

Gauss-Southwell type descent methods for low-rank matrix optimization

Guillaume Olikier, André Uschmajew, Bart Vandereycken

Angaben zur Veröffentlichung / Publication details:

Olikier, Guillaume, André Uschmajew, and Bart Vandereycken. 2025.
“Gauss-Southwell type descent methods for low-rank matrix optimization.”
Journal of Optimization Theory and Applications 206 (1): 6.
<https://doi.org/10.1007/s10957-025-02682-9>.

Nutzungsbedingungen / Terms of use:

CC BY 4.0

Dieses Dokument wird unter folgenden Bedingungen zur Verfügung gestellt: / This document is made available under these conditions:

CC-BY 4.0: Creative Commons: Namensnennung

Weitere Informationen finden Sie unter: / For more information see:

<https://creativecommons.org/licenses/by/4.0/deed.de>





Gauss–Southwell Type Descent Methods for Low-Rank Matrix Optimization

Guillaume Olikier¹ · André Uschmajew² · Bart Vandereycken³

Received: 17 July 2023 / Accepted: 21 March 2025
© The Author(s) 2025

Abstract

We consider gradient-related methods for low-rank matrix optimization with a smooth cost function. The methods operate on single factors of the low-rank factorization and share aspects of both alternating and Riemannian optimization. Two possible choices for the search directions based on Gauss–Southwell type selection rules are compared: one using the gradient of a factorized non-convex formulation, the other using the Riemannian gradient. While both methods provide gradient convergence guarantees that are similar to the unconstrained case, numerical experiments on a quadratic cost function indicate that the version based on the Riemannian gradient is significantly more robust with respect to small singular values and the condition number of the cost function. As a side result of our approach, we also obtain new convergence results for the alternating least squares method.

Keywords Low-rank matrix optimization · First-order method · Riemannian optimization · Convergence analysis · Critical point

Mathematics Subject Classification 65K05 · 65K10 · 65F55 · 90C26 · 90C30

Communicated by Edouard Pauwels.

✉ André Uschmajew
andre.uschmajew@uni-a.de

Guillaume Olikier
guillaume.olikier@inria.fr

Bart Vandereycken
bart.vandereycken@unige.ch

¹ Epione Project Team, Université Côte d’Azur and Inria, 2004 route des Lucioles, BP 93, 06902 Sophia Antipolis Cedex, France

² Institute of Mathematics & Centre for Advanced Analytics and Predictive Sciences, University of Augsburg, 86159 Augsburg, Germany

³ Section of Mathematics, University of Geneva, 1211 Geneva, Switzerland

1 Introduction

We consider a differentiable function $f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ and the optimization problem

$$\min_{X \in \mathbb{R}^{m \times n}} f(X) \quad \text{s.t.} \quad \text{rank}(X) \leq k, \quad (1.1)$$

where m, n , and k are positive integers such that $k < \min\{m, n\}$. The constraint set

$$\mathcal{M}_{\leq k} := \{X \in \mathbb{R}^{m \times n} : \text{rank}(X) \leq k\} \quad (1.2)$$

is non-convex, which makes it difficult to find the global minimum even if the function f would be convex. Nevertheless, very efficient first-order algorithms for finding critical points can be implemented based on alternating minimization.

At the core, these algorithms are derived from the fact that the rank constraint in (1.1) can be eliminated by using the bilinear factorization

$$X = LR^\top, \quad L \in \mathbb{R}^{m \times k}, \quad R \in \mathbb{R}^{n \times k}$$

of a matrix of rank at most k . Problem (1.1) is hence equivalent to the unconstrained minimization problem

$$\min_{\substack{L \in \mathbb{R}^{m \times k} \\ R \in \mathbb{R}^{n \times k}}} g(L, R) \quad \text{with} \quad g(L, R) := f(LR^\top). \quad (1.3)$$

In the terminology of the recent work [9], the map $\mathbb{R}^{m \times k} \times \mathbb{R}^{n \times k} \rightarrow \mathbb{R}^{m \times n} : (L, R) \mapsto LR^\top$ is a smooth lift of the determinantal variety $\mathcal{M}_{\leq k}$ [5, Lecture 9] and g is the lifted cost function. The class of alternating optimization methods then consists of algorithms of the following general form:

- (1) Given $X_\ell = L_\ell R_\ell^\top$ in factorized form, choose whether to update the L - or R -factor.
- (2) Depending on the choice in (1) find an update $\hat{L}_{\ell+1}$ or $\hat{R}_{\ell+1}$ and set $X_{\ell+1} := \hat{L}_{\ell+1} R_\ell^\top$ or $X_{\ell+1} := L_\ell \hat{R}_{\ell+1}^\top$, respectively.
- (3) Change the representation: choose $L_{\ell+1}$ and $R_{\ell+1}$ such that $X_{\ell+1} = L_{\ell+1} R_{\ell+1}^\top$.

In this paper we consider two different methods of this type (Algorithms 1 and 2 below), where the search directions in step (2) are gradient-related in a suitable sense and step (1) is based on a Gauss–Southwell type selection rule.

The most obvious search directions are the negative partial gradients $-\nabla_L g(L_\ell, R_\ell)$ and $-\nabla_R g(L_\ell, R_\ell)$. Unfortunately, this straightforward choice suffers from the non-uniqueness of the representation $X = LR^\top$ in a twofold way. First, the resulting search direction in the image space $\mathbb{R}^{m \times n}$ depends on the particular chosen representation $X_\ell = L_\ell R_\ell^\top$ (see the formulas (3.5) below), which is conceptually unsatisfying. Second, the Lipschitz constant for the partial gradients depends on the chosen representation as well, which in theory may pose non-uniform step-size restrictions for ensuring convergence. And indeed, it can be observed in practice that this has a large influence on the convergence speed. Both problems, however, can be overcome

by using a so-called balanced factorization in step (3) of the above pseudocode for removing the non-uniqueness (see Sect. 3.1), as has been proposed in several works; see, for example, [10, 11, 28, 29] where balancing is indirectly promoted via regularization. While such a regularization may have other potential benefits, we omit it here in order to not change the cost function and instead apply explicit balancing after every step, which comes at negligible additional cost.

Regarding the choice of the factor to be updated in step (1), the cyclic Gauss–Seidel type update rule is the most common choice and optimizes L and R alternately. For example, variants of the alternating gradient descent method (using the partial gradients of g) have been proposed and analyzed for low-rank matrix completion in [23] and [25]. However, instead of a cyclic rule, in this work we advocate for using Gauss–Southwell type update rules in which the block variable with the largest partial gradient is selected for the next update. This selection rule originates from classic iterative solvers for linear systems [21] and has also been well known for nonlinear optimization ever since; see, e.g., [13, 15, 18] for convex problems and [17] for a general overview. The general advantage of the Gauss–Southwell rule is that the norm of the selected block gradient is at least a fraction of the norm of the full gradient of g , which allows to deduce convergence rate estimates for obtaining ε -critical points by simply applying the descent property of gradient descent (see (2.3)) for the partial gradient steps. However, in the context of low-rank optimization, one known issue, also observed for other update rules, then still remains: the convergence rate in terms of function values or distances to solutions severely deteriorates for ill-conditioned solutions (accumulation points with a large ratio of the largest and k th largest singular value), one intrinsic reason being that the gradient of the function g is not a robust measure of criticality on the manifold of rank- k matrices.

As an alternative to the partial gradients of g , we will therefore also consider a different way of obtaining search directions, which is motivated from a Riemannian optimization perspective. The bilinear nature of the parametrization of the set $\mathcal{M}_{\leq k}$ entails that at the smooth points X of that set (the matrices of rank exactly k), the *tangent space* is decomposed into two overlapping linear spaces $T_1(X) + T_2(X)$, each of which contains X and is itself entirely contained in $\mathcal{M}_{\leq k}$. It is therefore natural to use the orthogonal projections of the negative gradient $-\nabla f(X_\ell)$ onto $T_1(X_\ell)$ or $T_2(X_\ell)$ as search directions. These search directions are *independent* of the chosen factorization $X_\ell = L_\ell R_\ell^\top$ and actually correspond to partial gradients of g when the factor not to be updated is chosen to have orthonormal columns, which can be achieved by changing the representation using QR factorization after every update. Such a variant of alternating optimization including QR factorization is usually implemented anyway for reasons of numerical stability. As an alternative interpretation, taking orthogonal projections of $-\nabla f(X)$ onto $T_1(X)$ and $T_2(X)$ as search directions turns out to be mathematically equivalent to taking so-called *scaled* (or *preconditioned*) variants $-\nabla_L g(L, R) \cdot (R^\top R)^{-1}$ and $-\nabla_R g(L, R) \cdot (L^\top L)^{-1}$ of the partial gradients of g for updating L and R ; see (3.12) below. Such scaled gradients have been considered for example in [14, 25, 26]. Notably, in [14, §4] they are derived from choosing a suitable quotient geometry for the manifold of fixed-rank matrices. Corresponding optimization methods based on cyclic or simultaneous update rules are called *scaled alternating steepest descent* (ScaledASD) in [25] and *scaled gradient descent* (ScaledGD) in [26].

In these works, the accelerating effect and/or robustness with respect to small singular values is clearly confirmed; see also [12, Remark 7]. While [25, Theorem 4.1] proves stationarity of accumulation points of ScaledASD for a non-convex formulation of matrix completion under minimal assumptions, but without rate, [26, Theorem 11] provides local linear convergence rates for ScaledGD close to global minima of general cost functions with rank-restricted convexity properties in noiseless scenarios, including matrix sensing under restricted isometry assumptions or matrix completion under incoherence assumptions.

The observation that the scaled partial gradients of $g(L, R)$ correspond to orthogonal projections of $\nabla f(X)$ onto linear subspaces within $\mathcal{M}_{\leq k}$, although not difficult, could be new and gives these methods an interpretation as Riemannian optimization methods on the embedded rank- k manifold. In particular, as we show in the present paper, when combined with a Gauss–Southwell rule, this viewpoint allows to analyze the convergence of the corresponding method (presented in Algorithm 2) as a gradient-related Riemannian descent method. In fact, the search directions then satisfy an angle condition with the negative *Riemannian gradient* (here: the projection of $-\nabla f(X)$ onto the full tangent space) and hence the convergence rate to an ε -Riemannian critical point can be estimated in a quite simple way based on standard results. This is carried out in Sect. 3.2, the main results being a global $O(1/\sqrt{\ell})$ convergence rate for the Riemannian gradient. In terms of the above references, Algorithm 2 may be equivalently interpreted as a Gauss–Southwell variant of scaled gradient descent, and while our convergence result seems to be new in this general form, we do not claim (and did not test) that a Gauss–Southwell rule for scaled gradients leads to better performance than other variants. Instead, our focus in this work is on the simple convergence analysis that it admits, despite being a Riemannian method. We note that a method of this type has already been proposed as a retraction-free Riemannian method in [22, Algorithm 4].

Having two possible choices of gradient-related search directions, it is then natural to ask which of the two is better. In terms of the rates, it turns out that the Riemannian version needs the same (estimated) number of steps for making the Riemannian gradient of f smaller than ε as the balanced factorized version needs for making the gradient of g smaller than ε . However, as we show in Sect. 3.3, the first condition implies the second (up to a small constant) but not the other way around. In this sense, we will confirm as one of the main messages of the paper that the Riemannian search directions are much more powerful than expected. Our numerical experiments in Sect. 4 indeed confirm that it converges faster and exhibits stability with respect to small singular values of X , which is completely in line with aforementioned results on scaled gradient descent in [26].

One important thing to note about both of our methods is that accumulation points can fail to be critical points of the original problem (1.1), despite the Riemannian gradients converging to zero during the iteration. Although it is rarely observed in practice, such a behavior, termed an “apocalypse” in [8], can occur when a rank drop happens at the limit; cf. Remark 3.3 below. For the Riemannian version of our algorithm, it has been recently proposed how to exclude such undesirable points as accumulation points of the algorithms by disregarding the k th singular value if it is below a prescribed threshold [20, Algorithm 3], while preserving the convergence

rates for the Riemannian gradient as stated in our Corollary 3.5. In our work, we will largely ignore this subtlety (and assume full rank k at accumulation points) in order to keep things conceptually simple.

There are two other sets of results in this work. In Sect. 3.4, we address the fact that the algebraic (i.e., sublinear) convergence rates for gradients obtained from a standard analysis based on Lipschitz continuity of the gradients are usually overly pessimistic. Indeed, in practice one observes a linear convergence rate. For the Riemannian version of our Gauss–Southwell approach we provide such a result in Sect. 3.4 provided one of the accumulation points has a positive-definite Riemannian Hessian. This is a standard assumption for obtaining linear rates and for instance also implied by the assumptions of [26, Theorem 11] on local linear convergence of ScaledGD. Our proof is similar to the linear convergence proof for the Riemannian gradient method in [2, Theorem 4.5.6]. In Sect. 3.5 we show how our arguments can be adapted for obtaining seemingly novel convergence estimates for the popular alternating least squares (ALS) algorithm. In fact, for two blocks as in low-rank matrix optimization, the cyclic rule of ALS coincides with a Gauss–Southwell rule, since one of the partial gradients is always zero by construction. Our global $O(1/\sqrt{\ell})$ rate for the Riemannian gradients in ALS as presented in Corollary 3.9 is insofar remarkable, as most of the available results for alternating minimization only deal with convex problems (which allows for better rates; see, e.g., [6] for an overview). Our arguments, however, share some similarity with the results in [3, Sect. 3.1] on general non-convex functions, but are more specific to the low-rank optimization scenario by considering scaled and Riemannian gradients.

