

Some results concerning rank-one truncated steepest descent directions in tensor spaces

André Uschmajew
Hausdorff Center for Mathematics &
Institute for Numerical Simulation
University of Bonn
53115 Bonn, Germany
Email: uschmajew@ins.uni-bonn.de

Abstract—The idea of finding low-rank solutions to matrix or tensor optimization tasks by greedy rank-one methods has been showing itself repeatedly in the literature. The simplest method, and often a central building block in accelerated methods, consists in performing updates along low-rank approximations of the negative gradient. This is convenient as it does increase the rank in a prescribed manner per step, and also because it allows for a somewhat surprisingly simple convergence analysis. The main point is that in a tensor product space of finite dimension, the best rank-one approximation of a tensor has a guaranteed minimal overlap with the tensor itself. Thus rank-one approximations of anti-gradients provide descent directions. This key concept can also be used in Hilbert space, if the rank growth of the approximation sequence can be balanced with convergence speed. This work presents a conceptual review of this approach, and also provides some new insights.

I. INTRODUCTION

We consider an optimization task in a tensor product space $\mathbb{R}^{n_1 \times \dots \times n_d}$ of the form

$$\text{minimize } J(\mathbf{X}). \quad (1)$$

We assume that the function $J: \mathbb{R}^{n_1 \times \dots \times n_d} \rightarrow \mathbb{R}$ is smooth, convex, and admits at least one solution \mathbf{X}_* that can be approximated more or less exactly by a tensor of rank at most r , that is,

$$\mathbf{X}_* \approx \sum_{j=1}^r u_j^{(1)} \circ \dots \circ u_j^{(d)}, \quad (2)$$

where \circ denotes the outer vector product. The main examples we have in mind are convex quadratic problems, for instance, the *best-approximation problem* in Frobenius norm,

$$\text{minimize } \frac{1}{2} \|\mathbf{X}_* - \mathbf{X}\|_F^2, \quad (3)$$

a smooth version of the *tensor completion* problem,

$$\text{minimize } \frac{1}{2} \|\mathcal{P}_\Omega(\mathbf{X}_* - \mathbf{X})\|_F^2, \quad (4)$$

where \mathcal{P}_Ω is the orthogonal projection onto the subspace of tensors which are zero at indices not in Ω , and, more generally, *energy (semi-)norm minimization* for high-dimensional linear system,

$$\text{minimize } \frac{1}{2} \langle \mathbf{X}, \mathcal{A}(\mathbf{X}) \rangle_F - \langle \mathbf{X}, \mathbf{B} \rangle_F, \quad (5)$$

where \mathcal{A} is a linear, symmetric, and positive semidefinite operator on $\mathbb{R}^{n_1 \times \dots \times n_d}$, and $\mathbf{B} = \mathcal{A}(\mathbf{X}_*)$.

In this work, we do not focus on a specific problem, but on the convergence of a certain class of greedy-type methods to construct low-rank approximations to the solution of such problems, which turned out to be useful in many applications. The simple idea is the following: given some approximate solution \mathbf{X}_k one wishes to find an improvement \mathbf{X}_{k+1} of increased rank. A possible strategy is to take a low-rank approximation of some descent direction, perform a line-search in this direction, and find \mathbf{X}_{k+1} in this way. A popular choice in case of quadratic problems as the ones above is to take rank-one approximations of residuals in suitable norms [1], [2], [3], that is, the anti-gradient $-\nabla J(\mathbf{X}_k)$. Using the best rank-one approximation ratio as in [4], one obtains an interpretation as a perturbed steepest descent method and can deduce a rate of convergence. This main insight is also the basis for obtaining convergence rates for related methods which use additional acceleration, like in [5], [6], [7], [8], [9].

As a contribution, we present convergence results for *semidefinite* quadratic problems, and also a new application of this idea regarding the low-rank approximability of solutions to Kronecker-structured matrix equations in Hilbert space [10]. Specifically, Theorem 5 below extends the convergence results from [3] on proper generalized decomposition for elliptic systems in Hilbert space by providing a rate of convergence for tensor structured semidefinite problems. For simplicity, we will focus on rank-one increments. W.l.o.g. we now silently assume $1 \leq n_1 \leq \dots \leq n_d$.

