

**Max-Planck-Institut
für Mathematik
in den Naturwissenschaften
Leipzig**

**Local convergence of alternating
low-rank optimization methods with
overrelaxation**

(revised version: June 2022)

by

*Ivan V. Oseledets, Maxim V. Rakhuba, and
André Uschmajew*

Preprint no.: 29

2021



Local convergence of alternating low-rank optimization methods with overrelaxation

Ivan V. Oseledets* Maxim V. Rakhuba† André Uschmajew‡

Abstract

The local convergence of alternating optimization methods with overrelaxation for low-rank matrix and tensor problems is established. The analysis is based on the linearization of the method which takes the form of an SOR iteration for a positive semidefinite Hessian and can be studied in the corresponding quotient geometry of equivalent low-rank representations. In the matrix case, the optimal relaxation parameter for accelerating the local convergence can be determined from the convergence rate of the standard method. This result relies on a version of Young's SOR theorem for positive semidefinite 2×2 block systems.

1 Introduction

We consider a low-rank matrix optimization problem of the form

$$\min_{\text{rank}(X) \leq k} f(X), \quad (1.1)$$

where $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is a smooth function on the space of real $m \times n$ matrices. It will be mostly assumed that f is strongly convex. This generic problem appears in a large number of applications, where low-rank matrices serve as nonlinear model classes, such as in matrix recovery, or are employed for reducing numerical complexity when dealing with large-scale matrices.

Since the constraint set admits the explicit parametrization $X = UV^\top$, the problem can be rewritten as

$$\min_{U \in \mathbb{R}^{m \times k}, V \in \mathbb{R}^{n \times k}} F(U, V) = f(UV^\top). \quad (1.2)$$

One of the basic methods for solving (1.2) is the alternating optimization (AO) method, which optimizes the factor matrices U and V in an alternating manner. Conceptually, ignoring the question of unique solvability of subproblems, the method looks as follows:

$$\begin{aligned} U_{\ell+1} &= \operatorname{argmin}_U F(U, V_\ell), \\ V_{\ell+1} &= \operatorname{argmin}_V F(U_{\ell+1}, V). \end{aligned} \quad (1.3)$$

While this is certainly a standard approach from the viewpoint of nonlinear optimization, where such a scheme is also known as nonlinear Gauss-Seidel method, it is worth emphasizing that the special structure of low-rank problems is particularly amenable to it. This is due to the bilinearity of the parametrization UV^\top , which turns the subproblems of (1.3) into optimization problems for the same initial function f , but on lower dimensional linear subspaces. Therefore, when f is a quadratic function, this method is called the alternating least squares (ALS) method.

*Skolkovo Institute of Science and Technology, Moscow, 121205 Russia

†HSE University, Moscow, 109028 Russia

‡Max Planck Institute for Mathematics in the Sciences, 04103 Leipzig, Germany

While the study of global convergence of the AO method (1.3) is usually difficult, its local convergence properties are well-understood [16, 14, 13]. The local analysis is based on the fact that the linearized version of the method at a critical point (U_*, V_*) takes the form of a block Gauss-Seidel method for the Hessian $\nabla^2 F(U_*, V_*)$. Due to the intrinsic overparametrization of rank- k matrices by the representation UV^\top , the Hessian is at best positive semidefinite, but never positive definite. The Gauss-Seidel error iteration matrix is not contractive on the null space of the Hessian, but it turns out that this problem can be overcome by passing to the corresponding quotient geometry of equivalent low-rank representations $X = UV^\top$. This is possible thanks to an invariance of the AO method under changes of the representation. In fact, this invariance allows one to regard the method (1.3) as a well-defined iteration on the set of rank- k matrices.

In this work, we consider the acceleration of the local convergence of the AO method (1.3) by means of overrelaxation. This is a classic idea in nonlinear optimization; see, e.g., [15, 9, 8, 3] to mention some early works. Several variants of such acceleration have been proposed for low-rank matrix problems, e.g., for matrix completion [21, 19]. The basic overrelaxation method that we consider has already been proposed in [1] for the more general low-rank tensor train format and in the matrix case reads as follows:

$$\begin{aligned} U_{\ell+1} &= (1 - \omega)U_\ell + \omega \operatorname{argmin}_U F(U, V_\ell), \\ V_{\ell+1} &= (1 - \omega)V_\ell + \omega \operatorname{argmin}_V F(U_{\ell+1}, V). \end{aligned} \tag{1.4}$$

Here $\omega > 0$ is a relaxation parameter, which sometimes is also called a shift. It can be observed numerically that a suitable choice of the shift significantly improves the convergence speed.

Our goal is to study the local convergence of this iteration for low-rank optimization in a similar spirit as for the plain AO method (1.3), which corresponds to the case $\omega = 1$. This will be done in the main section 2. The linearization of (1.4) at a critical point (U_*, V_*) leads to a 2×2 block SOR method for the Hessian $\nabla^2 F(U_*, V_*)$. Using the fact that for $0 < \omega < 2$ and a positive semidefinite Hessian with positive definite block diagonal the SOR error iteration is contractive on any subspace complementary to the null space of the Hessian, we obtain local convergence results for this range of ω . This result is stated in Theorem 2.3.

It is then natural to ask for the optimal shift ω achieving the fastest local convergence rate, which requires to minimize the spectral radius of the SOR error iteration matrix. For positive definite 2×2 block systems this can be achieved using a well-known theorem of Young. It is however possible to adjust the arguments to the positive semidefinite case, as will be done in Lemma 2.5. This yields the expected, yet not entirely trivial, formula for the asymptotically optimal shift in terms of the convergence rate of the standard AO method with $\omega = 1$. The result is stated in Theorem 2.7. In practice this means that the optimal shift can be estimated adaptively and at practically zero cost from the observed convergence rate of the standard method.

Of course, overrelaxation can also be applied to AO methods for low-rank tensor optimization. In section 3 we focus on the low-rank tensor train format as in [1]. Like low-rank matrix factorization, the tensor train decomposition is subject to an intrinsic overparametrization which can be described by a simple group action in parameter space, but leads to formally semidefinite Hessians in critical points. By passing to suitable quotient spaces, the local convergence of the method can be established in essentially the same way as for low-rank matrices (Theorem 3.1). However, a main difference to the matrix case is that the formula for the optimal shift cannot be made rigorous under reasonable assumptions, although it can still serve as a useful heuristic.

In the last section 4 we report on some numerical experiments that illustrate the advantage of using shifts in low-rank AO methods, and validate our theoretical findings regarding the optimal shift in the matrix case. We also demonstrate the adaptive procedure for choosing an almost optimal shift based on the observed convergence rate of the standard method.

2 Alternating optimization with relaxation for low-rank matrices

In this main section of the paper we first formalize the basic AO iteration (1.3) for low-rank matrix problems and recall some of its basic properties. We then proceed to the method with overrelaxation, establish its local convergence and determine the optimal shift parameter.

2.1 Standard AO method

Consider the scheme (1.3) and assume f to be strongly convex. Then the first argmin is uniquely defined if $\text{rank}(V_\ell) = k$, since it corresponds to minimizing the strongly convex function f on a linear subspace of $\mathbb{R}^{m \times n}$ which is the image of the injective linear map $U \mapsto UV_\ell^\top$. Likewise, the second argmin is well-defined if $\text{rank}(U_{\ell+1}) = k$ and returns the unique minimum of f on the linear subspace $V \mapsto U_{\ell+1}V^\top$. Therefore, in some open and dense subsets both argmins define smooth maps \hat{S}_1 and \hat{S}_2 , respectively, such that

$$U_{\ell+1} = \hat{S}_1(V_\ell), \quad V_{\ell+1} = \hat{S}_2(U_{\ell+1}). \quad (2.1)$$

One full update of the method then takes the form of a fixed point iteration

$$\begin{pmatrix} U_{\ell+1} \\ V_{\ell+1} \end{pmatrix} = S \begin{pmatrix} U_\ell \\ V_\ell \end{pmatrix} := \begin{pmatrix} \hat{S}_1(V_\ell) \\ \hat{S}_2(\hat{S}_1(V_\ell)) \end{pmatrix}. \quad (2.2)$$

The map S is well-defined and smooth on any open subset of

$$\mathcal{D} = \{(U, V) \in \mathbb{R}^{m \times k} \times \mathbb{R}^{n \times k} : V \text{ and } \hat{S}_1(V) \text{ have full column rank } k\}.$$

In particular, any critical point (U_*, V_*) of F in (1.2) for which U_* and V_* have full column rank belongs to \mathcal{D} and is a fixed point of S . To see this, note that $U \mapsto F(U, V_*) = f(UV_*^\top)$ is strongly convex since $U \mapsto UV_*^\top$ is injective. Since U_* is a critical point of that function, it is the global minimum and hence $\hat{S}_1(V_*) = U_*$. The argument for V_* is the same. Such a critical point of F possesses an open neighborhood in \mathcal{D} in which S is well-defined and smooth. Conversely, any fixed point $(U_*, V_*) \in \mathcal{D}$ of S must be a critical point of F since it implies that U_* is the global minimum of $U \mapsto F(U, V_*)$ and V_* is the global minimum of $V \mapsto F(U_*, V)$. Hence the partial gradients $\nabla_U F(U_*, V_*)$ and $\nabla_V F(U_*, V_*)$ are both zero.

