

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

**Kronecker product approximation of
operators in spectral norm via
alternating SDP**

by

*Mareike Dressler, André Uschmajew, and
Venkat Chandrasekaran*

Preprint no.: 21

2022



KRONECKER PRODUCT APPROXIMATION OF OPERATORS IN SPECTRAL NORM VIA ALTERNATING SDP

MAREIKE DRESSLER, ANDRÉ USCHMAJEV, AND VENKAT CHANDRASEKARAN

ABSTRACT. The decomposition or approximation of a linear operator on a matrix space as a sum of Kronecker products plays an important role in matrix equations and low-rank modeling. The approximation problem in Frobenius norm admits a well-known solution via the singular value decomposition. However, the approximation problem in spectral norm, which is more natural for linear operators, is much more challenging. In particular, the Frobenius norm solution can be far from optimal in spectral norm. We describe an alternating optimization method based on semidefinite programming to obtain high-quality approximations in spectral norm, and we present computational experiments to illustrate the advantages of our approach.

1. INTRODUCTION AND PROBLEM STATEMENT

Let $T: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ be a linear operator acting on the space of real $m \times n$ matrices. Every such operator can be decomposed into a sum of tensor product operators $A_j \otimes B_j$ such that

$$(1.1) \quad T(X) = \sum_{j=1}^r A_j X B_j^T,$$

for some matrices $A_j \in \mathbb{R}^{m \times m}$ and $B_j \in \mathbb{R}^{n \times n}$. By choosing a suitable $mn \times mn$ matrix representation of the linear operator T , the above decomposition can be written as

$$(1.2) \quad T = \sum_{j=1}^r A_j \otimes B_j,$$

where now $A_j \otimes B_j$ is the standard Kronecker product of matrices A_j and B_j . In the following we will not distinguish between these two interpretations of T .

The smallest r needed for a decomposition (1.2) to exist is often called the *Kronecker rank* of T . The Kronecker rank has the (sharp) upper bound

$$r \leq \min(m^2, n^2)$$

and a minimal decomposition can be found by applying a rank revealing decomposition to a suitable reshape \widehat{T} of T into an $m^2 \times n^2$ matrix [VLP93, KMP00], see section 2. In [KMP00], the Kronecker rank has also been called the Sylvester index of T .

As already noted in [VLP93, KMP00], linear operators with small Kronecker rank play an important role in numerical linear algebra and scientific computing. In particular, solving linear matrix equations such as

$$(1.3) \quad T(X) = Y$$

2020 *Mathematics Subject Classification*. Primary: 47A58, 90C22; Secondary: 65F45.

Key words and phrases. Operator approximation, Kronecker product, semidefinite programming, alternating optimization, matrix equations.

is numerically challenging if sought matrices X are large. Compared to a general operator T with an unstructured $mn \times mn$ representation T , an operator of the form (1.1) has the advantage that, if r is not too large, its evaluation $T(X)$ on a matrix X can be computed by just a few left and right matrix multiplications. This allows to implement iterative methods for solving equations such as (1.3) at (relatively) small computational cost.

There is, however, a second, in a sense even more important reason, why operators with small Kronecker rank are of interest in matrix equations: in practice, it can often be observed that the solution X to an equation (1.3), where T has small Kronecker rank and the right-hand side Y is a low-rank matrix, can itself be well approximated by low-rank matrices. Thus a low-rank model for $X \approx UV^\top$ is then justified which in turn makes the application of T via left and right multiplications even more efficient. This is the basis for many efficient low-rank solvers for matrix equations; see [Sim16] for an overview. The numerical observation can be rigorously proven for certain matrix equations, most notably for Sylvester-type equations

$$T(X) = AX + XB^\top = Y,$$

when A and B are symmetric positive definite matrices. It can be shown that the inverse T^{-1} of this operator can be approximated in operator norm by a sum of k Kronecker products with exponential (or subexponential) error decay in k . Specifically, such an approximation can be obtained using an exponential sum approximation of T^{-1} [Gra04]. Since the application of such an approximate inverse increases the rank of Y by a factor of at most k , the solution X can be well approximated by low-rank matrices. This even works for tensor versions of Sylvester-type operators. We refer to [BM02, Tyr04, HK06a, HK06b, OST09, Hac19] for applications of Kronecker product approximations of operators and low-rank tensor calculus in scientific computing.

Motivated by the above considerations, we are interested in the general question of *approximating* a matrix $T \in \mathbb{R}^{mn \times mn}$ or its inverse by an operator of low Kronecker rank in operator norm. Given T and a rank bound k , we hence consider the optimization problem

$$(1.4) \quad \min_{A_j, B_j} \left\| T - \sum_{j=1}^k A_j \otimes B_j \right\|_2.$$

Here $\|\cdot\|_2$ denotes the spectral norm of $mn \times mn$ matrices, which corresponds to the operator norm on $\mathbb{R}^{m \times n}$ with respect to the Frobenius norm. In particular,

$$\left\| T - \sum_{j=1}^k A_j \otimes B_j \right\|_2 = \max_{\|X\|_F=1} \left\| T(X) - \sum_{j=1}^k A_j X B_j^\top \right\|_F.$$

We will also consider the approximation of inverse operators T^{-1} using a different cost function (section 3.4).

It is known that if in (1.4) we replace the spectral norm by the Frobenius norm, the approximation problem admits (in principle) a closed-form solution via a singular value decomposition (SVD) of a reshaped $m^2 \times n^2$ matrix \widehat{T} , which was already mentioned before. This method has been worked out in [VLP93] and will be explained in section 2. The main underlying reason that such an approach works for the Frobenius norm is that this norm is invariant under reshaping a matrix. For the spectral norm, this is, however, not the case. So while indeed the SVD also provides a best low-rank approximation in spectral norm of \widehat{T} as an operator from \mathbb{R}^{n^2} to \mathbb{R}^{m^2} , it will not provide the optimal low Kronecker rank

approximation for T in spectral norm. To our knowledge, the problem (1.4) has not been addressed in the literature in this general form.

In this work, we consider an alternating optimization method for solving (1.4) based on semidefinite programming. The underlying idea is based on the well-known fact that the best approximation of a given matrix in spectral norm by an element from an affine linear matrix subspace can be found by solving a semidefinite program (SDP), and is therefore computable in polynomial-time to a desired accuracy. Due to the bilinearity of the Kronecker product, an alternating optimization approach for the unknown matrices A_j and B_j in problem (1.4) leads to a sequence of approximation problems on linear matrix subspaces, which therefore can be solved via SDPs. Furthermore, by adding regularization one can even guarantee that the solutions to the subproblems are unique and bounded, which plays a role in our convergence analysis.