II. RANK-ONE APPROXIMATION RATIO

A main ingredient in the considered methods are rank-one approximations to a given tensor \mathbf{X} . Let \mathcal{R}_1 denote the set of tensors of the form $u^{(1)} \circ \dots \circ u^{(d)}$ (rank *at most* one). This set is closed [11], therefore we can consider a function $\mathbf{R}_1(\mathbf{X})$ that assigns to \mathbf{X} a *best rank-one approximation* \mathbf{R}_1 , that is, a best approximation to \mathbf{X} by an element in \mathcal{R}_1 . It serves a theoretical purpose, and is not available in practice for true tensors ($d \geq 3$) [12]. To obtain an approximation to $\mathbf{R}_1(\mathbf{X})$ one may try to solve the nonlinear optimization task

$$\text{minimize}_{u^{(1)}, \dots, u^{(d)}} \frac{1}{2} \|\mathbf{X} - u^{(1)} \circ \dots \circ u^{(d)}\|_F^2.$$

For instance, one can use iterative methods like alternating least squares (see references in [13]), also known as the higher-order power method (HOPM) [14], [15]. This method will at least converge toward a single critical point, as was shown only recently [16], [17]. A practical alternative is to use the *higher-order SVD* (HOSVD) truncation of \mathbf{X} to multilinear rank $(1, \dots, 1)$ as a surrogate for the best rank-one approximation, or at least as an initial guess for the HOPM [18], [15]. When the tensor \mathbf{X} is very large, one will need some additional structure (sparsity, low-rank, ...) to apply these methods efficiently.

In principle, it will be sufficient to assign to a tensor \mathbf{X} a rank-one tensor $\mathbf{Y} \in \mathcal{R}_1$ that satisfies an angle condition

$$\langle \mathbf{X}, \mathbf{Y} \rangle_F \geq \alpha \|\mathbf{X}\|_F \|\mathbf{Y}\|_F,$$

with $\alpha > 0$ independent of \mathbf{X} . The role of the best rank-one approximation $\mathbf{Y} = \mathbf{R}_1(\mathbf{X})$ is that it always guarantees this. To state this assertion more precisely, one introduces the *best rank-one approximation ratio* [4]

$$\begin{aligned} \mu_{n_1, \dots, n_d} &= \min_{\mathbf{X} \neq 0} \max_{\mathbf{Y} \in \mathcal{R}_1} \frac{\langle \mathbf{X}, \mathbf{Y} \rangle_F}{\|\mathbf{X}\|_F \|\mathbf{Y}\|_F} \\ &= \min_{\mathbf{X} \neq 0} \frac{\langle \mathbf{X}, \mathbf{R}_1(\mathbf{X}) \rangle_F}{\|\mathbf{X}\|_F \|\mathbf{R}_1(\mathbf{X})\|_F}. \end{aligned} \quad (6)$$

Note that the second equality follows from the fact that the set \mathcal{R}_1 is a double-cone. So by standard least-squares considerations it is clear that for a given \mathbf{X} , the problem of finding an element of minimal distance in \mathcal{R}_1 , and a direction of maximum overlap are equivalent. The following observation has been stated in [4].

Lemma 1. *Recalling $n_1 \leq \dots \leq n_d$, it holds*

$$\mu_{n_1, \dots, n_d} \geq \frac{1}{\sqrt{n_1 \cdots n_{d-1}}}.$$

Proof. Let $e_1^{(j)}, \dots, e_{n_j}^{(j)}$ be an orthonormal basis of \mathbb{R}^{n_j} . Then one can expand an arbitrary \mathbf{X} as

$$\begin{aligned} \mathbf{X} &= \sum_{i_1=1}^{n_1} \cdots \sum_{i_d=1}^{n_d} c_{i_1 \dots i_d} e_{i_1}^{(1)} \circ \cdots \circ e_{i_d}^{(d)} \\ &= \sum_{i_1=1}^{n_1} \cdots \sum_{i_{d-1}=1}^{n_{d-1}} e_{i_1}^{(1)} \circ \cdots \circ e_{i_{d-1}}^{(d-1)} \circ \left(\sum_{i_d=1}^{n_d} c_{i_1 \dots i_d} e_{i_d}^{(d)} \right). \end{aligned}$$

The second representation is an orthogonal expansion (w.r.t. Frobenius inner product) into $n_1 \cdots n_{d-1}$ terms from \mathcal{R}_1 . Therefore, the term with largest norm must provide at least the asserted overlap. \square

The bound $\frac{1}{\sqrt{n_1 \cdots n_{d-1}}}$ deteriorates with the dimension of the tensor space, which will play a role later on. It is unknown whether it is sharp [4].

