

ALTERNATING LEAST SQUARES AS MOVING SUBSPACE CORRECTION*

IVAN V. OSELEDETS[†], MAXIM V. RAKHUBA[‡], AND ANDRÉ USCHMAJEV[§]

Abstract. In this note we take a new look at the local convergence of alternating optimization methods for low-rank matrices and tensors. Our abstract interpretation as sequential optimization on moving subspaces yields insightful reformulations of some known convergence conditions that focus on the interplay between the contractivity of classical multiplicative Schwarz methods with overlapping subspaces and the curvature of low-rank matrix and tensor manifolds. While the verification of the abstract conditions in concrete scenarios remains open in most cases, we are able to provide an alternative and conceptually simple derivation of the asymptotic convergence rate of the two-sided block power method of numerical algebra for computing the dominant singular subspaces of a rectangular matrix. This method is equivalent to an alternating least squares method applied to a distance function. The theoretical results are illustrated and validated by numerical experiments.

Key words. ALS, nonlinear Gauss–Seidel method, low-rank approximation, local convergence

AMS subject classifications. 15A99, 65F10, 65F30, 65K10

DOI. 10.1137/17M1148712

1. Introduction. Consider a real-valued function $F(x)$, where $x = (\xi_1, \dots, \xi_N)$ is a tuple of vectors $\xi_i \in \mathbb{R}^{n_i}$. The *alternating optimization* (AO) or *block coordinate descent* methods try to solve the problem

$$\min F(x) = \min F(\xi_1, \dots, \xi_N)$$

by alternating between updates of single (block) variables ξ_i while fixing all the other ξ_j , $j \neq i$:

$$\xi_i \leftarrow \operatorname{argmin}_{\xi \in \mathbb{R}^{n_i}} F(\xi_1, \dots, \xi_{i-1}, \xi, \xi_{i+1}, \dots, \xi_N).$$

In other words, such an update is a minimization of F on the affine linear manifold $x + T_i$ with the linear subspaces

$$(1.1) \quad T_i = \{0\} \times \dots \times \{0\} \times \mathbb{R}^{n_i} \times \{0\} \times \dots \times \{0\}.$$

If F is smooth enough, the minimization in every substep may be replaced by finding a critical point of F on $x + T_i$. The method is then also known under the name *nonlinear (block) Gauss–Seidel method*.

*Received by the editors September 21, 2017; accepted for publication (in revised form) October 1, 2018; published electronically December 11, 2018.

<http://www.siam.org/journals/sinum/56-6/M114871.html>

Funding: The work of the first and second authors was supported by the Ministry of Education and Science of the Russian Federation (grant 14.756.31.0001). The work of the second author was performed while he was a junior research scientist at Skolkovo Institute of Science and Technology.

[†]Skolkovo Institute of Science and Technology, Skolkovo Innovation Center, 143026 Moscow, Russia (i.oseldets@skoltech.ru).

[‡]Seminar for Applied Mathematics, ETH Zurich, Rämistrasse 101, 8092 Zurich, Switzerland (maksim.rakhuba@sam.math.ethz.ch).

[§]Hausdorff Center for Mathematics & Institute for Numerical Simulation, University of Bonn, 53115 Bonn, Germany. Current address: Max Planck Institute for Mathematics in the Sciences, 04103 Leipzig, Germany (uschmajew@mis.mpg.de).

Such an approach is effective if optimization on the hyperplanes $x + T_i$ is easy, for instance, because it is of lower dimension, or because F takes a simple form on it. Obviously, the hyperplanes $x + T_i$ are changing during this process, as they depend on x . The subspaces T_i , however, do not change with x . Based on this, the linearization of this method around a critical point of F corresponds to a classical Gauss–Seidel (successive relaxation) method applied to a quadratic model of F and, hence, its local convergence to a critical point can be shown under suitable assumptions on the Hessian in the critical point; see [12] or section 2.2.1.

There are too many areas of application of AO to mention here. In this paper, we wish to focus on multilinear optimization. This includes low-rank matrix and tensor approximation. Here the scenario is slightly more structured. Let us explain this using the example of low-rank matrix optimization. Assume we are given a function $f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ on the space of real $m \times n$ matrices, and we wish to minimize it subject to the constraint $\text{rank}(X) \leq k$. Then it is natural to use the parametrization $X = UV^T$ with $U \in \mathbb{R}^{m \times k}$, $V \in \mathbb{R}^{n \times k}$, and attempt solving

$$\min F(U, V) := f(UV^T)$$

via AO between U and V :

$$(1.2) \quad U \leftarrow \underset{\hat{U} \in \mathbb{R}^{m \times k}}{\text{argmin}} f(\hat{U}V^T), \quad V \leftarrow \underset{\hat{V} \in \mathbb{R}^{n \times k}}{\text{argmin}} f(U\hat{V}^T).$$

The easiest example to consider is the Euclidean distance function

$$f(X) = \frac{1}{2} \|X - B\|_F^2 = \frac{1}{2} \sum_{i,j} (x_{ij} - b_{ij})^2$$

to a given matrix B . In this case the AO strategy reads

$$(1.3) \quad U \leftarrow \underset{\hat{U} \in \mathbb{R}^{m \times k}}{\text{argmin}} \frac{1}{2} \|\hat{U}V^T - B\|_F^2, \quad V \leftarrow \underset{\hat{V} \in \mathbb{R}^{n \times k}}{\text{argmin}} \frac{1}{2} \|U\hat{V}^T - B\|_F^2,$$

and is called the *alternating least squares* algorithm. It will be discussed in detail in sections 3.3 and 3.2.

An alternative viewpoint, however, which is the starting point for the present work, is that in terms of the initial function f , the AO procedure (1.2) amounts to a sequence of optimization problems

$$(1.4) \quad X \leftarrow \underset{X \in T_1(X)}{\text{argmin}} f(X), \quad X \leftarrow \underset{X \in T_2(X)}{\text{argmin}} f(X),$$

on *varying linear subspaces*

$$(1.5) \quad T_1(X) = \{Y \in \mathbb{R}^{m \times n} : \text{row}(Y) \subseteq \text{row}(X)\},$$

respectively,

$$(1.6) \quad T_2(X) = \{Y \in \mathbb{R}^{m \times n} : \text{col}(Y) \subseteq \text{col}(X)\}.$$

Here row and col denote the row and column space of a matrix. To be precise, one should emphasize that the update rules (1.2) and (1.4) are only equivalent as long as all constructed matrices retain full possible rank k . Also note that $X \in T_1(X)$ and

$X \in T_2(X)$, hence, we can formally see (1.4) as minimizations on affine subspaces $X + T_1(X)$ and $X + T_2(X)$ as for the classical AO method.

The point we wish to make is that the formulation (1.4) is a more appropriate viewpoint on the AO method (1.2), since it is intrinsically invariant under different choices of U and V in the bilinear parametrization $X = UV^T$, which, at least formally, is highly nonunique. To see this more clearly, let us compare the two following pseudocodes:

<hr/> <p>Algorithm 1: Low-rank AO, vanilla.</p> <hr/> <p>Input: $U_0 \in \mathbb{R}^{m \times k}$, $V_0 \in \mathbb{R}^{n \times k}$.</p> <p>for $\ell = 0, 1, 2, \dots$ do</p> <div style="display: inline-block; vertical-align: middle;"> $U_{\ell+1} := \operatorname{argmin}_{\hat{U} \in \mathbb{R}^{m \times k}} f(\hat{U}V_\ell^T)$ $V_{\ell+1} := \operatorname{argmin}_{\hat{V} \in \mathbb{R}^{n \times k}} f(U_{\ell+1}\hat{V}^T)$ </div> <p>end</p> <hr/>	<hr/> <p>Algorithm 2: Low-rank AO with QR.</p> <hr/> <p>Input: $U_0 \in \mathbb{R}^{m \times k}$, $V_0 \in \mathbb{R}^{n \times k}$.</p> <p>for $\ell = 0, 1, 2, \dots$ do</p> <div style="display: inline-block; vertical-align: middle;"> $U \leftarrow \operatorname{argmin}_{\hat{U} \in \mathbb{R}^{m \times k}} f(\hat{U}V_\ell^T), \quad U = Q_1 R_1$ $V \leftarrow \operatorname{argmin}_{\hat{V} \in \mathbb{R}^{n \times k}} f(Q_1 \hat{V}^T), \quad V = Q_2 R_2$ $U_{\ell+1} := Q_1 R_1^T, \quad V_{\ell+1} := Q_2$ </div> <p>end</p> <hr/>
--	---

The algorithm on the right uses QR decompositions of the factors U and V in order to keep the low-rank representations stable, which is generally advised in practice as it keeps the argmin problems better conditioned. This is easily seen for the least squares problems (1.3) that arise for the squared Frobenius distance: if, e.g., the output matrix V of a previous step is badly conditioned, then the linear operator $U \mapsto UV^T$ is comparably badly conditioned, and the exact solution $U = BV(V^T V)^{-1}$ of the next least squares problem $\min_U \|UV^T - B\|_F$ may be difficult to compute accurately since the matrix $V^T V$ needs to be inverted (assuming it is invertible). If, on the other hand, $V = QR$ is first replaced by its q-factor, $\hat{V} = Q$, then the next update is just $\hat{U} = B\hat{V}$. However, it is easy to check that $U = BQR^{-T}$ and, therefore, $UV^T = BQQ^T = \hat{U}\hat{V}^T$. In other words, in both cases the same matrix is computed, but the second strategy does not require matrix inversion, just the computation of a QR decomposition.

At first, Algorithm 2 appears considerably harder to analyze than Algorithm 1, which is a plain AO method. A closer inspection, however, reveals that this is not true in the case that the solutions to the minimizations problems are unique (for instance, if all matrices retain rank k and f is strictly convex), since then in both algorithms the same sequences of low-rank matrices $X_\ell = U_\ell V_\ell^T$ are constructed when starting from the same initialization. The underlying reason is that replacing U by its QR-factor Q does not change the column space, and replacing V by its QR-factor does not change the row space. Hence the subspaces of $\mathbb{R}^{m \times n}$ over which the argmins are taken are the same in both algorithms. The superiority of the “subspace viewpoint” compared to the “representation viewpoint” lies in realizing this theoretical equivalence of both algorithms, although numerically they may still behave quite differently.

The above example of low-rank matrix optimization via AO generalizes to the scenario where we are given a multilinear map

$$\tau: V_1 \times \dots \times V_d \rightarrow \mathbf{V}$$

mapping from d linear spaces V_1, \dots, V_d to a space \mathbf{V} , and wish to optimize a function

$$(1.7) \quad F(\xi_1, \dots, \xi_d) = f(\tau(\xi_1, \dots, \xi_d)).$$

For instance the tasks of computing approximations to tensors in low-rank canonical polyadic (CP), tensor train, or (hierarchical) Tucker formats are of this type; see [16, 13, 11, 5].

The aim of this paper is to subsume previous local convergence analysis of AO for multilinear optimization [16, 13] (see section 5 for an overview on related work) into a transparent theorem that reduces to the subspaces correction method for the linearized problem at a fixed point. Furthermore, in section 3 we apply our framework to derive in a new way the (known) convergence rate of a two-sided block power method for computing the dominant k -dimensional singular subspaces of a matrix, by relating this power method to AO for the distance function.

Unfortunately, our techniques are currently not yet in a shape that would allow for substantially new insights in the analysis of alternating least squares for low-rank approximation of higher-order tensors like, e.g., for local convergence analysis of the *higher-order power method* (AO for best rank-one approximation). One reason is that we are lacking an analogous statement to Theorem 2.5 that relates the local contractivity to the spectral radius of a single product of operators. Yet we believe our general setup of AO with moving subspaces will be useful for the tensor case as well in the future. Some references to known results on AO for tensors are given at the end of the paper.

