



Some algorithms for maximum volume and cross approximation of symmetric semidefinite matrices

Stefano Massei¹

Received: 16 November 2020 / Accepted: 5 April 2021 / Published online: 22 April 2021
© The Author(s) 2021

Abstract

Various applications in numerical linear algebra and computer science are related to selecting the $r \times r$ submatrix of maximum volume contained in a given matrix $A \in \mathbb{R}^{n \times n}$. We propose a new greedy algorithm of cost $\mathcal{O}(n)$, for the case A symmetric positive semidefinite (SPSD) and we discuss its extension to related optimization problems such as the maximum ratio of volumes. In the second part of the paper we prove that any SPSP matrix admits a cross approximation built on a principal submatrix whose approximation error is bounded by $(r + 1)$ times the error of the best rank r approximation in the nuclear norm. In the spirit of recent work by Cortinovis and Kressner we derive some deterministic algorithms, which are capable to retrieve a quasi optimal cross approximation with cost $\mathcal{O}(n^3)$.

Keywords Maximum volume · Cross approximation · Symmetric positive definite matrices

1 Introduction

Given $A \in \mathbb{R}^{n \times n}$ and $r \in \mathbb{N}$, this work is mainly concerned with the selection of row and column subsets of indices $I, J \subset \{1, \dots, n\}$ of cardinality r with one of the following features:

- (i) $A(I, J)$ is a *maximum volume* submatrix that is

$$\mathcal{V}(A(I, J)) = \max_{|\hat{I}|=|\hat{J}|=r} \mathcal{V}(A(\hat{I}, \hat{J})), \quad \mathcal{V}(A(I, J)) := |\det(A(I, J))|,$$

Communicated by Lothar Reichel.

✉ Stefano Massei
s.massei@tue.nl ; massei.stef@gmail.com

¹ Centre for Analysis, Scientific Computing and Applications (CASA), TU Eindhoven, Eindhoven, Netherlands

(ii) given another matrix $B \in \mathbb{R}^{n \times n}$, (I, J) is a maximum point of

$$\frac{\mathcal{V}(A(I, J))}{\mathcal{V}(B(I, J))} = \max_{|\hat{I}|=|\hat{J}|=r} \frac{\mathcal{V}(A(\hat{I}, \hat{J}))}{\mathcal{V}(B(\hat{I}, \hat{J}))},$$

(iii) $A_{IJ} := A(:, J)A(I, J)^{-1}A(I, :)$ is a *quasi optimal cross approximation*, i.e., it verifies

$$\|A - A_{IJ}\| \leq p(r) \cdot \min_{\text{rk}(C)=r} \|A - C\|,$$

for a low-degree polynomial $p(\cdot)$ and a matrix norm $\|\cdot\|$.

A connection between problems (i) and (iii) is given by a result of Goreinov and Tyrtyshnikov [16], which says that if $A(I, J)$ has maximum volume then the cross approximation A_{IJ} satisfies the bound

$$\|A - A_{IJ}\|_{\max} \leq (r + 1)\sigma_{r+1}(A), \tag{1}$$

with $\sigma_k(\cdot)$ indicating the k -th singular value and $\|\cdot\|_{\max}$ denoting the maximum magnitude among the entries of the matrix argument. We remark that, in general being a quasi optimal cross approximation does not imply any connection between the volume of $A(I, J)$ and the maximum volume. Indeed, while (i) is an NP hard problem, it has been recently shown that a quasi optimal approximation with respect to the Frobenius norm always exists [33] and can be found in polynomial time [6].

Maximum volume. Problem (i) finds application in a varied range of fields that highlight how the maximum volume concept is multifaceted. For instance, identifying the optimal nodes for polynomial interpolation on a given domain, the so called *Fekete points*, can be recast as selecting the maximum volume submatrix of Vandermonde matrices on suitable discretization meshes [29]. In the optimal experimental design of linear regression models, it is of interest to select the subset of experiments, which is influenced the least by the noise in the measurements. To pursue this goal, the *D-optimality* criterion suggests to look at the covariance matrix of the model and find its principal subblock of maximum volume [22]. Other fields where (i) arises are rank revealing factorizations [17,18], preconditioning [1] and tensor decompositions [25].

Finding a submatrix with either exact or approximate maximum volume are both NP hard problems [5,31]. Despite this downside there has been quite some effort in the development of efficient heuristic algorithms for volume maximization. A central tool for our discussion is one of these methods: the *Adaptive Cross Approximation (ACA)* [2,20]. ACA is typically presented as a low-rank matrix approximation algorithm but it can be interpreted as a greedy method for maximizing the volume. When used for low-rank approximation, ACA is equivalent to a Gaussian elimination process with rook pivoting, and it returns an incomplete LU factorization. In particular, the approximant computed by ACA is of the form in (1) although there is no clear relation between the maximum volume submatrix and the submatrix selected by ACA. On the other hand, the latter can be used as starting guess for procedures that “locally maximize” the volume, e.g., [15,24]. These algorithms guarantee that the volume of the submatrix

that they return can not be increased with a small cardinality change of either its row or column index set. See also [26] for an analysis of these techniques.

In many situations the matrix A is *symmetric positive semidefinite (SPSD)*. For instance, this setting arises in kernel-based interpolation [13], low-rank approximation of covariance matrices [20,23] and discretization of operators involving convolution with a positive semidefinite kernel function [3]. The SPSPD structure comes with a major benefit: the submatrix of maximum volume is always attained for a principal submatrix [7]. Although this does not cure the NP hardness of the task, it reduces significantly the search space by adding the constraint $I = J$.

In Sect. 2.2 we propose a new efficient procedure for the local maximization of the volume over the set of principal submatrices. More specifically, our algorithm returns an $r \times r$ principal submatrix whose volume is maximal over the set of principal submatrices that can be obtained with the replacement of one of the selected indices. Implementation details and complexity analysis are discussed in Sect. 2.2.2. Numerical tests are reported in Sect. 2.4.

Maximum ratio of volumes. To the best of our knowledge, there is no reference to problem (ii) in the literature and there are no direct links with either (i) or (iii) when generic matrices A, B are considered. Nevertheless, we might think at the following situation: suppose that A is SPSPD, B is banded and symmetric positive definite and that we want to compute a cross approximation of $E := T_B^{-\top} A T_B^{-1}$ —where T_B indicates the Cholesky factor of B —without forming E . Since E is SPSPD it would make sense to apply ACA with diagonal pivoting. However, this requires to evaluate the diagonal of E , which is as expensive as forming the whole matrix. Our idea is to replace the diagonal pivoting with the solution of (ii) as heuristic strategy for finding a cross approximation for E .

Indeed, the Binet-Cauchy theorem tells us that a principal minor of E satisfies

$$\begin{aligned} \det(E(J, J)) &= \sum_{|H|=|K|=r} \det(T_B^{-\top}(J, H)) \det(A(H, K)) \det(T_B^{-1}(K, J)) \\ &= \det(T_B^{-\top}(J, J)) \det(A(J, J)) \det(T_B^{-1}(J, J)) \\ &\quad + \sum_{(H,K) \neq (J,J)} \det(T_B^{-\top}(J, H)) \det(A(H, K)) \det(T_B^{-1}(K, J)). \end{aligned}$$

If B is banded and well conditioned, then T_B is banded and the magnitude of the entries of T_B^{-1} decays exponentially with the distance from the main diagonal [9]. Under these assumptions we might have

$$\det(E(J, J)) \approx \det(T_B^{-\top}(J, J)) \det(A(J, J)) \det(T_B^{-1}(J, J)) \approx \frac{\det(A(J, J))}{\det(B(J, J))}. \tag{2}$$

Based on this argument we propose to select J via a greedy algorithm for (ii) and return $E_J := E(:, J)E(J, J)E(J, :)$ as approximation of E . Note that, forming the factors of E_J only requires to solve r linear systems with T_B and to compute r matrix vector products with A .

In Sect. 2.3 we describe how to extend the ACA based techniques for addressing (i) to deal with (ii). We conclude by testing the approximation property of the approach in Sect. 2.4.

Quasi optimal cross approximations. In contrast to the typical robustness of ACA and its simple formulation, very little can be said a priori on the quality of the cross approximation that it returns. Even for structured cases, a priori bounds for the approximation error contain factors that grow exponentially with r [20,21], with the only exception of the doubly diagonally dominant case [7].

Recently, Zamarshkin and Osinsky proved in [33] the existence of quasi optimal cross approximations with respect to the Frobenius norm by means of a probabilistic method. Derandomizing the proof of this result, Cortinovis and Kressner have shown in [6] how to design an algorithm that finds a quasi optimal cross approximation in polynomial time.

In Sect. 3.1 we describe how to modify the technique used in [33] to prove that for an SPSD matrix A there exists a quasi optimal cross approximation with respect to the nuclear norm which is built on a principal submatrix, i.e., $I = J$. This is of particular interest in uncertainty quantification: if A is the covariance matrix of a Gaussian process, then the nuclear norm of the error bounds the Wasserstein distance with respect to another Gaussian process that can be efficiently sampled [23].

In Sects. 3.2–3.3 we propose two algorithms, obtained with the method of conditional expectations, which are able to retrieve quasi optimal cross approximations of SPSD matrices in polynomial time. We conclude by discussing the algorithmic implementation and reporting, in Sect. 3.4, numerical experiments illustrating the performances of the methods.

Notation. In this work we use Matlab-like notation for denoting the submatrices. The identity matrix of dimension n is indicated with Id_n and we use e_j to denote the j -th column of the identity matrix, whose dimension will be clear from the context. The symbols $\|\cdot\|_*$, $\|\cdot\|_F$ indicate the nuclear and Frobenius norm, respectively.

