

A SUBLINEAR-TIME RANDOMIZED ALGORITHM FOR COLUMN AND ROW SUBSET SELECTION BASED ON STRONG RANK-REVEALING QR FACTORIZATIONS*

ALICE CORTINOVIS[†] AND LEXING YING[‡]

Abstract. In this work, we analyze a sublinear-time algorithm for selecting a few rows and columns of a matrix for low-rank approximation purposes. The algorithm is based on an initial uniformly random selection of rows and columns, followed by a refinement of this choice using a strong rank-revealing QR factorization. We prove bounds on the error of the corresponding low-rank approximation (more precisely, the CUR approximation error) when the matrix is a perturbation of a low-rank matrix that can be factorized into the product of matrices with suitable incoherence and/or sparsity assumptions.

Key words. low-rank approximation, column subset selection, strong rank-revealing QR, randomized algorithm, uniform sampling, incoherence

MSC codes. 65F30, 15A23

DOI. 10.1137/24M164063X

1. Introduction. Many large-scale matrices arising in applications have a low numerical rank (see, e.g., [25]), and while the truncated singular value decomposition gives a way to construct the *best* low-rank approximation with respect to all unitarily invariant norms (by the Eckart–Young–Mirsky theorem [11, 19]), this is often too expensive to compute. For this reason, different types of low-rank approximation strategies have been analyzed in the literature—for example, approximations constructed from some rows and columns of the matrix. This work is concerned with the analysis of a randomized algorithm that selects suitable rows and columns.

Let $A \in \mathbb{R}^{n \times m}$ be the matrix we want to approximate. Let us denote by $I \in \{1, \dots, n\}^\ell$ and $J \in \{1, \dots, m\}^\ell$ ordered index sets that correspond to rows and columns of A , respectively, for some $\ell \ll \min\{m, n\}$, and let us denote by $A(I, \cdot) \in \mathbb{R}^{\ell \times m}$ and $A(\cdot, J) \in \mathbb{R}^{n \times \ell}$ the submatrices of A corresponding to the rows indexed by I and the columns indexed by J , respectively. An approximation of A using these rows and columns has the form

$$(1.1) \quad A \approx A(\cdot, J)MA(I, \cdot)$$

for some matrix $M \in \mathbb{R}^{\ell \times \ell}$. The choice of M that minimizes the low-rank approximation error $\|A - A(\cdot, J)MA(I, \cdot)\|_F$ in the Frobenius norm (given a choice of I and J) is the orthogonal projection $M = A(\cdot, J)^\dagger A A(I, \cdot)^\dagger$ (see, e.g., [22]), where \dagger denotes the Moore–Penrose pseudoinverse of a matrix. The resulting approximation is usually called a “CUR approximation” [10, 20, 21]. In this case, the whole matrix A is needed in order to compute M . A less expensive choice for M is $M = A(I, J)^\dagger$, or a regularized version of this quantity, where $A(I, J)$ denotes the $\ell \times \ell$ matrix formed

*Received by the editors February 21, 2024; accepted for publication (in revised form) by A. K. Saibaba August 28, 2024; published electronically January 3, 2025.

<https://doi.org/10.1137/24M164063X>

[†]Department of Mathematics, Stanford University, Stanford, CA 94305 USA (alicecor@stanford.edu).

[‡]Department of Mathematics and ICME, Stanford University, Stanford, CA 94305 USA (lexing@stanford.edu).

by the intersection of the rows indexed by I and the columns indexed by J . This is called “cross approximation” [1] or “double-sided interpolatory decomposition” [9], because the resulting low-rank approximation $A(:, J)A(I, J)^\dagger A(I, :)$ coincides with A in the selected rows and columns.

Row and column selection strategies. The quality of the low-rank approximation, that is, the norm of the error matrix

$$A - A(:, J)MA(I, :),$$

depends on the choice of rows and columns. The CUR approximation error can be easily related to the column subset selection error and the row subset selection error:

$$\|A - A(:, J)A(:, J)^\dagger A\|, \quad \|A - AA(I, :)^{\dagger}A(I, :)\|.$$

If we consider the spectral norm, which we denote by $\|\cdot\|$, it is possible [14] to choose an index set J of cardinality ℓ such that

$$(1.2) \quad \|A - A(:, J)A(:, J)^\dagger A\| \leq \sqrt{1 + \ell(m - \ell)}\sigma_{\ell+1}(A),$$

where $\sigma_{\ell+1}(A)$ denotes the $(\ell + 1)$ th singular value of A and is the error of the best rank- ℓ approximation in the spectral norm.

In practice, the strategy for choosing rows and columns depends on the properties and the size of the matrix. In the literature, several deterministic and randomized strategies for selecting rows and columns for CUR approximation have been developed. If it is feasible to look at all the entries of the matrix, one option is the Adaptive Cross Approximation algorithm [1], a deterministic greedy algorithm that gives good results in practice on a wide range of matrices but can fail for a few others (see, e.g., [7]) and takes time $O(nm\ell)$ for a rank- ℓ approximation. We refer to the reader to, e.g., [6, 10, 20, 21], for discussions on deterministic techniques for column subset selection and CUR approximation.

In this paper, we are concerned with the case in which the matrix is so large that it is not possible (or it is impractical) to store it in memory or compute all its entries: randomized algorithms for sampling columns and rows come to the rescue. Among the strategies proposed in the literature, there are uniform sampling [4], sampling according to the norm of columns and rows of A [12], leverage score sampling [5], volume sampling [6, 8], pivoting on random sketches [9, 26], and the discrete empirical interpolation method [21].

Uniform sampling and rank-revealing factorizations. The simplest method, uniform sampling, gives good low-rank approximations in many cases of interest. In [23], it was shown that if A is a rank- k matrix that admits a low-rank decomposition with *incoherent* factors (see Definition 3.1), uniform sampling of rows and columns allows one to recover the matrix. This favorable property can be extended to matrices that have low *numerical* rank. More specifically, Chiu and Demanet [4] proved that if $A \approx XZY^T$ for some rank- k matrices $X \in \mathbb{R}^{n \times k}$, $Y \in \mathbb{R}^{m \times k}$, $Z \in \mathbb{R}^{k \times k}$, where X and Y have orthonormal columns, if X and Y are incoherent, then selecting ℓ rows and columns of A uniformly at random gives a good low-rank approximation of A if ℓ is chosen to be mildly larger than the target rank k .

When the matrix A does not satisfy these incoherence assumptions, a strategy that has been proposed in the literature is to first *sketch* the matrix A by multiplying it with a random matrix (tall and thin), perform a rank-revealing algorithm on the

sketch to select row indices, and then do an analogous procedure for selecting the column indices; see, e.g., [9, 29]. However, for a nonsparse matrix A , computing a sketch like the ones considered in [9, 29] cannot be done in sublinear time. In the (partial) absence of incoherence, uniform column and row sampling can still be an excellent ingredient for constructing a low-rank approximation of A in sublinear time, and related ideas have appeared multiple times in the literature. In [4, section 3.2], the case in which only the factor X (from the approximate low-rank decomposition of A) is incoherent is discussed: in this situation, a uniformly random row selection is followed by a strong rank-revealing QR (sRRQR) factorization of these rows to identify suitable columns. The resulting low-rank approximation is provably close to optimal. A similar, more heuristic, approach was suggested in [17, Algorithm 2.2]: starting from index sets I and J of rows and columns sampled uniformly at random, one alternately considers the rows $A(I, \cdot)$ and the columns $A(\cdot, J)$ and applies a rank-revealing algorithm (QR with column pivoting) to get a new—hopefully better—set of columns and rows, respectively. While this latter algorithm has been seen to perform very well in practice [17], the available theoretical results are not able to explain its behavior beyond the assumptions of [4]. An analogous strategy that uses sRRQR instead of QR with column pivoting has been investigated in the recent paper [27]. While the focus of our work is the case in which the matrix A is a small additive perturbation of a low-rank matrix, the combination of random uniform sampling with refinement step(s) has also been studied in different, but related, settings, such as, for instance, in the matrix completion literature (see, e.g., [3]) and in the robust PCA literature, where the assumption is that A is the sum of a low-rank matrix and a sparse error matrix; see, e.g., [2, 15].