By passing from the initial constrained problem (1.1) to the factorized problem (1.2), we formally introduced an ambiguity arising from the fact that the factorization $X = UV^\top$ of a rank- k matrix is not unique. In particular, $X = UAA^{-1}V^\top$ for any invertible $k \times k$ matrix A so that the function F has level sets of at least dimension k^2 (when U, V have full column rank). Therefore, a fixed point $(U_*, V_*) \in \mathcal{D}$ of S is never locally unique. However, this issue of non-uniqueness is only a formal one since one is ultimately interested in the sequence of generated matrices $X_\ell = U_\ell V_\ell^\top$. Assuming $\text{rank}(X_\ell) = k$ for all ℓ , this sequence is not affected by any reparametrization $(U_\ell, V_\ell) \rightarrow (U_\ell A_\ell, V_\ell A_\ell^{-T})$ with invertible matrices A_ℓ during the iteration. This is due to the invariance properties

$$\begin{aligned} \hat{S}_1(VA^{-T}) &= \hat{S}_1(V)A, \\ \hat{S}_2(UA) &= \hat{S}_2(U)A^{-T} \end{aligned} \quad (2.3)$$

of the maps \hat{S}_1 and \hat{S}_2 , which hold whenever U and V have full column rank. To see this, let $U_+ = \hat{S}_1(V)$ and $\hat{U}_+ = \hat{S}_1(VA^{-T})$. Then by construction U_+V^\top and $\hat{U}_+A^{-1}V^\top$ are the unique minimizers of f on the linear subspaces $\{UV^\top : U \in \mathbb{R}^{m \times k}\}$ and $\{UA^{-1}V^\top : U \in \mathbb{R}^{m \times k}\}$, respectively. Obviously both spaces are equal, hence $U_+V^\top = \hat{U}_+A^{-1}V^\top$. Since V has full column rank, we obtain $U_+ = \hat{U}_+A^{-1}$, the first identity in (2.3). The argument for \hat{S}_2 is analogous.

The above invariance of the AO method allows us to interpret it as a method

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

in the full matrix space $\mathbb{R}^{m \times n}$, or more precisely on the subset of matrices of rank at most k . This viewpoint has been taken in [13] and will be helpful in this work, too. From an algorithmic perspective, the AO viewpoint (2.2) is more useful since it operates on the smaller matrices U and V instead of the full matrix X . Furthermore, the invariance with respect to the described change of parametrization allows for a robust implementation of the AO method by orthogonalizing the columns of U_ℓ and V_ℓ after every partial update, without affecting the generated sequence X_ℓ of matrices. This method is a special case of Algorithm 1 below with $\omega = 1$.

2.2 Overrelaxation

Instead of (2.1) we now consider the more general update rule with a shift,

$$\begin{aligned} U_{\ell+1} &= (1 - \omega)U_\ell + \omega \hat{S}_1(V_\ell), \\ V_{\ell+1} &= (1 - \omega)V_\ell + \omega \hat{S}_2(U_{\ell+1}), \end{aligned} \tag{2.4}$$

which corresponds to (1.4). For $\omega = 1$, this iteration equals the standard AO method (2.1). By defining the map

$$S_\omega \begin{pmatrix} U \\ V \end{pmatrix} = (1 - \omega) \begin{pmatrix} U \\ V \end{pmatrix} + \omega \begin{pmatrix} \hat{S}_1(V) \\ \hat{S}_2((1 - \omega)U + \omega \hat{S}_1(V)) \end{pmatrix}, \tag{2.5}$$

we can write (2.4) as a nonlinear fixed point iteration

$$\begin{pmatrix} U_{\ell+1} \\ V_{\ell+1} \end{pmatrix} = S_\omega \begin{pmatrix} U_\ell \\ V_\ell \end{pmatrix}. \tag{2.6}$$

The map S_ω is well-defined and smooth on any open subset of

$$\mathcal{D}_\omega = \{(U, V) \in \mathbb{R}^{m \times k} \times \mathbb{R}^{n \times k} : V \text{ and } (1 - \omega)U + \omega \hat{S}_1(V) \text{ have full column rank}\}.$$

Note that $(U_*, V_*) \in \mathcal{D}_\omega$ is a fixed point of S_ω if and only if $(U_*, V_*) \in \mathcal{D}$ and (U_*, V_*) is a fixed point of S . In particular, any critical point (U_*, V_*) of F in (1.2) such that U_* and V_* have full column rank belongs to \mathcal{D}_ω and is a fixed point of S_ω . Moreover, any such critical point possesses an open neighborhood in \mathcal{D}_ω such that S_ω is well-defined and smooth on this neighborhood. Again, the converse is also true, that is, a fixed point $(U_*, V_*) \in \mathcal{D}_\omega$ of S_ω is a critical point of F .

The iteration (2.6) exhibits the same invariance under changes of representation as the standard AO method. Let $(U, V) \in \mathcal{D}_\omega$ and $(U_+, V_+) = S_\omega(U, V)$, then from (2.5) and (2.3) one verifies

$$S_\omega \begin{pmatrix} UA \\ VA^{-T} \end{pmatrix} = \begin{pmatrix} U_+A \\ V_+A^{-T} \end{pmatrix}. \tag{2.7}$$

Therefore, the generated sequence $X_\ell = U_\ell V_\ell^\top$ is essentially (if $\text{rank}(X_\ell) = k$ for all ℓ) invariant under changes of representation during the iteration. In particular, QR decomposition can be used in numerical implementation for keeping the argmin problems well-conditioned. The resulting method is denoted in Algorithm 1.

Since the goal of this work is a local convergence analysis of the fixed point iteration (2.6), it is important to observe that the invariance under the group action also carries over to the asymptotic linear convergence rate of the method, which depends on the eigenvalues of the

Algorithm 1: Low-rank AO with overrelaxation and QR

Data: $V_0 \in \mathbb{R}^{n \times k}$, relaxation parameter ω
for $\ell = 0, 1, 2, \dots$ **do**
 $U \leftarrow \operatorname{argmin}_{\hat{U} \in \mathbb{R}^{m \times k}} F(\hat{U}, V_\ell)$
 $U \leftarrow (1 - \omega)U_\ell + \omega U, \quad U = Q_1 R_1$
 $V \leftarrow \operatorname{argmin}_{\hat{V} \in \mathbb{R}^{n \times k}} F(Q_1, \hat{V})$
 $V \leftarrow (1 - \omega)V_\ell R_1^\top + \omega V, \quad V = Q_2 R_2$
 $U_{\ell+1} := Q_1 R_2^\top, \quad V_{\ell+1} := Q_2$
end

derivative $S'_\omega(U_*, V_*)$ at a fixed point $(U_*, V_*) \in \mathcal{D}_\omega$. To see this invariance it is convenient to introduce the corresponding group action θ_A of $\operatorname{GL}(k)$ acting on $\mathbb{R}^{m \times k} \times \mathbb{R}^{n \times k}$ via

$$A \mapsto \theta_A \cdot \begin{pmatrix} U \\ V \end{pmatrix} = \begin{pmatrix} UA \\ VA^{-T} \end{pmatrix}. \quad (2.8)$$