2 Maximizing the volume and the ratio of volumes

Given $r \in \mathbb{N}$, an SPSD matrix $A \in \mathbb{R}^{n \times n}$ and a symmetric positive definite matrix $B \in \mathbb{R}^{n \times n}$, the ultimate goal of this section is to discuss some numerical methods for dealing with the following optimization problems:

$$\max_{\widehat{C} \subset \{1, \dots, n\}, |\widehat{C}|=r} \mathcal{V}(A(\widehat{C}, \widehat{C})), \quad (3)$$

$$\max_{\widehat{C} \subset \{1, \dots, n\}, |\widehat{C}|=r} \frac{\mathcal{V}(A(\widehat{C}, \widehat{C}))}{\mathcal{V}(B(\widehat{C}, \widehat{C}))}. \quad (4)$$

When $B = \text{Id}_n$, (4) reduces to (3); moreover (3) corresponds to the maximum volume problem because for an SPSD matrix, the maximum is attained at a principal submatrix [7]. We start by recalling a well known greedy strategy to deal with (3), the so-called

Adaptive Cross Approximation (ACA) [20]. Then, we will see how to generalize ACA for addressing (4).

2.1 Adaptive cross approximation

The selection of high volume submatrices of A is intimately related with the low-rank approximation of A . The link is the *cross approximation* [2,32], which associates with a given subset of indices $J = \{j_1, \dots, j_r\}$, or equivalently with an invertible submatrix $A(J, J)$, the rank r matrix approximation¹

$$A_J := A(:, J)A(J, J)^{-1}A(J, :). \tag{5}$$

Cross approximations are attractive because to build A_J only requires a partial evaluation of the entries of A , which is crucial when considering large scale matrices. Moreover, since the residual matrix $R_J := A - A_J$ is SPSD, the approximation error can be cheaply estimated as

$$\text{trace}(R_J) = \|R_J\|_* \geq \|R_J\|_F \geq \|R_J\|_2 \geq \frac{\text{trace}(R_J)}{n}. \tag{6}$$

When J is a maximum point of (3), A_J yields a quasi optimal approximation error with respect to the maximum norm [16]. However, solving (3) is NP hard which paves the way to the use of heuristic approaches such as ACA.

The ACA method selects J with a process analogous to Gaussian elimination with complete pivoting. The algorithm begins by choosing $j_1 = \arg \max_j A_{jj}$ and computes $R_{J_1} = A - A(:, j_1)A_{j_1 j_1}^{-1}A(j_1, :)$. Then, the procedure is iterated on the residual matrices R_{J_i} , $i = 1, \dots, r - 1$ in order to retrieve r indices. The elements $(R_{J_i})_{j_{i+1} j_{i+1}}$ correspond to the first r pivots selected by the Gaussian elimination with complete pivoting on the matrix A , and we have the identity

$$\det(A(J, J)) = \prod_{i=0}^{r-1} (R_{J_i})_{j_{i+1} j_{i+1}}, \tag{7}$$

where $R_{J_0} := A$. In particular, (7) explains that each step of ACA augments the set of selected indices by following a greedy strategy with respect to the volume of the selected submatrix. The whole procedure is reported in Algorithm 1. Note that, if one stores the vectors u_1, \dots, u_r , then only the diagonal and the columns j_1, \dots, j_r , of A , need to be evaluated. The efficient implementation of the algorithm replaces the computation of the residual matrix at line 8 with the update of the diagonal of R_J . Computing $R_J(:, j_k) = A(:, j_k) - U_{k-1}U_{k-1}(j_k, :)^T$, $U_{k-1} := [u_1, \dots, u_{k-1}]$, only requires a partial access to A as well. In case the matrix A is not formed explicitly and its entries are evaluated with a given handle function, Algorithm 1 requires $\mathcal{O}(rn)$

¹ Cross approximation is generally associated with two subsets of indices, one for the rows and one for the columns of the submatrix. In view of the positive definiteness of A we restrict to principal submatrices.

storage and its computational cost is $\mathcal{O}((r + c_A)rn)$ where c_A denotes the cost of evaluating one entry of A .

Algorithm 1 ACA for (3)

```

1: procedure ACA( $A, r$ )
2:   Set  $R_J := A, J := \emptyset$ 
3:   for  $k := 1, 2, \dots, r$  do
4:      $j_k := \arg \max_j (R_J)_{jj}$ 
5:      $J \leftarrow J \cup \{j_k\}$ 
6:     if  $k < r$  then
7:        $u_k := R_J(:, j_k) / \sqrt{(R_J)_{j_k j_k}}$ 
8:        $R_J \leftarrow R_J - u_k u_k^\top$ 
9:     end if
10:  end for
11:  return  $J$ 
12: end procedure

```

Algorithm 2 ACA for (4)

```

1: procedure ACA_RATIO( $A, B, r$ )
2:   Set  $R_J^{(A)} := A, R_J^{(B)} := B, J := \emptyset$ 
3:   for  $k := 1, 2, \dots, r$  do
4:      $j_k := \arg \max_j (R_J^{(A)})_{jj} / (R_J^{(B)})_{jj}$ 
5:      $J \leftarrow J \cup \{j_k\}$ 
6:     if  $k < r$  then
7:        $u_k^{(A)} := R_J^{(A)}(:, j_k) / \sqrt{(R_J^{(A)})_{j_k j_k}}$ 
8:        $R_J^{(A)} \leftarrow R_J^{(A)} - u_k^{(A)} (u_k^{(A)})^\top$ 
9:        $u_k^{(B)} := R_J^{(B)}(:, j_k) / \sqrt{(R_J^{(B)})_{j_k j_k}}$ 
10:       $R_J^{(B)} \leftarrow R_J^{(B)} - u_k^{(B)} (u_k^{(B)})^\top$ 
11:    end if
12:  end for
13:  return  $J$ 
14: end procedure

```

2.2 Local maximization

Let us suppose that a certain index set $J = \{j_1, \dots, j_r\}$ is given. Inspired by [15], we would like to know whether the volume of $A(J, J)$ is *locally optimal*, in the sense that it cannot be increased with the replacement of just one of the indices in J . Practically, this requires to check that:

$$\frac{\det(A(\widehat{J}, \widehat{J}))}{\det(A(J, J))} \leq 1, \quad \forall \widehat{J} : |J \cap \widehat{J}| = r - 1, \quad |\widehat{J}| = r. \tag{8}$$

For the low-rank approximation problem in the maximum norm, a locally optimal determinant is sufficient to reach a quasi optimal accuracy.

Lemma 1 *Let $A \in \mathbb{R}^{n \times n}$ be an SPSD matrix and let J be an index set such that condition (8) is verified. Then*

$$\|A - A_J\|_{\max} \leq (r + 1)\sigma_{r+1}(A).$$

Proof When $n = r + 1$ the submatrix $A(J, J)$ has maximum volume and we get the claim simply applying the result of Goreinov and Tyrtyshnikov (equation (1)). For $n > r + 1$, we remark that each diagonal entry of the residual matrix $(R_J)_{hh}$ is equal to the Schur complement of $A(J, J)$ in $A(\tilde{J}, \tilde{J})$, for $\tilde{J} = J \cup \{h\}$. In view of (8), $A(J, J)$ is the maximum volume $r \times r$ submatrix of $A(\tilde{J}, \tilde{J})$ that implies

$$(R_J)_{hh} \leq (r + 1)\sigma_{r+1}(A(\tilde{J}, \tilde{J})) \leq (r + 1)\sigma_{r+1}(A).$$

Since R is SPSD, $(r + 1)\sigma_{r+1}(A)$ also bounds its max norm. □

In the following sections we describe an efficient procedure to iteratively increase $\mathcal{V}(A(J, J))$ based on the evaluation of the $r(n - r)$ ratios in (8). An algorithm for the analogous, yet simpler, task when the index replacement affects only the row or the column index set has been proposed in [15].

2.2.1 Updating the determinant

Let us remark that each $A(\hat{J}, \hat{J})$ in (8) is a rank-2 modification of the matrix $A(J, J)$. More precisely, if the index set \hat{J} is obtained by replacing $j_i \in J$ with $h \in \{1, \dots, n\} \setminus J$, then

$$A(\hat{J}, \hat{J}) = A(J, J) + U W U^T$$

where

$$U = [e_i \mid A(J, h) - A(J, j_i)], \quad W = \begin{bmatrix} A_{hh} + A_{j_i j_i} - 2A_{hj_i} & 1 \\ 1 & 0 \end{bmatrix},$$

and e_i indicates the i -th vector of the canonical basis. Applying the matrix determinant lemma yields

$$\frac{\det(A(\hat{J}, \hat{J}))}{\det(A(J, J))} = \det(W^{-1}) \det(W^{-1} + U^T A(J, J)^{-1} U),$$

with

$$W^{-1} = \begin{bmatrix} 0 & 1 \\ 1 & 2A_{hj_i} - A_{hh} - A_{j_i j_i} \end{bmatrix}, \quad \det(W^{-1}) = -1.$$

By denoting with $D := A(J, J)^{-1}$, $B := A(:, J)D$ and with $C := BA(J, :)$, we have that

$$U^T A(J, J)^{-1} U = \begin{bmatrix} D_{ii} & B_{hi} - 1 \\ B_{hi} - 1 & [B(h, :) - B(j_i, :)][A(J, h) - A(J, j_i)] \end{bmatrix}$$

where we have used the identities

$$\begin{aligned} [A(h, J) - A(j_i, J)]A(J, J)^{-1} &= B(h, :) - B(j_i, :), \\ [B(h, :) - B(j_i, :)]e_i &= B_{hi} - 1. \end{aligned}$$