Contributions. In this work, we consider and analyze a variation of the column and row selection strategies described above. Namely, given a matrix A , we first select some rows uniformly at random, then perform an sRRQR factorization on these rows to get columns, and add some more columns selected uniformly at random to get the final column index set. An analogous strategy is used to select a row index set. The procedure is summarized in Algorithm 2.1. We prove error bounds on the low-rank approximation given by Algorithm 2.1 under the assumption that A is a perturbation of a low-rank matrix XZY^T , for some $X \in \mathbb{R}^{n \times k}$, $Z \in \mathbb{R}^{k \times k}$, $Y \in \mathbb{R}^{m \times k}$, where some of the columns of X and Y are sparse, and the other ones are incoherent. The precise statement is in Corollary 3.13. In addition, we show that if there is no perturbation, that is, the matrix A has rank exactly k , then Algorithm 2.1 exactly recovers A with high probability.

We also consider a second algorithm (Algorithm 4.1), inspired by [17, Algorithm 2.2], that refines the choice of rows and columns iteratively. Under the same assumptions on A , Theorem 4.6 gives an error bound for the low-rank approximation corresponding to one iteration of this method.

Both Algorithm 2.1 and Algorithm 4.1—for the selection of row and column index sets—run in sublinear time with respect to the size of the matrix A . The construction of the corresponding CUR approximation, however, requires access to the full matrix A or to a black-box computation of matrix-vector products with A .

Outline. The paper is organized as follows. In section 2, we recall the notion of a strong rank-revealing QR factorization and describe the randomized sRRQR algorithm for column and row subset selection, together with some conceptual examples. In section 3.1, we report some existing results on incoherence and projection onto column spaces that will be useful for our analysis. The error analysis for Algorithm 2.1

is presented in sections 3.2 and 3.3. In section 4, we discuss the iterative version of the randomized sRRQR algorithm (Algorithm 4.1) and show an error bound. Section 5 illustrates some numerical examples, and section 6 concludes our paper with a discussion of the strengths and limitations of the obtained results.

2. A randomized sRRQR algorithm for column/row selection. We start this section by recalling the definition of an sRRQR factorization and some consequences. We explain the randomized sRRQR algorithm and show its behavior on some simple example matrices.

2.1. Strong rank-revealing QR factorization.

DEFINITION 2.1 (strong rank-revealing QR factorization [14]). *A rank- k strong rank-revealing QR (sRRQR) factorization of $A \in \mathbb{R}^{n \times m}$ is given by*

$$AP = Q \begin{bmatrix} A_k & B_k \\ 0 & C_k \end{bmatrix},$$

where $P \in \mathbb{R}^{m \times m}$ is a permutation matrix, $Q \in \mathbb{R}^{n \times n}$ is an orthogonal matrix, $A_k \in \mathbb{R}^{k \times k}$, $B_k \in \mathbb{R}^{k \times (m-k)}$, and $C_k \in \mathbb{R}^{(n-k) \times (m-k)}$, for which the following conditions hold for some $\eta \geq 1$:

- (a) $\sigma_i(A_k) \geq \frac{\sigma_i(A)}{\sqrt{1+\eta k(m-k)}}$ for $1 \leq i \leq k$;
- (b) $\sigma_j(C_k) \leq \sigma_{k+j}(A) \sqrt{1+\eta k(m-k)}$ for $1 \leq j \leq \min\{n, m\} - k$;
- (c) $\|A_k^{-1} B_k\|_{\max} \leq \sqrt{\eta}$.

Here, $\|\cdot\|_{\max}$ denotes the Chebyshev norm of a matrix, that is, the maximum absolute value of its entries. Intuitively, if the singular values $\sigma_{k+1}(A), \dots, \sigma_m(A)$ of A are small, this means that the columns of A corresponding to the first k columns of P well describe the range of *all* the columns of A . More precisely, let $J \in \{1, \dots, n\}^k$ be the index set corresponding to the first k columns of P . Then

$$\|A - A(:, J)A(:, J)^\dagger A\| = \|C_k\| \leq \sigma_{k+1}(A) \sqrt{1+\eta k(n-k)}.$$

The paper [14] contains an algorithm for obtaining an sRRQR factorization of rank k , which runs in time $\mathcal{O}(mnk \log_\eta m)$ for $\eta > 1$; see the discussion in [14, section 4.4]. In our error analysis in sections 3 and 4, for ease of notation, we assume $\eta = 1$, but all the results can be easily modified to account for a small fixed constant η (e.g., $\eta = 2$).

We choose to work with the sRRQR factorization because of its strong theoretical guarantees. Indeed, while pivoted LU or pivoted QR factorizations with column pivoting are slightly less computationally expensive and are often good strategies for finding the numerical rank of a matrix, there are pathological examples for which these algorithms fail to select a suitable set of columns.

2.2. The algorithm. We consider the following procedure for selecting rows and columns of a matrix $A \in \mathbb{R}^{n \times m}$ in sublinear time. We start from ℓ_0 randomly selected rows of A , corresponding to an index set I_0 . We perform the sRRQR algorithm on the submatrix $A(I_0, :)$ $\in \mathbb{R}^{\ell_0 \times m}$ with parameter ℓ_a and select a column index set J made of the corresponding columns selected by sRRQR and another ℓ_b columns selected uniformly at random. Now, to select the rows I , the same procedure is performed on the matrix A^T . The cost of the algorithm is $\mathcal{O}((m+n)\ell^2 \log_\eta(\max\{n, m\}))$, where we set $\ell = \max\{\ell_0, \ell_a, \ell_b\}$; in particular, for small values of ℓ this is sublinear in the size of A . The number of chosen rows and columns should be slightly larger than the target rank for the low-rank approximation for this algorithm to make sense.

Algorithm 2.1. Randomized sRRQR for row and column selection.

Input: Matrix A , number of indices ℓ_0, ℓ_a, ℓ_b

Output: Row index set I and column index set J of cardinality $\ell_a + \ell_b$

- 1: Select ℓ_0 rows of A uniformly at random (index set I_0)
 - 2: Select ℓ_a columns of $A(I, :)$ by sRRQR algorithm (index set J_a)
 - 3: Select another ℓ_b columns of A uniformly at random (index set J_b)
 - 4: Select ℓ_0 rows of A^T uniformly at random (index set J_0)
 - 5: Select ℓ_a columns of $A^T(J_0, :)$ by sRRQR algorithm (index set I_a)
 - 6: Select another ℓ_b columns of A^T uniformly at random (index set I_b)
 - 7: Return the column index set $J = (J_a, J_b)$ and the row index set $I = (I_a, I_b)$
-

Remark 2.2. Instead of choosing ℓ_a a priori, a more practical implementation would run sRRQR with a fixed accuracy and eliminate the “less interesting” rows and columns.

