quotient manifolds, such as the Grassmann, and their usage in optimization can be found in [2, 4]. The projection onto the horizontal space is given by

$$\text{Proj}_{\text{H}_U \text{St}(N, r)} : \text{T}_U \text{St}(N, r) \rightarrow \text{H}_U \text{St}(N, r) : H \mapsto (\text{Id} - UU^\top) H. \quad (4.13)$$

The horizontal space is equipped with the usual inner product

$$\langle H_1, H_2 \rangle_U = \text{trace}(H_1^\top H_2) \quad \text{for all } H_1, H_2 \in \text{H}_U \text{St}(N, r). \quad (4.14)$$

The norm of a tangent vector to the Grassmann is given by the norm of its horizontal lift. Hence we understand the notation $\|H_\mathcal{U}\|_F$ for $H_\mathcal{U} \in \text{T}_\mathcal{U} \text{Grass}(N, r)$ as $\|H_\mathcal{U}\|_F = \|H_U\|_F$, where H_U is the horizontal lift of $H_\mathcal{U}$. Let us call qf , the mapping that sends a matrix to the Q factor of its (reduced) QR decomposition with $Q \in \text{St}(N, r)$ and R upper triangular with positive diagonal entries. To move away from $\mathcal{U} \in \text{Grass}(N, r)$ in the direction $H \in \text{H}_U \text{St}(N, r)$, we use the following retraction

$$\text{Retr}_\mathcal{U} : \text{H}_U \text{St}(N, r) \rightarrow \text{Grass}(N, r) : H \mapsto \text{range}(\text{qf}(U + H)). \quad (4.15)$$

We are now in a position to present the Riemannian trust-region algorithm [2, Ch.7]. This is an extension of the classical trust-region methods [10] to Riemannian manifolds.

Riemannian trust-region (RTR) At each iterate, RTR builds a local model of the function. The method sequentially minimizes this model under a ball constraint that prevents undesirably large steps where the model does not accurately represent the function. The trust-region subproblem takes the following form around $z_k \in \mathcal{M}$

$$\min_{\substack{\eta \in \text{T}_{z_k} \mathcal{M} \\ \|\eta\|_{z_k} \leq \Delta_k}} \hat{m}_{z_k}(\eta) := f(z_k) + \langle \eta, \text{grad}f(z_k) \rangle_{z_k} + \frac{1}{2} \langle \eta, H_k[\eta] \rangle_{z_k}, \quad (4.16)$$

where $H_k: T_{z_k}\mathcal{M} \rightarrow T_{z_k}\mathcal{M}$ is a symmetric operator on $T_{z_k}\mathcal{M}$, Δ_k is the trust-region radius and the model $\hat{m}_{z_k}: T_{z_k}\mathcal{M} \rightarrow \mathbb{R}$ is a quadratic approximation of the *pullback* $\hat{f}_{z_k} = f \circ \text{Retr}_{z_k}$, defined on the tangent space at $z_k \in \mathcal{M}$.

First-order Riemannian trust-region When the Hessian of the cost function is expensive to compute or not available altogether, one can use a first-order model and set $H_k = 0$. This method will be very similar to gradient descent. But the trust-region is used to ensure global convergence, as opposed to a line-search.

Second-order Riemannian trust-region When the true Hessian of the cost function is available, the classical second-order trust-region method is obtained with $H_k = \text{Hess}f(z_k)$. It is also possible to use an approximation of the true Hessian for H_k .

Algorithm 1 Riemannian trust-region (RTR) [6]

```

1: Given:  $z_0 \in \mathcal{M}$  and  $0 < \Delta_0 < \bar{\Delta}$ ,  $\varepsilon_g > 0$ ,  $\varepsilon_H > 0$  and  $0 < \rho' < 1/4$ 
2: Init:  $k = 0$ 
3: while true do
4:   if  $\|\text{grad}f(z_k)\| > \varepsilon_g$  then
5:     Obtain  $\eta_k \in T_{z_k}\mathcal{M}$  satisfying A5
6:   else if  $\varepsilon_H < \infty$  then
7:     if  $\lambda_{\min}(H_k) < -\varepsilon_H$  then
8:       Obtain  $\eta_k \in T_{z_k}\mathcal{M}$  satisfying A6
9:     else
10:      return  $z_k$ 
11:    end if
12:   else
13:    return  $z_k$ 
14:   end if
15:    $z_k^+ = \text{Retr}_{z_k}(\eta_k)$ 
16:    $\rho = f(z_k) - f(z_k^+) / (\hat{m}_{z_k}(0) - \hat{m}_{z_k}(\eta_k))$ 
17:   if  $\rho < 1/4$  then
18:      $\Delta_{k+1} = \Delta_k/4$ 
19:   else if  $\rho > 3/4$  and  $\|\eta_k\| = \Delta_k$  then
20:      $\Delta_{k+1} = \min(2\Delta_k, \bar{\Delta})$ 
21:   else
22:      $\Delta_{k+1} = \Delta_k$ 
23:   end if
24:   if  $\rho > \rho'$  then
25:      $z_{k+1} = z_k^+$ 
26:   end if
27:    $k = k + 1$ 
28: end while

```

We apply RTR, as described in Algorithm 1, to problem (3.1) or (3.5). If a first-order critical point is sought, set $\varepsilon_H = \infty$. The second-order version of RTR provably converges to second-order critical points for any initialization under a weak decrease condition in the subproblems and satisfies global worst-case complexity bounds matching their unconstrained counterparts, as was shown in [6]. The local convergence rate is quadratic for an appropriate choice of parameters [2, Chap.7]. We note that there is no guarantee on the quality of the stationary point, due to the

nonconvexity. Nonetheless, we see in Section 9 that the method performs very well in practice for nonlinear matrix recovery. We introduce the following assumptions.

A2. *There exists $f^* > -\infty$ such that $f(x) \geq f^*$ for all $x \in \mathcal{M}$.*

The cost functions of problems (3.1) and (3.5) are nonnegative, so A2 is satisfied throughout this paper. We also state a regularity assumption on the gradient and Hessian of the pullback which was introduced in [6]. In Appendix A, we detail how these conditions relate to the smoothness of the Riemannian derivatives and discuss the practicality of these assumptions for problem (3.5).

A3 (Lipschitz gradient of the pullback). *There exists $L_g \geq 0$ such that for all $z = (X, \mathcal{U}) \in \mathcal{M}$, the pullback $\hat{f}_z = f \circ \text{Retr}_z$ has Lipschitz continuous gradient with constant L_g , that is, for all $\eta \in \mathbb{T}_z \mathcal{M}$, it holds that*

$$\left| \hat{f}_z(\eta) - [f(z) + \langle \eta, \text{grad} f(z) \rangle] \right| \leq \frac{L_g}{2} \|\eta\|^2. \quad (4.17)$$

A4 (Lipschitz Hessian of the pullback). *There exists $L_H \geq 0$ such that, for all $z = (X, \mathcal{U}) \in \mathcal{M}$, the pullback $\hat{f}_z = f \circ \text{Retr}_z$ has Lipschitz continuous Hessian with constant L_H , that is, for all $\eta \in \mathbb{T}_z \mathcal{M}$, it holds that*

$$\left| \hat{f}_z(\eta) - \left[f(z) + \langle \eta, \text{grad} f(z) \rangle + \frac{1}{2} \langle \eta, \nabla^2 \hat{f}_z(0_z)[\eta] \rangle \right] \right| \leq \frac{L_H}{6} \|\eta\|^3. \quad (4.18)$$

Algorithm 1 is flexible in that it does not specify how the subproblems are solved. We discuss the implementation of RTR in Section 9. For the complexity results, the following decreases in the model for first- and second-order steps are required.

A5. *There exists $c_2 > 0$ such that all first-order steps η_k satisfy*

$$\hat{m}_k(0_{z_k}) - \hat{m}_k(\eta_k) \geq c_2 \min \left(\Delta_k, \frac{\varepsilon_g}{c_0} \right) \varepsilon_g. \quad (4.19)$$

A6. *There exists $c_3 > 0$ such that all second-order steps η_k satisfy*

$$\hat{m}_k(0_{z_k}) - \hat{m}_k(\eta_k) \geq c_3 \Delta_k^2 \varepsilon_H. \quad (4.20)$$

A7. *There exists $c_0 \geq 0$ such that, for all first-order steps, $\|H_k\| \leq c_0$ and H_k is radially linear, that is, for all $\alpha \geq 0$ and $\eta \in \mathbb{T}_{z_k} \mathcal{M}$, it holds $H_k[\alpha\eta] = \alpha H_k[\eta]$.*

A8. *There exists $c_1 \geq 0$ such that, for all second-order steps,*

$$\left| \left\langle \eta_k, \left(\nabla^2 \hat{f}_{z_k}(0_{z_k}) - H_k \right) [\eta_k] \right\rangle \right| \leq \frac{c_1 \Delta_k}{3} \|\eta_k\|^2.$$

In addition, for all second-order steps, H_k is linear and symmetric.

Define the following constants

$$\lambda_g = \frac{1}{4} \min\left(\frac{1}{c_0}, \frac{c_2}{L_g + c_0}\right) \quad \text{and} \quad \lambda_H = \frac{3}{4} \frac{c_3}{L_H + c_1}. \quad (4.21)$$

The following (sharp) worst-case bound for Riemannian trust-region was recently established.

Theorem 4.1 (Global complexity of RTR [6]). *Under A3, A5, A7 and assuming $\varepsilon_g \leq \frac{\Delta_0}{\lambda_g}$, Algorithm 1 produces an iterate z_{N_1} satisfying $\|\text{grad}f(z_{N_1})\| \leq \varepsilon_g$ with*

$$N_1 \leq \mathcal{O}(1/\varepsilon_g^2). \quad (4.22)$$

Furthermore, if $\varepsilon_H < \infty$ then under additionally A4, A6, A8 and assuming $\varepsilon_g \leq \frac{c_2}{c_3} \frac{\lambda_H}{\lambda_g^2}$ and $\varepsilon_H \leq \frac{c_2}{c_3 \lambda_g}$, Algorithm 1 also produces an iterate z_{N_2} satisfying $\text{grad}f(z_{N_2}) \leq \varepsilon_g$ and $\lambda_{\min}(H_{N_2}) \geq -\varepsilon_H$ with

$$N_1 \leq N_2 \leq \mathcal{O}\left(\frac{1}{\varepsilon^2 \varepsilon_H}\right). \quad (4.23)$$

5 An alternating minimization algorithm

In this section, we propose an alternating minimization algorithm to solve (3.1) or (3.5) (Algorithm 2). This comes from the natural separation of the variables into two blocks X and \mathcal{U} , yielding two distinct minimization subproblems. Alternating minimization type methods have been very popular in recent years to solve large-scale nonconvex problems [3, 40]. This is due to their good practical performances and ease of implementation, as often one or both of the subproblems have a closed form solution. Strictly speaking, this is still a Riemannian optimization approach, as all iterates will be feasible for the constraints. But this section describes a two-block coordinate minimization, whereas the previous section was considering full block variants.

We set the initial guess X_0 as any solution of the underdetermined linear system $\mathcal{A}(X) = b$ and \mathcal{U}_0 as the span of the r leading singular vectors of $\Phi(X_0)$. The framework is as follows, for $k \geq 0$:

With \mathcal{U}_k fixed, solve

$$X_{k+1} = \begin{cases} \operatorname{argmin}_{X \in \mathbb{R}^{n \times s}} & \|\Phi(X) - P_{\mathcal{U}_k} \Phi(X)\|_{\mathbb{F}}^2 \\ & \mathcal{A}(X) = b. \end{cases} \quad (5.1)$$

With X_{k+1} fixed, solve

$$\mathcal{U}_{k+1} = \begin{cases} \operatorname{argmin}_U & \|\Phi(X_{k+1}) - P_U \Phi(X_{k+1})\|_{\mathbb{F}}^2 \\ & U \in \text{Grass}(N, r). \end{cases} \quad (5.2)$$

This separation of the variables takes advantage of the fact that problem (5.2), even though nonconvex, is solved to global optimality by computing the r leading left singular vectors of the matrix $\Phi(X_{k+1})$. The result is a consequence of the celebrated Eckart-Young-Mirsky theorem, which gives the best rank r approximation in Frobenius norm of a matrix by the r leading terms of the singular value decomposition [14, 28]. In particular, let $\Phi(X_k) = \sum_{i=1}^{\min(N,s)} \sigma_i u_i v_i^\top$, then

$$\mathcal{U}_{k+1} = \text{span}(u_1, \dots, u_r) \quad (5.3)$$

is a global minimizer of (5.2)¹. This truncated singular value decomposition (SVD) is denoted by `truncate_svd` in Algorithm 2. The singular vectors can be computed to high accuracy in polynomial time [38].

Problem (5.1) is in general hard to solve to global optimality. The difficulty comes from the nonconvexity of the cost function, which is due to Φ . One can choose from a variety of first-or second-order methods to find an approximate first-or second-order critical point. We will present the merits of both possibilities in Section 9.

i) First-order version of alternating minimization When only gradient information is available, a first-order method will be used to minimize subproblem (5.1). For the sake of illustration, in Algorithm 2 we present a projected gradient descent with line search for (5.1). The gradient of the cost function is projected onto the null space of \mathcal{A} . This ensures that the iterates remain in the feasible set $\mathcal{L}_{\mathcal{A},b}$. The line search is a classical backtracking with an Armijo condition for sufficient decrease. Variants in the line search or even constant step sizes are possible.

ii) Second-order version of alternating minimization In the subproblem (5.1), it is possible to use second-order methods to speed up the convergence and reach a higher accuracy. We apply RTR on the affine manifold $\mathcal{L}_{\mathcal{A},b}$ with the Hessian of the cost function $\nabla_{XX}^2 f(X_k, \mathcal{U}_k)$ in the model.

Algorithm 2 details the alternating minimization where gradient descent with an Armijo line search, a standard inexact procedure in nonconvex optimization, is applied in the subproblem (5.1). The Armijo line search is described in Algorithm 3.

iii) Accuracy of the subproblems solution Algorithm (2) alternatively solves the subproblems (5.1) and (5.2). For the solution of (5.1), there is no incentive to solve it to very high accuracy early on in the run of the algorithm, as we could still be far from convergence and the variable \mathcal{U} might still change a lot. At iterations k , we use the following stopping criterion for some $\varepsilon_{x,k} > 0$,

$$\|\text{grad}_X f(X_{k+1}, \mathcal{U}_k)\|_{\text{F}} \leq \varepsilon_{x,k}. \quad (5.4)$$

We propose the two following strategies for the choice of $\varepsilon_{x,k}$,

$$\varepsilon_{x,k} = \varepsilon_x \text{ for all } k, \quad (5.5)$$

or

$$\varepsilon_{x,k} = \max(\varepsilon_x, \theta \|\text{grad}_X f(X_k, \mathcal{U}_k)\|_{\text{F}}) \text{ for some user-chosen } 0 < \theta < 1. \quad (5.6)$$

To solve (5.2), it is possible to use a randomized SVD procedure. The randomized SVD is a stochastic algorithm that approximately computes the singular value decomposition of a matrix that exhibits a low-rank pattern [25]. The matrix must be of low rank or have a fast decay in its singular values for the random SVD to be accurate. As the iterates X_k converge towards the solution M , the matrix $\Phi(X_k)$, for which we have to compute an SVD, becomes low-rank and therefore it is natural to use a randomized SVD in Algorithm 2. In the first iterations, for a random starting point of the algorithm, the feature matrix is not expected to be low-rank. In those case, the random SVD should not be used. When the matrix $\Phi(X_k)$ is only approximately

¹Note that the solution need not be unique, in the case where $\sigma_r = \sigma_{r+1}$.