The remainder of this work is structured as follows. In Sect. 2 we review basic facts on the gradient method for λ -smooth functions on affine subspaces, which will serve as the main tool. Section 3 is the core section of this work. Section 3.1 develops the algorithm based on partial gradients of the function g and balancing (Algorithm 1). The main results on gradient convergence for various step-size rules are Theorem 3.1 and Corollary 3.2. In Sect. 3.2, the algorithm based on the decomposition of the Riemannian gradient (Algorithm 2) is developed in an analogous fashion. Main results are Theorem 3.4 and Corollary 3.5. In Sect. 3.3 both approaches are compared; corresponding numerical experiments are presented in Sect. 4. The results regarding local linear convergence and ALS are given in Sects. 3.4 and 3.5, respectively.

2 Recap of Gradient Descent Methods

Let us first briefly recall some main results on the convergence of the gradient method for λ -smooth functions. While the results are very well known (see, e.g., [16]), we present them here for the case of a differentiable function $\phi: T \rightarrow \mathbb{R}$ defined on an affine subspace T of a Euclidean vector space. In our applications later, ϕ will be the restriction of the cost functions f and g , respectively, to certain subspaces. Given $x \in T$, one step of the gradient method in the space T takes the form

$$x_+ := x - h\nabla\phi(x). \quad (2.1)$$

Here $h > 0$ is a step size that can be chosen either fixed or adaptively at each iteration. The ultimate goal of the iteration is to find a point $x^* \in T$ satisfying $\nabla\phi(x^*) = 0$.

The easiest analysis of the gradient descent step (2.1) assumes that the function ϕ is λ -smooth, which means that $\nabla\phi$ is Lipschitz continuous with constant $\lambda \geq 0$:

$$\|\nabla\phi(x) - \nabla\phi(y)\| \leq \lambda\|x - y\| \text{ for all } x, y \in T.$$

The key estimate implied by this property is the following. If ϕ is λ -smooth on the line segment connecting x and y , then

$$\phi(x) - \phi(y) \geq \langle \nabla\phi(x), x - y \rangle - \frac{\lambda}{2}\|x - y\|^2. \quad (2.2)$$

The proof is straightforward from a Taylor expansion with integral remainder; see, e.g., [16, Lemma 1.2.3].

If applicable, (2.2) yields the following estimate for one step of the gradient method:

$$\phi(x) - \phi(x_+) \geq \left(h - \frac{\lambda}{2}h^2\right) \|\nabla\phi(x)\|^2. \quad (2.3)$$

The factor in front of $\|\nabla\phi(x)\|^2$ is positive for step sizes $h = \frac{2\alpha}{\lambda}$ with $\alpha \in (0, 1)$, leading to

$$\phi(x) - \phi(x_+) \geq \frac{2}{\lambda}\alpha(1 - \alpha)\|\nabla\phi(x)\|^2.$$

The “optimal” step size that maximizes the right-hand side in this estimate is

$$h^* := \frac{1}{\lambda}$$

and yields

$$\phi(x) - \phi(x_+) \geq \frac{1}{2\lambda}\|\nabla\phi(x)\|^2. \quad (2.4)$$

The “optimal” step size h^* has the disadvantage that it requires knowing a conservative estimate for λ , which may lead to h^* being unnecessarily small. There are several adaptive strategies for addressing this. For instance, when ϕ is a convex quadratic function, it is possible to find a so-called “exact” step size that minimizes the function $h \mapsto \phi(x - h\nabla\phi(x))$. Using this step size we obviously obtain the same estimate as (2.4), given that the λ -smoothness condition is satisfied.

A more general strategy is line search. Here the Armijo step-size rule is among the most well known: given $\gamma \in (0, 1)$ find a step size $h > 0$ satisfying the inequality

$$\phi(x) - \phi(x_+) \geq h\gamma\|\nabla\phi(x)\|^2. \quad (2.5)$$

To avoid an unnecessarily small h , one typically uses a backtracking procedure by checking step sizes of the form

$$h = \bar{h}, \bar{h}\beta, \bar{h}\beta^2, \dots$$

in decreasing order with some fixed $\bar{h} \in (0, \infty)$ and $\beta \in (0, 1)$. By (2.3), the Armijo condition (2.5) will be satisfied once $h - \frac{\lambda}{2}h^2 \geq h\gamma$, or equivalently,

$$h \leq \frac{2}{\lambda}(1 - \gamma). \tag{2.6}$$

Therefore, the backtracking procedure will terminate with $h = \bar{h}\beta^\ell$ for some

$$\ell \leq \max \left\{ 0, \left\lceil \ln \left(\frac{2}{\bar{h}\lambda}(1 - \gamma) \right) / \ln \beta \right\rceil \right\}.$$

To avoid the rounding above, we note that $h \geq \min\{\frac{2}{\lambda}\beta(1 - \gamma), \bar{h}\}$ since we stop backtracking as soon as (2.6) is satisfied. The resulting estimate (2.5) is

$$\phi(x) - \phi(x_+) \geq \min \left\{ \frac{2}{\lambda}\beta\gamma(1 - \gamma), \bar{h}\gamma \right\} \|\nabla\phi(x)\|^2.$$

For later reference we emphasize that in this point-wise estimate for the Armijo step size it is in principle possible to replace a global smoothness constant λ with a restricted smoothness constant on the line segment connecting x and $x - \bar{h}\nabla f(x)$, that is, with a constant $\lambda_x \leq \lambda$ satisfying

$$\|\nabla\phi(x) - \nabla\phi(x - t\bar{h}\nabla\phi(x))\| \leq \lambda_x t\bar{h}\|\nabla\phi(x)\| \tag{2.7}$$

for all $0 \leq t \leq 1$. In this sense the Armijo step-size rule with backtracking automatically detects the locally optimal smoothness constant. We will make use of this fact in the proof of Theorem 3.8 further below.

3 Block Gradient Descent for Low-Rank Matrices

In this main section we show how to apply the above framework of gradient descent on subspaces to low-rank matrix optimization problems. We first consider the problem in the factorized form (1.3). Here the goal is to minimize the unconstrained function

$$g(L, R) := f(LR^\top)$$

for $L \in \mathbb{R}^{m \times k}$ and $R \in \mathbb{R}^{n \times k}$. The most obvious first-order method based on alternating optimization consists in taking gradient steps along the L or R block only, that is, taking $-\nabla_L g(L, R)$ or $-\nabla_R g(L, R)$ as a search direction. Depending on which block variable is chosen, we then update the point $X = LR^\top$ by using either

$$L_+ := L - h\nabla_L g(L, R), \quad R_+ := R, \tag{3.1}$$

or

$$L_+ := L, \quad R_+ := R - h\nabla_R g(L, R). \tag{3.2}$$

Observe that the first update (3.1) is in fact a gradient step for the function g when restricted to

$$\hat{T}_R := \{(\hat{L}, R) : \hat{L} \in \mathbb{R}^{m \times k}\} \quad (3.3)$$

with R fixed. Likewise, the second update (3.2) is a gradient step for g restricted to

$$\hat{T}_L := \{(L, \hat{R}) : \hat{R} \in \mathbb{R}^{n \times k}\} \quad (3.4)$$

with L fixed. Since \hat{T}_R and \hat{T}_L are *affine* subspaces of $\mathbb{R}^{m \times k} \times \mathbb{R}^{n \times k}$, we can analyze these updates using the results from Sect. 2 by taking $\phi := g|_{\hat{T}_R}$ and $\phi := g|_{\hat{T}_L}$, respectively.

To this end, let us estimate the Lipschitz constants of $\nabla_L g(L, R)$ and $\nabla_R g(L, R)$ on these affine subspaces. Note that with $X = LR^\top$ the partial gradients are given as

$$\nabla_L g(L, R) = \nabla f(X)R, \quad \nabla_R g(L, R) = \nabla f(X)^\top L. \quad (3.5)$$

We assume that f is λ -smooth for the Frobenius norm $\|\cdot\|_F$ on $\mathbb{R}^{m \times n}$. Then

$$\|\nabla_L g(L_1, R) - \nabla_L g(L_2, R)\|_F \leq \lambda \|L_1 R^\top - L_2 R^\top\|_F \|R\|_2 \leq \lambda \|R\|_2^2 \|L_1 - L_2\|_F \quad (3.6)$$

for all $L_1, L_2 \in \mathbb{R}^{m \times k}$, that is, the function $L \mapsto \nabla_L g(L, R)$ has a Lipschitz constant $\lambda \|R\|_2^2$ on \hat{T}_R . Here $\|R\|_2$ denotes the spectral norm. Similarly, we have the estimate

$$\|\nabla_R g(L, R_1) - \nabla_R g(L, R_2)\|_F \leq \lambda \|L\|_2^2 \|R_1 - R_2\|_F \quad (3.7)$$

for all $R_1, R_2 \in \mathbb{R}^{n \times k}$, that is, we obtain the Lipschitz constant $\lambda \|L\|_2^2$.

Observe, however, the factors L and R in the representation of $X = LR^\top$ are not unique, and so the Lipschitz constants depend on the chosen representation. In order to get independent constants some additional normalization will be necessary. In the following we discuss two approaches. The first one, discussed in the next section, is based on balancing the singular values among the factors. The second approach, discussed afterwards, uses orthogonalization and leads to a Riemannian interpretation.

3.1 Balanced Factorized Version

The low-rank representation $X = LR^\top$ is said to be *balanced* if L and R have the same singular values. Given an economy-sized SVD $X = U\Sigma V^\top$, this is accomplished by choosing

$$L := U\Sigma^{1/2}, \quad R := V\Sigma^{1/2}.$$

For a balanced representation, the estimates (3.6) and (3.7) become

$$\|\nabla_L g(L_1, R) - \nabla_L g(L_2, R)\|_F \leq \lambda \|X\|_2 \|L_1 - L_2\|_F \quad (3.8)$$

and

$$\|\nabla_R g(L, R_1) - \nabla_R g(L, R_2)\|_F \leq \lambda \|X\|_2 \|R_1 - R_2\|_F, \tag{3.9}$$

that is, we get the Lipschitz constant $\lambda \|X\|_2$ for both partial gradients.

As mentioned above, our goal is to apply the results from Sect. 2 to bound the decrease in function value $f(X) - f(X_+)$ after one gradient step in either L or R direction. This will result in a lower bound in terms of $\|\nabla_L g(L, R)\|_F$ or $\|\nabla_R g(L, R)\|_F$ using suitable step sizes. If we now choose the block variables L and R according to which partial gradient is larger, we can also bound this function value decrease by the (Euclidean) norm of the full gradient, which will be denoted as

$$\|\nabla g(L, R)\|_F := \sqrt{\|\nabla_L g(L, R)\|_F^2 + \|\nabla_R g(L, R)\|_F^2}.$$

Such a selection rule is called a Gauss–Southwell rule in the context of block coordinate descent (BCD) methods.

To repeat the process, the factors need to be balanced after every update. This requires to compute an SVD of a $k \times k$ matrix in every step, which is cheap if k is small. The final algorithm is displayed in Algorithm 1.

Algorithm 1 Balanced block gradient descent with Gauss–Southwell rule

Input: $L_0 \in \mathbb{R}^{m \times k}$, $R_0 \in \mathbb{R}^{n \times k}$ balanced

1: **for** $\ell = 0, 1, 2, \dots$ **do**

2: Compute $G := \nabla f(L_\ell R_\ell^\top)$ and partial gradients:

$$(G_1, G_2) := \nabla g(L_\ell, R_\ell) = (GR_\ell, G^\top L_\ell).$$

3: Perform a gradient step according to Gauss–Southwell rule with suitable step h_ℓ :

$$(L_+, R_+) := \begin{cases} (L_\ell - h_\ell G_1, R_\ell) & \text{if } \|G_1\|_F \geq \|G_2\|_F, \\ (L_\ell, R_\ell - h_\ell G_2) & \text{else.} \end{cases}$$

4: Compute the compact QR factorizations $Q_1 S_1 = L_+$ and $Q_2 S_2 = R_+$, the SVD factorization $U \Sigma V^\top = S_1 S_2^\top$, and balance:

$$L_{\ell+1} := Q_1 U \Sigma^{1/2}, \quad R_{\ell+1} := Q_2 V \Sigma^{1/2}.$$

5: **end for**

We can make the following convergence statements for this method, depending on the type of step-size rule that is applied.

Theorem 3.1 *Assume that $f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is λ -smooth and that the constrained sublevel set $N_0 := \{X \in \mathcal{M}_{\leq k} : f(X) \leq f(L_0 R_0^\top)\}$ is contained in a ball $\{X \in \mathbb{R}^{m \times n} : \|X\|_2 \leq \rho\}$. Depending on the chosen step size in the ℓ th step, the iterates of Algorithm 1 satisfy*

$$g(L_\ell, R_\ell) - g(L_{\ell+1}, R_{\ell+1}) \geq \frac{\vartheta_\ell}{2\lambda\rho} \|\nabla g(L_\ell, R_\ell)\|_F^2, \tag{3.10}$$

where:

- $\vartheta_\ell := 2\alpha_\ell(1 - \alpha_\ell)$ in case a fixed step size $h_\ell = \frac{2\alpha_\ell}{\lambda\rho}$ with $\alpha_\ell \in (0, 1)$ is used;
- $\vartheta_\ell := \frac{1}{2}$ in case the step size is obtained by exact line search (if well defined);
- $\vartheta_\ell := \min \{2\beta\gamma(1 - \gamma), \lambda\rho\gamma\bar{h}\}$ in case the step size is obtained from backtrackinging $h_\ell = \bar{h}, \bar{h}\beta, \bar{h}\beta^2, \dots$ until the Armijo condition

$$g(L_\ell, R_\ell) - g(L_{\ell+1}, R_{\ell+1}) \geq \gamma h_\ell \|G_{1/2}\|_F^2$$

is fulfilled (here $\beta, \gamma \in (0, 1)$ and $G_{1/2}$ denotes G_1 or G_2 depending on the block selected in this step). The backtracking procedure terminates with $h_\ell = \bar{h}\beta^i$ for some $i \leq \max \left\{ 0, \left\lceil \ln \left(\frac{2}{\bar{h}\lambda\rho} (1 - \gamma) \right) / \ln \beta \right\rceil \right\}$.