Remark 2.3. While sampling row and column index sets as in Algorithm 2.1 takes sublinear time in the matrix dimension, the process of computing the low-rank approximation

$$(2.1) \quad A(:, J)A(:, J)^\dagger AA(I, :)^{\dagger} A(I, :)$$

exactly is fast only if the matrix A admits a fast matrix-vector multiplication routine. If this is not the case, the time $\mathcal{O}(mnl)$ needed to compute (2.1) may be too much in practice. A possible heuristic is then to sample some larger index sets \tilde{I} and \tilde{J} and take the middle matrix, which should be $A(:, J)^\dagger AA(I, :)^{\dagger}$, to be

$$A(\tilde{I}, J)^\dagger A(\tilde{I}, \tilde{J})A(I, \tilde{J})^\dagger$$

instead, which is faster to compute; see Step 6 in Algorithm 2.2 from [17].

2.3. When is there hope for Algorithm 2.1 to work?. We begin the discussion by trying to understand in which cases it is reasonable to hope that Algorithm 2.1 gives good rows and columns of A . We present some examples to illustrate its behavior.

When the matrix A can be decomposed as the product of $XZY^T + E$ for some matrices $X \in \mathbb{R}^{n \times k}$ and $Y \in \mathbb{R}^{m \times k}$ with orthonormal columns and some $Z \in \mathbb{R}^{k \times k}$ and $E \in \mathbb{R}^{n \times m}$ with $\|E\|$ small, then it was proven in [4] that if X is incoherent, the first two lines of Algorithm 2.1 give provably good columns. Algorithm 2.1 works on a larger class of matrices, as we will show in Theorem 3.8, Theorem 3.10, and Corollary 3.13 below.

On the other hand, there are cases in which Algorithm 2.1 is not able to select a good row and column subset—for instance, when the matrix A is approximately a block matrix, and some blocks have a small size. The following is a simple conceptual example.

Example 2.4. Consider a matrix $A \in \mathbb{R}^{n \times n}$ that has entries $a_{11} = 1$, $a_{ij} = 1$ for $2 \leq i, j \leq n$, and 0 otherwise. The matrix has rank 2, but it is very unlikely for Algorithm 2.1 to end up including the first row and column in the selected index sets I and J for small values of ℓ_0 and large values of n . Indeed, the initial selection of ℓ_0 rows with probability $(1 - \frac{1}{n})^{\ell_0}$ will not include the first row, and since the first column will be zero in the selected submatrix, the first column will not be included in

of all the entries of X is more or less the same. We informally say that X is incoherent if its coherence is small.

In our analysis, we will use the following result, which states that randomly selecting rows from an incoherent matrix gives us a well-conditioned submatrix with high probability.

THEOREM 3.2 ([24, Lemma 3.4]). *Let $X \in \mathbb{R}^{n \times k}$ be a matrix with orthonormal columns, which is μ -coherent. Let $\ell \geq \alpha k \mu$ for some $\alpha > 0$, and let $I \in \{1, \dots, n\}^\ell$ be a subset of row indices chosen uniformly at random. Then*

$$\mathbb{P} \left(\|X(I, :)\dagger\| \geq \sqrt{\frac{n}{(1-\delta)\ell}} \right) \leq k \left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}} \right)^\alpha \quad \text{for any } \delta \in [0, 1)$$

and

$$\mathbb{P} \left(\|X(I, :)\| \geq \sqrt{\frac{(1+\delta')\ell}{n}} \right) \leq k \left(\frac{e^{\delta'}}{(1+\delta')^{1+\delta'}} \right)^\alpha \quad \text{for any } \delta' \geq 0.$$

COROLLARY 3.3. *Let $B \in \mathbb{R}^{n \times k}$ be a matrix, and let X be an orthonormal basis of the column space of B . Let us assume that X is μ -coherent. Let $\ell \geq \alpha k \mu$ for some $\alpha > 0$, and let $I \in \{1, \dots, n\}^\ell$ be a subset of row indices of X chosen uniformly at random. Then the submatrix $B(I, :)$ has full column rank with probability at least $1 - ke^{-\alpha}$.*

Proof. This follows from the first part of Theorem 3.2 by taking the limit $\delta \rightarrow 1$. \square

We will also use the following bound for the approximation of a matrix A by the projection on a subset of its columns.

THEOREM 3.4 ([4, Corollary 2.4]). *Let $A \in \mathbb{C}^{n \times m}$ and let $U = [U_1 \ U_2]$ be a unitary matrix such that U_1 has k columns. Let J be an index set of $\ell \geq k$ indices. Assume that $U_1(J, :)$ has full column rank. Then*

$$\|A - A(:, J)A(:, J)^\dagger A\| \leq \|U_1(J, :)\dagger\| \cdot \|AU_2U_2(J, :)^T\| + \|AU_2\|.$$

The following linear algebra results will be helpful as well.

THEOREM 3.5 ([16, Exercise 7.3.P16]). *For any matrices F and G (of compatible size) we have*

- (a) $\sigma_i(FG) \leq \|F\|\sigma_i(G)$ for all i , and
- (b) $\sigma_{i+j-1}(F+G) \leq \sigma_i(F) + \sigma_j(G)$ for all i and j .

LEMMA 3.6. *Let $X = \begin{bmatrix} A & B \\ 0 & C \end{bmatrix}$ be a block-triangular matrix where A , B , and C have full column rank. Assume that $\sigma_{\min}(A) \geq \delta_A$, $\|B\|_2 \leq b$, and $\sigma_{\min}(C) \geq \delta_C$. Then*

$$\sigma_{\min}(X) \geq \frac{\delta_A \delta_C}{b} \left(1 + \frac{\delta_A^2 + \delta_C^2}{b^2} \right)^{-1/2}.$$

Proof. We have that $\sigma_{\min}(X)^{-1} = \sqrt{\|(X^T X)^{-1}\|_2}$, and $(X^T X)^{-1} = \begin{bmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{bmatrix}$, where

$$\begin{aligned} X_{11} &= (A^T A)^{-1} + (A^T A)^{-1} A^T B (B^T B + C^T C - B^T A (A^T A)^{-1} A^T B)^{-1} B^T A (A^T A)^{-1}, \\ X_{12} &= -(A^T A)^{-1} A^T B (B^T B + C^T C - B^T A (A^T A)^{-1} A^T B)^{-1}, \\ X_{21} &= -(B^T B + C^T C - B^T A (A^T A)^{-1} A^T B)^{-1} B^T A (A^T A)^{-1}, \\ X_{22} &= (B^T B + C^T C - B^T A (A^T A)^{-1} A^T B)^{-1}; \end{aligned}$$

see, e.g., [18, Theorem 2.1]. We have that $\|(A^T A)^{-1}\| \leq \delta_A^{-2}$ and $\|(C^T C)^{-1}\| \leq \delta_C^{-2}$. Moreover, since $(A^T A)^{-1} A^T$ is the pseudoinverse of A , we have $\|(A^T A)^{-1} A^T\| = \|A(A^T A)^{-1}\| \leq \sigma_{\min}(A)^{-1} \leq \delta_A^{-1}$. Now note that

$$B^T B + C^T C - B^T A(A^T A)^{-1} A^T B = C^T C + B^T (I - A(A^T A)^{-1} A^T) B \succeq C^T C,$$

since $B(I - A(A^T A)^{-1} A^T) B^T$ is symmetric positive semidefinite. Therefore,

$$\|(B^T B + C^T C - B^T A(A^T A)^{-1} A^T B)^{-1}\| \leq \|(C^T C)^{-1}\| \leq \delta_C^{-2}.$$