Putting all pieces together we get

$$W^{-1} + U^T A(J, J)^{-1} U = \begin{bmatrix} D_{ii} & B_{hi} \\ B_{hi} & C_{hh} - A_{hh} \end{bmatrix}.$$

Then, we might think at the following greedy scheme for increasing the volume of a starting submatrix $A(J, J)$:

1. Compute the Cholesky decomposition $R^T R = A(J, J)$, $\mathcal{O}(r^3)$,
2. Retrieve the quantities D_{ii} by solving $R^T R x = e_i, i = 1, \dots, r$, $\mathcal{O}(r^3)$,
3. Compute $B = A(:, J)(R^T R)^{-1}$, $\mathcal{O}((r + c_A)rn)$,
4. Compute $C_{hh} \forall h \in \{1, \dots, n\} \setminus J$, $\mathcal{O}(r(n - r))$,
5. Compute $\mathcal{V}_{hi} := \left| \det \left(\begin{bmatrix} D_{ii} & B_{hi} \\ B_{hi} & C_{hh} - A_{hh} \end{bmatrix} \right) \right| \forall j_i \in J, \forall h \in \{1, \dots, n\} \setminus J$, $\mathcal{O}(r(n - r))$,
6. Identify $\mathcal{V}_{\hat{h}\hat{i}} = \max_{h,i} \mathcal{V}_{hi}$. If $\mathcal{V}_{\hat{h}\hat{i}} > 1 + \text{tol}$ —for a prescribed tolerance tol —then update J by replacing $j_{\hat{i}}$ with \hat{h} and repeat the procedure. Otherwise stop the iteration.

We will discuss possible improvements to this algorithm in the next section.

2.2.2 Updating the quantities B, C and D

The previously sketched procedure requires, whenever the index set J is updated, to recompute the quantities B, C and D . Here, we explain how to leverage the old information to decrease the iteration cost. In the following, we assume that the new index J_{new} is obtained by replacing $j_i \in J_{\text{old}}$ with the index $h \in \{1, \dots, n\} \setminus J_{\text{old}}$.

The new matrix D is the inverse of a rank-2 modification of the old D , therefore it can be obtained with the Woodbury identity:

$$D_{\text{new}} \leftarrow D_{\text{old}} - \underbrace{\begin{bmatrix} e_i^T A(J_{\text{old}}, J_{\text{old}})^{-1} \\ B(h, :) - B(j_i, :) \end{bmatrix}^T \begin{bmatrix} D_{hh} & B_{hi} \\ B_{hi} & C_{hh} - A_{hh} \end{bmatrix} \begin{bmatrix} e_i^T A(J_{\text{old}}, J_{\text{old}})^{-1} \\ B(h, :) - B(j_i, :) \end{bmatrix}}_{\Delta D}. \tag{9}$$

The decomposition $R_{\text{new}}^T R_{\text{new}} = A(J_{\text{new}}, J_{\text{new}})$, can be computed with cost $\mathcal{O}(r^2)$ by rewriting $U W U^T = \tilde{u}_1 \tilde{u}_1^T - \tilde{u}_2 \tilde{u}_2^T$, i.e., as the difference of two rank-1 SPSD

matrices, and performing a rank-1 update and a rank-1 downdate of the old Cholesky factor [30, Chapter 4, Section 3]. For instance, these routines are implemented in the Matlab command `cholupdate`.

The new matrix B is also a low-rank correction of the old B , given by

$$B_{\text{new}} \leftarrow B_{\text{old}} + \underbrace{[A(:, h) - A(:, j_i)]e_i^\top (R_{\text{new}}^\top R_{\text{new}})^{-1} + A(:, J_{\text{old}})\Delta D}_{\Delta B}. \quad (10)$$

Performing the updates of D and B with (9) and (10), respectively, brings down the iteration cost to $\mathcal{O}(r^2 + (r + c_A)n)$, apart from the first iteration which remains $\mathcal{O}(r^3 + (r + c_A)rn)$. The procedure is reported in Algorithm 3.

Since the use of the Woodbury identity is sometimes prone to numerical instabilities, e.g. when the selected submatrix is nearly singular, we may switch off the updating mechanism by setting the boolean variable `do_update` to false at line 3.

Algorithm 3 Local maximization of the volume

```

1: procedure LOCAL_MAXVOL( $A, J, \text{tol}$ )
2:   Set vol_ratio :=  $+\infty, k := 1$ 
3:   Set do_update := true/false ▷ Enable or disable the update of  $B$  and  $D$ 
4:   while vol_ratio >  $1 + \text{tol}$  do
5:     if  $k == 1$  or do_update == false then
6:        $R = \text{chol}(A(J, J))$ 
7:        $B = A(:, J)(R^\top R)^{-1}$ 
8:       Compute  $D_{ii}$  by solving  $R^\top R x = e_i, \forall i = 1, \dots, r$ 
9:     else
10:       $R \leftarrow \text{cholupdate}(R, U, W)$ 
11:      Update  $D$  via (9)
12:      Update  $B$  via (10)
13:    end if
14:     $C_{hh} \leftarrow B(h, :)A(J, h), \forall h \in \{1, \dots, n\} \setminus J$ 
15:     $\mathcal{V}_{hi} = \left| \det \begin{pmatrix} D_{ii} & B_{hi} \\ B_{hi} & C_{hh} - A_{hh} \end{pmatrix} \right| \forall j_i \in J, \forall h \in \{1, \dots, n\} \setminus J$ 
16:     $[\hat{h}, \hat{i}] \leftarrow \arg \max_{h,i} \mathcal{V}_{h,i}$ 
17:    vol_ratio  $\leftarrow \mathcal{V}_{\hat{h}\hat{i}}$ 
18:     $J \leftarrow J \cup \{\hat{h}\} \setminus \{j_i\}, k \leftarrow k + 1$ 
19:  end while
20:  return  $J$ 
21: end procedure

```

Finally, we remark that updating the diagonal elements of C with the relation

$$C_{\text{new}} \leftarrow C_{\text{old}} + B_{\text{old}}e_i[A(h, :) - A(j_i, :)] + \Delta B A(J_{\text{new}}, :),$$

would reduce the cost of line 14 in Algorithm 3 of a factor r . However, since this does not change the complexity of the iteration and requires to store additional intermediate quantities, it is not incorporated in our implementation.

2.2.3 A new algorithm for the maximum volume of SPSD matrices

Quite naturally, we propose to apply Algorithm 3 to the index set returned by Algorithm 1 as heuristic method for solving (3). The resulting procedure is ensured to return a locally optimal principal submatrix of A — in the sense of Sect. 2.2 — whose volume is larger or equal than the one returned by ACA. For completeness, we report the method in Algorithm 4.

By denoting with it the number of iterations performed by LOCAL_MAXVOL, we have that the computational cost of Algorithm 4 is $\mathcal{O}((r + c_A)(r + it)n)$.

We also show that it is possible to provide an upper bound for it that does not depend on n . Finding the maximum volume submatrix of an SPSD matrix is in one to one correspondence with selecting the columns of maximum volume in its Cholesky factor T_A such that $A = T_A^T T_A$ [7, Section 2.1.1]. In particular, the greedy algorithm for column selection, i.e. the partial QR with column pivoting, executed on T_A returns the same index set J_{aca} identified by $ACA(A, r)$, as they are both based on greedy unit augmentations of the index set. Moreover, the volume of $T_A(:, J_{aca})$ is at least $(r!)^{-1}$ times the maximum volume achievable with a subset of r columns [5, Theorem 11] and is equal to the square root of $\det(A(J_{aca}, J_{aca}))$. Then, we have

$$\det(A(J_{aca}, J_{aca})) \geq \frac{\det(A(J_{best}, J_{best}))}{(r!)^2},$$

where $A(J_{best}, J_{best})$ denotes the maximum volume $r \times r$ submatrix. This means that when calling LOCAL_MAXVOL in Algorithm 3, the volume cannot be increased more than a factor $(r!)^2$. Since each iteration of LOCAL_MAXVOL increases the volume of at least a factor $1 + tol$, this yields the following bound on its number of iterations:

$$(1 + tol)^{it} \leq (r!)^2 \implies it \leq 2 \frac{\log(r!)}{\log(1 + tol)}.$$

Finally, by means of the Stirling's approximation, we get $it \leq 2 \frac{(r+1) \log(r) - r + 1}{\log(1+tol)} = \mathcal{O}(r \log(r))$.

Algorithm 4

```

1: procedure MAXVOL( $A, r, tol$ )
2:    $J = ACA(A, r)$ 
3:    $J \leftarrow LOCAL\_MAXVOL(A, J, tol)$ 
4: end procedure

```

Algorithm 5

```

1: procedure MAXVOL_RATIO( $A, B, r, tol$ )
2:    $J = ACA\_RATIO(A, B, r)$ 
3:    $J \leftarrow LOCAL\_MAXVOL\_RATIO(A, B, J, tol)$ 
4: end procedure

```

2.3 Algorithms for maximizing the ratio of volumes

Let $J = \{j_1, \dots, j_r\}$ be the index set at the current iteration of either Algorithms 1 or 3. The two algorithms compute the gain factor $\det(A(\widehat{J}, \widehat{J})) / \det(A(J, J))$ for all the modifications $\widehat{J} \in \mathcal{J}_{aca}$ and $\widehat{J} \in \mathcal{J}_{lmvol}$, respectively, where

$$\begin{aligned} \mathcal{J}_{aca} &= \{\widehat{J} \subset \{1, \dots, n\} : J \subset \widehat{J}, |\widehat{J}| = r + 1\}, \\ \mathcal{J}_{lmvol} &= \{\widehat{J} \subset \{1, \dots, n\} : |J \cap \widehat{J}| = r - 1, |\widehat{J}| = r\}. \end{aligned}$$

Therefore, Algorithm 1 and 3 can be adapted for the ratio of volume problem (4) with the following idea: run in parallel the procedure for the matrices A, B and then identify the maximum ratio of gain factors

$$\frac{\det(A(\widehat{J}, \widehat{J})) \det(B(J, J))}{\det(A(J, J)) \det(B(\widehat{J}, \widehat{J}))} \quad \forall \widehat{J} \in \mathcal{J}_{aca} \text{ or } \forall \widehat{J} \in \mathcal{J}_{lmvol}.$$

For instance, the extension of ACA to (4) looks for $\arg \max_j (R_J^{(A)})_{jj} / (R_J^{(B)})_{jj}$ when choosing the next pivot element; see Algorithm 2. Analogously, the version of Algorithm 3 which deals with the ratio of volumes, identifies the pair of indices (h, i) which maximizes $\mathcal{V}_{hi}^{(A)} / \mathcal{V}_{hi}^{(B)}$. We refer to the latter with LOCAL_MAXVOL_RATIO and — due to its length — we refrain to write its pseudocode. Finally, the extension of Algorithm 4 to (4) is reported in Algorithm 5.