In this notation (2.7) reads

$$S_\omega \left(\theta_A \cdot \begin{pmatrix} U \\ V \end{pmatrix} \right) = \theta_A \cdot S_\omega \left(\begin{pmatrix} U \\ V \end{pmatrix} \right). \quad (2.9)$$

For fixed A , (2.8) defines an invertible linear map θ_A on $\mathbb{R}^{m \times k} \times \mathbb{R}^{n \times k}$ with $\theta_A^{-1} = \theta_{A^{-1}}$. Differentiating both sides of (2.9) it then follows that

$$S'_\omega \left(\begin{pmatrix} U_* A \\ V_* A^{-T} \end{pmatrix} \right) = \theta_A \cdot S'_\omega \left(\begin{pmatrix} U_* \\ V_* \end{pmatrix} \right) \cdot \theta_A^{-1}. \quad (2.10)$$

This shows that $S'_\omega(U_* A, V_* A^{-T})$ has the same eigenvalues as $S'_\omega(U_*, V_*)$ and allows us to study the local convergence rate of the iteration (2.6) at any particular fixed point (U_*, V_*) .

As for the standard AO method, the invariance property allows for an interpretation of the method (2.6) as an iteration

$$X_{\ell+1} = \mathbf{S}_\omega(X_\ell) \quad (2.11)$$

on the manifold

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

where $\mathbf{S}_\omega : \mathcal{O} \subseteq \mathcal{M}_k \rightarrow \mathbb{R}^{m \times n}$ is defined through

$$\mathbf{S}_\omega(X) = \tau(S_\omega(U, V)), \quad X = UV^\top, \quad (2.12)$$

with the map

$$\tau(U, V) = UV^\top.$$

Here the domain of definition \mathcal{O} of \mathbf{S}_ω should be contained in the image of $\mathcal{D} \cap \mathcal{D}_\omega$ under τ . In particular, let $(U_*, V_*) \in \mathcal{D}$ be a fixed point of the map S (the standard AO method), that is, a critical point of F . Then \mathbf{S}_ω is well-defined and smooth in some open neighborhood $\mathcal{O} \subseteq \mathcal{M}_k$ of $X_* = U_* V_*^\top$ and X_* is a fixed point of \mathbf{S}_ω . This manifold viewpoint will be useful in the local convergence analysis conducted in the next section.

2.3 Local convergence

Let $X_* = U_* V_*^\top \in \mathcal{M}_k$ be a fixed point of \mathbf{S}_ω . Then \mathbf{S}_ω locally maps to \mathcal{M}_k and thus the derivative $\mathbf{S}'_\omega(X_*)$ maps the tangent space $T_{X_*} \mathcal{M}_k$ to itself. This provides the following local convergence criterion.

Proposition 2.1. *Let f be strongly convex and (U_*, V_*) be a critical point of F in (1.2) with U_*, V_* having full column rank k . Then (U_*, V_*) is a fixed point of S_ω and $X_* = U_* V_*^\top \in \mathcal{M}_k$ is a fixed point of \mathbf{S}_ω . Let $\mathbf{S}'_\omega(X_*) : T_{X_*} \mathcal{M}_k \rightarrow \mathbb{R}^{m \times n}$ denote the derivative of \mathbf{S}_ω at X_* , and $\mathbf{P}_{X_*} : \mathbb{R}^{m \times n} \rightarrow T_{X_*} \mathcal{M}_k$ the tangent space projection. If for the spectral radius*

$$\rho_\omega = \rho(\mathbf{P}_{X_*} \mathbf{S}'_\omega(X_*)) < 1, \quad (2.13)$$

then for $X_0 = U_0 V_0^\top$ close enough to X_* the iterates $X_\ell = U_\ell V_\ell^\top$ generated by Algorithm 1 converge to X_* at an asymptotic linear rate ρ_ω .

To study the convergence criterion in more detail we investigate $\mathbf{S}'_\omega(X_*)$. For this we repeat some well-known computations. We first consider the map S_ω in parameter space. From (2.5), its derivative at (U_*, V_*) takes the form of a block matrix

$$S'_\omega \begin{pmatrix} U_* \\ V_* \end{pmatrix} = (1 - \omega) \begin{pmatrix} I & 0 \\ 0 & I \end{pmatrix} + \omega \begin{pmatrix} 0 & \hat{S}'_1(V_*) \\ (1 - \omega)\hat{S}'_2(U_*) & \omega\hat{S}'_2(U_*)\hat{S}'_1(V_*) \end{pmatrix},$$

where we have used that $(1 - \omega)U_* + \omega\hat{S}'_1(V_*) = U_*$. Setting

$$L = \begin{pmatrix} 0 & 0 \\ \hat{S}'_2(U_*) & 0 \end{pmatrix}, \quad R = \begin{pmatrix} 0 & \hat{S}'_1(V_*) \\ 0 & 0 \end{pmatrix}$$

one then verifies the identity

$$S'_\omega \begin{pmatrix} U_* \\ V_* \end{pmatrix} = (I - \omega L)^{-1} [(1 - \omega)I + \omega R]. \quad (2.14)$$

This linear operator can be interpreted as a block SOR error iteration matrix for the Hessian

$$H = \begin{pmatrix} \nabla_{UU}F(U_*, V_*) & \nabla_{UV}F(U_*, V_*) \\ \nabla_{VU}F(U_*, V_*) & \nabla_{VV}F(U_*, V_*) \end{pmatrix}$$

of F at (U_*, V_*) , written in 2×2 block form. To see this, consider the usual decomposition

$$H = D + E + E^\top$$

with

$$D = \begin{pmatrix} \nabla_{UU}F(U_*, V_*) & 0 \\ 0 & \nabla_{VV}F(U_*, V_*) \end{pmatrix}, \quad E = \begin{pmatrix} 0 & 0 \\ \nabla_{VU}F(U_*, V_*) & 0 \end{pmatrix}.$$

Assuming that the block diagonal part D is invertible and differentiating the equations

$$\nabla_U F(\hat{S}_1(V), V) = 0, \quad \nabla_V F(U, \hat{S}_2(U)) = 0 \quad (2.15)$$

(which implicitly define \hat{S}_1 and \hat{S}_2), one finds that

$$\hat{S}'_1(V_*) = -[\nabla_{UU}F(U_*, V_*)]^{-1} \nabla_{UV}F(U_*, V_*)$$

and

$$\hat{S}'_2(U_*) = -[\nabla_{VV}F(U_*, V_*)]^{-1} \nabla_{VU}F(U_*, V_*).$$

In other words, $L = -D^{-1}E$, $R = -D^{-1}E^\top$, and

$$L + R = I - D^{-1}H.$$

Using the expressions for L and R in (2.14), one obtains the alternative formula

$$S'_\omega \begin{pmatrix} U_* \\ V_* \end{pmatrix} = T_\omega := I - N_\omega^{-1}H, \quad N_\omega = \frac{1}{\omega}D + E. \quad (2.16)$$

We see from (2.16) that $S'_\omega(U_*, V_*)$ equals the error iteration matrix T_ω for the two-block SOR method for H . It is well-known that T_ω has spectral radius less than one if $0 < \omega < 2$ and H is positive definite. However, the latter is never the case here. Since ∇F is constantly zero on the orbit $\theta_A \cdot (U_*, V_*)$ (this follows from the chain rule by differentiating $F = F \circ \theta_A$ for fixed A), the Hessian at critical points $(U_*, V_*) \in \mathcal{D}$ has at least a k^2 -dimensional kernel $\ker H$ containing the tangent space to the orbit. On $\ker H$ the matrix T_ω acts as identity. However, if $0 < \omega < 2$, D is positive definite and H is at least positive semidefinite, then by classic results it still holds that T_ω is a contraction on any invariant subspace complementary to $\ker H$; see, e.g., [20, Sec. 3] or [5, Corollary 2.1]. Specifically, as follows from [20], the space

$$\mathcal{W}_\omega = N_\omega^{-1}(\ker H)^\perp \quad (2.17)$$

is an invariant subspace of T_ω , which splits the parameter space into a direct sum¹

$$\mathbb{R}^{m \times k} \times \mathbb{R}^{n \times k} = \ker H \oplus \mathcal{W}_\omega, \quad (2.18)$$

and T_ω is a contraction on \mathcal{W}_ω .