2. Abstract setup. To generalize our two motivating examples, we consider a C^1 function $f: \mathbf{V} \rightarrow \mathbf{V}$ on a Hilbert space \mathbf{V} . To every $x \in \mathbf{V}$ we attach a closed subspace $T(x)$ of \mathbf{V} . Further, we assume that we are given a possibly overlapping partition

$$T(x) = T_1(x) + \cdots + T_d(x)$$

into d closed subspaces $T_i(x)$. Then we define d maps

$$\mathbf{P}_i: \mathbf{V} \rightarrow \mathcal{L}(\mathbf{V}), \quad i = 1, \dots, d,$$

such that for every $x \in \mathbf{V}$ the linear operator $\mathbf{P}_i(x)$ is the orthogonal projection onto the space $T_i(x)$. Correspondingly, we let $\mathbf{P}(x)$ be the orthogonal projector on $T(x)$.

Next, let \mathbf{S}_i , $i = 1, \dots, d$, be (nonlinear) operators on \mathbf{V} such that $y = \mathbf{S}_i(x)$ satisfies

$$(2.1) \quad y \in x + T_i(x), \quad \mathbf{P}_i(x)\nabla f(y) = 0.$$

It means that \mathbf{S}_i maps x to a relative critical point of f on the hyperplane $x + T_i(x)$. If, for instance, f is strictly convex and coercive, then such an operator \mathbf{S}_i is uniquely defined and corresponds to minimizing f on $x + T_i(x)$.

AO on moving hyperplanes corresponds to an iteration of the form

$$(2.2) \quad x_{\ell+1} = \mathbf{S}(x_\ell) := (\mathbf{S}_d \circ \cdots \circ \mathbf{S}_1)(x_\ell).$$

In the following we consider points $\bar{x} \in \mathbf{V}$ which are fixed points of every \mathbf{S}_i , that is,

$$\bar{x} = \mathbf{S}_i(\bar{x}), \quad i = 1, \dots, d.$$

Then \bar{x} is obviously a fixed point of \mathbf{S} . We further note that if \bar{x} is a fixed point of every \mathbf{S}_i , then $\nabla f(\bar{x})$ is orthogonal to $T(\bar{x})$. The converse is also true under mild assumptions, for instance, by ensuring that on every hyperplane $x + T_i(x)$ there exists only one point y satisfying (2.1) (e.g., f is strictly convex with bounded sublevel sets),

or, if this is not the case, by requiring that y (2.1) should be chosen as close as possible to x .

In this work, we wish to investigate the local convergence properties of the fixed point iteration (2.2) under the assumption that all \mathbf{P}_i , all \mathbf{S}_i , and also \mathbf{P} are continuously (Fréchet) differentiable mappings in a neighborhood of \bar{x} . Without going into detail, we mention that for optimization tasks in low-rank tensor formats as mentioned in the context of (1.7) such smoothness assumptions are typically ensured if the fixed point has maximal feasible rank. For instance, local analysis of AO for low-rank matrices UV^T as considered in (1.2) will require the factors U and V to have full column rank k ; cf. section 3.

The local contractivity around \bar{x} is governed by the spectral properties of the derivative $\mathbf{S}'(\bar{x})$. By the chain rule,

$$(2.3) \quad \mathbf{S}'(\bar{x}) = \prod_{i=d}^1 \mathbf{S}'_i(\bar{x}).$$

The derivatives $\mathbf{S}'_i(\bar{x})$ are computed in the next section. Some preliminary properties, however, are obtained by differentiating the equation

$$\mathbf{P}_i(x)(\mathbf{S}_i(x) - x) = \mathbf{S}_i(x) - x.$$

It gives the relation

$$\mathbf{P}'_i(x; h)(\mathbf{S}_i(x) - x) + \mathbf{P}_i(x)(\mathbf{S}'_i(x)h - h) = \mathbf{S}'_i(x)h - h$$

for all $h \in \mathbf{V}$. Here, $\mathbf{P}'_i(x; h) \in \mathcal{L}(\mathbf{V})$ denotes the application of the derivative of $\mathbf{P}_i(x)$ at x to h . Hence, in a fixed-point $\bar{x} = \mathbf{S}_i(\bar{x})$, it holds

$$(2.4) \quad \mathbf{S}'_i(\bar{x})h = \mathbf{P}_i(\bar{x})\mathbf{S}'_i(\bar{x})h + (\mathbf{I} - \mathbf{P}_i(\bar{x}))h.$$

This equation is interesting as it shows the following.

PROPOSITION 2.1. *Assume \mathbf{P}_i and \mathbf{S}_i are continuously differentiable around a fixed point $\bar{x} = \mathbf{S}(\bar{x})$.*

- (i) *The subspaces $T_i(\bar{x})$ and $T(\bar{x})$ are both invariant subspaces of $\mathbf{S}'_i(\bar{x})$.*
- (ii) *The restriction of $\mathbf{S}'_i(\bar{x})$ to the orthogonal complement $T_i(\bar{x})^\perp$ has all its singular values bounded from below by one, and equals the identity on $T_i(\bar{x})^\perp$ if and only if $T_i(\bar{x})^\perp$ is also an invariant subspace of $\mathbf{S}'_i(\bar{x})$.*
- (iii) *The subspace $T(\bar{x})$ is an invariant subspace of $\mathbf{S}'(\bar{x})$, that is, it holds*

$$\mathbf{S}'(\bar{x})\mathbf{P}(\bar{x}) = \mathbf{P}(\bar{x})\mathbf{S}'(\bar{x})\mathbf{P}(\bar{x}).$$

The restriction of $\mathbf{S}'(\bar{x})$ to the orthogonal complement $T(\bar{x})^\perp$ has all its singular values bounded from below by one, and equals the identity on $T(\bar{x})^\perp$ if and only if $T(\bar{x})^\perp$ is also an invariant subspace of $\mathbf{S}'(\bar{x})$.

Proof. Ad (i). Obviously, by (2.4), $h \in T_i(\bar{x})$ is mapped to $T_i(\bar{x})$. On the other hand, since $T_i(\bar{x})$ is a subspace of $T(\bar{x})$, the element $(\mathbf{I} - \mathbf{P}_i(\bar{x}))h$ in (2.4) belongs to $T(\bar{x})$ for every $h \in T(\bar{x})$. Hence $T(\bar{x})$ is also an invariant subspace of $\mathbf{S}'_i(\bar{x})$.

Ad (ii). Equation (2.4) shows that $\mathbf{S}'_i(\bar{x})h = h$ for all $h \in T_i(\bar{x})^\perp$ if and only if $\mathbf{P}_i(\bar{x})\mathbf{S}'_i(\bar{x})h = 0$ for such h , which is equivalent to $T_i(\bar{x})^\perp$ being an invariant subspace of $\mathbf{S}'_i(\bar{x})$. In any case, it holds, by orthogonality of both terms in (2.4), that

$\|\mathbf{S}'_i(\bar{x})h\|^2 \geq \|h\|^2$ for all $h \in T'_i(\bar{x})^\perp$, which shows that the singular values of the restriction to that space cannot be smaller than one.

Ad (iii). Since $T(\bar{x})$ is an invariant subspace of every $\mathbf{S}'_i(\bar{x})$ by (i), it follows from the chain rule (2.3) that it is also an invariant subspace of $\mathbf{S}'(\bar{x})$. Even more, an induction based on (2.4) shows that

$$(\mathbf{I} - \mathbf{P}(\bar{x}))\mathbf{S}'(\bar{x}) = (\mathbf{I} - \mathbf{P}(\bar{x}))(\mathbf{I} - \mathbf{P}_d(\bar{x})) \cdots (\mathbf{I} - \mathbf{P}_1(\bar{x})).$$

So, since the $T_i(\bar{x})$ are subspaces of $T(\bar{x})$, it holds that

$$(\mathbf{I} - \mathbf{P}(\bar{x}))\mathbf{S}'(\bar{x})(\mathbf{I} - \mathbf{P}(\bar{x})) = \mathbf{I} - \mathbf{P}(\bar{x}).$$

This implies the assertion as in (ii). □

The proposition shows that we can only hope for contractivity of the map $\mathbf{S}'(\bar{x})$ on its invariant subspace $T(\bar{x})$. Therefore, in what follows, we focus on the spectral radius of $\mathbf{S}'(\bar{x})\mathbf{P}(\bar{x}) = \mathbf{P}(\bar{x})\mathbf{S}'(\bar{x})\mathbf{P}(\bar{x})$. Concerning the convergence of our fixed point iteration $x_{\ell+1} = \mathbf{S}(x_\ell)$, the contractivity of $\mathbf{S}'(\bar{x})\mathbf{P}(\bar{x})$ turns out to be sufficient for local linear convergence in two notable cases: (i) the subspace $T(\bar{x})$ equals the whole space \mathbf{V} as is the case for classical AO with the subspaces given by (1.1); or (ii) the iterates (x_ℓ) lie on a smooth submanifold \mathcal{M} and $T(\bar{x})$ is the tangent space to that manifold at \bar{x} . This scenario is often encountered in low-rank matrix and tensor optimization via AO, when the objective function (1.7) is considered and the image of the multilinear map τ is locally a manifold. We will demonstrate this for low-rank matrix approximation in section 3.

2.1. Computation of derivatives. We recall that $f: \mathbf{V} \rightarrow \mathbb{R}$ is said to be twice continuously (Fréchet) differentiable in a neighborhood of $\bar{x} \in \mathbf{V}$, if for every x in that neighborhood there exists a bounded linear form $f'(x)$ on \mathbf{V} and a bounded bilinear form $f''(x)$ on $\mathbf{V} \times \mathbf{V}$, which both depend continuously on x , such that

$$f(x + h) = f(x) + f'(x)h + \frac{1}{2}f''(x)(h, h) + o(\|h\|^2).$$

The bilinear forms $f''(x)$ are necessarily symmetric; see, e.g., [4, section (8.12.2)]. Hence, since \mathbf{V} is a Hilbert space, there exist elements $\nabla f(x) \in \mathbf{V}$ (gradient) and unique bounded self-adjoint operators $\mathbf{A}(x)$ (Hessian) on \mathbf{V} , both depending continuously on x , such that

$$f(x + h) = f(x) + \langle \nabla f(x), h \rangle + \frac{1}{2}\langle h, \mathbf{A}(x)h \rangle + o(\|h\|^2)$$

for all x in a neighborhood of \bar{x} . Note that $x \mapsto \mathbf{A}(x)$ is the (Fréchet) derivative of the map $x \mapsto \nabla f(x)$. For brevity the following shorthand notation will be used for the rest of the paper:

$$\bar{\mathbf{P}}_i := \mathbf{P}_i(\bar{x}), \quad \bar{\mathbf{P}} := \mathbf{P}(\bar{x}), \quad \bar{\mathbf{A}} := \mathbf{A}(\bar{x}), \quad \bar{\mathbf{B}}_i := (\bar{\mathbf{P}}_i \bar{\mathbf{A}} \bar{\mathbf{P}}_i)^{-1}.$$

The inverse operator \mathbf{B}_i is obtained by considering $\bar{\mathbf{P}}_i \bar{\mathbf{A}} \bar{\mathbf{P}}_i$ as an operator on $T_i(\bar{x})$.

To obtain a formula for $\mathbf{S}'(\bar{x})$, we differentiate each \mathbf{S}_i separately. The derivatives $\mathbf{S}'_i(\bar{x})$ are given as follows.