Thus, we can get upper bounds on the four blocks of $(X^T X)^{-1}$:

$$\begin{aligned} \|X_{11}\| &\leq \delta_A^{-2} (1 + b^2 \delta_C^{-2}), \\ \|X_{12}\|, \|X_{21}\| &\leq \delta_A^{-1} \delta_C^{-2} b, \\ \|X_{22}\| &\leq \delta_C^{-2}. \end{aligned}$$

Using the inequality $\|(X^T X)^{-1}\|^2 \leq \|X_{11}\|^2 + \|X_{12}\|^2 + \|X_{21}\|^2 + \|X_{22}\|^2$ we obtain

$$\begin{aligned} \|(X^T X)^{-1}\|^2 &\leq \delta_A^{-4} (1 + b^2 \delta_C^{-2})^2 + \delta_C^{-4} + 2b^2 \delta_A^{-2} \delta_C^{-4} \\ &\leq \delta_A^{-4} (1 + b^2 \delta_C^{-2})^2 + \delta_C^{-4} + 2\delta_A^{-2} (1 + b^2 \delta_C^{-2})^2 \delta_C^{-2} \\ &= (\delta_A^{-2} (1 + b^2 \delta_C^{-2}) + \delta_C^{-2})^2 = (b^2 \delta_A^{-2} \delta_C^{-2} + \delta_A^{-2} + \delta_C^{-2})^2, \end{aligned}$$

from which we obtain

$$\sigma_{\min}(X) \geq \frac{1}{\sqrt{b^2 \delta_A^{-2} \delta_C^{-2} + \delta_A^{-2} + \delta_C^{-2}}} = \frac{\delta_A \delta_C}{b} \left(1 + \frac{\delta_A^2 + \delta_C^2}{b^2}\right)^{-1/2},$$

which is the desired result. \square

3.2. Analysis for low-rank matrices. In this section, we consider the case in which A has a rank exactly k , and we want to find ℓ rows and columns that span *exactly* the range of A and A^T . Throughout this section, we will assume that the implementation of the sRRQR algorithm, when given a zero residual matrix, selects one of the remaining pivots uniformly at random.

Assumption 3.7. We assume that $A = XZY^T$ for some $X \in \mathbb{R}^{n \times k}$, $Z \in \mathbb{R}^{k \times k}$, and $Y \in \mathbb{R}^{m \times k}$ that have rank exactly k . Let us assume that $Z = \text{blkdiag}(Z_1, Z_2, Z_3)$, and that X and Y can be partitioned as

$$X = [X_1 \ X_2 \ X_3], \quad Y = [Y_1 \ Y_2 \ Y_3],$$

with $X_1 \in \mathbb{R}^{n \times k_1}$, $Y_1 \in \mathbb{R}^{m \times k_1}$, $X_2 \in \mathbb{R}^{n \times k_2}$, $Y_2 \in \mathbb{R}^{m \times k_2}$, $X_3 \in \mathbb{R}^{n \times k_3}$, $Y_3 \in \mathbb{R}^{m \times k_3}$, $Z_1 \in \mathbb{R}^{k_1 \times k_1}$, $Z_2 \in \mathbb{R}^{k_2 \times k_2}$, $Z_3 \in \mathbb{R}^{k_3 \times k_3}$ for some $k_1 + k_2 + k_3 = k$. Let us further assume that $[X_1 \ X_2]$ and $[Y_1 \ Y_3]$ have orthonormal columns and are μ -coherent and that each column of X_3 and Y_2 is sparse and has at most β nonzero elements. See also Figure 1 for a visual illustration.

THEOREM 3.8. *If Assumption 3.7 holds and we take*

$$\ell_0 = \alpha \mu k, \quad \ell_a = k, \quad \ell_b = \alpha \mu k$$

for some $\alpha > 0$, then Algorithm 2.1 returns sets of rows I and columns J such that

$$A = A(:, J)A(:, J)^\dagger AA(I, :)^\dagger A(I, :) = A(:, J)A(I, J)^\dagger A(I, :)$$

with probability at least

$$1 - \left(\frac{\beta \ell k}{n} + \frac{\beta \ell k}{m} + 4ke^{-\alpha} \right),$$

where $\ell = \max\{\ell_0, \ell_a, \ell_b\} = \alpha \mu k$.

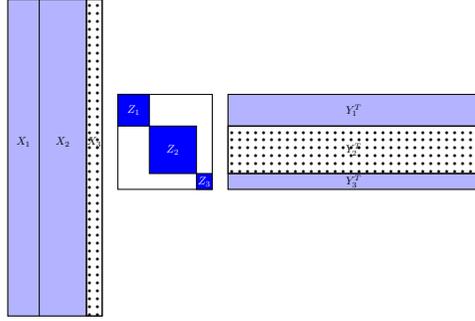


FIG. 1. Illustration of the partitioning of the factorization of the matrix A from Assumption 3.7. The dotted matrices are sparse, the light blue matrices are incoherent, and the dark blue matrices can be anything. (Color available online.)

Proof. Let us prove that lines 1–3 in Algorithm 2.1 select a column index set J such that $A(:, J)$ spans the range of X . In the first step of the algorithm, we select $\ell_0 \geq \alpha\mu(k_1 + k_2)$ rows uniformly at random. By Corollary 3.3, with probability at least $1 - ke^{-\alpha} = 1 - p_1$ the submatrix

$$[X_1(I_0, :) \quad X_2(I_0, :)]$$

has rank $k_1 + k_2$. By the sparsity assumption, with a probability of at least

$$(3.1) \quad 1 - k_3 \left(1 - \left(1 - \frac{\beta}{n} \right)^{\ell_0} \right) = 1 - p_2 \geq 1 - \frac{\beta\ell_0 k_3}{n}$$

the submatrix $X_3(I_0, :)$ is made of zeros. Indeed, for each column of X_3 , by selecting ℓ_0 rows uniformly at random (without replacement), we have a probability of at least $(1 - \frac{\beta}{n})^{\ell_0}$ of sampling only zero entries, and a probability of at most $1 - (1 - \frac{\beta}{n})^{\ell_0}$ of selecting at least one nonzero entry; the bound (3.1) follows from a union bound on the k_3 columns of X_3 . Therefore, using again a union bound, with probability at least $1 - p_1 - p_2$, the matrix $A(I_0, :)$ has rank exactly $k_1 + k_2$.

The sRRQR algorithm gives an index set J_a of cardinality $k_1 + k_2$. From the rank-revealing property, we have that the top $(k_1 + k_2) \times (k_1 + k_2)$ matrix $A(I, J_a)$ has a full rank (is invertible), and therefore

$$(3.2) \quad [Y_1(J_a, :) \quad Y_2(J_a, :)] \in \mathbb{R}^{(k_1+k_2) \times (k_1+k_2)}$$

is also invertible. The index set J_b is chosen uniformly at random from $\{1, \dots, n\}$. We can partition

$$Y(J, :) = \begin{bmatrix} Y_1(J_a, :) & Y_2(J_a, :) & Y_3(J_a, :) \\ Y_1(J_b, :) & Y_2(J_b, :) & Y_3(J_b, :) \end{bmatrix}.$$

The matrix $Y_2(J_a, :)$ has rank k_2 because it is a submatrix with k_2 columns of the invertible matrix (3.2). With probability at least $1 - ke^{-\alpha} = 1 - p_3$, by Corollary 3.3 the submatrix $[Y_1(J_b, :) \quad Y_3(J_b, :)]$ has full rank $k_1 + k_3$. With probability at least

$$1 - k_2 \left(1 - \left(1 - \frac{\beta}{m} \right)^{\ell_b} \right) = 1 - p_4 \leq 1 - \frac{\beta\ell_b k_2}{m}$$

the matrix $Y_2(J_b, \cdot)$ is made of zeros. Therefore, with probability at least $1 - p_3 - p_4$, the matrix $Y(J, \cdot)$ has rank k . This implies that the matrix

$$A(:, J) = XZY(J, \cdot)^T$$

also has full rank, so it spans the range of the columns of A . In a completely analogous way, one can show that, with probability at least

$$1 - 2ke^{-\alpha} - \frac{\beta\ell_0k_2}{m} - \frac{\beta\ell_bk_3}{n},$$

the matrix $A(I, \cdot)$ spans the range of the rows of A . Overall, this implies the exactness of the CUR approximation $A = A(:, J)A(:, J)^\dagger AA(I, \cdot)^\dagger A(I, \cdot)$ with probability at least

$$1 - 4ke^{-\alpha} - \frac{\beta\ell k}{n} - \frac{\beta\ell k}{m}.$$

Moreover, it implies that the matrix $A(I, J)$ has rank k , so $A = A(:, J)A(I, J)^\dagger A(I, \cdot)$. \square

3.3. Analysis for numerically low-rank matrices. We now consider the case in which A is a perturbation of order ε of a rank- k matrix of the form described in Assumption 3.7.