Algorithm 2 Alternating minimization scheme for Problem (3.1) or (3.5)

1: **Given:** The sensing matrix $A \in \mathbb{R}^{m \times ns}$, measurements $b \in \mathbb{R}^m$, tolerances $\varepsilon_u \geq 0, \varepsilon_x \geq 0$,
an estimation of $r = \text{rank}(\Phi(M))$.

2: Set $k = 0$

3: Find X_0 that satisfies $AX_0 = b$

4: $U_0 = \text{truncate_svd}(\Phi(X_0))$ ▷ Equation (5.3)

5: **while** $\|\text{grad}_X f(X_k, \mathcal{U}_k)\|_F > \varepsilon_x$ or $\|\text{grad}_{\mathcal{U}} f(X_k, \mathcal{U}_k)\|_F > \varepsilon_u$ **do**

6: Set $X_k^{(0)} = X_k, i = 0$

7: Choose $\varepsilon_{x,k}$ using Equation (5.5) or (5.6)

8: **while** $\|\text{grad}_X f(X_k^{(i)}, \mathcal{U}_k)\| > \varepsilon_{x,k}$ **do**

9: $\text{grad}_X f(X_k^{(i)}, \mathcal{U}_k) = \text{P}_{\text{TL}_{\mathcal{A},b}}(\nabla_X f(X_k^{(i)}, \mathcal{U}_k))$ ▷ Equation (4.8)

10: $\alpha_k^{(i)} = \text{Armijo}((X_k^{(i)}, \mathcal{U}_k), -\text{grad}_X f(X_k^{(i)}, \mathcal{U}_k))$ ▷ Algorithm 3

11: $X_k^{(i+1)} = X_k^{(i)} - \alpha_k^{(i)} \text{grad}_X f(X_k^{(i)}, \mathcal{U}_k)$

12: $i = i + 1$

13: **end while**

14: $X_{k+1} = X_k^{(i)}$

15: **if** $\|\text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_k)\| \leq \varepsilon_u$ **then**

16: $U_{k+1} = U_k$

17: **else**

18: $U_{k+1} = \text{truncate_svd}(\Phi(X_{k+1}))$

19: **end if**

20: $k = k + 1$

21: **end while**

22: **Output:** (X_k, \mathcal{U}_k) such that $\|\text{grad} f(X_k, \mathcal{U}_k)\| \leq \varepsilon_u + \varepsilon_x$.

Algorithm 3 Armijo(z_k, d_k): Line search with Armijo condition

INPUT: Function f and gradient $\text{grad}_X f$, current iterate $(X_k^{(i)}, \mathcal{U}_k)$ and a descent direction d_k such that $\langle \text{grad}_X f(X_k^{(i)}, \mathcal{U}_k), d_k \rangle < 0$, a sufficient decrease coefficient $\beta \in]0, 1[$, initial step $\alpha_0 > 0$ and $\tau \in]0, 1[$.

OUTPUT: Step size $\alpha_k^{(i)}$.

1: Set $\alpha = \alpha_0$.

2: **while** $f(X_k^{(i)} + \alpha d_k, \mathcal{U}_k) > f(X_k^{(i)}, \mathcal{U}_k) + \beta \alpha \langle \text{grad}_X f(X_k^{(i)}, \mathcal{U}_k), d_k \rangle$ **do**

3: $\alpha = \tau \alpha$.

4: **end while**

5: Set $\alpha_k^{(i)} = \alpha$.

low-rank, we can apply power iterations to make the singular values decrease faster. This will make the randomized decomposition more costly, but will improve the accuracy of the computed SVD.

Our strategy is as follows, choose two parameters $0 < \tau_1 \ll \tau_2 < 1$. As long as $f(X_{k+1}, \mathcal{U}_k) > \tau_2$, use an exact SVD algorithm, without randomization. When $\tau_1 < f(X_{k+1}, \mathcal{U}_k) \leq \tau_2$ we know that the energy of $\Phi(X_{k+1})$ is mostly contained in the span of \mathcal{U}_k which has dimension r . So we are justified in using a randomized SVD, which we start up with a step of the power method to improve the accuracy. When $f(X_{k+1}, \mathcal{U}_k) \leq \tau_1$, the matrix $\Phi(X_{k+1})$ is even closer to low-rank and we no longer need to apply a power iteration before computing the randomized SVD.

6 Convergence of the alternating minimization algorithm

In this section we present convergence results for the alternating minimization Algorithm 2. We consider a first-order version where subproblem (5.1) is minimized with the gradient descent method and the Armijo backtracking line-search (Algorithm 3). We will first show asymptotic convergence of the gradient norms to zero. We also give a worst-case global complexity bound on the number of iterations necessary to achieve a small gradient from an arbitrary initial starting point. Note that we chose the Armijo linesearch for the sake of example. Minor adjustments of the proof below allow to prove similar results for other minimization methods in subproblem (5.1).

6.1 Assumptions

A9. *There exist constants L_x and L_u (which are both independent of X and \mathcal{U}) such that for all $z = (X, \mathcal{U}) \in \mathcal{M}$, the pullback $\hat{f}_z = f \circ \text{Retr}_z$ has a Lipschitz continuous gradient in X and \mathcal{U} , with constants L_x and L_u respectively. That is, for all $(\eta_x, \eta_u) \in \mathbb{T}_{(X, \mathcal{U})}\mathcal{M}$,*

$$|f \circ \text{Retr}(\eta_x, 0) - [f(X, \mathcal{U}) + \langle \text{grad}_X f(X, \mathcal{U}), \eta_x \rangle]| \leq \frac{L_x}{2} \|\eta_x\|^2 \quad (6.1)$$

and

$$|f \circ \text{Retr}(0, \eta_u) - [f(X, \mathcal{U}) + \langle \text{grad}_{\mathcal{U}} f(X, \mathcal{U}), \eta_u \rangle]| \leq \frac{L_u}{2} \|\eta_u\|^2. \quad (6.2)$$

In words, this means that, in each variable, the pullback is well approximated by its first-order Taylor approximation.

Remark 6.1. *Note that if A3 holds, then A9 holds with $L_x = L_u = L_g$.*

Let us discuss, under which conditions on the kernel one can ensure that A3 or A9 are satisfied for the cost function of (3.5). The following discussion requires to use the exponential map as the retraction. The exponential map follows geodesics along the manifold in directions prescribed by tangent vectors. Using the exponential map on \mathcal{M} is not a restriction, as the exponential map on the Grassmann manifold is computable [1] and the exponential map on $\mathcal{L}_{\mathcal{A}, b}$ is the identity.

Proposition 6.1. *Consider the cost function of (3.5) and assume that the retraction being used is the exponential map. If $\text{DK}(X, X)$ is Lipschitz continuous over $\mathcal{L}_{\mathcal{A}, b}$, then (6.1) holds where L_x is the Lipschitz constant of $\text{DK}(X, X)$. If $\|\text{K}(X, X)\|_{\text{F}} \leq M$ for all $X \in \mathcal{L}_{\mathcal{A}, b}$, condition (6.2) holds with $L_u = 2M$.*

Proof. See Appendix A. □

The conditions listed in Proposition 6.1 on the kernel and its derivatives are not always satisfied or can be difficult to verify. For instance, the Gaussian kernel K_G is bounded above on $\mathcal{L}_{A,b}$, but the monomial kernel K_d is not for any degree $d \geq 1$. For the Gaussian kernel, the map $DK_G(X)$ is always Lipschitz continuous on $\mathcal{L}_{A,b}$. Whereas for the monomial kernel, the map $DK_d(X)$ is Lipschitz continuous for $d \leq 2$, and only locally Lipschitz continuous for $d \geq 3$.

Fortunately, the picture is much simpler if the sequence of iterates $(X_k)_{k \in \mathbb{N}}$ generated by Algorithm 1 or 2 is contained in a bounded set. This ensures that we can find Lipschitz constants such that the bounds of A3, A4 and A9 hold at every iterate of the algorithm (and trial points if any), which is all that is needed in the convergence analysis.

Proposition 6.2. *Consider the cost function of either (3.1) or (3.5) and apply Algorithm 2 or Algorithm 1 with the exponential map as the retraction. If the convex hull of the sequence of iterates $(X_k)_{k \in \mathbb{N}}$ and the trial points of the algorithm is a bounded set, then (4.17), (4.18) and (6.2)-(6.1) hold at every iterate $(X_k, \mathcal{U}_k)_{k \in \mathbb{N}}$ and trial points of the algorithm.*

Proof. See appendix A. □

6.2 Global convergence results

We now carry on with the convergence analysis of the alternating minimization algorithm. The next lemma is an adaptation of the classical descent lemma for the SVD step.

Lemma 6.3 (Descent lemma based on [6, Theorem 4]). *Let A2 and A9 hold for $f : \mathcal{M} \rightarrow \mathbb{R}$. Then, for any $k \geq 0$, the iterates generated by Algorithm 2 satisfy*

$$f(X_k, \mathcal{U}_k) - f(X_{k+1}, \mathcal{U}_{k+1}) \geq \frac{1}{2L_u} \|\text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_k)\|_F^2, \quad (6.3)$$

where L_u is the Lipschitz constant of the gradient of the pullback (A9).

Proof. See Appendix B. □

Throughout this section we use the following notation. Let the number of gradient steps between X_k and X_{k+1} be $n_k \geq 0$ and the intermediate iterates,

$$X_k = X_k^{(0)}, X_k^{(1)}, X_k^{(2)}, \dots, X_k^{(n_k)} = X_{k+1}.$$

The next lemma gives upper and lower bounds on the step sizes given by the Armijo linesearch. This is an adaptation of a standard argument for linesearch methods [30] where the constraint $\mathcal{A}(X) = b$ is added.

Lemma 6.4. *Under A9, for the direction $-\text{grad}_X f(X_k^{(i)}, \mathcal{U}_k) \in T_{X_k} \mathcal{L}_{A,b}$, the linesearch Algorithm 3 produces a step size $\alpha_k^{(i)}$ that satisfies*

$$\underline{\alpha} := \min \left\{ \alpha_0, \frac{2\tau(1-\beta)}{L_x} \right\} \leq \alpha_k^{(i)} \leq \alpha_0 \quad (6.4)$$

and produces the following decrease

$$f(X_k^{(i)}, \mathcal{U}_k) - f(X_k^{(i+1)}, \mathcal{U}_k) \geq \beta \alpha \left\| \text{grad}_X f(X_k^{(i)}, \mathcal{U}_k) \right\|_F^2, \quad (6.5)$$

where $X_k^{(i+1)} = X_k^{(i)} - \alpha_k^{(i)} \text{grad}_X f(X_k^{(i)}, \mathcal{U}_k)$.

Proof. See Appendix B. □

We are now ready to prove global convergence of the alternating minimization algorithm.

Theorem 6.5 (Global convergence for Alternating minimization). *Let A9 hold for $f : \mathcal{M} \rightarrow \mathbb{R}$ from (3.1) or (3.5). Let $\varepsilon_x = 0$, $\varepsilon_u = 0$ and use Equation (5.6) to set $\varepsilon_{x,k}$, for any starting point $(X_0, \mathcal{U}_0) \in \mathcal{M}$, Algorithm 2 produces a sequence $(X_k, \mathcal{U}_k)_{k \in \mathbb{N}}$ such that*

$$\lim_{k \rightarrow \infty} \|\text{grad}f(X_k, \mathcal{U}_k)\|_{\text{F}} = 0. \quad (6.6)$$

Proof. First note that f is bounded below by $f_* = 0$. For any $k \geq 0$,

$$\|(\text{grad}_X f(X_k, \mathcal{U}_k), \text{grad}_{\mathcal{U}} f(X_k, \mathcal{U}_k))\| \leq \|\text{grad}_X f(X_k, \mathcal{U}_k)\|_{\text{F}} + \|\text{grad}_{\mathcal{U}} f(X_k, \mathcal{U}_k)\|_{\text{F}} \quad (6.7)$$

$$\leq \|\text{grad}_X f(X_k, \mathcal{U}_k)\|_{\text{F}} \quad (6.8)$$

since $\varepsilon_u = 0$. Given that each step is non-increasing,

$$\begin{aligned} f(X_k, \mathcal{U}_k) - f(X_{k+1}, \mathcal{U}_{k+1}) &\geq f(X_k, \mathcal{U}_k) - f(X_{k+1}, \mathcal{U}_k) \\ &\geq f(X_k, \mathcal{U}_k) - f(X_k^{(1)}, \mathcal{U}_k) \\ &\geq \beta \alpha_k^{(0)} \|\text{grad}_X f(X_k, \mathcal{U}_k)\|_{\text{F}}^2 \\ &\geq \beta \underline{\alpha} \|\text{grad}_X f(X_k, \mathcal{U}_k)\|_{\text{F}}^2, \end{aligned}$$

where we used Lemma 6.4 about Armijo steps. Summing over all iterations gives a telescopic sum on the left-hand side. For any $\bar{k} \geq 0$,

$$f(X_0, \mathcal{U}_0) - f_* \geq f(X_0, \mathcal{U}_0) - f(X_{\bar{k}}, \mathcal{U}_{\bar{k}}) \geq \beta \underline{\alpha} \sum_{k=0}^{\bar{k}} \|\text{grad}_X f(X_k, \mathcal{U}_k)\|_{\text{F}}^2. \quad (6.9)$$

The series is convergent since it is bounded independently of \bar{k} . Therefore

$$\lim_{k \rightarrow \infty} \|\text{grad}_X f(X_k, \mathcal{U}_k)\|_{\text{F}} = 0. \quad (6.10)$$

We have $\|\text{grad}_{\mathcal{U}} f(X_k, \mathcal{U}_k)\| = 0$ for all $k \geq 0$ since $\varepsilon_u = 0$. This corresponds to taking exact SVDs. Taking $k \rightarrow \infty$ in equation (6.8) gives convergence of the gradient norms to zero (6.6). □

Theorem 6.6 (Global complexity for Alternating minimization). *Let A9 hold for $f : \mathcal{M} \rightarrow \mathbb{R}$ from (3.1) or (3.5). Let $\varepsilon_x > 0$, $\varepsilon_u > 0$, and use Equation (5.5) or (5.6) to set $\varepsilon_{x,k}$. For any starting point $z_0 = (X_0, \mathcal{U}_0) \in \mathcal{M}$, Algorithm 2 produces a sequence $(X_k, \mathcal{U}_k)_{k \in \mathbb{N}}$ such that*

$$\left\| \left(\text{grad}_X f(X_k, \mathcal{U}_k), \text{grad}_{\mathcal{U}} f(X_k, \mathcal{U}_k) \right) \right\|_{\text{F}} \leq \varepsilon_x + \varepsilon_u \quad (6.11)$$

is achieved using at most N_{grad} gradient steps and N_{svd} singular value decompositions with

$$N_{\text{grad}} \leq \frac{(f(z_0) - f_*)}{\underline{\alpha} \beta \varepsilon_x^2} \quad \text{and} \quad N_{\text{svd}} \leq \frac{2L_u(f(z_0) - f_*)}{\varepsilon_u^2}, \quad (6.12)$$

where $\underline{\alpha} = \min \{ \alpha_0, 2\tau(1 - \beta)/L_x \}$ is a constant depending on parameters of the line search (5.1).