At this point we can exploit that we are actually interested in the convergence of the products $X_\ell = U_\ell V_\ell^\top$. Under the assumption that $\ker H$ equals the tangent space to the orbit $\theta_A \cdot (U_*, V_*)$, any complementary subspace, such as \mathcal{W}_ω , satisfies the properties of a so-called horizontal space for the quotient manifold structure of \mathcal{M}_k . For us, this means the following.

Proposition 2.2. *Assume $X_* = U_* V_*^\top$ has rank k , H is positive semidefinite, $\dim(\ker H) = k^2$ and a decomposition (2.18) holds. Then the map $\tau(U, V) = UV^\top$ is a local diffeomorphism between a (relative) neighborhood of (U_*, V_*) in $(U_*, V_*) + \mathcal{W}_\omega$ and a neighborhood of X_* in the embedded submanifold $\mathcal{M}_k \subseteq \mathbb{R}^{m \times n}$.*

Proof. The proof can be given without particular reference to quotient manifolds, but assuming knowledge that \mathcal{M}_k is a smooth embedded submanifold of dimension $mk + nk - k^2$ [7, Ex. 8.14], and τ is a local submersion on \mathcal{M}_k in a neighborhood of (U_*, V_*) (it is not difficult to verify that $\tau'(U_*, V_*)$ has rank $mk + nk - k^2$). Then since τ is constant on the θ_A -orbit of (U_*, V_*) , its derivative vanishes on the tangent space to that orbit at (U_*, V_*) , which is of dimension k^2 . We already noted that $\ker H$ contains that tangent space, so if $\dim(\ker H) = k^2$, then $\tau'(U_*, V_*)$ vanishes on $\ker H$. Hence, due to (2.18), $\tau'(U_*, V_*)$ must be a bijection between \mathcal{W}_ω and the tangent space $T_{X_*} \mathcal{M}_k$. The assertion follows by the inverse function theorem. \square

We are now in the position to formulate a local convergence result for the iteration (2.11).

Theorem 2.3. *Let f be strongly convex and (U_*, V_*) be a critical point of F in (1.2) with U_*, V_* having full column rank k . Assume that the Hessian $H = \nabla^2 F(U_*, V_*)$ is positive semidefinite and $\dim(\ker H) = k^2$. Fix $0 < \omega < 2$. Then for $X_0 = U_0 V_0^\top$ close enough (this may depend on ω) to $X_* = U_* V_*^\top$ Algorithm 1 is well-defined and the iterates $X_\ell = U_\ell V_\ell^\top$ converge to X_* at an asymptotic linear rate*

$$\rho_\omega = \limsup_{\ell \rightarrow \infty} \|X_\ell - X_*\|^{1/\ell} < 1,$$

where ρ_ω is the spectral radius of T_ω on \mathcal{W}_ω .

¹Obviously, \mathcal{W}_ω is an invariant subspace of T_ω and has the correct dimension. To see that it is complementary to $\ker H$, note that any $x \in \mathcal{W}_\omega$ satisfies $N_\omega x \in (\ker H)^\perp$. If $x \in \ker H$, one verifies $0 = \langle x, N_\omega x \rangle = (\frac{1}{\omega} - \frac{1}{2}) \langle x, Dx \rangle$, which under the given assumptions implies $x = 0$.

The convergence rate ρ_ω is determined in the next section. We stated the above result separately because its proof can be easily generalized to alternating optimization methods for low-rank tensor formats that admit a similar invariance under a group action. This will be outlined for the tensor train format in section 3.

Proof. Since V_* has full column rank, the linear map $U \mapsto UV_*^\top$ is injective and hence the restricted map $U \mapsto F(U, V_*) = f(UV_*^\top)$ is strongly convex. Therefore, $\nabla_{UU}F(U_*, V_*)$ is positive definite. Likewise, $\nabla_{VV}F(U_*, V_*)$ is positive definite, so that the block diagonal part D of H is positive definite. As a result, the decomposition (2.18) of the parameter space applies and T_ω is a contraction on its invariant subspace \mathcal{W}_ω . In a neighborhood of X_* in \mathcal{M}_k the map \mathbf{S}_ω in (2.12) can be written as

$$\mathbf{S}_\omega = \tau \circ S_\omega \circ \tau^{-1},$$

where we have restricted τ to the affine subspace $(U_*, V_*) + \mathcal{W}_\omega$. Therefore, by chain rule,

$$\mathbf{S}'_\omega(X_*) = [\tau'(U_*, V_*)] \circ T_\omega \circ [\tau'(U_*, V_*)]^{-1}. \quad (2.19)$$

By Proposition 2.2, the derivative $\tau'(U_*, V_*)$ is an isomorphism between \mathcal{W}_ω and $T_{X_*}\mathcal{M}_k$. Due to (2.19), this implies that the convergence criterion (2.13) in Proposition 2.1 is satisfied. \square

Remark 2.4. The assumptions that H is positive semidefinite and $\dim(\ker H) = k^2$ already imply by themselves that D is positive definite. Indeed, as noted in the proof of Proposition 2.2, $\dim(\ker H) = k^2$ means that $\ker H$ equals the tangent space to the θ_A -orbit of (U_*, V_*) , which however does not contain elements of the form $(U, 0)$ or $(0, V)$ (this can be seen from (2.8)). This allows to define a nonlinear SOR process in a neighborhood of such a critical point based on the implicit definitions (2.15) of \hat{S}_1 and \hat{S}_2 even when f is not strongly convex; cf. [8, Thm. 10.3.5].

To get a better intuition for the assumptions in the theorem, it is useful to write the Hessian as a bilinear form

$$\nabla^2 F(U_*, V_*)[h, h] = \langle \tau'(U_*, V_*)[h], \nabla^2 f(X_*) \cdot \tau'(U_*, V_*)[h] \rangle + \langle \nabla f(X_*), \tau''(U_*, V_*)[h, h] \rangle,$$

where $h = (\delta U, \delta V)$. If f is strictly convex, then the first term is nonnegative and equal to zero if and only if $\tau'(U_*, V_*)[h] = 0$, that is, if h is in the tangent space to the θ_A -orbit at (U_*, V_*) . Thus, an important situation in which the assumptions of the theorem are satisfied is when $\nabla f(X_*) = 0$, that is, when $X_* = U_* V_*^\top$ is a global minimum of f . In Section 4.1 we conduct some numerical experiments for a matrix completion problem (4.1) admitting such a global minimum $X_* = U_* V_*^\top$ with $\nabla f(X_*) = 0$. In this application, however, f is only convex, but not strictly convex. Then in order to satisfy the assumptions of the theorem at X_* one would need that the tangent space $T_{X_*}\mathcal{M}_k$ (the image of $\tau'(U_*, V_*)$) does not intersect the null space of $\nabla^2 f(X_*) = P_\Omega$, but we will not investigate this condition in detail.

2.4 Asymptotically optimal relaxation

It is well-known that under certain assumptions the relaxation parameter ω in the linear SOR method can be optimized using a theorem of Young; see, e.g., [22, Sec. 6.2] or [2, Sec. 4.6.2]. This theory is usually presented for positive definite systems. However, for 2×2 block systems it is possible to adjust the arguments to the positive semidefinite case.