Assumption 3.9. Assume that $A = XZY^T + E$ for some $X \in \mathbb{R}^{n \times k}$, $Z \in \mathbb{R}^{k \times k}$, $Y \in \mathbb{R}^{m \times k}$, and $E \in \mathbb{R}^{n \times m}$. Assume that $\text{rank}(XZY^T) = k$ and $\|E\| \leq \varepsilon$. Let us assume that $Z = \text{blkdiag}(Z_1, Z_2, Z_3)$, and that we can partition X and Y as

$$X = [X_1 \ X_2 \ X_3], \quad Y = [Y_1 \ Y_2 \ Y_3],$$

with $X_1 \in \mathbb{R}^{n \times k_1}$, $Y_1 \in \mathbb{R}^{m \times k_1}$, $X_2 \in \mathbb{R}^{n \times k_2}$, $Y_2 \in \mathbb{R}^{m \times k_2}$, $X_3 \in \mathbb{R}^{n \times k_3}$, $Y_3 \in \mathbb{R}^{m \times k_3}$, $Z_1 \in \mathbb{R}^{k_1 \times k_1}$, $Z_2 \in \mathbb{R}^{k_2 \times k_2}$, $Z_3 \in \mathbb{R}^{k_3 \times k_3}$, for some $k_1 + k_2 + k_3 = k$. Let us further assume that the block matrices $[X_1 \ X_2]$ and $[Y_1 \ Y_3]$ have orthonormal columns and are μ -coherent, and that X_3 and Y_2 have orthonormal columns and each column is sparse and has at most β nonzero elements.

A factorization of the form of Assumption 3.9 could be obtained from a truncated singular value decomposition of A , but the decomposition XZY^T is more general as it does not require full orthonormality of the X and Y factors, or the fact that Z is diagonal. Moreover, the fact that X_3 and Y_2 have orthonormal columns is needed only marginally in our analysis; see Remark 3.11 below. While, in general, it is not easy to show that a given matrix satisfies Assumption 3.9, our goal in this work is to extend the class of matrices for which it can be shown that a combination of random uniform sampling and rank-revealing factorizations provably gives a good low-rank approximation. Note that Assumption 3.9 extends the assumptions in [4, section 3.2] by allowing some of the factors in the (approximate) decomposition to be sparse.

The first result in this section, Theorem 3.10, provides an error bound for the column subset selection error given by the index set J selected in lines 1–3 of Algorithm 2.1.

THEOREM 3.10. *If Assumption 3.9 holds and we take*

$$\ell_0 = \alpha\mu(k_1 + k_2), \quad \ell_a = k_1 + k_2, \quad \ell_b = \alpha\mu(k_1 + k_3)$$

for some $\alpha > 0$, then Algorithm 2.1 returns a column index J such that

$$(3.3) \quad \|A - A(:, J)A(:, J)^\dagger A\| \leq \varepsilon \left(1 + \sqrt{2}\delta_A^{-1}\delta_C^{-1} (1 + \delta_A^2 + \delta_C^2)^{1/2}\right),$$

where

$$\delta_A^{-1} = \sqrt{\frac{1+\delta}{1-\delta}} \cdot \frac{\sigma_1(XZY^T)\sqrt{km}\|Y^\dagger\|}{\sigma_k(XZY^T) - 2\varepsilon\sqrt{\frac{nmk}{(1-\delta)\ell}}} \quad \text{and} \quad \delta_C^{-1} = \sqrt{\frac{m}{(1-\delta)\ell}},$$

and $\ell = \max\{\ell_0, \ell_a, \ell_b\}$, with a probability of at least

$$(3.4) \quad p = 1 - k \left(\frac{\beta\ell}{n} + \frac{\beta\ell}{m} + \left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}} \right)^\alpha + \left(\frac{e^\delta}{(1+\delta)^{1+\delta}} \right)^\alpha \right) \quad \text{for all } \delta \in (0, 1).$$

The general idea of the proof of Theorem 3.10 is similar to the proof of Theorem 3.8, but in this case, instead of showing that the selected submatrices have the right rank, we also need to prove that they are not too ill-conditioned.

Proof. We divide the proof into several steps.

Step 1. A lower bound for $\sigma_{k_1+k_2}(A(I_0, :))$. Let us consider line 1 in Algorithm 2.1, that is, the selection of ℓ_0 rows uniformly at random, corresponding to the index set I_0 . We have

$$A(I_0, :) = X_1(I_0, :)Z_1Y_1^T + X_2(I_0, :)Z_2Y_2^T + X_3(I_0, :)Z_3Y_3^T + E(I_0, :).$$

With probability at least $1 - k_3(1 - (1 - \frac{\beta}{n})^{\ell_0}) = 1 - p_1$, the submatrix $X_3(I_0, :)$ is made of zeros; the argument is the same as in the proof of Theorem 3.8. By Theorem 3.2, with probability at least $1 - (k_1 + k_2)(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}})^\alpha = 1 - p_2$, we have

$$\left\| [X_1(I_0, :), X_2(I_0, :)]^\dagger \right\| \leq \sqrt{\frac{n}{(1-\delta)\ell_0}}.$$

Therefore, with probability at least $1 - p_1 - p_2$, we have

$$\begin{aligned} \sigma_{k_1+k_2}(A(I_0, :)) &\geq \sigma_{k_1+k_2}(A(I_0, :) - E(I_0, :)) - \sigma_1(E(I_0, :)) \\ &= \sigma_{k_1+k_2}(X_1(I_0, :)Z_1Y_1^T + X_2(I_0, :)Z_2Y_2^T) - \|E(I_0, :)\| \\ &\geq \sigma_{k_1+k_2}(\tilde{X}(I_0, :)\tilde{Z}\tilde{Y}^T) - \varepsilon, \end{aligned}$$

where the first inequality follows from Theorem 3.5(b) applied to $F = A(I_0, :)$ and $G = -E(I_0, :)$ with $i = k_1 + k_2$ and $j = 1$, we used that $\|E(I_0, :)\| \leq \|E\| \leq \varepsilon$, and we denoted

$$\tilde{X} = [X_1 \quad X_2], \quad \tilde{Z} = \begin{bmatrix} Z_1 \\ Z_2 \end{bmatrix}, \quad \tilde{Y} = [Y_1 \quad Y_2].$$