Corollary 3.2 Under the conditions of Theorem 3.1, let (L_ℓ, R_ℓ) be the iterates of Algorithm 1 using any combination of the considered step-size rules. However, if step sizes $h_\ell = \frac{2\alpha_\ell}{\lambda\rho}$ from the first step-size rule are used infinitely often, assume (for this subsequence) $\inf_\ell \alpha_\ell > 0$ and $\sup_\ell \alpha_\ell < 1$. Then the generated sequence $g(L_\ell, R_\ell) = f(L_\ell R_\ell^\top)$ of function values is monotonically decreasing and converges to some value $f_* \geq \min_{X \in \mathcal{M}_{\leq k}} f(X)$. Moreover, $\nabla g(L_\ell, R_\ell) \rightarrow 0$ and the sequence (L_ℓ, R_ℓ) has at least one accumulation point. Every accumulation point (L_*, R_*) satisfies $g(L_*, R_*) = f_*$ and $\nabla g(L_*, R_*) = 0$. For every nonnegative integer j it holds that

$$\min_{0 \leq \ell \leq j} \|\nabla g(L_\ell, R_\ell)\|_F \leq \left(\frac{2\lambda\rho}{\vartheta} \cdot \frac{f(X_0) - f_*}{j + 1} \right)^{1/2}, \tag{3.11}$$

where $\vartheta := \inf_\ell \vartheta_\ell > 0$. In particular, given $\varepsilon > 0$ the algorithm returns an ε -critical point, i.e., a point (L, R) such that $\|\nabla g(L, R)\|_F \leq \varepsilon$, after at most $\lceil 2\lambda\rho(f(X_0) - f_*)\vartheta^{-1}\varepsilon^{-2} - 1 \rceil$ iterations.

Note that $0 < \vartheta_\ell \leq \frac{1}{2}$ in all cases of Theorem 3.1 and hence $0 < \vartheta \leq \frac{1}{2}$ in Corollary 3.2.

Remark 3.3 Before giving the proofs, we make some remarks. The first is that λ -smoothness is in principle only required on the sublevel set N_0 of the starting point. In addition, the exact line search is well defined since N_0 is bounded.

The second remark is that in the version with exact line search one could, as an a posteriori estimate, replace for given $\varepsilon > 0$ and large enough $\ell = \ell(\varepsilon)$ the constant ρ in (3.10) with $\max \|X_*\|_2 + \varepsilon$ where $\max \|X_*\|_2$ is the maximum spectral norm among all accumulation points. (For the other two versions the accumulation points depend on ρ through step sizes.)

Finally, as mentioned in the introduction, the result $\nabla g(L_*, R_*) = 0$ does not necessarily imply that $X_* = L_* R_*^\top$ is a critical point of f on the variety $\mathcal{M}_{\leq k}$ (in the sense that the tangent cone to $\mathcal{M}_{\leq k}$ at X_* contains no descent direction for f). By [9, Proposition 2.8], we know that the implication holds if $\text{rank}(L_*) = \text{rank}(R_*) = k$ but may not hold otherwise. If the implication does not hold, X_* is only critical for f on the manifold of matrices with fixed rank $s := \text{rank}(X_*) < k$. Such a triplet

$(X_*, (L_\ell R_\ell^\top)_{\ell \in \mathbb{N}}, f)$ is called an apocalypse in [8]. An apocalypse-free version of Algorithm 2 based on the original method from [22, Algorithm 4] has been recently proposed in [20, Algorithm 3].

Proof of Theorem 3.1. Let (L_ℓ, R_ℓ) be balanced such that $X_\ell = L_\ell R_\ell^\top \in N_0$. As noted above, the single updates are gradient descent steps for the function g restricted to the subspaces \hat{T}_{L_ℓ} and \hat{T}_{R_ℓ} , respectively, defined in (3.3) and (3.4). By (3.8) and (3.9), these restrictions are $(\lambda\rho)$ -smooth and the results of Sect. 2 allow us to conclude that

$$f(X_\ell) - f(X_{\ell+1}) = g(L_\ell, R_\ell) - g(L_{\ell+1}, R_{\ell+1}) \geq \frac{\vartheta_\ell}{\lambda\rho} \|G_{1/2}\|_{\mathbb{F}}^2$$

with ϑ_ℓ determined by the step-size rule as specified in the statement of the theorem. Since $\|\nabla g(L_\ell, R_\ell)\|_{\mathbb{F}}^2 = \|G_1\|_{\mathbb{F}}^2 + \|G_2\|_{\mathbb{F}}^2$, the maximum block selected by the Gauss–Southwell rule satisfies $\|G_{1/2}\|_{\mathbb{F}}^2 \geq \frac{1}{2} \|\nabla g(L_\ell, R_\ell)\|_{\mathbb{F}}^2$. This shows (3.10) after induction since $X_{\ell+1} = L_{\ell+1} R_{\ell+1}^\top \in N_0$ by the above estimates. \square

Proof of Corollary 3.2. Since the sequence $g(L_\ell, R_\ell) = f(X_\ell)$ is monotonically decreasing and bounded from below (since N_0 is bounded), it converges to $f_* := \inf_{\ell \in \mathbb{N}} f(X_\ell)$. Thus, for every $j \in \mathbb{N}$, it holds that

$$\begin{aligned} f(X_0) - f_* &\geq f(X_0) - f(X_{j+1}) = \sum_{\ell=0}^j f(X_\ell) - f(X_{\ell+1}) \\ &\geq \sum_{\ell=0}^j \frac{\vartheta}{2\lambda\rho} \|\nabla g(L_\ell, R_\ell)\|_{\mathbb{F}}^2, \end{aligned}$$

using the estimates from Theorem 3.1. Since j is arbitrary, this shows that $\|\nabla g(L_\ell, R_\ell)\|_{\mathbb{F}} \rightarrow 0$. On the other hand, there must be at least one $0 \leq \ell \leq j$ satisfying

$$\frac{\vartheta}{2\lambda\rho} \|\nabla g(L_\ell, R_\ell)\|_{\mathbb{F}}^2 \leq \frac{f(X_0) - f_*}{j + 1}.$$

This shows (3.11). Finally, since N_0 is bounded and $\|L_\ell\|_2 = \|R_\ell\|_2 = \|X_\ell\|_2^{1/2}$, the sequence (L_ℓ, R_ℓ) is bounded and thus has at least one accumulation point (L_*, R_*) . By continuity, $g(L_*, R_*) = f_*$ and $\nabla g(L_*, R_*) = 0$. \square

3.2 Embedded Riemannian Version

As derived in (3.5), the search directions in a single step of the “vanilla” block gradient descent method for the function $g(L, R) := f(LR^\top)$ depend on the chosen factorization $X = LR^\top$. In Algorithm 1 this ambiguity is resolved using a balancing, which is a reasonable approach, but arguably still somewhat arbitrary. In particular, it leads to the appearance of the spectral norm of X in the Lipschitz constants (see (3.8) and (3.9)) and thus in the convergence estimates (through the constant ρ). If instead of

balancing one could choose the factor that is currently *not* optimized to have orthonormal columns, this artefact would disappear (see (3.6) and (3.7)). In a cyclic update scheme for the factors this can be easily achieved and has been considered in many works. However, for the Gauss–Southwell rule it is not so clear how to combine it with orthogonalization. And even so, the partial gradients of g would still depend on the choice of the orthogonal factor, which is not unique.

A natural solution to these issues with the factorized approach can be obtained by taking a subspace viewpoint in the actual ambient space $\mathbb{R}^{m \times n}$. For this, besides the low-rank factorization $X = LR^\top$ we also consider in the sequel an SVD-like decomposition $X = USV^\top$ where $S \in \mathbb{R}^{k \times k}$ and U and V have pairwise orthonormal columns. Assume $\text{rank}(X) = k$. We now consider the subspace of matrices whose row space is contained in the row space of X ,

$$T_V := \{\hat{L}V^\top : \hat{L} \in \mathbb{R}^{m \times k}\} \subseteq \mathbb{R}^{m \times n},$$

and likewise for the column space,

$$T_U := \{U\hat{R}^\top : \hat{R} \in \mathbb{R}^{n \times k}\} \subseteq \mathbb{R}^{m \times n}.$$

Observe that T_V and T_U are *linear* subspaces contained in $\mathcal{M}_{\leq k}$ and both of them contain X . They are the respective images of the affine subspaces \hat{T}_R and \hat{T}_L from (3.3) and (3.4) under the map $\mathbb{R}^{m \times k} \times \mathbb{R}^{n \times k} \rightarrow \mathbb{R}^{m \times n} : (\hat{L}, \hat{R}) \mapsto \hat{L}\hat{R}^\top$. (When $\text{rank}(X) < k$ some subtleties regarding the choice of U and V occur.) However, contrary to \hat{T}_R and \hat{T}_L , the subspaces T_V and T_U are uniquely determined by X (in case $\text{rank}(X) = k$) and do not depend on any fixed factorization: even if V is not uniquely defined from $X = USV^\top$, any equivalent choice is of the form VQ for some orthogonal Q since $\text{rank}(X) = k$, and obviously it then holds $T_{VQ} = T_V$. The same observation holds for U and T_U .

Any optimization method based on sequential optimization of factors L and R in $X = LR^\top$ implicitly selects a search direction in the linear space T_V or T_U . For instance, by (3.5), the negative partial gradient $-\nabla_L g(L, R) = -\nabla f(X)R$ can be identified with the negative gradient of the restriction of g to \hat{T}_R , namely $(-\nabla f(X)R, R)$, which is an element of \hat{T}_R whose corresponding element in T_V is $-\nabla f(X)RR^\top$. In line with what was noted above, this particular direction depends on the choice of R (and L) in the decomposition $X = LR^\top$.

Gradient-related search directions that are independent from the factorization can be obtained based on information of the gradient $\nabla f(X)$ in the ambient space only. A natural choice is to consider the best approximations of $-\nabla f(X)$ in T_V or T_U w.r.t. the Frobenius inner product. These can be easily computed as the orthogonal projections of $-\nabla f(X)$ at $X = USV^\top$ onto T_V or T_U using the formulas

$$\mathcal{P}_V(Z) = ZVV^\top, \quad \mathcal{P}_U(Z) = UU^\top Z.$$

This yields the following type of iteration in the set $\mathcal{M}_{\leq k}$:

(1) Given $X = USV^\top \in \mathcal{M}_{\leq k}$, choose an update direction:

$$P_1 := \mathcal{P}_V(\nabla f(X)) = \nabla f(X)VV^\top \quad \text{or} \quad P_2 := \mathcal{P}_U(\nabla f(X)) = UU^\top \nabla f(X).$$

(2) Perform the step: $X_+ := X - hP_{1/2}$.

In terms of a factorization $X = LR^\top$, the two possible updates X_+ in this iteration can be expressed using pseudo-inverses and (3.5) as (assuming $\text{rank}(X) = k$)

$$\begin{aligned} X_+ &:= LR^\top - h\nabla f(X)RR^\dagger = [L - h\nabla_L g(L, R) \cdot (R^\top R)^{-1}]R^\top \\ \text{or} \quad X_+ &:= LR^\top - hLL^\dagger \nabla f(X) = L[R - h\nabla_R g(L, R) \cdot (L^\top L)^{-1}]^\top. \end{aligned}$$

As announced in the introduction, this shows that the Gauss–Southwell updates implicitly perform the updates

$$L_+ := L - h\nabla_L g(L, R) \cdot (R^\top R)^{-1} \quad \text{or} \quad R_+ := R - h\nabla_R g(L, R) \cdot (L^\top L)^{-1} \quad (3.12)$$

for the factors L and/or R , which coincide with the well-known *scaled gradient descent* update formulas; see, e.g., [26, (3)].

Obviously, the projected gradients P_1 and P_2 equal the gradients of the restriction of f to the linear subspaces T_V and T_U , respectively. Furthermore, if f is λ -smooth in Frobenius norm, then so is its restriction to any linear subspace. This means we can apply the abstract results recorded in Sect. 2 to analyze one step of such a method and obtain an estimate

$$f(X) - f(X_+) \geq \frac{\vartheta}{\lambda} \|P_{1/2}\|_F^2, \quad (3.13)$$

where the value of ϑ depends on the step-size rule and is specified in Theorem 3.4 below.

The estimate (3.13) indicates that a Gauss–Southwell rule could be useful in deciding whether P_1 or P_2 should be selected for an update direction. This has been proposed in [22, Algorithm 4] in a slightly more elaborate setup for dealing with rank-deficient iterates. Indeed, we can motivate the Gauss–Southwell rule from the perspective of Riemannian optimization as follows. It is well known that if $X = USV^\top$ has rank k , then the linear space

$$T_{U,V} := T_U + T_V$$

can be identified with the tangent space to the embedded submanifold \mathcal{M}_k of rank- k matrices at X . From a geometric perspective, our goal in low-rank optimization methods is to drive the Riemannian gradient of f to zero. When taking the embedded metric the following definition applies.

Definition The *Riemannian gradient* of f at $X \in \mathbb{R}^{m \times n}$, denoted by $\text{grad } f(X)$, is the orthogonal projection of $\nabla f(X)$ onto the tangent space to $\mathcal{M}_{\text{rank}(X)}$ at X .