III. CONVERGENCE OF TRUNCATED GRADIENT DESCENT

A. General convergence statement

The method of rank-one truncated gradient descent for solving (1) we consider here is

$$\mathbf{X}_{k+1} = \mathbf{X} + \alpha_k \mathbf{R}_1(-\nabla J(\mathbf{X}_k)). \quad (7)$$

We assume that the step lengths are chosen such that the Wolfe conditions are satisfied; cf. [19, Eq. (3.6)]. Then one step of (7) provides descent while increasing the rank at most by one. Hence the iteration can be regarded as a greedy rank-one algorithm. By the standard results, see for example [19, Theorem 3.2], $\nabla J(\mathbf{X}_k) \rightarrow 0$, which specifically implies $\mathbf{X}_k \rightarrow \mathbf{X}_*$ in case of strictly convex J . The key point here is that step sizes satisfying the Wolfe conditions do really exist [19, Lemma 3.1], since, by Lemma 1, the search directions $\mathbf{R}_1(-\nabla J(\mathbf{X}_k))$ retain sufficient overlap with $-\nabla J(\mathbf{X}_k)$, the cosine of the subspace angle being at least $1/\sqrt{n_1 \cdots n_{d-1}}$. Moreover, when J is strongly convex, an exact line-search, that is, minimization of $\alpha \mapsto J(\mathbf{X}_k + \alpha \mathbf{R}_1(-\nabla J(\mathbf{X}_k)))$ will provide such a step size.

B. Rate of convergence for quadratic problems

A linear rate of convergence can be derived for strongly convex functions with exact line-search [19]. To obtain this result it is more or less sufficient to study the case of quadratic J , that is, problem (5). However, in the quadratic case we can even work with positive semidefinite \mathcal{A} if we confine ourselves to measure convergence in the energy (semi-)norm

$$\|\mathbf{X}\|_{\mathcal{A}} = \sqrt{\langle \mathbf{X}, \mathcal{A}(\mathbf{X}) \rangle_F}.$$

The corresponding (pseudo-)inner product is $\langle \mathbf{X}, \mathbf{Y} \rangle_{\mathcal{A}} = \langle \mathbf{X}, \mathcal{A}(\mathbf{Y}) \rangle_F$. For the quadratic problem (5), $\nabla J(\mathbf{X}_n) = \mathcal{A}(\mathbf{X}) - \mathbf{B}$ is called the residual. Also note that

$$J(\mathbf{X}) = \frac{1}{2} \|\mathbf{X} - \mathbf{X}_*\|_{\mathcal{A}}^2 + J(\mathbf{X}^*) \quad (8)$$

for any \mathbf{X}^* from the possibly affine linear space of solution.

A contraction estimate for perturbed steepest descent is derived from the following logic. Letting λ_{\min}^+ denote the minimal positive eigenvalue of \mathcal{A} , one checks that

$$\begin{aligned} \|\nabla J(\mathbf{X}_k)\|_F^2 &= \|\mathcal{A}^{1/2}(\mathbf{X}_k - \mathbf{X}_*)\|_{\mathcal{A}}^2 \\ &\geq \lambda_{\min}^+ \|\mathcal{A}^{1/2}(\mathbf{X}_k - \mathbf{X}_*)\|_F^2 \\ &= \lambda_{\min}^+ \|\mathbf{X}_k - \mathbf{X}_*\|_{\mathcal{A}}^2 \end{aligned}$$

(the inequality holds since $\mathcal{A}^{1/2}$ maps into the invariant subspace of \mathcal{A}). Therefore, a relation

$$\langle -\nabla J(\mathbf{X}_k), \mathbf{Y} \rangle_F \geq \cos \theta_k \|\nabla J(\mathbf{X}_k)\|_F \|\mathbf{Y}\|_F \quad (9)$$

will imply

$$|\langle \mathbf{X}_k - \mathbf{X}_*, \mathbf{Y} \rangle_{\mathcal{A}}| \geq \frac{\cos \theta_k}{\sqrt{\kappa_{\text{eff}}}} \|\mathbf{X}_k - \mathbf{X}_*\|_{\mathcal{A}} \|\mathbf{Y}\|_{\mathcal{A}},$$

where $\kappa_{\text{eff}} = \|\mathcal{A}\|/\lambda_{\min}^+$ is the effective condition of \mathcal{A} on its range (which notably equals one for the problems (3) and (4)). Hence, writing the \mathcal{A} -orthogonal projection of $\mathbf{X}_k - \mathbf{X}_*$ onto the span of \mathbf{Y} by $\mathbf{X}_{k+1} - \mathbf{X}_*$ (which is unique if \mathbf{Y} is not in the null space of \mathcal{A}), it holds

$$\|\mathbf{X}_{k+1} - \mathbf{X}_*\|_{\mathcal{A}} \leq \left(1 - \frac{\cos^2 \theta_k}{\kappa_{\text{eff}}}\right)^{1/2} \|\mathbf{X}_k - \mathbf{X}_*\|_{\mathcal{A}}. \quad (10)$$

It remains to validate that \mathbf{X}_{k+1} is equivalently obtained by steepest descent as

$$\mathbf{X}_{k+1} = \mathbf{X}_k + \alpha_k \mathbf{Y} \quad (11)$$

when α_k is chosen by minimizing $\alpha \mapsto J(\mathbf{X}_k + \alpha \mathbf{Y})$. However this follows easily when using the formula (8).