Lemma 2.5. *Let $H = D + E + E^\top \in \mathbb{R}^{p \times p}$ be a positive semidefinite 2×2 block matrix with positive definite block diagonal D and such that $\frac{1}{\omega}D + E$ is invertible for any $0 < \omega < 2$. Assume $q = \dim(\ker H) < p/2$. Let $\sigma(I - D^{-1}H)$ denote the spectrum of $I - D^{-1}H$, then*

$$\beta := \max\{|\mu| : \mu \in \sigma(I - D^{-1}H) \setminus \{\pm 1\}\} < 1.$$

The matrix $T_\omega = I - N_\omega^{-1}H$, where $N_\omega = \frac{1}{\omega}D + E$, induces a decomposition (2.18) into two invariant subspaces and the spectral radius ρ_ω of T_ω on \mathcal{W}_ω equals

$$\rho_\omega = \begin{cases} 1 - \omega + \frac{1}{2}\omega^2\beta^2 + \omega\beta\sqrt{1 - \omega + \frac{1}{4}\omega^2\beta^2} & \text{if } 0 < \omega \leq \omega_{\text{opt}}, \\ \omega - 1 & \text{if } \omega_{\text{opt}} \leq \omega < 2, \end{cases}$$

where

$$\omega_{\text{opt}} = \frac{2}{1 + \sqrt{1 - \beta^2}} > 1. \quad (2.20)$$

The value of ρ_ω is minimal for $\omega = \omega_{\text{opt}}$. It holds that $\beta^2 = \rho_1$ is the spectral radius for the standard AO method with $\omega = 1$ (on its invariant subspace \mathcal{W}_1).

Proof. The decomposition (2.18) into invariant subspaces has already been verified (see Footnote 1). We follow the arguments in the proof of [2, Thm. 4.27]. There it is shown that the eigenvalues μ of $I - D^{-1}H$ and λ of T_ω are related via

$$\lambda = 1 - \omega + \frac{1}{2}\omega^2\mu^2 \pm \omega\mu\sqrt{1 - \omega + \frac{1}{4}\omega^2\mu^2}. \quad (2.21)$$

Indeed, note that under the given assumptions $I - D^{-1}H = -D^{-1}(E + E^\top)$ is a two-cyclic matrix and has only real eigenvalues, of which the nonzero ones come in pairs $\pm\mu$. By (2.21), both μ and $-\mu$ create a same pair of eigenvalues λ . Eigenvalues $\mu = 1$ of $I - D^{-1}H$ must belong to eigenvectors in $\ker H$. Therefore, $\mu = 1$ and $\mu = -1$ both have multiplicity q . They yield eigenvalues $\lambda = 1$ and $\lambda = (1 - \omega)^2$ of T_ω . Since eigenvectors of T_ω with $\lambda = 1$ must belong to $\ker H$, we conclude that the restriction of T_ω to the invariant subspace \mathcal{W}_ω has an eigenvalue $(1 - \omega)^2$ and its other eigenvalues are generated from formula (2.21) with $|\mu| \neq 1$. Since $2q < p$ such μ must exist. Rewriting the eigenvalue equation $(I - D^{-1}H)x = \mu x$ in the two ways

$$Hx = (1 - \mu)Dx, \quad (2D - H)x = (1 + \mu)Dx,$$

and using a special property of 2×2 block matrices that H and $2D - H$ have the same eigenvalues, we obtain $|\mu| \leq 1$ since both H and $2D - H$ are positive semidefinite and D is positive definite. This shows $\beta < 1$.

Consider eigenvalues μ of $I - D^{-1}H$ with $|\mu| < 1$. If $\omega \geq \omega_{\text{opt}}$, then for such μ the expression under the square root in formula (2.21) is always negative. Hence they generate pairs of conjugate complex eigenvalues λ , but one verifies that they all have the same modulus $|\lambda| = |1 - \omega|$, independent from $|\mu|$. Clearly $|1 - \omega| > (1 - \omega)^2$ so that the asserted formula for ρ_ω is proven for $\omega \geq \omega_{\text{opt}}$. When $0 < \omega < \omega_{\text{opt}}$ the expression under the square root in (2.21) may be negative or not. If it is negative, we have already seen that $|\lambda| = |1 - \omega|$ is generated. If it is nonnegative, which in particular is the case for $\mu = \pm\beta$, the corresponding λ with the larger absolute value is

$$\lambda = 1 - \omega + \frac{1}{2}\omega^2\mu^2 + \omega|\mu|\sqrt{1 - \omega + \frac{1}{4}\omega^2\mu^2}$$

(since the sum before \pm in (2.21) then is nonnegative, too). This expression is maximized for $\mu = \pm\beta$ and also is then larger than $|1 - \omega|$ on the interval $0 < \omega < \omega_{\text{opt}}$. The statements of the lemma follow. \square

Remark 2.6. In the setting of the lemma one always has $q = \dim(\ker H) \leq p/2$, but (if p is even) equality $q = 2p$ could in principle hold. It is then interesting to note that in this case $\rho_\omega = (1 - \omega)^2$, which is minimized for $\omega = 1$, yielding a superlinear convergence rate. However, this case is not relevant in the context of our work, where $p = km + kn$ with a rank $k < \min(m, n)$. In the following theorem we assume $q = k^2$ so that $q < p/2$ is satisfied.

Applying Lemma 2.5 in the context of Theorem 2.3 immediately provides our main result on the asymptotically optimal choice of the shift ω for Algorithm 1.

Theorem 2.7. *Let f be strongly convex and (U_*, V_*) be a critical point of F in (1.2) with U_*, V_* having full column rank $k < \min(m, n)$. Assume that the Hessian $H = \nabla^2 F(U_*, V_*)$ is positive semidefinite and $\dim(\ker H) = k^2$. Let $\rho_1 < 1$ be the asymptotic linear convergence rate of the standard AO method with $\omega = 1$. Fix $0 < \omega < 2$. Then for $X_0 = U_0 V_0^\top$ close enough (this may depend on ω) to $X_* = U_* V_*^\top$ Algorithm 1 is well-defined and the iterates $X_\ell = U_\ell V_\ell^\top$ converge to X_* at an asymptotic linear rate $\rho_\omega < 1$ given in Lemma 2.5 with $\beta^2 = \rho_1$. The optimal asymptotic rate is achieved for*

$$\omega_{\text{opt}} = \frac{2}{1 + \sqrt{1 - \beta^2}}.$$

In practice, the simplest approach for approximating ω_{opt} adaptively is by running the standard method with $\omega = 1$ and estimating $\beta^2 \approx \rho_1$ based on its numerically observed convergence rate. The efficiency of this approach will be illustrated in section 4.

3 Low-rank tensor problems

Clearly, the nonlinear SOR method can be applied to functions with more than two block variables. In low-rank tensor optimization one frequently considers problems of the form

$$\min F(U^1, \dots, U^D) = f(\tau(U^1, \dots, U^D)), \quad (3.1)$$

where now f is a smooth function on a tensor space $\mathbb{R}^{n_1 \times \dots \times n_d}$, and $\tau: \mathcal{V}_1 \times \dots \times \mathcal{V}_D \rightarrow \mathbb{R}^{n_1 \times \dots \times n_d}$ is a multilinear map parametrizing a low-rank tensor format. Such a problem is amenable to alternating optimization since the update for a single block variable U^μ is just an optimization problem for the function f , but on a linear subspace of $\mathbb{R}^{n_1 \times \dots \times n_d}$.

As an important example we mention optimization in the tensor train (TT) format [11]. Here $D = d$ and

$$\tau: \mathcal{V} := \mathbb{R}^{n_1 \times k_1} \times \mathbb{R}^{k_1 \times n_2 \times k_2} \times \dots \times \mathbb{R}^{k_{d-2} \times n_{d-1} \times k_{d-1}} \times \mathbb{R}^{k_{d-1} \times n_d} \rightarrow \mathbb{R}^{n_1 \times \dots \times n_d}$$

is defined via

$$X = \tau(U^1, \dots, U^d) \Leftrightarrow X(i_1, \dots, i_d) = U^1(i_1, :) U^2(:, i_2, :) \cdots U^d(:, i_d), \quad (3.2)$$

which are matrix products of corresponding slices in the so called TT cores $U^\mu \in \mathbb{R}^{k_{\mu-1} \times n_\mu \times k_\mu}$ (one fixes $k_0 = k_1 = 1$). The minimal possible values (k_1, \dots, k_{d-1}) , which determine the sizes of the TT cores, such that such a decomposition is possible are called the TT-ranks of tensor X . Alternating optimization methods form the basis for the majority of computational methods in the tensor train format [4].

