

MATHICSE Technical Report

Nr. 13.2014

February 2014 (NEW 04.2015)



Convergence results for projected line-search methods on varieties of low-rank matrices via Lojasiewicz inequality

Reinhold Schneider, André Uschmajew

CONVERGENCE RESULTS FOR PROJECTED LINE-SEARCH METHODS ON VARIETIES OF LOW-RANK MATRICES VIA ŁOJASIEWICZ INEQUALITY

REINHOLD SCHNEIDER[†] AND ANDRÉ USCHMAJEW[‡]

Abstract. The aim of this paper is to derive convergence results for projected line-search methods on the real-algebraic variety $\mathcal{M}_{\leq k}$ of real $m \times n$ matrices of rank at most k . Such methods extend Riemannian optimization methods, which are successfully used on the smooth manifold \mathcal{M}_k of rank- k matrices, to its closure by taking steps along gradient-related directions in the tangent cone, and afterwards projecting back to $\mathcal{M}_{\leq k}$. Considering such a method circumvents the difficulties which arise from the nonclosedness and the unbounded curvature of \mathcal{M}_k . The pointwise convergence is obtained for real-analytic functions on the basis of a Łojasiewicz inequality for the projection of the antigradient to the tangent cone. If the derived limit point lies on the smooth part of $\mathcal{M}_{\leq k}$, i.e. in \mathcal{M}_k , this boils down to more or less known results, but with the benefit that asymptotic convergence rate estimates (for specific step-sizes) can be obtained without an a priori curvature bound, simply from the fact that the limit lies on a smooth manifold. At the same time, one can give a convincing justification for assuming critical points to lie in \mathcal{M}_k : if X is a critical point of f on $\mathcal{M}_{\leq k}$, then either X has rank k , or $\nabla f(X) = 0$.

Key words. Convergence analysis, line-search methods, low-rank matrices, Riemannian optimization, steepest descent, Łojasiewicz gradient inequality, tangent cones

AMS subject classifications. 65K06, 40A05, 26E05, 65F30, 15B99, 15A83,

1. Introduction. This paper is concerned with line-search algorithms for low-rank matrix optimization. Let $k \leq \min(m, n)$. The set

$$\mathcal{M}_k = \{X \in \mathbb{R}^{m \times n} : \text{rank}(X) = k\}$$

of real rank- k matrices is a smooth submanifold of $\mathbb{R}^{m \times n}$. Thus, in order to approach a solution of

$$\min_{X \in \mathcal{M}_k} f(X), \tag{1.1}$$

where $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is continuously differentiable, one can use the algorithms known from Riemannian optimization, the simplest being the steepest descent method

$$X_{n+1} = R(X_n, \alpha_n \Pi_{T_{X_n} \mathcal{M}_k}(-\nabla f(X_n))). \tag{1.2}$$

Here, $\Pi_{T_{X_n} \mathcal{M}_k}$ is the orthogonal projection on the tangent space at X_n , $\alpha_n \geq 0$ is a step-size, and R is a *retraction*, which takes vectors from the affine tangent plane back to the manifold [2, 46]. Riemannian optimization on \mathcal{M}_k (and other matrix manifolds) has become an important tool for low-rank approximation in several applications, e.g. solutions of matrix equations such as Lyapunov equations, model reduction in machine learning, low-rank matrix completion, and others; see, for instance, [13, 37, 38, 45, 51, 50]. Typically, methods more sophisticated than steepest descent, such as nonlinear conjugate gradients, Newton's method, or line-search along

[†]Institut für Mathematik, Technische Universität Berlin, 10623 Berlin, Germany (schneidr@math.tu-berlin.de).

[‡]MATHICSE-ANCHP, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland. Current address: Hausdorff Center for Mathematics & Institute for Numerical Simulation, University of Bonn, 53115 Bonn, Germany (uschmajew@ins.uni-bonn.de).

geodesics, are employed. However, in most cases, convergence results of such line-search methods require the search directions to be sufficiently gradient-related.

An alternative interpretation of the projected gradient method (1.2) is that of a discretized gradient flow satisfying the Dirac–Frenkel variational principle, i.e., of the integration of the ODE

$$\dot{X}(t) = \Pi_{T_{X(t)}\mathcal{M}_k}(-\nabla f(X(t)))$$

using Euler’s explicit method with some step-size strategy. Therefore, our studies are also related to the growing field of dynamical low-rank approximation of ODEs [21, 41, 33] that admit a strict Lyapunov function.

The convergence analysis of sequences in \mathcal{M}_k is hampered by the fact that this manifold is not closed in the ambient space $\mathbb{R}^{m \times n}$. The manifold properties break down at the boundary which consists of matrices of rank less than k . It might happen that a minimizing sequence for (1.1) needs to cross such a singular point or even converge to it. Also, the effective domain of definition of a smooth retraction can become tiny at points close to singularities, leading to too small allowed step-sizes in theory. Even if these objections pose no serious problems in practice, they make it difficult to derive a priori convergence statements without making unjustified assumptions on the smallest singular values or adding regularization; cf. [25, 21, 22, 34, 50].

It certainly would be more convenient to optimize and analyze on the closure

$$\mathcal{M}_{\leq k} = \{X \in \mathbb{R}^{m \times n} : \text{rank}(X) \leq k\} \quad (1.3)$$

of \mathcal{M}_k , which is a real-algebraic variety. In many applications one will be satisfied with any solution of the problem

$$\min_{X \in \mathcal{M}_{\leq k}} f(X). \quad (1.4)$$

There is no principal difficulty in devising line-search methods on $\mathcal{M}_{\leq k}$. First, in singular points, one has to use search directions in the tangent cone (instead of tangent space), for instance, a projection of the antigradient¹ on the tangent cone. The tangent cones of $\mathcal{M}_{\leq k}$ are explicitly known [13], and projecting on them is easy (see Theorem 3.2 and Corollary 3.3). Second, one needs a “retraction” that maps from the affine tangent cone back to $\mathcal{M}_{\leq k}$, a very natural choice being a metric projection

$$R(X_n + \Xi) \in \underset{Y \in \mathcal{M}_{\leq k}}{\text{argmin}} \|Y - (X_n + \Xi)\|_F \quad (1.5)$$

(here in Frobenius norm), which can be calculated using singular value decomposition. The aim of this paper is to develop convergence results for such a method based on a Lojasiewicz inequality for the projected antigradient.

Convergence analysis of gradient flows based on the Lojasiewicz gradient inequality [32], or on the more general Lojasiewicz–Kurdyka inequality [26, 8, 9], has attracted much attention in nonlinear optimization during recent years [1, 4, 5, 6, 7, 12, 27, 28, 29, 30, 31, 36, 40, 52]. In part, this interest seems to have been triggered by the paper [1], where the following theorem was proved.

THEOREM. *Let $f : \mathbb{R}^N \rightarrow \mathbb{R}$ be continuously differentiable, and let $(x_n) \subseteq \mathbb{R}^N$ be a sequence of iterates satisfying the strong descent conditions*

$$\begin{aligned} f(x_{n+1}) - f(x_n) &\leq -\sigma \|\nabla f(x_n)\| \|x_{n+1} - x_n\| \quad (\text{for some } \sigma > 0), \\ f(x_{n+1}) = f(x_n) &\Rightarrow x_{n+1} = x_n. \end{aligned} \quad (1.6)$$

¹We use this terminology for the negative gradient $-\nabla f$ throughout the paper.

Assume also that the sequence possesses a cluster point x^* that satisfies the Lojasiewicz gradient inequality; i.e., there exist $\theta > 0$ and $\Lambda > 0$ such that

$$|f(y) - f(x^*)|^{1-\theta} \leq \Lambda \|\nabla f(y)\| \quad (1.7)$$

for all y in some neighborhood of x^* . Then x^* is the limit of the sequence (x_n) .

It is possible to obtain a stronger result if a small step-size safeguard of the form $\|x_{n+1} - x_n\| \geq \kappa \|\nabla f(x_n)\|$ (for some $\kappa > 0$) can be assumed. Not only can one then conclude that the limit x^* is a critical point of f , but the asymptotic convergence rate in terms of the Lojasiewicz parameters θ and Λ also can be estimated along lines developed, e.g., in [4, 7, 30, 36]. No second-order information is required, but a linear convergence rate can only be established when $\theta = 1/2$, which in general cannot be checked in advance. The most notable class of functions satisfying the Lojasiewicz gradient inequality in every point are real-analytic functions. Therefore, this type of results can be applied to classical line-search algorithms in \mathbb{R}^N when minimizing a real-analytic function using an angle condition for the search directions and Wolfe conditions for the step-size selection [1].

The theory can be generalized to gradient flows on real-analytic manifolds. Lageman [28] considered descent iterations on Riemannian manifolds via families of local parametrizations, with retracted line-search methods like (1.2) being a special case of his setting. Convergence results were obtained by making regularity assumptions on the used family of parametrizations. Merlet and Nguyen [36] considered a discrete projected θ -scheme for integrating an ODE on a smooth embedded submanifold. They proved the existence of step-sizes ensuring convergence to a critical point via Lojasiewicz gradient inequality by assuming a *uniform* bound on the second-order terms in the Taylor expansion of the metric projection, i.e., a curvature bound for the manifold. The main problem with the noncompact submanifold \mathcal{M}_k , without which our work would be unnecessary, is that such an assumption is unjustified. The second-order term in the metric projection on \mathcal{M}_k scales like the inverse of the smallest singular value [21, 3], which gets arbitrarily large in the case when the iterates approach the boundary of \mathcal{M}_k . However, such a uniform bound for the projection is not needed if one is willing to sacrifice some more information on the constants in convergence rate estimates: if a gradient projection method (x_n) on a smooth manifold *is known* to converge to *some* point of it, one will have *some* curvature bound in the vicinity of the limit.

Therefore, our plan is this: via a version of the Lojasiewicz inequality for projections of the antigradients on tangent cones we prove that the iterates of a line-search method on $\mathcal{M}_{\leq k}$ with a particular choice of step-sizes *do* converge. This would not be possible, or would at least be much more involved, for a line-search method formally designed on \mathcal{M}_k for the reasons mentioned below. Once the existence of a limit is established we may assume it to lie in \mathcal{M}_k , in order to deduce that it is a critical point and to estimate the convergence rate. Due to the following insight (repeated as Corollary 3.4), such a full-rank assumption on the limit can be regarded as very natural, or even necessary in some cases, when aiming at critical-point convergence.

THEOREM. *Let $k \leq \min(m, n)$, and let $X^* \in \mathcal{M}_{\leq k}$ be a critical point of (1.4) (see section 2.1). Then either $\text{rank}(X^*) = k$ or $\nabla f(X^*) = 0$.*

Accordingly, it will be typically impossible to prove convergence to a rank-deficient critical point by a method which (in regular points) only “sees” projections of the gradient on tangent spaces. We therefore emphasize again that in our paper line-search methods on $\mathcal{M}_{\leq k}$ are *not* considered for the purpose of detecting or correcting an

overestimated target rank, and are also not capable of doing so. Instead, the idea behind this work can be summarized as follows: a line-search method on $\mathcal{M}_{\leq k}$ can deal with singular iterates in theory, although in the most likely cases it will not generate a single one in a real computation. Thus, in the end it will not differ from a line-search method on \mathcal{M}_k as used in practice, thereby establishing its convergence.

The problem of correct rank estimation remains but has been recently successfully addressed using rank-increasing algorithms [37, 47], in which the target rank is successively increased during the process. It turns out that our concept of gradient projection on the tangent cones of $\mathcal{M}_{\leq k}$ is also useful in justifying and understanding such rank-increasing strategies from a theoretical perspective; see [49] for further explanation.

Contributions and outline. The paper has two parts: in section 2 abstract convergence statements for line-search methods on closed sets are established. In section 3 these are applied to line-search methods on $\mathcal{M}_{\leq k}$ with real-analytic cost function. The following list highlights the results.

- Theorem 2.3 states an abstract convergence result for descent methods on closed sets $\mathcal{M} \subseteq \mathbb{R}^N$ under the assumption of a Łojasiewicz inequality for the projections of the antigradient to the tangent cones. As it follows more or less known lines, the proof is provided in the appendix.
- In section 2.4 we define line-search schemes using gradient-related search directions on tangent cones (Algorithm 1). The step-sizes are selected by backtracking to satisfy an Armijo-like rule. Our notion of a retraction (Definition 2.4) is tailored to tangential projections on algebraic varieties: in every fixed tangent direction it needs to be a first-order approximation of the identity, but in contrast to a smooth retraction on a smooth manifold, it is not required to have a uniform bound on the second-order terms. The main result is Corollary 2.11: if f is real-analytic, then any cluster point of the sequence generated by the method, in whose neighborhood \mathcal{M} forms a real-analytic submanifold, must be its limit, and a critical point of the problem.
- Section 3.1 is devoted to the tangent cones of $\mathcal{M}_{\leq k}$. We give a much shorter derivation of their structure (Theorem 3.2) compared to [13]. The projection on the tangent cone is a simple and feasible operation (Algorithm 2). When $\text{rank}(X) < k$, the norm of this projection can be estimated from below by the norm of the antigradient itself (Corollary 3.3). This implies the above a priori statement on the rank of critical points (Corollary 3.4).
- Finally, in sections 3.3 and 3.4 we consider two concrete line-search methods on $\mathcal{M}_{\leq k}$: the classical steepest descent method with projection (Algorithm 3) and a retraction-free method using search directions which do not leave $\mathcal{M}_{\leq k}$ (Algorithm 4). If f is real-analytic, pointwise convergence for both methods is guaranteed, but only when the limit has full rank k can one conclude that it is a critical point; see Theorems 3.9 and 3.10. We compare both algorithms for a toy example of matrix completion.