Such an approach has computational limitations when m or n become large, as it requires solving rather large scale SDPs. However, as a proof of concept for our numerical experiments demonstrate that applying our algorithm can provide significantly better low Kronecker rank approximations to operators T than the SVD method in certain settings.

The paper is organized as follows. In section 2, we review the SVD method for computing low Kronecker rank approximations and provide an example showcasing that this method can be far from optimal with regard to the spectral norm. Section 3 presents the alternating SDP approach for solving the spectral norm approximation problem (1.4). We also discuss a regularized version of the problem in section 3.2, which leads to the general form of Algorithm 2. Moreover, we briefly touch upon different updating choices for the variables in the alternating SDP approach (section 3.3) and how the problem of approximating inverses could be handled (section 3.4). In section 4, we provide some convergence statements based on results for bi-convex optimization from the literature. Finally, in section 5, we present computational experiments illustrating the benefit of our proposed approach. We conclude with a brief discussion and an outlook on potential future work, see section 6.

2. THE SVD METHOD

The SVD solution to the Kronecker product approximation problem has been proposed in [VLP93], see also [VL00]. It works by rearranging the matrix $T \in \mathbb{R}^{mn \times mn}$ into another matrix $\hat{T} \in \mathbb{R}^{m^2 \times n^2}$. Specifically, from a block partition

$$T = \begin{bmatrix} T_{11} & T_{12} & \cdots \\ T_{21} & \ddots & \\ \vdots & & \end{bmatrix}$$

with m^2 blocks $T_{\mu\nu} \in \mathbb{R}^{n^2}$ one constructs the matrix \hat{T} with rows $\text{vec}(T_{\mu\nu})^\top$ in the ordering

$$\hat{T} = \begin{bmatrix} -\text{vec}(T_{11})^\top - \\ -\text{vec}(T_{21})^\top - \\ \vdots \\ -\text{vec}(T_{12})^\top - \\ \vdots \end{bmatrix}.$$

It is then easy to see that the Kronecker product decomposition (1.2) of T is equivalent to a decomposition

$$\widehat{T} = \sum_{j=1}^r \text{vec}(A_j) \text{vec}(B_j)^\top$$

into rank-one matrices. This shows that a (usual) low-rank decomposition or approximation of \widehat{T} yields decompositions or approximations of T by sums of Kronecker products. Clearly the transformation $T \leftrightarrow \widehat{T}$ is an isometry between $\mathbb{R}^{mn \times mn}$ and $\mathbb{R}^{m^2 \times n^2}$ in Frobenius norm, since it only rearranges entries of T .

The required low-rank approximation of \widehat{T} can be achieved for instance using singular value decomposition (SVD). This provides the method displayed as Algorithm 1, where $\text{mat}(\cdot)$ denotes the inverse of the $\text{vec}(\cdot)$ operation. In the algorithm, we included an optional step for optimizing the scaling of the truncated SVD with respect to spectral norm approximation. It could be solved using semidefinite programming similar to the methods derived in section 3.

Algorithm 1 SVD method [VLP93]

Input: Matrix $T \in \mathbb{R}^{mn \times mn}$, dimensions m, n , approximation rank k

Output: Kronecker rank k approximation T_{svd} of T

Construct the matrix \widehat{T} from T .

Compute SVD: $\widehat{T} = U \Sigma V^\top = \sum_{j=1}^r \sigma_j u_j v_j^\top$.

Set $A_j = \text{mat}(\sigma_j^{1/2} u_j)$, $B_j = \text{mat}(\sigma_j^{1/2} v_j)$.

Set $\alpha_1 = \dots = \alpha_k = 1$.

Optional: Find $(\alpha_1, \dots, \alpha_k) = \underset{\alpha_j \in \mathbb{R}}{\text{argmin}} \left\| T - \sum_{j=1}^k \alpha_j A_j \otimes B_j \right\|_2$. ▷ Optimal scaling

return $T_{\text{svd}} := \sum_{j=1}^k \alpha_j A_j \otimes B_j$

It is now interesting to observe that the approximation $\sum_{j=1}^k A_j \otimes B_j$ of T obtained by the above procedure is optimal in Frobenius norm, that is, it solves a version of problem (1.4) in Frobenius norm instead of the spectral norm. The reason being that $\sum_{j=1}^k \sigma_j u_j v_j^\top$ is an optimal rank- k approximation of \widehat{T} in Frobenius norm and the transformation $T \leftrightarrow \widehat{T}$ is an isometry. More generally, the truncated SVD of \widehat{T} actually is optimal in any unitarily invariant norm on $\mathbb{R}^{m^2 \times n^2}$, in particular, it is the best approximation in spectral norm on $\mathbb{R}^{m^2 \times n^2}$. Therefore, $\sum_{j=1}^k A_j \otimes B_j$ is also an optimal approximation of T in the norm defined via

$$\|T\|_2 := \|\widehat{T}\|_2$$

on $\mathbb{R}^{mn \times mn}$. In general, this norm (and other possible ones obtained in this way) does not equal the spectral norm in $\mathbb{R}^{mn \times mn}$ though.

Correspondingly, the SVD approach will usually not lead to optimal solutions of the initial approximation problem for T in spectral norm, even if the scaling is optimized. It is in fact not so difficult to construct such counterexamples.

Example 2.1. Let

$$(2.1) \quad T = \sigma_1 A_1 \otimes B_1 + A_2 \otimes B_2,$$

such that $\langle A_1, A_2 \rangle = 0$, $\langle B_1, B_2 \rangle = 0$, $\|A_1\|_F = \|A_2\|_F = \|B_1\|_F = \|B_2\|_F = 1$, and $\sigma_1 > 1$. Here $\langle \cdot, \cdot \rangle$ denotes the Frobenius inner product. Assume the goal is to find a Kronecker rank one approximation, i.e., $k = 1$. Note that by construction (2.1) already corresponds to an SVD of \widehat{T} . Hence Algorithm 1 proposes $T_{\text{svd}} = \alpha_1 \sigma_1 A_1 \otimes B_1$ as an approximation to T . The error is

$$T - T_{\text{svd}} = \sigma_1(1 - \alpha_1)A_1 \otimes B_1 + A_2 \otimes B_2.$$