Proof. Note that f is bounded below by $f_* = 0$. Define N_{iter} as the number of iterations performed by Algorithm 2, the smallest k such that $\|\text{grad}_X f(X_k, \mathcal{U}_k)\| \leq \varepsilon_x$ and $\|\text{grad}_{\mathcal{U}} f(X_k, \mathcal{U}_k)\| \leq \varepsilon_u$. Let N_{svd} be the number of SVDs that have to be performed to get $\|\text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_k)\| \leq \varepsilon_u$, at which point the algorithm would return without performing another computation. For any $k \leq N_{svd}$, from Lemma 6.3 we have

$$f(X_k, \mathcal{U}_k) - f(X_{k+1}, \mathcal{U}_{k+1}) \geq \frac{1}{2L_u} \|\text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_k)\|_{\text{F}}^2 \geq \frac{1}{2L_u} \varepsilon_u^2. \quad (6.13)$$

Summing from $k = 0$ to N_{svd} gives,

$$f(z_0) - f_* \geq f(z_0) - f(z_{N_{svd}}) \geq \sum_{k=0}^{N_{svd}} \frac{\varepsilon_u^2}{2L_u} = \frac{\varepsilon_u^2 N_{svd}}{2L_u}. \quad (6.14)$$

Hence, this bounds the number of SVDs to ensure $\|\text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_k)\| \leq \varepsilon_u$, as

$$N_{svd} \leq 2L_u \frac{(f(z_0) - f_*)}{\varepsilon_u^2}. \quad (6.15)$$

For $0 \leq i \leq n_k - 1$, we have $\|\text{grad}_X f(X_k^{(i)}, \mathcal{U}_k)\|_{\text{F}}^2 \geq \varepsilon_{x,k}^2$ by definition since the stopping criterion is $\|\text{grad}_X f(X_k^{(n_k)}, \mathcal{U}_k)\|_{\text{F}} \leq \varepsilon_{x,k}$. Combined with the Armijo decrease this gives

$$f(X_k^{(i)}, \mathcal{U}_k) - f(X_k^{(i+1)}, \mathcal{U}_k) \geq \alpha_k^{(i)} \beta \|\text{grad}_X f(X_k^{(i)}, \mathcal{U}_k)\|_{\text{F}}^2 \geq \alpha_k^{(i)} \beta \varepsilon_{x,k}^2. \quad (6.16)$$

We sum these bounds for the n_k gradient steps from X_k to X_{k+1} ,

$$\sum_{i=0}^{n_k-1} \left[f(X_k^{(i)}, \mathcal{U}_k) - f(X_k^{(i+1)}, \mathcal{U}_k) \right] \geq \sum_{i=0}^{n_k-1} \alpha_k^{(i)} \beta \varepsilon_{x,k}^2. \quad (6.17)$$

Using that the step sizes $\alpha_k^{(i)}$ are bounded below by $\underline{\alpha} = \min\{\alpha_0, 2\tau(1-\beta)/L_x\}$ (Lemma 6.4),

$$f(X_k, \mathcal{U}_k) - f(X_{k+1}, \mathcal{U}_k) \geq n_k \underline{\alpha} \beta \varepsilon_{x,k}^2 \quad \text{for all } k \geq 0. \quad (6.18)$$

The SVD is nonincreasing, meaning $f(X_k, \mathcal{U}_k) - f(X_{k+1}, \mathcal{U}_{k+1}) \geq f(X_k, \mathcal{U}_k) - f(X_{k+1}, \mathcal{U}_k)$. This yields,

$$f(X_k, \mathcal{U}_k) - f(X_{k+1}, \mathcal{U}_{k+1}) \geq n_k \underline{\alpha} \beta \varepsilon_{x,k}^2 \geq n_k \underline{\alpha} \beta \varepsilon_x^2 \quad \text{for all } k \leq N_{iter}, \quad (6.19)$$

as both (5.5) and (5.6) satisfy $\varepsilon_{x,k} \geq \varepsilon_x$. We sum once again over the iterations,

$$f(X_0, \mathcal{U}_0) - f_* \geq f(X_0, \mathcal{U}_0) - f(X_{N_{iter}+1}, \mathcal{U}_{N_{iter}+1}) \geq \sum_{k=0}^{N_{iter}} n_k \underline{\alpha} \beta \varepsilon_x^2. \quad (6.20)$$

We conclude that

$$\frac{(f(z_0) - f_*)}{\underline{\alpha} \beta \varepsilon_x^2} \geq \sum_{k=0}^{N_{iter}} n_k =: N_{grad}. \quad (6.21)$$

□

A similar algorithm using fixed step sizes for the update of X will also converge, provided the step sizes are small enough.

Corollary 6.7. *If the Armijo linesearch in Algorithm 2 is replaced by a gradient descent with constant stepsizes α satisfying $\alpha < \frac{2}{L_x}$, Algorithm 2 also converges*

$$\lim_{k \rightarrow \infty} \left\| \left(\text{grad}_X f(X_k, \mathcal{U}_k), \text{grad}_{\mathcal{U}} f(X_k, \mathcal{U}_k) \right) \right\|_{\mathbb{F}} = 0. \quad (6.22)$$

We also have the worst case bound

$$N_{grad} \leq \frac{L_x(f_0 - f_*)}{\alpha \varepsilon_x^2}. \quad (6.23)$$

Proof. We derive the usual descent lemma from Lipschitz continuity of the gradient. This gives

$$f(X_k - \alpha \text{grad}_X f(X_k, \mathcal{U}_k), \mathcal{U}_k) \leq f(X_k, \mathcal{U}_k) - \alpha \|\text{grad}_X f(X_k, \mathcal{U}_k)\|_{\mathbb{F}}^2 + \alpha^2 L_x / 2 \|\text{grad}_X f(X_k, \mathcal{U}_k)\|_{\mathbb{F}}^2 \quad (6.24)$$

which simplifies to

$$f(X_k, \mathcal{U}_k) - f(X_k^{(1)}, \mathcal{U}_k) \geq (\alpha - \alpha^2 L_x / 2) \|\text{grad}_X f(X_k, \mathcal{U}_k)\|_{\mathbb{F}}^2. \quad (6.25)$$

This bound replaces the Armijo decrease of Equation (6.5). The rest of the proofs from Theorems 6.5 and 6.6 holds verbatim with stepsize α for every iteration. Note that for $\alpha > 0$, the factor $(\alpha - \alpha^2 L_x / 2)$ is positive for $\alpha < 2/L_x$ and is maximized at $\alpha = 1/L_x$. \square

7 Convergence of the iterates using the Kurdyka-Lojasiewicz property

This section proves convergence of the sequence of iterates to a unique stationary point for a simplified version of the alternating minimization scheme. This section considers an algorithm where only one gradient step is performed in between the truncated SVDs (Algorithm 4). This is analogue to the algorithm described in [20] which does not provide theoretical convergence guarantees. Our observations indicate that Algorithm 4 is expected to behave similarly to Algorithm 2 in the limit. Asymptotically, there is usually only one gradient step needed between two truncated SVDs. It is only in the early iterations that Algorithm 2 differs by making several gradient steps in between SVDs. For the purpose of this theoretical section, we will assume that the singular value decompositions in Algorithm 4 are exact and not approximated or randomized. This corresponds to setting $\varepsilon_u = 0$ in Algorithm 2. This section is written using the notation of a feature matrix Φ as in problem (3.1), but the results apply similarly to problem (3.5) if one assumes that the Lipschitz condition A11 applies to a kernel K instead of Φ .

Algorithm 4 A simple alternating minimization scheme for Problem (3.1) or (3.5)

- 1: **Given:** The sensing matrix $A \in \mathbb{R}^{m \times ns}$, measurements $b \in \mathbb{R}^m$, a tolerance $\varepsilon_x > 0$, an estimation of $r = \text{rank}(\Phi(M))$.
 - 2: Set $k = 0$
 - 3: Find X_0 that satisfies $AX_0 = b$.
 - 4: $U_0 = \text{truncate_svd}(\Phi(X_0))$ ▷ Equation (5.3)
 - 5: **while** $\|\text{grad}_X f(X_k, \mathcal{U}_k)\|_F > \varepsilon_x$ **do**
 - 6: $\text{grad}_X f(X_k, \mathcal{U}_k) = \text{P}_{\text{TL}_{\mathcal{A},b}}(\nabla_X f(X_k, \mathcal{U}_k))$ ▷ Equation (4.7)
 - 7: $\alpha_k = \text{Armijo}((X_k, \mathcal{U}_k), -\text{grad}_X f(X_k, \mathcal{U}_k))$ ▷ Algorithm 3
 - 8: $X_{k+1} = X_k - \alpha_k \text{grad}_X f(X_k, \mathcal{U}_k)$
 - 9: $U_{k+1} = \text{truncate_svd}(\Phi(X_{k+1}))$ ▷ exact SVD, not randomized
 - 10: **end while**
 - 11: **Output:** (X_k, \mathcal{U}_k) such that $\|\text{grad}f(X_k, \mathcal{U}_k)\|_F \leq \varepsilon_x$.
-

Let us define a distance on the manifold \mathcal{M} (Equation (4.2)).

Definition 7.1 (Distance on \mathcal{M}). *Given two subspaces $\mathcal{U}_1, \mathcal{U}_2 \in \text{Grass}(N, r)$, the canonical angles θ_i for $i = 1, \dots, r$ are defined as $\theta_i = \cos^{-1}(\sigma_i)$ where σ_i are the r singular values of $U_1^\top U_2$, with $\text{range}(U_1) = \mathcal{U}_1$ and $\text{range}(U_2) = \mathcal{U}_2$. For all $\mathcal{U}_1, \mathcal{U}_2$ in $\text{Grass}(N, r)$ define $\text{dist}(\mathcal{U}_1, \mathcal{U}_2) := \sqrt{\sum_{i=1}^r \sin^2 \theta_i}$. This gives a distance on \mathcal{M} ,*

$$\text{dist}\left((X_1, \mathcal{U}_1), (X_2, \mathcal{U}_2)\right) := \sqrt{\|X_1 - X_2\|_F^2 + \sum_{i=1}^r \sin^2 \theta_i} \quad (7.1)$$

for all $(X_1, \mathcal{U}_1), (X_2, \mathcal{U}_2)$ in \mathcal{M} .

We will prove finite length of the sequence of iterates in this metric on \mathcal{M} . A more mainstream approach to define the distance between \mathcal{U}_1 and \mathcal{U}_2 on the Grassmann would be to use $\|\Theta\|_F$ instead of $\|\sin \Theta\|_F$, where $\Theta = \text{diag}(\theta_i)$ is the diagonal matrix containing the principal angles. We do so because the distance (7.1) makes it easier to derive perturbation bounds for the SVD and is equivalent to the usual distance.

The following assumption ensures a useful non-degeneracy of the spectrum of the feature matrix.

A10 (Gap between the singular values). *There exists $\delta > 0$ such that, the accumulation points of the sequence generated by Algorithm 4 satisfy*

$$\sigma_r(\Phi(X)) - \sigma_{r+1}(\Phi(X)) \geq \delta > 0. \quad (7.2)$$

This property ensures that the minimizer of the function $f(X, \cdot) : \text{Grass}(N, r) \rightarrow \mathbb{R}$ is well defined, i.e., that its truncated SVD is unique. As $\sigma_{r+1}(\Phi(X)) \geq 0$, this assumption also implies that

$$\sigma_r(\Phi(X)) \geq \delta > 0$$

In particular it means that we cannot overestimate the rank. If the true rank is $r-1$, then $\sigma_r = 0$ and the assumption does not hold. Let us stress that this is an artefact of the convergence proof and does not imply poor practical performance of the algorithm when the rank is overestimated. We investigate this in the numerics Section 9.4.5. We will need Assumption 10 to derive a Lipschitz continuity result on the SVD. We now show the two main lemmas (7.1 and 7.5), inspired by [3].

Lemma 7.1 (Gradient lower bound on iterates gap). *Assume that Algorithm 4 generates a bounded sequence of iterates. Then, there exists $\rho_2 > 0$ such that, for all $k \in \mathbb{N}$,*

$$\|\text{grad}f(X_{k+1}, \mathcal{U}_{k+1})\|_{\text{F}} \leq \rho_2 \text{dist}\left((X_{k+1}, \mathcal{U}_{k+1}), (X_k, \mathcal{U}_k)\right) \quad (7.3)$$

with $\rho_2 := 2(L_g + 1/\underline{\alpha})$ for some $L_g \geq 0$.

Proof. The expression

$$X_{k+1} = X_k - \alpha_k \text{grad}_X f(X_k, \mathcal{U}_k) \quad (7.4)$$

implies

$$\text{grad}_X f(X_{k+1}, \mathcal{U}_{k+1}) = (X_k - X_{k+1})/\alpha_k + \text{grad}_X f(X_{k+1}, \mathcal{U}_{k+1}) - \text{grad}_X f(X_k, \mathcal{U}_k) \quad (7.5)$$

Define the set $\tilde{S} = \text{cl}(\text{conv}((X_k)_{k \in \mathbb{N}}))$, the closure of the convex hull of the sequence of iterates, and $S = \tilde{S} \times \text{Grass}(N, r)$. We show that the vector field $\text{grad}f|_S: S \rightarrow \text{TM}$ is L_g -Lipschitz continuous in the sense of Definition A.1 for some $L_g \geq 0$. Since S is bounded and the Hessian is continuous, there exists $L_g \geq 0$ such that $\|\text{Hess}f(x)\| \leq L_g$ for all $x \in S$. By Proposition A.2, $\text{grad}f|_S$ is L_g -Lipschitz continuous. Using the triangular inequality and the fact that $\underline{\alpha}$ is a lower bound of α_k for all k gives

$$\begin{aligned} \|\text{grad}_X f(X_{k+1}, \mathcal{U}_{k+1})\|_{\text{F}} &\leq \|X_{k+1} - X_k\|_{\text{F}}/\underline{\alpha} + \|\text{grad}_X f(X_{k+1}, \mathcal{U}_{k+1}) - \text{grad}_X f(X_k, \mathcal{U}_k)\|_{\text{F}} \\ &\leq \text{dist}\left((X_{k+1}, \mathcal{U}_{k+1}), (X_k, \mathcal{U}_k)\right)/\underline{\alpha} + L_g \text{dist}\left((X_{k+1}, \mathcal{U}_{k+1}), (X_k, \mathcal{U}_k)\right) \\ &\leq (1/\underline{\alpha} + L_g) \text{dist}\left((X_{k+1}, \mathcal{U}_{k+1}), (X_k, \mathcal{U}_k)\right). \end{aligned}$$

This gives (7.3) recalling that, since $\text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_{k+1}) = 0$,

$$\|\text{grad}f(X_{k+1}, \mathcal{U}_{k+1})\|_{\text{F}} = \|\text{grad}_X f(X_{k+1}, \mathcal{U}_{k+1})\|_{\text{F}}. \quad \square$$

Further auxiliary results are needed.