By denoting with it the number of iterations performed by LOCAL_MAXVOL_RATIO, we have that the computational cost of Algorithm 5 is $\mathcal{O}((r + c_A + c_B)(r + it)n)$, where c_B indicates the cost of evaluating one entry of B .

2.4 Numerical results

Algorithms 1–5 have been implemented in Matlab version R2020a and all the numerical tests in this work have been executed on a Laptop with the dual-core Intel Core i7-7500U 2.70 GHz CPU, 256 KB of level 2 cache, and 16 GB of RAM. The parameter tol used in Algorithm 4 and 5 has been set to $5 \cdot 10^{-2}$ for all the experiments reported in this manuscript. In the numerical tests involving the test matrix A_3 and Algorithm 3 the updating mechanism has been switched off by setting do_update to false. Everywhere else, do_update has been set to true.

The code is freely available at <https://github.com/numpi/max-vol>.

Test matrices Let us define five SPSD matrices $A_1, A_2, A_3, A_4, A_5 \in \mathbb{R}^{n \times n}$ which are involved in the numerical experiments that we are going to present:

- $(A_1)_{ij} := \exp(-0.3 |i - j|/n)$,
- $(A_2)_{ij} := \min\{i, j\}$,
- $(A_3)_{ij} := \frac{1}{i+j-1}$ (Hilbert matrix),
- $A_4 := \text{trid}(1, 1, 1) \otimes \text{ld}_6 + \text{ld}_{\frac{n}{6}} \otimes \text{trid}(-0.34, 1.7, -0.34)$,
- $A_5 := Q \text{diag}(d) Q^\top$, $d_i := \rho^{i-1}$, $\rho \in (0, 1)$, and Q is the eigenvector matrix of $\text{trid}(-1, 2, -1)$,

with \otimes indicating the Kronecker product. The aforementioned test matrices are representative of various singular values distributions. A_1, A_2 have a subexponential decay, A_3, A_5 have an exponential decay and A_4 , taken from [19], is banded and well conditioned. We also indicate with $T_4^\top T_4 = A_4$ the Cholesky factorization of A_4 . When running the numerical algorithms, the matrices A_1, A_2, A_3 and A_4 are provided as function handles. Instead, the matrix A_5 is formed explicitly.

Test 1. As first experiment we run Algorithm 1 and 4 on A_1, A_2, A_3 , by setting $n = 1020$ and varying the size r of the sought submatrix. For the matrices A_1, A_2 we let r to range in $\{1, \dots, 100\}$. When experimenting on A_3 we consider $r \in \{1, \dots, 20\}$ because of the small numerical rank of the Hilbert matrix. We measure the timings required by the two methods and the gain factor $|\det(A(J_{\max\text{vol}}, J_{\max\text{vol}}))/\det(A(J_{\text{aca}}, J_{\text{aca}}))|$ which Algorithm 4 provides with respect to Algorithm 1. From the results reported in Fig. 1, we see that the costs of both algorithms scale quadratically with respect to the parameter r . For small values of r MAXVOL struggles to increase the volume of the submatrix returned by ACA. This happens more often and more consistently for larger values of r . We mention that disabling the updates based on the Woodbury identity generally increases of about 20% the timings of Algorithm 4 for this test.

Test 2. The second numerical test considers maximizing the ratio of volumes (4). We keep $n = 1020$ and we run Algorithm 2 and 5 using A_1, A_2, A_3 , as numerator and A_4 as denominator. The time consumption as the size r of the submatrix increases is reported Fig. 2. Also in this case, quadratic complexity with respect to r is observed for the computational cost. The gain factor $|\det(A(J_{\max\text{vol_ratio}}, J_{\max\text{vol_ratio}}))/\det(A(J_{\text{aca_ratio}}, J_{\text{aca_ratio}}))|$ is shown as well in the bottom right part of Fig. 2.

Test 3. Let us test the computational cost of ACA, MAXVOL, ACA_RATIO and MAXVOL_RATIO as the size of the target matrices increases. We fix $r = 40$ and we let $n = 1020 \cdot 2^t, t = 0, \dots, 10$. Then, we run ACA, MAXVOL on A_1 and MAXVOL, ACA_RATIO on the pair (A_1, A_4) . The timings reported in Fig. 3 confirm that the computational time scales linearly with respect to n .

Test 4. Finally, we test the quality of the cross approximations returned by ACA_RATIO and MAXVOL_RATIO. More specifically, we compute the approximation error $\|E_i - (E_i)_J\|_2, i = 1, 2, 3, 5$, with $E_i := (T_4^\top)^{-1} A_i T_4^{-1}, n = 1020$ and J chosen as either $J_{\text{aca_ratio}}$ or $J_{\max\text{vol_ratio}}$. In Fig. 4 we compare the error curves, as r increases, of the cross approximations with the ones associated with the truncated SVD, which represents the best attainable scenario. We see that the decay rate of the error of ACA_RATIO is pretty similar to the one of the truncated SVD. MAXVOL_RATIO performs also well on the matrices which have a fast decay of the singular values, i.e., A_3, A_5 . However, its convergence deteriorates for the matrices A_1 and A_2 and the associated error is worse than the one of ACA_RATIO. It turns out that in these cases the approximation given in (2) is less accurate and the submatrix of $(T_4^\top)^{-1} A_i T_4^{-1}$ corresponding to $J_{\text{aca_ratio}}$ has a larger volume than the one corresponding to $J_{\max\text{vol_ratio}}$.




Fig. 1 Timings of Algorithm 1 and 4 on the test matrices A_1 (top-left), A_2 (top-right), A_3 (bottom-left) and measured gain factors (bottom-right)

3 Quasi optimal cross approximation in the nuclear norm

Adaptive cross approximation has a much lower cost than computing the truncated SVD for the low-rank matrix approximation, although the latter provides an optimal solution, in any unitarily invariant norm. Empirically, ACA typically returns an approximant that is close, in terms of the associated approximation error, to the truncated SVD. However, it appears difficult to ensure this property theoretically, e.g., see the quite pessimistic bounds in [7,20,21]. On the other hand, there are some recent results about cross approximations with quasi optimal approximation error.

Zamarashkin and Osinsky proved in [33, Theorem 1] that, given $A \in \mathbb{C}^{m \times n}$ of rank k , $\forall r = 1, \dots, k$ there exist $I = \{i_1, \dots, i_r\} \subset \{1, \dots, m\}$ and $J = \{j_1, \dots, j_r\} \subset \{1, \dots, n\}$, such that $A(I, J)$ is invertible and

$$\|A - A_{IJ}\|_F \leq (r + 1) \sqrt{\sum_{s \geq r+1} \sigma_s^2}, \quad A_{IJ} := A(:, J)A(I, J)^{-1}A(I, :). \quad (11)$$

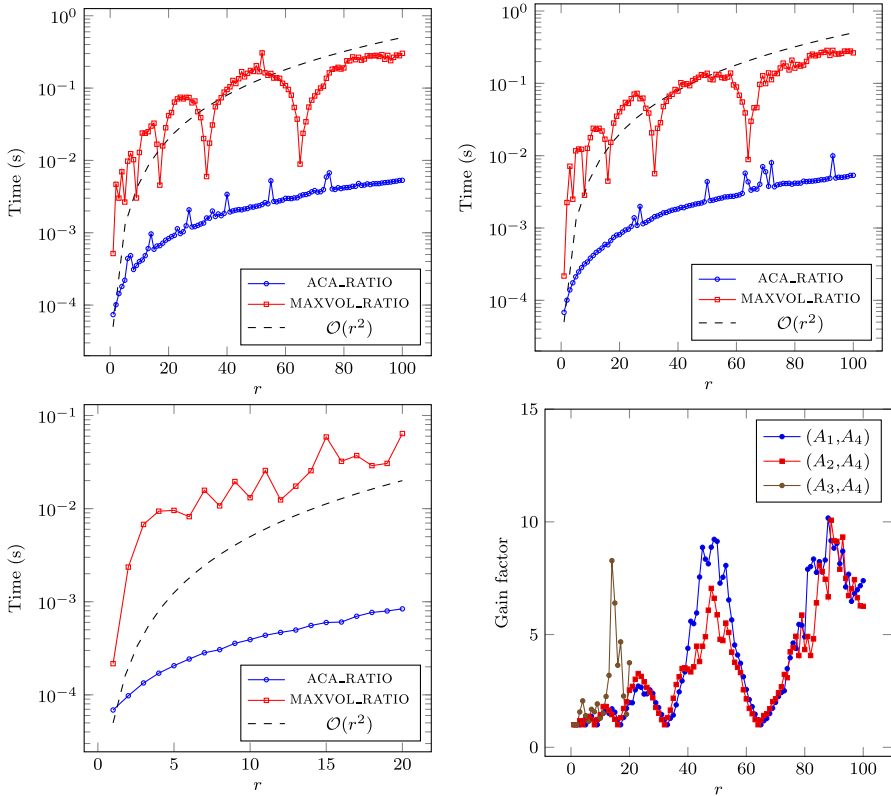


Fig. 2 Timings of Algorithm 2 and 5 on the test matrices (A_1, A_4) (top-left), (A_2, A_4) (top-right), (A_3, A_4) (bottom-left) and measured gain factors (bottom-right)