We now choose $A_2 = aa^\top$ and $B_2 = bb^\top$ to be rank-one matrices (the symmetry is not essential), and A_1, B_1 to be matrices with $A_1 a = 0$ and $B_1 b = 0$. We then have for $X = ab^\top$ that

$$(T - T_{\text{svd}})(ab^\top) = \sigma_1(1 - \alpha_1)(A_1 a)(B_1 b)^\top + aa^\top ab^\top bb^\top = ab^\top,$$

for any choice of α_1 since $a^\top a = b^\top b = 1$. This shows that

$$\|T - T_{\text{svd}}\|_2 \geq 1.$$

(Actually, when α_1 is picked optimally as assumed in Algorithm 1, one also has $\|T - T_{\text{svd}}\|_2 \leq \|A_2 \otimes B_2\|_2 = 1$, hence equality.) On the other hand, using $A_2 \otimes B_2$ as an approximation, would lead to

$$\|T - A_2 \otimes B_2\|_2 = \|\sigma_1 A_1 \otimes B_1\|_2 = \sigma_1 \|A_1\|_2 \|B_1\|_2$$

using properties of the Kronecker product. If we take A_1 and B_1 to be of rank $m - 1$ and $n - 1$, respectively, with identical (nonzero) singular values $\frac{1}{\sqrt{m-1}}$ and $\frac{1}{\sqrt{n-1}}$, we obtain

$$(2.2) \quad \|T - A_2 \otimes B_2\|_2 = \sigma_1 \frac{1}{\sqrt{(m-1)(n-1)}}.$$

For σ_1 close to one and m, n large this is much smaller than one, demonstrating that the SVD solution can be rather bad. Note that in this construction (assuming $\sigma_1 \frac{1}{\sqrt{(m-1)(n-1)}} < 1$), one can show that $\|T\|_2 = 1$, so the above estimates measure the relative error. The SVD solution is hence not better than taking the zero approximation.

3. ALTERNATING SDP APPROACH

In this main section we describe an approach for solving the problem (1.4) based on alternating optimization and semidefinite programming.

3.1. Basic idea. We use the fact that a minimization of the operator norm on an affine linear space of matrices can be turned into an SDP, see [BV04, Section 4.6.3]. This is based on the observation that for a matrix an inequality $\|S\|_2 \leq \tau$ for the spectral norm is equivalent with $S^\top S \preceq \tau^2 I$, and I being the identity matrix. Hence the problem (1.4) can be first reformulated as

$$\begin{aligned} & \min_{\tau \geq 0, A_j, B_j} \tau \\ \text{s.t.} \quad & \left(T - \sum_{j=1}^k A_j \otimes B_j \right)^\top \left(T - \sum_{j=1}^k A_j \otimes B_j \right) \preceq \tau^2 I. \end{aligned}$$

By applying Schur complement this turns into

$$(3.1) \quad \begin{aligned} & \min_{\tau, A_j, B_j} \tau \\ \text{s.t.} \quad & M(\tau, A, B) \succeq 0, \end{aligned}$$

with

$$M(\tau, A, B) := \begin{bmatrix} \tau I & T - \sum_{j=1}^k A_j \otimes B_j \\ (T - \sum_{j=1}^k A_j \otimes B_j)^\top & \tau I \end{bmatrix}.$$

Note that the condition $\tau \geq 0$ has been dropped in this second version since it is automatically implied by the positive semidefiniteness of $M(\tau, A, B)$.

Since $M(\tau, A, B)$ is affine linear in each of the (block) variables τ , $A = (A_1, \dots, A_k)$, and $B = (B_1, \dots, B_k)$, we have achieved that the optimization problem (3.1) for each single of these block variables, when the others are kept fixed, is an SDP. In fact, the problem for the auxiliary variable τ can be subsumed into the other ones: we have an SDP in (τ, A) when B is fixed, or in (τ, B) when A is fixed. As SDPs can be solved in polynomial-time to a desired accuracy, this suggests an alternating optimization strategy to tackle the initial problem (1.4) through a sequence of SDPs: we first fix $B = (B_1, \dots, B_k)$ and optimize $A = (A_1, \dots, A_k)$ (and τ), then we fix A to find the optimal B (and τ). We describe a generalization of this procedure in the sequel with additional regularization parameters in Algorithm 2. By construction the subproblems for A and B always admit at least one solution, since from (1.4) it follows, that they are equivalent with finding a best approximation of T in operator norm on the *linear* subspaces $\{\sum_{j=1}^k A_j \otimes B_j : A_1, \dots, A_k \in \mathbb{R}^{m \times m}\}$ and $\{\sum_{j=1}^k A_j \otimes B_j : B_1, \dots, B_k \in \mathbb{R}^{n \times n}\}$ of $\mathbb{R}^{mn \times mn}$, respectively.

3.2. Adding regularization. While the subproblems for A and B in the alternating optimization approach outlined above are guaranteed to admit optimal solutions, neither uniqueness nor boundedness of these solutions seem to be easy to guarantee in general. As we will see in section 4, this poses a problem in the convergence analysis of alternating block minimization problems. Both issues, however, can be addressed by adding regularization to the approach outlined in section 3.1, albeit at the expense of larger SDPs.

For convenience, we define

$$F(A, B) := \left\| T - \sum_{j=1}^k A_j \otimes B_j \right\|_2,$$

where again $A = (A_1, \dots, A_k)$ and $B = (B_1, \dots, B_k)$. Instead of (1.4) we then consider

$$(3.2) \quad \min F_{\lambda, \mu}(A, B) = F(A, B) + \lambda \sum_{j=1}^k \|A_j\|_F^2 + \mu \sum_{j=1}^k \|B_j\|_F^2$$

with regularization parameters $\lambda, \mu \geq 0$. If $\lambda > 0$, then for fixed B the function $A \mapsto F_{\lambda, \mu}(A, B)$ is strictly convex and coercive. A similar statement holds for the function $B \mapsto F_{\lambda, \mu}(A, B)$ in the case $\mu > 0$. Therefore, if λ and μ are positive, the subproblems now admit unique solutions. Moreover, the sequences of solutions can easily be bounded (since the function values of $F_{\lambda, \mu}$ are decreasing and F is bounded from below).

We claim that in the regularized problem the subproblems for A and B (and likewise for single A_j and B_j) can again be turned into SDPs as follows. Assume B is fixed and A

should be optimized. Introducing slack variables $\gamma_0, \gamma_1, \dots, \gamma_k$ we first rewrite (3.2) as

$$\begin{aligned} & \min_{\tau, \gamma_0, \gamma_1, \dots, \gamma_k, A} \quad \tau \\ \text{s.t.} \quad & F(A, B) \leq \gamma_0, \\ & \lambda \|A_j\|_F^2 \leq \gamma_j, \quad j = 1, \dots, k, \\ & \gamma_0 + \gamma_1 + \dots + \gamma_k \leq \tau. \end{aligned}$$

As in (3.1), the first constraint $F(A, B) \leq \gamma_0$ can be turned into $M(\gamma_0, A, B) \succeq 0$. The constraints $\lambda \|A_j\|_F^2 \leq \gamma_j$ are equivalent to

$$(3.3) \quad N_\lambda(\gamma_j, A_j) := \begin{bmatrix} \gamma_j & \sqrt{\lambda} \text{vec}(A_j) \\ \sqrt{\lambda} \text{vec}(A_j)^\top & I \end{bmatrix} \succeq 0,$$

due to Schur complement. Finally note that the last constraint is $\tau - \gamma_0 - \dots - \gamma_k \geq 0$. We hence arrive at the following SDP:

$$(3.4) \quad \begin{aligned} & \min_{\tau, \gamma_0, \gamma_1, \dots, \gamma_k, A} \quad \tau \\ \text{s.t.} \quad & M(\gamma_0, A, B) \succeq 0, \\ & N_\lambda(\gamma_j, A_j) \succeq 0, \quad j = 1, \dots, k, \\ & \tau - \gamma_0 - \dots - \gamma_k \geq 0. \end{aligned}$$

Note that the constraints automatically imply $\gamma_0, \gamma_1, \dots, \gamma_k \geq 0$, and thus also $\tau \geq 0$.