Currently, results are restricted to finite-dimensional spaces, and one has to expect that hidden constants and provable convergence rates deteriorate with the problem size. This is to be expected from a black-box tool like the Łojasiewicz inequality which cannot be easily extended to infinite dimension; cf. [17, 19]. The limitation to finite dimension has been disregarded in related works as well [12, 30]. On the other hand, Vandereycken [50], for instance, observed more or less dimension-independent convergence rates for matrix completion of synthetic data using nonlinear CG.

2. Convergence of gradient methods via Łojasiewicz inequality. Let $\mathcal{D} \subseteq \mathbb{R}^N$ be open, and $f : \mathcal{D} \rightarrow \mathbb{R}$. Throughout the paper—, unless something else is stated, we assume at least that

$$f \text{ is continuously differentiable and bounded below.} \quad (\mathbf{A0})$$

Together with f we consider the minimization problem

$$\min_{x \in \mathcal{M}} f(x) \quad (2.1)$$

on a *closed* subset $\mathcal{M} \subset \mathcal{D}$ and assume it to have a solution. By $\|\cdot\|$ we denote the usual Euclidean norm on \mathbb{R}^N .

2.1. Optimality condition. We recall the necessary first-order optimality conditions for problem (2.1) and introduce some further notation.

Let $x \in \mathcal{M}$. The *tangent cone* (also called *contingent cone*) at x is

$$T_x \mathcal{M} = \{\xi \in \mathbb{R}^N : \exists(x_n) \subseteq \mathcal{M}, (a_n) \subseteq \mathbb{R}^+ \text{ s.t. } x_n \rightarrow x, a_n(x_n - x) \rightarrow \xi\}; \quad (2.2)$$

see, e.g., [15, 44]. It is a closed cone. Since it is in general not convex, a metric projection onto $T_x \mathcal{M}$ may not be uniquely defined. However, if we let $y \in \mathbb{R}^N$, then any $z \in T_x \mathcal{M}$ with $\|y - z\| = \text{dist}_{\|\cdot\|}(y, T_x \mathcal{M})$ is an orthogonal projection in the sense that

$$\|z\|^2 = \|y\|^2 - \|y - z\|^2 = \|y\|^2 - \text{dist}_{\|\cdot\|}(y, T_x \mathcal{M})^2. \quad (2.3)$$

Specifically, the norm of any such projection of the antigradient $-\nabla f(x)$ onto $T_x \mathcal{M}$ will be denoted by

$$g^-(x) = \sqrt{\|\nabla f(x)\|^2 - \text{dist}_{\|\cdot\|}(-\nabla f(x), T_x \mathcal{M})^2}.$$

An equivalent characterization, which resembles the norm of a restricted linear operator, is

$$g^-(x) = \max_{\substack{\xi \in T_x \mathcal{M} \\ \|\xi\| \leq 1}} -\nabla f(x)^\top \xi, \quad (2.4)$$

and the maximum is achieved if and only if ξ is a best approximation of $-\nabla f(x)$ in $T_x \mathcal{M}$, which then must have norm $\|\xi\| = g^-(x)$.

The *polar tangent cone*

$$T_x^\circ \mathcal{M} = \{y \in \mathbb{R}^N : y^\top \xi \leq 0 \text{ for all } \xi \in T_x \mathcal{M}\}$$

is always a closed convex cone. It equals the cone $\hat{N}_x \mathcal{M}$ of regular normal vectors at x [44, Definition 6.3 and Proposition 6.5]. The necessary first-order optimality condition for x^* to be a relative local minimum of f on \mathcal{M} is (see [15, Theorem 1] or [44, Theorem 6.12])

$$-\nabla f(x^*) \in T_{x^*}^\circ \mathcal{M} = \hat{N}_{x^*} \mathcal{M}. \quad (2.5)$$

Points with this property are called *critical*. By (2.4), x^* is critical if and only if

$$g^-(x^*) = 0.$$

This is the optimality condition we shall use in this paper.

In the case that $T_x \mathcal{M}$ is a linear space, $T_x^\circ \mathcal{M}$ is its orthogonal complement, $g^-(x)$ is the norm of the orthogonal projection of $\nabla f(x)$, and everything that has been said becomes quite evident. Moreover, if \mathcal{M} is a differentiable manifold and ∇f continuous, then g^- is continuous on \mathcal{M} . In general, it is not.

2.2. Łojasiewicz inequality. Our convergence results for line-search methods fundamentally rely on assuming the following property at a cluster point.

DEFINITION 2.1. *We say that $x \in \mathcal{M}$ satisfies a Łojasiewicz inequality for the projected antigradient if there exist $\delta > 0$, $\Lambda > 0$, and $\theta \in (0, 1/2]$ such that for all $y \in \mathcal{M}$ with $\|y - x\| < \delta$ it holds that*

$$|f(y) - f(x)|^{1-\theta} \leq \Lambda g^-(y). \quad (\mathbf{L})$$

As shown in the original work by Łojasiewicz [32, p. 92],² the classical, unconstrained Łojasiewicz gradient inequality

$$|f(y) - f(x)|^{1-\theta} \leq \Lambda \|\nabla f(y)\| \quad (2.6)$$

holds for the important class of real-analytic functions f . For this class, we can prove (\mathbf{L}) in the case that \mathcal{M} is locally the image of a real-analytic map (parametrization), for instance, as it is the case for the set $\mathcal{M}_{\leq k}$ (see Theorem 3.8). Basically, we need only apply the chain rule.

PROPOSITION 2.2. *Let f be real-analytic, $\mathcal{M} \subseteq \mathcal{D}$, and $x \in \mathcal{M}$. Assume there exists $M > 0$, an open set $\mathcal{N} \subseteq \mathbb{R}^M$, $t_0 \in \mathcal{N}$, and a (componentwise) real-analytic map $\tau : \mathcal{N} \rightarrow \mathbb{R}^N$ such that*

- (i) $\tau(\mathcal{N}) \subseteq \mathcal{M}$, $x = \tau(t_0)$, and
- (ii) *the image of every open neighborhood of t_0 (within \mathcal{N}) under τ contains a relatively open neighborhood of x within \mathcal{M} (in the induced topology of \mathbb{R}^N).*

Then (\mathbf{L}) holds at x .

Remark. A special case arises when \mathcal{M} is, at least in a neighborhood of x , an M -dimensional real-analytic submanifold of \mathbb{R}^N . In the terminology used in [24, Definition 2.7.1] this means that there exist such a map τ as in (ii) mapping *any* open subset of \mathcal{N} onto a relatively open subset of \mathcal{M} and having a derivative of rank M in any point.

Proof. The composition $f \circ \tau$ is real-analytic on \mathcal{N} [24, Proposition 2.2.8] and therefore satisfies the classical Łojasiewicz gradient inequality (2.6) in some open neighborhood $\tilde{\mathcal{N}} \subseteq \mathcal{N}$ of t_0 , that is,

$$|f(\tau(t)) - f(\tau(t_0))|^{1-\theta} \leq \tilde{\Lambda} \|\nabla(f \circ \tau)(t)\| \leq \tilde{\Lambda} \|\nabla\tau(t) \cdot \nabla f(\tau(t))\| \quad (2.7)$$

for all $t \in \tilde{\mathcal{N}}$. As τ maps on \mathcal{M} , it is easy to show that the derivative $\tau'(t) = [\nabla\tau(t)]^\top$ maps onto the tangent cone $T_{\tau(t)}\mathcal{M}$ at $\tau(t)$. Using (2.4), we deduce

$$\begin{aligned} \|\nabla\tau(t) \cdot \nabla f(\tau(t))\| &= \max_{\|h\|=1} -\nabla f(\tau(t))^\top \tau'(t)h \\ &\leq \max_{\|h\|=1} g^-(\tau(t)) \|\tau'(t)h\| = g^-(\tau(t)) \|\tau'(t)\|. \end{aligned} \quad (2.8)$$

Without loss of generality, we can assume that $\|\tau'(t)\| \leq C$ for all $t \in \tilde{\mathcal{N}}$. Combining (2.7) and (2.8) then proves (\mathbf{L}) for $\Lambda = \tilde{\Lambda}C$, since by (ii) there exists a $\delta > 0$ such that every $y \in \mathcal{M}$ with $\|y - x\| < \delta$ can be written as $y = \tau(t)$ with $t \in \tilde{\mathcal{N}}$. \square

The actual values of θ and Λ may depend on x and are typically not known. We will see below that the strongest convergence statements are achieved when $\theta = \frac{1}{2}$.

²In this original reference the statement is that there exists $\theta \in (0, 1)$ satisfying the inequality, but the proof shows that $\theta \in (0, 1/2]$.

The two generic cases in which this happens in the unconstrained version (2.6) (for δ small enough) are $\nabla f(x) \neq 0$ or $\nabla f(x) = 0$ with positive definite Hessian. Of these two cases, only the second is of interest, and we could make a similar statement in the setting of Proposition 2.2 by assuming the Hessian of $f \circ \tau$ to be positive definite at t_0 . However, this does not seem to be very useful. First, it is not clear how such an assumption could be related to more concrete conditions on f and \mathcal{M} . Second, in the context of line-search methods we shall consider below, the case of a positive definite Hessian at a cluster point could likely be treated by a more constructive local convergence analysis. In summary, the value of θ will remain unknown in our subsequent results.

2.3. General convergence theorem. Here we state a meta convergence theorem. Consider some iteration $(x_n) \subseteq \mathcal{M}$ that is intended to solve (2.1). Throughout the paper we will use the shorthand

$$f_n = f(x_n), \quad \nabla f_n = \nabla f(x_n), \quad g_n^- = g^-(x_n), \quad \text{and} \quad T_n \mathcal{M} = T_{x_n} \mathcal{M}.$$

Using this notation, we make the following assumptions.

- **Primary descent condition:** There exists $\sigma > 0$ such that for large enough n it holds that

$$f_{n+1} - f_n \leq -\sigma g_n^- \|x_{n+1} - x_n\|. \quad (\mathbf{A1})$$

- **Stationary condition:** For large enough n it holds that

$$g_n^- = 0 \quad \Rightarrow \quad x_{n+1} = x_n. \quad (\mathbf{A2})$$

- **Asymptotic small step-size safeguard:** There exists $\kappa > 0$ such that for large enough n it holds that

$$\|x_{n+1} - x_n\| \geq \kappa g_n^-. \quad (\mathbf{A3})$$

In combination with a Łojasiewicz inequality (**L**), these assumptions imply a fairly strong convergence result.

THEOREM 2.3. *Under assumptions (A1)–(A2), if there exists a cluster point x^* of the sequence (x_n) satisfying (L), it is actually its limit point. Further if (A3) holds, then the convergence rate can be estimated by*

$$\|x_n - x^*\| \lesssim \begin{cases} e^{-cn} & \text{if } \theta = \frac{1}{2} \text{ (for some } c > 0), \\ n^{-\frac{\theta}{1-2\theta}} & \text{if } 0 < \theta < \frac{1}{2}. \end{cases}$$

Moreover, $g_n^- \rightarrow 0$.

This theorem is an adaption of similar results scattered throughout the literature. Up to replacing the usual gradient by the projected antigradient, assumptions (A1), (A2) and (L) are the same as in [1, Theorem 3.2] and are sufficient to prove the convergence of the sequence (x_n) if it is bounded. (A2) is a natural technical requirement to the algorithm for not moving in the critical-point set and is typically satisfied if the iteration is gradient-related. Adding assumption (A3) does not only guarantee that the g_n^- tend to zero, but it allows us to estimate the convergence rate along known lines, e.g., [4, 30]. However, as (A3) is required here only for n larger than some unknown n_0 , one cannot determine the constants behind \lesssim explicitly (a constant depending on n_0 may be deduced).

Corresponding results for smooth manifolds have been obtained in [28, 29, 36]. In this context, we should remark that the ambient norm $\|x_{n+1} - x_n\|$, as used in **(A1)** and **(A3)**, is not necessarily a natural measure of distance on \mathcal{M} , but is particularly appropriate when the restriction to \mathcal{M} is motivated to reduce the complexity of a minimization problem in \mathbb{R}^N , as is typically the case for low-rank optimization.

Although no changes in the known arguments besides replacing $\|\nabla f\|$ by g^- are required, we give a proof of Theorem 2.3 in the appendix to keep the paper self-contained.

We emphasize that the theorem only states $g_n^- \rightarrow 0$. The question of when this actually implies $g^-(x^*) = 0$ is delicate, and simple counterexamples can be constructed. A sufficient condition would be $T_{x^*}\mathcal{M} \subset \liminf_{n \rightarrow \infty} T_n\mathcal{M}$ in the sense of set convergence (see, e.g., [44]). Unfortunately, this will usually not hold in the singular points of $\mathcal{M}_{\leq k}$ when approached by a sequence of full-rank matrices (Theorem 3.2). Later, we will be forced to make some smoothness assumptions on a neighborhood of x^* .