An AO method with relaxation for (3.1) takes the form

$$\begin{aligned} U_{\ell+1}^1 &= (1 - \omega)U_\ell^1 + \omega \hat{S}_1(U_\ell^2, \dots, U_\ell^d), \\ &\vdots \\ U_{\ell+1}^\mu &= (1 - \omega)U_\ell^\mu + \omega \hat{S}_\mu(U_{\ell+1}^1, \dots, U_{\ell+1}^{\mu-1}, U_\ell^{\mu+1}, \dots, U_\ell^d), \\ &\vdots \\ U_{\ell+1}^d &= (1 - \omega)U_\ell^d + \omega \hat{S}_d(U_{\ell+1}^1, \dots, U_{\ell+1}^{d-1}), \end{aligned} \quad (3.3)$$

where the \hat{S}_μ return minimizers (or critical points) of the restricted functions $U^\mu \mapsto F(\dots, U^\mu, \dots)$ with the other block variables being fixed. In this form the method has been proposed for the tensor train format in [1]. Under suitable assumptions such an iteration defines a smooth map S_ω from an open subset of $\mathcal{V}_1 \times \dots \times \mathcal{V}_D$ to $\mathbb{R}^{n_1 \times \dots \times n_d}$ for which a similar fixed point analysis as in the matrix case can be conducted. In the following we sketch this for the tensor train format, but the ideas can be applied to general tree tensor network formats such as the Tucker or hierarchical Tucker format. We will make use of several well-known properties of the tensor train format, in particular the quotient manifold structure of tensors of fixed TT-rank and the orbital invariance of AO methods. Most of the related details can be found in [14] and [17].

For the tensor train format (3.2) we assume that $\mathbf{k} = (k_1, \dots, k_{d-1})$ is chosen such that tensors of TT-rank \mathbf{k} exist. Then in fact on a dense and open subset \mathcal{V}' of \mathcal{V} the map τ maps to tensors of fixed TT-rank \mathbf{k} . For convenience we will use the notation $\mathbf{U} = (U^1, \dots, U^d)$ for the elements in \mathcal{V} . As in the matrix case, τ in (3.2) is invariant under a group action, namely,

$$\mathcal{G} = \text{GL}(k_1) \times \dots \times \text{GL}(k_{d-1}) \ni \mathbf{A} = (A_1, \dots, A_{d-1}) \mapsto \theta_{\mathbf{A}} \cdot \mathbf{U},$$

which inserts the product $A_\mu A_\mu^{-1}$ between the matrices $U^\mu(:, i_\mu, :)$ and $U^{\mu+1}(:, i_{\mu+1}, :)$ in (3.2), that is, the slices of the TT cores are transformed according to

$$U^\mu(:, i_\mu, :) \rightarrow A_{\mu-1}^{-1} U^\mu(:, i_\mu, :) A_\mu \quad (3.4)$$

(here $A_0 = A_d = 1$). The corresponding restriction of τ to the quotient manifold \mathcal{V}'/\mathcal{G} is a diffeomorphism onto the set $\mathcal{M}_{\mathbf{k}}$ of tensors of fixed TT-rank \mathbf{k} , which is an embedded submanifold of $\mathbb{R}^{n_1 \times \dots \times n_d}$ of dimension $\dim(\mathcal{M}_{\mathbf{k}}) = \dim(\mathcal{V}) - \dim(\mathcal{G})$. Notably, let $\mathbf{U} \in \mathcal{V}'$, then $\tau'(\mathbf{U}) = 0$ on the tangent space of the orbit $\theta_{\mathbf{A}} \cdot \mathbf{U}$ at \mathbf{U} . On any complementary subspace \mathcal{W} to that tangent space, $\tau'(\mathbf{U})$ is a bijection from \mathcal{W} to the tangent space of $\mathcal{M}_{\mathbf{k}}$ at $\tau(\mathbf{U})$.

Assume again that f is smooth and strongly convex. Then any critical point \mathbf{U}_* of the function $F = f \circ \tau$ that lies in \mathcal{V}' is a fixed point of the iteration (3.3) since the restricted linear maps $U^\mu \mapsto \tau(U_*^1, \dots, U^\mu, \dots, U_*^d)$ are injective so that the corresponding restriction $U^\mu \mapsto F(U_*^1, \dots, U^\mu, \dots, U_*^d)$ is strongly convex. Moreover, the whole process is well-defined in some neighborhood of (the orbit of) \mathbf{U}_* where it can be written as

$$\mathbf{U}_{\ell+1} = S_\omega(\mathbf{U}_\ell)$$

with a smooth map S_ω . A key observation to make is that the maps $\hat{S}_1, \dots, \hat{S}_d$ in (3.3) that realize the updates of single TT cores exhibit an analogous compatibility with the group action as in (2.3) for the matrix case, namely

$$\hat{S}_\mu(\theta_{\mathbf{A}} \cdot \mathbf{U}) = A_{\mu-1}^{-1} \hat{S}_\mu(\mathbf{U}) A_\mu,$$

where the matrix product is understood slice-wise as in (3.4) (and we slightly abused notation since \hat{S}_μ does not depend on U^μ). It entails a corresponding invariance

$$S_\omega(\theta_{\mathbf{A}} \cdot \mathbf{U}) = \theta_{\mathbf{A}} \cdot S_\omega(\mathbf{U}) \quad (3.5)$$

of a full update loop, in analogy to (2.9). This allows us to regard (3.3) as a well-defined iteration

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

on the manifold $\mathcal{M}_{\mathbf{k}}$, at least locally in a neighborhood of $\tau(\mathbf{U}_*)$. From a practical viewpoint, the invariance (3.5) admits to change the tensor train representation in every substep of (3.3) in order to make the restricted linear maps $U^\mu \mapsto \tau(\dots, U^\mu, \dots)$ orthogonal and improve numerical stability. We refer to [4, 14] for details on orthogonalization of substeps.

Based on these similarities to the matrix case, one can proceed in almost the same way as in section 2. Let $\mathbf{U}_* \in \mathcal{V}'$ be a critical point of F , that is, $\nabla F(\mathbf{U}_*) = 0$. Due to the orbital invariance of F , the Hessian $H = \nabla^2 F(\mathbf{U}_*)$ has a kernel of dimension at least $\dim(\mathcal{G})$ since it contains the tangent space to the orbit at \mathbf{U}_* . However, in the block decomposition

$$H = D + E + E^\top \quad (3.6)$$

into a block diagonal part D (corresponding to the block variables U^1, \dots, U^d) and lower block triangular part E , the block matrix D is positive definite since f is strongly convex.² The derivative of S_ω then again takes the form of an SOR error iteration matrix

$$T_\omega = I - N_\omega^{-1}H, \quad N_\omega = \frac{1}{\omega}D + E,$$

similar to (2.16); see, e.g., [8, Thm. 10.3.4 & 10.3.5] for the derivation. For $0 < \omega < 2$ a decomposition $\mathcal{V} = \ker H \oplus \mathcal{W}_\omega$ as in (2.18) applies and T_ω is a contraction on the invariant subspace \mathcal{W}_ω if H is positive semidefinite. Using the same proof as for Theorem 2.3 we obtain the analogous local convergence result for tensor train optimization. Recall that we assume that \mathbf{k} is properly chosen so that τ maps the open and dense subset \mathcal{V}' to the manifold $\mathcal{M}_{\mathbf{k}}$.

Theorem 3.1. *Let $\mathbf{U}_* \in \mathcal{V}'$ be a critical point of function F in (3.1) where f is strongly convex. Assume that the Hessian $H = \nabla^2 F(\mathbf{U}_*)$ is positive semidefinite and $\dim(\ker H) = \dim(\mathcal{G}) = k_1^2 + \dots + k_{d-1}^2$. Fix $0 < \omega < 2$. Then for $X_0 = \tau(\mathbf{U}_0)$ close enough (this may depend on ω) to $X_* = \tau(\mathbf{U}_*)$ the iteration (3.3) is well-defined and the iterates $X_\ell = \tau(\mathbf{U}_\ell)$ converge to X_* at an asymptotic linear rate*

$$\rho_\omega = \limsup_{\ell \rightarrow \infty} \|X_\ell - X_*\|^{1/\ell} < 1,$$

where ρ_ω is the spectral radius of T_ω on \mathcal{W}_ω .