When A is fixed and we aim to optimize for B , we have to solve the analogous problem

$$(3.5) \quad \begin{aligned} & \min_{\tau, \delta_0, \delta_1, \dots, \delta_k, B} \quad \tau \\ \text{s.t.} \quad & M(\delta_0, A, B) \succeq 0, \\ & N_\mu(\delta_j, B_j) \succeq 0, \quad j = 1, \dots, k, \\ & \tau - \delta_0 - \dots - \delta_k \geq 0. \end{aligned}$$

Here we have slightly abused notation since $N_\mu(\delta_j, B_j)$ might be of different size than defined in (3.3). Of course, when λ or μ are zero, the corresponding constraints in (3.4) and (3.5) can be simply omitted and the problem reduces to the basic idea explained in section 3.1.

The resulting alternating SDP algorithm (ASDP) is summarized as Algorithm 2.

Algorithm 2 Alternating SDP method (ASDP)

Input: Operator T , dimensions m, n , initial $B_0 = (B_{01}, \dots, B_{0k})$, approximation rank k , regularization parameters $\lambda, \mu \geq 0$, number of iterations N_{outer}

Output: Kronecker rank k approximation T_{sdp} of T in spectral norm

$B \leftarrow B_0$

for N_{outer} iterations **do**

 Update A by solving the SDP (3.4) with current B fixed.

 Update B by solving the SDP (3.5) with current A fixed.

end for

return $T_{\text{sdp}} := \sum_{j=1}^k A_j \otimes B_j$

3.3. Alternative update rules. In Algorithm 2 the block variables $A = (A_1, \dots, A_k)$ and $B = (B_1, \dots, B_k)$ in problem (3.2) are each fully updated at a time. It is of course possible to devise different update strategies by optimizing for single matrices A_j and B_j , which may be computationally beneficial. In particular, by writing

$$T - \sum_{j=1}^k A_j \otimes B_j = \left(T - \sum_{i \neq j} A_i \otimes B_i \right) - A_j \otimes B_j$$

we could update a single pair (A_j, B_j) by calling Algorithm 2 with $\tilde{T} = T - \sum_{i \neq j} A_i \otimes B_i$ and $k = 1$ with a number N_{inner} of inner iterations. This process can then be iterated over $j = 1, \dots, k$, for instance in a cyclic manner. The “previous” values for A_j or B_j could be reused as starting guesses.

Now depending on the choice of N_{inner} such a scheme allows for different interpretations. When $N_{\text{inner}} = 1$, it is an alternating optimization method for the ordering $(A_1, B_1, A_2, B_2, \dots, A_k, B_k)$ of block variables. When N_{inner} is large, we can imagine the inner cycle as computing a best Kronecker rank-one approximation $A_j \otimes B_j$ to the current residual \tilde{T} by means of alternating optimization (although it is not guaranteed to find the global optimum). One loop over $j = 1, \dots, k$ can then be interpreted as a “greedy” method for finding a best approximation of Kronecker rank k .

However, in our numerical experiments, we found no substantial advantage of these variations compared to Algorithm 2 for the considered scenarios, and therefore we do not present numerical results for them in section 5.

3.4. Approximation of inverse operators. As outlined in the introduction, in several applications one is actually interested in the approximation of the inverse T^{-1} of a given operator T by sums of Kronecker products. Since the goal is usually to find such an approximation without forming T^{-1} , we cannot apply the above methods directly. A natural way is to consider the modified optimization problem

$$(3.6) \quad \min_{A_j, B_j} \left\| I - T \cdot \left(\sum_{j=1}^k A_j \otimes B_j \right) \right\|_2$$

instead, and tackle it via alternating optimization. Here \cdot is the matrix product. Note that doing this for the Frobenius norm would lead to an alternating least squares algorithm. For the spectral norm, we rely once again on semidefinite programming.

The derivation of an alternating SDP method for (3.6) is almost analogous to the one developed above. Instead of (3.1), we now have to solve

$$(3.7) \quad \begin{aligned} \min_{\tau, A_j, B_j} \quad & \tau \\ \text{s.t.} \quad & M_{\text{inv}}(\tau, A, B) \succeq 0, \end{aligned}$$

with

$$(3.8) \quad M_{\text{inv}}(\tau, A, B) := \begin{bmatrix} \tau I & I - T \cdot (\sum_{j=1}^k A_j \otimes B_j) \\ (I - T \cdot (\sum_{j=1}^k A_j \otimes B_j))^{\top} & \tau I \end{bmatrix}.$$

Since M_{inv} is affine linear in τ , A , and B , this problem can again be tackled via alternating semidefinite programs.

It is also possible to include regularization in the same way as in (3.2). The resulting SDP problems in the alternating optimization approach read the same as (3.4) and (3.5), except with $M(\tau, A, B)$ being replaced with $M_{\text{inv}}(\tau, A, B)$.

Note that in such an implicit formulation to approximate the inverse one does in principle not need T in the full format, but only the ability to apply T to operators of the form $\sum_{j=1}^k A_j \otimes B_j$ several times. While without any structure this poses a potential computational bottleneck, it can be handled in the situation when T itself is a sum of Kronecker products,

$$T = \sum_{J=1}^r C_J \otimes D_J.$$

Then $M_{\text{inv}}(\tau, A, B)$ becomes

$$(3.9) \quad M_{\text{inv}} = \begin{bmatrix} \tau I & I - \sum_{J=1}^r \sum_{j=1}^k (C_J A_j) \otimes (D_J B_j) \\ (I - \sum_{J=1}^r \sum_{j=1}^k (C_J A_j) \otimes (D_J B_j))^{\top} & \tau I \end{bmatrix}.$$

If k and r are small, the double sums can be efficiently computed.