Fig. 3 Computational times of the algorithms as n increases for $r = 40$. On the left ACA and MAXVOL have been run on the matrix A_1 . On the right ACA_RATIO and MAXVOL_RATIO have been run on the pair of matrices (A_1, A_4)

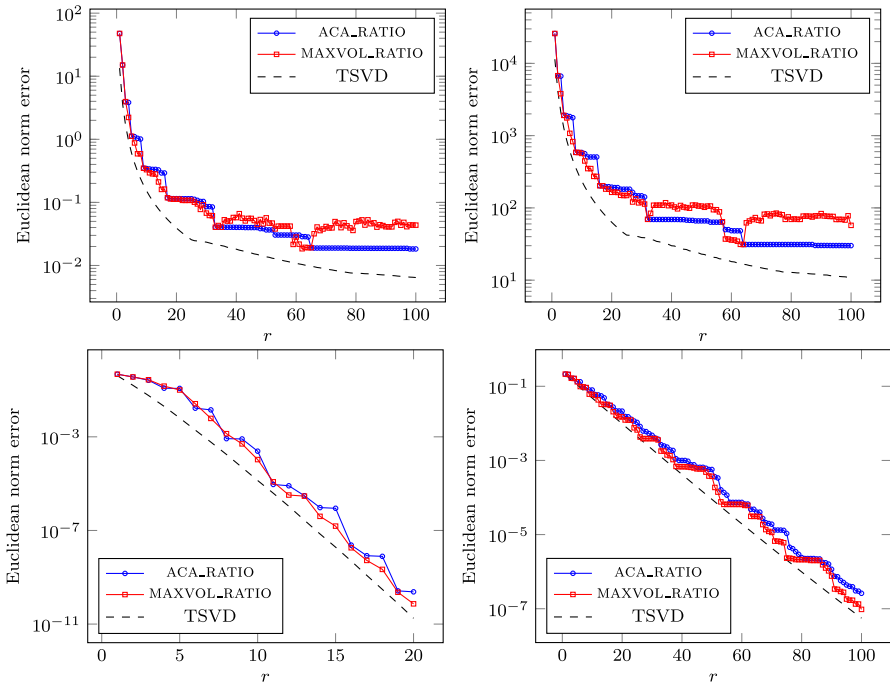


Fig. 4 Approximation of $(T_4^T)^{-1}A_iT_4^{-1}$ for $i = 1$ (top-left), $i = 2$ (top-right), $i = 3$ (bottom-left), and $i = 5$ (bottom-right), by means of the cross approximations associated with the outcome of Algorithm 2 and 5. All plots report the lower bound given by the error of the truncated SVD. The size of the matrices is $n = 1020$

The authors of [33] uses a probabilistic argument: they define the probability measure

$$\mathbb{P}(A(I, J)) = \frac{\mathcal{V}(A(I, J))^2}{\sum_{|\hat{I}|=|\hat{J}|=r} \mathcal{V}(A(\hat{I}, \hat{J}))^2}$$

on the set of $r \times r$ submatrices of A . Then, they show that $\mathbb{E}[\|A - A_{IJ}\|_F] \leq (r + 1)\sqrt{\sum_{s \geq r+1} \sigma_s^2}$, which implies that there exists at least one choice of I, J that verifies (11).

Cortinovis and Kressner proposed in [6] a polynomial time algorithm to find I and J such that A_{IJ} is quasi optimal with respect to the Frobenius norm. Their approach, inspired by [12], is based on the derandomization of the result by Zamarashkin and Osinsky with the method of conditional expectations. More precisely, let $t \leq r$ and assuming to have already selected the first $t - 1$ indices $\{i_1, \dots, i_{t-1}\}, \{j_1, \dots, j_{t-1}\}$ of I and J , the pair (i_t, j_t) is chosen as the one which minimizes

$$\mathbb{E}[\|A - A_{IJ}\|_F \mid i_1, \dots, i_t, j_1, \dots, j_t]. \tag{12}$$

Incrementally selecting all the indices with this criteria ensures that (I, J) identifies a cross approximation which verifies (11). Interestingly, (12) can be shown to be $(r -$

$t + 1$) times the ratio of two consecutive coefficients in the characteristic polynomial of the symmetrized residual matrix $R_{I_t, J_t} := (A - A_{I_t, J_t})(A - A_{I_t, J_t})^*$, with $I_t := \{i_1, \dots, i_t\}$ and $J_t := \{j_1, \dots, j_t\}$. The algorithm in [6] computes the coefficients of the characteristic polynomial of R_{I_t, J_t} for all possible choices of i_t and j_t by updating the characteristic polynomial of $R_{I_{t-1}, J_{t-1}}$; then, it chooses the pair of indices which minimizes the aforementioned ratio.

In the next section, we analyze what can be achieved with cross approximations built on principal submatrices, when A is SPSD.

3.1 Existence result

In view of [7, Theorem 1] it is tempting to replace a symmetric choice of indices $I = J$ in (11) when A is SPSD. However, such error bound it is not true in general and it is not possible to get rid of the dependency on n in the multiplicative constant. For instance, consider $A = E + \epsilon \cdot \text{Id}$ for a small $\epsilon > 0$ and with E denoting the matrix of all ones; then, for the rank 1 approximation of A , the error of the truncated SVD is $\epsilon\sqrt{n - 1}$ while the one associated with any symmetric cross approximation is approximately $\epsilon(n - 1)$. The following result shows that a quasi optimal error in the nuclear norm can be obtained by restricting the search space to principal submatrices. In view of the previous remark, this yields a sharp quasi optimal error in the Frobenius norm, with a constant increased by a factor $\sqrt{n - r}$.

Theorem 1 *Let $A \in \mathbb{R}^{n \times n}$ be SPSD of rank k and $r \in \{1, \dots, k\}$. Then, there exists a subset of indices $J^* \subset \{1, \dots, n\}$, $|J^*| = r$ such that $A(J^*, J^*)$ is invertible and*

$$\|A - A_{J^*}\|_* \leq (r + 1) \cdot \sum_{s \geq r+1} \sigma_s(A), \quad \text{and}$$

$$\|A - A_{J^*}\|_F \leq \sqrt{n - r} \cdot (r + 1) \cdot \sqrt{\sum_{s \geq r+1} \sigma_s(A)^2}. \tag{13}$$

Before going into the proof of Theorem 1, let us state and prove some properties regarding the volume of principal submatrices.

Lemma 2 *Let $A \in \mathbb{R}^{n \times n}$ be SPSD and $J := \{j_1, \dots, j_r\} \subset \{1, \dots, n\}$ such that $A(J, J)$ is invertible. Then:*

- (i) $\|A - A_J\|_* = \sum_{|\hat{J}|=r+1, J \subset \hat{J}} \frac{\mathcal{V}(A(\hat{J}, \hat{J}))}{\mathcal{V}(A(J, J))}$,
- (ii) $\sum_{|J|=r} \mathcal{V}(A(J, J)) = \sum_{1 \leq j_1 < \dots < j_r \leq n} \sigma_{j_1}(A) \cdots \sigma_{j_r}(A)$,
- (iii) for $t \in \{1, \dots, r\}$ and $J_1 := \{j_1, \dots, j_t\} \subset J$

$$\sum_{j_{t+1}, \dots, j_r} \mathcal{V}(A(J, J)) = \mathcal{V}(A(J_1, J_1)) \cdot (r - t)! \cdot c_{n-r+t}(A - A_{J_1}),$$

where $(-1)^{n-r+t} c_{n-r+t}(A - A_{J_1})$ indicates the coefficient which multiplies z^{n-r+t} in the characteristic polynomial of $A - A_{J_1}$.

Proof

(i) Let us remark that in the particular case $J = \{1, \dots, n - 1\}$ we have

$$A = \begin{bmatrix} A(J, J) & b \\ b^\top & d \end{bmatrix}, \quad A - A_J = \begin{bmatrix} 0 & 0 \\ 0 & d - b^\top A(J, J)^{-1} b \end{bmatrix}$$

and specifically:

$$\|A - A_J\|_* = \|A - A_J\|_F = \frac{\mathcal{V}(A)}{\mathcal{V}(A(J, J))}, \tag{14}$$

where the second equality has been proved in [33, Lemma 1]. If J is generic and A is SPSD, then $A - A_J$ is SPSD and its nuclear norm is the sum of its diagonal entries which are all Schur complements of the form given in (14); this yields (i).

(ii) The volume of a principal submatrix of an SPSD matrix corresponds to its determinant so that $\sum_{|J|=r} \mathcal{V}(A(J, J))$ is equal $c_{n-r}(A)$. Since the singular values of an SPSD matrix A are equal to its eigenvalues we have $c_{n-r}(A) = \sum_{1 \leq j_1 < \dots < j_r \leq n} \sigma_{j_1}(A) \cdots \sigma_{j_r}(A)$.