Note that \tilde{X} has orthonormal columns, \tilde{Z} is invertible, \tilde{Y} has full column rank, and $\tilde{X}(I_0, :)$ also has full column rank (with probability at least $1 - p_2$). We can further bound

$$\begin{aligned} \sigma_{k_1+k_2}(\tilde{X}(I_0, :)\tilde{Z}\tilde{Y}^T) &= \|(\tilde{X}(I_0, :)\tilde{Z}\tilde{Y}^T)^\dagger\|^{-1} = \|(\tilde{Y}^T\tilde{Z})^\dagger\tilde{X}(I_0, :)\|^{-1} \\ &\geq \|(\tilde{Y}^T\tilde{Z})^\dagger\|^{-1}\|\tilde{X}(I_0, :)\|^{-1} \geq \sigma_{k_1+k_2}(\tilde{Z}\tilde{Y}^T)\sqrt{\frac{(1-\delta)\ell_0}{n}} \\ &= \sigma_{k_1+k_2}(\tilde{X}\tilde{Z}\tilde{Y}^T)\sqrt{\frac{(1-\delta)\ell_0}{n}} \geq \sigma_k(XZY^T)\sqrt{\frac{(1-\delta)\ell_0}{n}}, \end{aligned}$$

where the last inequality is obtained from Theorem 3.5(b) applied to $F = X_1 Z_1 Y_1^T + X_2 Z_2 Y_2^T$ and $G = X_3 Z_3 Y_3^T$ with $i = k_1 + k_2$ and $j = k_3 + 1$. Therefore, we can conclude that

$$(3.5) \quad \sigma_{k_1+k_2}(A(I_0, :)) \geq \sigma_k(XZY^T) \sqrt{\frac{(1-\delta)\ell_0}{n}} - \varepsilon$$

with probability at least $1 - p_1 - p_2$.

Step 2. The matrix $A(I_0, J_a)$ is well-conditioned. Let us now consider the column index set $J = (J_a, J_b)$ where J_a contains the ℓ_a indices from sRRQR, and J_b contains the new ℓ_b indices selected uniformly at random. Since J_a has been selected using sRRQR on $A(I_0, :)$, we have that

$$(3.6) \quad \sigma_{k_1+k_2}(A(I_0, J_a)) \geq \frac{\sigma_{k_1+k_2}(A(I_0, :))}{\sqrt{1 + (k_1 + k_2)(m - k_1 - k_2)}} \geq \frac{\sigma_{k_1+k_2}(A(I_0, :))}{\sqrt{km}},$$

with probability at least $1 - p_1 - p_2$.

Step 3. The matrix $[Y_1(J_a, :), Y_2(J_a, :)]$ is well-conditioned. With probability at least $1 - p_1 - p_2$, we have that

$$\begin{aligned} A(I_0, J_a) &= [X_1(I_0, :), X_2(I_0, :)] \begin{bmatrix} Z_1 & \\ & Z_2 \end{bmatrix} [Y_1(J_a, :), Y_2(J_a, :)]^T + E(I_0, J_a) \\ &= [X_1(I_0, :), X_2(I_0, :)] \begin{bmatrix} Z_1 & \\ & Z_2 \end{bmatrix} [Y_1, Y_2]^T \left([Y_1, Y_2]^T \right)^\dagger [Y_1(J_a, :), Y_2(J_a, :)]^T \\ &\quad + E(I_0, J_a). \end{aligned}$$

We have that

$$(3.7)$$

$$\begin{aligned} &\sigma_{k_1+k_2}([Y_1(J_a, :), Y_2(J_a, :)]) \\ &\geq \frac{\sigma_{k_1+k_2}(A(I_0, J_a) - E(I_0, J_a))}{\|A(I_0, :) - E(I_0, :)\| \cdot \|[X_1(I_0, :), X_2(I_0, :)]\| \cdot \|Y^\dagger\|} \\ &\geq \sqrt{\frac{n}{(1+\delta)\ell_0}} \cdot \frac{\sigma_{k_1+k_2}(A(I_0, J_a)) - \varepsilon}{\sigma_1(X(I_0, :))ZY^T} \cdot \|Y^\dagger\| \geq \sqrt{\frac{n}{(1+\delta)\ell_0}} \cdot \frac{\sigma_{k_1+k_2}(A(I_0, :)) - \varepsilon\sqrt{km}}{\sqrt{km} \cdot \sigma_1(XZY^T) \cdot \|Y^\dagger\|} \\ &\geq \sqrt{\frac{1-\delta}{1+\delta}} \cdot \frac{\sigma_k(XZY^T) - 2\varepsilon\sqrt{\frac{nmk}{(1-\delta)\ell_0}}}{\sigma_1(XZY^T)\sqrt{km}\|Y^\dagger\|}, \end{aligned}$$

where we used Theorem 3.5(a) for the first inequality, the second inequality holds with probability at least $1 - (k_1 + k_3)\left(\frac{e^\delta}{(1+\delta)^{1+\delta}}\right)^\alpha = 1 - p_3$ by Theorem 3.2, because I_0 was selected uniformly at random, the third inequality follows from (3.6), and the fourth follows from (3.5). Therefore, (3.7) holds with probability at least $1 - p_1 - p_2 - p_3$.

Step 4. The matrix $Y(J, :)$ is well-conditioned. We can partition

$$Y(J, :) = \begin{bmatrix} Y_1(J_a, :), & Y_2(J_a, :), & Y_3(J_a, :) \\ Y_1(J_b, :), & Y_2(J_b, :), & Y_3(J_b, :) \end{bmatrix} = \begin{bmatrix} Y_1(J_a, :), & Y_2(J_a, :), & Y_3(J_a, :) \\ Y_1(J_b, :), & 0 & Y_3(J_b, :) \end{bmatrix},$$

where the second equality holds with a probability of at least

$$1 - p_4 = 1 - k_2 \left(1 - \left(1 - \frac{\beta}{m} \right)^{\ell_b} \right),$$

for the same reason as (3.1) in the proof of Theorem 3.8. In the rest of this proof, we assume that all equalities and inequalities above hold, which happens with probability at least $1 - p_1 - p_2 - p_3 - p_4$. We wish to apply Lemma 3.6 with $A = Y_2(J_a, :)$, $B = [Y_1(J_a, :)\ Y_3(J_a, :)]$, and $C = [Y_1(J_b, :)\ Y_3(J_b, :)]$. Using Theorem 3.2 again, we have that

$$\left\| [Y_1(J_b, :)\ Y_3(J_b, :)]^\dagger \right\| \leq \sqrt{\frac{m}{(1-\delta)\ell_b}}$$

with probability at least $1 - (k_1 + k_3)\left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right)^\alpha = 1 - p_4$. We have that

$$\left\| [Y_1(J_a, :)\ Y_3(J_a, :)] \right\| \leq 1$$

because this is a submatrix of a matrix with orthonormal columns. Furthermore, we have that

$$\sigma_{k_2}(Y_2(J_a, :)) \geq \sigma_{k_1+k_2}([Y_1(J_a, :)\ Y_2(J_a, :)])$$

and the right-hand side can be bounded by (3.7). Therefore, the assumptions of Lemma 3.6 are satisfied with

$$\delta_A^{-1} = \sqrt{\frac{1+\delta}{1-\delta}} \cdot \frac{\sigma_1(XZY^T)\sqrt{km}\|Y^\dagger\|}{\sigma_k(XZY^T) - 2\varepsilon\sqrt{\frac{nmk}{(1-\delta)\ell_0}}}, \quad \delta_C^{-1} = \sqrt{\frac{m}{(1-\delta)\ell_b}}, \quad b=1,$$

so

$$(3.8) \quad \|Y(J_1, :)\dagger\| \leq \delta_A^{-1}\delta_C^{-1}(1 + \delta_A^2 + \delta_C^2)^{1/2}.$$