This is of course more or less classic theory, but has found new applications in the context of the low-rank truncations of residuals as considered here. In fact, for the iteration (7) we obtain from Lemma 1 the following rate of convergence.

Theorem 2. *For the quadratic, semidefinite problem (5), the truncated gradient descent method (7) satisfies*

$$\begin{aligned} \|\mathbf{X}_k - \mathbf{X}_*\|_{\mathcal{A}} &\leq \left(1 - \frac{\mu_{n_1, \dots, n_d}^2}{\kappa_{\text{eff}}}\right)^{k/2} \|\mathbf{X}_0 - \mathbf{X}_*\|_{\mathcal{A}} \\ &\leq \left(1 - \frac{1}{\kappa_{\text{eff}} \cdot n_1 \cdots n_{d-1}}\right)^{k/2} \|\mathbf{X}_0 - \mathbf{X}_*\|_{\mathcal{A}}. \end{aligned}$$

This result has been previously obtained in [4] for the pure greedy rank-one approximation algorithm, which consists in applying (7) to (3), and in [7] for solving the semidefinite completion problem (4) by rank-one matrix pursuit. In both cases, it holds $\kappa_{\text{eff}} = 1$. One should stress, however, that for (4) the convergence of the residual to zero, as follows from the above theorem, does not imply that a matrix of lowest rank is found [7]. Only the error on the observed entries will converge to zero. To achieve low-rank recovery, additional optimization on fixed rank manifolds is necessary, such as GECO [5], Riemannian pursuit [8], [9], or the subspace improvement methods in [7]. Typically, such improvements result in either the same or even accelerated convergence, as can be analyzed in a fairly general context [20]. This remark also applies to the AMEn algorithm [6] for low-rank approximation in tensor train format.

The contraction factor in (10) is not optimal. This is well known for the unperturbed steepest descent ($\theta_k = 0$ in (9)). Kantorovich [21] proved that the contraction factor for the squared energy error is at most $\left(\frac{\kappa_{\text{eff}} - 1}{\kappa_{\text{eff}} + 1}\right)^2$, which is roughly the square of the rate in (10) (when $\theta_k = 0$). It requires a more sophisticated analysis based on the Kantorovich inequality. Munthe-Kaas¹ [22] was able to modify the result for the perturbed steepest descent method in the case that \mathcal{A} is positive definite:

$$\|\mathbf{X}_{k+1} - \mathbf{X}_*\|_{\mathcal{A}} \leq \left(\frac{\tilde{\kappa}_k - 1}{\tilde{\kappa}_k + 1}\right) \|\mathbf{X}_k - \mathbf{X}_*\|_{\mathcal{A}}, \quad (12)$$

where $\tilde{\kappa}_k = \left(\frac{1 + \sin \theta_k}{1 - \sin \theta_k}\right) \kappa$. Here we write κ instead of κ_{eff} since \mathcal{A} is assumed to be positive definite. However, by the usual logic of restricting the analysis to the invariant subspace of \mathcal{A} , one can show that the same contraction factor in energy semi-norm is achieved in the semidefinite case, too. To see this, we compare the update (11) with

$$\tilde{\mathbf{X}}_{k+1} = \mathcal{P}_{\mathcal{A}}(\mathbf{X}_k) + \tilde{\alpha}_k \mathcal{P}_{\mathcal{A}}(\mathbf{Y}) = \mathcal{P}_{\mathcal{A}}(\mathbf{X}_k + \tilde{\alpha}_k \mathbf{Y}), \quad (13)$$

¹We found the reference to Munthe-Kaas' result in [6].

where $\mathcal{P}_{\mathcal{A}}$ is the orthogonal projection on the range of \mathcal{A} , and $\tilde{\alpha}_k$ is chosen to minimize $\alpha \mapsto J(\mathcal{P}_{\mathcal{A}}(\mathbf{X}_k + \alpha \mathbf{Y}))$. Since $J(\mathbf{X}) = J(\mathcal{P}_{\mathcal{A}}(\mathbf{X}))$, the result is of course the same, that is, $\tilde{\alpha}_k = \alpha_k$. But then it is also true that $\tilde{\mathbf{X}}_{k+1} = \mathcal{P}_{\mathcal{A}}(\mathbf{X}_{k+1})$. Now note that (13) is a perturbed steepest descent step in the invariant subspace of the operator \mathcal{A} , on which it is positive definite with condition κ_{eff} (the unique minimizer in this space is $\mathcal{P}_{\mathcal{A}}(\mathbf{X}_*)$). Therefore, the result (12) of Munthe-Kaas applies correspondingly when using the angle $\tilde{\theta}_k$ between $-\nabla J(\mathbf{X}_k)$ and $\mathcal{P}_{\mathcal{A}}(\mathbf{Y})$. For this angle, it holds