(iii) Let us denote $B := A - A_{J_1}$ and $J_2 := J \setminus J_1$. Since $B(J_2, J_2)$ is the Schur complement of $A(J, J)$ with respect to $A(J_1, J_1)$ we have $\mathcal{V}(A(J, J)) = \mathcal{V}(A(J_1, J_1))\mathcal{V}(B(J_2, J_2))$ so that

$$\begin{aligned} \sum_{j_{t+1}, \dots, j_r} \mathcal{V}(A(J, J)) &= \mathcal{V}(A(J_1, J_1)) \sum_{j_{t+1}, \dots, j_r} \mathcal{V}(B(J_2, J_2)) \\ &= \mathcal{V}(A(J_1, J_1)) \cdot (r - t)! \cdot c_{n-r+t}(B), \end{aligned}$$

where the factor $(r - t)!$ accounts the repetitions in the choice of J_2 .

□

Proof (Proof of Theorem 1) Let us denote by Ω_r the set of $r \times r$ principal submatrices of A . We show that $(r + 1) \cdot \sum_{t \geq r+1} \sigma_t(A)$ is larger than the expected value of the cross approximation error, with respect to the following probability distribution on Ω_r :

$$\mathbb{P}(A(J, J)) = \gamma \cdot \mathcal{V}(A(J, J)), \quad \gamma := \frac{1}{\sum_{B \in \Omega_r} \mathcal{V}(B)}.$$

Indeed, we have:

$$\begin{aligned} \mathbb{E}[\|A - A_J\|_*] &= \sum_{|J|=r} \mathbb{P}(A(\widehat{J}, J)) \|A - A_J\|_* \\ \text{Lemma 2 - (i)} &= \sum_{|J|=r} \sum_{|\widehat{J}|=r+1, J \subset \widehat{J}} \mathbb{P}(A(J, J)) \frac{\mathcal{V}(A(\widehat{J}, \widehat{J}))}{\mathcal{V}(A(J, J))} \\ &= \gamma \sum_{|\widehat{J}|=r+1} \sum_{|J|=r, J \subset \widehat{J}} \mathcal{V}(A(\widehat{J}, \widehat{J})) \\ &= \gamma(r+1) \sum_{|\widehat{J}|=r+1} \mathcal{V}(A(\widehat{J}, \widehat{J})) \\ \text{Lemma 2 - (ii)} &= \gamma(r+1) \sum_{1 \leq j_1 < \dots < j_{r+1} \leq n} \sigma_{j_1}(A) \cdots \sigma_{j_{r+1}}(A) \\ &= \gamma(r+1) \sum_{1 \leq j_1 < \dots < j_r \leq n} \sigma_{j_1}(A) \cdots \sigma_{j_r}(A) \sum_{j_{r+1} > j_r} \sigma_{j_{r+1}}(A) \\ &\leq \gamma(r+1)(\sigma_{r+1}(A) + \dots + \sigma_n(A)) \\ &\quad \sum_{1 \leq j_1 < \dots < j_r \leq n} \sigma_{j_1}(A) \cdots \sigma_{j_r}(A) \\ \text{Lemma 2 - (ii)} &= (r+1)(\sigma_{r+1}(A) + \dots + \sigma_n(A)), \end{aligned}$$

where we used that once \widehat{J} is fixed, there are $r + 1$ possible choices for J .

Finally, we have

$$\|A - A_J^*\|_F \leq \|A - A_J^*\|_* \leq (r+1) \sum_{s \geq r+1} \sigma_s(A) \leq \sqrt{n-r}(r+1) \sqrt{\sum_{s \geq r+1} \sigma_s(A)^2},$$

where the last inequality follows from the Cauchy–Schwarz inequality. □

3.2 Derandomizing Theorem 1

Following the approach in [6], we obtain a deterministic algorithm for computing a cross approximation, which verifies (13), by derandomizing Theorem 1. In order to do so, we need to determine the conditional expectation of the cross approximation error, with respect to a partial choice of the indices in J .

Theorem 2 *Let $A \in \mathbb{R}^{n \times n}$ be SPSD and $J_t := \{j_1, \dots, j_t\} \subset \{1, \dots, n\}$ such that $A(J_t, J_t)$ is invertible, then*

$$\mathbb{E}(\|A - A_J\|_* \mid j_1, \dots, j_t) = (r - t + 1) \frac{c_{n-r+t-1}(A - A_{J_t})}{c_{n-r+t}(A - A_{J_t})}.$$

Proof

$$\begin{aligned}
 \mathbb{E}(\|A - A_J\|_* \mid j_1, \dots, j_t) &= \sum_{j_{t+1}, \dots, j_r} \|A - A_J\|_* \mathbb{P}(A(J, J) \mid j_1, \dots, j_t) \\
 &= \sum_{j_{t+1}, \dots, j_r} \|A - A_J\|_* \frac{\mathbb{P}(A(J, J))}{\mathbb{P}(A(J_t, J_t))} \\
 &= \sum_{j_{t+1}, \dots, j_r} \|A - A_J\|_* \frac{\mathcal{V}(A(J, J))}{\sum_{j_{t+1}, \dots, j_{r+1}} \mathcal{V}(A(J, J))} \\
 \text{Lemma 2 - (i)} &= \frac{\sum_{j_{t+1}, \dots, j_{r+1}} \mathcal{V}(A(\{J, j_{r+1}\}, \{J, j_{r+1}\}))}{\sum_{j_{t+1}, \dots, j_r} \mathcal{V}(A(J, J))} \\
 \text{Lemma 2 - (iii)} &= (r - t + 1) \frac{c_{n-r+t-1}(A - A_{J_t})}{c_{n-r+t}(A - A_{J_t})}.
 \end{aligned}$$

□

Theorem 2 suggests to design an iterative scheme that in each step computes the characteristic polynomial of $A - A_{J_t}$ for all the possible choices of the last index j_t and select the one which minimizes $\frac{c_{n-r+t-1}(A - A_{J_t})}{c_{n-r+t}(A - A_{J_t})}$. Interpreting $A - A_{J_t}$ as a rank-1 modification of $A - A_{J_{t-1}}$, we may look at the problem of updating the coefficients of the characteristic polynomial under a rank-1 change of the matrix. Since stable procedures, such as the *Summation Algorithm* [28, Algorithm 1], compute the characteristic polynomial from the eigenvalues, our task boils down to updating the eigenvalues of an SPSD matrix and in turn to computing the eigenvalues of a real diagonal matrix minus a rank-1 symmetric matrix. The latter can be transformed into a symmetric tridiagonal eigenvalue problem with a standard bulge chasing procedure [14, Section 5] and finally solved with Cuppen’s divide and conquer method [8]. Both tridiagonalization and Cuppen’s method require $\mathcal{O}(n^2)$ flops.

The *certified cross approximation (CCA)* obtained from the derandomization of Theorem (1) is reported in Algorithm 6. Note that all the operations inside the inner loop have at most a quadratic cost and computing the eigendecomposition at line 4 is cubic. Therefore, the asymptotic computational cost is $\mathcal{O}(rn^3)$.

3.3 Updating the characteristic polynomial via trace of powers

Each iteration of Algorithm 6 requires to update the eigendecomposition of the residual matrix, resulting in a computational cost $\mathcal{O}(rn^3)$. Here we discuss how, in principle, to reduce the complexity to $\mathcal{O}(r^2n^\omega)$ where $2 < \omega < 3$ is the exponent of the computational complexity of the matrix-matrix multiplication. The idea is that, since we need to update only a (small) portion of the characteristic polynomial we may avoid to deal with the eigendecomposition.

The coefficients of the characteristic polynomial of a matrix A can be expressed with the so called *Plemelj-Smithies formula* [27, Theorem XII 1.108]

Algorithm 6 Certified cross approximation for SPSD matrices

```

1: procedure CCA( $A, r$ )
2:   Set  $R := A, J := \emptyset$ .
3:   for  $t = 1 \dots, r$  do
4:     Compute the eigendecomposition  $R = Q \Lambda Q^\top$ 
5:     Compute the characteristic polynomial of  $R$  via [28, Algorithm 1]
6:      $\text{min\_ratio} \leftarrow \infty$ 
7:     for  $j \in \{1 \dots, n\} \setminus J$  do
8:        $u_j = R(:, j) / \sqrt{R(j, j)}, \tilde{u}_j = Q^\top u_j$ 
9:       Reduce  $\Lambda - \tilde{u}_j \tilde{u}_j^\top$  to a tridiagonal matrix  $T$  via bulge chasing
10:      Compute the eigenvalues of  $T$  with Cuppen’s method
11:      Compute the characteristic polynomial of  $R - u_j u_j^\top$  via [28, Algorithm 1]
12:       $\text{ratio} \leftarrow \frac{c_{n-r+t-1}(R - u_j u_j^\top)}{c_{n-r+t}(R - u_j u_j^\top)}$ 
13:      if  $\text{ratio} < \text{min\_ratio}$  then
14:         $\text{min\_ratio} \leftarrow \text{ratio}, j^* \leftarrow j$ 
15:      end if
16:    end for
17:     $R \leftarrow R - u_{j^*} u_{j^*}^\top$ 
18:     $J \leftarrow J \cup \{j^*\}$ 
19:  end for
20: end procedure

```

$$c_{n-k}(A) = \frac{(-1)^k}{k!} \det \left(\underbrace{\begin{bmatrix} \text{trace}(A) & k-1 & & & \\ \text{trace}(A^2) & \text{trace}(A) & k-2 & & \\ \vdots & \ddots & \ddots & \ddots & \\ \vdots & \ddots & \ddots & \ddots & 1 \\ \text{trace}(A^k) & \dots & \dots & \text{trace}(A^2) & \text{trace}(A) \end{bmatrix}}_{T_{(k)}} \right), \quad (15)$$

so that

$$\frac{c_{n-(k+1)}(A)}{c_{n-k}(A)} = -\frac{1}{k+1} \frac{\det(T_{(k+1)})}{\det(T_{(k)})}. \quad (16)$$