Note that in this definition the fixed-rank manifold $\mathcal{M}_{\text{rank}(X)}$ to which the Riemannian gradient relates is implicitly determined by X . Hence, when $\text{rank}(X) = k$, it holds

$$\text{grad } f(X) = \mathcal{P}_{U,V}(\nabla f(X)),$$

where $\mathcal{P}_{U,V}$ is the orthogonal projector onto $T_{U,V}$, given by

$$\mathcal{P}_{U,V}(Z) = ZVV^\top + UU^\top Z - UU^\top ZVV^\top,$$

with

$$\|\mathcal{P}_{U,V}(Z)\|_F^2 = \|ZVV^\top\|_F^2 + \|UU^\top Z\|_F^2 - \|UU^\top ZVV^\top\|_F^2.$$

Applied to $\nabla f(X)$, this then implies

$$\|\mathcal{P}_{U,V}(\nabla f(X))\|_F^2 \leq \|\nabla f(X)VV^\top\|_F^2 + \|UU^\top \nabla f(X)\|_F^2 = \|P_1\|_F^2 + \|P_2\|_F^2,$$

which leads to the important observation

$$\max\{\|P_1\|_F^2, \|P_2\|_F^2\} \geq \frac{1}{2} \|\mathcal{P}_{U,V}(\nabla f(X))\|_F^2.$$

As a result, employing a Gauss–Southwell rule will allow to turn the estimate (3.13) into

$$f(X) - f(X_+) \geq \frac{\vartheta}{2\lambda} \|\text{grad } f(X)\|_F^2. \quad (3.14)$$

Based on this, a method with Gauss–Southwell rule allows for convergence statements regarding the Riemannian gradient without requiring any curvature bounds. The result is stated in Theorem 3.4 below.

Some subtleties arise when $X = USV^\top$ has rank less than k but the matrix S has size $k \times k$. This can happen in our algorithm since the sizes of the decomposition are fixed. It is not difficult to show that the corresponding tangent space to $\mathcal{M}_{\text{rank}(X)}$ at X is then contained in the space $T_{U,V}$. Hence we still get the relation

$$\|\mathcal{P}_{U,V}(\nabla f(X))\|_F \geq \|\text{grad } f(X)\|_F$$

so that (3.14) is valid for all $X = USV^\top$ when using the Gauss–Southwell rule.

While we have derived the method in the ambient space $\mathbb{R}^{m \times n}$, its practical implementation is of course done in the factorized form to keep benefit of the low-rank representation and avoid forming potentially huge matrices. Note in this context that we can exploit $\|\nabla f(X)VV^\top\|_F = \|\nabla f(X)V\|_F$ and $\|UU^\top \nabla f(X)\|_F = \|U^\top \nabla f(X)\|_F$ for deciding the Gauss–Southwell rule. In this form the algorithm is outlined in Algorithm 2, where $\text{qr}(Y) = (Q, R)$ denotes a QR decomposition of $Y = QR$ with k columns. For every iterate of rank k , a step of Algorithm 2 is equivalent to a step of [22, Algorithm 4].

Algorithm 2 Riemannian block gradient descent with Gauss–Southwell rule

Input: SVD-like decomposition $X_0 = U_0 S_0 V_0^T$ where $U_0 \in \mathbb{R}^{m \times k}$ and $V_0 \in \mathbb{R}^{n \times k}$, $U_0^T U_0 = V_0^T V_0 = I_k$.

- 1: **for** $\ell = 0, 1, 2, \dots$ **do**
- 2: Compute gradient $G := \nabla f(U_\ell S_\ell V_\ell^T)$ and

$$(\hat{P}_1, \hat{P}_2) := (GV_\ell, G^T U_\ell).$$

- 3: Perform a gradient step according to the Gauss–Southwell rule with suitable step size h_ℓ :
- 4: **if** $\|\hat{P}_1\|_F \geq \|\hat{P}_2\|_F$ **then**
- 5:

$$(U_{\ell+1}, S_{\ell+1}) := \text{qr}(U_\ell S_\ell - h_\ell \hat{P}_1), \quad V_{\ell+1} := V_\ell,$$

- 6: **else**
- 7:

$$U_{\ell+1} := U_\ell, \quad (V_{\ell+1}, S_{\ell+1}) := \text{qr}(V_\ell S_\ell - h_\ell \hat{P}_2).$$

- 8: **end if**
- 9: **end for**

The convergence statements for this method parallel those in Theorem 3.1. The difference is that they relate to the projected (Riemannian) gradients $\mathcal{P}_{U,V}(\nabla f(X))$ instead of $\nabla g(L, R)$, and only depend on the Lipschitz constant of $\nabla f(X)$.

Theorem 3.4 *Assume $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is λ -smooth and bounded from below. Depending on the chosen step size in the ℓ th step, the iterates $X_\ell = U_\ell S_\ell V_\ell^T$ of Algorithm 2 satisfy*

$$f(X_\ell) - f(X_{\ell+1}) \geq \frac{\vartheta_\ell}{2\lambda} \|\mathcal{P}_{U_\ell, V_\ell}(\nabla f(X_\ell))\|_F^2 \geq \frac{\vartheta_\ell}{2\lambda} \|\text{grad } f(X_\ell)\|_F^2, \quad (3.15)$$

where:

- $\vartheta_\ell := 2\alpha_\ell(1 - \alpha_\ell)$ in case a step size $h_\ell = \frac{2\alpha_\ell}{\lambda}$ with $\alpha_\ell \in (0, 1)$ is used;
- $\vartheta_\ell := \frac{1}{2}$ in case the step size is obtained by exact line search (if well defined);
- $\vartheta_\ell := \min \{2\beta\gamma(1 - \gamma), \lambda\gamma\bar{h}\}$ in case the step size is obtained from backtracking $h_\ell = \bar{h}, \bar{h}\beta, \bar{h}\beta^2, \dots$ until the Armijo condition

$$f(X_\ell) - f(X_{\ell+1}) \geq \gamma h_\ell \|P_{1/2}\|_F^2$$

is fulfilled (here $\beta, \gamma \in (0, 1)$ and $P_{1/2}$ denotes P_1 or P_2 depending on the chosen block). The backtracking procedure terminates with $h_\ell = \bar{h}\beta^i$ for some $i \leq \max \left\{ 0, \left\lceil \ln \left(\frac{2}{\bar{h}\lambda} (1 - \gamma) \right) / \ln \beta \right\rceil \right\}$.

Corollary 3.5 *Under the conditions of Theorem 3.4, let (X_ℓ) be the iterates of Algorithm 2 using any combination of the considered step-size rules. However, if step sizes $h_\ell = \frac{2\alpha_\ell}{\lambda}$ from the first step-size rule are used infinitely often, assume (for*

this subsequence) $\inf_{\ell} \alpha_{\ell} > 0$ and $\sup_{\ell} \alpha_{\ell} < 1$. Then the generated sequence $f(X_{\ell})$ of function values is monotonically decreasing and converges to some value $f_* \geq \inf_{X \in \mathcal{M}_{\leq k}} f(X)$. Moreover, $\mathcal{P}_{U_{\ell}, V_{\ell}}(\nabla f(X_{\ell})) \rightarrow 0$ and $\text{grad } f(X_{\ell}) \rightarrow 0$. Every accumulation point X_* of the sequence (X_{ℓ}) satisfies $f(X_*) = f_*$ and $\text{grad } f(X_*) = 0$, which means that X_* is a critical point of f on the manifold $\mathcal{M}_{k'}$ where $k' := \text{rank}(X_*) \leq k$. For every nonnegative integer j it holds that

$$\min_{0 \leq \ell \leq j} \|\text{grad } f(X_{\ell})\|_{\text{F}} \leq \min_{0 \leq \ell \leq j} \|\mathcal{P}_{U_{\ell}, V_{\ell}}(\nabla f(X_{\ell}))\|_{\text{F}} \leq \left(\frac{2\lambda}{\vartheta} \cdot \frac{f(X_0) - f_*}{j+1} \right)^{1/2},$$

where $\vartheta := \inf_{\ell} \vartheta_{\ell} > 0$. In particular, given $\varepsilon > 0$ the algorithm returns a point with $\|\text{grad } f(X_{\ell})\|_{\text{F}} \leq \varepsilon$ after at most $\lceil 2\lambda(f(X_0) - f_*)\vartheta^{-1}\varepsilon^{-2} - 1 \rceil$ iterations.

Note that $0 < \vartheta_{\ell} \leq \frac{1}{2}$ in all cases of Theorem 3.4 and hence $0 < \vartheta \leq \frac{1}{2}$ in Corollary 3.5.

Based on the considerations above, the proofs of the theorem and the corollary are analogous to the proofs of Theorem 3.1 and Corollary 3.5 and are therefore omitted. There is one important detail to note. The function $X \mapsto \|\text{grad } f(X)\|_{\text{F}}$ is not continuous but lower semicontinuous, which basically follows from the fact that the rank in a limit of matrices can only drop, and the tangent space at the limit is contained in the limit of tangent spaces. As a result, $X_{\ell} \rightarrow X_*$ and $\|\text{grad } f(X_{\ell})\|_{\text{F}} \rightarrow 0$ implies $\text{grad } f(X_*) = 0$. For more discussion on limiting behavior of tangent and normal cones for low-rank matrices, see [7, 19].

3.3 Comparison of Convergence Guarantees

We now compare the obtained convergence statements from Theorems 3.1 and 3.4. For this, we first note how the norms of the Riemannian gradient and the gradient of $g(L, R) := f(LR^{\top})$ can be compared. In a balanced representation $X = LR^{\top} = (U\Sigma^{1/2})(V\Sigma^{1/2})^{\top}$ we have, by (3.5), that

$$\begin{aligned} \|\nabla g(L, R)\|_{\text{F}}^2 &= \|\nabla f(X)V\Sigma^{1/2}\|_{\text{F}}^2 + \|\Sigma^{1/2}U^{\top}\nabla f(X)\|_{\text{F}}^2 \\ &\geq \sigma_k(X)(\|\nabla f(X)V\|_{\text{F}}^2 + \|U^{\top}\nabla f(X)\|_{\text{F}}^2) \\ &\geq \sigma_k(X)\|\text{grad } f(X)\|_{\text{F}}^2, \end{aligned} \tag{3.16}$$

where $\sigma_k(X)$ is the k th singular value of X . We also have a reverse estimate:

$$\begin{aligned} \|\nabla g(L, R)\|_{\text{F}}^2 &= \|\nabla f(X)V\Sigma^{1/2}\|_{\text{F}}^2 + \|\Sigma^{1/2}U^{\top}\nabla f(X)\|_{\text{F}}^2 \\ &\leq 2\|X\|_2\|\text{grad } f(X)\|_{\text{F}}^2. \end{aligned} \tag{3.17}$$

Let us take the Riemannian gradient as a reference for the comparison. For achieving a target accuracy

$$\|\text{grad } f(X_{\ell})\|_{\text{F}} \leq \varepsilon,$$

Corollary 3.5 asserts that Algorithm 2 needs at most

$$\ell_\varepsilon := \lceil 2\lambda(f(X_0) - f_*)\vartheta^{-1}\varepsilon^{-2} - 1 \rceil$$

iterations. By (3.17), it implies

$$\|\nabla g(L_{\ell_\varepsilon}, R_{\ell_\varepsilon})\|_F \leq \sqrt{2\|X_{\ell_\varepsilon}\|_2} \varepsilon \leq \sqrt{2\rho} \cdot \varepsilon,$$

where ρ is an upper bound for the spectral norm $\|X\|_2$ in the sublevel set $N_0 := \{X \in \mathcal{M}_{\leq k} : f(X) \leq f(L_0 R_0^\top)\}$. Note that under comparable assumptions (assuming the same lower bound $\vartheta := \inf_\ell \vartheta_\ell$ and using the same ρ), Corollary 3.2 guarantees after the same number ℓ_ε of iterations of Algorithm 1 only a slightly better accuracy

$$\|\nabla g(L_\ell, R_\ell)\|_F \leq \sqrt{\rho} \cdot \varepsilon. \quad (3.18)$$

In this sense, Algorithm 2 automatically ensures almost the same convergence guarantee for the balanced gradient $\nabla g(L, R)$ as when using Algorithm 1.

The converse is not true, at least not from the estimates. Suppose the balanced version in Algorithm 1 is used. After at most ℓ_ε steps it reaches the accuracy (3.18). By (3.16), it only implies

$$\|\text{grad } f(X_{\ell_\varepsilon})\|_F \leq \sqrt{\frac{\rho}{\sigma_k(X_{\ell_\varepsilon})}} \cdot \varepsilon.$$

This estimate is not robust to the occurrence of small singular values during the iteration. In fact, when starting close enough to a potential limit point X_* (of rank k) one obtains a prefactor roughly of the size $\sqrt{\frac{\|X_*\|_2}{\sigma_k(X_*)}} = \sqrt{\frac{\sigma_1(X_*)}{\sigma_k(X_*)}}$, so it depends on the (relative) condition number of X_* . Note again that Algorithm 2 would ensure $\|\text{grad } f(X_\ell)\|_F \leq \varepsilon$ without any condition on singular values.

In the sense of this comparison we may conclude that the Riemannian gradient $\text{grad } f(X_\ell)$ is a better target quantity, since it provides estimates for $\nabla g(L_\ell, R_\ell)$ independent of small singular values, which themselves are related to the local curvature of the rank- k manifold; see, e.g., [4, Sect. 4]. Consequently, the embedded subspace version of Algorithm 2 is a better and more stable choice for selecting alternating steepest descent directions with the Gauss–Southwell rule, since it directly aims at reducing the Riemannian gradient. The numerical experiments presented in Sect. 4 confirm this intuition.