We point out that the implicit formulation (3.6) for approximating inverses comes at a price in case of badly conditioned operators. On the one hand, (3.6) gives us full control on the relative approximation error of T^{-1} as follows: assume we have achieved

$$\left\| I - T \cdot \left(\sum_{j=1}^k A_j \otimes B_j \right) \right\|_2 \leq \varepsilon,$$

then

$$\frac{\left\| T^{-1} - \sum_{j=1}^k A_j \otimes B_j \right\|_2}{\|T^{-1}\|_2} \leq \frac{\|T^{-1}\|_2 \left\| I - T \cdot \left(\sum_{j=1}^k A_j \otimes B_j \right) \right\|_2}{\|T^{-1}\|_2} \leq \varepsilon.$$

On the other hand, when the goal is to compute approximate solutions of linear matrix equations such as (1.3) by replacing T^{-1} with its approximation, then the absolute error or the relative error with respect to $\|T\|_2$ would be of more relevance. We have

$$\frac{\left\| T^{-1} - \sum_{j=1}^k A_j \otimes B_j \right\|_2}{\|T\|_2} \leq \frac{\|T^{-1}\|_2}{\|T\|_2} \varepsilon,$$

so this error estimate deteriorates with a bad condition number. Since the case of ill conditioned operators appears frequently in numerical analysis, this may put some additional limitations on this approach. In section 5.3 we conduct some numerical experiments for approximating inverses based on formulation (3.6).

4. CONVERGENCE

Algorithm 2 realizes a block coordinate optimization method for the cost function

$$F_{\lambda, \mu}(A, B) = \left\| T - \sum_{j=1}^k A_j \otimes B_j \right\|_2 + \lambda \sum_{j=1}^k \|A_j\|_F^2 + \mu \sum_{j=1}^k \|B_j\|_F^2$$

in problem (3.2), by sequentially setting block variables to restricted global minima:

$$(4.1) \quad \begin{aligned} A^{(\ell+1)} &\leftarrow \underset{A}{\operatorname{argmin}} F_{\lambda,\mu}(A, B^{(\ell)}), \\ B^{(\ell+1)} &\leftarrow \underset{B}{\operatorname{argmin}} F_{\lambda,\mu}(A^{(\ell+1)}, B). \end{aligned}$$

The convergence of such block coordinate methods is in general not easy to deduce and requires specific assumptions. For differentiable cost functions it can be shown that if the subproblems are guaranteed to admit unique minima, then all accumulation points of the iterates must be critical points; see, e.g., [Ber16, Proposition 3.7.1] or [Tse01]. However, $F_{\lambda,\mu}$ above is not a differentiable function, since the spectral norm is not. The convergence of (4.1) in such a scenario has been less studied to our knowledge, an exception being the work [GPK07] on continuous biconvex problems on which we rely in what follows.

Clearly, $F_{\lambda,\mu}$ is biconvex, which means that the restriction to one of the block variables A or B is always a convex function. When $\lambda > 0$ and $\mu > 0$, then these restricted functions are even strictly convex and coercive, which means that the subproblems have unique minima. For clarity, we note again in the case $\lambda = 0$ or $\mu = 0$ the solvability of the substeps is still ensured, since, e.g., the update of A realizes a best approximation of T in operator norm on the linear subspace of all $\sum_{j=1}^k A_j \otimes B_j$ with the B being fixed. However, we cannot guarantee uniqueness in this case.

Fixed points (A^*, B^*) of the procedure (4.1) are characterized by the properties

$$F_{\lambda,\mu}(A^*, B^*) \leq F_{\lambda,\mu}(A, B^*) \quad \text{and} \quad F_{\lambda,\mu}(A^*, B^*) \leq F_{\lambda,\mu}(A^*, B)$$

for all A and B , respectively. Such points are called *partial optima* (of the function $F_{\lambda,\mu}$). In [GPK07, Theorem 4.9] it is shown that for continuous biconvex functions any accumulation point (A^*, B^*) of a sequence $(A^{(\ell)}, B^{(\ell)})$ generated by (4.1) is indeed a partial optimum (i.e., a fixed point of (4.1)) under the following assumptions:

- (i) the iterates $(A^{(\ell)}, B^{(\ell)})$ are bounded;
- (ii) at every accumulation point (A^*, B^*) it holds that either $A \mapsto F_{\lambda,\mu}(A, B^*)$ or $B \mapsto F_{\lambda,\mu}(A^*, B)$ have a unique minimizer.

In addition, if in (ii) both $A \mapsto F_{\lambda,\mu}(A, B^*)$ and $B \mapsto F_{\lambda,\mu}(A^*, B)$ have unique minimizers (at every accumulation point), then one has $(A^{(\ell+1)}, B^{(\ell+1)}) - (A^{(\ell)}, B^{(\ell)}) \rightarrow 0$. We conclude that if regularization is present we have the following convergence result for Algorithm 2.

Theorem 4.1. *Assume $\lambda > 0$ and $\mu > 0$. Then the sequence $(A^{(\ell)}, B^{(\ell)})$ generated by Algorithm 2 possesses at least one accumulation point (A^*, B^*) . Further, every accumulation point is a partial optimum of $F_{\lambda,\mu}$ and achieves the same function value $F_{\lambda,\mu}(A^*, B^*)$. It holds $(A^{(\ell+1)}, B^{(\ell+1)}) - (A^{(\ell)}, B^{(\ell)}) \rightarrow 0$.*

Proof. Property (ii) above is fulfilled (for both subproblems) by strict convexity and coercivity. Property (i) also holds, since by construction the sequence of $F_{\lambda,\mu}(A^{(\ell)}, B^{(\ell)})$ is monotonically decreasing, but $\left\| T - \sum_{j=1}^k A_j^{(\ell)} \otimes B_j^{(\ell)} \right\|_2$ is bounded from below by zero. Hence the term $\lambda \sum_{j=1}^k \|A_j^{(\ell)}\|_F^2 + \mu \sum_{j=1}^k \|B_j^{(\ell)}\|_F^2$, and with that all $A_j^{(\ell)}$ and $B_j^{(\ell)}$, must remain bounded. As a result there is at least one accumulation point and [GPK07, Theorem 4.9] is applicable. The monotonicity of $F_{\lambda,\mu}(A^{(\ell)}, B^{(\ell)})$ further implies that all accumulation points achieve the same function value. \square

We point out that in the unregularized case ($\lambda = 0$ and/or $\mu = 0$) we do not only have a potential lack of uniqueness in the subproblems but are also unable to even ensure boundedness of the iterates, although we do not have a counterexample at hand.