2.4. Retracted line-search methods. For line-search methods in \mathbb{R}^N it is well known how to obtain convergence results based on the Lojasiewicz gradient inequality [1]. Here we consider projected gradient flows on a set \mathcal{M} .

2.4.1. Retractions. Following [2], a retracted line-search method on a smooth manifold \mathcal{M} has the general form

$$x_0 \in \mathcal{M}, \quad x_{n+1} = R(x_n, \alpha_n \xi_n), \quad (2.9)$$

where ξ_n are tangent vectors at x_n , $\alpha_n \geq 0$, and $R : T\mathcal{M} \rightarrow \mathcal{M}$ is a *smooth retraction* [46]. This means that R is a C^∞ map which takes pairs (x, ξ_x) from the tangent bundle $T\mathcal{M}$ (which represent vectors $x + \xi_x$ on the affine tangent plane at x) back to the manifold, and has the property of being a first-order approximation of the exponential map, that is, its derivative at $(x, 0)$ with respect to ξ_x is the identity on $T_x\mathcal{M}$:

$$\lim_{T_x\mathcal{M} \ni \xi_x \rightarrow 0} \frac{\|R(x, \xi_x) - (x + \xi_x)\|}{\|\xi_x\|} = 0 \quad (2.10)$$

for all $x \in \mathcal{M}$. However, since we do not want to restrict ourselves to smooth manifolds, we make the following, more general definition.

DEFINITION 2.4 (retraction). *Let \mathcal{M} be closed. A map*

$$R : \bigcup_{x \in \mathcal{M}} \{x\} \times T_x\mathcal{M} \rightarrow \mathcal{M}$$

(where now $T_x\mathcal{M}$ is the tangent cone) will be called a retraction if for any fixed $x \in \mathcal{M}$ and $\xi_x \in T_x\mathcal{M}$ it holds that $\alpha \mapsto R(x, \alpha\xi_x)$ is continuous on $[0, \infty)$, and

$$\lim_{\alpha \rightarrow 0^+} \frac{R(x, \alpha\xi_x) - (x + \alpha\xi_x)}{\alpha} = 0. \quad (2.11)$$

The existence of such a retraction has implications for the regularity of the set \mathcal{M} . It is equivalent to the (one-sided) differentiability of the map $\alpha \mapsto \text{dist}_{\|\cdot\|}(x + \alpha\xi_x, \mathcal{M})$ in zero. This is, for instance, the case for real-algebraic varieties like $\mathcal{M}_{\leq k}$, and follows from the fact that for every tangent vector ξ_x to an algebraic variety, there exists an analytic arc $\gamma : [0, \epsilon) \rightarrow \mathcal{M}$ such that $\xi_x = \dot{\gamma}(0)$ [43, Proposition 2].

By (2.11), $R(x + \alpha\xi_x)$ is better than a first-order approximation of $x + \alpha\xi_x$ for very small $\alpha > 0$. In particular, for any fixed ξ_x and $\epsilon > 0$, (2.11) implies that

$$(1 - \epsilon)\alpha\|\xi_x\| \leq \|R(x, \alpha\xi_x) - x\| \leq (1 + \epsilon)\alpha\|\xi_x\| \quad \text{for sufficiently small } \alpha. \quad (2.12)$$

It means that a (small enough) step made in the tangent cone is neither increased nor decreased too much by the retraction, which obviously is of importance in analyzing a line-search method like (2.9). In what follows, we assume that we have a general upper bound for arbitrary steps:

$$\|R(x, \xi_x) - x\| \leq M\|\xi_x\| \quad \text{for all } x \in \mathcal{M} \text{ and } \xi_x \in T_x\mathcal{M}. \quad (2.13)$$

This imposes no serious restriction.

Since \mathcal{M} is assumed to be closed, a natural choice for R , though practically not always the most convenient, is the best approximation of $x + \xi_x$ in the Euclidean ambient norm (metric projection), that is,

$$R(x, \xi_x) \in \underset{y \in \mathcal{M}}{\operatorname{argmin}} \|y - (x + \xi_x)\|. \quad (2.14)$$

By the remarks above, this defines a valid retraction, for example, on closed real-algebraic varieties (cf. (3.9)) with $M = 2$ in (2.13). For the variety $\mathcal{M}_{\leq k}$ of bounded rank matrices one even can take $M = 1 + 2^{-1/2}$ (Proposition 3.6).

2.4.2. Angle condition. To obtain such strong convergence results as we have in mind, one naturally has to guarantee that the search directions ξ_n in (2.9) remain sufficiently gradient-related. We call $\xi_n \in T_n\mathcal{M}$ a *descent direction* if $\nabla f_n^T \xi_n < 0$.

DEFINITION 2.5 (angle condition). *Given $x_n \in \mathcal{M}$ and $\omega \in (0, 1]$, $\xi_n \in T_n\mathcal{M}$ is said to satisfy the ω -angle condition if*

$$\nabla f_n^T \xi_n \leq -\omega g_n^- \|\xi_n\|. \quad (2.15)$$

An equivalent statement is that the inner product between $-\nabla f_n / \|\nabla f_n\|$ and $\xi_n / \|\xi_n\|$ is at least $\omega g_n^- / \|\nabla f_n\|$.

For clarity, we emphasize the following.

PROPOSITION 2.6. *Any Euclidean best approximation*

$$\xi_n \in \underset{\xi \in T_n\mathcal{M}}{\operatorname{argmin}} \|\nabla f_n - \xi\|$$

of $-\nabla f_n$ on $T_n\mathcal{M}$ satisfies the ω -angle condition with $\omega = 1$. Moreover, with this choice, $\xi_n = 0$ if and only if $g_n^- = 0$.

Proof. As discussed in section 2.1, it holds in this case that $g_n^- = \|\xi_n\| = \sqrt{\|\nabla f_n\|^2 - \|\nabla f_n + \xi_n\|^2}$, which implies $\nabla f_n^T \xi_n = -g_n^- \|\xi_n\|$. \square

2.4.3. Armijo point. Given $x_n \in \mathcal{M}$ and a descent direction $\xi_n \in T_n\mathcal{M}$, we will have to pick a step-size α_n small enough to satisfy **(A1)**. It should, however, be as large as possible in order to hopefully guarantee **(A3)**.

DEFINITION 2.7 (Armijo point). *Let $\xi_n \in T_n\mathcal{M}$ be a descent direction at $x_n \in \mathcal{M}$, $\bar{\beta}_n > 0$, and $\beta, c \in (0, 1)$. The number*

$$\alpha_n = \max\{\beta^m \bar{\beta}_n : m \in \mathbb{N} \cup \{0\}, f(R(x_n, \beta^m \bar{\beta}_n \xi_n)) - f_n \leq c \beta^m \bar{\beta}_n \nabla f_n^T \xi_n\} \quad (2.16)$$

is called the Armijo point for $x_n, \xi_n, \bar{\beta}_n, \beta, c$.

This will be our choice for the step-size α_n in all subsequent algorithms. The importance of the Armijo point lies in the fact that in principle it can be found in finitely many steps using backtracking. To see that the maximum in (2.16) is not taken over the empty set, we introduce another important point:

$$\bar{\alpha}_n = \min\{\alpha > 0 : f(R(x_n, \alpha\xi_n)) - f_n = c\alpha\nabla f_n^\top \xi_n\}. \quad (2.17)$$

Then the following relations hold.

PROPOSITION 2.8. *Assume **(A0)**. Let $\xi_n \in T_n\mathcal{M}$ be a descent direction at $x_n \in \mathcal{M}$, and $\beta, c \in (0, 1)$. Then $\bar{\alpha}_n > 0$ exists, i.e., the minimum in (2.17) is not taken over the empty set, and $f(R(x_n, \alpha\xi_n)) - f_n \leq c\alpha\nabla f_n^\top \xi_n$ for all $\alpha \in [0, \bar{\alpha}_n]$. The Armijo point α_n defined by (2.16) satisfies*

$$\begin{aligned} \alpha_n &\geq \beta\bar{\alpha}_n && \text{if } \bar{\beta}_n > \bar{\alpha}_n, \\ \alpha_n &= \bar{\beta}_n && \text{if } \bar{\beta}_n \leq \bar{\alpha}_n. \end{aligned}$$

Proof. For convenience, let $\hat{R}(\alpha) = R(x_n, \alpha\xi_n)$ and $F(\alpha) = f_n + c\alpha\nabla f_n^\top \xi_n$. We have to show that $f(\hat{R}(\alpha))$ is strictly smaller than $F(\alpha)$ for sufficiently small $\alpha > 0$. By Taylor's theorem and (2.11),

$$\begin{aligned} f(\hat{R}(\alpha)) &= f_n + \nabla f_n^\top (\hat{R}(\alpha) - x_n) + o(\|\hat{R}(\alpha) - x_n\|) \\ &= f_n + \nabla f_n^\top (\alpha\xi_n + o(\alpha)) + o(\|\hat{R}(\alpha) - x_n\|) \\ &= f_n + c\alpha\nabla f_n^\top \xi_n + (1-c)\alpha\nabla f_n^\top \xi_n + o(\alpha) + o(\|\hat{R}(\alpha) - x_n\|), \end{aligned}$$

where $o(h)$ denotes a quantity with $o(h)/h = 0$ for $h \rightarrow 0^+$. By (2.12), the ratio $\|\hat{R}(\alpha) - x_n\|/\alpha$ converges to $\|\xi_n\|$ for $\alpha \rightarrow 0^+$, which implies $o(\|\hat{R}(\alpha) - x_n\|) = o(\alpha)$. As desired, it now follows that

$$\frac{1}{\alpha}(f(\hat{R}(\alpha)) - f_n - c\alpha\nabla f_n^\top \xi_n) = (1-c)\nabla f_n^\top \xi_n + \frac{o(\alpha)}{\alpha}$$

is negative for small enough α . Since $\alpha \mapsto \hat{R}(\alpha)$ is continuous and bounded below by **(A0)**, while $F(\alpha)$ is not, the smallest positive intersection point $\bar{\alpha}_n$ must exist. The assertions on α_n are immediate. \square

The role of the parameter $\bar{\beta}_n$ in Definition 2.7 is to adjust the initial length of the search direction ξ_n on which no assumptions have been made. When $\bar{\beta}_n\|\xi_n\|$ is too small, one has no chance to establish a minimum step-size safeguard like **(A3)**. The restriction we make is

$$\bar{\beta}_n \geq \min\left(\frac{g_n^-}{\|\xi_n\|}, \bar{\alpha}_n\right). \quad (2.18)$$

To achieve this, one needs to either calculate g_n^- , or increase the value of $\bar{\beta}_n$ until $f(R(x_n, \bar{\beta}\xi_n)) \geq f_n + c\bar{\beta}\nabla f_n^\top \xi_n$ holds.

2.4.4. Convergence results. The algorithm we analyze is formalized as Algorithm 1. By Propositions 2.6 and 2.8, all steps are feasible. We first assert that the mere convergence of the produced sequence (x_n) , when assuming the Łojasiewicz inequality, is guaranteed by Theorem 2.3.

Algorithm 1: Gradient-related projection method with line-search

Input: Starting point $x_0 \in \mathcal{M}$, $\omega \in (0, 1]$, $\beta \in (0, 1)$, $c \in (0, 1)$.

- 1 **for** $n=0, 1, 2, \dots$ **do**
- 2 Choose $\xi_n \in T_n \mathcal{M}$ satisfying (2.15), but choose $\xi_n = 0$ only when $g_n^- = 0$;
- 3 Choose $\bar{\beta}_n \geq \min(g_n^- / \|\xi_n\|, \bar{\alpha}_n)$; find Armijo point α_n for $x_n, \xi_n, \bar{\beta}_n, \beta, c$;
- 4 Form the next iterate

$$x_{n+1} = R(x_n, \alpha_n \xi_n).$$

5 **end**

COROLLARY 2.9. *Assume (A0). The sequence (x_n) produced by Algorithm 1 satisfies (A1) with $\sigma = \omega c M^{-1}$ (M being the constant from (2.13)) and (A2). Consequently, if a cluster point x^* exists and satisfies the Lojasiewicz gradient inequality (L), then $\lim_{n \rightarrow \infty} x_n = x^*$.*

Proof. Property (A1) follows immediately from (2.15) and (2.13); (A2) holds by construction. \square

Obviously, it is not necessary to choose the Armijo step-size to obtain this result; it suffices to have $f(R(x_n, \alpha_n \xi_n)) - f_n \leq c \alpha_n \nabla f_n^\top \xi_n$. The choice of the Armijo point becomes important, however, when one also aims for (A3) and the convergence rate estimate in Theorem 2.3. To proceed in this direction, we were not able to avoid imposing additional regularity assumptions on the retraction in the limit point.

THEOREM 2.10. *In the situation of Corollary 2.9, assume further that*

- (i) $\alpha_n \xi_n \rightarrow 0$, and
- (ii) *there exists a constant $C > 0$ such that for all sequences $(\hat{\xi}_n)$ with $\hat{\xi}_n \in T_n \mathcal{M}$ and $\hat{\xi}_n \rightarrow 0$ it holds that*

$$\limsup_{n \rightarrow \infty} \frac{\|R(x_n, \hat{\xi}_n) - (x_n + \hat{\xi}_n)\|}{\|\hat{\xi}_n\|^2} \leq C. \quad (2.19)$$

Assume further that f is bounded below on the whole of \mathbb{R}^N , and that there exists an open (in \mathbb{R}^N) neighborhood \mathcal{N} of x^ such that*

$$\|\nabla f(x) - \nabla f(y)\| \leq L \|x - y\| \quad \text{for all } x, y \in \mathcal{N}. \quad (2.20)$$

Then (A3) holds (with a generally unknown constant κ). Consequently, $g_n^- \rightarrow 0$, and the convergence rate estimates in Theorem 2.3 apply.