3.4 Linear Convergence Rate

The algebraic convergence rates for the gradients obtained in Theorems 3.1 and 3.4 are deduced under minimal assumptions but are pessimistically slow. Instead, in practice one often observes a linear convergence rate. This is natural since the Gauss–Southwell approach leads to search directions that remain uniformly gradient-related. Indeed, in

the factorized version of Algorithm 1 the search directions $G_{1/2}$ satisfy the angle condition

$$\langle G_{1/2}, \nabla g(L, R) \rangle_{\mathbb{F}} \geq \frac{1}{\sqrt{2}} \|G_{1/2}\|_{\mathbb{F}} \|\nabla g(L, R)\|_{\mathbb{F}}$$

(here $G_{1/2}$ technically is of the form $(G_1, 0)$ or $(0, G_2)$) as well as the scaling condition

$$\frac{1}{\sqrt{2}} \|\nabla g(L, R)\|_{\mathbb{F}} \leq \|G_{1/2}\|_{\mathbb{F}} \leq \|\nabla g(L, R)\|_{\mathbb{F}}.$$

For the embedded Riemannian version in Algorithm 2 we have the analogous properties

$$\langle P_{1/2}, \text{grad } f(X) \rangle_{\mathbb{F}} \geq \frac{1}{\sqrt{2}} \|P_{1/2}\|_{\mathbb{F}} \|\text{grad } f(X)\|_{\mathbb{F}}$$

and

$$\frac{1}{\sqrt{2}} \|\text{grad } f(X)\|_{\mathbb{F}} \leq \|P_{1/2}\|_{\mathbb{F}} \leq \|\text{grad } f(X)\|_{\mathbb{F}}.$$

Provided the cost function is sufficiently smooth, a descent method with such strongly gradient-related search directions can be expected to be locally linearly convergent to critical points with positive-definite Hessian. However, it is not that straightforward in our setting. For the factorized version one is faced with the problem that the cost function $g(L, R) := f(LR^{\top})$ does not have isolated critical points due to the non-uniqueness of the factorization LR^{\top} . Hence the Hessian will not be positive-definite. In our version of Algorithm 1 this non-uniqueness is removed by balancing. While a local convergence analysis for Algorithm 1 could likely be carried out in the corresponding quotient structure of balanced representations, we omit it here and will focus on Algorithm 2 instead.

A framework for pointwise convergence of gradient-related descent methods on conic varieties such as $\mathcal{M}_{\leq k}$ has been developed in [22] assuming a Łojasiewicz-type inequality for the projected gradient onto tangent cones. It notably features a result for a variant of Algorithm 2 that also handles rank-deficient iterates [22, Algorithm 4 & Theorem 3.10]. The required inequality for a critical point $X_* \in \mathcal{M}_{\leq k}$ with full rank k would read

$$|f(X) - f(X_*)|^{1-\theta} \leq C \|\text{grad } f(X)\|_{\mathbb{F}}$$

for all $X \in \mathcal{M}_{\leq k}$ in a neighborhood of X_* , where $\theta \in (0, 1/2]$ and $C > 0$ are constants. It is also noted in [22] that $\theta = 1/2$ implies an asymptotic linear convergence rate, which is satisfied at smooth points when the Riemannian Hessian of f at X_* is positive-definite (this is proved in the lemma below; for the concept of the Riemannian Hessian we refer to [2, Chapter 5]). Since the algorithm considered in [22] slightly differs from Algorithm 2 and the provided convergence proof via the Łojasiewicz

inequality is somewhat very indirect, we present here a more direct local convergence analysis for the case of a positive-definite Riemannian Hessian. Since we already have sufficient-decrease properties as stated in Theorem 3.4—thanks to our restriction to λ -smooth cost functions—the local analysis turns out to be much easier than in the general case.

The following lemma restates the well-known fact that a Łojasiewicz-type inequality with $\theta = 1/2$ holds at critical points with positive-definite (Riemannian) Hessian.

Lemma 3.6 *Let \mathcal{M} be a smooth d -dimensional embedded Riemannian manifold in \mathbb{R}^N , $f: \mathcal{M} \rightarrow \mathbb{R}$ a twice continuously differentiable function, and $x_* \in \mathcal{M}$ a critical point of f (i.e., $\text{grad } f(x_*) = 0$) at which the Riemannian Hessian $H(x_*)$ is positive-definite (on $T_{x_*}\mathcal{M}$). Let $\lambda_{\min}^* > 0$ be its smallest eigenvalue. Then for any $\varepsilon > 0$ there exists a neighborhood $\mathcal{N} \subseteq \mathcal{M}$ of x_* such that*

$$0 < f(x) - f(x_*) \leq \left(\frac{1}{2\lambda_{\min}^*} + \varepsilon \right) \|\text{grad } f(x)\|^2$$

for all $x \in \mathcal{N}$ with $x \neq x_*$.

Proof We follow the reasoning in [2, Theorem 4.5.6] based on local charts. Choose a local chart $\psi: \mathbb{R}^d \rightarrow \mathcal{M}$ for x_* such that ψ is a diffeomorphism between a neighborhood of zero in \mathbb{R}^d and a neighborhood of x_* in \mathcal{M} , that is, $x_* = \psi(0)$ and $f(x) = \hat{f}(y)$ in these neighborhoods where $\hat{f} := f \circ \psi$. The function \hat{f} satisfies $\nabla \hat{f}(0) = [\psi'(0)]^\top \cdot \text{grad } f(x_*) = 0$. By replacing the diffeomorphism $\psi(y)$ with $\psi(Ty)$ if necessary, where T is a suitable invertible linear transformation, we can assume that $\psi'(0): \mathbb{R}^d \rightarrow T_{x_*}\mathcal{M}$ is an orthogonal map, that is, $[\psi'(0)]^\top = [\psi'(0)]^{-1}$. Then the Hessian of \hat{f} at zero,

$$\hat{H} := \nabla^2 \hat{f}(0) = [\psi'(0)]^\top \cdot H(x_*) \cdot [\psi'(0)],$$

has the same eigenvalues as the Riemannian Hessian $H(x_*)$. In particular it is positive-definite so that zero is an isolated local minimizer of \hat{f} . Therefore $f(x) - f(x_*) = \hat{f}(y) - \hat{f}(0) > 0$ for y close enough (but not equal) to zero. By Taylor expansion it holds that

$$\hat{f}(y) - \hat{f}(0) = \frac{1}{2}y^\top \hat{H}y + o(\|y\|^2), \tag{3.19}$$

as well as $\nabla \hat{f}(y) = \hat{H}y + o(\|y\|)$, that is,

$$y = \hat{H}^{-1} \nabla \hat{f}(y) + o(\|y\|). \tag{3.20}$$

The relation (3.20) in particular implies that for given $0 < \varepsilon' < 1$ and $\|y\|$ small enough one has

$$(1 - \varepsilon')\|y\| \leq \|\hat{H}^{-1} \nabla \hat{f}(y)\| \leq (1 + \varepsilon')\|y\|.$$

Using this when inserting (3.20) into (3.19) allows to conclude that

$$\begin{aligned}\hat{f}(y) - \hat{f}(0) &= \frac{1}{2} \nabla \hat{f}(y)^\top \hat{H}^{-1} \nabla \hat{f}(y) + o(\|\nabla \hat{f}(y)\|^2) \\ &\leq \frac{1}{2\lambda_{\min}^*} \|\nabla \hat{f}(y)\|^2 + o(\|\nabla \hat{f}(y)\|^2).\end{aligned}$$

Since $\psi'(0)$ is orthogonal, we have $\|\nabla \hat{f}(y)\| = \|[\psi'(y)]^\top \text{grad } f(x)\| \leq (1 + \varepsilon'') \|\text{grad } f(x)\|$ for any given $\varepsilon'' > 0$ provided $\|y\|$ is small enough. This implies the assertion. \square

We now state a linear convergence result for Algorithm 2 for λ -smooth functions under the assumption that one of the accumulation points has full rank k and a positive-definite Riemannian Hessian with respect to the smooth embedded submanifold

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

The result applies to any combination of step-size rules considered in Theorem 3.4, but requires additional assumptions for fixed step sizes (as before) and for exact step sizes in order to ensure suitable lower (in case of the former) and upper (for the latter) bounds on the step sizes.

Theorem 3.7 *Assume that $f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is λ -smooth and bounded from below. Let (X_ℓ) be the iterates of Algorithm 2 using any combination of the step-size rules considered in Theorem 3.4. If step sizes $h_\ell = \frac{2\alpha_\ell}{\lambda}$ are used for infinitely many iterates, assume (for this subsequence) $\inf_\ell \alpha_\ell > 0$ and $\sup_\ell \alpha_\ell < 1$, implying $\vartheta := \inf_\ell \vartheta_\ell > 0$ for the ϑ_ℓ defined in Theorem 3.4. If exact line search is used infinitely often, assume in addition that the constrained sublevel set $N_0 := \{X \in \mathcal{M}_{\leq k} : f(X) \leq f(X_0)\}$ is bounded.*

Suppose that (X_ℓ) has an accumulation point X_ of rank k such that f is twice continuously differentiable in a neighborhood of X_* and the Riemannian Hessian $H(X_*)$ of $f|_{\mathcal{M}_k}$ at X_* is positive-definite. Then X_* is the unique limit point of (X_ℓ) , an isolated local minimizer of $f|_{\mathcal{M}_k}$, and the convergence is R -linear with a rate*

$$\limsup_{\ell \rightarrow \infty} \|X_\ell - X_*\|^{1/\ell} \leq \left(1 - \vartheta \frac{\lambda_{\min}^*}{\lambda}\right)^{1/2} < 1,$$

where λ_{\min}^* is the smallest eigenvalue of $H(X_*)$ and satisfies $\lambda_{\min}^* \leq \lambda$. For the function values it holds that

$$\limsup_{\ell \rightarrow \infty} [f(X_\ell) - f(X_*)]^{1/\ell} \leq 1 - \vartheta \frac{\lambda_{\min}^*}{\lambda}.$$

Before giving the proof, we note that the asserted inequality $\lambda_{\min}^* \leq \lambda$ holds for all Riemannian manifolds that are invariant under multiplication by positive real numbers,

because the restriction of the Riemannian Hessian to one-dimensional subspaces corresponding to rays within the manifold just equals the restriction of the Euclidean Hessian to that subspace. Hence λ_{\min}^* is bounded by the largest eigenvalue of the Euclidean Hessian, which in turn is bounded by the Lipschitz constant λ of the gradient.

In fact, since in the Gauss–Southwell rule we only deal with linear subspaces within the tangent spaces of the low-rank manifold, a somewhat similar logic allows to obtain a refinement of Theorem 3.7, which applies when only Armijo backtracking or exact line search are used infinitely often. The proofs of both theorems will be given subsequently.

Theorem 3.8 *Under the conditions of Theorem 3.7 assume that, except for finitely many steps, exact line search or Armijo backtracking is used. Let $\nabla^2 f(X_*)$ be the Euclidean Hessian of f at X_* . Then*

$$\limsup_{\ell \rightarrow \infty} \|X_\ell - X_*\|^{1/\ell} \leq \left(1 - \vartheta^* \frac{\lambda_{\min}^*}{\lambda^*}\right)^{1/2} < 1,$$

where $\lambda^* \leq \lambda$ is the maximum between the operator norms (induced by the Frobenius norm) of the two projected Hessians $\mathcal{P}_{V_*} \nabla^2 f(X_*) \mathcal{P}_{V_*}$ and $\mathcal{P}_{U_*} \nabla^2 f(X_*) \mathcal{P}_{U_*}$, and

$$\vartheta^* := \min\{2\beta\gamma(1 - \gamma), \lambda^* \gamma \bar{h}\}.$$

For the function values it holds that

$$\limsup_{\ell \rightarrow \infty} [f(X_\ell) - f(X_*)]^{1/\ell} \leq 1 - \vartheta^* \frac{\lambda_{\min}^*}{\lambda^*}.$$

If only exact line search is used (up to finitely many steps), the above estimates hold with

$$\vartheta^* = \frac{1}{2}.$$

Note that λ^* can be bounded from above by more natural quantities such as the largest eigenvalue of $\mathcal{P}_{U_*, V_*} \nabla^2 f(X_*) \mathcal{P}_{U_*, V_*}$ (i.e., the restriction of $\nabla^2 f(X_*)$ to the tangent space $T_{X_*} \mathcal{M}_k$), the largest eigenvalue of the Riemannian Hessian of $f|_{\mathcal{M}_k}$ at X_* , or simply the largest eigenvalue of $\nabla^2 f(X_*)$.