Lemma 7.2 (Wedin's theorem [35]). *Let $Y, \check{Y} \in \mathbb{R}^{N \times s}$ with singular value decompositions*

$$Y = \sum_{i=1}^{\min(N,s)} \sigma_i u_i (v_i)^\top \quad \text{and} \quad \check{Y} = \sum_{i=1}^{\min(N,s)} \check{\sigma}_i \check{u}_i (\check{v}_i)^\top,$$

with $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{\min(N,s)}$ and similarly for \check{Y} . If there exists $\delta > 0$ such that

$$\min_{\substack{1 \leq i \leq r \\ r+1 \leq j \leq \min(N,s)}} |\check{\sigma}_i - \sigma_j| \geq \delta \quad (7.6)$$

and

$$\check{\sigma}_r \geq \delta,$$

then

$$\|\sin \Theta\|_{\text{F}}^2 \leq \frac{2 \|\check{Y} - Y\|_{\text{F}}^2}{\delta^2} \quad (7.7)$$

with Θ the matrix of the principal angles between $[u_1 \ u_2 \ \dots \ u_r]$ and $[\check{u}_1 \ \check{u}_2 \ \dots \ \check{u}_r]$.

The following lemma is a direct consequence of Wedin's theorem.

Lemma 7.3. Let $Y, \check{Y} \in \mathbb{R}^{N \times s}$. Consider the singular value decomposition of $Y = \sum_{i=1}^{\min(N,s)} \sigma_i u_i v_i^\top$, with $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{\min(N,s)}$. Let us also write $U_r := [u_1 \ u_2 \ \dots \ u_r]$, a matrix whose columns span the left principal subspace associated to the r largest singular values. Similarly, $\check{Y} = \sum_{i=1}^{\min(N,s)} \check{\sigma}_i \check{u}_i \check{v}_i^\top$, with $\check{\sigma}_1 \geq \check{\sigma}_2 \geq \dots \geq \check{\sigma}_{\min(N,s)}$. Let us also write $\check{U}_r := [\check{u}_1 \ \check{u}_2 \ \dots \ \check{u}_r]$. If there exists $\delta > 0$ such that $\sigma_r - \sigma_{r+1} \geq \delta$ and $\check{\sigma}_r - \check{\sigma}_{r+1} \geq \delta$, then

$$\text{dist}(\check{U}_r, U_r)^2 \leq \frac{2}{\delta^2} \|\check{Y} - Y\|_{\text{F}}^2,$$

where $\text{dist}(U_r, \check{U}_r) = \sqrt{\sum_{i=1}^r \sin(\theta_i)^2}$ (with θ_i the principal angles between U_r and \check{U}_r) is the distance between the subspaces U_r and \check{U}_r .

Proof. The result follows from the $\sin \Theta$ bound (7.7) in Wedin's theorem. Let us verify the assumptions. From the assumptions we know that $\sigma_r \geq \delta$ and $\check{\sigma}_r \geq \delta$. If Wedin's theorem does not apply, condition (7.6) is not satisfied and neither is it satisfied with the roles of Y and \check{Y} reversed. In that case, since there exists no $\delta > 0$ such that (7.6) holds, one must have $\sigma_i = \check{\sigma}_j$, for some $i \leq r, j \geq r+1$, and $\check{\sigma}_l = \sigma_m$, for $l \leq r, m \geq r+1$. However, since the singular values are ordered decreasingly, this gives:

$$\sigma_m \leq \sigma_i = \check{\sigma}_j \leq \check{\sigma}_l = \sigma_m,$$

which implies that there exists $i \leq r$ and $m \geq r+1$ such that

$$\sigma_m = \sigma_i = \check{\sigma}_j = \check{\sigma}_l.$$

This is a contradiction with $\sigma_r - \sigma_{r+1} \geq \delta$ and $\check{\sigma}_r - \check{\sigma}_{r+1} \geq \delta$. Therefore, these conditions guarantee that Wedin's theorem applies. \square

In the next lemma we combine the previous bound with the Lipschitz continuity of Φ .

A11 (Lipschitz continuity of the features). For Problem (3.1), there exists $L_\Phi \geq 0$ such that for any X_k, X_{k+1} produced by Algorithm 4, $\|\Phi(X_{k+1}) - \Phi(X_k)\|_{\text{F}} \leq L_\Phi \|X_{k+1} - X_k\|_{\text{F}}$. For Problem (3.5), there exists $L_K \geq 0$ such that $\|\mathbf{K}(X_{k+1}, X_{k+1}) - \mathbf{K}(X_k, X_k)\|_{\text{F}} \leq L_K \|X_{k+1} - X_k\|_{\text{F}}$.

If we assume that the sequence $(X_k)_{k \in \mathbb{N}}$ is bounded, which we do in the main result of this section (Theorem 7.8), then it is sufficient for the features and kernel to be locally Lipschitz in order for A11 to hold. We also note that if the sublevel set $\{(X, \mathcal{U}) \in \mathcal{M} : f(X, \mathcal{U}) \leq f(X_0, \mathcal{U}_0)\}$ is bounded, then the iterates are contained in a bounded set since Algorithm 4 is a descent method.

Lemma 7.4. Let A10 and A11 hold. It follows that

$$\text{dist}(\mathcal{U}_k, \mathcal{U}_{k+1})^2 \leq \frac{2L_\Phi^2}{\delta^2} \|X_{k+1} - X_k\|_{\text{F}}^2.$$

Proof. By definition of \mathcal{U}_k , Lemma 7.3 ensures that

$$\text{dist}(\mathcal{U}_k, \mathcal{U}_{k+1})^2 \leq \frac{2}{\delta^2} \|\Phi(X_{k+1}) - \Phi(X_k)\|_{\text{F}}^2. \quad (7.8)$$

Indeed, $\mathcal{U}_k = \text{truncate-svd}(\Phi(X_k))$ is composed of the r first left singular vectors of $\Phi(X_k)$. They are uniquely defined due to A10. The result then follows from the Lipschitz continuity of Φ . \square

This lemma allows us to show the following crucial result.

Lemma 7.5 (Sufficient decrease property). *Assume that A10 and A11 hold at X_k and X_{k+1} . Then, there exists $\rho_1 > 0$, independent of k , such that the iterates of Algorithm 4 satisfy*

$$f(X_k, \mathcal{U}_k) - f(X_{k+1}, \mathcal{U}_{k+1}) \geq \rho_1 \text{dist}\left((X_k, \mathcal{U}_k), (X_{k+1}, \mathcal{U}_{k+1})\right)^2. \quad (7.9)$$

Proof. From the Armijo decrease of Lemma 6.4,

$$f(X_k, \mathcal{U}_k) - f(X_{k+1}, \mathcal{U}_k) \geq \frac{\beta}{\alpha_0} \|X_k - X_{k+1}\|_{\mathbb{F}}^2, \quad (7.10)$$

where α_0 is the largest step allowed by the backtracking. Set $M^2 = 2L_{\mathbb{Q}}^2/\delta^2$. Using that $f(X_{k+1}, \mathcal{U}_{k+1}) \leq f(X_{k+1}, \mathcal{U}_k)$, we get

$$f(\mathcal{U}_k, X_k) - f(\mathcal{U}_{k+1}, X_{k+1}) \geq f(\mathcal{U}_k, X_k) - f(\mathcal{U}_k, X_{k+1}) \quad (7.11)$$

$$\geq \frac{\beta}{\alpha_0} \|X_{k+1} - X_k\|_{\mathbb{F}}^2 \quad (7.12)$$

$$= \frac{\beta}{\alpha_0(1+M^2)} (1+M^2) \|X_{k+1} - X_k\|_{\mathbb{F}}^2 \quad (7.13)$$

$$\geq \frac{\beta}{\alpha_0(1+M^2)} \left(\|X_{k+1} - X_k\|_{\mathbb{F}}^2 + \text{dist}^2(\mathcal{U}_k, \mathcal{U}_{k+1}) \right) \quad (7.14)$$

where the last inequality comes from Lemma 7.4. This establishes (7.9) with $\rho_1 := \frac{\beta}{\alpha_0(1+M^2)}$. \square

We now show convergence of the gradient norms to zero for Algorithm 4.

Corollary 7.6 (Global convergence for Algorithm 4). *Set $\varepsilon_x = 0$, for any starting point $z_0 = (X_0, \mathcal{U}_0) \in \mathcal{M}$, Algorithm 4 applied to (3.1) or (3.5) produces a sequence $(X_k, \mathcal{U}_k)_{k \in \mathbb{N}}$ such that*

$$\lim_{k \rightarrow \infty} \|\text{grad}f(X_k, \mathcal{U}_k)\|_{\mathbb{F}} = 0. \quad (7.15)$$

Proof. Using Lemmas 7.1 and 7.5 gives,

$$f(X_k, \mathcal{U}_k) - f(X_{k+1}, \mathcal{U}_{k+1}) \geq \rho_1 \text{dist}\left((X_k, \mathcal{U}_k), (X_{k+1}, \mathcal{U}_{k+1})\right)^2 \quad (7.16)$$

$$\geq \rho_1/\rho_2 \|\text{grad}f(X_{k+1}, \mathcal{U}_{k+1})\|_{\mathbb{F}}^2. \quad (7.17)$$

The telescopic sum is bounded above by $f(z_0)$ independently of k since f is nonnegative, which ensures $\lim_{k \rightarrow \infty} \|\text{grad}f(X_{k+1}, \mathcal{U}_{k+1})\|_{\mathbb{F}} = 0$. \square

The bounds in Lemmas 7.1 and 7.5 are standard and hold for most descent methods. The values ρ_1, ρ_2 depend on the specifics of the algorithm used [3].

We now define the Kurdyka-Lojasiewicz inequality on Riemannian manifolds, which was already introduced in [26].

Definition 7.2 (The Kurdyka-Lojasiewicz inequality). *A locally Lipschitz function $f : \mathcal{M} \rightarrow \mathbb{R}$ satisfies the Kurdyka-Lojasiewicz inequality at $x \in \mathcal{M}$ iff there exist $\eta \in (0, \infty)$, a neighbourhood $V \subset \mathcal{M}$ of x , and a continuous concave function $\kappa : [0, \eta] \rightarrow [0, \infty[$ such that*

- $\kappa(0) = 0$,
- κ is continuously differentiable on $(0, \eta)$,

- $\kappa' > 0$ on $(0, \eta)$,
- For every $y \in V$ with $f(x) < f(y) < f(x) + \eta$, we have

$$\kappa'(f(y) - f(x)) \|\text{grad}f(y)\| \geq 1.$$

If f satisfies the KL inequality at every point $x \in \mathcal{M}$ we call f a KL function.

Lemma 7.7. Let $\{a_k\}_{k \in \mathbb{N}}$ be a sequence of nonnegative numbers. If $\sum_{k=1}^{\infty} \frac{a_k^2}{a_{k-1}}$ converges, then

$\sum_{k=1}^{\infty} a_k$ converges as well.

Proof. This is a standard result. A proof can be found in [13, Lemma 4.1]. \square

Theorem 7.8. Assume that Algorithm 4 is applied to problem (3.1) or (3.5), for case study 2.1 or 2.3, generates a bounded sequence $(X_k, \mathcal{U}_k)_{k \in \mathbb{N}}$. If A10 and A11 hold, then, the sequence has finite length, that is

$$\sum_{k=1}^{\infty} \text{dist}\left((X_k, \mathcal{U}_k), (X_{k+1}, \mathcal{U}_{k+1})\right) < \infty. \quad (7.18)$$

Therefore $(X_k, \mathcal{U}_k)_{k \in \mathbb{N}}$ converges to a unique point (X_*, \mathcal{U}_*) , which is a critical point of f on \mathcal{M} .

Proof. For case-studies 2.1 and 2.3, the feature map or kernel is an algebraic or exponential function. These functions are known to be KL functions [3], so the cost function f is a KL function (Definition 7.2). For convenience, we write $z_k = (X_k, \mathcal{U}_k)$. Since the sequence $(z_k)_{k \in \mathbb{N}}$ is bounded, there is a subsequence $(z_{k_q})_{q \in \mathbb{N}}$ which converges to some $\bar{z} \in \mathcal{M}$. Let $\omega(z_0)$ denote the set of limit points for some starting point z_0 . The set $\omega(z_0)$ is bounded by assumption and clearly closed, therefore it is compact. We want to show that $\omega(z_0)$ is a singleton, i.e. $\omega(z_0) = \{\bar{z}\}$. The function f is continuous, which implies $\lim_{q \rightarrow \infty} f(X_{k_q}, \mathcal{U}_{k_q}) = f(\bar{z})$. Since $f(z_{k_q})_{q \in \mathbb{N}}$ is non-increasing, the function f is also constant on $\omega(z_0)$. Since f is a KL function, for every point $z \in \omega(z_0)$, there exists a neighbourhood V_z of z and a continuous concave function $\kappa_z : [0, \eta_z] \rightarrow [0, \infty[$ of class C^1 on $]0, \eta_z[$ with $\kappa_z(0) = 0$, $\kappa'_z > 0$ on $]0, \eta_z[$ such that, for all $y \in V(z)$ with $f(z) < f(y) < f(z) + \eta_z$, we have

$$\kappa'_z(f(y) - f(z)) \|\text{grad}f(y)\|_{\mathbb{F}} \geq 1. \quad (7.19)$$

By compactness of $\omega(z_0)$, we find a finite number of points $\bar{z}_1, \dots, \bar{z}_p \in \omega(z_0)$ such that $\cup_{i=1}^p V_{\bar{z}_i}$ covers $\omega(z_0)$. We choose $\varepsilon > 0$, such that $V := \{z \in \mathcal{M} : \text{dist}(z, \omega(z_0)) < \varepsilon\}$ is contained in $\cup_{i=1}^p V_{\bar{z}_i}$. Then, we set $\eta = \min_{i=1, \dots, p} \eta_{\bar{z}_i}$, $\kappa'(t) = \max_{i=1, \dots, p} \kappa'_{\bar{z}_i}(t)$ and $\kappa(t) = \int_0^t \kappa'(\tau) d\tau$. We claim that for every $z \in \omega(z_0)$, and $y \in V$, with $f(z) < f(y) < f(z) + \eta$, we have

$$\kappa'(f(y) - f(z)) \|\text{grad}f(y)\|_{\mathbb{F}} \geq 1. \quad (7.20)$$

Indeed, there exists some \bar{z}_i such that $y \in V_{\bar{z}_i}$. Then, from the definition of κ' and the fact that f is constant on $\omega(z_0)$,

$$\kappa'(f(y) - f(z)) \|\text{grad}f(y)\|_{\mathbb{F}} \geq \kappa'_{\bar{z}_i}(f(y) - f(\bar{z}_i)) \|\text{grad}f(y)\|_{\mathbb{F}} \geq 1. \quad (7.21)$$