We discuss these two conditions after the proof.

Proof. We can assume $g_n^- > 0$ for all n , since otherwise the sequence becomes stationary. Then we have $\|\alpha_n \xi_n\| > 0$ for all n . We have to show that $\liminf_{n \rightarrow \infty} \|x_{n+1} - x_n\| / g_n^- > 0$. We do this by showing that the assumption $\liminf_{n \rightarrow \infty} \|x_{n+1} - x_n\| / g_n^- = 0$ leads to a contradiction. In the following we consider a subsequence which converges to the limes inferior, but for notational convenience we assume that

$$\lim_{n \rightarrow \infty} \frac{\|x_{n+1} - x_n\|}{g_n^-} = 0. \quad (2.21)$$

Fix $m \in (0, 1)$. As $\alpha_n \xi_n \rightarrow 0$, (2.19) implies that for large enough n we will have

$$\|\alpha_n \xi_n\| \leq \|x_{n+1} - x_n - \alpha_n \xi_n\| + \|x_{n+1} - x_n\| \leq C \|\alpha_n \xi_n\|^2 + \|x_{n+1} - x_n\|.$$

We consider n so large that $m \leq (1 - C\|\alpha_n \xi_n\|)$ or, after rearranging,

$$m\|\alpha_n \xi_n\| \leq \|x_{n+1} - x_n\|. \quad (2.22)$$

Since \mathcal{N} is open and $x_n \rightarrow x^* \in \mathcal{N}$ and $\hat{\xi}_n \rightarrow 0$, it also holds that

$$x_n + z \in \mathcal{N} \quad \text{for all } z \text{ with } \|z\| \leq M\beta^{-1}\|\hat{\xi}_n\| \quad (2.23)$$

if only n is large enough. Hence we may assume, without loss of generality, that (2.22) and (2.23) hold for all n . Now we distinguish the iterates by two disjoint cases: $\bar{\beta}_n \leq \bar{\alpha}_n$ and $\bar{\beta}_n > \bar{\alpha}_n$. In the first case, we have $\alpha_n = \bar{\beta}_n$ by Proposition 2.8, which by the choice of $\bar{\beta}_n$ in the algorithm according to (2.18) gives

$$\|x_{n+1} - x_n\| \geq m\|\alpha_n \xi_n\| \geq mg_n^-.$$

Assumption (2.21) implies that this happens only for finitely many n . Let us therefore assume that the second case $\bar{\beta}_n > \bar{\alpha}_n$ always occurs. In this case, Proposition 2.8 states that $\bar{\alpha}_n \leq \beta^{-1}\alpha_n$. Hence, by (2.22) and (2.21),

$$\lim_{n \rightarrow \infty} \frac{\|\bar{\alpha}_n \xi_n\|}{g_n^-} \leq \lim_{n \rightarrow \infty} \frac{m^{-1}\beta^{-1}\|x_{n+1} - x_n\|}{g_n^-} = 0. \quad (2.24)$$

We now show that (2.24) leads to a contradiction by mimicking arguments that are used to prove existence of step-sizes satisfying the strong Wolfe conditions in linear spaces, e.g., [39, Lemma 3.1]. Let again $\hat{R}(\alpha) = R(x_n, \alpha \xi_n)$. By the mean value theorem, there exists $\vartheta \in (0, 1)$ such that $z = \vartheta(\hat{R}(\bar{\alpha}_n) - x_n)$ satisfies

$$(\hat{R}(\bar{\alpha}_n) - x_n)^\top \nabla f(x_n + z) = f(\hat{R}(\bar{\alpha}_n)) - f_n = c\bar{\alpha}_n \xi_n^\top \nabla f_n, \quad (2.25)$$

where the second equality holds by definition (2.17). By (2.13), $\|z\| \leq M\|\bar{\alpha}_n \xi_n\| \leq M\beta^{-1}\|\alpha_n \xi_n\|$ so that $x_n + z \in \mathcal{N}$ by (2.23). Using (2.20), Cauchy–Schwarz, the definition of z , the reverse triangle inequality, and the angle condition (2.15), we can estimate:

$$\begin{aligned} \|z\| \|\hat{R}(\bar{\alpha}_n) - x_n\| &\geq L^{-1} \|\nabla f_n - \nabla f(x_n + z)\| \|\hat{R}(\bar{\alpha}_n) - x_n\| \\ &\geq L^{-1} |\nabla f_n^\top (\hat{R}(\bar{\alpha}_n) - x_n) - c\bar{\alpha}_n \nabla f_n^\top \xi_n| \\ &\geq L^{-1} ((1 - c)\omega g_n^- \|\bar{\alpha}_n \xi_n\| - |\nabla f_n^\top (\hat{R}(\bar{\alpha}_n) - (x_n + \bar{\alpha}_n \xi_n))|). \end{aligned}$$

Since we have $\|\bar{\alpha}_n \xi_n\| \geq M^{-1} \|\hat{R}(\bar{\alpha}_n) - x_n\| \geq M^{-1} \|z\|$, we arrive at

$$\frac{\|\bar{\alpha}_n \xi_n\|}{g_n^-} \geq M^{-2} L^{-1} \left((1 - c)\omega - \frac{|\nabla f_n^\top (\hat{R}(\bar{\alpha}_n) - (x_n + \bar{\alpha}_n \xi_n))|}{g_n^- \|\bar{\alpha}_n \xi_n\|} \right). \quad (2.26)$$

By assumption, $\|\bar{\alpha}_n \xi_n\| \leq \beta^{-1}\|\alpha_n \xi_n\| \rightarrow 0$. Since ∇f is continuous, it follows from Cauchy–Schwarz, (2.19), and (2.24) that

$$\lim_{k \rightarrow \infty} \frac{|\nabla f_n^\top (\hat{R}(\bar{\alpha}_n) - (x_n + \bar{\alpha}_n \xi_n))|}{g_n^- \|\bar{\alpha}_n \xi_n\|} \leq \lim_{k \rightarrow \infty} \frac{\|\nabla f(x^*)\| C \|\bar{\alpha}_n \xi_n\|^2}{g_n^- \|\bar{\alpha}_n \xi_n\|} = 0.$$

Therefore, (2.26) yields

$$\liminf_{n \rightarrow \infty} \frac{\|\bar{\alpha}_n \xi_n\|}{g_n^-} \geq M^{-2} L^{-1} (1 - c)\omega > 0,$$

in contradiction to (2.24). \square

Property (2.19) in Theorem 2.10 holds, for instance, if \mathcal{M} is locally a smooth submanifold in a neighborhood of the limit x^* of (x_n) , and R is a smooth retraction in that neighborhood (one can locally bound the second derivatives of $\xi_x \mapsto R(x, \xi_x)$). On the other hand, the condition $\alpha_n \xi_n \rightarrow 0$ is reasonable, but cannot be removed in general.³ When R is concretely specified, the situation can change. Considering the relevant example where a metric projection is used as retraction, it turns out that $x_{n+1} - x_n \rightarrow 0$ in combination with $x_n \rightarrow x^*$ automatically implies $\alpha_n \xi_n \rightarrow 0$ in the smooth case. This leads to the following powerful corollary of Theorem 2.10.

COROLLARY 2.11. *Let f be real-analytic and bounded below. Assume the metric projection (2.14) (the choice of norm does not matter here) is used as retraction in Algorithm 1. Further assume a cluster point x^* of the sequence (x_n) produced by Algorithm 1 exists, satisfies the Lojasiewicz gradient inequality **(L)**, and possesses an open neighborhood $\mathcal{O} \subseteq \mathbb{R}^N$ such that $\mathcal{M} \cap \mathcal{O}$ is a smooth embedded submanifold of \mathbb{R}^N . Then **(A1)**–**(A3)** hold. Consequently, $\lim_{n \rightarrow \infty} x_n = x^*$ with a rate of convergence as indicated in Theorem 2.3, and $\lim_{n \rightarrow \infty} g_n^- = g^-(x^*) = 0$.*

Remark. More concisely one may assume that $\mathcal{M} \cap \mathcal{O}$ is a real-analytic submanifold of \mathbb{R}^N [24, Definition 2.7.1]. Then the assumption on the validity of **(L)** is superfluous; cf. the remark following Proposition 2.2.

Proof. By Corollary 2.9, $x_n \rightarrow x^*$ and $\lim_{n \rightarrow \infty} g_n^- = g^-(x^*)$ (since on a smooth manifold g^- is a continuous function). For completeness, we now sketch the more or less elementary arguments that $\alpha_n \xi_n \rightarrow 0$ and (2.19) hold. Then Theorem 2.10 applies (the local Lipschitz condition for the gradient follows from the analyticity assumption).

There exists a local diffeomorphism ϕ from a neighborhood of $0 \in T_{x^*} \mathcal{M}$ (which is a linear space now) to \mathcal{M} such that for large enough n we can write $x^* = \phi(0)$, $x_n = \phi(y_n)$, and $T_n \mathcal{M} = \text{ran}(\phi'(y_n))$. The optimality condition for $x_{n+1} = R(x_n + \alpha_n \xi_n)$ when it is the orthogonal projection of $x_n + \alpha_n \xi_n$ is that the error is orthogonal on the tangent space at x_{n+1} , i.e.,

$$0 = \eta^\top \phi'(y_{n+1})^\top (x_{n+1} - (x_n + \alpha_n \xi_n)) \quad \text{for all } \eta \in T_{x^*} \mathcal{M}.$$

As $x_{n+1} - x_n \rightarrow 0$, this implies

$$0 = \lim_{n \rightarrow \infty} \alpha_n \phi'(y_{n+1})^\top \xi_n.$$

Since the smallest singular value of $\phi'(y_{n+1})^\top$ can be uniformly bounded below for n large enough (the limit $\phi'(0)^\top$ has full rank), it follows that $\alpha_n \xi_n \rightarrow 0$. Further, for any $\hat{\xi}_n = \phi'(y_n) \hat{\eta}_n$ we have by the best approximation property of R and Taylor's theorem that

$$\|R(x_n + \hat{\xi}_n) - (x_n + \hat{\xi}_n)\| \leq \|\phi(y_n + \hat{\eta}_n) - (\phi(y_n) + \phi'(y_n) \hat{\eta}_n)\| \leq \|\phi''(x_n + \vartheta_n \hat{\xi}_n)\| \|\hat{\eta}_n\|^2$$

for some $\vartheta_n \in (0, 1)$. If $\hat{\xi}_n \rightarrow 0$ for $n \rightarrow \infty$, then it follows that

$$\limsup_{n \rightarrow \infty} \frac{\|R(x_n + \hat{\xi}_n) - (x_n + \hat{\xi}_n)\|}{\|\hat{\xi}_n\|^2} \leq \|\phi''(0)\| \|(\phi'(0))^{-1}\|,$$

since ϕ'' is continuous in zero. \square

³The reason is that our requirements on R are only of a local kind. Consider, as a counterexample, the exponential retraction on a unit sphere; i.e., $R(x, \xi)$ is the endpoint of the arc of length $\|\xi_x\|$ in the great circle from x in direction ξ . Then for $\|\xi_x\| \rightarrow 2\pi$ it holds that $x - R(x, \xi_x) \rightarrow 0$.

3. Results for matrix varieties of bounded rank. The space $\mathbb{R}^m \otimes \mathbb{R}^n \cong \mathbb{R}^{m \times n} \cong \mathbb{R}^{mn}$ becomes a Euclidean space when equipped with the Frobenius inner product $\langle X, Y \rangle_{\mathbb{F}} = \text{trace}(X^{\top}Y)$. The corresponding norm and distance function are denoted by $\|\cdot\|_{\mathbb{F}}$ and $\text{dist}_{\mathbb{F}}$, respectively. Points in this space will now be denoted by X instead of x , tangent vectors by Ξ instead of ξ . Mainly to save space, we prefer in this paper the subspace and tensor product notation over explicit matrix representations. However, if we use the latter (as in the definition of the inner product), then it is with respect to some fixed orthonormal bases in \mathbb{R}^m and \mathbb{R}^n . For example, writing $X \in \mathcal{U} \otimes \mathcal{V}$ in $\mathbb{R}^m \otimes \mathbb{R}^n$ would mean in $\mathbb{R}^{m \times n}$ that $X = USV^{\top}$ for some matrices U, S, V with $\text{ran}(U) = \mathcal{U}$ and $\text{ran}(V) = \mathcal{V}$. By $\Pi_{\mathcal{S}}$ we denote the orthogonal projection onto a subspace \mathcal{S} . Then $(\Pi_{\mathcal{U}} \otimes \Pi_{\mathcal{V}})X$ corresponds to $UU^{\top}XVV^{\top}$, where U and V are orthonormal basis representations for \mathcal{U} and \mathcal{V} , respectively.