Proof of Theorem 3.7. As in the proof of Lemma 3.6 we consider a local chart $X = \psi(y)$ for \mathcal{M}_k at X_* , that is, ψ is a diffeomorphism between a neighborhood of zero in \mathbb{R}^d and a neighborhood of $\psi(0) = X_*$ in \mathcal{M}_k . Let \hat{H} be the Hessian of $f \circ \psi$ at zero, which is positive-definite. The expression $\|y\|_{\hat{H}} := (\frac{1}{2} y^\top \hat{H} y)^{1/2}$ hence defines a norm on \mathbb{R}^d . We can assume that ψ is a diffeomorphism on an open ball $\|y\|_{\hat{H}} < r$ around zero. Let $\varepsilon > 0$. By the lemma, there exists $r_\varepsilon \in (0, r]$ such that

$$f(\psi(y)) - f(X_*) \leq \left(\frac{1}{2\lambda_{\min}^*} + \varepsilon\right) \|\text{grad } f(\psi(y))\|_{\mathbb{F}}^2 \text{ for all } \|y\|_{\hat{H}} < r_\varepsilon. \tag{3.21}$$

Let $\delta > 0$. Then the Taylor expansion as in (3.19) implies that there is $r_\delta \in (0, r_\varepsilon]$ such that

$$(1 - \delta)\|y\|_{\hat{H}}^2 \leq |f(\psi(y)) - f(\psi(0))| \leq (1 + \delta)\|y\|_{\hat{H}}^2 \quad \text{for all } \|y\|_{\hat{H}} < r_\delta.$$

The iterates of Algorithm 2 satisfy

$$\|X_{\ell+1} - X_\ell\|_F = h_\ell \|P_{1/2}\|_F \leq h_\ell \|\text{grad } f(X_\ell)\|_F. \quad (3.22)$$

Hence $X_{\ell+1} - X_\ell \rightarrow 0$ since $\text{grad } f(X_\ell) \rightarrow 0$ by Corollary 3.5 and the h_ℓ are bounded (for the exact line search one uses here that N_0 is bounded). This implies that there exist $\ell_0 \in \mathbb{N}$ and $r'_\delta \in (0, r_\delta]$ such that for all $\ell \geq \ell_0$ the property $X_\ell = \psi(y_\ell)$ with $\|y_\ell\|_{\hat{H}} < r'_\delta$ implies $X_{\ell+1} = \psi(y_{\ell+1})$ with $\|y_{\ell+1}\|_{\hat{H}} < r_\delta$, and hence both inequalities

$$f(X_\ell) - f(X_*) \leq (1 + \delta)\|y_\ell\|_{\hat{H}}^2 \quad (3.23)$$

and

$$f(X_{\ell+1}) - f(X_*) \geq (1 - \delta)\|y_{\ell+1}\|_{\hat{H}}^2 \quad (3.24)$$

are applicable, where we use $f(X_\ell) \geq f(X_{\ell+1}) \geq f(X_*)$ by construction. Since X_* is an accumulation point, such ℓ exist. Consider one such ℓ . By (3.21), it satisfies

$$f(X_\ell) - f(X_*) \leq \left(\frac{1}{2\lambda_{\min}^*} + \varepsilon \right) \|\text{grad } f(X_\ell)\|^2. \quad (3.25)$$

Combining this inequality with the main estimates (3.15) in Theorem 3.4 results in

$$\begin{aligned} f(X_{\ell+1}) - f(X_*) &\leq f(X_\ell) - f(X_*) - \frac{\vartheta_\ell}{2\lambda} \|\text{grad } f(X_\ell)\|^2 \\ &\leq \left(1 - \frac{\vartheta_\ell}{\frac{\lambda}{\lambda_{\min}^*} + 2\lambda\varepsilon} \right) (f(X_\ell) - f(X_*)). \end{aligned}$$

From (3.23) and (3.24) we conclude

$$\|y_{\ell+1}\|_{\hat{H}}^2 \leq \left(\frac{1 + \delta}{1 - \delta} \right) \left(1 - \frac{\vartheta_\ell}{\frac{\lambda}{\lambda_{\min}^*} + 2\lambda\varepsilon} \right) \|y_\ell\|_{\hat{H}}^2.$$

Since $\vartheta := \inf_\ell \vartheta_\ell > 0$, we can now assume ℓ_0 is large enough and δ is small enough such that the product of both parentheses is bounded by some value $q < 1$ for all $\ell \geq \ell_0$. For the particular ℓ under consideration it then follows $\|y_{\ell+1}\|_{\hat{H}} \leq \sqrt{q}\|y_\ell\|_{\hat{H}} < r'_\delta$, so that the argument can be inductively repeated. This implies $y_\ell \rightarrow 0$ and hence $X_\ell \rightarrow X_*$. The estimates above then imply

$$\limsup_{\ell \rightarrow \infty} (f(X_\ell) - f(X_*))^{1/\ell} \leq 1 - \frac{\vartheta}{\frac{\lambda}{\lambda_{\min}^*} + 2\lambda\varepsilon} \quad (3.26)$$

and

$$\limsup_{\ell \geq \ell_0} \|y_\ell\|^{1/\ell} \leq \sqrt{\left(\frac{1+\delta}{1-\delta}\right) \left(1 - \frac{\vartheta}{\frac{\lambda}{\lambda_{\min}^*} + 2\lambda\varepsilon}\right)}. \tag{3.27}$$

Finally, note that

$$\limsup_{\ell \rightarrow \infty} \|X_\ell - X_*\|_F^{1/\ell} = \limsup_{\ell \geq \ell_0} \|\psi(y_\ell) - \psi(0)\|_F^{1/\ell} \leq \limsup_{\ell \rightarrow \infty} \|y_\ell\|_F^{1/\ell}$$

since ψ is locally Lipschitz continuous. As δ and ε can be taken arbitrarily small, this proves the theorem. □

Proof of Theorem 3.8. We have already shown that X_ℓ converges to X_* . Moreover, $\text{grad } f(X_\ell) \rightarrow 0$ by Theorem 3.4. We can now refer to the remark at the end of Sect. 2: assuming $X_{\ell+1}$ is obtained from X_ℓ using the Armijo backtracking, then in the estimate in (3.15) we can replace λ with a restricted smoothness constant λ_ℓ of the projected gradient on the line segment connecting X_ℓ and $X_\ell - \bar{h}P_{1/2,\ell}$ as in (2.7). For concreteness, assume that $P_{1,\ell}$ is used. Let ϕ be the restriction of f to T_{V_ℓ} . Then $\nabla\phi(X) = \mathcal{P}_{V_\ell}\nabla f(X)$ and $\nabla^2\phi(X) = \mathcal{P}_{V_\ell}\nabla^2 f(X)\mathcal{P}_{V_\ell}$. Let $0 \leq t \leq 1$. Using a mean value inequality for the vector-valued function $\xi \mapsto \nabla\phi(X_\ell) - \nabla\phi(X_\ell - \xi t\bar{h}P_{1,\ell})$, $\xi \in [0, 1]$, gives that

$$\begin{aligned} \|\nabla\phi(X_\ell) - \nabla\phi(X_\ell - t\bar{h}P_{1,\ell})\|_F &\leq \max_{\xi \in [0,1]} \|\nabla^2\phi(X_\ell - \xi t\bar{h}P_{1,\ell}) \cdot t\bar{h}P_{1,\ell}\|_F \\ &\leq \max_{\xi \in [0,1]} \|\nabla^2\phi(X_\ell - \xi\bar{h}P_{1,\ell})\|_{F \rightarrow F} \cdot t\bar{h}\|P_{1,\ell}\|_F. \end{aligned}$$

Here, $\|\cdot\|_{F \rightarrow F}$ is the operator norm induced by the Frobenius norm. Note that since $X_\ell \rightarrow X_*$ and $P_{1,\ell} \rightarrow 0$, the function ϕ is indeed differentiable on the line segment in question for ℓ large enough by assumption. Referring to (2.7), it means that in that step we can replace λ in (3.15) with

$$\lambda_\ell := \max_{\xi \in [0,1]} \|\nabla^2\phi(X_\ell - \xi\bar{h}P_{1,\ell})\|_{F \rightarrow F} = \max_{\xi \in [0,1]} \|\mathcal{P}_{V_\ell}\nabla^2 f(X_\ell - \xi\bar{h}P_{1,\ell})\mathcal{P}_{V_\ell}\|_{F \rightarrow F}.$$

If $P_{2,\ell}$ is used, we take

$$\lambda_\ell := \max_{\xi \in [0,1]} \|\mathcal{P}_{U_\ell}\nabla^2 f(X_\ell - \xi\bar{h}P_{2,\ell})\mathcal{P}_{U_\ell}\|_{F \rightarrow F}$$

instead. Let $\eta > 0$. By continuity, we conclude from $X_\ell \rightarrow X_*$ and $P_{1/2,\ell} \rightarrow 0$ that for ℓ large enough and in case the Armijo step size is used we can replace (3.15) with

$$f(X_\ell) - f(X_{\ell+1}) \geq \frac{\vartheta_\eta}{2\lambda_\eta} \|\text{grad } f(X_\ell)\|_F^2, \tag{3.28}$$

where

$$\lambda_\eta := (1 + \eta) \max \left\{ \|\mathcal{P}_{V_*}\nabla^2 f(X_*)\mathcal{P}_{V_*}\|_{F \rightarrow F}, \|\mathcal{P}_{U_*}\nabla^2 f(X_*)\mathcal{P}_{U_*}\|_{F \rightarrow F} \right\} = (1 + \eta)\lambda^*$$

and

$$\vartheta_\eta := \min \{2\beta\gamma(1 - \gamma), \lambda_\eta\gamma\bar{h}\}.$$

Obviously, for exact line search the same estimate (3.28) must then also hold. If it is combined with the gradient inequality (3.25) as in the proof of Theorem 3.7, one arrives at estimates like (3.26) and (3.27), but with λ replaced with λ_η and ϑ replaced with ϑ_η , which leads to the improved rates, since η can be taken arbitrarily small.

If only exact line search is used for ℓ large enough, then the Armijo parameters in the above analysis are only auxiliary quantities. We can then optimize ϑ_η by choosing $\gamma := \frac{1}{2}$, β arbitrarily close to one, and \bar{h} large enough, so that ϑ_η is arbitrarily close to $\frac{1}{2}$. \square

3.5 Alternating Least Squares

Interestingly, our results on the gradient-related algorithms based on Gauss–Southwell rule allow to make some conclusions on another popular algorithm used in low-rank optimization, the *alternating least squares* (ALS) algorithm, in the context of λ -smooth functions. Ignoring details on implementation, this algorithm operates as follows.

Algorithm 3 Alternating least squares (ALS)

Input: Low-rank decomposition $X_0 = U_0 V_0^\top$ with $U_0 \in \mathbb{R}^{m \times k}$, $V_0 \in \mathbb{R}^{n \times k}$

- 1: **for** $\ell = 0, 1, 2, \dots$ **do**
- 2: Replace V_ℓ with a QR-factor: $V_\ell \leftarrow \text{qf}(V_\ell)$.
- 3: Find

$$U_{\ell+1} \in \underset{U \in \mathbb{R}^{m \times k}}{\text{argmin}} f(UV_\ell^\top).$$

- 4: Replace $U_{\ell+1}$ with a QR-factor: $U_{\ell+1} \leftarrow \text{qf}(U_{\ell+1})$.
- 5: Find

$$V_{\ell+1} \in \underset{V \in \mathbb{R}^{n \times k}}{\text{argmin}} f(U_{\ell+1}V^\top).$$

6: **end for**

The name alternating least squares stems from the fact that this algorithm is usually applied to convex quadratic cost functions f such as L_2 -losses, which makes it possible to solve the minimization problems in the substeps via least squares problems. Note that the orthogonalization steps in lines 2 and 4 are not needed from a theoretical perspective, but ensure the numerical stability of the subproblems in practice.

Let $X_\ell = U_\ell V_\ell^\top$. There are two simple but crucial observations for relating the ALS algorithm to our previous results. The first is that by construction it trivially follows a Gauss–Southwell rule for selecting the next block variable to be updated. In fact, before the first half-step, which finds an update $X_{\ell+1/2} = U_{\ell+1} V_\ell^\top$ in the subspace T_{V_ℓ} , the projected gradient $P_{2,\ell} = \mathcal{P}_{U_\ell} \nabla f(X_\ell)$ onto T_{U_ℓ} is just zero (for

$\ell \geq 1$), since by the result of the previous step X_ℓ is the global minimizer on T_{U_ℓ} . Hence the Gauss–Southwell rule (trivially) also would select the subspace T_{V_ℓ} for the next update, and a corresponding gradient step in T_{V_ℓ} would satisfy the estimates in (3.15), but even without the factor $1/2$ since $\|P_{2,\ell}\|_F = \|\text{grad } f(X_\ell)\|_F$.

The second observation is that the decrease in function value for the ALS half-step is better than in any of the gradient methods, since $X_{\ell+1/2}$ is the global minimizer on T_{V_ℓ} . Hence the best of the estimates (3.15), which is the one with the exact line search, is valid for the ALS half-step. Analogous considerations apply to the second half-step. Relabelling the iterates of Algorithm 3 such that

$$\tilde{X}_{2\ell} = X_\ell, \quad \tilde{X}_{2\ell+1} = X_{\ell+1/2},$$

we obtain from these considerations that

$$f(\tilde{X}_\ell) - f(\tilde{X}_{\ell+1}) \geq \frac{1}{2\lambda} \|\text{grad } f(\tilde{X}_\ell)\|_F^2, \tag{3.29}$$

that is, according to the previous notation we have $\vartheta_\ell = \vartheta = 1$ for all half-steps. In a sense, this approach of obtaining estimates by comparison to a gradient method can be seen as an instance of the accelerated line-search methods as considered in [1].

The estimate of the form (3.29) is the basis for the statements on the gradient convergence such as in Corollary 3.5. This allows to formulate the corresponding conclusion for the sequence \tilde{X}_ℓ , essentially by using $\vartheta = 1$ in Corollary 3.5. If we switch to the original ALS sequence $X_\ell = \tilde{X}_{2\ell}$ of Algorithm 3 consisting of two half-steps, the following result is obtained.