PROPOSITION 2.2. *Assume that \mathbf{P}_i and \mathbf{S}_i are continuously differentiable in a neighborhood of a fixed point \bar{x} , and that f is twice continuously differentiable around*

\bar{x} . Then $\mathbf{P}'_i(\bar{x}; h)\nabla f(\bar{x}) \in T_i(\bar{x})$. If the linear operator $\bar{\mathbf{P}}_i\bar{\mathbf{A}}\bar{\mathbf{P}}_i$ is invertible on $T_i(\bar{x})$, then

$$(2.5) \quad \mathbf{S}'_i(\bar{x})h = h - \bar{\mathbf{B}}_i\bar{\mathbf{P}}_i\bar{\mathbf{A}}h - \bar{\mathbf{B}}_i\mathbf{P}'_i(\bar{x}; h)\nabla f(\bar{x}).$$

In particular,

$$\mathbf{S}'_i(\bar{x}) = -\bar{\mathbf{B}}_i\mathbf{P}'_i(\bar{x}; h)\nabla f(\bar{x}) \quad \text{on } T_i(\bar{x}).$$

Proof. Differentiating the equation $\mathbf{P}_i(x)\nabla f(\mathbf{S}_i(x)) = 0$ yields

$$(2.6) \quad \mathbf{P}'_i(x; h)\nabla f(\mathbf{S}_i(x)) + \mathbf{P}_i(x)\mathbf{A}(x)\mathbf{S}'_i(x)h = 0$$

for all variations $h \in \mathbf{V}$. Splitting the term of interest $\mathbf{S}'_i(x)h$ in (2.6) into its parts on $T_i(\bar{x})$ and the orthogonal complement, we get

$$\mathbf{P}_i(x)\mathbf{A}(x)\mathbf{P}_i(x)\mathbf{S}'_i(x)h = -\mathbf{P}_i(x)\mathbf{A}(x)(\mathbf{I} - \mathbf{P}_i(x))\mathbf{S}'_i(x)h - \mathbf{P}'_i(x; h)\nabla f(\mathbf{S}_i(x)).$$

At a fixed point, we can use (2.4). Therefore

$$\bar{\mathbf{P}}_i\bar{\mathbf{A}}\bar{\mathbf{P}}_i\mathbf{S}'_i(\bar{x})h = -\bar{\mathbf{P}}_i\bar{\mathbf{A}}(\mathbf{I} - \bar{\mathbf{P}}_i)h - \mathbf{P}'_i(\bar{x}; h)\nabla f(\bar{x}).$$

This equation shows that $\mathbf{P}'_i(\bar{x}; h)\nabla f(\bar{x})$ lies in $T_i(\bar{x})$. Assuming further that $\bar{\mathbf{P}}_i\bar{\mathbf{A}}\bar{\mathbf{P}}_i$ has an inverse $\bar{\mathbf{B}}_i$ on $T_i(\bar{x})$, we get

$$\bar{\mathbf{P}}_i\mathbf{S}'_i(\bar{x})h = \bar{\mathbf{P}}_i h - \bar{\mathbf{B}}_i\bar{\mathbf{P}}_i\bar{\mathbf{A}}h - \bar{\mathbf{B}}_i\mathbf{P}'_i(\bar{x}; h)\nabla f(\bar{x}).$$

Using (2.4) once more, one arrives at

$$\begin{aligned} \mathbf{S}'_i(\bar{x})h &= (\mathbf{I} - \bar{\mathbf{P}}_i)\mathbf{S}'_i(\bar{x})h + \bar{\mathbf{P}}_i\mathbf{S}'_i(\bar{x})h \\ &= (\mathbf{I} - \bar{\mathbf{P}}_i)h + \bar{\mathbf{P}}_i h - \bar{\mathbf{B}}_i\bar{\mathbf{P}}_i\bar{\mathbf{A}}h - \bar{\mathbf{B}}_i\mathbf{P}'_i(\bar{x}; h)\nabla f(\bar{x}), \end{aligned}$$

which is (2.5). □

It will be useful to simplify notation. We denote

$$(2.7) \quad \bar{\mathbf{P}}_i^{\bar{\mathbf{A}}} := \bar{\mathbf{B}}_i\bar{\mathbf{P}}_i\bar{\mathbf{A}}.$$

If $\bar{\mathbf{A}}$ is a positive definite operator, then $\bar{\mathbf{B}}_i$ is always well defined and $\bar{\mathbf{P}}_i^{\bar{\mathbf{A}}}$ allows an interpretation as the $\bar{\mathbf{A}}$ -orthogonal projection onto subspace $T_i(\bar{x})$, that is, an orthogonal projection with respect to the inner product $(x, y) \mapsto \langle x, \bar{\mathbf{A}}y \rangle$.¹

Further, we define the linear operator $\bar{\mathbf{N}}_i$ on \mathbf{V} such that

$$(2.8) \quad \bar{\mathbf{N}}_i h := \mathbf{P}'_i(\bar{x}; h)\nabla f(\bar{x})$$

for all h . With this notation, and under the assumptions of Proposition 2.2, $\mathbf{S}'_i(\bar{x})$ can be conveniently written as

$$\mathbf{S}'_i(\bar{x}) = (\mathbf{I} - \bar{\mathbf{P}}_i^{\bar{\mathbf{A}}}) - \bar{\mathbf{B}}_i\bar{\mathbf{N}}_i.$$

The formula for $\mathbf{S}'(\bar{x})$ is now obtained by the chain rule. For later reference we formulate it as a theorem.

THEOREM 2.3. *Assume that all \mathbf{P}_i and \mathbf{S}_i are continuously differentiable in a neighborhood of a fixed point \bar{x} , and that f is twice continuously differentiable around \bar{x} . Assume all $\bar{\mathbf{B}}_i = (\bar{\mathbf{P}}_i\bar{\mathbf{A}}\bar{\mathbf{P}}_i)^{-1}$ exist on $T_i(\bar{x})$. Then*

$$(2.9) \quad \mathbf{S}'(\bar{x}) = \prod_{i=d}^1 \mathbf{S}'_i(\bar{x}) = \prod_{i=d}^1 [(\mathbf{I} - \bar{\mathbf{P}}_i^{\bar{\mathbf{A}}}) - \bar{\mathbf{B}}_i\bar{\mathbf{N}}_i].$$

¹To see it, observe (we omit the subscript i) that $\bar{\mathbf{B}}\bar{\mathbf{P}} = \bar{\mathbf{P}}\bar{\mathbf{B}}\bar{\mathbf{P}}$ is self-adjoint and therefore $\langle \bar{\mathbf{P}}^{\bar{\mathbf{A}}}x, \bar{\mathbf{A}}(\mathbf{I} - \bar{\mathbf{P}}^{\bar{\mathbf{A}}})x \rangle = \langle \bar{\mathbf{B}}\bar{\mathbf{P}}\bar{\mathbf{A}}x, (\bar{\mathbf{A}} - \bar{\mathbf{A}}\bar{\mathbf{B}}\bar{\mathbf{P}}\bar{\mathbf{A}})x \rangle = \langle x, (\bar{\mathbf{A}}\bar{\mathbf{B}}\bar{\mathbf{P}}\bar{\mathbf{A}} - \bar{\mathbf{A}}\bar{\mathbf{B}}\bar{\mathbf{P}}\bar{\mathbf{A}}\bar{\mathbf{B}}\bar{\mathbf{P}}\bar{\mathbf{A}})x \rangle = 0$ for all $x \in \mathbf{V}$, since $\bar{\mathbf{P}}\bar{\mathbf{A}}\bar{\mathbf{B}}\bar{\mathbf{P}} = \bar{\mathbf{P}}\bar{\mathbf{A}}\bar{\mathbf{P}}\bar{\mathbf{B}}\bar{\mathbf{P}} = \bar{\mathbf{P}}$. Hence $\bar{\mathbf{P}}^{\bar{\mathbf{A}}}x$ is $\bar{\mathbf{A}}$ -orthogonal to $(\mathbf{I} - \bar{\mathbf{P}}^{\bar{\mathbf{A}}})x$.

2.2. Curvature free cases ($\bar{\mathbf{N}}_i = \mathbf{0}$). An easy case to investigate is when all $\bar{\mathbf{N}}_i = \mathbf{0}$, since in this case we obtain the formula

$$\mathbf{S}'(\bar{x}) = \prod_{i=d}^1 (\mathbf{I} - \bar{\mathbf{P}}_i^{\bar{\mathbf{A}}}),$$

which is well known from the theory of subspace correction methods for the solution of linear systems, specifically the *multiplicative Schwarz method*. The following statement is obtained from the standard results on the multiplicative Schwarz method (see, e.g., [20, Theorem 4.2] for the Hilbert space case), by restricting everything to the subspace $T(\bar{x})$ and considering the equivalent $\bar{\mathbf{A}}$ -inner-product $(x, y)_{\bar{\mathbf{A}}} = (x, \bar{\mathbf{A}}y)$. We recall that all subspaces $T_i(x)$ and $T(x)$ have been assumed to be closed, which is important; cf. Theorem 4.6 in [20].

THEOREM 2.4. *Assume all $\bar{\mathbf{N}}_i = \mathbf{0}$ and $\bar{\mathbf{A}}$ is positive definite on $T(\bar{x})$.² Then $\|\bar{\mathbf{P}}\mathbf{S}'(\bar{x})\bar{\mathbf{P}}\|_{\bar{\mathbf{A}}} < 1$. In particular, $\rho(\mathbf{S}'(\bar{x})\bar{\mathbf{P}}) = \rho(\bar{\mathbf{P}}\mathbf{S}'(\bar{x})\bar{\mathbf{P}}) < 1$.*

The case $\bar{\mathbf{N}}_i = \mathbf{0}$ considered here arises in two notable cases.

2.2.1. Locally constant subspaces. If the subspaces $T_i(x)$ are the same for all x in a neighborhood of \bar{x} , then $\mathbf{P}'_i(\bar{x}) = \mathbf{0}$. This case occurs in the classical nonlinear Gauss–Seidel method discussed in the introduction, where the subspaces T_i are fixed and do not depend on x at all. Hence, in this case Theorem 2.4 simply recovers the well-known fact that the local convergence rate of the nonlinear (block) Gauss–Seidel method equals the rate of the linear block Gauss–Seidel method with the Hessian as the system matrix; cf., e.g., [12].

2.2.2. Zero gradient. The operators $\bar{\mathbf{N}}_i$ are also zero in fixed points satisfying $\nabla f(\bar{x}) = \mathbf{0}$. An interesting scenario for this situation is low-rank optimization where an unconstrained critical point lies on a considered manifold of low-rank matrices or tensors. This scenario is presented in section 3; see, in particular, Lemma 3.1.

2.3. A nontrivial example including curvature ($\bar{\mathbf{N}}_i \neq \mathbf{0}$). A case with $\bar{\mathbf{N}}_i \neq \mathbf{0}$, but allowing for considerable simplification, is obtained for $d = 2$ when $T(\bar{x}) = T_1(\bar{x}) + T_2(\bar{x})$ can be decomposed into its intersection and two other $\bar{\mathbf{A}}$ -orthogonal parts. This case occurs for problems of low-rank best approximation. In these cases, f is a quadratic function with Hessian equal to identity matrix: $\bar{\mathbf{A}} = \mathbf{I}$; see section 3 below.

THEOREM 2.5. *In addition to the assumptions of Theorem 2.3, suppose the following two conditions hold:*

- (i) $\bar{\mathbf{P}}_1^{\bar{\mathbf{A}}}$ and $\bar{\mathbf{P}}_2^{\bar{\mathbf{A}}}$ commute,³
- (ii) $\bar{\mathbf{N}}_i = \mathbf{0}$ on $T_i(\bar{x})$ for $i = 1, 2$.