Step 5. The columns are a good subset for low-rank approximation. In order to bound $\|A - A(:, J)A(:, J)^\dagger A\|$ we apply Theorem 3.4 to the matrix A , the unitary matrix $U = [U_1\ U_2]$, where U_1 is an orthonormal basis of Y (which has rank k) and U_2 is an orthogonal complement of Y , and the index set J . We have that $\|AU_2\| = \|EU_2\| \leq \varepsilon$, and $\|U_2(J, :)\| \leq 1$. To bound $\|U_1(J, :)\dagger\|$ let us consider the skinny QR decomposition $Y = U_1R$ for some invertible matrix $R \in \mathbb{R}^{k \times k}$; then

$$\|U_1(J, :)\dagger\| = \|R \cdot Y(J, :)\dagger\| \leq \|Y(J, :)\dagger\| \cdot \|Y\| \leq \sqrt{2}\delta_A^{-1}\delta_C^{-1}(1 + \delta_A^2 + \delta_C^2)^{1/2},$$

where the last equality follows from (3.8) and the fact that Y has two blocks of orthonormal columns. Applying Theorem 3.4 then gives

$$(3.9) \quad \|A - A(:, J)A(:, J)^\dagger A\| \leq \varepsilon(1 + \sqrt{2}\delta_A^{-1}\delta_C^{-1}(1 + \delta_A^2 + \delta_C^2)^{1/2}).$$

Step 6. The failure probability. The success probability p can be bounded using a union bound on all the probabilistic results used in this proof: we have

$$\begin{aligned} p &\geq 1 - (p_1 + p_2 + p_3 + p_4) \\ &\geq 1 - k_3 \left(1 - \left(1 - \frac{\beta}{n}\right)^{\ell_0}\right) - (k_1 + k_2) \left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right)^\alpha \\ &\quad - (k_1 + k_3) \left(\frac{e^\delta}{(1+\delta)^{1+\delta}}\right)^\alpha - k_2 \left(1 - \left(1 - \frac{\beta}{m}\right)^{\ell_b}\right) \\ &\geq \left(1 - \frac{\beta\ell_0 k_3}{n}\right) + \left(1 - \frac{\beta\ell_b k_2}{m}\right) - k \left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right)^\alpha - k \left(\frac{e^\delta}{(1+\delta)^{1+\delta}}\right)^\alpha, \end{aligned}$$

which concludes the proof of the theorem. \square

Remark 3.11. The fact that X_3 and Y_2 have orthonormal columns in Assumption 3.9, which does not appear in Assumption 3.7, is used in Step 5 of the proof to bound $\|Y\| \leq \sqrt{2}$. The proof of Theorem 3.10 would also work for nonorthonormal X_3 and Y_2 , as long as the matrices X and Y have full column rank; the only difference would be an extra factor $\|Y\|$ in the right-hand side of the bound. We also remark that, while X and Y are not required to have (overall) orthonormal columns in Assumption 3.9, the condition number of Y does come into play in the bound of Theorem 3.10.

Let us provide a simplified bound below that “hides” all the constants of Theorem 3.10 that distract from the essence of the result.

COROLLARY 3.12. *In the same assumptions as Theorem 3.10, we have that*

$$\|A - A(:, J)A(:, J)^\dagger A\| \leq \mathcal{O}\left(\varepsilon m \sqrt{\frac{k}{\ell}} \cdot \frac{\sigma_1(XZY^T)}{\sigma_k(XZY^T)} \cdot \|Y^\dagger\|\right)$$

with a probability of at least

$$1 - k \left(\frac{\beta\ell}{n} + \frac{\beta\ell}{m} + c_1^\alpha + c_2^\alpha \right)$$

for some fixed constants $c_1, c_2 \in (0, 1)$.

For example, if we fix $\delta = 4/5$, the success probability (3.4) gives $c_1 \approx 0.62$ and $c_2 \approx 0.78$. For this corollary (and Theorem 3.10) to make sense, we need to have $\varepsilon m \sqrt{k/\ell} \ll \sigma_k(XZY^T)$; that is, the magnitude of the perturbation E needs to be small compared to the smallest “relevant” singular value of A .

As mentioned in the introduction, there is a well-known way to relate the error of row/column subset selection to the low-rank approximation error; see, e.g., [6, 10, 20, 21]. In particular, for any index sets I and J , we have

$$\begin{aligned} & \|A - A(:, J)A(:, J)^\dagger AA(I, :)^{\dagger} A(I, :)\|_2 \\ & \leq \|A - A(:, J)A(:, J)^\dagger A\| + \|A(:, J)A(:, J)^\dagger (A - AA(I, :)^{\dagger} A(I, :))\| \\ & \leq \|A - A(:, J)A(:, J)^\dagger A\| + \|A(:, J)A(:, J)^\dagger\| \cdot \|A - AA(I, :)^{\dagger} A(I, :)\| \\ (3.10) \quad & = \|A - A(:, J)A(:, J)^\dagger A\| + \|A - AA(I, :)^{\dagger} A(I, :)\|, \end{aligned}$$

where the last equality follows from the fact that $A(:, J)A(:, J)^\dagger$ is an orthogonal projection.

COROLLARY 3.13. *Given a matrix A satisfying the assumptions of Theorem 3.10 and given the output index sets I and J from Algorithm 2.1 with $\ell_0 = \ell_b = \alpha\mu k$ for some $\alpha > 0$ and $\ell_a = k$, we have that*

$$\begin{aligned} & \|A - A(:, J)A(:, J)^\dagger AA(I, :)^{\dagger} A(I, :)\| \\ & \leq \mathcal{O}\left(\varepsilon(m+n) \sqrt{\frac{k}{\ell}} \cdot \frac{\sigma_1(XZY^T)}{\sigma_k(XZY^T)} \cdot (\|X^\dagger\| + \|Y^\dagger\|)\right) \end{aligned}$$

with a probability of at least

$$1 - 2k \left(\frac{\beta\ell}{n} + \frac{\beta\ell}{m} + c_1^\alpha + c_2^\alpha \right).$$

Algorithm 4.1. Iterative randomized sRRQR for row and column selection.

Input: Matrix $A \in \mathbb{R}^{n \times m}$, iterations H , numbers $\ell_0, \ell_h^{\text{sRRQR, col}}, \ell_h^{\text{new, col}}, \ell_h^{\text{sRRQR, row}}, \ell_h^{\text{new, row}}$,

Output: Row index set I_H and column index set J_H

- 1: Select ℓ_0 row indices I_0 uniformly at random
 - 2: **for** $h = 1, \dots, H$ **do**
 - 3: Apply sRRQR algorithm to $A(I_{h-1}, :)$ with parameter $\ell_h^{\text{sRRQR, col}}$
 - 4: Build column index set J_h from the $\ell_h^{\text{sRRQR, col}}$ indices given by sRRQR factorization and $\ell_h^{\text{new, col}}$ new indices selected uniformly at random
 - 5: Apply sRRQR algorithm to $A(:, J_h)^T$ with parameter $\ell_h^{\text{sRRQR, row}}$
 - 6: Build row index set I_h from the $\ell_h^{\text{sRRQR, row}}$ indices given by sRRQR factorization and $\ell_h^{\text{new, row}}$ new indices selected uniformly at random
 - 7: **end for**
-

Proof. The bound follows from the inequality (3.10) and Theorem 3.10 applied to A and A^T . The bound on the failure probability follows from a union bound of the failure probabilities given by Theorem 3.10 on A and A^T . \square

4. Iterative randomized sRRQR algorithm. In this section, we discuss a variation of Algorithm 2.1 that is inspired by [17, Algorithm 2.2]. We first present the algorithm (Algorithm 4.1) and then detail its relation to Algorithm 2.1 and the literature.