Let us also briefly discuss whether partial optima of $F_{\lambda,\mu}$ are critical points. Clearly, if $F_{\lambda,\mu}$ is differentiable at a partial optimum (A^*, B^*) , then $\nabla F_{\lambda,\mu}(A^*, B^*) = 0$. This condition boils down to $\|T - \sum_{j=1}^k A_j \otimes B_j\|_2$ being differentiable in (A^*, B^*) . It is well known that the matrix spectral norm function

$$S \mapsto \|S\|_2 = \max_{\|x\|_2=1} \|Sx\|_2$$

is differentiable at such S for which the maximum on the right-hand side is achieved for a unique x . This follows from a general result on max-functions [Cla75]. Alternatively, let $s_1 \geq s_2 \geq \dots \geq 0$ denote the singular values of S , then the spectral norm $\|S\|_2 = s_1$ is (continuously) differentiable at S if $s_1 > s_2$ [SS90]. This leads to the following statement: if at a partial optimum (A^*, B^*) of $F_{\lambda,\mu}$ the matrix $S = T - \sum_{j=1}^k A_j \otimes B_j$ has a unique largest singular value, then (A^*, B^*) is a critical point of $F_{\lambda,\mu}$ in the sense that $\nabla F_{\lambda,\mu}(A^*, B^*) = 0$. Unfortunately, at the moment we do not have any alternative characterization of this property in terms of structural properties of T .

5. NUMERICAL EXPERIMENTS

In this section, we present results of computational experiments showcasing the potential advantages of the proposed alternating SDP approach. The algorithms have been coded in MATLAB. The SDPs are implemented using the YALMIP toolbox [Löf04] with the SDP solver SeDuMi [Stu99]. An advantage of this toolbox is that the SDPs can be parsed almost directly in the given forms such as (3.4) and (3.5). However, solving those SDPs becomes computationally expensive in larger dimensions and no attempts for a more efficient implementation have been made. This explains the small values for m and n in the following experiments. One should also keep in mind that the nonconvex and nonsmooth optimization task (1.4) is far from trivial even in small dimensions. Our goal here is a proof of concept that the alternating SDP approach allows for better approximations in spectral norm than methods based on Frobenius norm.

5.1. Illustration of Example 2.1. As a first experiment we simulate the construction given in Example 2.1 in which the SVD solution is particularly poor. We take $m = n$ and construct

$$(5.1) \quad T = \sigma_1 A_1 \otimes B_1 + A_2 \otimes B_2$$

with $\sigma_1 = 1.9$ and

$$A_1 = B_1 = \frac{1}{\sqrt{m-1}} \begin{pmatrix} I_{m-1} & 0 \\ 0 & 0 \end{pmatrix}, \quad A_2 = B_2 = \begin{pmatrix} 0_{m-1} & 0 \\ 0 & 1 \end{pmatrix}.$$

Here I_{m-1} and 0_{m-1} are the $(m-1) \times (m-1)$ identity and zero matrix, respectively. The goal is to solve (1.4) with $k = 1$, that is, to minimize $\|T - A \otimes B\|_2$.

As explained in Example 2.1, the SVD method in Algorithm 1 will select a solution $T_{\text{svd}} = \alpha_1 \sigma_1 A_1 \otimes A_2$, but the error will be constant, $\|T - T_{\text{svd}}\|_2 = 1$. Note that for $m = n \geq 3$ this is also the relative error, since $\|T\|_2 = 1$ then. On the other hand, an optimal approximation should yield an error less than $\|T - A_2 \otimes B_2\|_2 = \sigma_1 / (m-1)$.




FIGURE 1. Results of Algorithm 1 (SVD method) and Algorithm 2 (ASDP) with $k = 1$ for the operator T in (5.1) with $m = n = 3, \dots, 10$. The y -axis shows the error $\|T - A \otimes B\|_2$ for the computed approximation $A \otimes B$. Note that $\|T\|_2 = 1$.

In Figure 1 we compare the achieved error $\|T - A \otimes B\|_2$ of the SVD method computed by Algorithm 1 (dashed line) with ASDP solutions obtained with Algorithm 2 for different values $m = 3, \dots, 10$. Three setups for Algorithm 2 were tested: For the red curve (cross markers) and yellow curve (square markers) no regularization has been applied, whereas in the purple curve (triangle markers) we took $\lambda = \mu = 0.1$. The red curve uses the SVD solution as an initial guess (for B). For the yellow and purple curve we used a random initial guess (the same for both). In all cases only $N_{\text{outer}} = 5$ iterations were performed.

As expected, the SVD solution produces a constant error $\|T - A \otimes B\|_2 = 1$. The other curves follow a predicted algebraic decay. As a curiosity, the purple curve, for which regularization was used, coincides exactly with the error bound $\sigma_1/(m-1)$. In fact, due to the diagonal structure of T any $A \otimes B = \beta A_2 \otimes B_2$ with $1 - \beta \leq \sigma_1/(m-1)$ achieves this approximation error, and the algorithm indeed returned such solutions with $1 - \beta = \sigma_1/(m-1)$ because with regularization it also aims at A and B with small Frobenius norms. We did not investigate the effect in detail, but it occurs for a certain range of λ and μ . The red and yellow curves in Figure 1, obtained without regularization, produce even better approximations and are on top of each other. This shows that in this example the choice of the initial guess does not seem to have a big influence and may even suggest that the obtained solutions could be globally optimal.

5.2. Approximation of random operators. In this experiment we apply the algorithms to a matrix $T \in \mathbb{R}^{mn \times mn}$ with random Gaussian entries (normalized to $\|T\|_2 = 1$) and for a sequence of target Kronecker ranks $k = 1, 2, \dots, \min(m^2, n^2)$. Recall that with $k = \min(m^2, n^2)$ an exact decomposition exists, so the error should be zero.

Figure 2 shows computed approximation errors $\|T - \sum_{j=1}^k A_j \otimes B_j\|_2$ for the case $m = 4$, $n = 5$, and $k = 1, 2, \dots, 16$. The ASDP curves are obtained using the same setup as in section 5.1 (again $N_{\text{outer}} = 5$), except that for the purple curve (triangle markers) the regularization parameters were adapted to $\lambda = \mu = 0.1/k$. Decreasing the regularization with k is necessary to account for the effect that the number of penalized terms $\|A_j\|_F$ and $\|B_j\|_F$ grows, while the norms $\|T - \sum_{j=1}^k A_j \otimes B_j\|_2$ remain bounded (and are even




FIGURE 2. Comparison of Algorithm 1 (SVD method) and Algorithm 2 (ASDP) for a random $T \in \mathbb{R}^{mn \times mn}$ with $m = 4$, $n = 5$, scaled to $\|T\|_2 = 1$. The y -axis shows the error $\|T - \sum_{j=1}^k A_j \otimes B_j\|_2$ of computed approximations for $k = 1, 2, \dots, 16$.

intended to decrease). Hence for larger k the regularization terms would dominate in the optimization problem, and indeed in our experiments we observed that using the same λ and μ for all k did not lead to a descending approximation error.