$$\begin{aligned} \cos \tilde{\theta}_k &= \frac{\langle -\nabla J(\mathbf{X}_k), \mathcal{P}_{\mathcal{A}}(\mathbf{Y}) \rangle_F}{\|\nabla J(\mathbf{X}_k)\|_F \|\mathcal{P}_{\mathcal{A}}(\mathbf{Y})\|_F} \\ &\geq \frac{\langle -\nabla J(\mathbf{X}_k), \mathbf{Y} \rangle_F}{\|\nabla J(\mathbf{X}_k)\|_F \|\mathbf{Y}\|_F} = \cos \theta_k, \end{aligned}$$

since $\mathcal{P}_{\mathcal{A}}$ is an orthogonal projection and $\nabla J(\mathbf{X}_k)$ is in the range of \mathcal{A} . In summary, we arrive at

$$\begin{aligned} \|\mathbf{X}_{k+1} - \mathbf{X}_*\|_{\mathcal{A}} &= \|\mathcal{P}_{\mathcal{A}}(\mathbf{X}_{k+1}) - \mathcal{P}_{\mathcal{A}}(\mathbf{X}_*)\|_{\mathcal{A}} \\ &= \|\tilde{\mathbf{X}}_{k+1} - \mathcal{P}_{\mathcal{A}}(\mathbf{X}_*)\|_{\mathcal{A}} \\ &\leq \left(\frac{\tilde{\kappa}_{k, \text{eff}} - 1}{\tilde{\kappa}_{k, \text{eff}} + 1}\right) \|\mathcal{P}_{\mathcal{A}}(\mathbf{X}_k) - \mathcal{P}_{\mathcal{A}}(\mathbf{X}_*)\|_{\mathcal{A}} \\ &= \left(1 - \frac{2}{\tilde{\kappa}_{k, \text{eff}} + 1}\right) \|\mathbf{X}_k - \mathbf{X}_*\|_{\mathcal{A}}. \quad (14) \end{aligned}$$

The contraction factor in (14) looks considerably better than the one in (10) as it features no square root. However, a closer inspection reveals that this is only true for small θ_k . For $\theta_k \rightarrow \pi/2$ the improvement quickly becomes negligible. Figure 1 shows the ratios between the logarithms of both factors for different κ as functions of $\theta \in [0, \pi/2)$. Since, by Lemma 1, the provable bound $\cos \theta \geq 1/\sqrt{n_1 \cdots n_{d-1}}$ for rank-one truncations of the gradient is quite small already for tensors of moderate size and dimension, we stick to the simpler rate (10) in the following.

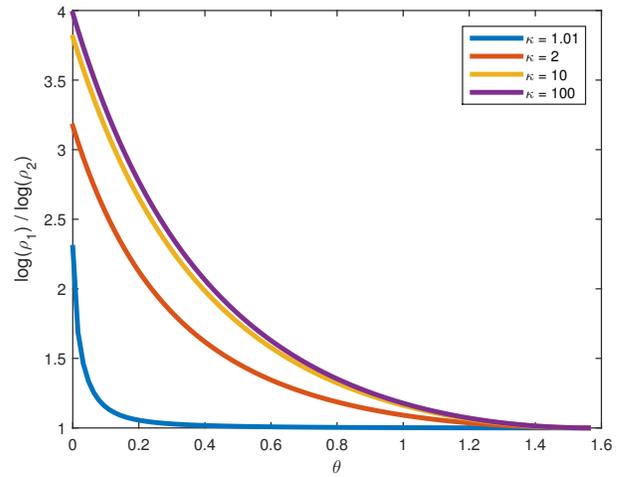


Fig. 1. Ratio $\log(\rho_1)/\log(\rho_2)$ of contraction factors $\rho_1 = \left(1 - \frac{\cos^2 \theta}{\kappa}\right)^{1/2}$ and $\rho_2 = 1 - \frac{2}{\left(\frac{1 + \sin \theta}{1 - \sin \theta}\right) \kappa + 1}$ as function of $\theta \in [0, \pi/2)$ for some κ .

IV. KRONECKER STRUCTURED LINEAR EQUATIONS IN HILBERT SPACE

In the case that $\nabla J(\mathbf{X}_k)$ has itself low rank, the estimation of $\cos \theta_k$ by μ_{n_1, \dots, n_d} can be too rough. Better results may be obtained from a modification of Lemma 1.