In this section we apply the above convergence theory for line-search methods to the real-algebraic variety $\mathcal{M}_{\leq k}$ of matrices with rank at most k (see (1.3)). We consider the problem

$$\min_{X \in \mathcal{M}_{\leq k}} f(X), \quad (3.1)$$

where, as before, $f : \mathbb{R}^{m \times n} \supseteq \mathcal{D} \rightarrow \mathbb{R}$ is continuously differentiable and bounded below. In fact, in the end we will assume that f is real-analytic to ensure the Łojasiewicz gradient inequality.

3.1. Tangent cone and optimality. Here and in the following, we suppose that

$$\text{rank}(X) = s \leq k, \quad \mathcal{U} = \text{ran}(X), \quad \mathcal{V} = \text{ran}(X^{\top}).$$

The following is well known; see, e.g., [18, 21, 50].

THEOREM 3.1. *The set \mathcal{M}_s of rank- s matrices is a smooth submanifold of dimension $(m+n-s)s$. It is dense and relatively open in $\mathcal{M}_{\leq s}$. The tangent space of \mathcal{M}_s at X is*

$$T_X \mathcal{M}_s = (\mathcal{U} \otimes \mathcal{V}) \oplus (\mathcal{U}^{\perp} \otimes \mathcal{V}) \oplus (\mathcal{U} \otimes \mathcal{V}^{\perp}). \quad (3.2)$$

The orthogonal projector on $T_X \mathcal{M}_s$ is hence given by

$$\Pi_{T_X \mathcal{M}_s} = \Pi_{\mathcal{U}} \otimes \Pi_{\mathcal{V}} + \Pi_{\mathcal{U}^{\perp}} \otimes \Pi_{\mathcal{V}} + \Pi_{\mathcal{U}} \otimes \Pi_{\mathcal{V}^{\perp}} = \Pi_{\mathcal{U}} \otimes I + I \otimes \Pi_{\mathcal{V}} - \Pi_{\mathcal{U}} \otimes \Pi_{\mathcal{V}}, \quad (3.3)$$

and it holds that

$$T_X \mathcal{M}_s \oplus (\mathcal{U}^{\perp} \otimes \mathcal{V}^{\perp}) = \mathbb{R}^m \otimes \mathbb{R}^n. \quad (3.4)$$

In fact, \mathcal{M}_s is even a real-analytic submanifold; see Lemma 3.7.

Our main task is to investigate the tangent cones of $\mathcal{M}_{\leq k}$ in points with $s < k$. The tangent cone $T_X \mathcal{M}_{\leq k}$ clearly contains $T_X \mathcal{M}_s$, but, in case $s < k$, also contains rays that arise when approaching X by a matrix of rank at most k but larger than s .

THEOREM 3.2. (see [13]) *Let $X \in \mathcal{M}_{\leq k}$, $\text{rank}(X) = s$. The tangent cone of $\mathcal{M}_{\leq k}$ at X is*

$$T_X \mathcal{M}_{\leq k} = T_X \mathcal{M}_s \oplus \{\Xi_{k-s} \in \mathcal{U}^{\perp} \otimes \mathcal{V}^{\perp} : \text{rank}(\Xi_{k-s}) \leq k-s\}.$$

Proof. To prove the “ \supseteq ” part, let Ξ be an element from the set on the right side of the equality. Then $\Xi = \Xi_s + \Xi_{k-s}$ with $\Xi_s \in T_X \mathcal{M}_s$, and $\text{rank}(\Xi_{k-s}) \leq k - s$. There exist a sequence $(Y_n) \subseteq \mathcal{M}_s$ and a sequence $(a_n) \subseteq \mathbb{R}^+$ such that $Y_n \rightarrow X$, and $a_n(Y_n - X) = \Xi_s$. One can assume $a_n \rightarrow \infty$. Then $X_n = Y_n + a_n^{-1} \Xi_{k-s}$ is a sequence in $\mathcal{M}_{\leq k}$ which converges to X , and $a_n(X_n - X)$ converges to Ξ , which proves $\Xi \in T_X \mathcal{M}_{\leq k}$.

To prove the reverse inclusion “ \subseteq ”, assume $\Xi = \lim_{n \rightarrow \infty} a_n(X - X_n)$, $X_n \rightarrow X$ in $\mathcal{M}_{\leq k}$, and $(a_n) \subseteq \mathbb{R}^+$. In the orthogonal decomposition

$$a_n(X_n - X) = \Pi_{T_X \mathcal{M}_s} a_n(X_n - X) + (\Pi_{\mathcal{U}^\perp} \otimes \Pi_{\mathcal{V}^\perp}) a_n X_n,$$

both terms have to converge separately. Denote their limits by Ξ_s and Ξ_{k-s} , respectively. Then obviously $\Xi = \Xi_s + \Xi_{k-s}$ with $\Xi_s \in T_X \mathcal{M}_s$ and $\Xi_{k-s} \in \mathcal{U}^\perp \otimes \mathcal{V}^\perp$. Since $(\Pi_{\mathcal{U}} \otimes \Pi_{\mathcal{V}}) X_n \rightarrow (\Pi_{\mathcal{U}} \otimes \Pi_{\mathcal{V}}) X = X$, and since the set of rank- s matrices is relatively open in $\mathcal{U} \otimes \mathcal{V}$, $\text{rank}((\Pi_{\mathcal{U}} \otimes \Pi_{\mathcal{V}}) X_n) = s$ for large enough n . Consequently, since $\text{rank}(X_n) \leq k$ for all n , it must hold that $\text{rank}((\Pi_{\mathcal{U}^\perp} \otimes \Pi_{\mathcal{V}^\perp}) a_n X_n) = \text{rank}((\Pi_{\mathcal{U}^\perp} \otimes \Pi_{\mathcal{V}^\perp}) X_n) \leq k - s$ for large enough n . It follows from the semicontinuity of matrix rank that $\text{rank}(\Xi_{k-s}) \leq k - s$. \square

Remark. In the recent paper [13] the tangent cones of $\mathcal{M}_{\leq k}$ have been previously derived, but in contrast to (2.2) are defined via analytic curves as

$$T_X \mathcal{M}_{\leq k} = \{\dot{\gamma}(0) : \gamma \text{ is an analytic curve with } \gamma(0) = X \text{ and } \gamma(t) \in \mathcal{M}_{\leq k} \text{ for } t \geq 0\}. \quad (3.5)$$

As shown in [43, Proposition 2], both definitions are equivalent. Up to an additional normalization constraint, the authors of [13] essentially prove Theorem 3.2 using definition (3.5), which together with our proof provides a direct verification that both definitions are equivalent. As mentioned in [13], when using definition (3.5), the “ \subseteq ” part in Theorem 3.2 follows from known results on the existence of analytic “singular value decomposition paths” [10]. We can easily modify our argument above to prove the “ \supseteq ” part for (3.5) by choosing an analytic curve γ_s in \mathcal{M}_s (possible by Lemma 3.7) such that $\Xi_s = \dot{\gamma}_s(0)$, and put $\gamma(t) = \gamma_s(t) + t \Xi_{k-s}$, which is an analytic curve in $\mathcal{M}_{\leq k}$ with $\dot{\gamma}(0) = \Xi = \Xi_s + \Xi_{k-s}$. The proof of “ \supseteq ” given in [13] seems more involved than is probably necessary, since the well-known structure of $T_X \mathcal{M}_s$ is not exploited.

Remark. In [35], formulas for normal cones of $\mathcal{M}_{\leq k}$ have been derived. They do not imply the formula for the tangent cone in singular points X with $s < k$. The reverse, however, is true. In view of (2.5), we can rephrase Corollary 3.4 below by stating that the regular normal cone at such X contains only zero. This then implies that the general normal cone [44, Definition 6.3] at X is the union of all limits of subspaces $(T_{X_n} \mathcal{M}_{\leq k})^\perp = \mathcal{U}_n^\perp \otimes \mathcal{V}_n^\perp$ with $X_n \rightarrow X$ and $\text{rank}(X_n) = k$. Consequently, the singular points of $\mathcal{M}_{\leq k}$ are also not regular in the sense of Clarke [44, Definition 6.4].

Now that we know the structure of the tangent cone in rank-deficient points, we can calculate the projection of the antigradient on it. This turns out to be easy. Moreover, the tangent cone in such points is so “large” that the projection on it carries over astonishingly much information. In fact, it generates all of $\mathbb{R}^{m \times n}$.

COROLLARY 3.3. *Let $X \in \mathcal{M}_{\leq k}$, $\text{rank}(X) = s$. Any $G \in T_X \mathcal{M}_{\leq k}$ satisfying $\| -\nabla f(X) - G \|_F = \text{dist}_F(-\nabla f(X), T_X \mathcal{M}_{\leq k})$ has the form*

$$G = \Pi_{T_X \mathcal{M}_s}(-\nabla f(X)) + \Xi_{k-s}, \quad (3.6)$$

where Ξ_{k-s} is a best rank- $(k - s)$ approximation of $(\Pi_{\mathcal{U}^\perp} \otimes \Pi_{\mathcal{V}^\perp})(-\nabla f(X)) = -\nabla f(X) - \Pi_{T_X \mathcal{M}_s}(-\nabla f(X))$ in the Frobenius norm. (Obviously, $\Xi_{k-s} \in \mathcal{U}^\perp \otimes \mathcal{V}^\perp$

then.) Moreover,

$$g^-(X) = \|G\|_{\mathbb{F}} \geq \sqrt{\frac{k-s}{\min(m-s, n-s)}} \|\nabla f(X)\|_{\mathbb{F}}. \quad (3.7)$$

Proof. The form of G is clear from Theorem 3.2 by orthogonality considerations. We prove the norm estimate. The square of the Frobenius norm of a matrix is the sum of its squared singular values. A best rank- $(k-s)$ approximation of a matrix in the Frobenius norm is obtained by truncating its singular value decomposition up to the largest $k-s$ singular values. As $\dim(\mathcal{U}^\perp) = m-s$ and $\dim(\mathcal{V}^\perp) = n-s$, the matrix $(\Pi_{\mathcal{U}^\perp} \otimes \Pi_{\mathcal{V}^\perp})(-\nabla f(X))$ has at most $\min(m-s, n-s)$ nonzero singular values. We conclude that

$$\|\Xi_{k-s}\|_{\mathbb{F}}^2 \geq \frac{k-s}{\min(m-s, n-s)} \|(\Pi_{\mathcal{U}^\perp} \otimes \Pi_{\mathcal{V}^\perp})\nabla f(X)\|_{\mathbb{F}}^2.$$

Since $\Xi_{k-s} \in \mathcal{U}^\perp \otimes \mathcal{V}^\perp$, (3.6) and (3.4) now show that

$$\begin{aligned} \|G\|_{\mathbb{F}}^2 &= \|\Pi_{T_X \mathcal{M}_s}(-\nabla f(X))\|_{\mathbb{F}}^2 + \|\Xi_{k-s}\|_{\mathbb{F}}^2 \\ &\geq \frac{k-s}{\min(m-s, n-s)} (\|\Pi_{T_X \mathcal{M}_s}(-\nabla f(X))\|_{\mathbb{F}}^2 + \|(\Pi_{\mathcal{U}^\perp} \otimes \Pi_{\mathcal{V}^\perp})(-\nabla f(X))\|_{\mathbb{F}}^2) \\ &= \frac{k-s}{\min(m-s, n-s)} \|\nabla f(X)\|_{\mathbb{F}}^2, \end{aligned}$$

as asserted. \square

The estimate (3.7) allows us to make a remarkable a priori statement about critical points of differentiable functions on $\mathcal{M}_{\leq k}$.

COROLLARY 3.4. *Let $k \leq \min(m, n)$, and let $X^* \in \mathcal{M}_{\leq k}$ be a critical point of (3.1) in the sense $g^-(X^*) = 0$. Then either $\text{rank}(X^*) = k$ or $\nabla f(X^*) = 0$.*

Conceptually similar statements have been made in [11, Proposition 4] and [20, Theorem 7] for optimization tasks on the set of positive semidefinite matrices. As an illustration consider the following.

COROLLARY 3.5. *Let $k \leq \min(m, n)$. Assume that $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is strictly convex and coercive and its unique minimizer on $\mathbb{R}^{m \times n}$ has rank larger than or equal to k . Then any relative local minimizer of f on $\mathcal{M}_{\leq k}$ has rank k .*

In light of these results, it is not surprising that we will have to make the assumption $\text{rank}(X^*) = k$ in our convergence results below in order to conclude $g^-(X^*) = 0$. It is not an artifact of the used techniques. Instead, Corollary 3.4 tells us that it will be normally impossible to find a rank-deficient critical point by a projected gradient method that most of the time moves on \mathcal{M}_k , since on \mathcal{M}_k the projection of the antigradient contains much less information.

We finish with a practical remark. When the matrices are large, one will only be able to work with sparse or low-rank representations of all involved quantities. In particular, $\nabla f(X)$ needs to allow for a sparse or a low-rank representation. If $s \ll \min(m, n)$, the calculation of $\Pi_{\mathcal{U}^\perp} \otimes \Pi_{\mathcal{V}^\perp}(-\nabla f(X)) = -\nabla f(X) - \Pi_{T_X \mathcal{M}_s}(-\nabla f(X))$ is then feasible using the second representation of $\Pi_{T_X \mathcal{M}_s}$ in (3.3). With some effort one can even exploit the low-rank structure of $\Pi_{T_X \mathcal{M}_s}(-\nabla f(X))$ to calculate an approximate singular value decomposition of the difference without explicitly assembling it. The huge projector $\Pi_{\mathcal{U}^\perp} \otimes \Pi_{\mathcal{V}^\perp}$ should never be formed. The final rank of tangent vectors itself is not larger than $2s + (k-s) = k+s$, which can be seen from the decomposition. We summarize the procedure as Algorithm 2.