Corollary 3.9 *Assume that $f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is λ -smooth and bounded from below. Let $(X_\ell) = (U_\ell V_\ell^\top)$ be the iterates of Algorithm 3. Then the generated sequence $f(X_\ell)$ of function values is monotonically decreasing and converges to some value $f_* \geq \inf_{X \in \mathcal{M}_{\leq k}} f(X)$. Moreover, $\mathcal{P}_{U_\ell, V_\ell}(\nabla f(X_\ell)) \rightarrow 0$ and $\text{grad } f(X_\ell) \rightarrow 0$. Every accumulation point X_* of the sequence (X_ℓ) satisfies $f(X_*) = f_*$ and $\text{grad } f(X_*) = 0$, which means that X_* is a critical point of f on the manifold $\mathcal{M}_{k'}$ where $k' := \text{rank}(X_*) \leq k$. For every nonnegative integer j it holds that*

$$\min\{\|\text{grad } f(X_\ell)\|_F, \|\text{grad } f(X_{\ell+1/2})\|_F : 0 \leq \ell \leq j\} \leq \left(2\lambda \cdot \frac{f(X_0) - f_*}{2j + 1}\right)^{1/2}.$$

In particular, given $\varepsilon > 0$ the algorithm returns a point satisfying $\|\text{grad } f(X_\ell)\|_F \leq \varepsilon$ or $\|\text{grad } f(X_{\ell+1/2})\|_F \leq \varepsilon$ after at most $\lceil \lambda(f(X_0) - f_)\varepsilon^{-2} - \frac{1}{2} \rceil$ iterations.*

It is interesting to note that (3.29) can be applied in a slightly different way, as it implies the two separate estimates

$$f(X_\ell) - f(X_{\ell+1}) \geq f(X_{\ell+1/2}) - f(X_{\ell+1}) \geq \|\text{grad } f(X_{\ell+1/2})\|_F^2$$

and

$$f(X_\ell) - f(X_{\ell+1}) \geq f(X_\ell) - f(X_{\ell+1/2}) \geq \|\text{grad } f(X_\ell)\|_F^2,$$

which under the assumptions of Corollary 3.9 lead to the statement

$$\min_{0 \leq \ell \leq j} \max \left\{ \|\text{grad } f(X_\ell)\|_F, \|\text{grad } f(X_{\ell+1/2})\|_F \right\} \leq \left(2\lambda \cdot \frac{f(X_0) - f_*}{j+1} \right)^{1/2}.$$

In this form the result is quite analogous to [3, Theorem 3.2], which (as a special case) provides global rates for the (Euclidean) gradients in the alternating minimization method for non-convex functions with two block variables. In our case, this would be the function $g(L, R)$, however the assumptions in [3] require that the partial gradients (or at least one of them) satisfy uniform Lipschitz conditions, which by (3.8) and (3.9) cannot be guaranteed due to the non-uniqueness of the LR factorization. We fixed this by replacing the partial gradients with their scaled or Riemannian counterparts. Apart from that, the reasoning of [3] is actually similar to ours, as it is also based on comparing the single steps in alternating minimization to partial gradient descent. Finally note that as in Sect. 3.2 the Riemannian gradients in our results for ALS could in fact be replaced by the projection of $\nabla f(X_\ell)$ onto T_{U_ℓ, V_ℓ} , which in theory provides slightly more precise information in case of rank-deficient iterates or accumulation points.

The next natural step is to apply the local linear convergence analysis as conducted for Theorems 3.7 and 3.8 to ALS. Here one principal difficulty is to ensure $X_{\ell+1} - X_\ell \rightarrow 0$, which is a crucial ingredient in the proof of Theorem 3.7. For the gradient-related search directions as considered there this follows from $\text{grad } f(X_\ell) \rightarrow 0$ together with (3.22), but for ALS it seems to require additional assumptions (or modifications of the algorithm). One such assumption is that the cost function f is strongly convex, that is, there exists $\mu > 0$ such that

$$f(Y) \geq f(X) + \langle \nabla f(X), Y - X \rangle_F + \frac{\mu}{2} \|Y - X\|_F^2$$

for all $X, Y \in \mathbb{R}^{m \times n}$. This still covers a large field of applications.

Theorem 3.10 *Assume that $f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is λ -smooth and strongly convex. Let $(X_\ell) = (U_\ell V_\ell^\top)$ be the iterates of Algorithm 3. Assume that (X_ℓ) has an accumulation point X_* with full rank k such that f is twice continuously differentiable in a neighborhood of X_* and the Riemannian Hessian $H(X_*)$ of $f|_{\mathcal{M}_k}$ at X_* is positive-definite. Then X_* is the unique limit point of (X_ℓ) , an isolated local minimizer of $f|_{\mathcal{M}_k}$, and the convergence is R -linear with a rate*

$$\limsup_{\ell \rightarrow \infty} \|X_\ell - X_*\|^{1/\ell} \leq 1 - \frac{\lambda_{\min}^*}{\lambda^*} < 1,$$

where λ_{\min}^* is the smallest eigenvalue of $H(X_*)$ and λ^* is the maximum between the operator norms (induced by the Frobenius norm) of the two projected Hessians $\mathcal{P}_{V_*} \nabla^2 f(X_*) \mathcal{P}_{V_*}$ and $\mathcal{P}_{U_*} \nabla^2 f(X_*) \mathcal{P}_{U_*}$. For the function values it holds that

$$\limsup_{\ell \rightarrow \infty} [f(X_\ell) - f(X_*)]^{1/\ell} \leq \left(1 - \frac{\lambda_{\min}^*}{\lambda^*} \right)^2.$$

We note again that λ^* can be bounded from above by the largest eigenvalue of $\mathcal{P}_{U_*, v_*} \nabla f(X_*) \mathcal{P}_{U_*, v_*}$, the largest eigenvalue of the Riemannian Hessian of $f|_{\mathcal{M}_k}$ at X_* , or simply the largest eigenvalue of $\nabla^2 f(X_*)$.

Proof Instead of (X_ℓ) we consider as above the sequence (\tilde{X}_ℓ) of half-steps. It satisfies

$$\begin{aligned} f(\tilde{X}_\ell) - f(\tilde{X}_{\ell+1}) &\geq \langle \nabla f(\tilde{X}_{\ell+1}), \tilde{X}_\ell - \tilde{X}_{\ell+1} \rangle_F + \frac{\mu}{2} \|\tilde{X}_\ell - \tilde{X}_{\ell+1}\|_F^2 \\ &= \frac{\mu}{2} \|\tilde{X}_\ell - \tilde{X}_{\ell+1}\|_F^2, \end{aligned}$$

since by construction $\tilde{X}_{\ell+1}$ minimizes f on a linear subspace that contains $\tilde{X}_\ell - \tilde{X}_{\ell+1}$ (hence $\langle \nabla f(\tilde{X}_{\ell+1}), \tilde{X}_\ell - \tilde{X}_{\ell+1} \rangle_F = 0$). The sequence $f(\tilde{X}_\ell)$ is monotonically decreasing and converges to some f_* by Corollary 3.9. In particular, $f(\tilde{X}_\ell) - f(\tilde{X}_{\ell+1}) \rightarrow 0$, which by the above inequality implies $\|\tilde{X}_\ell - \tilde{X}_{\ell+1}\|_F \rightarrow 0$. From here on one can first repeat the arguments in the proof of Theorem 3.7 for the sequence (\tilde{X}_ℓ) (one has $\vartheta_\ell = 1$ by (3.29)), which in particular prove that X_* is the unique limit point of the sequence.

One can then proceed precisely as in the proof of Theorem 3.8 for obtaining the refined estimate

$$f(\tilde{X}_\ell) - f(\tilde{X}_{\ell+1}) \geq \frac{\vartheta_\eta}{\lambda_\eta} \|\text{grad } f(\tilde{X}_\ell)\|_F^2$$

similar to (3.28) with a local smoothness constant, which holds for ℓ large enough for a projected gradient step from \tilde{X}_ℓ when using the Armijo backtracking, and hence by construction it also holds for the ALS substep as it maximizes the left-hand side in the corresponding subspace. Note that the above estimate differs from (3.28) by the factor $1/2$, which again accounts for the fact that $\|P_{1/2, \ell}\|_F = \|\text{grad } f(\tilde{X}_\ell)\|_F$ by the properties of ALS. Since λ_η is arbitrarily close to λ^* and ϑ_η can be chosen arbitrarily close to $1/2$ using suitable Armijo parameters, one obtains as in the aforementioned proofs (essentially replacing ϑ_η with $2\vartheta_\eta \approx 1$) the R-linear rates $\limsup_{\ell \rightarrow \infty} \|\tilde{X}_\ell - X_*\|^{1/\ell} \leq \left(1 - \frac{\lambda_{\min}^*}{\lambda^*}\right)^{1/2}$ and $\limsup_{\ell \rightarrow \infty} [f(\tilde{X}_\ell) - f(X_*)]^{1/\ell} \leq 1 - \frac{\lambda_{\min}^*}{\lambda^*}$. Noting that $X_\ell = \tilde{X}_{2\ell}$ gives the result. \square

4 Numerical Experiments

We present the results of some numerical experiments that highlight the differences between Algorithms 1 and 2. As objective function, we take the strongly convex quadratic function

$$f(X) := \frac{1}{2} \text{trace}(X^\top \mathcal{A}(X)) - \text{trace}(X^\top B)$$

with a prescribed condition number $\kappa_{\mathcal{A}}$. The linear, symmetric positive-definite operator $\mathcal{A}: \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$ is random with n logarithmically spaced eigenvalues between

$\lambda_{\min} = 1/\kappa_{\mathcal{A}} > 0$ and $\lambda_{\max} = 1$. The unique minimizer X_* of f is constructed randomly but with a prescribed rank k and prescribed effective condition number $\kappa_{X_*} = \sigma_1(X_*)/\sigma_k(X_*)$. Its singular values are also logarithmically spaced. The matrix B is determined from $\nabla f(X^*) = \mathcal{A}(X_*) - B = 0$. Since f is quadratic, an exact line search can be easily calculated in closed form. Other line searches will be tested and explained further below.

It has been reported in a number of works [25–27] that ALS (and ALS-like methods that solve the subproblems in steps 3 and 5 of Algorithm 3 exactly) are often not competitive compared with (Riemannian) gradient methods like Algorithms 1 and 2. A general observation is that while ALS converges in less steps than gradient methods, the computational cost for solving these subproblems is not offset when the rank is sufficiently large. A precise comparison would be very problem dependent though. We will not compare the computational costs of Algorithms 1 and 2 with that of Algorithm 3.

We first remark that Algorithms 1 and 2 have similar computational costs. Being gradient-based methods, they both require the computation of $\nabla f(X)$ for a rank- k matrix X . As other dominant cost per step, both algorithms require k matrix–vector products with ∇f and $(\nabla f)^\top$. In large-scale applications, it is important to exploit the low-rank structures¹ since otherwise these computations will dominate all others. Next, both algorithms compute two compact QR factorizations of matrices of same size. Since Algorithm 2 does not perform explicit balancing based on a small $k \times k$ SVD, it is slightly faster per step. As we will see in the numerical experiment later, Algorithm 2 also converges faster for ill-conditioned problems.

Influence of balancing and orthogonalization. We report on the behavior of Algorithms 1 and 2 for $n = 30$ and rank $k = 8$ in the above setting. Each algorithm was started in the same set of 100 random initial points X_0 . We also included the results for a variant of Algorithm 1 without balancing, where line 4 in Algorithm 1 simply reads $L_{\ell+1} = L_+$ and $R_{\ell+1} = R_+$.

In Fig. 1, we show the results for $\kappa_{\mathcal{A}} = 20$ and $\kappa_{X_*} = 5$, and with exact line search. It is immediately clear that balancing (Algorithm 1) or orthogonalization (Algorithm 2) leads to a substantial improvement in convergence speeds. In addition, Algorithm 2, which uses orthogonalization (or can be seen as a Riemannian block method), is considerably faster and has less variance than Algorithm 1, which only balances.

Influence of condition numbers κ_{X_} and $\kappa_{\mathcal{A}}$.* As next experiment, we vary the condition numbers κ_{X_*} and $\kappa_{\mathcal{A}}$. We no longer report on the version of Algorithm 1 without balancing since it is clearly inferior. The size $n = 30$ and rank $k = 8$ are the same as above, and we take again 100 random realizations.

In Fig. 2, the condition number $\kappa_{\mathcal{A}} = 20$ of f is kept the same but the effective condition number κ_{X_*} of the rank $k = 8$ minimizer X_* is varied from 1.1 to 100. We clearly see that the Riemannian Algorithm 2 is mostly unaffected by changes in κ_{X_*} , whereas the convergence for balancing (Algorithm 1) deteriorates with increasing κ_{X_*} . Quite remarkably, for $\kappa_{X_*} = 1$ (not shown) both algorithms have essentially the same convergence behavior. For $\kappa_{X_*} = 1.1$ already, Algorithm 1 is slightly slower than

¹ Consider for example the case of a gradient that has low operator rank: $\nabla f(X) = \sum_{i=1}^{\ell} A_i X B_i$. If $X = GH^\top$, an efficient evaluation of $\nabla f(X)G$ is $\sum_{i=1}^{\ell} (A_i G)(H^\top B_i G)$.