Then

$$\rho(\mathbf{S}'(\bar{x})\bar{\mathbf{P}}) = \rho(\bar{\mathbf{B}}_2\bar{\mathbf{N}}_2\bar{\mathbf{B}}_1\bar{\mathbf{N}}_1\bar{\mathbf{P}}).$$

Proof. When $\bar{\mathbf{P}}_1^{\bar{\mathbf{A}}}$ and $\bar{\mathbf{P}}_2^{\bar{\mathbf{A}}}$ commute, it is easily verified that the operator $(\mathbf{I} - \bar{\mathbf{P}}_1^{\bar{\mathbf{A}}})\bar{\mathbf{P}}$ maps to $T_2(\bar{x})$. By (2.9) and assumption (ii), it then holds that

$$\mathbf{S}'(\bar{x})\bar{\mathbf{P}} = -(\mathbf{I} - \bar{\mathbf{P}}_2^{\bar{\mathbf{A}}} - \bar{\mathbf{B}}_2\bar{\mathbf{N}}_2)\bar{\mathbf{B}}_1\bar{\mathbf{N}}_1\bar{\mathbf{P}} = -(\mathbf{I} - \bar{\mathbf{P}}_2^{\bar{\mathbf{A}}} - \bar{\mathbf{B}}_2\bar{\mathbf{N}}_2)\bar{\mathbf{P}}\bar{\mathbf{B}}_1\bar{\mathbf{N}}_1\bar{\mathbf{P}}.$$

²It means that there exists $m > 0$ such that $\langle x, \bar{\mathbf{A}}x \rangle \geq m\|x\|^2$ for all $x \in T(\bar{x})$.

³This condition is equivalent to the fact that the $\bar{\mathbf{A}}$ -orthogonal projector on $T(\bar{x})$ allows the two decompositions $\bar{\mathbf{P}}^{\bar{\mathbf{A}}} = \bar{\mathbf{P}}_1^{\bar{\mathbf{A}}} + \bar{\mathbf{P}}_2^{\bar{\mathbf{A}}} - \bar{\mathbf{P}}_2^{\bar{\mathbf{A}}}\bar{\mathbf{P}}_1^{\bar{\mathbf{A}}}$ and $\bar{\mathbf{P}}^{\bar{\mathbf{A}}} = \bar{\mathbf{P}}_1^{\bar{\mathbf{A}}} + \bar{\mathbf{P}}_2^{\bar{\mathbf{A}}} - \bar{\mathbf{P}}_1^{\bar{\mathbf{A}}}\bar{\mathbf{P}}_2^{\bar{\mathbf{A}}}$.

It is a well-known fact that the spectral radius of the product of two operators is invariant under the order of factors. Thus, by the above formula, the spectral radius of $\mathbf{S}'(\bar{x})\bar{\mathbf{P}}$ is the same as the spectral radius of $-\bar{\mathbf{B}}_1\bar{\mathbf{N}}_1\bar{\mathbf{P}}(\mathbf{I} - \bar{\mathbf{P}}_2^{\mathbf{A}} - \bar{\mathbf{B}}_2\bar{\mathbf{N}}_2)\bar{\mathbf{P}} = \bar{\mathbf{B}}_1\bar{\mathbf{N}}_1\bar{\mathbf{P}}\bar{\mathbf{B}}_2\bar{\mathbf{N}}_2\bar{\mathbf{P}}$. Here we have used that $(\mathbf{I} - \bar{\mathbf{P}}_2^{\mathbf{A}})\bar{\mathbf{P}}$ maps to $T_1(\bar{x})$ by (i). Changing the order of factors again, we obtain the result. \square

Remark 2.6. An even stronger result is obviously obtained when again $\bar{\mathbf{N}}_i = 0$ on the whole space $T(\bar{x})$ as in section 2.2. Then $\mathbf{S}'(\bar{x})\bar{\mathbf{P}} = 0$ and we expect superlinear convergence (given sufficient smoothness of \mathbf{S}). This happens, for instance, when $\nabla f(\bar{x}) = 0$. If, additionally, f is quadratic, then the sequential solution of $\mathbf{P}_i(\bar{x})\nabla f(y) = 0$ on $T_i(\bar{x})$ provides a critical point on the whole space $T(\bar{x})$ after only one sweep through $i = 1, 2$ (due to orthogonal residuals). Of course, the condition (i) in Theorem 2.5 is very strong when $\bar{\mathbf{A}}$ is not the identity operator, as it implies that we are given a possibly overlapping, but otherwise $\bar{\mathbf{A}}$ -orthogonal splitting of the space $T(\bar{x})$.

3. AO for low-rank matrices. We return to the AO method (1.2) for solving the problem

$$(3.1) \quad \min_{\text{rank}(X) \leq k} f(X)$$

for a function $f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$, as outlined in the introduction. We first give an overview of what the abstract setup developed above looks like in this case. We then deal with the alternating least squares method for quadratic functions f , and its relation to power iterations in the case that the Hessian of the function is the identity operator. By $\langle X, Y \rangle_F = \sum_{i,j} x_{ij}y_{ij}$ we denote the Frobenius inner product of two matrices, and by $\|X\|_F$ the corresponding induced norm.

Starting from an initial guess $X_0 = U_0V_0^T$ of rank k , the method produces a sequence $X_\ell = U_\ell V_\ell^T$ of matrices of rank at most k by minimizing the function $f(UV^T)$ with respect to U and V only in an alternating manner. As long as the matrices U_ℓ and V_ℓ remain of rank k , this method is equivalently described as AO on the varying subspaces defined in (1.5) and (1.6).⁴ Using the projections

$$(3.2) \quad Z \mapsto \mathbf{P}_1(X)[Z] := ZX^+X, \quad Z \mapsto \mathbf{P}_2(X)[Z] := XX^+Z,^5$$

these subspaces can also be written as

$$\begin{aligned} T_1(X) &= \{Y \in \mathbb{R}^{m \times n} : Y = \mathbf{P}_1(X)[Y]\}, \\ T_2(X) &= \{Y \in \mathbb{R}^{m \times n} : Y = \mathbf{P}_2(X)[Y]\}. \end{aligned}$$

Here we recall that the Moore–Penrose inverse of $X \in \mathbb{R}^{m \times n}$ is defined as $X^+ = V\Sigma^{-1}U^T \in \mathbb{R}^{n \times m}$, where $X = U\Sigma V^T$ is a “slim” singular value decomposition of X involving only the positive singular values. It is then obvious that the $\mathbf{P}_i(X)$ are projections whose ranges are the subspaces $T_i(X)$ as defined in (1.5). We also see that XX^+ and X^+X are themselves orthogonal projections in \mathbb{R}^m and \mathbb{R}^n , respectively.

⁴When the rank drops, some formal subtleties appear. In the alternating subspace method the rank can only decrease, but never increase again, whereas in the AO method for U and V the size of the blocks is not changed, and even if, say, U with rank less than k is fixed, the minimizer for V is then not unique and a full-rank matrix V could be selected.

⁵Note that different from previous notation in \mathbb{R}^n we now use square brackets in $\mathbf{P}(X)[Z]$ to describe the linear action of $\mathbf{P}(X)$ on Z in order to avoid confusion with matrix multiplication.

Since the Frobenius inner product of two matrices can be computed column- or row-wise, it then easily follows that the $\mathbf{P}_i(X)$ are in fact orthogonal projections with respect to this inner product.

It is well known that for every k the set

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

is a smooth embedded submanifold of $\mathbb{R}^{m \times n}$ of codimension $(m - k)(n - k)$ [10, Example 8.14]. It can further be shown that the space

$$T(X) = T_1(X) + T_2(X)$$

is the tangent space to that manifold at $X \in \mathcal{M}_k$. Therefore, if \bar{X} has rank k and is a fixed point of a (locally) smooth map $\mathbf{S}: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ satisfying

$$(3.3) \quad \text{rank}(\mathbf{S}(X)) \leq \text{rank}(X)$$

for all X , then the condition $\rho(\mathbf{S}'(\bar{X})\bar{\mathbf{P}}) < 1$ is sufficient for R-linear convergence

$$(3.4) \quad \limsup_{\ell \rightarrow \infty} \|X_\ell - \bar{X}\|^{1/\ell} \leq \rho(\mathbf{S}'(\bar{X})\bar{\mathbf{P}})$$

(in any norm, since we are now in a finite-dimensional setting) of an iteration

$$X_{\ell+1} = \mathbf{S}(X_\ell)$$

with starting guess X_0 of rank k close enough to \bar{X} .⁶

From (3.2) it is obvious that $\mathbf{P}_1(X)$ and $\mathbf{P}_2(X)$ commute. Correspondingly,

$$(3.5) \quad \mathbf{P}(X) = \mathbf{P}_1(X) + \mathbf{P}_2(X) - \mathbf{P}_1(X)\mathbf{P}_2(X) = \mathbf{P}_1(X) + \mathbf{P}_2(X) - \mathbf{P}_2(X)\mathbf{P}_1(X)$$

is the orthogonal projection on $T(X)$.

3.1. Derivatives of projections. The reader will have noticed that the mappings $X \mapsto \mathbf{P}_i(X)$ as defined in (3.2) are not differentiable on $\mathbb{R}^{m \times n}$ unless $\text{rank}(X) = \min(m, n)$, since the map $X \mapsto X^+$ is not. To resolve this potential conflict to the theory developed above, we can formally extend the projections to smooth maps in a neighborhood of \mathcal{M}_k . Indeed, let \mathcal{D} be the open set of all matrices whose k th singular value is strictly larger than the $(k + 1)$ th one (such a matrix necessarily has rank at least k). To every $X \in \mathcal{D}$ we attach the orthogonal projections $U(X)U(X)^T$ and $V(X)V(X)^T$ onto the subspaces spanned by the dominant left, respectively, right singular vectors gathered as columns in the matrices $U(X) \in \mathbb{R}^{m \times k}$, respectively, $V(X) \in \mathbb{R}^{n \times k}$. These maps are smooth on \mathcal{D} . The projections \mathbf{P}_i from (3.2) can hence be extended to smooth maps on \mathcal{D} via

$$(3.6) \quad \mathbf{P}_1(X)[Z] = ZV(X)V(X)^T, \quad \mathbf{P}_2(X)[Z] = U(X)U(X)^T Z,$$

which coincide with (3.2) if $X \in \mathcal{M}_k$. The formula (3.5) remains valid.

⁶Let us prove this. If \bar{X} is a fixed point, then, by continuity, $\mathbf{S}(X)$ is close to \bar{X} when X is close to \bar{X} . Hence under the given assumptions, $\text{rank}(\mathbf{S}(X)) = k$ for all X with $\text{rank}(X) = k$ that are close enough to \bar{X} (by semicontinuity of rank). Therefore, \mathbf{S} can be locally regarded as a map between smooth submanifolds of \mathcal{M}_k , $\mathbf{S}'(\bar{X})$ maps the tangent space $T(\bar{X})$ into itself, and the sufficiency of the condition $\rho(\mathbf{S}'(\bar{X})) < 1$ on $T(\bar{X})$ for local contractivity follows in the same way as in linear space using differential calculus on manifolds.

We now present the derivatives of \mathbf{P}_1 and \mathbf{P}_2 at points $X \in \mathcal{M}_k$. We first consider directional derivatives $\mathbf{P}'_i(X; H)$ for $H \in T(X)$, which is the tangent space to \mathcal{M}_k at X . The Moore–Penrose pseudoinverse is a smooth map on manifolds of constant rank, and its Riemannian derivative at $X \in \mathcal{M}_k$ is given by

$$(3.7) \quad DX^+[H] = -X^+HX^+ + X^+(X^+)^T H^T(I - XX^+) + (I - X^+X)H^T(X^+)^T X^+$$

with $H \in T(X)$; see [7]. Hence, we compute from (3.2) that

$$(3.8) \quad \begin{aligned} \mathbf{P}'_1(X; H)[Z] &= ZX^+H + Z \cdot DX^+[H] \cdot X \\ &= ZX^+H - ZX^+HX^+X + Z(I - X^+X)H^T(X^+)^T \end{aligned}$$

for $H \in T(X)$. Here we have used $(I - XX^+)X = 0$ and $(X^+)^T X^+X = (X^+)^T$. Correspondingly,

$$(3.9) \quad \begin{aligned} \mathbf{P}'_2(X; H)[Z] &= HX^+Z + X \cdot DX^+[H] \cdot Z \\ &= HX^+Z - XX^+HX^+Z + (X^+)^T H^T(I - XX^+)Z \end{aligned}$$

for $H \in T(X)$, since $X(I - X^+X) = 0$ and $XX^+(X^+)^T = (X^+)^T$.