Lemma 3. *Let the tensor \mathbf{X} have rank at most r . Then*

$$\langle \mathbf{X}, \mathbf{R}_1(\mathbf{X}) \rangle_F \geq \frac{1}{\sqrt{r^{d-1}}} \|\mathbf{X}\|_F \|\mathbf{R}_1(\mathbf{X})\|_F.$$

Proof. The assumption implies that \mathbf{X} lies in a tensor product subspace which is isomorphic to $\mathbb{R}^{r \times \dots \times r}$. In this space, we can apply Lemma 1 to obtain the result. \square

By construction, the iterates \mathbf{X}_k of algorithm (7) have rank at most $k + \text{rank}(\mathbf{X}_0)$. Consequently, Lemma 3 will be useful for problems, where a low rank of \mathbf{X}_k implies a comparably low rank of $\nabla J(\mathbf{X}_k)$. Our main point now is that this property holds for quadratic problems (5) arising from linear equations with Kronecker-structured operators

$$\mathcal{A} = \sum_{j=1}^R A_j^{(1)} \otimes \dots \otimes A_j^{(d)}, \quad (15)$$

and low-rank hand right hand side

$$\mathbf{B} = \sum_{j=1}^s b_j^{(1)} \circ \dots \circ b_j^{(d)}. \quad (16)$$

Lemma 4. *Assuming (15) and (16), if $\text{rank}(\mathbf{X}_k) \leq k$, then*

$$\text{rank}(\mathcal{A}(\mathbf{X}_k) - \mathbf{B}) \leq Rk + s.$$

Proof. An operator $A^{(1)} \otimes \dots \otimes A^{(d)}$ acts on rank-one tensors $u^{(1)} \circ \dots \circ u^{(d)}$ via $(A^{(1)}u^{(1)}) \circ \dots \circ (A^{(d)}u^{(d)})$. The claim is now trivial by superposition. \square

A. *Matrix case $d = 2$*

When $d = 2$, the Kronecker-structured equation is a linear matrix equation

$$\left(\sum_{j=1}^R A_j^{(1)} \otimes A_j^{(2)} \right) (\mathbf{X}) = \sum_{j=1}^R A_j^{(1)} \mathbf{X} (A_j^{(2)})^T = \mathbf{B}, \quad (17)$$

which contains notable special cases like the Sylvester and Lyapunov equations. For these equations it is known that for low-rank right-hand side \mathbf{B} , the solution \mathbf{X} will have super-algebraically decaying singular values [23]. In the general case of an operator (15) this is not so clear. Of course, when $R = 1$ and $A_1^{(1)} \otimes A_1^{(2)}$ is invertible, then the solution of (17) has the same rank as \mathbf{B} . But can we say anything about the case $R = 2$, for instance? In Figure 2 we plotted the singular values of a solution for $R = 2$ and s when all $A_j^{(i)}$ and $b^{(i)}$ have been randomly generated. It suggests an exponential decay of the singular values of the solution, however, with slower exponent for growing space dimension. In the infinite dimensional Hilbert space setting, we expect a sub-exponential or algebraic rate of decay. We now attempt to obtain such an asymptotic statement in a constructive way by using the

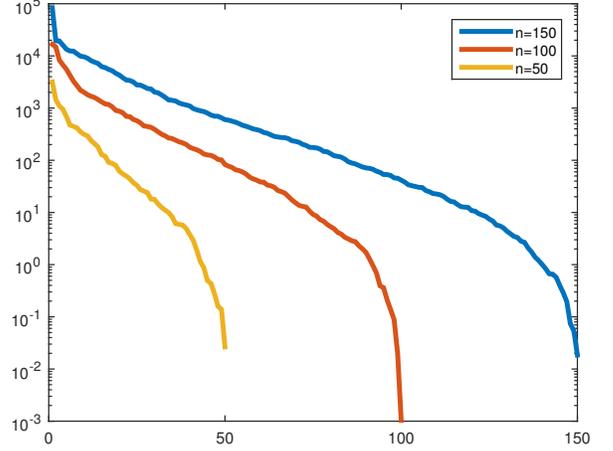


Fig. 2. Singular values of solutions to randomly generated matrix equations (17) of Kronecker rank $R = 2$ in $\mathbb{R}^{n \times n}$, with $s = \text{rank}(\mathbf{B}) = 1$.

method of rank-one truncated gradient descent. The logic here is that the decay of singular values of the solution is directly related to its approximability by matrices \mathbf{X}_k of growing rank k , e.g., by the sequence produced with the algorithm (7).