Equation (15) says that for updating the $(n - k)$ -th coefficient of the characteristic polynomial it is sufficient to update the trace of the first k powers of A and to compute the determinant of a $k \times k$ matrix. Interestingly, if $\text{trace}(A), \dots, \text{trace}(A^k)$ are known then the quantities $\text{trace}(A - uu^\top), \dots, \text{trace}((A - uu^\top)^k)$, for a vector $u \in \mathbb{R}^n$, can be computed with a Krylov projection method. More specifically, we have the following property [4, Theorem 3.2]:

$$(A - uu^\top)^k - A^k \in \mathcal{K}_k(A, u) := \text{span}(u, Au, \dots, A^{k-1}u).$$

Let H_k and $\tilde{H}_k := H_k - \|u\|_2 e_1 e_1^\top$ be the orthogonal projections of A and $A - uu^\top$ on $\mathcal{K}_k(A, u)$, then it holds

$$\text{trace}((A - uu^\top)^j) - \text{trace}(A^j) = \text{trace}(\tilde{H}_k^j) - \text{trace}(H_k^j), \quad j = 1, \dots, k. \quad (17)$$

Hence, to update the traces of the first k powers of A we may perform k steps of the Arnoldi method to get \tilde{H}_k, H_k , compute the trace of their powers (via their eigenvalues) and, finally, evaluate (17).

Updating the traces for a single low-rank modification costs $\mathcal{O}(k \cdot \text{matvec}(A) + k^2 n)$; so a procedure that naively applies this computation for the $\mathcal{O}(n)$ low-rank modifications still provides a cubic iteration cost — with respect to n — unless $\text{matvec}(A)$ has a subquadratic cost. In the case $\mathcal{O}(\text{matvec}(A)) = \mathcal{O}(n^2)$, we propose to carry on the Arnoldi step simultaneously for all the $\mathcal{O}(n)$ low-rank modifications $u_i u_i^\top$. More specifically, if $u_{i(h)}$ denotes the h -th vector computed by the Arnoldi process for $\mathcal{K}_k(A, u_i)$, then we perform all the Arnoldi steps together by computing the matrix-matrix multiplication $A \cdot [u_{1(h)} | \dots | u_{n-k+1(h)}]$. Theoretically, this yields the iteration cost $\mathcal{O}(kn^\omega)$. This has also practical benefits because of the use of highly optimized BLAS 3 operations. The procedure for updating the trace of powers is reported in Algorithm 8; the certified cross approximation method (CCA2) that relies on Algorithm 8 is reported in Algorithm 7.

Algorithm 7 CCA via trace of powers

```

1: procedure CCA2( $A, r$ )
2:   Compute  $\mathbf{t} = (t_j)_{j=1, \dots, r+1}$ ,  $t_j = \text{trace}(A^j)$  via  $r$  matrix-matrix multiplications
3:   Set  $R := A$ ,  $J := \emptyset$ 
4:   for  $k := 1, 2, \dots, r$  do
5:      $u_{j_h} = R(:, j_h) / \sqrt{R(j_h, j_h)}$ ,  $j_h \notin J$ 
6:      $U = [u_{j_1} | \dots | u_{j_{r+1}}]$ 
7:      $T \leftarrow \text{UPDATE\_TRACES}(A, \mathbf{t}, U, r - k + 2)$ 
8:     Set  $\text{min\_ratio} = \infty$ 
9:     for  $j_h \in \{1, \dots, n\} \setminus J$  do
10:       $r_{j_h} \leftarrow \frac{c_{n-(r+1)}(R_J - u_{j_h} u_{j_h}^\top)}{c_{n-r}(R_J - u_{j_h} u_{j_h}^\top)}$  via (16)
11:      if  $r_{j_h} < \text{min\_ratio}$  then
12:         $\text{min\_ratio} \leftarrow r_{j_h}$ ,  $h^* \leftarrow h$ 
13:      end if
14:    end for
15:     $\mathbf{t} \leftarrow T(h^*, 1 : r - k + 1)$ 
16:     $R \leftarrow R - u_{j_{h^*}} u_{j_{h^*}}^\top$ 
17:     $J \leftarrow J \cup \{j_{h^*}\}$ 
18:  end for
19:  return  $J$ 
20: end procedure

```

Unfortunately, Algorithm 7 suffers from the numerical instability of evaluating the determinant in (15). More specifically, when the matrix $T_{(k)}$ becomes nearly singular the use of standard techniques provide small singular values, which are accurate only in an absolute sense. Methods that guarantee relative accuracy for singular values

apply only to particular classes of matrices [10,11]; $T_{(k)}$ does not belong to any of such classes. On top of that, we often observe that the matrix $T_{(k)}$ becomes nearly singular quite fast as k increases; typically for k above 10 the computed ratio (16) has no reliable digits. In the next section we propose a strategy to partially circumvent this problem.

Algorithm 8 Update the trace of powers

```

1: procedure UPDATE_TRACES( $A, \mathbf{t}, U, k$ )
2:   Set  $V_j = U(:, j)/\|U(:, j)\|_2$ , ▷  $U \in \mathbb{R}^{n \times s}$ 
3:   for  $h = 1, \dots, k$  do
4:      $U \leftarrow A \cdot U$ 
5:     for  $j = 1, \dots, s$  do
6:        $H_j(1 : h, h) = V_j^\top U(:, j)$ 
7:        $U(:, j) \leftarrow U(:, j) - V_j H_j(1 : h, h)$ 
8:        $H_j(h, h + 1) = \|U(:, j)\|_2$ 
9:        $U(:, j) \leftarrow U(:, j)/H_j(h, h + 1)$ 
10:       $V_j \leftarrow [V_j, U(:, j)]$ 
11:    end for
12:  end for
13:  for  $j = 1, \dots, s$  do
14:     $H_j \leftarrow H_j(1 : k, :)$ 
15:     $\tilde{\sim} = \text{eig}(H_j)$ 
16:     $\tilde{\sim} = \text{eig}(H_j - \|u\|_2 e_1 e_1^\top)$ 
17:     $\hat{\mathbf{t}} = (\hat{t}_j)_{j=1, \dots, k}, \quad t_j = \sum_{h=1}^n (\tilde{\lambda}_h^j - \lambda_h^j)$ 
18:     $T(j, :) = \mathbf{t} + \hat{\mathbf{t}}$ 
19:  end for
20:  return  $T$ 
21: end procedure

```

3.3.1 A restarted algorithm

In view of the instability issues related to evaluating (16), we propose to combine Algorithm 7 with a restarting mechanism. Let us assume that the rank of the sought cross approximation is r and that $\bar{r} < r$ is a small value for which (16) can be computed with a sufficient accuracy. We might think at forming the index set J by the incremental application of Algorithm 7 with input parameter \bar{r} . This means that we first compute a certified cross approximation of rank \bar{r} of A . Then, we add to the latter a certified cross approximation of rank \bar{r} of the residual matrix, and so on and so for. The procedure stops when we reach an index set J of cardinality r . We call this method *quasi certified cross approximation* (QUASI_CCA) and we report its pseudocode in Algorithm 9. The asymptotic cost of QUASI_CCA is r/\bar{r} times the one of CCA2 for a submatrix of size $\bar{r} \times \bar{r}$, that is $\mathcal{O}(\bar{r}rn^\omega)$. Even though the cross approximation returned by Algorithm 9 is not guaranteed to verify (13), it is usually the case, as we will see in the numerical results.

Algorithm 9

```

1: procedure QUASI_CCA( $A, r, \bar{r}$ )
2:   Set  $J = \emptyset$ 
3:   while  $r > 0$  do
4:      $\hat{J} = \text{CCA2}(A, \min\{r, \bar{r}\})$ 
5:      $A \leftarrow A - A\hat{J}$ 
6:      $J \leftarrow J \cup \hat{J}$ 
7:      $r = r - \bar{r}$ 
8:   end while
9:   return  $J$ 
10: end procedure

```

3.4 Numerical results

Let us compare the performances of Algorithm 6 and 9 on the test matrices A_1, A_2, A_3, A_5 introduced in Sect. 2.4. The bulge chasing procedure used in Algorithm 6 has been implemented in Fortran and is called via a MEX interface. When executing Algorithm 9, the parameter \bar{r} has been set to 5.

Test 5. We set $n = 100$, $\rho = 0.85$ and we measure the nuclear norm of the cross approximation error, $\|A - A_J\|_*$, obtained with CCA and QUASI_CCA as the parameter r increases. The results are shown in Fig. 5, where we also report the upper bound provided by Theorem 1 and the lower bound $g(r) := \sum_{j \geq r+1} \sigma_j$, corresponding to the approximation error of the truncated SVD (TSVD). We see that, on all examples, the accuracy of CCA and QUASI_CCA is really close and often the convergence curves are not distinguishable. In addition, in the examples where the decay of the singular values is slow we notice that Theorem 1 tends to be pessimistic and the accuracy of CCA and QUASI_CCA is very close to the one of the TSVD.

Test 6. Finally, we test the computational cost of the proposed numerical procedure. We fix $r = 20$, $\rho = 0.85$ and we run Algorithm 6 and 9 on A_5 for $n \in \{50, 100, 200, 400, 800, 1600\}$. The timings, reported in Fig. 6, confirm the cubic complexity with respect to n of Algorithm 6. Although the complexity of the implementation of QUASI_CCA is cubic as well (no fast matrix multiplication algorithm has been implemented), it results in a significant gain of computational time due to the more intense use of BLAS 3 operations.