Algorithm 2: Calculate the projection of $-\nabla f(X)$ on $T_X \mathcal{M}_{\leq k}$

Input: Antigradient $F = -\nabla f(X)$ at $X \in \mathcal{M}_{\leq k}$.

Output: Projection $G \in T_X \mathcal{M}_{\leq k}$ with $\|F - G\|_F = \text{dist}_F(F, T_X \mathcal{M}_{\leq k})$.

- 1 Find orthonormal bases U and V for $\text{ran}(X)$ and $\text{ran}(X^\top)$, respectively;
- 2 Calculate the projection on $T_X \mathcal{M}_s$:

$$\Xi_s = UU^\top F + FVV^\top - UU^\top FVV^\top;$$

- 3 Calculate best rank- $(k - s)$ approximation of the difference:

$$\Xi_{k-s} \in \underset{\text{rank}(Y) \leq k-s}{\text{argmin}} \|F - \Xi_s - Y\|_F;$$

- 4 Output:

$$G = \Xi_s + \Xi_{k-s}, \quad g^-(X) = \|G\|_F = \sqrt{\|\Xi_s\|_F^2 + \|\Xi_{k-s}\|_F^2}.$$

3.2. Retraction by best low-rank approximation. As a retraction we choose the best approximation by a matrix of rank at most k in the Frobenius norm, i.e.,

$$R(X, \Xi_X) \in \underset{Y \in \mathcal{M}_{\leq k}}{\text{argmin}} \|Y - (X + \Xi_X)\|_F. \quad (3.8)$$

It can be explicitly calculated using singular value decomposition. In unlikely events, (3.8) is set-valued, but we can assume that a specific choice is made by fixing deterministic singular value decomposition and truncation algorithms. The particular choice does not matter. We emphasize once more that Definition 2.4 is indeed fulfilled: let $\Xi \in T_X \mathcal{M}_{\leq k}$; then by [43, Proposition 2] there exists an analytic arc $\gamma: [0, \epsilon) \rightarrow \mathcal{M}_{\leq k}$ such that $\dot{\gamma}(0) = \Xi$. Hence,

$$\lim_{\alpha \rightarrow 0^+} \frac{\|R(X, \alpha \Xi_X) - (X + \Xi_X)\|_F}{\alpha} \leq \lim_{\alpha \rightarrow 0^+} \frac{\|\gamma(\alpha) - (X + \dot{\gamma}(0))\|_F}{\alpha} = 0. \quad (3.9)$$

We have the following nice estimate, which provides $M = 1 + 2^{-1/2}$ in (2.13).

PROPOSITION 3.6. *The above retraction satisfies*

$$\|R(X, \Xi_X) - (X + \Xi_X)\|_F \leq \frac{1}{\sqrt{2}} \|\Xi_X\|_F \quad \text{for all } X \in \mathcal{M}_{\leq k} \text{ and } \Xi_X \in T_X \mathcal{M}_{\leq k}.$$

Proof. The matrices $X + (\Pi_{\mathcal{U}} \otimes I)\Xi_X = (\Pi_{\mathcal{U}} \otimes I)(X + \Xi_X)$ and $X + (I \otimes \Pi_{\mathcal{V}})\Xi_X = (I \otimes \Pi_{\mathcal{V}})(X + \Xi_X)$ both have rank at most s . Thus, by Theorem 3.2,

$$X + (\Pi_{\mathcal{U}} \otimes I)\Xi_X + (\Pi_{\mathcal{U}^\perp} \otimes \Pi_{\mathcal{V}^\perp})\Xi_X$$

and

$$X + (I \otimes \Pi_{\mathcal{V}})\Xi_X + (\Pi_{\mathcal{U}^\perp} \otimes \Pi_{\mathcal{V}^\perp})\Xi_X$$

both have rank not larger than k . Considering them as possible candidates for a best approximation $R(X, \Xi_X)$ of $X + \Xi_X$ by a matrix of rank at most k , we obtain the

desired bound

$$\|R(X + \Xi_X) - (X + \Xi_X)\|_{\mathbb{F}}^2 \leq \min(\|(\Pi_{\mathcal{U}^\perp} \otimes \Pi_{\mathcal{V}})\Xi_X\|_{\mathbb{F}}^2, \|(\Pi_{\mathcal{U}} \otimes \Pi_{\mathcal{V}^\perp})\Xi_X\|_{\mathbb{F}}^2) \leq \frac{1}{2}\|\Xi_X\|_{\mathbb{F}}^2,$$

where we have made use of the orthogonal decompositions (3.2) and (3.4). \square

We conclude that $\|R(X, \Xi_X) - X\|_{\mathbb{F}}^2 \rightarrow 0$ automatically implies $\Xi_X \rightarrow 0$.

3.3. Lojasiewicz inequality and convergence result. To apply the convergence results of section 2.4, we will show that the Lojasiewicz gradient inequality **(L)** holds for real-analytic functions in every point of $\mathcal{M}_{\leq k}$. The aim is to apply Proposition 2.2. As a first step, the following lemma implies that the smooth submanifolds \mathcal{M}_s are indeed real-analytic submanifolds in the sense of [24, Definition 2.7.1].

LEMMA 3.7. *Let $0 < s \leq \min(m, n)$. The (componentwise) real-analytic map*

$$(U, V) \mapsto UV^T$$

is a submersion from the open subset $\{(U, V) \in \mathbb{R}^{m \times s} \times \mathbb{R}^{n \times s} : \text{rank}(U) = \text{rank}(V) = s\}$ of $\mathbb{R}^{m \times s} \times \mathbb{R}^{n \times s}$ onto \mathcal{M}_s .

Proof. The openness of the domain of definition and the surjectivity are clear. The derivative at (U, V) is the linear map $(\delta U, \delta V) \mapsto \delta UV^T + U\delta V^T$. As both U and V have rank s , one verifies that it has no nontrivial kernel in the $(m+n-s)s$ -dimensional subspace of all $(\delta U, \delta V)$ satisfying $U^T\delta U = 0$. Hence the derivative has rank at least $(m+n-s)s = \dim(\mathcal{M}_s)$ (see Theorem 3.1), which already proves the claim. \square

THEOREM 3.8. *Let $\mathcal{D} \subseteq \mathbb{R}^{m \times n}$ be open, $\mathcal{M}_{\leq k} \subset \mathcal{D}$, and $f: \mathcal{D} \rightarrow \mathbb{R}$ be real-analytic. Then the Lojasiewicz gradient inequality **(L)** holds at any point $X \in \mathcal{M}_{\leq k}$.*

Proof. Let $X \in \mathcal{M}_{\leq k}$, $\text{rank}(X) = s$. We assume $s > 0$; otherwise the proof requires some obvious notational modifications. There exist matrices $U_0 \in \mathbb{R}^{m \times s}$ and $V_0 \in \mathbb{R}^{n \times s}$, both of rank s , such that $X = U_0V_0^T$. We consider the map

$$\begin{aligned} \tau: \mathcal{N} = \mathbb{R}^{m \times s} \times \mathbb{R}^{n \times s} \times \mathbb{R}^{m \times (k-s)} \times \mathbb{R}^{n \times (k-s)} &\rightarrow \mathbb{R}^{m \times n}, \\ (U, V, U_{k-s}, V_{k-s}) &\mapsto UV^T + U_{k-s}V_{k-s}^T. \end{aligned}$$

This map is obviously real-analytic, its image is $\mathcal{M}_{\leq k}$, and $\tau(U_0, V_0, 0, 0) = X$. We need to prove property (ii) in Proposition 2.2. Let $\tilde{\mathcal{N}} \subseteq \mathcal{N}$ be an open neighborhood of $(U_0, V_0, 0, 0)$. We may assume that this neighborhood is so small such that for all $(U, V, U_{k-s}, V_{k-s}) \in \tilde{\mathcal{N}}$ it holds that $\text{rank}(U) = \text{rank}(V) = s$. Lemma 3.7 then implies that the map $(U, V, U_{k-s}, V_{k-s}) \mapsto UV^T$ is a submersion from $\tilde{\mathcal{N}}$ on the smooth manifold \mathcal{M}_s , and as such an open map [14, §(16.7.5)]. Consequently, for small enough $\delta > 0$ we can claim that $\tau(\tilde{\mathcal{N}})$ contains *all* matrices of the form $Y_s + Y_{k-s}$ with $\text{rank}(Y_s) = s$, $\text{rank}(Y_{k-s}) \leq k-s$, and $\|Y_s - X\|_{\mathbb{F}} < 2\delta$. By semicontinuity of rank, we can choose δ so small that $\text{rank}(Y) \geq s$ for all $Y \in B_\delta(X)$, where $B_\delta(X)$ denotes the open ball in $\mathbb{R}^{m \times n}$ of radius δ (in Frobenius norm) with center X . Let Y_s denote the best rank- s approximation (in Frobenius norm) of $Y \in B_\delta(X) \cap \mathcal{M}_{\leq k}$. As Y_s is obtained by truncating a singular value decomposition of Y , we have $Y = Y_s + Y_{k-s}$ with $\text{rank}(Y_{k-s}) \leq k-s$, and

$$\|Y_s - X\|_{\mathbb{F}} \leq \|Y_s - Y\|_{\mathbb{F}} + \|Y - X\|_{\mathbb{F}} \leq \|X - Y\|_{\mathbb{F}} + \|Y - X\|_{\mathbb{F}} < 2\delta,$$

the second inequality holding since $\text{rank}(X) = s$. By the previous considerations, this implies $Y \in \tau(\tilde{\mathcal{N}})$. We thus have shown that $\tau(\tilde{\mathcal{N}})$ contains the relatively open set $B_\delta(X) \cap \mathcal{M}_{\leq k}$. \square

We now have collected all requirements to apply Theorem 2.10 or Corollary 2.11. For concreteness, we consider a particular algorithm where the search direction equals the projected antigradient and the retraction is obtained by best rank- k approximation. It is denoted as Algorithm 3.

Algorithm 3: Projected steepest descent with line-search on $\mathcal{M}_{\leq k}$

Input: Starting guess $X_0 \in \mathcal{M}_{\leq k}$, $\beta, c \in (0, 1)$.

```

1 for  $n=0, 1, 2, \dots$  do
2   Calculate a projection  $G_n$  of  $-\nabla f(X_n)$  on  $T_{X_n}\mathcal{M}_{\leq k}$  using Algorithm 2;
3   Choose  $\tilde{\beta}_n \geq 1$ , and find Armijo point  $\alpha_n$  for  $X_n, G_n, \tilde{\beta}_n, \beta, c$ ;
4   Set  $X_{n+1}$  to be a best approximation (with respect to Frobenius norm) of
    $X_n + \alpha_n G_n$  of rank at most  $k$ .
5 end

```

THEOREM 3.9. *Let f be real-analytic and bounded below. If the sequence (X_n) generated by Algorithm 3 possesses a cluster point X^* , then it is its limit. If further $\text{rank}(X^*) = k$, then $g^-(X^*) = 0$, and the convergence rate estimates of Theorem 2.3 apply.*

Proof. The convergence of the sequence follows from Theorem 3.8, Proposition 2.6, and Corollary 2.9. Due to Theorem 3.1, the rest is an instance of Corollary 2.11. \square

3.4. A method without retraction. It is possible to have a gradient-related search direction Ξ_n such that $X_n + \alpha\Xi_n \in \mathcal{M}_{\leq k}$ for all α . The idea is the same as in the proof of Proposition 3.6. By (3.6) and (3.2), a projection G of $-\nabla f_n$ consists of four, mutually orthogonal parts:

$$G_n = (\Pi_{\mathcal{U}} \otimes \Pi_{\mathcal{V}})(-\nabla f_n) + (\Pi_{\mathcal{U}^\perp} \otimes \Pi_{\mathcal{V}})(-\nabla f_n) + (\Pi_{\mathcal{U}} \otimes \Pi_{\mathcal{V}^\perp})(-\nabla f_n) + \Xi_{k-s,n},$$

with $\text{rank}(\Xi_{k-s,n}) \leq k - s$. Consider the two possible partial projections

$$G_n^{(1)} = (\Pi_{\mathcal{U}} \otimes \Pi_{\mathcal{V}})(-\nabla f_n) + (\Pi_{\mathcal{U}} \otimes \Pi_{\mathcal{V}^\perp})(-\nabla f_n) + \Xi_{k-s,n} = (\Pi_{\mathcal{U}} \otimes I)(-\nabla f_n) + \Xi_{k-s,n} \quad (3.10)$$

and

$$G_n^{(2)} = (\Pi_{\mathcal{U}} \otimes \Pi_{\mathcal{V}})(-\nabla f_n) + (\Pi_{\mathcal{U}^\perp} \otimes \Pi_{\mathcal{V}})(-\nabla f_n) + \Xi_{k-s,n} = (I \otimes \Pi_{\mathcal{V}})(-\nabla f_n) + \Xi_{k-s,n}. \quad (3.11)$$

Both are elements of the tangent cone at X_n and satisfy $\text{rank}(X_n + \alpha G_n^{(i)}) \leq k$ for all α , $i = 1, 2$. Assume that $\|G_n^{(1)}\|_{\mathbb{F}} \geq \|G_n^{(2)}\|_{\mathbb{F}}$. Then, by orthogonality arguments, $\|G_n^{(1)}\|_{\mathbb{F}}^2 \geq \frac{1}{2}\|G_n\|_{\mathbb{F}}^2$, and

$$\langle \nabla f_n, G_n^{(1)} \rangle_{\mathbb{F}} = \|G_n^{(1)}\|_{\mathbb{F}}^2 \geq \frac{1}{\sqrt{2}}\|G_n\|_{\mathbb{F}}\|G_n^{(1)}\|_{\mathbb{F}} = \frac{1}{\sqrt{2}}g_n^-\|G_n^{(1)}\|_{\mathbb{F}}.$$

Thus $G_n^{(1)}$ satisfies the ω -angle condition with $\omega = \frac{1}{\sqrt{2}}$. If $\|G_n^{(1)}\|_{\mathbb{F}} \leq \|G_n^{(2)}\|_{\mathbb{F}}$, then $G_n^{(2)}$ satisfies this angle condition.