Theorem 5. *Assume that \mathcal{A} and \mathbf{B} are of the form (15) and (16), respectively, and that \mathcal{A} is positive semidefinite (which is the case when all $A_j^{(i)}$ are positive semidefinite). Let (\mathbf{X}_k) denote the sequence generated by algorithm (7) for the matrix equation (17) when starting with $\mathbf{X}_0 = 0$. Then for all k it holds $\text{rank}(\mathbf{X}_k) \leq k$ and*

$$\|\mathbf{X}_k - \mathbf{X}_*\|_{\mathcal{A}} \leq \left(\frac{1 + s/R}{k + 1 + s/R} \right)^{\frac{1}{2\kappa_{\text{eff}}R}} \|\mathbf{X}_*\|_{\mathcal{A}}.$$

As mentioned, in finite-dimensional space, this sub-linear convergence rate looks worse than the linear rate in Theorem 2. What makes Theorem 5 non-trivial is that the dimensions n_1 and n_2 do not enter. In fact, the result and the following proof are also true in tensor products of Hilbert spaces, and therefore improve on [2], [3], where mere convergence without a rate is proved (albeit for general \mathcal{A} without Kronecker structure).

Proof of Theorem 5. As the iteration starts with $\mathbf{X}_0 = 0$, it holds $\text{rank}(\mathbf{X}_k) \leq k$ by construction, and therefore $\text{rank}(\nabla J(\mathbf{X}_k)) \leq Rk + s$ by Lemma 4. Combining Lemma 3 with (10) we have that

$$\begin{aligned} \|\mathbf{X}_k - \mathbf{X}_*\|_{\mathcal{A}} &\leq \left(1 - \frac{1}{\kappa_{\text{eff}}(Rk + s)} \right)^{1/2} \|\mathbf{X}_{k-1} - \mathbf{X}_*\|_{\mathcal{A}} \\ &\leq \|\mathbf{X}_*\|_{\mathcal{A}} \cdot \prod_{\ell=1}^k \left(1 - \frac{1}{\kappa_{\text{eff}}(R\ell + s)} \right)^{1/2}. \end{aligned}$$

The assertion now follows from the estimate

$$\begin{aligned} \ln \prod_{\ell=1}^k \left(1 - \frac{1}{\kappa_{\text{eff}}(R\ell + s)}\right)^{1/2} &= \frac{1}{2} \sum_{\ell=1}^k \ln \left(1 - \frac{1}{\kappa_{\text{eff}}(R\ell + s)}\right) \\ &\leq \frac{1}{2} \sum_{\ell=1}^k \frac{-1}{\kappa_{\text{eff}}(R\ell + s)} \\ &\leq \frac{-1}{2\kappa_{\text{eff}}R} \ln \left(\frac{k+1+s/R}{1+s/R}\right), \end{aligned}$$

where the second inequality is obtained as usual from $\sum_{\ell=1}^k (\ell + s/R)^{-1} \geq \int_1^{k+1} (x + s/R)^{-1} dx$. \square

B. Tensor case $d \geq 3$

It is interesting to note that, unfortunately, the same argument does not work for tensors of higher order. Lemma 3 will provide the energy error reduction

$$\prod_{\ell=1}^k \left(1 - \frac{1}{\kappa_{\text{eff}}(R\ell + s)^{d-1}}\right)^{1/2}$$

after k steps. When $d \geq 3$, this expression does not converge to zero for $k \rightarrow \infty$.

What we would need are improved estimates for the overlap of rank k tensors with rank one tensors. The proof of Lemma 3 relying on an embedding into $\mathbb{R}^{k \times \dots \times k}$ appears a little bit crude (although, as mentioned earlier, even then it is not known whether the bound sharp [4]). An example illustrates this: when \mathbf{X} is a sum of k orthogonal rank-one tensors, then of course the maximum cosine of an angle with a rank-one tensor is at least $1/\sqrt{k}$ like in the matrix case. In general, when the rank-one terms form a linear independent set of sufficiently good condition, we can give the same bound up to a factor. What would be needed are properties which ensure this during the iteration when building the rank-one truncated steepest descent sequence (7), but we have to leave this as an open problem.

We note that it is nevertheless possible to obtain an algebraic decay rate for solutions of Kronecker-structured equations with elliptic bounded operators in Hilbert space by balancing the linear convergence rate of the untruncated steepest descent with the exponential rank growth, see [10].

V. CONCLUSION

We have reviewed how the convergence of rank truncated steepest descent methods, which appeared in many versions for different low-rank matrix and tensor optimization tasks during the last years, is related to the best rank-one approximation ratio in the tensor space. We derived convergence statements for semidefinite quadratic problems in energy seminorm. Finally, we proved algebraic decay rates for solutions of Kronecker-structured matrix equations which are independent of the space dimension, by estimating the error in the rank-one truncated steepest descent algorithm based on the slow rank growth of residuals.