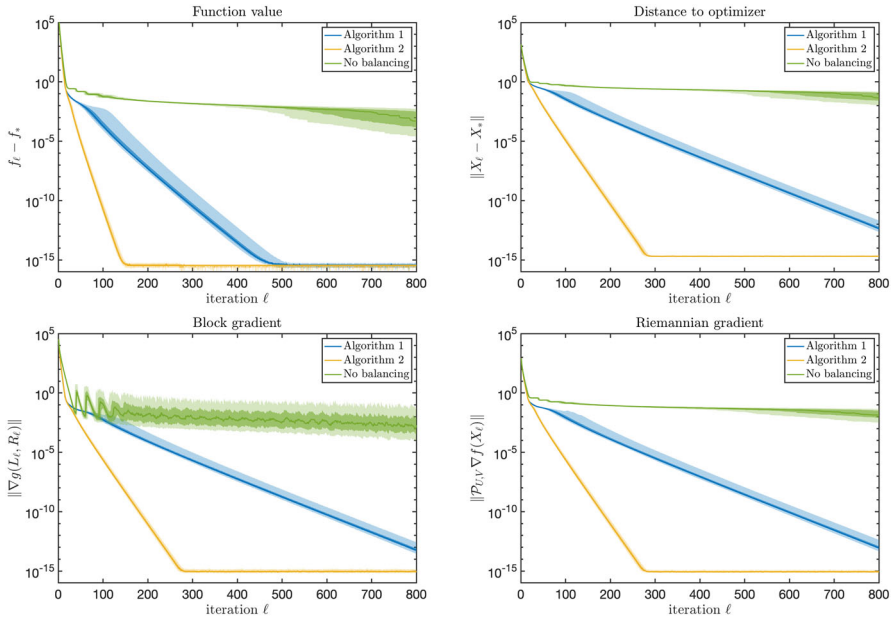


Fig. 1 Strongly convex quadratic function with condition number $\kappa_{\mathcal{A}} = 20$. The effective condition number of the minimizer X_* is $\kappa_{X_*} = 5$. The shaded areas correspond to the 50 and 90 percentiles of the relevant quantity in each panel

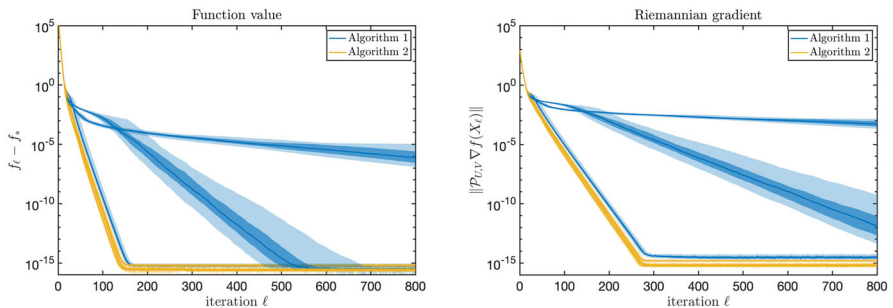


Fig. 2 Same setting as Fig. 1 but now with several values of $\kappa_{X_*} \in \{1.1, 10, 100\}$. The shaded areas correspond from left to right to increasing values of κ_{X_*}

Algorithm 2. We can thus conclude that for this problem Algorithm 1 is not robust with respect to the smallest singular values of X_* . Algorithm 2 however is robust.

As related experiment, we now keep κ_{X_*} fixed but vary the condition number $\kappa_{\mathcal{A}}$ of f . In this case, it is normal that any algorithm that uses only first-order information will converge slower with higher $\kappa_{\mathcal{A}}$. Instead of a figure, Table 1 lists the iteration numbers needed to achieve a reduction of the relative error in function value below a fixed tolerance of $\tau := 10^{-10}$. In particular, for 100 random realizations of f with

Table 1 Number of iterations ℓ_τ needed to reduce the relative error in function value to at least 10^{-10} . The first number shows the median over 100 random realizations and the second number is half the width of the 90 percentile interval

	$\kappa_{\mathcal{A}}$	1	10	1000
$\kappa_{X_*} = 1$	Alg. 1	12 ± 1.5	25 ± 4.5	98 ± 30
	Alg. 2	3 ± 0	18 ± 1	85 ± 6.5
$\kappa_{X_*} = 5$	Alg. 1	20 ± 4.5	45 ± 10	142 ± 23
	Alg. 2	3 ± 0	18 ± 1	81 ± 6

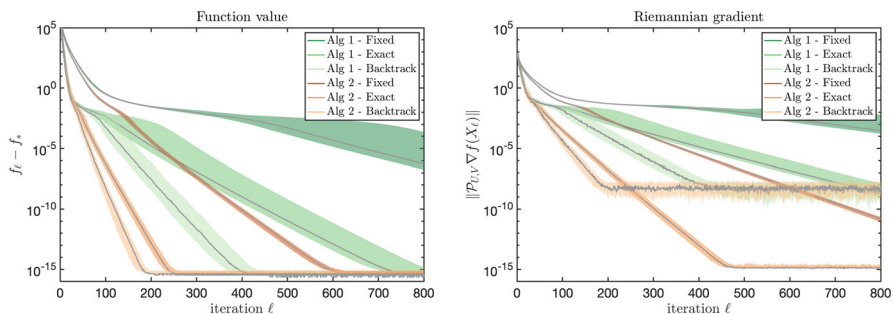


Fig. 3 Influence of the line search for a strongly convex quadratic function with fixed condition numbers $\kappa_{\mathcal{A}} = 40$ and $\kappa_{X_*} = 5$. The colored shaded areas correspond to the 90 percentile data over 20 random realizations. The grey lines indicate the median values

the same κ_{X_*} , we compute the median of

$$\ell_\tau := \min \left\{ \ell \in \mathbb{N} : \frac{f(X_\ell) - f(X_*)}{f(X_0) - f(X_*)} < \tau \right\}.$$

Influence of line search. Since the experiments from above all used exact line searches, we also tested fixed step sizes and Armijo backtracking. The experimental results are visible in Fig. 3.

For the Armijo step-size rule, we use backtracking with the following parameters² as defined in Theorems 3.1 and 3.4: initial step size $\bar{h} := 10$, step reduction factor $\beta := 0.5$, and sufficient-decrease factor $\gamma := 10^{-4}$. For the exact step size, we use $\alpha_\ell := 0.5$ in Theorems 3.1 and 3.4. In addition, for the radius ρ needed in Theorem 3.1 we take $\|X_\ell\|_2$. While this is not practical, it gives the best performance for Algorithm 1.

As expected, the fixed step-size rules lead to the slowest convergence since they are based on the worst-case global Lipschitz constant. Rather surprising, the Armijo step-size rule has the fastest convergence in terms of iterations. However, it needs on average about 10 function evaluations per iteration (but this can likely be improved with more advanced line searches based on interpolation and warm starting). On the other hand, the exact line search is only feasible for some problems like quadratic functions. Moreover, Armijo backtracking cannot decrease the Riemannian gradient much beyond the square root $\sqrt{\varepsilon_{\text{mach}}}$ of the machine precision. This is a well-known

² The experimental results were not very sensitive to the choice of these parameters as long as \bar{h} was sufficiently large.

problem due to roundoff error when testing the sufficient-decrease condition in finite precision (but it could be mitigated to some extent by the techniques from [24]). Finally, we see that for each choice of line search, Algorithm 1 is slower than Algorithm 2.

In summary, our experiments confirm several indications from the theoretical analysis that the Riemannian version of the Gauss–Southwell selection rule as realized by Algorithm 2 is superior to the factorized version based on balancing (Algorithm 1). This concerns the convergence rates as well as robustness against small singular values.

Acknowledgements We thank P.-A. Absil for valuable discussion on several parts of this work. We are also grateful to the anonymous referees for their valuable feedback. The work of G.O. was supported by the Fonds de la Recherche Scientifique – FNRS and the Fonds Wetenschappelijk Onderzoek – Vlaanderen under EOS Project no 30468160 and by the Fonds de la Recherche Scientifique – FNRS under Grant No. T.0001.23. The work of A.U. was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - Projektnummer 506561557. The work of B.V. was supported by the Swiss National Science Foundation (Grant Number 192129).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Absil, P.A., Gallivan, K.A.: Accelerated line-search and trust-region methods. *SIAM J. Numer. Anal.* **47**(2), 997–1018 (2009). <https://doi.org/10.1137/08072019X>
2. Absil, P.A., Mahony, R., Sepulchre, R.: *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton (2008). <https://doi.org/10.1515/9781400830244>
3. Beck, A.: On the convergence of alternating minimization for convex programming with applications to iteratively reweighted least squares and decomposition schemes. *SIAM J. Optim.* **25**(1), 185–209 (2015). <https://doi.org/10.1137/13094829X>
4. Feppon, F., Lermusiaux, P.F.J.: A geometric approach to dynamical model order reduction. *SIAM J. Matrix Anal. Appl.* **39**(1), 510–538 (2018). <https://doi.org/10.1137/16M1095202>
5. Harris, J.: *Algebraic Geometry: A First Course*. Springer, New York (1992). <https://doi.org/10.1007/978-1-4757-2189-8>
6. Hong, M., Wang, X., Razaviyayn, M., Luo, Z.Q.: Iteration complexity analysis of block coordinate descent methods. *Math. Program.* **163**(1–2), 85–114 (2017). <https://doi.org/10.1007/s10107-016-1057-8>
7. Hosseini, S., Luke, D.R., Uschmajew, A.: Tangent and normal cones for low-rank matrices. In: *Non-smooth Optimization and Its Applications*, pp. 45–53. Birkhäuser/Springer, Cham (2019). https://doi.org/10.1007/978-3-030-11370-4_3
8. Levin, E., Kileel, J., Boumal, N.: Finding stationary points on bounded-rank matrices: a geometric hurdle and a smooth remedy. *Math. Program.* **199**(1–2), 831–864 (2023). <https://doi.org/10.1007/s10107-022-01851-2>
9. Levin, E., Kileel, J., Boumal, N.: The effect of smooth parametrizations on nonconvex optimization landscapes. *Math. Program.* **209**(1–2), 63–111 (2025). <https://doi.org/10.1007/s10107-024-02058-3>
10. Li, Q., Zhu, Z., Tang, G.: The non-convex geometry of low-rank matrix optimization. *Inf. Inference* **8**(1), 51–96 (2019). <https://doi.org/10.1093/imaiai/iay003>
11. Li, X., Zhu, Z., So, A.M.C., Vidal, R.: Nonconvex robust low-rank matrix recovery. *SIAM J. Optim.* **30**(1), 660–686 (2020). <https://doi.org/10.1137/18M1224738>

12. Luo, Y., Li, X., Zhang, A.R.: On geometric connections of embedded and quotient geometries in Riemannian fixed-rank matrix optimization. *Math. Oper. Res.* **49**(2), 782–825 (2024). <https://doi.org/10.1287/moor.2023.1377>
13. Luo, Z.Q., Tseng, P.: On the convergence of the coordinate descent method for convex differentiable minimization. *J. Optim. Theory Appl.* **72**(1), 7–35 (1992). <https://doi.org/10.1007/BF00939948>
14. Mishra, B., Adithya Apuroop, K., Sepulchre, R.: A Riemannian geometry for low-rank matrix completion. [arXiv:1211.1550](https://arxiv.org/abs/1211.1550) (2012)
15. Nesterov, Y.: Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM J. Optim.* **22**(2), 341–362 (2012). <https://doi.org/10.1137/100802001>
16. Nesterov, Y.: *Lectures on Convex Optimization*. Springer, Cham (2018). <https://doi.org/10.1007/978-3-319-91578-4>
17. Nutini, J., Laradji, I., Schmidt, M.: Let’s make block coordinate descent converge faster: Faster greedy rules, message-passing, active-set complexity, and superlinear convergence. *J. Mach. Learn. Res.* **23**(131), 1–74 (2022). <https://jmlr.org/papers/v23/18-045.html>
18. Nutini, J., Schmidt, M., Laradji, I., Friedlander, M., Koepke, H.: Coordinate descent converges faster with the Gauss–Southwell rule than random selection. In: *Proceedings of the 32nd International Conference on Machine Learning*, pp. 1632–1641. PMLR (2015). <https://proceedings.mlr.press/v37/nutini15.html>
19. Oliker, G., Absil, P.A.: On the continuity of the tangent cone to the determinantal variety. *Set-Valued Var. Anal.* **30**(2), 769–788 (2022). <https://doi.org/10.1007/s11228-021-00618-9>
20. Oliker, G., Absil, P.A.: An apocalypse-free first-order low-rank optimization algorithm with at most one rank reduction attempt per iteration. *SIAM J. Matrix Anal. Appl.* **44**(3), 1421–1435 (2023). <https://doi.org/10.1137/22M1518256>
21. Saad, Y.: Iterative methods for linear systems of equations: a brief historical journey. In: *75 Years of Mathematics of Computation*, pp. 197–215. Amer. Math. Soc., [Providence], RI (2020). <https://doi.org/10.1090/conm/754/15141>
22. Schneider, R., Uschmajew, A.: Convergence results for projected line-search methods on varieties of low-rank matrices via Łojasiewicz inequality. *SIAM J. Optim.* **25**(1), 622–646 (2015). <https://doi.org/10.1137/140957822>
23. Sun, R., Luo, Z.Q.: Guaranteed matrix completion via non-convex factorization. *IEEE Trans. Inf. Theory* **62**(11), 6535–6579 (2016). <https://doi.org/10.1109/TIT.2016.2598574>
24. Sutti, M., Vandereycken, B.: Riemannian multigrid line search for low-rank problems. *SIAM J. Sci. Comput.* **43**(3), A1803–A1831 (2021). <https://doi.org/10.1137/20M1337430>
25. Tanner, J., Wei, K.: Low rank matrix completion by alternating steepest descent methods. *Appl. Comput. Harmon. Anal.* **40**(2), 417–429 (2016). <https://doi.org/10.1016/j.acha.2015.08.003>
26. Tong, T., Ma, C., Chi, Y.: Accelerating ill-conditioned low-rank matrix estimation via scaled gradient descent. *J. Mach. Learn. Res.* **22**(150), 1–63 (2021). <https://jmlr.org/papers/v22/20-1067.html>
27. Vandereycken, B.: Low-rank matrix completion by Riemannian optimization. *SIAM J. Optim.* **23**(2), 1214–1236 (2013). <https://doi.org/10.1137/110845768>
28. Zhu, Z., Li, Q., Tang, G., Wakin, M.B.: Global optimality in low-rank matrix optimization. *IEEE Trans. Signal Process.* **66**(13), 3614–3628 (2018). <https://doi.org/10.1109/TSP.2018.2835403>
29. Zhu, Z., Li, Q., Tang, G., Wakin, M.B.: The global optimization geometry of low-rank matrix optimization. *IEEE Trans. Inform. Theory* **67**(2), 1308–1331 (2021). <https://doi.org/10.1109/tit.2021.3049171>