For $\eta > 0$ given above, there exists k_0 such that for all $k > k_0$,

$$f(z_k) < f(\bar{z}) + \eta. \quad (7.22)$$

By definition of the accumulation points, there exists k_1 such that for all $k > k_1$,

$$\text{dist}(z_k, \omega(z_0)) < \varepsilon. \quad (7.23)$$

Since $\sigma_r(\Phi(X)) > \sigma_{r+1}(\Phi(X))$ for any X such that $(X, \mathcal{U}) \in \omega(z_0)$ (A10), by continuity of the singular values, there exists $\bar{\delta} > 0$ such that for all points z_k satisfying $\text{dist}(z_k, \omega(z_0)) < \bar{\delta}$, we have $\sigma_r(\Phi(X_k)) > \sigma_{r+1}(\Phi(X)_k)$. Again by definition, there exists k_2 such that for all $k > k_2$,

$$\text{dist}(z_k, \omega(z_0)) < \bar{\delta}. \quad (7.24)$$

For $k > l = \max\{k_0, k_1, k_2\}$, we have

$$\kappa'(f(z_k) - f(\bar{z})) \|\text{grad}f(z_k)\|_{\mathbb{F}} \geq 1. \quad (7.25)$$

Using $\|\text{grad}f(z_{k+1})\|_{\mathbb{F}} \leq \rho_2 \text{dist}(z_k, z_{k+1})$ (Equation (7.3)), gives

$$\kappa'(f(z_k) - f(\bar{z})) \geq \frac{1}{\rho_2 \text{dist}(z_{k-1}, z_k)}. \quad (7.26)$$

Concavity of κ gives

$$\kappa(f(z_k) - f(\bar{z})) - \kappa(f(z_{k+1}) - f(\bar{z})) \geq \kappa'(f(z_k) - f(\bar{z})) (f(z_k) - f(z_{k+1})). \quad (7.27)$$

Since $k > l \geq k_2$, we have that $\rho_1 \text{dist}^2(z_k, z_{k+1}) \leq f(z_k) - f(z_{k+1})$ by Equation (7.9),

$$\kappa(f(z_k) - f(\bar{z})) - \kappa(f(z_{k+1}) - f(\bar{z})) \geq \frac{1}{\rho_2 \text{dist}(z_{k-1}, z_k)} \rho_1 \text{dist}^2(z_k, z_{k+1}), \quad (7.28)$$

and so

$$\frac{\text{dist}^2(z_k, z_{k+1})}{\text{dist}(z_{k-1}, z_k)} \leq \frac{\rho_2}{\rho_1} \kappa(f(z_k) - f(\bar{z})) - \kappa(f(z_{k+1}) - f(\bar{z})) \quad (7.29)$$

For any $N > l$, we sum (7.29) for all $l \leq k \leq N$, using that the right hand side is a telescopic sum,

$$\begin{aligned} \sum_{k \geq l}^N \frac{\text{dist}^2(z_k, z_{k+1})}{\text{dist}(z_{k-1}, z_k)} &\leq \sum_{k \geq l}^N \frac{\rho_2}{\rho_1} \left[\kappa(f(z_k) - f(\bar{z})) - \kappa(f(z_{k+1}) - f(\bar{z})) \right] \\ &\leq \frac{\rho_2}{\rho_1} \left[\kappa(f(z_l) - f(\bar{z})) - \kappa(f(z_N) - f(\bar{z})) \right] \\ &\leq \frac{\rho_2}{\rho_1} \left[\kappa(f(z_l) - f(\bar{z})) - \kappa(f(\bar{z}) - f(\bar{z})) \right] \\ &= \frac{\rho_2}{\rho_1} \kappa(f(z_l) - f(\bar{z})), \end{aligned} \quad (7.30)$$

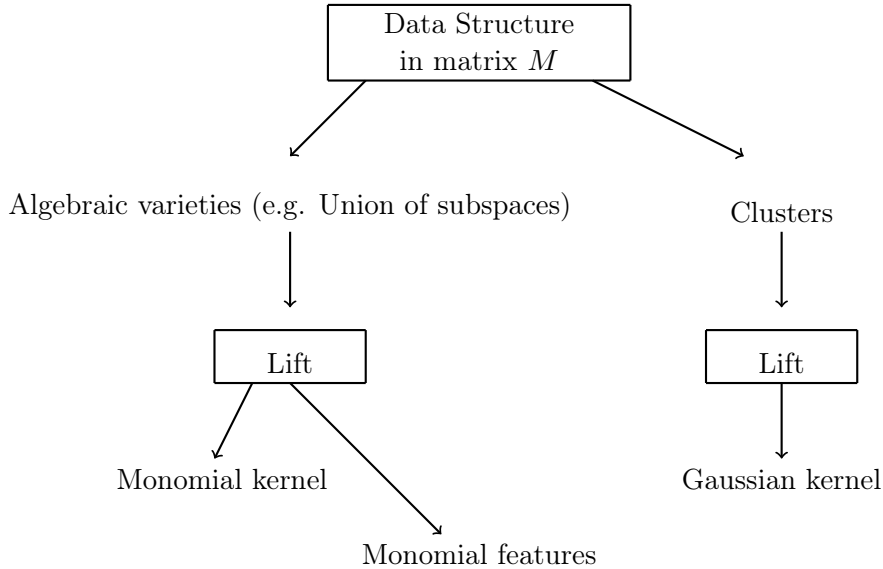
where we used that $f(\bar{z}) \leq f(z_N)$, κ is increasing and $\kappa(0) = 0$. Letting $N \rightarrow \infty$ in (7.30), we deduce that the left-hand side of (7.30) converges. By Lemma 7.7, $\sum_{k \geq l}^{\infty} \text{dist}(z_k, z_{k+1})$ also converges and therefore

$$\sum_{k=1}^{\infty} \text{dist}(z_k, z_{k+1}) < \infty. \quad (7.31)$$

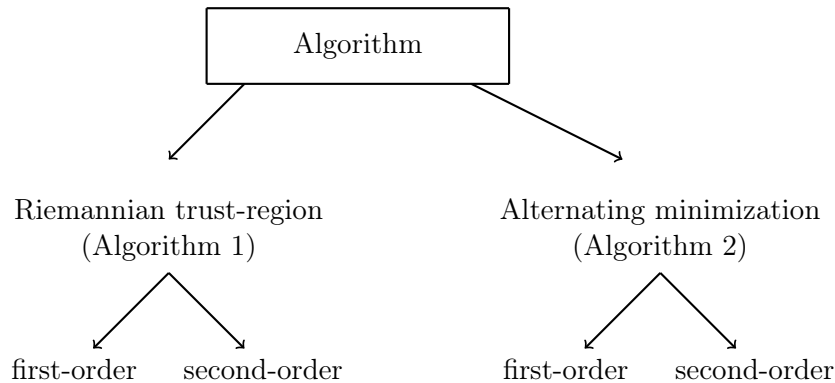
This concludes the proof and shows finite length of the sequence of iterates, which implies convergence of the Cauchy sequence $(X_k, \mathcal{U}_k)_{k \in \mathbb{N}}$ to a unique point (X_*, \mathcal{U}_*) . \square

8 Framework for nonlinear matrix recovery

We summarize the different components of the nonlinear matrix recovery problem. The matrix to be completed must be lifted to a higher dimensional space. This can be done through a kernel, in which case one solves problem (3.5), or a matrix of features, in which case one solves (3.1). When the matrix M to be recovered follows an algebraic variety model, one should use the monomial features or kernel (case study 2.1). When the data is scattered in clusters, the Gaussian kernel must be used as lifting (case study 2.3).



In addition, one needs to choose an algorithm to solve the problem formulation (3.1) or (3.5). We propose two families of algorithms: Alternating minimization (Algorithm 2) and Riemannian trust region (Algorithm 1). Each of these algorithms have first- and second-order variants, depending on whether they use the Hessian of the cost function.



9 Numerical experiments

In this section we validate our approach with numerical results on randomly generated test problems. We also compare the performances of the different algorithms we propose.

9.1 Implementation of the algorithms

Let us describe the implementation of the different algorithms and variants that are considered. **Altmin1** is a first-order version of alternating minimization (Algorithm 2) which uses gradient descent with Armijo linesearch to solve subproblem (5.1). It uses the monomial kernel (Equation (2.4)). The degree of the kernel that gives the best results is almost always $d = 2$. We set the constant $c = 1$ in the monomial kernel. In **Altmin2**, a second-order trust region method using the exact Hessian is applied to the minimization of (5.1). This is the only difference with **Altmin1**. The default values for the parameters of Algorithm 2 and the Gaussian and monomial kernels are presented in the table below.

Parameter	Default value	Parameter	Default value
$\varepsilon_x, \varepsilon_u$	10^{-6}	c in Equation (2.4)	1
$\varepsilon_{x,k}$	Equation (5.5)	α_0 in Algorithm 3	2
σ in Equation (2.10)	2.5	τ in Algorithm 3	0.5
d in Equation (2.4)	2	β in Algorithm 3	10^{-4}

Our code is available at <https://github.com/flgoyens/nonlinear-matrix-recovery> in both Matlab and Python. We use the Manopt [7] and Pymanopt [37] libraries for optimization on manifold solvers. The Riemannian trust-region **RTR2**, which implements Algorithm 1, is the corresponding Manopt solver for optimization on manifolds. We used a second-order version with the Hessian in the model and the default parameters of the solver. The maximum number of iterations is set at 500 for **RTR2**. In Manopt, the subproblems are solved with a truncated conjugate gradient method and the final termination criterion is only a first-order condition (the norm of the gradient) which we set at 10^{-6} for **RTR2**. Pymanopt uses automatic differentiation and does not require to compute the derivatives by hand, while the Manopt uses finite differences if the Hessian is not given as an input.

9.2 Test problems

We describe the set of parameters that we want to vary and test the dependence of each algorithm with respect to these parameters.

Union of subspaces Case study 2.2 depends on the following parameters: ambient dimension n , number of subspaces, dimension of each subspace, number of points on each subspace. To generate a random union of subspaces, we place the same number of points on each subspace and take subspaces of the same dimension. We calculate a basis for a random subspace and generate each point on that subspace by taking a random combination of the columns of that basis.

Clusters For case study 2.3, the parameters defining a point cloud divided in clusters in \mathbb{R}^n are: the number of clusters, the number of points in each cluster and the standard deviation σ_c of each cluster. We first generate random centres in \mathbb{R}^n . We then add to each centre a cluster of points with multivariate Gaussian distribution with zero mean and covariance $\sigma_c^2 \text{Id}$ with $\sigma_c = 0.5$.

9.3 Testing methodology

Throughout we say that an algorithm successfully recovers the matrix $M \in \mathbb{R}^{n \times s}$ if it returns a matrix X^{output} such that the root mean square error (RMSE) is below 10^{-3} ,

$$\text{RMSE}(M, X^{output}) := \|X^{output} - M\|_{\mathbb{F}} / \sqrt{ns} \leq 10^{-3}. \quad (9.1)$$

Our goal is to test the ability of our methods to recover the original matrix M . We measure the performance against an increase in difficulty of the problem for several parameters. Parameters that increase the difficulty of the recovery include:

1. Reducing the number of measurements m ;
2. Increasing the rank in the feature space.

In the case of unions of subspaces, for a fixed number of points, the rank of $\Phi_d(M)$ depends on the number and the dimension of the subspaces, as indicated by Proposition 2.1. For clusters, the rank increases with the number of clusters. The undersampling ratio is defined as

$$\delta = \frac{m}{ns}, \quad (9.2)$$

it is the number of measurements over the number of entries in M . We present phase transition results to numerically show which geometries can be recovered and which undersampling ratios are needed. Typical phase transition plots for matrix completion vary the undersampling ratio and the rank of the matrix [36]. For union of subspaces, the rank of the feature space is difficult to control, therefore we vary the number and dimension of the subspaces. For each value of the varying parameter, we generate 10 random matrices M that follow the desired structure. We try to recover each with varying δ from 0.1 to 0.9 for a random initial guess. If the RMSE is below 10^{-3} in the maximum number of iterations allowed by the algorithm, we consider the recovery to be successful. The phase transition plots record which of the 10 random problems is solved for each configuration. In Figures 4 through 9 the grayscale indicates the proportion of problems solved, with white = 100% of instances solved and black = 0%.

9.4 Numerical results

9.4.1 Comparing the performance of RTR and Alternating minimization

Figure 3 compares the performance of RTR2 (Algorithm 1 using a second-order Taylor model), `Altmin1` and `Altmin2` which are first- and second-order alternating minimization (Algorithm 2). We chose a problem of matrix completion over a union of subspaces. We find that RTR2 has a local quadratic rate of convergence, which makes it the method of choice if we want to recover M to high accuracy. Both `Altmin1` and `Altmin2` make faster progress during the early iterations; thus these methods should be considered if the required accuracy is low. We also observed that, in general, the distance to the solution M is of the same order of magnitude than the gradient norm. That is, using an algorithm, such as RTR2, which terminates with a smaller gradient norm yields a greater accuracy for the recovery. We noticed that the first-order methods, as well as the second-order `Altmin2`, typically stall numerically when the gradient norm is around 10^{-7} , but that is not the case for RTR2.

9.4.2 Recovery of unions of subspaces

We now illustrate how the parameters at play affect the recovery for data that follows a union of subspaces model.