We note that the ASDP method is able to significantly improve the relative approximation error compared to the SVD method, even when randomly initialized. This last point is particularly relevant when the computation of an SVD of the $m^2 \times n^2$ matrix \hat{T} needs to be avoided.

5.3. Inverses of operators with small Kronecker rank. Following the considerations in section 3.4 we present some experimental results for approximating inverses T^{-1} of operators T which themselves have small Kronecker rank. This is based on the modified cost function (3.6) to which we also add regularization terms. The alternating SDP algorithm for this problem is almost identical to Algorithm 2, the (formally) only change being that $M(\tau, A, B)$ in the subproblems (3.4) and (3.5) is replaced with $M_{\text{inv}}(\tau, A, B)$ in (3.8). We note that in the experiments we did not implement the more efficient representation (3.9) of M_{inv} , but treated T as an unstructured operator in (3.8) instead.

Two scenarios are considered. The first is a Lyapunov operator

$$(5.2) \quad T = L \otimes I + I \otimes L,$$

where L is positive definite. This is a special case of more general Sylvester-type operators $T = L_1 \otimes I + I \otimes L_2$ with L_1 and L_2 positive definite. Such types of operators play an important role in matrix equations and numerical analysis, and it is well known that their inverses admit highly accurate approximations in spectral norm by operators of low Kronecker rank [BM02, Gra04]. Specifically, spectral approximation of T^{-1} by sums of exponentials leads to operators of the form $S_k = \sum_{j=1}^k w_j \exp(-t_j L_1) \otimes \exp(-t_j L_2)$, where the parameters w_j and t_j can be chosen such that the error to T^{-1} satisfies

$$\|T^{-1} - S_k\|_2 \leq C e^{-ck}.$$




FIGURE 3. Numerical results for approximating the inverse of the Lyapunov operator (5.2) with $m = n = 10$ and approximation ranks $k = 1, 2, \dots, 8$. Solid lines correspond to implicit methods based on formulation (3.6). For comparison, dashed lines apply Algorithms 1 and 2 directly to T^{-1} .

Here the optimal choice of the parameters, as well as the resulting constants c and C , depend on the spectral bounds of L_1 and L_2 and the constants deteriorate with a growing condition number. If only a positive lower bound on the smallest eigenvalues of L_1 and L_2 is known (and the largest eigenvalues are potentially unbounded), one still has a subexponential but superalgebraic convergence rate $\|T^{-1} - S_k\|_2 \leq C e^{-c\sqrt{k}}$ with (different) constants independent from the upper bound on the spectrum. We refer to [Hac19, Section 9.8.2].

In the numerical experiment we take T of the form (5.2) with $m = n = 10$ and $L = \frac{1}{m-1} \text{tridiag}(-1, 2, -1)$ (tridiagonal matrix), which corresponds to a finite-difference discretization of a (negative) second derivative. The results are shown in Figure 3. Depicted are the computed relative errors $\frac{\|T^{-1} - \sum_{j=1}^k A_j \otimes B_j\|_2}{\|T^{-1}\|_2}$ for several algorithms and approximation ranks $k = 1, 2, \dots, 8$. The red (cross markers) and yellow (square markers) solid lines are the results of an ASDP method for solving (3.6) without regularization ($\lambda = \mu = 0$). The only difference between them is that for the yellow curve a random initialization of the B_j has been taken, whereas for the red curve the B_j were initialized with the solution computed by an alternating least squares (ALS) method for the corresponding problem $\min_{A_j, B_j} \|T - \sum_{j=1}^k A_j \otimes B_j\|_F$ in Frobenius norm. This ALS solution is shown as solid blue curve (circle markers). For comparison, we include the results from the SVD and ASDP methods (Algorithms 1 and 2) when applied directly to the inverse operator T^{-1} , which we explicitly computed for this purpose. In all methods $N_{\text{outer}} = 5$.

Since the plot is in semilogarithmic scale, the results verify the fact that T^{-1} is extremely well approximable in spectral norm by sums of Kronecker products. No big differences between the algorithms can be identified. While the implicit approach (3.6) clearly works, somewhat surprisingly the ALS method based on the Frobenius norm provides slightly better results. We did not investigate whether this could be caused by internal termination criteria in the SDP solver, but it also does not violate the theory since the approximation of T^{-1} in (3.6) is only an implicit one. We note that in this example $\|T^{-1}\|_2 \approx 55.5$.

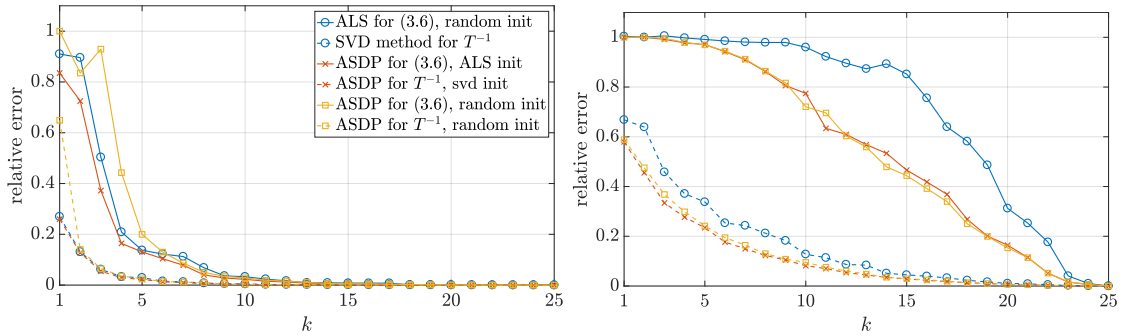


FIGURE 4. Numerical results for approximating the inverse of T in (5.3) with $m = n = 5$ and approximation ranks $k = 1, 2, \dots, 25$. In the left plot C_J, B_J are positive definite, in the right one general. The legend applies to both plots.

Therefore the absolute error $\|T^{-1} - \sum_{j=1}^k A_j \otimes B_j\|_2$ is larger (at most) by this factor than the depicted curves, see the discussion in section 3.4.

In the second scenario, we apply the same algorithms to an operator

$$(5.3) \quad T = C_1 \otimes D_1 + C_2 \otimes B_2 + C_3 \otimes D_3$$

of Kronecker rank 3. In this case we are not aware of a rigorous result on the approximability of T^{-1} by Kronecker products. We take $m = n = 5$ and generate the matrices C_J and D_J randomly, then we scale T to $\|T\|_2 = 1$. We distinguish, however, two cases: in the first the C_J and D_J are symmetric positive definite (achieved by replacing them with $C_J C_J^T$ and $D_J D_J^T$), in the second they are not. Figure 4 shows the computed relative errors $\frac{\|T^{-1} - \sum_{j=1}^k A_j \otimes B_j\|_2}{\|T^{-1}\|_2}$ for $k = 1, 2, \dots, 25$ for both cases. Here we have set $N_{\text{outer}} = 10$ for the alternating optimization methods. The left plot indicates a better error decay in the positive definite case, perhaps even super algebraic. Here $\|T^{-1}\|_2 \approx 83.4$. In the right plot, the implicit methods (solid lines) do not quite capture the achievable error decay (dashed lines). However, compared to the left plot here the ASDP approach performs significantly better than the ALS method based on Frobenius norm. In this example $\|T^{-1}\|_2 \approx 203.4$.