4 Outlook

We have proposed several numerical methods for the solution of problems related to the selection of the maximum volume submatrix and the cross approximation of symmetric definite matrices.

We remark that, the idea used for deriving Algorithm 2 and 5 extends easily to combinatorial optimization problems of the form

$$\max_{J \subset \{1, \dots, n\}, |J|=r} f(\mathcal{V}(A_1(J, J)), \dots, \mathcal{V}(A_p(J, J)))$$

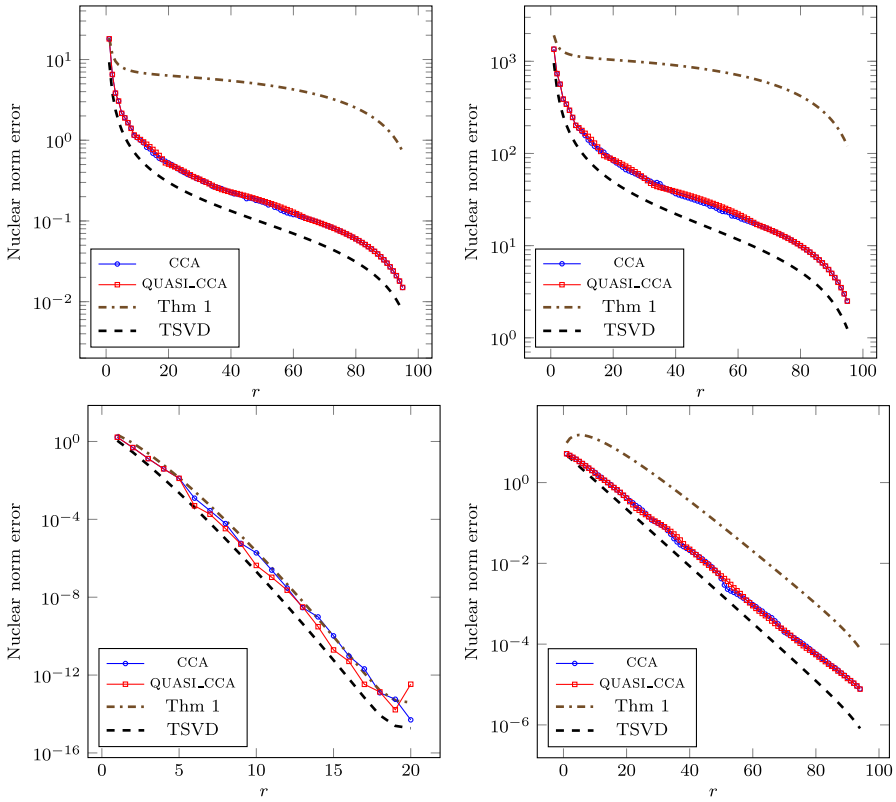



Fig. 5 Nuclear norm of the error associated with the cross approximations returned by Algorithm 6 and 9 on the test matrices A_1 (top-left), A_2 (top-right), A_3 (bottom-left) and A_5 (bottom-right). All plots report the upper bound provided by Theorem 1 and the lower bound given by the error of the truncated SVD

Fig. 6 Timings of Algorithm 6 and 9 on the test matrix A_5 for $r = 20$ and $n \in \{50, 100, 200, 400, 800, 1600\}$



for a multivariate function f and SPSD matrices A_1, \dots, A_p .

Also the second part of the manuscript can inspire some future works. For instance, the fact that the maximum volume submatrix of a diagonally dominant matrix is principal might suggest that a result analogous to Theorem 1 holds also for diagonally dominant matrices. However, it is not straightforward to adjust the proof of Theorem 1 to this case because we lose the connection between the sum of the volumes of the principal submatrices and the coefficients of the characteristic polynomial.

Another interesting point is to understand whether the ratio of determinants in (16) can be computed with high relative accuracy. This would pave the way to the use of CCA2 without incorporating any restart mechanisms.

Finally, in the case of large scale matrices one might derive new scalable algorithms for computing cross approximations by combining Algorithm 6–9 with heuristic techniques for reducing the dependence on n in the computational cost.

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

References

1. Arioli, M., Duff, I.S.: Preconditioning linear least-squares problems by identifying a basis matrix. *SIAM J. Sci. Comput.* **37**(5), S544–S561 (2015)
2. Bebendorf, M.: Approximation of boundary element matrices. *Numer. Math.* **86**(4), 565–589 (2000)
3. Bebendorf, M.: *Hierarchical Matrices*. Springer, New York (2008)
4. Beckermann, B., Kressner, D., Schweitzer, M.: Low-rank updates of matrix functions. *SIAM J. Matrix Anal. Appl.* **39**(1), 539–565 (2018)
5. Çivril, A., Magdon-Ismail, M.: On selecting a maximum volume sub-matrix of a matrix and related problems. *Theoret. Comput. Sci.* **410**(47–49), 4801–4811 (2009)
6. Cortinovis, A., Kressner, D.: Low-rank approximation in the Frobenius norm by column and row subset selection. *arXiv preprint arXiv:1908.06059* (2019)
7. Cortinovis, A., Kressner, D., Massei, S.: On maximum volume submatrices and cross approximation for symmetric semidefinite and diagonally dominant matrices. *Linear Algebra Appl.* **593**, 251–268 (2020)
8. Cuppen, J.J.: A divide and conquer method for the symmetric tridiagonal eigenproblem. *Numer. Math.* **36**(2), 177–195 (1980)
9. Demko, S., Moss, W.F., Smith, P.W.: Decay rates for inverses of band matrices. *Math. Comput.* **43**(168), 491–499 (1984)
10. Demmel, J., Gu, M., Eisenstat, S., Slapničar, I., Veselić, K., Drmač, Z.: Computing the singular value decomposition with high relative accuracy. *Linear Algebra Appl.* **299**(1–3), 21–80 (1999)
11. Demmel, J., Koev, P.: Accurate SVDs of weakly diagonally dominant M-matrices. *Numer. Math.* **98**(1), 99–104 (2004)
12. Deshpande, A., Rademacher, L.: Efficient volume sampling for row/column subset selection. In 2010 IEEE 51st Annual Symposium on Foundations Of Computer Science, pp. 329–338. IEEE, (2010)
13. Fasshauer, G.E., McCourt, M.J.: *Kernel-Based Approximation Methods Using Matlab*, vol. 19. World Scientific Publishing Company, Singapore (2015)
14. Golub, G.H.: Some modified matrix eigenvalue problems. *SIAM Rev.* **15**(2), 318–334 (1973)

15. Goreinov, S. A., Oseledets, I. V., Savostyanov, D. V., Tyrtshnikov, E. E., Zamarashkin, N. L.: How to find a good submatrix. In *Matrix Methods: Theory, Algorithms And Applications: Dedicated to the Memory of Gene Golub*, pp. 247–256. World Scientific, (2010)
16. Goreinov, S.A., Tyrtshnikov, E.E.: The maximal-volume concept in approximation by low-rank matrices. *Contemp. Math.* **280**, 47–52 (2001)
17. Gu, M., Eisenstat, S.C.: Efficient algorithms for computing a strong rank-revealing QR factorization. *SIAM J. Sci. Comput.* **17**(4), 848–869 (1996)
18. Gu, M., Miranian, L.: Strong rank revealing cholesky factorization. *Electron. Trans. Numer. Anal.* **17**, 76–92 (2004)
19. Haber, A., Verhaegen, M.: Sparse solution of the Lyapunov equation for large-scale interconnected systems. *Automatica* **73**, 256–268 (2016)
20. Harbrecht, H., Peters, M., Schneider, R.: On the low-rank approximation by the pivoted Cholesky decomposition. *Appl. Numer. Math.* **62**(4), 428–440 (2012)
21. Higham, N.J.: Accuracy and stability of numerical algorithms, vol. 80. SIAM (2002)
22. Kiefer, J.: Optimum experimental designs V, with applications to systematic and rotatable designs. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pp. 381–405. University of California Press, (1961)
23. Kressner, D., Latz, J., Massei, S., Ullmann, E.: Certified and fast computations with shallow covariance kernels. *arXiv preprint arXiv:2001.09187* (2020)
24. Mikhalev, A., Oseledets, I.V.: Rectangular maximum-volume submatrices and their applications. *Linear Algebra Appl.* **538**, 187–211 (2018)
25. Oseledets, I., Tyrtshnikov, E.: TT-cross approximation for multidimensional arrays. *Linear Algebra Appl.* **432**(1), 70–88 (2010)
26. Osinsky, A.: Rectangular maximum volume and projective volume search algorithms. *arXiv preprint arXiv:1809.02334* (2018)
27. Reed, M., Simon, B.: *Methods of modern mathematical physics, vol. 4 analysis of operators*, (1978)
28. Rehman, R., Ipsen, I.C.: Computing characteristic polynomials from eigenvalues. *SIAM J. Matrix Anal. Appl.* **32**(1), 90–114 (2011)
29. Sommariva, A., Vianello, M.: Computing approximate Fekete points by QR factorizations of vandermonde matrices. *Comput. Math. Appl.* **57**(8), 1324–1336 (2009)
30. Stewart, G. W.: *Matrix Algorithms: Volume 1: Basic Decompositions*. SIAM, (1998)
31. Summa, M. D., Eisenbrand, F., Faenza, Y., Moldenhauer, C.: On largest volume simplices and sub-determinants. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 315–323. SIAM, (2014)
32. Tyrtshnikov, E.: Incomplete cross approximation in the mosaic-skeleton method. *Computing* **64**(4), 367–380 (2000)
33. Zamarashkin, N., Osinsky, A.: On the existence of a nearly optimal skeleton approximation of a matrix in the Frobenius norm. In *Doklady Mathematics*, vol 97, pp. 164–166. Springer, (2018)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.