While so far everything looks conceptually almost identical to the matrix case, a major difference arises when proceeding to determine the optimal shift parameter ω . It is not clear whether Lemma 2.5 can be generalized. This is in fact already an issue with the linear SOR method with more than two blocks for positive definite systems, since certain conditions on the decomposition (3.6) of H are required in order to derive a formula like (2.21) for the eigenvalues of T_ω ; cf. [2, Sec. 4.6]. In the matrix case, the fact that $E + E^\top$ is two-cyclic makes this possible but for more than two block variables assuming such conditions on E does not appear very reasonable, especially when taking into account that the critical point (U_*, V_*) and hence its Hessian are not given a priori. Moreover, even if the formula (2.21) would apply, one would need that the eigenvalues of the Jacobi error iteration matrix $I - D^{-1}H$ have absolute value at most one, but for a block decomposition (3.6) with more than two blocks this does not follow from the positive definiteness of D alone. Thus, for the tensor train format the estimation of an optimal parameter ω_{opt} from formula (2.20) remains a heuristic.

4 Numerical experiments

In the last section we present some numerical experiments to illustrate the benefit of overrelaxation in low-rank optimization.

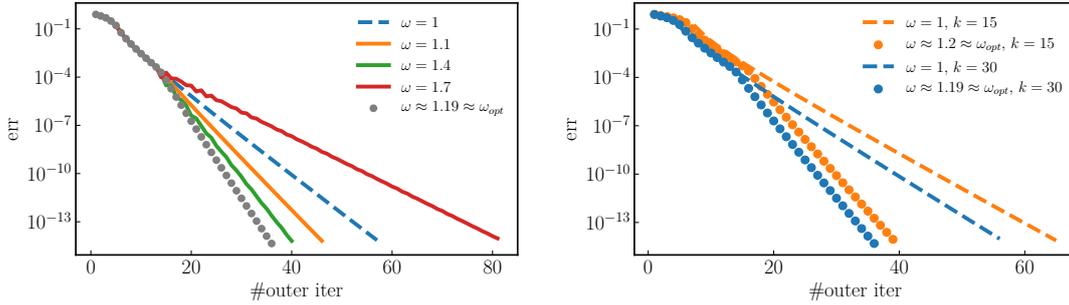


Figure 1: Relative residuals (4.2) of Algorithm 1 for the completion problem (4.1) with respect to the number of outer iterations using various shift parameters ω for rank $k = 30$ (left) and for rank values $k = 15, 30$ (right). The parameter $\omega = 1$ corresponds to the standard ALS method, while parameters $\omega \in (1, 2)$ represent the version of the iteration with overrelaxation. The $\omega \approx \omega_{\text{opt}}$ case corresponds to the choice (2.20) with β^2 estimated using (4.3).

4.1 Matrix completion problem

First, we apply the proposed AO overrelaxation scheme in Algorithm 1 to the following nonconvex formulation of a low-rank matrix completion problem:

$$\min_{U \in \mathbb{R}^{m \times k}, V \in \mathbb{R}^{n \times k}} F(U, V) = \frac{1}{2} \|P_{\Omega}(A - UV^{\top})\|_F^2. \quad (4.1)$$

Here Ω is a given set of index pairs, and the linear operator $P_{\Omega}: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ is defined as

$$(P_{\Omega}(X))_{ij} = \begin{cases} x_{ij}, & (i, j) \in \Omega, \\ 0, & \text{otherwise.} \end{cases}$$

In our experiments the set Ω consists of randomly generated index pairs. We set $m = n$ and choose A to be a random rank- k matrix, i.e., $A = U_* V_*^{\top}$, where $U_*, V_* \in \mathbb{R}^{n \times r}$ are random matrices with each element sampled from a standard Gaussian distribution. The number of sampled entries $|\Omega|$ is defined by an oversampling parameter $\text{OS} \geq 1$,

$$|\Omega| = \text{OS} \cdot (2nk - k^2),$$

since we want $|\Omega|$ to be larger than $2nk - k^2 = \dim(\mathcal{M}_k)$, which is the number of essential degrees of freedom for an $n \times n$ matrix of rank k . In the experiments $\text{OS} = 3$.

In Figure 1 we present the convergence plots for an experiment with $n = 2000$ for several choices of ω and different ranks k . We report the relative residuals

$$\mathbf{err}_{\ell} = \frac{\|P_{\Omega}(A - U_{\ell} V_{\ell}^{\top})\|_F}{\|P_{\Omega}(A)\|_F}, \quad (4.2)$$

where the sequence (U_{ℓ}, V_{ℓ}) is generated by Algorithm 1 with a shift parameter ω . The only difference with Algorithm 1 is that we always start with $\omega = 1$ (standard ALS) and only turn on the shift after the convergence has stabilized, in this experiment usually after 12 iterations. The optimal shift ω_{opt} from (2.20) depends on the convergence rate $\beta^2 = \rho_1$ of the standard AO method, which is estimated while running the iteration with $\omega = 1$ using

$$\beta^2 \approx \sqrt{\frac{\mathbf{err}_{\ell+2}}{\mathbf{err}_{\ell}}}. \quad (4.3)$$

²As in Remark 2.4, this also follows from the assumptions that H is positive semidefinite and $\dim(\ker H) = \dim(\mathcal{G})$, since the tangent space to the orbit at $\mathbf{U}_* \in \mathcal{V}'$ does not contain elements of the form $(0, \dots, 0, U^{\mu}, 0, \dots, 0)$.

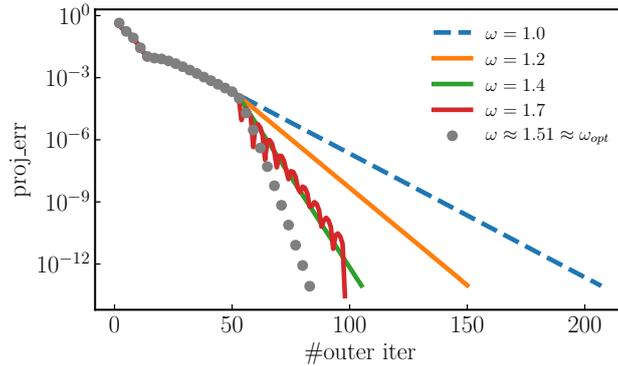


Figure 2: Relative residuals (4.5) of Algorithm 1 for the low-rank solution of a Lyapunov equation (4.4) with respect to the number of outer iterations using various shift parameters ω . Here $k = 2$. The parameter $\omega = 1$ corresponds to the standard ALS method, while parameters $\omega \in (1, 2)$ represent the version of the iteration with overrelaxation. The $\omega \approx \omega_{\text{opt}}$ case corresponds to the choice (2.20) with β^2 estimated using (4.6).

As expected, using overrelaxation accelerates the convergence of the ALS method if ω is chosen properly. We note that the additional computations arising from the overrelaxation scheme come in asymptotically negligible cost as compared with the basic ALS. In turn, the proposed approach leads to a significant reduction of the total number of iterations for achieving a high accuracy.

4.2 Low-rank solution of the Lyapunov equation

Consider the Lyapunov equation

$$AX + XA^\top = B, \quad A, B \in \mathbb{R}^{n \times n}, \quad (4.4)$$

where $X \in \mathbb{R}^{n \times n}$ is a matrix to be found. In case of a symmetric positive definite matrix A , equation (4.4) represents the optimality condition for the strongly convex optimization problem

$$\min_{X \in \mathbb{R}^{n \times n}} f(X) = \frac{1}{2} \langle AX + XA^\top, X \rangle_F - \langle B, X \rangle_F.$$

A rank- k approximation to the solution is, therefore, obtained by solving the problem

$$\min_{U, V \in \mathbb{R}^{n \times k}} F(U, V) = f(UV^\top)$$

instead, which is of the form (1.2). For this we employ the proposed overrelaxation algorithm.

In the experiment, we choose $A = (n + 1)^2 \text{tridiag}(-1, 2, -1)$, set $n = 256$ and generate the right-hand side $B = AX_* + X_*A^\top$ from a specified solution X_* . Specifically, we choose $k = 2$ and generate the second and the third singular values of X_* such that their ratio equals to 0.99, which is similar to experiments conducted in [13]. There it has been numerically observed that such a large ratio at the target singular value can lead to slow convergence of the standard ALS method.