Regarding directional derivatives $H \in T(X)^\perp$ (which will not be needed later on), we invite the reader to verify that $T(X)^\perp$ consists of all matrices of the form $H = (I - XX^+)E(I - X^+X)$ for some E , and that small perturbations of X along such directions do not change the dominant singular vectors. Hence $t \mapsto \mathbf{P}_i(X + tH)$ as defined in (3.6) is constant for small $|t|$ and so

$$\mathbf{P}'_1(X; H) = \mathbf{P}'_2(X; H) = 0 \quad \text{for } H \in T(X)^\perp.$$

Note that the formulas (3.8) and (3.9) also yield zero when applied to $H \in T(X)^\perp$ (since $X^+H = 0$ and $HX^+ = 0$) and can hence be used in general.

The formulas (3.8) and (3.9) can be considerably simplified when Z is orthogonal to the tangent space $T(X)$, since in this case $\mathbf{P}_1(X)[Z] = 0$, implying $ZX^+ = 0$, and $\mathbf{P}_2(X)[Z] = 0$, implying $X^+Z = 0$. For such Z , (3.8) and (3.9) become

$$\mathbf{P}'_1(X; H)[Z] = ZH^T(X^+)^T$$

and

$$\mathbf{P}'_2(X; H)[Z] = (X^+)^T H^T Z.$$

Note that since we need to derive the operators $\bar{\mathbf{N}}_i[H] = \mathbf{P}'_i(\bar{X}; H)[\nabla f(\bar{X})]$ defined in (2.8) at critical points \bar{X} of f on \mathcal{M}_k , where $\nabla f(\bar{X})$ is orthogonal to $T(\bar{X})$, this is indeed the case of interest. For reference we state this as a lemma.

LEMMA 3.1. *Let $X \in \mathcal{M}_k$ be a critical point of f on \mathcal{M}_k , that is, $\mathbf{P}(\bar{X})\nabla f(\bar{X}) = 0$. Then for the projections (3.2) it holds*

$$\bar{\mathbf{N}}_1[H] := \mathbf{P}'_1(\bar{X}; H)\nabla f(\bar{X}) = \nabla f(\bar{X})H^T(\bar{X}^+)^T$$

and

$$\bar{\mathbf{N}}_2[H] := \mathbf{P}'_2(\bar{X}; H)[\nabla f(\bar{X})] = (\bar{X}^+)^T H^T \nabla f(\bar{X})$$

for all $H \in T(\bar{X})$. In particular, $\bar{\mathbf{N}}_i = 0$ on $T_i(\bar{X})$.

Remark 3.2. Regarding the initial problem (3.1) on $\mathcal{M}_{\leq k}$, we remark that the “smoothness” assumption $\text{rank}(\bar{X}) = k$, which has been crucial in the above derivations, is plausible in most applications, except for very special or artificial cases. It has been shown in [14, Corollary 3.4] that critical points \bar{X} of (3.1), for example, local minima, satisfy either $\text{rank}(\bar{X}) = k$ or $\nabla f(\bar{X}) = 0$.

3.2. Alternating least squares algorithm. When f is a strictly convex quadratic function, the outlined method is known as the *alternating least squares* (ALS) method. Let us give formulas for this important special case in more detail.

For simplicity, we assume that $f(0) = 0$. Then f takes the form

$$(3.10) \quad f(X) = \frac{1}{2} \langle X, \mathbf{A}[X] \rangle_F - \langle X, B \rangle_F,$$

where \mathbf{A} is a symmetric positive definite linear operator on $\mathbb{R}^{m \times n}$, and $B \in \mathbb{R}^{m \times n}$. We have $\nabla f(X) = \mathbf{A}[X] - B$, and the Hessian at every point is the operator \mathbf{A} .

Minimizing the function f without constraint is equivalent to solving the linear matrix equation $\mathbf{A}[X] = B$. Let

$$X^* = \mathbf{A}^{-1}[B]$$

be the solution. Introducing the \mathbf{A} -norm

$$\|X\|_{\mathbf{A}} = \sqrt{\langle X, \mathbf{A}[X] \rangle_F}$$

on $\mathbb{R}^{m \times n}$, we can rewrite the function f as

$$f(X) = \frac{1}{2} \|X - X^*\|_{\mathbf{A}}^2 - \frac{1}{2} \|X^*\|_{\mathbf{A}}^2.$$

This shows that minimizing f on $\mathcal{M}_{\leq k}$ is equivalent to finding the best rank- k approximation(s) of the true solution X^* in \mathbf{A} -norm, and they serve as approximate low-rank solutions to the linear equation. The ALS algorithm tries to find such minima of f on $\mathcal{M}_{\leq k}$.

At a given iterate $X_\ell \in \mathcal{M}_{\leq k}$, the first step of ALS computes

$$\mathbf{S}_1(X_\ell) = \operatorname{argmin}_{X \in T_1(X_\ell)} f(X).$$

Since \mathbf{A} is positive definite, there is indeed a unique solution, and it is given as

$$(3.11) \quad \mathbf{S}_1(X) = (\mathbf{P}_1(X)\mathbf{A}\mathbf{P}_1(X))^{-1}[\mathbf{P}_1(X)[B]].$$

Here, as usual, $(\mathbf{P}_1(X)\mathbf{A}\mathbf{P}_1(X))^{-1}$ is understood as the inverse of the operator $\mathbf{P}_1(X)\mathbf{A}\mathbf{P}_1(X)$ on its invariant subspace $T_1(X)$. The map $X \mapsto \mathbf{S}_1(X)$ is differentiable on the manifold \mathcal{M}_k of rank- k matrices, since \mathbf{P}_1 is.

If $\mathbf{S}_1(X_\ell)$ has rank k ,⁷ then the next step of ALS computes

$$(3.12) \quad X_{\ell+1} = \mathbf{S}(X_\ell) := \mathbf{S}_2(\mathbf{S}_1(X_\ell)) = \operatorname{argmin}_{X \in T_2(\mathbf{S}_1(X_\ell))} f(X).$$

The solution map \mathbf{S}_2 is given as

$$(3.13) \quad \mathbf{S}_2(X) = (\mathbf{P}_2(X)\mathbf{A}\mathbf{P}_2(X))^{-1}[\mathbf{P}_2(X)[B]].$$

We repeat once more that $X \in T_1(X)$ and $X \in T_2(X)$ for every X , so we are in the abstract framework developed in section 2.

The original idea of AO for low-rank optimization is to operate on a (nonunique) factorization $X_\ell = U_\ell V_\ell^T$. In terms of these factors, more precisely, their vectorizations,

⁷If not, there are several options, but we ignore that case.

the ALS method becomes the algorithm displayed as Algorithm 3, where \mathbf{A} is to be understood as an $mn \times mn$ matrix and \otimes is the standard Kronecker product for matrices. As explained in the introduction, the QR decompositions are not mandatory in theory, but highly recommended in practice for numerical stability.

Algorithm 3: ALS algorithm for (3.10).

Input: $B \in \mathbb{R}^{m \times n}$, $V_0 \in \mathbb{R}^{n \times k}$.
for $\ell = 0, 1, 2, \dots$ **do**
 $\text{vec}(U) = ((V_\ell^T \otimes I) \mathbf{A} (V_\ell \otimes I))^{-1} \text{vec}(BV_\ell)$, $QR = \text{qr}(U)$
 $U := Q$
 $\text{vec}(V^T) = ((I \otimes U^T) \mathbf{A} (I \otimes U))^{-1} \text{vec}(U^T B)$, $QR = \text{qr}(V)$
 $V_{\ell+1} := Q$, $U_{\ell+1} := UR^T$
end

3.3. SVD block power method. As a special case, we now consider the quadratic function (3.10) with $\mathbf{A} = \mathbf{I}$. It corresponds to the task

$$(3.14) \quad \min_{\text{rank}(X) \leq k} \frac{1}{2} \|X - B\|_F^2$$

of computing a best rank- k approximation to matrix B in the Frobenius norm. Since $\mathbf{A} = \mathbf{I}$, we have $(\mathbf{P}_i(X) \mathbf{A} \mathbf{P}_i(X))^{-1} \mathbf{P}_i(X) = \mathbf{P}_i(X)$ for $i = 1, 2$, and hence the update formulas (3.11) and (3.13) for AO simplify to

$$(3.15) \quad \mathbf{S}_1(X) = \mathbf{P}_1(X)[B] = BX^+X, \quad \mathbf{S}_2(X) = \mathbf{P}_2(X)[B] = XX^+B.$$

The resulting ALS iteration $X_{\ell+1} = \mathbf{S}(X_\ell)$ becomes

$$(3.16) \quad X_{\ell+1/2} = BX_\ell^+X_\ell, \quad X_{\ell+1} = X_{\ell+1/2}X_{\ell+1/2}^+B.$$

Writing $X_\ell = U_\ell \Sigma_\ell V_\ell^T$, it is easily seen (and shown below) that the sequence generated by (3.16) is the same as in the *simultaneous orthogonal iteration*, which is a two-sided block power method for computing the dominant k left and right singular subspaces of B , displayed as Algorithm 4 (provided X_0 has the row space spanned by V_0).

Algorithm 4: Simultaneous orthogonal iteration.

Input: $B \in \mathbb{R}^{m \times n}$, $k \leq \min(m, n)$, $V_0 \in \mathbb{R}^{n \times k}$ such that $V_0^T V_0 = I$
for $\ell = 0, 1, 2, \dots$ **do**
 $QR = \text{qr}(BV_\ell)$ // tall QR decomposition
 $U_{\ell+1} := Q$
 $QR = \text{qr}(B^T U_{\ell+1})$ // tall QR decomposition
 $V_{\ell+1} := Q$, $S_{\ell+1} = R^T$
end

Let

$$(3.17) \quad B = \sum_{i=1}^{\min(m,n)} \sigma_i \bar{u}_i \bar{v}_i^T$$

be the SVD of B , that is, $\bar{u}_1, \bar{u}_2, \dots$ and $\bar{v}_1, \bar{v}_2, \dots$ are orthonormal systems in \mathbb{R}^m and \mathbb{R}^n , respectively, and $\sigma_1 \geq \sigma_2 \geq \dots \geq 0$. If $\sigma_k > \sigma_{k+1}$, then it can be shown that

the sequence $X_\ell = U_\ell S_\ell V_\ell^T$ generated in Algorithm 4 converges to the unique best rank- k approximation

$$(3.18) \quad \bar{X} = \sum_{i=1}^k \sigma_i \bar{u}_i \bar{v}_i^T$$

for almost every starting guess $X_0 = U_0 V_0^T$. In fact, the method produces the same subspaces as the corresponding orthogonal iterations for the symmetric matrices $B^T B$ and $B B^T$, respectively, whose eigenvalues are the σ_i^2 (zero may be a further eigenvalue). Hence, by well-known results, $U_\ell \rightarrow [\bar{u}_1, \dots, \bar{u}_k]$ and $V_\ell \rightarrow [\bar{v}_1, \dots, \bar{v}_k]$ in terms of subspaces for almost every initialization, with a convergence rate $O(\sigma_{k+1}^2/\sigma_k^2)$; see [3, 8]. As an application of our abstract framework we are able to obtain this (known) rate of convergence from the local convergence analysis of the ALS sequence (3.16).