This leads us to Algorithm 4, which contains no retraction steps. Still, it shares the nice abstract convergence features with the projected steepest descent, even with a slightly extended statement in singular points (convergence rate).

THEOREM 3.10. *Let f be real-analytic and bounded below. If the sequence (X_n) generated by Algorithm 4 possesses a cluster point X^* , then it is its limit, and the*

Algorithm 4: Descent method on $\mathcal{M}_{\leq k}$ without retraction

Input: Starting guess $X_0 \in \mathcal{M}_{\leq k}$, $\beta, c \in (0, 1)$.

- 1 **for** $n=0, 1, 2, \dots$ **do**
- 2 **if** $\|(\Pi_{\mathcal{U}} \otimes I)(-\nabla f(X_n))\|_{\mathbb{F}} \geq \|(I \otimes \Pi_{\mathcal{V}})(-\nabla f(X_n))\|_{\mathbb{F}}$ **then**
- 3 | Use $\Xi_n = G_n^{(1)}$ from (3.10);
- 4 **else**
- 5 | Use $\Xi_n = G_n^{(2)}$ from (3.11);
- 6 **end**
- 7 Choose $\bar{\beta}_n \geq \sqrt{2}$, and find Armijo point α_n for $X_n, \Xi_n, \bar{\beta}_n, \beta, c$;
- 8 Form the next iterate

$$X_{n+1} = X_n + \alpha_n \Xi_n.$$

- 9 **end**

convergence rate estimates of Theorem 2.3 apply. If further $\text{rank}(X^*) = k$, then $g^-(X^*) = 0$.

Proof. Since $\text{rank}(X_n + \alpha_n \Xi_n) \leq k$, we can formally write $X_{n+1} = R(X_n, \alpha_n \Xi_n)$ in the algorithm in order to get into the abstract framework (here R is again retraction by best low-rank approximation). Then the mere convergence of the sequence follows again from Theorem 3.8 and Corollary 2.9. The feature is now that (2.19) is trivially satisfied since R acts as identity; therefore the validity of convergence rate estimates follows from Theorem 2.10 even if the limit point is singular (the Lipschitz condition (2.20) follows from analyticity). We also have $g_n^-(X_n) \rightarrow 0$ from which we can conclude $g^-(X^*) = 0$ if g^- is continuous in X^* . But this is the case if $X^* \in \mathcal{M}_k$. \square

Since it does not leave the feasible set, Algorithm 4 is very elegant and saves some cost in every step of the backtracking to find the Armijo point. In applications, however, the retraction from rank (at most) $2k$ to rank k , as required in Algorithm 3, is typically much less expensive than, for instance, a function value evaluation or the projection of the gradient. We hence expect that the saved retractions will seldom compensate for the less gradient-related search directions.

We checked this with a toy example of matrix completion in a setup similar to [50], using *straightforward, comparably nonoptimized* MATLAB R2012b implementations of both algorithms (choosing $\beta = \frac{1}{2}$ and $c = 10^{-4}$) on a Linux workstation with six 3.2 GHz CPU cores and 6 GB of memory. The problem that was solved is

$$\min_{X \in \mathcal{M}_{\leq k}} \frac{1}{2} \|P_{\Omega}(A - X)\|_{\mathbb{F}}^2, \quad (3.12)$$

where P_{Ω} is the projector on a subset Ω of indices. The $n \times n$ matrix $A = UV^{\top}$ of rank r was generated by randomly generating the two $n \times r$ factor matrices U and V from a normal distribution. The size of Ω was chosen as $|\Omega| = \max(\text{OS} \cdot (2kn - k^2), n \log n)$, which corresponds to an oversampling rate of at least OS when assuming A to have rank k (cf. [50]), and Ω itself was drawn uniformly at random. As a starting guess we chose in all experiments a best rank- k approximation of the antigradient $-\nabla f(0) = -P_{\Omega}(A)$. In both Algorithms 3 and 4, this choice of starting guess is formally equivalent to starting with zero and performing an exact line-search in the very first step.

In the first test the rank of A was indeed set to be $r = k$, so that the global solution of (3.12) lies on the smooth part \mathcal{M}_k of $\mathcal{M}_{\leq k}$. For $n = 2000$, $k = 20$, and




FIG. 3.1. Application of Algorithms 3 and 4 to (3.12) with $A \in \mathbb{R}^{2000 \times 2000}$, $\text{rank}(A) = k$, for $k = 20$ (94.03% missing entries), and $k = 80$ (76.48% missing entries). Solid lines: relative errors (3.13) (full index set); dashed lines: relative errors (3.14) (sample index set).

OS = 3 (94.03% missing entries), the relative errors,

$$\frac{\|A - X_n\|_F}{\|A\|_F}, \quad (3.13)$$

as well as the relative errors on the visible index set,

$$\frac{\|P_\Omega(A - X_n)\|_F}{\|P_\Omega A\|_F} = \frac{\sqrt{2f(X_n)}}{\|P_\Omega A\|_F}, \quad (3.14)$$

are plotted in Figure 3.1. As one can see, Algorithm 4 is inferior to Algorithm 3 with respect to both number of iterations and computation time (the latter is plotted just to give an impression). One might think that the relative performance of Algorithm 4 improves for larger k . The plots for $k = 80$ do not support this hope (in this case only 76.48% entries are missing, which perhaps explains the faster error decay). On the other hand, we did not observe that the superiority of Algorithm 3 would become considerably more pronounced for larger matrices; plots looked very similar.

We repeated the same experiments with matrices A having rank $r = k/2$. Interestingly, Algorithm 4 now performed better than Algorithm 3, but both methods were unable to find a good approximation of the rank-deficient global solution A ; see Figure 3.2. In fact, we practically never encountered iterates whose k th singular value was less than 10^{-4} . This confirms our expectation that our approach via line-search methods on $\mathcal{M}_{\leq k}$ does not contribute to the problem of rank estimation. Yet it served as an elegant theoretical vehicle to prove the convergence of Riemannian line-search methods on \mathcal{M}_k in “almost every instance.” As indicated in the introduction, a possible synthesis are rank-increasing strategies which subsequently optimize on varieties $\mathcal{M}_{\leq s}$ for a growing sequence of s [37, 47, 49].

Of course, Algorithms 3 and 4 served here only as examples and are naturally inferior to more sophisticated line-search methods, such as the nonlinear CG methods used in [50], which use gradient information from previous iterates.




FIG. 3.2. Application of Algorithms 3 and 4 to (3.12) with $A \in \mathbb{R}^{2000 \times 2000}$, $\text{rank}(A) = k/2$, for $k = 20$ (94.03% missing entries), and $k = 80$ (76.48% missing entries). Solid lines: relative errors (3.13) (full index set); dashed lines: relative errors (3.14) (sample index set).

4. Conclusion. We extended available results on convergence of descent iterations on manifolds via Lojasiewicz gradient inequality to gradient-related line-search methods on the real-algebraic variety $\mathcal{M}_{\leq k}$ of real $m \times n$ matrices of rank at most k , by explicitly taking the tangent cones at singular points into consideration. This made it possible to overcome some theoretical difficulties arising from the nonclosedness and unbounded curvature that one faces in the convergence analysis of Riemannian optimization methods on the smooth manifold \mathcal{M}_k of rank- k matrices. So far, the results are applicable for real-analytic cost functions.

There is growing interest in treating low-rank tensor problems by Riemannian optimization, e.g., tensor completion [25] or dynamical tensor approximation [22, 34, 48]. It would be important and interesting to extend the results to tensor varieties of bounded subspace ranks, e.g., bounded Tucker ranks, hierarchical Tucker ranks, or tensor train ranks [23, 16, 42]. As these varieties take the form of intersections of low-rank matrix varieties [48], the results in this paper can likely be generalized in this direction.

Acknowledgments. We thank Pierre-Antoine Absil who suggested an improvement of Corollary 3.3, and Bart Vandereycken for useful hints to the literature.

Appendix. Proof of Theorem 2.3.

We can assume that $g_n^- > 0$ for all n since otherwise the sequence becomes stationary and there is nothing to prove. There will also be no loss of generality to assume that (A1) and (A2) hold for all n and that $f(x^*) = 0$. Then $0 \leq f(x^*) \leq f_n$ for all n , and the Lojasiewicz gradient inequality at x^* reads as

$$f(x)^{1-\theta} \leq \Lambda g^-(x) \tag{A.1}$$

whenever $\|x - x^*\| < \delta = \delta(x^*)$. Let $\epsilon \in (0, \delta]$, and assume $\|x_n - x^*\| < \delta$. Then,

by (A.1) and **(A1)**,

$$\|x_n - x_{n+1}\| \leq \frac{\Lambda}{\sigma} f_n^{\theta-1} (f_n - f_{n+1}).$$

Using the fact that for $\varphi \in [f_{n+1}, f_n]$ there holds $f_n^{\theta-1} \leq \varphi^{\theta-1} \leq f_{n+1}^{\theta-1}$, we can estimate

$$f_n^{\theta-1} (f_n - f_{n+1}) \leq \int_{f_{n+1}}^{f_n} \varphi^{\theta-1} d\varphi = \frac{1}{\theta} (f_n^\theta - f_{n+1}^\theta)$$

and thus obtain

$$\|x_n - x_{n+1}\| \leq \frac{\Lambda}{\sigma\theta} (f_n^\theta - f_{n+1}^\theta).$$

More generally, let $\|x_k - x^*\| < \epsilon \leq \delta$ for $n \leq k < m$; we get by this argument that

$$\|x_m - x_n\| \leq \sum_{k=n}^m \|x_{k+1} - x_k\| \leq \sum_{k=n}^m \frac{\Lambda}{\sigma\theta} (f_k^\theta - f_{k+1}^\theta) = \frac{\Lambda}{\sigma\theta} (f_n^\theta - f_m^\theta) \leq \frac{\Lambda}{\sigma\theta} f_n^\theta. \quad (\text{A.2})$$

Since x^* is an accumulation point, we can pick n so large that (recall that f is continuous and $f(x^*) = 0$)

$$\|x_n - x^*\| < \frac{\epsilon}{2} \quad \text{and} \quad \frac{\Lambda}{\sigma\theta} f_n^\theta < \frac{\epsilon}{2}.$$

Then (A.2) inductively implies $\|x_m - x^*\| < \epsilon$ for all $m \geq n$. This proves that x^* is the limit point of the sequence, and, by **(A3)**, $g_n^- \rightarrow 0$.