6. CONCLUSION AND OUTLOOK

The problem of approximating a linear operator by sums of Kronecker products in spectral norm is of interest in matrix equations and low-rank calculus. While for the Frobenius norm the approximation problem admits a solution using the singular value decomposition, Example 2.1 shows that the obtained approximation can be far from optimal in spectral norm and suggests that alternative methods should be studied. In this work, we propose to tackle the approximation problem in spectral norm directly utilizing alternating optimization, where the subproblems can be formulated as semidefinite programs and can therefore be solved in polynomial-time to a desired accuracy. The numerical experiments suggest that only a few iterations are necessary to obtain improved approximations using this approach.

This initial work on the subject could be extended in several directions. The presented approach requires the solution of rather high-dimensional semidefinite programs which is computationally expensive and potentially limits the practical applicability. A focus of future work could be to improve the efficiency by fully exploiting structural properties of

the operator T such as sparsity or low (Kronecker) rank in an implementation. In certain situations, it might also be possible to assume a low-rank model for the factor matrices A_j and B_j . In such cases one could appeal to the recent literature for (approximately) solving in an efficient manner large SDPs for which one expects low-rank solutions, see for example [BM03]. Note that such models, i.e., sums of Kronecker products of low-rank matrices, are of potential interest in quantum entanglement, see e.g. [Gur03, NQB14, LN20].

An extension of the results to complex matrices should be straightforward. We also mention that the alternating SDP approach could be generalized to problems on tensor spaces, that is, for finding low Kronecker-rank approximations $\sum_{j=1}^k A_j^{(1)} \otimes \cdots \otimes A_j^{(d)}$ to operators on tensor spaces.

REFERENCES

- [Ber16] D. P. Bertsekas, *Nonlinear programming*, third ed., Athena Scientific, Belmont, MA, 2016.
- [BM02] G. Beylkin and M. J. Mohlenkamp, *Numerical operator calculus in higher dimensions*, Proc. Natl. Acad. Sci. USA **99** (2002), no. 16, 10246–10251.
- [BM03] S. Burer and R. D. C. Monteiro, *A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization*, Math. Program. **95** (2003), no. 2, Ser. B, 329–357.
- [BV04] S. Boyd and L. Vandenberghe, *Convex optimization*, Cambridge University Press, Cambridge, 2004.
- [Cla75] F. H. Clarke, *Generalized gradients and applications*, Trans. Amer. Math. Soc. **205** (1975), 247–262.
- [GPK07] J. Gorski, F. Pfeuffer, and K. Klamroth, *Biconvex sets and optimization with biconvex functions: a survey and extensions*, Math. Methods Oper. Res. **66** (2007), no. 3, 373–407.
- [Gra04] L. Grasedyck, *Existence and computation of low Kronecker-rank approximations for large linear systems of tensor product structure*, Computing **72** (2004), no. 3-4, 247–265.
- [Gur03] L. Gurvits, *Classical deterministic complexity of Edmond’s problem and quantum entanglement*, Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing, ACM, New York, 2003, pp. 10–19.
- [Hac19] W. Hackbusch, *Tensor spaces and numerical tensor calculus*, second ed., Springer, Cham, 2019.
- [HK06a] W. Hackbusch and B. N. Khoromskij, *Low-rank Kronecker-product approximation to multi-dimensional nonlocal operators. I. Separable approximation of multi-variate functions*, Computing **76** (2006), no. 3-4, 177–202.
- [HK06b] ———, *Low-rank Kronecker-product approximation to multi-dimensional nonlocal operators. II. HKT representation of certain operators*, Computing **76** (2006), no. 3-4, 203–225.
- [KMP00] M. Konstantinov, V. Mehrmann, and P. Petkov, *On properties of Sylvester and Lyapunov operators*, Linear Algebra Appl. **312** (2000), no. 1-3, 35–71.
- [LN20] Y. Li and G. Ni, *Separability discrimination and decomposition of m -partite quantum mixed states*, Phys. Rev. A **102** (2020), no. 1, 012402, 8.
- [Löf04] J. Löfberg, *Yalmip: A toolbox for modeling and optimization in MATLAB*, In Proceedings of the CACSD Conference (Taipei, Taiwan), 2004, Available online at <https://yalmip.github.io/>.
- [NQB14] G. Ni, L. Qi, and M. Bai, *Geometric measure of entanglement and U -eigenvalues of tensors*, SIAM J. Matrix Anal. Appl. **35** (2014), no. 1, 73–87.
- [OST09] I. V. Oseledets, D. V. Savostyanov, and E. E. Tyrtyshnikov, *Linear algebra for tensor problems*, Computing **85** (2009), no. 3, 169–188.
- [Sim16] V. Simoncini, *Computational methods for linear matrix equations*, SIAM Rev. **58** (2016), no. 3, 377–441.
- [SS90] G. W. Stewart and J. G. Sun, *Matrix perturbation theory*, Academic Press, Inc., Boston, MA, 1990.
- [Stu99] J. F. Sturm, *Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones*, Optim. Methods Softw. **11/12** (1999), no. 1-4, 625–653.
- [Tse01] P. Tseng, *Convergence of a block coordinate descent method for nondifferentiable minimization*, J. Optim. Theory Appl. **109** (2001), no. 3, 475–494.

- [Tyr04] E. Tyrtyshnikov, *Kronecker-product approximations for some function-related matrices*, Linear Algebra Appl. **379** (2004), 423–437.
- [VL00] C. F. Van Loan, *The ubiquitous Kronecker product*, J. Comput. Appl. Math. **123** (2000), no. 1-2, 85–100.
- [VLP93] C. F. Van Loan and N. Pitsianis, *Approximation with Kronecker products*, Linear algebra for large scale and real-time applications (Leuven, 1992), Kluwer Acad. Publ., Dordrecht, 1993, pp. 293–314.

Mareike Dressler, School of Mathematics and Statistics, University of New South Wales, Sydney, NSW 2052, Australia.

André Uschmajew, Max Planck Institute for Mathematics in the Sciences, 04103 Leipzig, Germany.

Venkat Chandrasekaran, Department of Computing and Mathematical Sciences and Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125, USA.