Degree of the polynomial features Deciding which degree d to use in practice requires a careful choice. Previous works limit themselves to $d = 2$ and $d = 3$. This is understandable because the dimension of the features $N(n, d)$ increases exponentially with d , and so the dimension of the Grassmannian variable in (3.1) becomes too large and the problem becomes practically

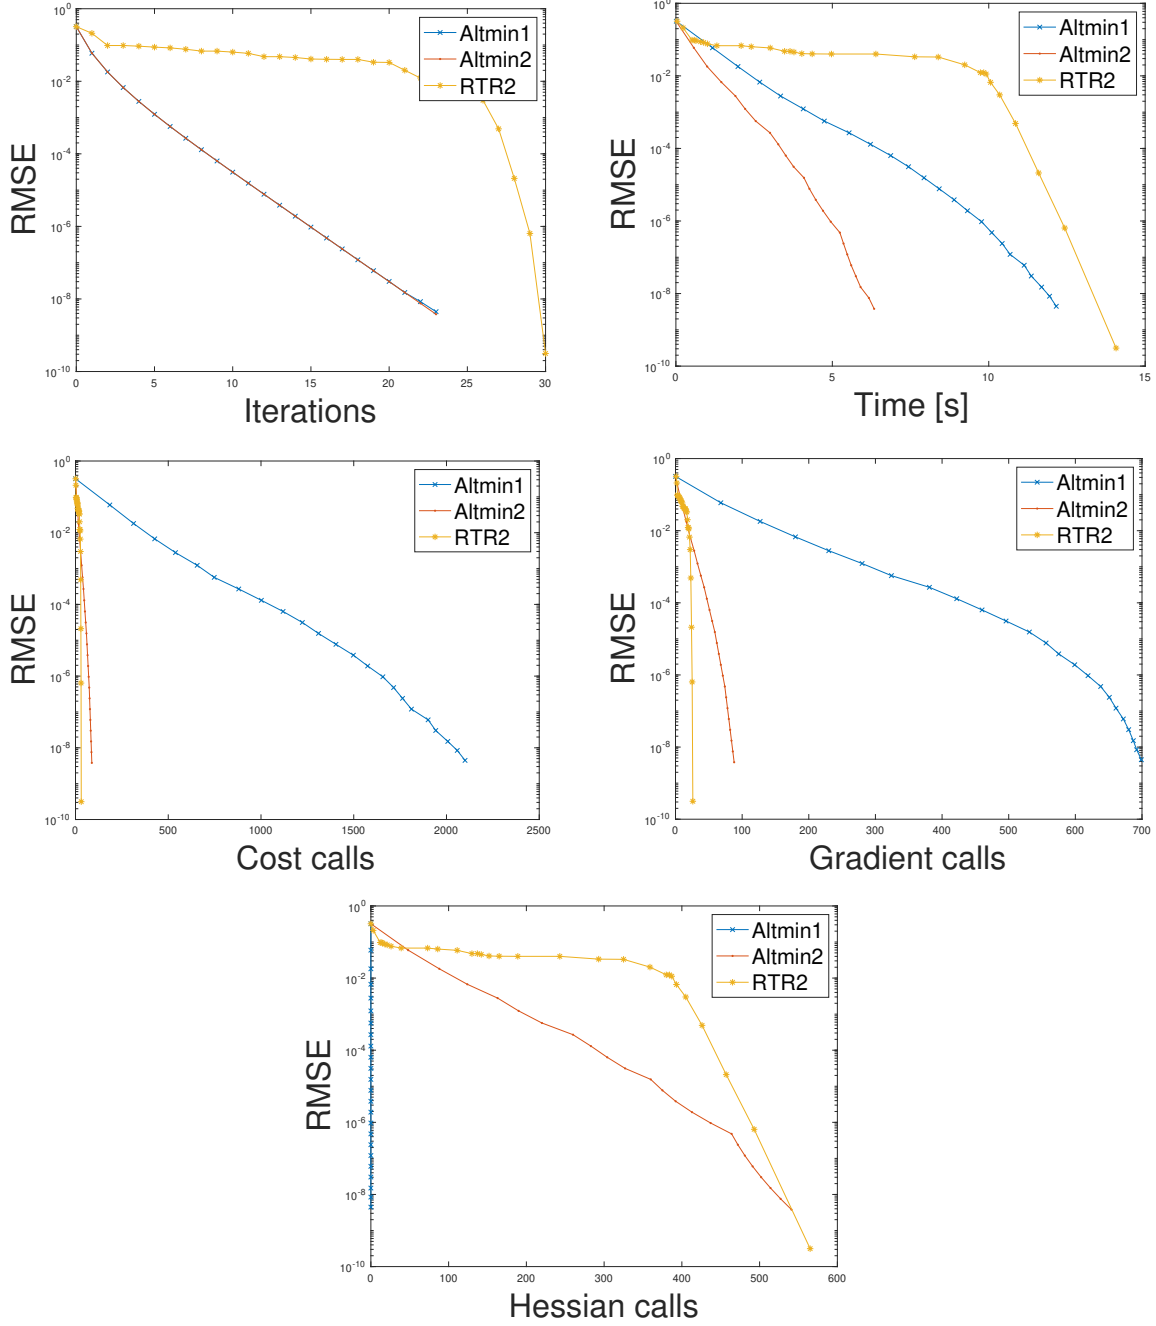


Figure 3: Comparing alternating minimization (first-order Altmin1 and second-order Altmin2) with the Riemannian trust-region algorithm (RTR2) for a union of subspaces recovery.

intractable for even moderate values of d and n . For example, $N(n = 20, d = 5) = 53130$ while $N(n = 20, d = 2) = 231$.

Therefore, the other natural option is to solve the kernel-based problem (3.5), where the dimension of the feature space $N(n, d)$ does not appear explicitly and the Grassmann has dimension $s \times r$. This is attractive because, a priori, s may not be as large as $N(n, d)$. However, there are important requirements on the number of samples needed to allow recovery. We need to ensure that $s \geq N - q$ where q is the number of linearly independent vectors v such that $v^\top \Phi_d(M) = 0$. That is, s needs to be large enough so that $\text{rank}(\Phi_d(M))$ is not limited by s but

the dimension of the variety. Some analysis in [32] shows that the number of points s needed to allow recovery increases exponentially with d . For that reason, if d is not small, it is not realistic to solve problems where s is large enough to enable recovery. The monomial basis is also known to be ill-conditioned for large degrees. This gives two obstacles to the performances of these algorithms when the degree increases.

In Figure 4, we solve the recovery problems using RTR2 with an increasing number of data points, and using monomial kernels of degree one, two and three to compare the recovery that is possible for each degree. In Figure 4(a), the degree used is $d = 1$. For $n = 15$, the dimension of the feature space is $N(15, 1) = 16$. For a large number of data points spread over 4 subspaces of dimension 2, the rank of the monomial kernel is 9. This explains why recovery is impossible when $s \leq 9$, since the kernel is not rank deficient at the solution M . In Figure 4(b), the degree used is $d = 2$. For $n = 15$, the dimension of the feature space $N(15, 2) = 136$. For a large number of data points spread over 4 subspaces of dimension 2, the rank of the monomial kernel is 21. This explains why recovery is impossible when $s \leq 21$, since the kernel is not rank deficient at the solution M . In Figure 4(c), the degree used is $d = 3$. For $n = 15$, the dimension of the feature space is $N(15, 3) = 816$. For a large number of data points spread over 4 subspaces of dimension 2, the rank of the monomial kernel is 37. This explains why recovery is impossible when $s \leq 37$, since the kernel is not rank deficient at the solution M . We notice that the recovery is still poor for $s \geq 37$, which is likely induced by a worse conditioning of the monomial embedding. In general, $d = 2$ seems to give the best results for the majority of data sets.

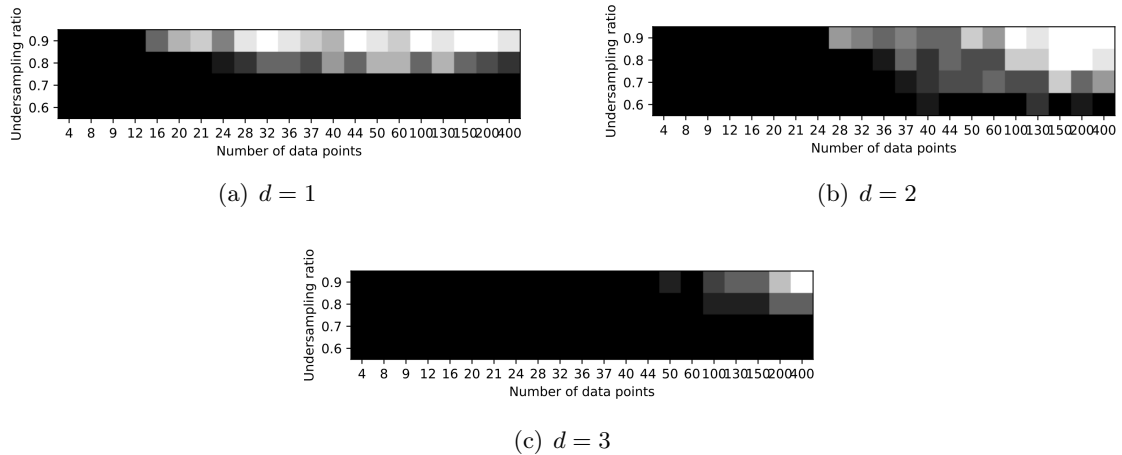


Figure 4: Phase transition for data belonging to a union of 4 subspaces of dimension 2 in \mathbb{R}^{15} for an increasing number of data points spread across the subspaces. Each square gives the proportion of problems solved over 50 randomly generated problems, white = 100% of instance recovered and black = 0%.

The dimension of the subspaces that we aim to recover plays a role in the possibility to recover. In Figure 5, we increase the dimension of the subspaces while the other parameters of the data remain fixed. For a fixed number of data points s , increasing the dimension of the subspaces increases the rank of the monomial features (see Proposition 2.1), and therefore, if the dimension of the subspaces becomes too large, the recovery is compromised. For 2 subspaces of dimension smaller than 4 in \mathbb{R}^{10} , we can observe good recovery depending on the undersampling ratio.

The same phenomenon is observed when the number of subspaces is increased for a fixed number of data points, see Figure 6.







Figure 5: Phrase transition for data belonging to a union of 2 subspaces of increasing dimension in \mathbb{R}^{10} with 20 points on each subspace. Each square gives the proportion of problems solved over 50 randomly generated problems, white = 100% of instance recovered and black = 0%, using the polynomial kernel of degree $d = 2$ for the embedding.



(a) Monomial kernel degree $d = 1$



(b) Monomial kernel degree $d = 2$



(c) Monomial kernel degree $d = 3$

Figure 6: Phrase transition for data belonging to a union of an increasing number of subspaces of dimension 2 in \mathbb{R}^{15} with 150 data points spread across the subspaces. Each square gives the proportion of problems solved over 10 randomly generated problems, white = 100% of instance recovered and black = 0%.

9.4.3 Clustering with missing data

In the case of clusters (case study 2.3), there is a noise inherent to the model because the kernel at the solution is only approximately low-rank. In fact, the matrix $K^G(M, M)$ has numerical full rank, but there is a big gap in the singular values. These matrices are notoriously difficult to recover in low rank matrix completion. For this reason the recovery error $\|M - X^*\|_F^2$ rarely converges to high accuracy and recovering the matrix up to 2 digits of accuracy is typical. This completed matrix X^* allows to do a clustering of the data points starting from missing entries. We are interested in determining when the matrix X^* has the same clustering as M . We use the Rand index to measure the compatibility of two different clusterings of the same set [33]. We can see in Figure 7 that for 5 clusters or less, the original clustering can be recovered even though up to 40% of the entries in the original matrix are missing.




Figure 7: Percentage of problems correctly clustered for different sampling rates and increasing number of clusters. 50 random instances generated in each case; white = 100% of instance correctly clustered and black = 0%. Clusters belong to \mathbb{R}^5 with 20 points in each cluster.

9.4.4 Robustness to measurement noise

In applications, it is common to assume some noise on the measurements, namely, $\langle A_i, M \rangle = b_i + \xi_i =: \tilde{b}_i$ where $\xi_i \in \mathbb{R}$ is some noise. In the following numerical test we generate white Gaussian noise, i.e. $\xi_i \sim \mathcal{N}(0, \sigma^2)$ for some variance σ^2 . The problem formulation then becomes Problem (3.2) with $\lambda > 0$ as the penalty parameter that should be tuned based on the noise level.

Estimating an appropriate value for λ without knowledge of the noise variance σ^2 is an intricate task. The solution of (3.2) for $\lambda \in [0, \infty[$ represents the trade-off curve between minimization of the rank residual and minimization of the residual on the linear measurements. In practical settings, a user may be able to determine which trade-off is more meaningful for a particular application. As a general strategy, we use a scheme which increases λ over successive calls to the solver, while warm starting each solve with the previous solution to (3.2). We have found that starting with the value $\lambda = 10^{-6}$ is satisfactory and we multiply λ by a factor 10 at each iteration. Figure 8 shows, for three different noise levels, the evolution of the solution of (3.2), labelled X^* , as the penalty parameter λ increases. We see that where the blue and red lines cross, the green line is still near its lowest point, that is, the solution is still minimizing the true measurement residual as well as for any other value of λ . This allows to recommend the simple strategy of choosing the value of λ where the values of the red and blue curves are the closest (which approximates the value for which they intersect). This choice gives equal weight

Standard deviation	$\ \mathcal{A}(X^*) - \tilde{b}\ $	$\ \mathcal{A}(X^*) - b\ $	$\ X^* - M\ $	$\ \mathcal{A}(M) - b\ $
$\sigma = 10^{-2}$	$2 \cdot 10^{-1}$	$7 \cdot 10^{-2}$	$8 \cdot 10^{-2}$	0.2
$\sigma = 10^{-3}$	$2 \cdot 10^{-2}$	$8 \cdot 10^{-3}$	$8 \cdot 10^{-3}$	0.02
$\sigma = 10^{-4}$	$2 \cdot 10^{-3}$	$8 \cdot 10^{-4}$	$9 \cdot 10^{-4}$	0.002

Table 1: Quality of the solution X^* for different levels of noise σ in the measurements.

to the rank minimization and satisfaction of the measurements.

Table 1 shows the accuracy of the solution X^* for that choice of λ . We see that both the infeasibility ($\|\mathcal{A}(X^*) - b\|$) and the distance to the solution ($\|X^* - M\|$) are proportional to the noise level and decreases with the later. This shows that the warm start scheme to find a good value for λ in conjunction with Problem (3.2) handles the presence of noise in the measurements very well.




Figure 8: Solutions for noisy problems as a function of the parameter λ on the horizontal axis.

9.4.5 Robustness to a bad estimate of the rank

In the case of a union of subspaces, the polynomial features are exactly low rank. Then, it is important to have an accurate upper bound on the rank. Recovery is sometimes possible if the upper bound is close to the correct value. If the estimated rank is less than the exact rank or much too large, recovery will normally fail. This intuition is guided by the cost function that we use. If the variable \mathcal{U} is artificially constrained to be the leading singular vectors of $\Phi(X)$, that

is $\mathcal{U} = \text{truncate-svd}(\Phi(X))$, then it can be substituted and the cost function in (3.1) simplifies to

$$\min_X \sum_{i=r+1}^{\min(N,s)} \sigma_i^2(\Phi(X)). \quad (9.3)$$

The cost function represents the energy in the tail of the singular value decomposition, where r is the estimation of the rank. In Figure 9(a), the data belongs to a union of 2 subspaces of dimension 2 in \mathbb{R}^{15} , with a total of $s = 150$ data points. With a kernel of degree 2, the dimension of the feature space is $N(15, 2) = 136$.




Figure 9: Impact of an incorrect estimate of the rank for the completion of a union of subspaces. Each square gives the proportion of problems solved over 50 randomly generated problems, white = 100% of instance recovered and black = 0%, using the polynomial kernel of degree $d = 2$ for the embedding.

9.4.6 Comparison with other methods

We compare the proposed methods `Altmin1` and `RTR2` with `VMC` (variety matrix completion) from [32], described in the related work section on page 3. We compare the methods on the recovery of a union of subspaces from a subset of entries, as `VMC` is designed for matrix completion. `RTR2` is a second-order method, while `Altmin1` and `VMC` are both first-order methods which are quite similar in spirit. They both alternate between truncated SVDs of the kernel matrix and some gradient steps, which are performed on different cost functions. `Altmin1` minimizes a smooth approximation of the Schatten p-norm (Equation (1.6)), while `Altmin1` minimizes Equation (5.1). `Altmin1` may perform several gradient steps between two SVD, while `VMC` performs a single gradient step between two SVD.

Figure 10 shows the decrease in root mean square error (RMSE) over time for the three methods on the completion of a matrix $M \in \mathbb{R}^{15 \times s}$ whose columns are contained in a union of two subspaces of dimension two. The total number of points (divided equally across each subspace) is taken as $s = \{100, 200, 400\}$. Figure 10(c) shows that `RTR2` clearly outperforms `VMC` and `Altmin1` in run-time for matrices with many columns (large s). For matrices with fewer columns (Figure 10(a)), `VMC` performs well in the early iterations in comparison with `RTR2`, but is consistently slower than `Altmin1`. Figure 10 indicates that for a comparable runtime, `VMC` performs many more iterations than `Altmin1` does. Each iteration of `VMC` is therefore faster to compute than an iteration of `Altmin1`, but they yield a smaller decrease in RMSE.