THEOREM 3.3. *Let B have singular values $\sigma_1 \geq \dots \geq \sigma_k > \sigma_{k+1} \geq \dots$, and the unique best rank- k approximation \bar{X} . Then the sequence $X_{\ell+1} = \mathbf{S}(X_\ell) = \mathbf{S}_2(\mathbf{S}_1(X_\ell))$ defined via (3.15), respectively, (3.16) (AO for problem (3.14)) is, in exact arithmetic, identical to the sequence $X_\ell = U_\ell S_\ell V_\ell^T$ generated by the simultaneous orthogonal iteration (Algorithm 4). With $\bar{\mathbf{P}} = \mathbf{P}(\bar{X})$ as before, it holds that*

$$(3.19) \quad \rho(\mathbf{S}'(\bar{X})\bar{\mathbf{P}}) = \left(\frac{\sigma_{k+1}}{\sigma_k} \right)^2.$$

Consequently, by (3.4), the sequence (X_ℓ) converges (for close enough starting guesses $X_0 \in \mathcal{M}_k$) R -linearly to \bar{X} at a rate

$$(3.20) \quad \limsup_{\ell \rightarrow \infty} \|X_\ell - \bar{X}\|^{1/\ell} \leq \left(\frac{\sigma_{k+1}}{\sigma_k} \right)^2$$

(in any norm). The convergence of the column and row spaces can be estimated correspondingly in the sense of the operator norm of projectors as

$$(3.21) \quad \limsup_{\ell \rightarrow \infty} \|\mathbf{P}_i(X_\ell) - \mathbf{P}_i(\bar{X})\|^{1/\ell} \leq \left(\frac{\sigma_{k+1}}{\sigma_k} \right)^2, \quad i = 1, 2.$$

Proof. We first show by induction that the methods are the same. If X_ℓ has the row space spanned by V_ℓ , then $X_{\ell+1/2}$ in (3.16) can be written $B V_\ell V_\ell^T$, which has the same column space as $B V_\ell$. Therefore, using $U_{\ell+1}$ from Algorithm 4, we get that $X_{\ell+1}$ from (3.16) equals $U_{\ell+1} U_{\ell+1}^T B = U_{\ell+1} \Sigma_{\ell+1} V_{\ell+1}^T$.

One may attempt to compute the spectral radius of $\mathbf{S}'(\bar{X})\bar{\mathbf{P}}$ from the explicit formulas (3.15) and (3.7), but it will be more elegant to invoke Theorem 2.5. Since $\bar{\mathbf{A}} = \mathbf{A} = \mathbf{I}$, the condition in item (i) of that theorem is obviously satisfied ($\mathbf{P}_1(X)$ and $\mathbf{P}_2(X)$ commute; see (3.2)). The condition (ii), that $\bar{\mathbf{N}}_i = 0$ on $T_i(\bar{X})$, is stated in Lemma 3.1. Taking into account further that the $\bar{\mathbf{B}}_i$ are identities, Theorem 2.5 yields the formula

$$(3.22) \quad \rho(\mathbf{S}'(\bar{X})\bar{\mathbf{P}}) = \rho(\bar{\mathbf{N}}_2 \bar{\mathbf{N}}_1 \bar{\mathbf{P}})$$

for the iteration (3.16). By Lemma 3.1,

$$\bar{\mathbf{N}}_2[\bar{\mathbf{N}}_1[H]] = (\bar{X}^+)^T \bar{X}^+ H (\nabla f(\bar{X}))^T \nabla f(\bar{X}) \quad \text{for } H \in T(\bar{X}).$$

Taking further into account that $\nabla f(\bar{X}) = \bar{X} - B$, this shows that

$$(3.23) \quad \bar{\mathbf{N}}_2 \bar{\mathbf{N}}_1 \bar{\mathbf{P}} = \bar{\mathbf{P}} \bar{\mathbf{N}}_2 \bar{\mathbf{N}}_1 \bar{\mathbf{P}} = \bar{\mathbf{P}} \cdot [(\bar{X}^+)^T \bar{X}^+ \otimes (\bar{X} - B)(\bar{X} - B)^T] \cdot \bar{\mathbf{P}},$$

where we use the Kronecker product operator notation (see (3.24)). By (3.17) and (3.18), the rank-one matrices $E_{i,j} = \bar{u}_i \bar{v}_j^T$, $i = 1, \dots, m$, $j = 1, \dots, n$, form an orthonormal system of eigenvectors of the operator $(\bar{X}^+)^T \bar{X}^+ \otimes (\bar{X} - B)(\bar{X} - B)^T$, corresponding to eigenvalues $\lambda_{i,j} = \sigma_j^2 / \sigma_i^2$ for $i \leq k$ and $j > k$, and $\lambda_{i,j} = 0$ otherwise. The largest of these eigenvalues is $\lambda_{k,k+1} = \sigma_{k+1}^2 / \sigma_k^2$. Since the corresponding eigenvector $E_{k,k+1} = \bar{u}_k \bar{v}_{k+1}^T$ belongs to $T_2(\bar{X}) \subseteq T(\bar{X})$ (see (3.2)), the formula (3.23) implies that $\lambda_{k,k+1}$ is also the largest (in modulus) eigenvalue of $\bar{\mathbf{N}}_2 \bar{\mathbf{N}}_1 \bar{\mathbf{P}}$, which proves the assertion (3.19).

Since $\text{rank}(\bar{X}) = k$, it follows that \mathbf{S} is a local contraction on the manifold \mathcal{M}_k in the neighborhood of \bar{X} , and the R-linear convergence rate of $\|X_\ell - \bar{X}\|_F$ is as asserted (see the explanations for (3.4) in section 3).

Let us show that (3.20) implies (3.21) for \mathbf{P}_1 . For all Z with $\|Z\|_F = 1$, we can estimate

$$\begin{aligned} \|(\mathbf{P}_1(X_\ell) - \mathbf{P}_1(\bar{X}))[Z]\|_F &= \|Z(X_\ell^+ X_\ell - \bar{X}^+ \bar{X})\|_F \\ &\leq \|X_\ell^+ X_\ell - \bar{X}^+ \bar{X}\|_2 \\ &= \|X_\ell^+(X_\ell - \bar{X}) + (X_\ell^+ - \bar{X}^+) \bar{X}\|_2 = O(\|X_\ell - \bar{X}\|_2), \end{aligned}$$

since $X_\ell \rightarrow \bar{X}$ on \mathcal{M}_k (implying that $\|X_\ell^+\|_2$ is bounded). □

Remark 3.4. When $\sigma_k > \sigma_{k+1} = 0$, that is $\text{rank}(B) = k$, the theorem yields $\rho(\mathbf{S}'(\bar{X})\bar{\mathbf{P}}) = 0$ which technically indicates superlinear convergence. In fact, this is a situation where Remark 2.6 applies: it holds $\nabla f(\bar{X}) = 0$ and, hence, $\bar{\mathbf{N}}_i = 0$. However, as it is known, and not difficult to see, the power method (3.16) initialized with the correct rank k will find $\bar{X} = B$ after only one sweep for almost every starting guess X_0 . The only condition is that $X_{1/2} = BX_0^+ X_0$ is of rank k which, in particular, is true for all X_0 in some neighborhood of $\bar{X} = B$.

3.4. Kronecker product operators. A main feature in the previous derivation of the local convergence rate of the block power method via the ALS analysis was the possibility of applying Theorem 2.5 for the computation of the spectral radius $\rho(\mathbf{S}'(\bar{X})\bar{\mathbf{P}})$, since for $\bar{\mathbf{A}} = \mathbf{A} = \mathbf{I}$ the $\bar{\mathbf{A}}$ -orthogonal projectors $\bar{\mathbf{P}}_1^{\bar{\mathbf{A}}}$ and $\bar{\mathbf{P}}_2^{\bar{\mathbf{A}}}$ commute. Note that $\mathbf{I} = I \otimes I$ is a Kronecker product of two identity matrices. To allow for at least a small generalization, we now investigate the case that $\bar{\mathbf{A}} = A_1 \otimes A_2$ is a Kronecker product of symmetric positive definite matrices. One can show that in this case the projectors $\bar{\mathbf{P}}_1^{\bar{\mathbf{A}}}$ and $\bar{\mathbf{P}}_2^{\bar{\mathbf{A}}}$ still commute and, hence, derive estimates for $\rho(\mathbf{S}'(\bar{X})\bar{\mathbf{P}})$ for this case based on Theorem 2.5. There is, however, a simpler way to analyze the ALS method for Kronecker product operators by reducing it to the block power method again.

Consider a quadratic function

$$f(X) = \frac{1}{2} \langle X, \mathbf{A}[X] \rangle_F - \langle X, B \rangle_F$$

on $\mathbb{R}^{m \times n}$, where the Hessian is a Kronecker product operator,

$$\bar{\mathbf{A}} = \mathbf{A} = A_2 \otimes A_1, \quad A_1 \in \mathbb{R}^{m \times m}, \quad A_2 \in \mathbb{R}^{n \times n},$$

by which we mean that

$$(3.24) \quad \mathbf{A}[X] = (A_2 \otimes A_1)[X] = A_1 X A_2^T.$$

Since \mathbf{A} should be symmetric positive definite, we assume that A_1 and A_2 are both symmetric positive definite.

We have already noted in section 3.2 that minimizing f subject to $\text{rank}(X) \leq k$ corresponds to finding the best rank- k approximation of the global minimum $X^* = \mathbf{A}^{-1}[B]$ in \mathbf{A} -norm. For the case that \mathbf{A} is a Kronecker product operator of the considered type, the best rank- k approximation in the \mathbf{A} -norm can be in principle computed directly via SVD. For this we rewrite

$$\begin{aligned} f(X) &= \frac{1}{2} \|\mathbf{A}^{1/2}[X] - \mathbf{A}^{-1/2}[B]\|_F^2 - \frac{1}{2} \|X^*\|_{\mathbf{A}}^2 \\ &= \frac{1}{2} \|A_1^{1/2} X A_2^{1/2} - A_1^{-1/2} B A_2^{-1/2}\|_F^2 - \frac{1}{2} \|X^*\|_{\mathbf{A}}^2. \end{aligned}$$

Since left or right multiplication by an invertible matrix does not change the rank, we can clearly see that the global minima of f on $\mathcal{M}_{\leq k}$ are given as

$$(3.25) \quad \bar{X} = A_1^{-1/2} \bar{Y} A_2^{-1/2},$$

where \bar{Y} is a best rank- k approximation of $A_1^{-1/2} B A_2^{-1/2}$, which can be computed through SVD. Obviously, there is a unique global minimum \bar{X} if and only if $A_1^{-1/2} B A_2^{-1/2}$ has a unique best rank- k approximation in the Frobenius norm.