To estimate the convergence rate, let $r_n = \sum_{k=n}^{\infty} \|x_{k+1} - x_k\|$. Then $\|x_n - x^*\| \leq r_n$, so it suffices to estimate the latter. By (A.2), (A.1), and **(A3)**, there exists $n_0 \geq 1$ such that for $n \geq n_0$ it holds that

$$r_n^{\frac{1-\theta}{\theta}} \leq \left(\frac{\Lambda}{\sigma\theta}\right)^{\frac{1-\theta}{\theta}} f_n^{1-\theta} \leq \left(\frac{\Lambda}{\sigma\theta}\right)^{\frac{1-\theta}{\theta}} \frac{\Lambda}{\kappa} \|x_{n+1} - x_n\| = \left(\frac{\Lambda}{\sigma\theta}\right)^{\frac{1-\theta}{\theta}} \frac{\Lambda}{\kappa} (r_n - r_{n+1}),$$

that is,

$$r_{n+1} \leq r_n - \nu r_n^{\frac{1-\theta}{\theta}} \quad (\text{A.3})$$

with $\nu = \left(\frac{\Lambda}{\sigma\theta}\right)^{\frac{\theta-1}{\theta}} \frac{\kappa}{\Lambda}$. Now, if $\theta = 1/2$, we get from (A.3) that $\nu \in (0, 1)$, and

$$r_n \leq r_{n_0} (1 - \nu)^{n-n_0} \left(e^{\ln(1-\nu)}\right)^n$$

for $n \geq n_0$. The case $0 < \theta < 1/2$ is more delicate. We follow Levitt [30]: put $p = \frac{\theta}{1-2\theta}$, $C \geq \max\left(\left(\frac{\nu}{p}\right)^{-p}, r_{n_0} n_0^{-p}\right)$, and $s_n = C n^{-p}$; then $s_{n_0} \geq r_{n_0}$, and

$$s_{n+1} = s_n (1 + n^{-1})^{-p} \geq s_n (1 - p n^{-1}) = s_n - \frac{p}{C^{1/p}} s_n^{\frac{p+1}{p}} \geq s_n - \nu s_n^{\frac{p+1}{p}} = s_n - \nu s_n^{\frac{1-\theta}{\theta}}$$

(the first inequality holding by convexity of x^{-p}). Using induction, it now follows from (A.3) that $r_n \leq s_n$ for all $n \geq n_0$, which finishes the proof. \square

REFERENCES

- [1] P.-A. ABSIL, R. MAHONY, AND B. ANDREWS, *Convergence of the iterates of descent methods for analytic cost functions*, SIAM J. Optim., 16 (2005), pp. 531–547.
- [2] P.-A. ABSIL, R. MAHONY, AND R. SEPULCHRE, *Optimization Algorithms on Matrix Manifolds*, Princeton University Press, Princeton, NJ, 2008.
- [3] P.-A. ABSIL AND J. MALICK, *Projection-like retractions on matrix manifolds*, SIAM J. Optim., 22 (2012), pp. 135–158.
- [4] H. ATTOUCH AND J. BOLTE, *On the convergence of the proximal algorithm for nonsmooth functions involving analytic features*, Math. Program., 116 (2009), pp. 5–16.
- [5] H. ATTOUCH, J. BOLTE, P. REDONT, AND A. SOUBEYRAN, *Proximal alternating minimization and projection methods for nonconvex problems: An approach based on the Kurdyka-Lojasiewicz inequality*, Math. Oper. Res., 35 (2010), pp. 438–457.
- [6] H. ATTOUCH, J. BOLTE, AND B.F. SVAITER, *Convergence of descent methods for semi-algebraic and tame problems: Proximal algorithms, forward-backward splitting, and regularized Gauss-Seidel methods*, Math. Program., 137 (2013), pp. 91–129.
- [7] J. BOLTE, A. DANILIDIS, AND A. LEWIS, *The Lojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems*, SIAM J. Optim., 17 (2007), pp. 1205–1223.
- [8] J. BOLTE, A. DANILIDIS, A. LEWIS, AND M. SHIOTA, *Clarke subgradients of stratifiable functions*, SIAM J. Optim., 18 (2007), pp. 556–572.
- [9] J. BOLTE, A. DANILIDIS, O. LEY, AND L. MAZET, *Characterizations of Lojasiewicz inequalities: Subgradient flows, talweg, convexity*, Trans. Amer. Math. Soc., 362 (2010), pp. 3319–3363.
- [10] A. BUNSE-GERSTNER, R. BYERS, V. MEHRMANN, AND N.K. NICHOLS, *Numerical computation of an analytic singular value decomposition of a matrix valued function*, Numer. Math., 60 (1991), pp. 1–39.
- [11] S. BURER AND R.D.C. MONTEIRO, *A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization*, Math. Program., 95 (2003), pp. 329–357.
- [12] E. CANCÈS, V. EHLACHER, AND T. LELIÈVRE, *Greedy algorithms for high-dimensional eigenvalue problems*, Constr. Approx., 40 (2014), pp. 387–423.
- [13] T.P. CASON, P.-A. ABSIL, AND P. VAN DOOREN, *Iterative methods for low rank approximation of graph similarity matrices*, Linear Algebra Appl., 438 (2013), pp. 1863–1882.
- [14] J. DIEUDONNÉ, *Treatise on Analysis. Vol. III*, Academic Press, New York, London, 1972.
- [15] M. GUIGNARD, *Generalized Kuhn–Tucker conditions for mathematical programming problems in a Banach space*, SIAM J. Control, 7 (1969), pp. 232–241.
- [16] W. HACKBUSCH AND S. KÜHN, *A new scheme for the tensor representation*, J. Fourier Anal. Appl., 15 (2009), pp. 706–722.
- [17] A. HARAUX AND M.A. JENDOUBI, *The Lojasiewicz gradient inequality in the infinite-dimensional Hilbert space framework*, J. Funct. Anal., 260 (2011), pp. 2826–2842.
- [18] U. HELMKE AND M.A. SHAYMAN, *Critical points of matrix least squares distance functions*, Linear Algebra Appl., 215 (1995), pp. 1–19.
- [19] S.-Z. HUANG, *Gradient Inequalities. With Applications to Asymptotic Behavior and Stability of Gradient-like Systems*, American Mathematical Society, Providence, RI, 2006.
- [20] M. JOURNÉE, F. BACH, P.-A. ABSIL, AND R. SEPULCHRE, *Low-rank optimization on the cone of positive semidefinite matrices*, SIAM J. Optim., 20 (2010), pp. 2327–2351.
- [21] O. KOCH AND C. LUBICH, *Dynamical low-rank approximation*, SIAM J. Matrix Anal. Appl., 29 (2007), pp. 434–454.
- [22] ———, *Dynamical tensor approximation*, SIAM J. Matrix Anal. Appl., 31 (2010), pp. 2360–2375.
- [23] T.G. KOLDA AND B.W. BADER, *Tensor decompositions and applications*, SIAM Rev., 51 (2009), pp. 455–500.
- [24] S.G. KRANTZ AND H.R. PARKS, *A Primer of Real Analytic Functions*, Birkhäuser Boston Boston, 2nd ed., 2002.
- [25] D. KRESSNER, M. STEINLECHNER, AND B. VANDEREYCKEN, *Low-rank tensor completion by Riemannian optimization*, BIT, 54 (2014), pp. 447–468.
- [26] K. KURDYKA, *On gradients of functions definable in o-minimal structures*, Ann. Inst. Fourier (Grenoble), 48 (1998), pp. 769–783.
- [27] C. LAGEMAN, *Konvergenz reell-analytischer gradientenähnlicher Systeme*, diploma thesis, Universität Würzburg, Würzburg, Germany, 2002. In German.
- [28] ———, *Convergence of Gradient-like Dynamical Systems and Optimization Algorithms*, PhD thesis, Universität Würzburg, Würzburg, Germany, 2007.
- [29] ———, *Pointwise convergence of gradient-like systems*, Math. Nachr., 280 (2007), pp. 1543–1558.
- [30] A. LEVITT, *Convergence of gradient-based algorithms for the Hartree-Fock equations*, ESAIM

- Math. Model. Numer. Anal., 46 (2012), pp. 1321–1336.
- [31] Z. LI, A. USCHMAJEV, AND S. ZHANG, *On convergence of the maximum block improvement method*, SIAM J. Optim., 25 (2015), pp. 210–233.
 - [32] S. LOJASIEWICZ, *Ensemble semi-analytique*. Note des cours, Institut des Hautes Etudes Scientifique, 1965.
 - [33] C. LUBICH AND I.V. OSELEDETS, *A projector-splitting integrator for dynamical low-rank approximation*, BIT, 54 (2014), pp. 171–188.
 - [34] C. LUBICH, T. ROHWEDDER, R. SCHNEIDER, AND B. VANDEREYCKEN, *Dynamical approximation by hierarchical Tucker and tensor-train tensors*, SIAM J. Matrix Anal. Appl., 34 (2013), pp. 470–494.
 - [35] D.R. LUKE, *Prox-regularity of rank constraints sets and implications for algorithms*, J. Math. Imaging Vision, 47 (2013), pp. 231–238.
 - [36] B. MERLET AND T.N. NGUYEN, *Convergence to equilibrium for discretizations of gradient-like flows on Riemannian manifolds*, Differential Integral Equations, 26 (2013), pp. 571–602.
 - [37] B. MISHRA, G. MEYER, F. BACH, AND R. SEPULCHRE, *Low-rank optimization with trace norm penalty*, SIAM J. Optim., 23 (2013), pp. 2124–2149.
 - [38] B. MISHRA, G. MEYER, S. BONNABEL, AND R. SEPULCHRE, *Fixed-rank matrix factorizations and Riemannian low-rank optimization*, Comput. Statist., 29 (2014), pp. 591–621.
 - [39] J. NOCEDAL AND S.J. WRIGHT, *Numerical Optimization*, Springer, New York, 2006.
 - [40] D. NOLL, *Convergence of non-smooth descent methods using the Kurdyka-Lojasiewicz inequality*, J. Optim. Theory Appl., 160 (2014), pp. 553–572.
 - [41] A. NONNENMACHER AND C. LUBICH, *Dynamical low-rank approximation: Applications and numerical experiments*, Math. Comput. Simulation, 79 (2008), pp. 1346–1357.
 - [42] I.V. OSELEDETS, *Tensor-train decomposition*, SIAM J. Sci. Comput., 33 (2011), pp. 2295–2317.
 - [43] D.B. O’SHEA AND L.C. WILSON, *Limits of tangent spaces to real surfaces*, Amer. J. Math., 126 (2004), pp. 951–980.
 - [44] R.T. ROCKAFELLAR AND R.J.-B. WETS, *Variational Analysis*, Springer-Verlag, Berlin, 1998.
 - [45] U. SHALIT, D. WEINSHALL, AND G. CHECHIK, *Online learning in the embedded manifold of low-rank matrices*, J. Mach. Learn. Res., 13 (2012), pp. 429–458.
 - [46] M. SHUB, *Some remarks on dynamical systems and numerical analysis*, in Dynamical Systems and Partial Differential Equations (Caracas, 1984), Univ. Simon Bolivar, Caracas, 1986, pp. 69–91.
 - [47] M. TAN, I. W. TSANG, L. WANG, B. VANDEREYCKEN, AND S. J. PAN, *Riemannian pursuit for big matrix recovery*, in Proceedings of the 31st International Conference on Machine Learning (ICML), vol. 32 of JMLR Workshop and Conference Proceedings, 2014, pp. 1539–1547.
 - [48] A. USCHMAJEV AND B. VANDEREYCKEN, *The geometry of algorithms using hierarchical tensors*, Linear Algebra Appl., 439 (2013), pp. 133–166.
 - [49] ———, *Line-search methods and rank increase on low-rank matrix varieties*, in Proceedings of the 2014 International Symposium on Nonlinear Theory and its Applications (NOLTA2014), 2014, pp. 52–55.
 - [50] B. VANDEREYCKEN, *Low-rank matrix completion by Riemannian optimization*, SIAM J. Optim., 23 (2013), pp. 1214–1236.
 - [51] B. VANDEREYCKEN AND S. VANDEWALLE, *A Riemannian optimization approach for computing low-rank solutions of Lyapunov equations*, SIAM J. Matrix Anal. Appl., 31 (2010), pp. 2553–2579.
 - [52] Y. XU AND W. YIN, *A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion*, SIAM J. Imaging Sci., 6 (2013), pp. 1758–1789.

Recent publications:

MATHEMATICS INSTITUTE OF COMPUTATIONAL SCIENCE AND ENGINEERING
Section of Mathematics
Ecole Polytechnique Fédérale
CH-1015 Lausanne

- 01.2014** GIOVANNI MIGLIORATI:
Multivariate Markov-type and Nicolskii-type inequalities for polynomials associated with downward closed multi-index sets
- 02.2014** FEDERICO NEGRI, ANDREA MANZONI, GIANLUIGI ROZZA:
Certified reduced basis method for parametrized optimal control problems governed by the Stokes equations
- 03.2014** CEDRIC EFFENBERGER, DANIEL KRESSNER:
On the residual inverse iteration for nonlinear eigenvalue problems admitting a Rayleigh functional
- 04.2014** TAKAHITO KASHIWABARA, CLAUDIA M. COLCIAGO, LUCA DEDÈ, ALFIO QUARTERONI:
Numerical Well-posedness, regularity, and convergence analysis of the finite element approximation of a generalized Robin boundary value problem
- 05.2014** BJÖRN ADLERBORN, BO KAGSTRÖM, DANIEL KRESSNER:
A parallel QZ algorithm for distributed memory HPC systems
- 06.2014** MICHELE BENZI, SIMONE DEPARIS, GWENOL GRANDPERRIN, ALFIO QUARTERONI:
Parameter estimates for the relaxed dimensional factorization preconditioner and application to hemodynamics
- 07.2014** ASSYR ABDULLE, YUN BAI:
Reduced order modelling numerical homogenization
- 08.2014** ANDREA MANZONI, FEDERICO NEGRI:
Rigorous and heuristic strategies for the approximation of stability factors in nonlinear parametrized PDEs
- 09.2014** PENG CHEN, ALFIO QUARTERONI:
A new algorithm for high-dimensional uncertainty quantification problems based on dimension-adaptive and reduced basis methods
- 10.2014** NATHAN COLLIER, ABDUL-LATEEF HAJI-ALI, FABIO NOBILE, ERIK VON SCHWERIN, RAÚL TEMPONE:
A continuation multilevel Monte Carlo algorithm
- 11.2014** LUKA GRUBISIC, DANIEL KRESSNER:
On the eigenvalue decay of solutions to operator Lyapunov equations
- 12.2014** FABIO NOBILE, LORENZO TAMELLINI, RAÚL TEMPONE:
Convergence of quasi-optimal sparse grid approximation of Hilbert-valued functions: application to random elliptic PDEs
- 13.2014** REINHOLD SCHNEIDER, ANDRÉ USCHMAJEW:
Convergence results for projected line-search methods on varieties of low-rank matrices via Lojasiewicz inequality