Figure 10: Comparison of the proposed `Altmin1` and `RTR2` methods with `VMC` from [32] on the recovery of a union of 2 subspaces of dimension 2 in \mathbb{R}^{15} with an under-sampling ratio of 0.9 and an increasing number of points s .

Figure 11 shows the runtime of the methods `VMC`, `Altmin1` and `RTR2` for an increasing ambient dimension n . Again, `Altmin1` is consistently faster than `VMC`; and we see that both first-order methods perform better than the second-order `RTR2` in the early iterations as n increases (Figure 11(c)).


Figure 12 compares the proportion of problems solved for a decreasing undersampling ratio (defined in Equation (9.2) as the ratio of observed entries over the size of the matrix to complete). The recovery rate indicates the proportion of problems solved over a set of 5 randomly generated problems of recovery of a matrix $M \in \mathbb{R}^{15 \times 100}$ whose columns belong to the union of two subspaces of dimension two. The methods `VMC` and `RTR2` perform slightly better than `Altmin1` at recovering the matrix M when the number of available entries decreases; though the difference is not significant.



(a) $n = 15$



(b) $n = 30$



(c) $n = 50$

Figure 11: Comparison of the proposed `Altmin1` and `RTR2` methods with `VMC` from [32] on the recovery of 100 data points belonging to a union of 2 subspaces of dimension 2 in \mathbb{R}^n for dimensions $n = \{15, 30, 50\}$ with an under-sampling ratio of 0.9.




Figure 12: Proportion of problems solved for decreasing under-sampling ratio

10 Conclusion

In this work, we study the problem of nonlinear matrix completion where one tries to recover a high rank matrix that exhibit low rank structure in a feature space. In terms of the use cases considered, in addition to the union of subspaces and algebraic varieties, we propose the use of the Gaussian kernel for clustering problems with missing data, which we believe is novel in the context of nonlinear matrix completion.

We propose a novel formulation for the nonlinear matrix completion problem using the Grassmann manifold, which is inspired from low-rank matrix completion techniques. We then show how Riemannian optimization and alternating minimization methods can be applied effectively to solve this optimization problem. The algorithms we propose, come with strong global convergence results to critical points and worst-case complexity guarantees. In addition, we show that the alternating minimization algorithm converges to a unique limit point using the Kurdyka-Lojasiewicz property.

We provide extensive numerical results that attest to the efficiency of the approach to recover high-rank matrices drawn from union of subspaces or clustered data. We note that the second-order Riemannian trust-region method allows to recover with high accuracy. We expose the difficulty of using polynomials of high degree in the monomial kernel, as they require an exponentially increasing number of sample points to allow recovery. Our approach proves to be efficient at clustering a data set despite the presence of missing entries and our approach also shows great robustness against the presence of noise in the measurements. Finally, we show that our algorithm greatly outperforms other code available online for nonlinear matrix completion.

Acknowledgement

The authors would like to thank Estelle Massart and Greg Ongie for interesting discussions and their helpful ideas.

References

- [1] P.-A. Absil, R. Mahony, and R. Sepulchre. Riemannian geometry of grassmann manifolds with a view on algorithmic computation. *Acta Applicandae Mathematica*, 80(2):199–220, 2004.
- [2] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2008.
- [3] J. Bolte, S. Sabach, and M. Teboulle. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1-2):459–494, July 2013.
- [4] N. Boumal. An introduction to optimization on smooth manifolds. To appear with Cambridge University Press, Jun 2022.
- [5] N. Boumal and P.-A. Absil. Low-rank matrix completion via preconditioned optimization on the grassmann manifold. *Linear Algebra and its Applications*, 475:200–239, 2015.
- [6] N. Boumal, P.-A. Absil, and C. Cartis. Global rates of convergence for nonconvex optimization on manifolds. *IMA Journal of Numerical Analysis*, 39(1):1–33, 2019.
- [7] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15:1455–1459, 2014.

- [8] P. Breiding, S. Kališnik, B. Sturmfels, and M. Weinstein. Learning algebraic varieties from samples. *Revista Matemática Complutense*, 31(3):545–593, 2018.
- [9] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, apr 2009.
- [10] A. R. Conn, N. I. Gould, and P. L. Toint. *Trust region methods*, volume 1. Siam, 2000.
- [11] D. Cox, J. Little, D. O’Shea, and M. Sweedler. Ideals, varieties, and algorithms. *American Mathematical Monthly*, 101(6):582–586, 1994.
- [12] M. A. Davenport and J. Romberg. An overview of low-rank matrix recovery from incomplete observations. *arXiv preprint arXiv:1601.06422*, 2016.
- [13] G. de Carvalho Bento, J. X. da Cruz Neto, and P. R. Oliveira. A new approach to the proximal point method: convergence on general riemannian manifolds. *Journal of Optimization Theory and Applications*, 168(3):743–755, 2016.
- [14] C. Eckart and G. Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.
- [15] A. Eftekhari, G. Ongie, L. Balzano, and M. B. Wakin. Streaming principal component analysis from incomplete data. *Journal of Machine Learning Research*, 20(86):1–62, 2019.
- [16] J. Fan and J. Cheng. Matrix completion by deep matrix factorization. *Neural Networks*, 98:34–41, 2018.
- [17] J. Fan and T. W. Chow. Non-linear matrix completion. *Pattern Recognition*, 77:378–394, 2018.
- [18] J. Fan and M. Udell. Online high rank matrix completion. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [19] J. Fan, C. Yang, and M. Udell. Robust non-linear matrix factorization for dictionary learning, denoising, and clustering, 2020.
- [20] J. Fan, Y. Zhang, and M. Udell. Polynomial matrix completion for missing data imputation and transductive learning, 2019.
- [21] J. Fan, M. Zhao, and T. W. S. Chow. Matrix completion via sparse factorization solved by accelerated proximal alternating linearized minimization. *IEEE Transactions on Big Data*, pages 1–1, 2018.
- [22] M. Fazel, H. Hindi, and S. Boyd. Rank minimization and applications in system theory. In *Proceedings of the 2004 American control conference*, volume 4, pages 3273–3278. IEEE, 2004.
- [23] M. Fornasier, H. Rauhut, and R. Ward. Low-rank matrix recovery via iteratively reweighted least squares minimization. *SIAM Journal on Optimization*, 21(4):1614–1640, 2011.
- [24] F. Goyens, S. Chretien, and C. Cartis. Smoothing of point clouds using riemannian optimization. *ICML Workshop Beyond first order methods in machine learning*, 2020.
- [25] N. Halko, P.-G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.

- [26] S. Hosseini. Convergence of nonsmooth descent methods via kurdyka-lojasiewicz inequality on riemannian manifolds. *Hausdorff Center for Mathematics and Institute for Numerical Simulation, University of Bonn (2015,(INS Preprint No. 1523))*, 2015.
- [27] P. Jain, R. Meka, and I. S. Dhillon. Guaranteed rank minimization via singular value projection. In *Advances in Neural Information Processing Systems*, pages 937–945, 2010.
- [28] L. Mirsky. Symmetric gauge functions and unitarily invariant norms. *The quarterly journal of mathematics*, 11(1):50–59, 1960.
- [29] K. Mohan and M. Fazel. Iterative reweighted algorithms for matrix rank minimization. *Journal of Machine Learning Research*, 13(Nov):3441–3473, 2012.
- [30] J. Nocedal and S. Wright. *Numerical optimization*. Springer Science & Business Media, 2006.
- [31] G. Ongie, D. Pimentel-Alarcón, L. Balzano, R. Willett, and R. D. Nowak. Tensor methods for nonlinear matrix completion. *SIAM Journal on Mathematics of Data Science*, 3(1):253–279, 2021.
- [32] G. Ongie, R. Willett, R. D. Nowak, and L. Balzano. Algebraic variety models for high-rank matrix completion. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2691–2700, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- [33] W. M. Rand. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical association*, 66(336):846–850, 1971.
- [34] A. Singer. From graph to manifold laplacian: The convergence rate. *Applied and Computational Harmonic Analysis*, 21(1):128–134, jul 2006.
- [35] G. W. Stewart. Perturbation theory for the singular value decomposition. Technical report, 1998.
- [36] J. Tanner and K. Wei. Normalized iterative hard thresholding for matrix completion. *SIAM Journal on Scientific Computing*, 35(5):S104–S125, jan 2013.
- [37] J. Townsend, N. Koep, and S. Weichwald. Pymanopt: A python toolbox for optimization on manifolds using automatic differentiation. *arXiv preprint arXiv:1603.03236*, 2016.
- [38] L. N. Trefethen and D. Bau III. *Numerical linear algebra*, volume 50. Siam, 1997.
- [39] B. Vandereycken. Low-rank matrix completion by riemannian optimization. *SIAM Journal on Optimization*, 23(2):1214–1236, jan 2013.
- [40] Z. Wen, W. Yin, and Y. Zhang. Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm. *Mathematical Programming Computation*, 4(4):333–361, 2012.

A Derivatives of cost functions

In this section we compute by hand the first- and second-order derivatives of the cost function that appears in the optimization problem (3.5). To compute the derivative of a matrix valued function, we write a Taylor expansion and identify the gradient by looking at first order terms. Before computing derivatives for a specific kernel, we look at the Lipschitz continuity properties of the gradient problem, for an arbitrary kernel.

A.1 Lipschitz properties

Let us consider the cost function of (3.5), defined over $\mathcal{M} = \mathcal{L}_{\mathcal{A},b} \times \text{Grass}(s, r)$. We work out conditions on the problem which ensure that the Riemannian gradient is Lipschitz continuous. Lipschitz continuity of a vector field on a smooth manifold is defined as follows using a vector transport.

Definition A.1. ([4, Definition 10.42]) *A vector field V on a connected manifold \mathcal{M} is L -Lipschitz continuous if, for all $x, y \in \mathcal{M}$ with $\text{dist}(x, y) < \text{inj}(x)$,*

$$\|\text{PT}_{0 \leftarrow 1}^\gamma V(y) - V(x)\| \leq L \text{dist}(x, y),$$

where $\gamma : [0, 1] \rightarrow \mathcal{M}$ is the unique minimizing geodesic connecting x to y and $\text{PT}_{0 \leftarrow 1}^\gamma$ denotes the parallel transport along γ .

Functions with Lipschitz continuous gradient exhibit the following regularity condition for the pullback $\hat{f} = f \circ \text{Retr}$, provided the retraction used is the exponential map, $\text{Retr} = \text{Exp}$.

Proposition A.1. ([4, Corollary 10.52]) *If $f : \mathcal{M} \rightarrow \mathbb{R}$ has L -Lipschitz continuous gradient, then*

$$|f(\text{Exp}_x(s)) - f(x) - \langle s, \text{grad}f(x) \rangle| \leq \frac{L}{2} \|s\|^2$$

for all (x, s) in the domain of the exponential map.

In order to show Lipschitz continuity of the gradient, we use both the definition, and the following proposition which uses an upper bound on the derivative of the gradient.

Proposition A.2. ([4, Corollary 10.45]) *If $f : \mathcal{M} \rightarrow \mathbb{R}$ is twice continuously differentiable on a manifold \mathcal{M} , then $\text{grad}f$ is L -Lipschitz continuous if and only if $\text{Hess}f(x)$ has operator norm bounded by L for all x , that is, if for all x we have*

$$\|\text{Hess}f(x)\| = \max_{\substack{x \in T_x \mathcal{M} \\ \|s\|=1}} \|\text{Hess}f(x)[s]\| \leq L.$$

We first compute the Euclidean gradient with respect to each variable,

$$\nabla_X f(X, \mathcal{W}) = \text{DK}(X)^* \text{P}_{\mathcal{W}^\perp} \tag{A.1}$$

and

$$\nabla_{\mathcal{W}} f(X, \mathcal{W}) = -2\text{K}(X)\mathcal{W}. \tag{A.2}$$

This naturally gives

$$\nabla_{\mathcal{W}\mathcal{W}}^2 f(X, \mathcal{W})[\Delta] = -2\text{K}(X)\Delta. \tag{A.3}$$

Proposition A.3. *Consider the cost function of (3.5) and assume that the retraction being used is the exponential map. If $\text{DK}(X)$ is Lipschitz continuous over $\mathcal{L}_{A,b}$, then (6.1) holds where L_x is the Lipschitz constant of $\text{DK}(X)$. If $\|\text{K}(X)\|_{\text{F}} \leq M$ for all $X \in \mathcal{L}_{A,b}$, condition (6.2) holds with $L_u = 2M$.*

Proof. For a given $X \in \mathcal{L}_{A,b}$, consider the function $f_{\mathcal{W}}(X, \cdot) : \text{Grass}(s, r) \rightarrow \mathbb{R}$. Its Riemannian Hessian is such that for $\Delta \in \text{T}_{\mathcal{W}}\text{Grass}(s, r)$, $\text{Hess}f_{\mathcal{W}}(\mathcal{W}, X)[\Delta] = -2\text{P}_{\mathcal{W}^\perp}\text{K}(x)\Delta$. Hence $\|\text{Hess}f_{\mathcal{W}}(\mathcal{W}, X)\| \leq \|k(X)\|_{\text{F}}$. Using [4] Corollary 10.45, the vector field $\text{grad}f_{\mathcal{W}}(X, \cdot)$ is Lipschitz continuous with constant $L_{\mathcal{W}} = 2\|k(x)\|_{\text{F}}$. If the kernel is upper-bounded for all $X \in \mathcal{L}_{A,b}$, the constant $L_{\mathcal{W}}$ is independent of X . This implies that (6.2) holds.