It should therefore not come as a surprise that the ALS method in this case will be equivalent to the block power method for the matrix $A_1^{-1/2} B A_2^{-1/2}$. To prove this, it will be convenient to have the ALS update formulas in explicit matrix notation at hand. Using a decomposition

$$X = USV^T, \quad U^T U = I_k, \quad V^T V = I_k,$$

the formulas (3.11) and (3.13) become

$$(3.26) \quad \mathbf{S}_1(X) = A_1^{-1} B V (V^T A_2 V)^{-1} V^T$$

and

$$(3.27) \quad \mathbf{S}_2(X) = U (U^T A_1 U)^{-1} U^T B A_2^{-1}.$$

THEOREM 3.5. *Let $\mathbf{A} = A_2 \otimes A_1$ with A_1 and A_2 being symmetric positive definite, and $B \in \mathbb{R}^{m \times n}$. Denote by $\varsigma_1 \geq \varsigma_2 \geq \dots$ the singular values of $C = A_1^{-1/2} B A_2^{-1/2}$. If then $\varsigma_k > \varsigma_{k+1}$, then for almost every starting point $X_0 \in \mathcal{M}_k$ the sequence $X_{\ell+1} = \mathbf{S}(X_\ell) = \mathbf{S}_2(\mathbf{S}_1(X_\ell))$ defined via (3.26), respectively, (3.27) (Algorithm 3) is well defined and converges to the unique global minimum $\bar{X} = A_1^{-1/2} \bar{Y} A_2^{-1/2}$ of the function f given by (3.10) on $\mathcal{M}_{\leq k}$, where \bar{Y} is the unique best rank- k approximation of C in the Frobenius norm. In fact, for almost every $X_0 \in \mathcal{M}_k$ it holds (in exact arithmetic) that*

$$X_\ell = A_1^{-1/2} Y_\ell A_2^{-1/2},$$

where Y_ℓ is the sequence generated by the SVD block power method (3.16) (Algorithm 4) applied to matrix C with starting point $Y_0 = A_1^{1/2} X_0 A_2^{1/2}$. The asymptotic R-linear convergence rate is estimated as

$$\limsup_{\ell \rightarrow \infty} \|X_\ell - \bar{X}\|^{1/\ell} = \limsup_{\ell \rightarrow \infty} \|Y_\ell - \bar{Y}\|^{1/\ell} \leq \left(\frac{\varsigma_{k+1}}{\varsigma_k} \right)^2$$

(in any norm).

Proof. We know that under the given assumptions on C the sequence Y_ℓ is well defined (that is, every half-step in (3.16) remains in \mathcal{M}_k) for almost every starting point $Y_0 \in \mathcal{M}_k$ and converges to \bar{Y} at an asymptotic R -linear rate $(\varsigma_{k+1}/\varsigma_k)^2$.

Let $X_\ell = A_1^{-1/2}Y_\ell A_2^{-1/2} \in \mathcal{M}_k$ be true for some ℓ . Then, by (3.26), we can write

$$\mathbf{S}_1(X_\ell) = A_1^{-1/2}CPA_2^{-1/2}$$

with

$$P_\ell = A_2^{1/2}V(V^T A_2 V)^{-1}V^T A_2^{1/2},$$

and the columns of V forming a basis for the row space of X_ℓ . We claim that $P = Y_\ell^+ Y_\ell$. To see this we note that P is symmetric and $P^2 = P$. Further, the null space of P obviously equals the orthogonal complement of the column space of $A_2^{1/2}V$. Hence P is the orthogonal projector on this subspace, which, however, equals the row space of $Y_\ell = A_1^{1/2}X_\ell A_2^{1/2}$. This shows $P = Y_\ell^+ Y_\ell$. It follows that

$$\mathbf{S}_1(X_\ell) = A_1^{-1/2}CY_\ell^+ Y_\ell A_2^{-1/2} = A_1^{-1/2}Y_{\ell+1/2}A_2^{-1/2},$$

where $Y_{\ell+1/2}$ is the next half-step from Y_ℓ in the block power method for C . The argument for the second half-step is analogous and the proof of the theorem is finished by induction. (Both inequalities in the asserted equality $\limsup \|X_\ell - \bar{X}\|^{1/\ell} = \limsup \|Y_\ell - \bar{Y}\|^{1/\ell}$ are immediate for a submultiplicative matrix norm.) \square

Remark 3.6. Analogously to Remark 3.4 we note that in the case $\varsigma_k > \varsigma_{k+1} = 0$ the theorem technically indicates a superlinear local convergence rate, while in reality the method will actually find the correct solution $\bar{X} = X^* = A_1^{-1}BA_2^{-1}$ in just one sweep for almost all X_0 (and, in particular, for all X_0 in a neighborhood of \bar{X}).

4. Numerical experiments. The goal of this section is to investigate the agreement between the theoretical estimates and the numerical behavior. The goal is to minimize the quadratic cost function (3.10) subject to $\text{rank}(X) \leq k$ using the ALS Algorithm 3. We consider four examples for the Hessian operator $\mathbf{A} = \bar{\mathbf{A}}$: the identity operator, a simple Kronecker product operator, a Laplace-like operator, and a random positive definite operator.

In all experiments, the initial guesses X_0 in Algorithm 3 have been chosen randomly. In the figures we depict lines corresponding to the theoretical rate of convergence $\rho(\mathbf{S}'(\bar{X})\bar{\mathbf{P}})$ by the color black, which has been computed numerically at the observed limit point \bar{X} by forming a matrix representation of the linear operator $\mathbf{S}'(\bar{X}) = [(\mathbf{I} - \bar{\mathbf{P}}_2^{\bar{\mathbf{A}}}) - \bar{\mathbf{B}}_2 \bar{\mathbf{N}}_2][(\mathbf{I} - \bar{\mathbf{P}}_1^{\bar{\mathbf{A}}}) - \bar{\mathbf{B}}_1 \bar{\mathbf{N}}_1]$ and solving a full eigenvalue problem to find the spectral radius. To assemble such a matrix representation of $\mathbf{S}'(\bar{X})$, we applied it successively to the (reshaped) columns of an $mn \times mn$ identity matrix using the formulas provided by Lemma 3.1.

In the plots, the theoretical rate $\rho(\mathbf{S}'(\bar{X})\bar{\mathbf{P}})$ is compared with the relative errors $\frac{\|X_\ell - \bar{X}\|}{\|\bar{X}\|}$ as well as with the relative norm of projected residuals $\frac{\|\mathbf{P}(X_\ell)[A[X_\ell] - B]\|}{\|\mathbf{P}(X_0)[A[X_0] - B]\|}$, which are the quantities of interest from the perspective of Riemannian optimization (since $A[X_\ell] - B = \nabla f(X_\ell)$). Moreover, the latter have the advantage that they can be monitored in practice during the iteration.

4.1. Case $\bar{\mathbf{A}} = \mathbf{I}$ and $\bar{\mathbf{A}} = \mathbf{A}_2 \otimes \mathbf{A}_1$. Consider the ALS method for problem (3.14), that is, minimizing the function

$$f(X) = \frac{1}{2}\|X - B\|_F^2$$

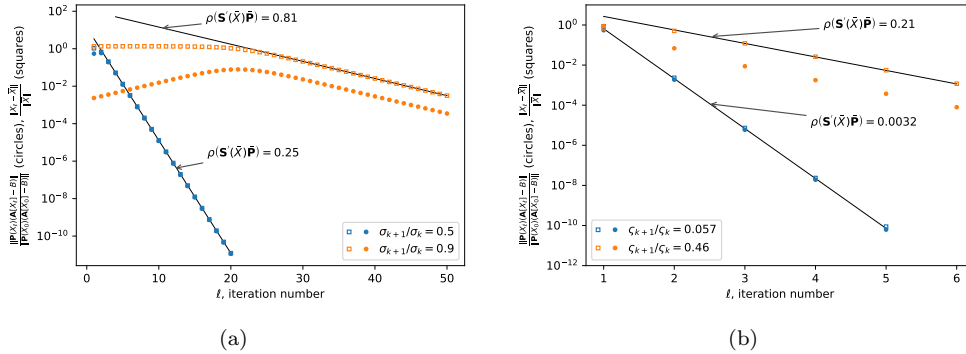


FIG. 1. (a) Relative errors in ALS w.r.t. the iteration number for $\bar{\mathbf{A}} = \mathbf{I}$, $k = 2$, $\sigma_k = 10^{-3}$. Black lines have slopes $\rho(\mathbf{S}'(\bar{X})\bar{\mathbf{P}}) = (\sigma_{k+1}/\sigma_k)^2$ for different σ_{k+1} , while colored dots represent the observed convergence. (b) Same experiment for $\bar{\mathbf{A}} = \mathbf{A}_2 \otimes \mathbf{A}_1$ with positive definite $\mathbf{A}_1, \mathbf{A}_2$. The ς_k are the singular values of $\mathbf{A}_1^{-1/2} \mathbf{B} \mathbf{A}_2^{-1/2}$. The value $\varsigma_k = 10^{-3}$ is kept fixed.

subject to $\text{rank}(X) \leq k$, where $B \in \mathbb{R}^{n \times n}$ is a given matrix with a predefined distribution of singular values. The goal is to find the best rank- k approximation \bar{X} of B .

Specifically, we consider $n = 50$, $k = 2$, set $\sigma_2(B) = 10^{-3}$, and test with different $\sigma_3(B)$'s. By Theorem 3.3, the ALS method in this case is locally convergent at with an asymptotic R -linear rate $(\sigma_{k+1}/\sigma_k)^2$ and, in fact, this bound is sharp.⁸ As illustrated in Figure 1(a), we observe close experimental agreement with this bound.

Note that if $\sigma_{k+1} = 0$, then the method (generically) converges in one iteration (so technically superlinear) since row and column spaces of B are found immediately (Remark 3.4).

Note that the other extreme case, when $\sigma_{k+1} = \sigma_k$, is not covered by our local convergence analysis, which does not necessarily mean absence of convergence to *some* best rank- k approximation. However, usually \bar{X} itself will not be a point of attraction of the block power method for all X_0 in the neighborhood. For instance, when $B = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ and $k = 1$, the matrix $\bar{X} = \frac{1}{2} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$ is a best rank-one approximation and a fixed point of the method. However, for $X_0 = \begin{bmatrix} \alpha & \beta \\ \alpha & \beta \end{bmatrix}$, the method becomes stationary after one sweep at $X_1 = \frac{1}{\alpha^2 + \beta^2} \begin{bmatrix} \alpha^2 & \alpha\beta \\ \alpha\beta & \beta^2 \end{bmatrix}$, which is also a best rank-one approximation of B , but equals \bar{X} only when $\alpha = \beta = 1/2$. At the same time, X_0 can be arbitrarily close to \bar{X} .

We also verify our theoretical result on Kronecker product operators. We randomly generate matrices $\tilde{A}_1, \tilde{A}_2 \in \mathbb{R}^{n \times n}$ and use the operator $\mathbf{A} = \bar{\mathbf{A}} = \mathbf{A}_2 \otimes \mathbf{A}_1$ with $\mathbf{A}_1 = \tilde{A}_1 \tilde{A}_1^T$, $\mathbf{A}_2 = \tilde{A}_2 \tilde{A}_2^T$ in the function f in (3.10). The global minimizer \bar{X} of f on $\mathcal{M}_{\leq k}$ is given by (3.25) and can be computed using SVD. By Theorem 3.5, the ALS algorithm should converge for almost any $X_0 \in \mathcal{M}_k$ to \bar{X} and the asymptotic R -linear rate is $(\varsigma_{k+1}/\varsigma_k)^2$, where ς_k are the singular values of $\mathbf{A}_1^{-1/2} \mathbf{B} \mathbf{A}_2^{-1/2}$. Figure 1(b) shows perfect agreement with the theoretical prediction. Here again we considered $k = 2$, and generated different right-hand sides B such that always $\varsigma_2 = 10^{-3}$.

4.2. More general symmetric positive definite $\bar{\mathbf{A}}$. We now go beyond Kronecker product operators. First, we consider an entirely random symmetric positive

⁸Using the classical linear algebra approach related to spectral decomposition and power method one should see that this rate is in fact attained for almost every starting guess.