Due to the fact that X_* cannot be approximated with high accuracy using $k = 2$, the function values $f(X)$ will not converge to zero and hence cannot be taken as an appropriate error measure. Instead, we compute the values

$$\text{proj_err}_\ell = \frac{\|\mathbf{P}_{X_\ell} \nabla f(X_\ell)\|_F}{\|\mathbf{P}_{X_\ell}(B)\|_F} = \frac{\|\mathbf{P}_{X_\ell}(AX_\ell + X_\ell A^\top - B)\|_F}{\|\mathbf{P}_{X_\ell}(B)\|_F}, \quad X_\ell = U_\ell V_\ell^\top, \quad (4.5)$$

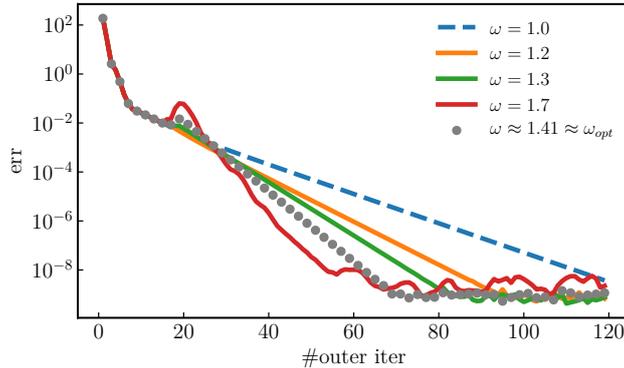


Figure 3: Maximal (within one ALS sweep) relative local residual with respect to the number of outer iterations of the QTT ALS method for solving a Lyapunov equation (4.4) using various shifts ω . The QTT ranks are $(4, \dots, 4)$. The parameter $\omega = 1$ corresponds to the standard QTT ALS method, while parameters $\omega \in (1, 2)$ represent the iteration with overrelaxation. The $\omega \approx \omega_{\text{opt}}$ case corresponds to (2.20) with β^2 estimated from the standard ALS method, even though for tensor problems there is no theoretical guarantee that this choice is actually close to optimal. Shifts are activated after 15 iterations.

where \mathbf{P}_{X_ℓ} denotes the orthogonal projection operator to the tangent space of the manifold \mathcal{M}_k of fixed rank- k matrices at X_ℓ ; see, e.g., [18]. This reflects the fact that the method can be regarded as a minimization method on that manifold. Similarly to (4.3), we can then use these values for approximating the optimal shift parameter ω_{opt} using (2.20) with β^2 estimated from

$$\beta^2 \approx \sqrt{\frac{\text{proj_err}_{\ell+2}}{\text{proj_err}_\ell}}. \quad (4.6)$$

In Figure 2, we plot the values of proj_err_ℓ against the number of outer iterations ℓ for several values of shifts ω , including the basic ALS and the approximated optimal shift. In all cases the shift is activated after 50 iterations. All considered shifts lead to convergence improvement with the shift that approximates the optimal one being the best.

4.3 Linear systems in the quantized tensor train format

Finally, we test our approach for solving linear systems in the tensor train format. In particular, we apply the so-called quantized tensor train (QTT) format [6, 10] to solve the equation (4.4) by fixing $n = 2^d$, $d = 12$, and by representing $X \in \mathbb{R}^{2^d \times 2^d}$ as order- $2d$ tensors in $\mathbb{R}^{2 \times \dots \times 2}$ using reshape in the lexicographical order. These tensors are then further restricted to the tensor train format with the TT-rank equal to $(4, 4, \dots, 4)$ (this choice of ranks led to a much slower convergence of the ALS method as compared to other rank values). The right-hand side B was selected to be a matrix of all ones, which trivially admits a QTT representation with all TT-ranks equal to one. Note that all computations were performed directly in the tensor train format, i.e., no full tensors were formed.

As an error measure err_ℓ we take the maximum relative norm of all local residuals within one sweep of the standard ALS [12]. Based on this error we estimate β^2 and use the same formula (4.3) for ω_{opt} , but as noted in section 3 there is no theoretical guarantee that this formula provides the optimal shift parameter in the tensor case. Nevertheless, the results from Figure 3 suggest that this choice leads to nearly the fastest convergence among the considered choices of shifts.

References

- [1] L. Grasedyck, M. Kluge, and S. Krämer. Variants of alternating least squares tensor completion in the tensor train format. *SIAM J. Sci. Comput.*, 37(5):A2424–A2450, 2015.
- [2] W. Hackbusch. *Iterative solution of large sparse systems of equations*. Springer, [Cham], second edition, 2016.
- [3] L. A. Hageman and T. A. Porsching. Aspects of nonlinear block successive overrelaxation. *SIAM J. Numer. Anal.*, 12:316–335, 1975.
- [4] S. Holtz, T. Rohwedder, and R. Schneider. The alternating linear scheme for tensor optimization in the tensor train format. *SIAM J. Sci. Comput.*, 34(2):A683–A713, 2012.
- [5] H. B. Keller. On the solution of singular and semidefinite linear systems by iteration. *J. Soc. Indust. Appl. Math. Ser. B Numer. Anal.*, 2:281–290, 1965.
- [6] B. N. Khoromskij. $O(d \log N)$ -quantics approximation of N - d tensors in high-dimensional numerical modeling. *Constr. Approx.*, 34(2):257–280, 2011.
- [7] J. M. Lee. *Introduction to smooth manifolds*. Springer-Verlag, New York, 2003.
- [8] J. M. Ortega and W. C. Rheinboldt. *Iterative solution of nonlinear equations in several variables*. Academic Press, New York-London, 1970.
- [9] J. M. Ortega and M. L. Rockoff. Nonlinear difference equations and Gauss-Seidel type iterative methods. *SIAM J. Numer. Anal.*, 3:497–513, 1966.
- [10] I. V. Oseledets. Approximation of $2^d \times 2^d$ matrices using tensor decomposition. *SIAM J. Matrix Anal. Appl.*, 31(4):2130–2145, 2009/10.
- [11] I. V. Oseledets. Tensor-train decomposition. *SIAM J. Sci. Comput.*, 33(5):2295–2317, 2011.
- [12] I. V. Oseledets and S. V. Dolgov. Solution of linear systems and matrix inversion in the TT-format. *SIAM J. Sci. Comput.*, 34(5):A2718–A2739, 2012.
- [13] I. V. Oseledets, M. V. Rakhuba, and A. Uschmajew. Alternating least squares as moving subspace correction. *SIAM J. Numer. Anal.*, 56(6):3459–3479, 2018.
- [14] T. Rohwedder and A. Uschmajew. On local convergence of alternating schemes for optimization of convex problems in the tensor train format. *SIAM J. Numer. Anal.*, 51(2):1134–1162, 2013.
- [15] S. Schechter. Iteration methods for nonlinear problems. *Trans. Amer. Math. Soc.*, 104:179–189, 1962.
- [16] A. Uschmajew. Local convergence of the alternating least squares algorithm for canonical tensor approximation. *SIAM J. Matrix Anal. Appl.*, 33(2):639–652, 2012.
- [17] A. Uschmajew and B. Vandereycken. The geometry of algorithms using hierarchical tensors. *Linear Algebra Appl.*, 439(1):133–166, 2013.
- [18] B. Vandereycken. Low-rank matrix completion by Riemannian optimization. *SIAM J. Optim.*, 23(2):1214–1236, 2013.
- [19] J. Wang, Y.-P. Wang, Z. Xu, and C.-L. Wang. Accelerated low rank matrix approximate algorithms for matrix completion. *Comput. Math. Appl.*, 77(2):334–341, 2019.
- [20] J. Weissinger. Verallgemeinerungen des Seidelschen Iterationsverfahrens. *Z. Angew. Math. Mech.*, 33:155–163, 1953. (German).
- [21] Z. Wen, W. Yin, and Y. Zhang. Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm. *Math. Program. Comput.*, 4(4):333–361, 2012.
- [22] D. M. Young. *Iterative solution of large linear systems*. Academic Press, New York-London, 1971.