We also analyze Lipschitz continuity of the vector field $\text{grad}f_X$.

$$\|\text{grad}_X f(X_1, \mathcal{W}) - \text{grad}_X f(X_2, \mathcal{W})\|_{\text{F}} = \|\text{P}_{T\mathcal{L}_{A,b}}(\nabla_X f(X_1, \mathcal{W}) - \nabla_X f(X_2, \mathcal{W}))\|_{\text{F}} \quad (\text{A.4})$$

$$\leq \|\nabla_X f(X_1, \mathcal{W}) - \nabla_X f(X_2, \mathcal{W})\|_{\text{F}} \quad (\text{A.5})$$

$$\leq \|(\text{DK}(X_1)^* - \text{DK}(X_2)^*)\text{P}_{\mathcal{W}^\perp}\|_{\text{F}} \quad (\text{A.6})$$

$$\leq \|\text{DK}(X_1)^* - \text{DK}(X_2)^*\|_2 \|\text{P}_{\mathcal{W}^\perp}\|_{\text{F}} \quad (\text{A.7})$$

$$\leq \|\text{DK}(X_1) - \text{DK}(X_2)\|_2. \quad (\text{A.8})$$

If $\text{DK}(X)$ is L_x -Lipschitz over $\mathcal{L}_{A,b}$, we can write

$$\|\text{grad}_X f(X_1, \mathcal{W}) - \text{grad}_X f(X_2, \mathcal{W})\|_{\text{F}} \leq L_x \|X_1 - X_2\|_{\text{F}} \quad (\text{A.9})$$

and $\text{grad}_X f(\cdot, \mathcal{W})$ is also L_x -Lipschitz, where the constant L_x is independent of $\mathcal{W} \in \text{Grass}(s, r)$. This implies that (6.1) holds. \square

Proposition A.4. *Consider the cost function of either (3.1) or (3.5) and apply Algorithm 2 or Algorithm 1 with the exponential map as the retraction. If the convex hull of the sequence of iterates $(X_k)_{k \in \mathbb{N}}$ and the trial points is a bounded set, then (4.17), (4.18) and (6.2)-(6.1) hold at every iterate $(X_k, \mathcal{U}_k)_{k \in \mathbb{N}}$ and trial points of the algorithm.*

Proof. Provided the kernel is a smooth function, the derivatives of the cost function are continuous. As a consequence of the Weierstrass theorem, the derivatives are bounded on the closure of the convex hull of the iterates, which is compact ($\text{Grass}(s, r)$ is compact). If the Hessian is bounded on the closure of convex hull of the iterates, the gradient is Lipschitz continuous on that set (Proposition A.2) and therefore A3 and A9 hold with the exponential map as the retraction (Proposition A.1). The continuity of the third-order derivatives implies A4 in a similar way. \square

A.2 Monomial kernel

We wish to find the Euclidean derivatives of the cost function in (3.5) for the monomial kernel defined in (2.4). First, we make the following developments. Up to first order in Δ_X ,

$$\text{K}_d(X + \Delta_X) = \text{K}_d(X) + d\text{K}_{d-1}(X) \odot (X^\top \Delta_X + \Delta_X^\top X) + \mathcal{O}(\Delta_X^2). \quad (\text{A.10})$$

For $\Delta_W = W^\perp B$, we have

$$\begin{aligned} \text{P}_{W+\Delta_W} &= WW^\top + W\Delta_W^\top + \Delta_W W^\top, \\ &= \text{P}_W + W\Delta_W^\top + \Delta_W W^\top + \mathcal{O}(\Delta_W^2). \end{aligned} \quad (\text{A.11})$$

Let us write

$$f(X, W) = \text{tr}(\text{K}_d(X) - \text{P}_W \text{K}_d(X)).$$

We find the gradient in X using direct computation,

$$\begin{aligned}\nabla_X f(X, W) &= \nabla_X \text{tr}(\mathbf{P}_{W^\perp} \mathbf{K}_d(X)), \\ &= 2dX (\mathbf{K}_{d-1}(X) \odot \mathbf{P}_{W^\perp}),\end{aligned}\tag{A.12}$$

since \mathbf{P}_{W^\perp} is symmetric. Quite naturally we find $\nabla_W f(X, W)$ with the expansion,

$$\begin{aligned}f(X, W + \Delta_W) &= \text{tr}(\mathbf{K}_d(X) - \mathbf{P}_{W+\Delta_W} \mathbf{K}_d(X)), \\ &= \text{tr}\left(\mathbf{K}_d(X) - (\mathbf{P}_W + W\Delta_W^\top + \Delta_W W^\top) \mathbf{K}_d(X)\right), \\ &= \text{tr}\left(\mathbf{P}_{W^\perp} \mathbf{K}_d(X) - (W\Delta_W^\top + \Delta_W W^\top) \mathbf{K}_d(X)\right), \\ &= \text{tr}(\mathbf{P}_{W^\perp} \mathbf{K}_d(X)) - \text{tr}\left((W\Delta_W^\top + \Delta_W W^\top) \mathbf{K}_d(X)\right), \\ &= f(X, W) - \text{tr}\left(W\Delta_W^\top \mathbf{K}_d(X)\right) - \text{tr}\left(\Delta_W W^\top \mathbf{K}_d(X)\right), \\ &= f(X, W) + \langle \Delta_W, -2\mathbf{K}_d(X)W \rangle.\end{aligned}\tag{A.13}$$

Ans so we observe $\nabla_W f(X, W) = -2\mathbf{K}_d(X)W$. Quickly we have $\nabla_W^2 f(X, W)[E] = -2\mathbf{K}_d(X)E$. Now we want to find the second derivative in X

$$\begin{aligned}\nabla_X f(X + \Delta_X, W) &= 2d(X + \Delta_X) (\mathbf{K}_{d-1}(X + \Delta_X) \odot \mathbf{P}_{W^\perp}) \\ &= 2d(X + \Delta_X) \left(\left[\mathbf{K}_{d-1}(X) + (d-1)\mathbf{K}_{d-2}(X) \odot (X^\top \Delta_X + \Delta_X^\top X) \right] \odot \mathbf{P}_{W^\perp} \right) \\ &= 2dX (\mathbf{K}_{d-1}(X) \odot \mathbf{P}_{W^\perp}) + 2d(d-1)X \left(\mathbf{K}_{d-2}(X) \odot (X^\top \Delta_X + \Delta_X^\top X) \odot \mathbf{P}_{W^\perp} \right) \\ &\quad + 2d\Delta_X (\mathbf{K}_{d-1}(X) \odot \mathbf{P}_{W^\perp}) + \mathcal{O}(\Delta_X^2) \\ &= \nabla_X f(X, W) + 2d(d-1)X \left(\mathbf{K}_{d-2}(X) \odot (X^\top \Delta_X + \Delta_X^\top X) \odot \mathbf{P}_{W^\perp} \right) \\ &\quad + 2d\Delta_X (\mathbf{K}_{d-1}(X) \odot \mathbf{P}_{W^\perp}) + \mathcal{O}(\Delta_X^2).\end{aligned}\tag{A.14}$$

And so we identify

$$\nabla_X^2 f(X, W)[\Delta_X] = 2d(d-1)X \left(\mathbf{K}_{d-2}(X) \odot (X^\top \Delta_X + \Delta_X^\top X) \odot \mathbf{P}_{W^\perp} \right) + 2d\Delta_X (\mathbf{K}_{d-1}(X) \odot \mathbf{P}_{W^\perp}).\tag{A.15}$$

Now we need the cross derivatives $\nabla_W \nabla_X f(X, W)[\Delta_W]$ and $\nabla_X \nabla_W f(X, W)[\Delta_W]$.

$$\begin{aligned}\nabla_W f(X + \Delta_X, W) &= -2\mathbf{K}_d(X + \Delta_X)W \\ &= -2 \left(\mathbf{K}_d(X) + d\mathbf{K}_{d-1}(X) \odot (X^\top \Delta_X + \Delta_X^\top X) \right) W \\ &= -2\mathbf{K}_d(X)W - 2d \left(\mathbf{K}_{d-1}(X) \odot (X^\top \Delta_X + \Delta_X^\top X) \right) W \\ &= \nabla_W f(X, W) - 2d \left(\mathbf{K}_{d-1}(X) \odot (X^\top \Delta_X + \Delta_X^\top X) \right) W.\end{aligned}\tag{A.16}$$

So

$$\nabla_X \nabla_W f(X, W)[\Delta_X] = -2d \left(\mathbf{K}_{d-1}(X) \odot (X^\top \Delta_X + \Delta_X^\top X) \right) W \in s \times r.$$

Similarly we find

$$\begin{aligned}\nabla_X f(X, W + \Delta_W) &= 2dX \left(\mathbf{K}_{d-1}(X) \odot \mathbf{P}_{(W+\Delta_W)^\perp} \right) \\ &= 2dX \left(\mathbf{K}_{d-1}(X) \odot (\mathbf{P}_{W^\perp} - W\Delta_W^\top - \Delta_W W^\top) \right) \\ &= 2dX (\mathbf{K}_{d-1}(X) \odot \mathbf{P}_{W^\perp}) - 2dX \left(\mathbf{K}_{d-1}(X) \odot (W\Delta_W^\top - \Delta_W W^\top) \right) \\ &= \nabla_X f(X, W) - 2dX \left(\mathbf{K}_{d-1}(X) \odot (W\Delta_W^\top - \Delta_W W^\top) \right).\end{aligned}\tag{A.17}$$

And so

$$\nabla_W \nabla_X f(X, W) [\nabla_W] = -2dX \left(K_{d-1}(X) \odot (W \Delta_W^\top - \Delta_W W^\top) \right) \in n \times s.$$

In the end

$$\nabla^2 f(X, W) \begin{bmatrix} \Delta_X \\ \Delta_W \end{bmatrix} = \begin{pmatrix} \nabla_X^2 f(X, W) [\Delta_X] + \nabla_W \nabla_X f(X, W) [\nabla_W] \\ \nabla_X \nabla_W f(X, W) [\Delta_X] + \nabla_W^2 f(X, W) [\Delta_W] \end{pmatrix} \quad (\text{A.18})$$

A.3 Gaussian kernel

Consider the Gaussian kernel defined in (2.10). A direct computation gives

$$\nabla_X f(X, W) = -\frac{2}{\sigma^2} X \left(\text{diag} \left(\text{sum}(K^G \odot P_{W^\perp}, 1) \right) - K^G \odot P_{W^\perp} \right) \quad (\text{A.19})$$

for K^G the Gaussian kernel and $\text{sum}(K^G \odot P_{W^\perp}, 1)$ is the vector with the sum of each column of the matrix $K^G \odot P_{W^\perp}$. And similarly to the monomial kernel above, we have

$$\nabla_W f(X, W) = -2K^G(X)W. \quad (\text{A.20})$$

We do not compute the Hessian by hand for the Gaussian kernel. We either use automatic differentiation or finite differences of the gradient in the algorithm.

B Proofs for Section 6 (Convergence of the alternating minimization algorithm)

Lemma B.1. (Descent lemma based on [6, Theorem 4]) *Let A2 and A9 hold for $f : \mathcal{M} \rightarrow \mathbb{R}$. Then, for any $k \geq 0$, the iterates produced by Algorithm 2 satisfy*

$$f(X_k, \mathcal{U}_k) - f(X_{k+1}, \mathcal{U}_{k+1}) \geq \frac{1}{2L_u} \|\text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_k)\|^2, \quad (\text{B.1})$$

where L_u is the Lipschitz constant of the gradient of the pullback (A9).

Proof. We follow the development of [6, Theorem 4]. By Lipschitz continuity of the gradient we have,

$$|f(X_{k+1}, \text{Retr}_{\mathcal{U}_k}(\eta)) - [f(X_{k+1}, \mathcal{U}_k) + \langle \text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_k), \eta \rangle]| \leq \frac{L_u}{2} \|\eta\|^2 \quad \forall \eta \in T_{\mathcal{U}} \text{Grass}. \quad (\text{B.2})$$

Let $\eta = -\frac{1}{L_u} \text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_k)$ and define $\mathcal{U}^+ = \text{Retr}_{\mathcal{U}_k} \left(\frac{-1}{L_u} \text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_k) \right)$, which gives

$$f(X_{k+1}, \mathcal{U}^+) \leq f(X_{k+1}, \mathcal{U}_k) + \langle \text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_k), \frac{-1}{L_u} \text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_k) \rangle \quad (\text{B.3})$$

$$\begin{aligned} &+ \frac{L_u}{2} \left\| \frac{-1}{L_u} \text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_k) \right\|^2 \\ &\leq f(X_{k+1}, \mathcal{U}_k) - \frac{1}{2L_u} \|\text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_k)\|^2. \end{aligned} \quad (\text{B.4})$$

This gives

$$f(X_{k+1}, \mathcal{U}_k) - f(X_{k+1}, \mathcal{U}^+) \geq \frac{1}{2L_u} \|\text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_k)\|^2. \quad (\text{B.5})$$

Using that the singular value decomposition step of Algorithm 2 finds the minimum of $f(X_{k+1}, \cdot)$ over $\text{Grass}(N, r)$, we have $f(X_{k+1}, \mathcal{U}_{k+1}) \leq f(X_{k+1}, \mathcal{U}^+)$. Each update of the variable X is non-increasing, that is, $f(X_k, \mathcal{U}_k) \geq f(X_{k+1}, \mathcal{U}_k)$. Hence, we can conclude

$$f(X_k, \mathcal{U}_k) - f(X_{k+1}, \mathcal{U}_{k+1}) \geq f(X_{k+1}, \mathcal{U}_k) - f(X_{k+1}, \mathcal{U}_{k+1}) \geq \frac{1}{2L_u} \|\text{grad}_{\mathcal{U}} f(X_{k+1}, \mathcal{U}_k)\|^2. \quad (\text{B.6})$$

□

Lemma B.2. *Under A9, for the direction $-\text{grad}_X f(X_k^{(i)}, \mathcal{U}_k) \in T_{X_k} \mathcal{L}_{\mathcal{A}, b}$, the linesearch Algorithm 3 produces a step size $\alpha_k^{(i)}$ that satisfies*

$$\underline{\alpha} := \min \left\{ \alpha_0, \frac{2\tau(1-\beta)}{L_x} \right\} \leq \alpha_k^{(i)} \leq \alpha_0 \quad (\text{B.7})$$

and produces the following decrease

$$f(X_k^{(i)}, \mathcal{U}_k) - f(X_k^{(i+1)}, \mathcal{U}_k) \geq \beta \alpha \left\| \text{grad}_X f(X_k^{(i)}, \mathcal{U}_k) \right\|_{\text{F}}^2, \quad (\text{B.8})$$

where $X_k^{(i+1)} = X_k^{(i)} - \alpha_k^{(i)} \text{grad}_X f(X_k^{(i)}, \mathcal{U}_k)$.

Proof. It is clear from the algorithm that $\alpha_k^{(i)} \leq \alpha_0$. The Armijo condition also ensures (B.8). For any $\alpha > 0$, Lipschitz continuity of the gradient gives

$$f\left(X_k^{(i)} - \alpha \text{grad}_X f(X_k^{(i)}, \mathcal{U}_k), \mathcal{U}_k\right) \leq f(X_k^{(i)}, \mathcal{U}_k) - \alpha \left\| \text{grad}_X f(X_k^{(i)}, \mathcal{U}_k) \right\|_{\text{F}}^2 \quad (\text{B.9})$$

$$+ \alpha^2 \frac{L_x}{2} \left\| \text{grad}_X f(X_k^{(i)}, \mathcal{U}_k) \right\|_{\text{F}}^2. \quad (\text{B.10})$$

Hence the Armijo condition (B.8) is satisfied whenever

$$-\alpha \left\| \text{grad}_X f(X_k^{(i)}, \mathcal{U}_k) \right\|_{\text{F}}^2 + \alpha^2 \frac{L_x}{2} \left\| \text{grad}_X f(X_k^{(i)}, \mathcal{U}_k) \right\|_{\text{F}}^2 \leq -\alpha \beta \left\| \text{grad}_X f(X_k^{(i)}, \mathcal{U}_k) \right\|_{\text{F}}^2, \quad (\text{B.11})$$

which simplifies to

$$\alpha \leq \frac{2(1-\beta)}{L_x} =: \alpha_{max}. \quad (\text{B.12})$$

If α_0 satisfies Armijo, then $\alpha_k^{(i)} = \alpha_0$. Otherwise, we have $\alpha_k^{(i)} = \tau \alpha_l$ where $\alpha_l > \alpha_{max}$ is the last α that does not satisfy Armijo and $\alpha_{l+1} = \tau \alpha_l$ satisfies Armijo. In this case we have

$$\alpha_k^{(i)} \geq \tau \alpha_{max} = \frac{2\tau(1-\beta)}{L_x}. \quad \square$$