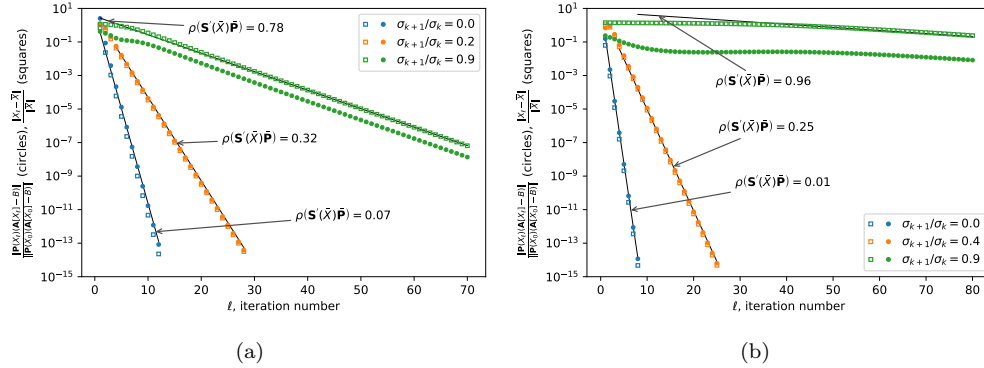


FIG. 2. Relative errors in ALS w.r.t. the iteration number for \mathbf{A} : (a) random symmetric positive definite matrix and (b) Laplace matrix. Experiments for $k = 2$ and predefined singular value $\sigma_k = 10^{-3}$ of the solution of $\mathbf{A}[X] = B$. Black lines have slopes corresponding to the theoretical convergence rate for different σ_{k+1} .

definite matrix

$$\mathbf{A} = \mathbf{R}^\top \mathbf{R} \in \mathbb{R}^{n^2 \times n^2},$$

where \mathbf{R} is a matrix with each element produced by the standard normal distribution. As another example, we take the highly structured matrix arising in the discretization of a two-dimensional Laplacian on uniform tensor product grid with zero Dirichlet boundary conditions:

$$\mathbf{A} = I_n \otimes D_n + D_n \otimes I_n \in \mathbb{R}^{n^2 \times n^2}, \quad D_n = (n + 1)^2 \text{tridiag}(-1, 2, -1)_{n \times n}.$$

Figure 2 displays experimental results for the ALS algorithm with $n = 50$ and $k = 2$. The matrix B has been chosen such that the solution of the matrix equation $\mathbf{A}[Y] = B$ (that is, the global minimizer of (3.10) without low-rank constraint) has a predefined distribution of singular values. Similarly to the previous experiments we set $\sigma_2(Y) = 10^{-3}$, while $\sigma_3(Y)$ varies and the goal is to find the best rank-2 approximation \tilde{X} . As we observe, numerical behavior is in close agreement with theoretical estimates. While the convergence rate $\rho(\mathbf{S}'(\tilde{X})\tilde{P})$ does not equal σ_{k+1}/σ_k as in the case $\tilde{\mathbf{A}} = \mathbf{I}$, it still seems related to this ratio for both choices of $\tilde{\mathbf{A}}$. Remarkably, $\rho(\mathbf{S}'(\tilde{X})\tilde{P})$ is considerably smaller than one even when σ_{k+1}/σ_k is close to one. A decisive question for future work would be for which combinations of $\tilde{\mathbf{A}}$ and B this can be rigorously shown.

Also note that, in contrast to the case of Kronecker product operators, there is no superlinear convergence in this experiment when $\sigma_{k+1} = 0$. In this situation the minimizer of (3.10) on the rank- k variety is the same as the global one, so the curvature-free case considered in section 2.2.2 (zero gradient) applies. Local linear convergence of the ALS method to this minimizer is then guaranteed by Theorem 2.4.

5. Conclusion. The goal of this paper was to derive transparent conditions for the local linear convergence of AO algorithms for multilinear and low-rank optimization, specifically the ALS algorithm, which reflect the underlying geometry and do not depend on the representation of low-rank tensors as in previous works. Due to multilinearity of the cost function, single optimization steps take place on linear subspaces, leading (in particular for quadratic cost functions) to an interpretation of AO as a nonlinear subspace correction method (with changing subspaces). Using a sufficiently general framework, a formula for the derivative of the nonlinear iteration

function can be obtained (Theorem 2.3), which displays the interplay of terms from the classic linear subspace correction method with the curvature of the underlying low-rank manifold and the gradient of the cost function in a clear way. The main task remains to show that the spectral radius of this derivative is less than one in applications of interest. This is true in low-rank optimization tasks where the global minimizer lies on the considered low-rank manifold. The case where this is not true is more subtle. For AO for low-rank matrices, the curvature terms can be considerably simplified, which allows for an alternative, analytic proof for the well-known convergence rate of the simultaneous orthogonal iteration for computing the dominant left and right singular subspaces of a matrix. While the main trick (Theorem 2.5) that was used to obtain this result may not apply in more general situations, we hope that our framework can be a useful starting point in future work for finding rigorous statements for the observed linear convergence of AO and ALS in other applications, like low-rank solutions of Lyapunov equations (cost function (3.10)) and low-rank tensor approximation.

Related work. In [16] and [13] the local convergence of the ALS algorithm has been analyzed for low-rank tensor approximation in the CP and tensor train formats, respectively, using the nonlinear Gauss–Seidel approach for a cost function of the form (1.7), e.g., using an explicit representation of low-rank tensors. To address the problem that the Hessian of this cost function cannot be positive definite due to nonuniqueness of tensor representations, equivalence classes of representations (level sets of the function τ in (1.7)) are introduced. Linear convergence is then established for the case that the null space of the Hessian equals the tangent space of the orbit of equivalent representations. The idea is certainly analogous to restricting the operator $\mathbf{S}'(\bar{x})$ to the subspace $T(\bar{x})$ as in the present paper, but we believe that our approach provides a much clearer picture by avoiding the unintuitive concept of equivalent representations. A formula

$$(5.1) \quad \langle h, \nabla F^2(\xi)h \rangle = \langle \tau'(\xi)h, \nabla^2 f(x)\tau'(\xi)h \rangle + \langle \nabla f(x), \tau''(\xi)[h, h] \rangle$$

for the Hessian at $x = \tau(\xi)$ is given in [13], which features the Hessian $\nabla^2 f(x)$ on the tangent space of the image of τ , and the interaction of curvature ($\tau''(x)$) and gradient ∇f as in our work (cf. the definition (2.8) of \mathbf{N}_i). In particular, it is concluded that local convergence is guaranteed if $\nabla f(\bar{x}) = 0$ under an injectivity assumption on $\tau'(\xi)$.

For optimization problems on manifolds, the interplay of global Hessian, gradient, and curvature as displayed in (5.1) is gathered in the important concept of the *Riemannian Hessian*. This is thoroughly discussed in [2]; see, in particular, section 6 therein. Similar to its role in smooth optimization in linear spaces, the positivity of the Riemannian Hessian ensures local (Riemannian) convexity and hence contractivity of many Riemannian optimization methods; see the book [1]. For manifolds of low-rank matrices, the curvature terms in this Hessian have been obtained in other works. Specifically, [18, Proposition 2.2] features a formula that makes the Kronecker-type interplay between \bar{X}^+ and $\nabla f(\bar{X})$ in the curvature (see Lemma 3.1) clearly visible, albeit for a special case of the cost function (3.10) related to matrix completion (with $\mathbf{A} = \mathbf{P}_\Omega$ being a projection on given entries Ω). In [9], the curvature term in the Riemannian Hessian is explicitly neglected to derive Riemannian Gauss–Newton-type methods on low-rank tensor manifolds.

In the works [11] and [5], convergence of the ALS method for low-rank tensor approximation has been investigated using alternative techniques. In particular, questions on cluster points and global convergence are addressed. Also, examples for

sublinear, linear, and superlinear convergence are presented in [5]. The references [21, 19, 17, 6] specifically deal with the convergence of the higher-order power method.

Finally, we mention the work [15], in which global convergence of a related method (called scaled alternating steepest descent) for matrix completion is investigated.

Acknowledgment. A main part of this work was developed while the authors enjoyed the hospitality of the Hausdorff Institute for Mathematics in Bonn during the trimester program on multiscale problems.

REFERENCES

- [1] P.-A. ABSIL, R. MAHONY, AND R. SEPULCHRE, *Optimization Algorithms on Matrix Manifolds*, Princeton University Press, Princeton, NJ, 2008, <https://doi.org/10.1515/9781400830244>.
- [2] P.-A. ABSIL, J. TRUMPF, R. MAHONY, AND B. ANDREWS, *All Roads Lead to Newton: Feasible Second-Order Methods for Equality-Constrained Optimization*, Technical report UCL-INMA-2009.024, Catholic University of Louvain, Louvain-la-Neuve, Brussels, 2009.
- [3] F. L. BAUER, *Das Verfahren der Treppeniteration und verwandte Verfahren zur Lösung algebraischer Eigenwertprobleme*, *Z. Angew. Math. Phys.*, 8 (1957), pp. 214–235.
- [4] J. DIEUDONNÉ, *Foundations of Modern Analysis*, Academic Press, New York, 1969.
- [5] M. ESPIG, W. HACKBUSCH, AND A. KHACHATRYAN, *On the Convergence of Alternating Least Squares Optimisation in Tensor Format Representations*, <https://arxiv.org/abs/1506.00062> (2015).
- [6] M. ESPIG AND A. KHACHATRYAN, *Convergence of Alternating Least Squares Optimisation for Rank-One Approximation to High Order Sensors*, <https://arxiv.org/abs/1503.05431> (2015).
- [7] G. H. GOLUB AND V. PEREYRA, *The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate*, *SIAM J. Numer. Anal.*, 10 (1973), pp. 413–432.
- [8] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 4th ed., Johns Hopkins University Press, Baltimore, MD, 2013.
- [9] D. KRESSNER, M. STEINLECHNER, AND B. VANDEREYCKEN, *Preconditioned low-rank Riemannian optimization for linear systems with tensor product structure*, *SIAM J. Sci. Comput.*, 38 (2016), pp. A2018–A2044.
- [10] J. M. LEE, *Introduction to Smooth Manifolds*, Springer, New York, 2003.
- [11] M. J. MOHLENKAMP, *Musings on multilinear fitting*, *Linear Algebra Appl.*, 438 (2013), pp. 834–852.
- [12] J. M. ORTEGA AND W. C. RHEINOLDT, *Iterative Solution of Nonlinear Equations in Several Variables*, Academic Press, New York, 1970.
- [13] T. ROHWEDDER AND A. USCHMAJEV, *On local convergence of alternating schemes for optimization of convex problems in the tensor train format*, *SIAM J. Numer. Anal.*, 51 (2013), pp. 1134–1162.
- [14] R. SCHNEIDER AND A. USCHMAJEV, *Convergence results for projected line-search methods on varieties of low-rank matrices via Lojasiewicz inequality*, *SIAM J. Optim.*, 25 (2015), pp. 622–646.
- [15] J. TANNER AND K. WEI, *Low rank matrix completion by alternating steepest descent methods*, *Appl. Comput. Harmon. Anal.*, 40 (2016), pp. 417–429.
- [16] A. USCHMAJEV, *Local convergence of the alternating least squares algorithm for canonical tensor approximation*, *SIAM J. Matrix Anal. Appl.*, 33 (2012), pp. 639–652.
- [17] A. USCHMAJEV, *A new convergence proof for the higher-order power method and generalizations*, *Pac. J. Optim.*, 11 (2015), pp. 309–321.
- [18] B. VANDEREYCKEN, *Low-rank matrix completion by Riemannian optimization*, *SIAM J. Optim.*, 23 (2013), pp. 1214–1236.
- [19] L. WANG AND M. T. CHU, *On the global convergence of the alternating least squares method for rank-one approximation to generic tensors*, *SIAM J. Matrix Anal. Appl.*, 35 (2014), pp. 1058–1072.
- [20] J. XU AND L. ZIKATANOV, *The method of alternating projections and the method of subspace corrections in Hilbert space*, *J. Amer. Math. Soc.*, 15 (2002), pp. 573–597.
- [21] T. ZHANG AND G. H. GOLUB, *Rank-one approximation to high order tensors*, *SIAM J. Matrix Anal. Appl.*, 23 (2001), pp. 534–550.