REFERENCES

- [1] A. Ammar, B. Mokdad, F. Chinesta, and R. Keunings, "A new family of solvers for some classes of multidimensional partial differential equations encountered in kinetic theory modeling of complex fluids," *J. Non-Newtonian Fluid Mech.*, vol. 139, no. 3, pp. 153–176, 2006.
- [2] A. Ammar, F. Chinesta, and A. Falcó, "On the convergence of a greedy rank-one update algorithm for a class of linear systems," *Arch. Comput. Methods Eng.*, vol. 17, no. 4, pp. 473–486, 2010.
- [3] A. Falcó and A. Nouy, "A proper generalized decomposition for the solution of elliptic problems in abstract form by using a functional Eckart-Young approach," *J. Math. Anal. Appl.*, vol. 376, no. 2, pp. 469–480, 2011.
- [4] L. Qi, "The best rank-one approximation ratio of a tensor space," *SIAM J. Matrix Anal. Appl.*, vol. 32, no. 2, pp. 430–442, 2011.
- [5] S. Shalev-Shwartz, A. Gonen, and O. Shamir, "Large-scale convex minimization with a low-rank constraint," in *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, 2011, pp. 329–336.
- [6] S. V. Dolgov and D. V. Savostyanov, "Alternating minimal energy methods for linear systems in higher dimensions," *SIAM J. Sci. Comput.*, vol. 36, no. 5, pp. A2248–A2271, 2014.
- [7] Z. Wang, M.-J. Lai, Z. Lu, W. Fan, H. Davulcu, and J. Ye, "Orthogonal Rank-One Matrix Pursuit for Low Rank Matrix Completion," *SIAM J. Sci. Comput.*, vol. 37, no. 1, pp. A488–A514, 2015.
- [8] M. Tan, I. W. Tsang, L. Wang, B. Vandereycken, and S. J. Pan, "Riemannian pursuit for big matrix recovery," in *Proceedings of the 31st International Conference on Machine Learning (ICML 2014)*, 2014, pp. 1539–1547.
- [9] A. Uschmajew and B. Vandereycken, "Line-search methods and rank increase on low-rank matrix varieties," in *Proceedings of the 2014 International Symposium on Nonlinear Theory and its Applications (NOLTA2014)*, 2014, pp. 52–55.
- [10] D. Kressner and A. Uschmajew, "On low-rank approximability of solutions to high-dimensional operator equations and eigenvalue problems," *arXiv:1406.7026*, 2014.
- [11] V. De Silva and L.-H. Lim, "Tensor rank and the ill-posedness of the best low-rank approximation problem," *SIAM J. Matrix Anal. Appl.*, vol. 30, no. 3, pp. 1084–1127, 2008.
- [12] C. J. Hillar and L.-H. Lim, "Most tensor problems are NP-hard," *J. ACM*, vol. 60, no. 6, pp. Art. 45, 39, 2013.
- [13] T. Kolda and B. Bader, "Tensor decompositions and applications," *SIAM Rev.*, vol. 51, no. 3, pp. 455–500, 2009.
- [14] L. De Lathauwer, P. Comon, B. De Moor, and J. Vandewalle, "High-order power method – Application in Independent Component Analysis," in *Proceedings of the 1995 International Symposium on Nonlinear Theory and its Applications (NOLTA'95)*, 1995, pp. 91–96.
- [15] L. De Lathauwer, B. De Moor, and J. Vandewalle, "On the best rank-1 and rank- (R_1, R_2, \dots, R_N) approximation of higher-order tensors," *SIAM J. Matrix Anal. Appl.*, vol. 21, no. 4, pp. 1324–1342 (electronic), 2000.
- [16] 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.*, vol. 35, no. 3, pp. 1058–1072, 2014.
- [17] A. Uschmajew, "A new convergence proof for the higher-order power method and generalizations," *arXiv:1407.4586*, 2014, to appear in *Pacific J. Optim.*
- [18] L. De Lathauwer, B. De Moor, and J. Vandewalle, "A multilinear singular value decomposition," *SIAM J. Matrix Anal. Appl.*, vol. 21, no. 4, pp. 1253–1278 (electronic), 2000.
- [19] J. Nocedal and S. Wright, *Numerical Optimization*. New York: Springer, 2006.
- [20] P.-A. Absil and K. A. Gallivan, "Accelerated line-search and trust-region methods," *SIAM J. Numer. Anal.*, vol. 47, no. 2, pp. 997–1018, 2009.
- [21] L. V. Kantorovich, "Functional analysis and applied mathematics," *Uspehi Matem. Nauk (N.S.)*, vol. 3, no. 6(28), pp. 89–185, 1948.
- [22] H. Munthe-Kaas, "The convergence rate of inexact preconditioned steepest descent algorithm for solving linear systems," Tech. Rep. Numerical Analysis Project Manuscript NA-87-04, 1987. [Online]. Available: <http://i.stanford.edu/pub/cstr/reports/na/m/87/04/NA-M-87-04.pdf>
- [23] L. Grasedyck, "Existence and computation of low Kronecker-rank approximations for large linear systems of tensor product structure," *Computing*, vol. 72, no. 3–4, pp. 247–265, 2004.