We start from ℓ_0 randomly selected rows of A , corresponding to an index set I_0 . We perform the sRRQR algorithm on the submatrix $A(I_0, :) \in \mathbb{R}^{\ell_0 \times m}$ and select a column index set J_1 made of the corresponding columns selected by sRRQR plus some other columns selected uniformly at random. Now we look at submatrix $A(:, J_1)$ and perform the sRRQR algorithm on its transpose. This gives us a new set of row indices, to which we add some more new rows, sampled uniformly at random, and merge them into the new row index set I_1 . This procedure can be iterated a couple of times by alternately optimizing the columns and the rows. The cost of each iteration is $\mathcal{O}((m+n)\ell^2 \log_\eta(\max\{n, m\}))$, where ℓ is the total number of indices we are considering; in particular, for small values of ℓ this is sublinear in the size of A . The number of chosen indices does not need to be the same in each step of the algorithm; see Algorithm 4.1 below. The number of rows and columns should be chosen to be slightly larger than k , the target rank for the low-rank approximation.

Remark 4.1. In Algorithm 4.1, some of the row/column indices of A that have been sampled get thrown away at each step; this has been done because it simplifies our analysis in the rest of the paper. In practice, one can use the extra information contained in those rows and columns and consider final index sets $I = I_0 \cup I_1 \cup \dots \cup I_H$ and $J = J_0 \cup J_1 \cup \dots \cup J_H$. The corresponding low-rank approximation will be slightly more expensive to compute but will also yield a smaller error. Indeed, if $I \subseteq I'$ and $J \subseteq J'$, then

$$\|A - A(:, J')A(:, J')^\dagger AA(I', :)^{\dagger} A(I', :)\| \leq \|A - A(:, J)A(:, J)^\dagger AA(I, :)^{\dagger} A(I, :)\|.$$

Some additional remarks are in order:

- Algorithm 4.1 is almost the same as Algorithm 2.2 in [17], with the difference that we use an sRRQR factorization instead of a pivoted QR factorization,

which allows us to obtain theoretical guarantees on the quality of the selected rows and columns.

- An iterative refinement of the initial choice of row and column indices has been shown to be effective in providing more accurate low-rank approximations in the setting of interpolatory decompositions. For example, the paper [27] considers an interpolatory decomposition for which the row and column index sets are obtained from Algorithm 4.1. A similar strategy was proposed in [13]: here, the authors employed a volume maximization technique, closely related to sRRQR factorizations.
- The first three steps of Algorithm 4.1 (lines 1, 3, 4) coincide with the first three lines of Algorithm 2.1; what changes here afterward is that, instead of restarting the sampling process from scratch to select the row indices, one uses the selected columns as a starting point. The advantage of Algorithm 4.1 is that the process can be iterated, and the quality of the low-rank approximation improves in many cases, in practice (see Example 5.2). Intuitively, performing more iterations of Algorithm 4.1 allows us to “see” more information from the matrix A due to the new rows and columns sampled at each step, and the sRRQR factorization allows us only to keep the “important” parts.

Let us revisit Examples 2.4 and 2.5 in the context of Algorithm 4.1. We postpone further examples to section 5.

Example 4.2. Let us consider again the matrix A from Example 2.4. Also, for Algorithm 4.1, it is very unlikely to end up selecting the first row and column, for moderately large values of n and small values of the ℓ -parameters. Indeed, the initial selection of ℓ_0 rows with probability $(1 - \frac{1}{n})^{\ell_0}$ will not include the first row, and since the first column will be zero in the selected submatrix, the first column will not be included in the column index set selected by sRRQR. At the next steps, again, the probability of *not* selecting the first row or column is high in this setting.

Example 4.3. Let us consider again the matrix from Example 2.5 and let us show that Algorithm 4.1 is able to recover A successfully. The row set I_0 , chosen uniformly at random, will most likely not include the first row—with probability $(1 - \frac{1}{n})^{\ell_0}$. However, when looking at the matrix $A(I_0, :)$, the sRRQR algorithm will select a set J_1 containing the first column, plus some other columns sampled uniformly at random. Now, when applying sRRQR on the column subset $A(:, J_1)$, we will get an index set I_1 that includes the first row. At this point, the index sets I_1 and J_1 are enough to recover the rank-2 matrix A .

4.1. Error analysis. We restrict the analysis of Algorithm 4.1 to its first full step. We provide a result that states that exact reconstruction happens with high probability when the matrix A is exactly low-rank and satisfies Assumption 3.7, and an error bound in the case in which A has been perturbed by error, similarly to what we have done for the analysis of Algorithm 2.1.

THEOREM 4.4. *If Assumption 3.7 holds and we take*

$$\ell_0 = \alpha\mu(k_1 + k_2), \ell_1^{\text{sRRQR, col}} = k_1 + k_2, \ell_1^{\text{new, col}} = \alpha\mu(k_1 + k_3), \\ \ell_1^{\text{sRRQR, row}} = k, \ell_1^{\text{new, row}} = 0,$$

for some $\alpha > 0$, then one iteration of Algorithm 4.1 returns sets of rows I_1 and columns J_1 such that

$$A = A(:, J_1)A(:, J_1)^\dagger AA(I_1, :)^{\dagger} A(I_1, :) = A(:, J_1)A(I_1, J_1)^\dagger A(I_1, :)$$

with a probability of at least

$$1 - \left(\frac{\beta \ell k}{n} + \frac{\beta \ell k}{m} + 2ke^{-\alpha} \right),$$

where $\ell = \max\{\ell_0, \ell_1^{\text{sRRQR, col}}, \ell_1^{\text{new, col}}, \ell_1^{\text{sRRQR, row}}\} \leq \alpha \mu k$.

Proof. Since lines 1–4 in Algorithm 4.1 coincide with lines 1–3 in Algorithm 2.1, the proof of Theorem 3.8 shows that with probability at least

$$1 - p_1 - p_2 - p_3 - p_4 \geq 1 - \left(\frac{\beta \ell k}{n} + \frac{\beta \ell k}{m} + 2ke^{-\alpha} \right),$$

we have that $A(:, J_1) = XZY(J_1, :)^T$ has full rank k . Line 5 in Algorithm 4.1 will then select a row-index set I_1 such that $A(I_1, J_1)$ is a rank- k matrix. Hence, we have that $A(I_1, :)$ spans the range of the rows of A , so the index sets I_1 and J_1 guarantee that the matrix A is reconstructed exactly when forming both the CUR approximation $A(:, J_1)A(:, J_1)^\dagger AA(I_1, :)^{\dagger} A(I_1, :)$ and the cross approximation $A(:, J_1)A(I_1, J_1)^\dagger A(I_1, :)$. \square

Let us now look at the perturbed case. Since the first half iteration of Algorithm 4.1 coincides with lines 1–3 of Algorithm 2.1, Theorem 3.10 gives a bound on the column subset selection error obtained using the index set J_1 obtained in the first half iteration. The following result allows us to obtain an upper bound for the low-rank approximation error of Algorithm 4.1 after one full iteration.

THEOREM 4.5 ([9, Theorem 1]). *Given a “column space approximator” $B \in \mathbb{R}^{n \times \ell}$ of A that has full column rank, let $C_1 \in \mathbb{R}^{\ell \times \ell}$ be a subset of ℓ linearly independent rows of B and let $C_2 \in \mathbb{R}^{(n-\ell) \times \ell}$ be the set formed by all other rows of B . Let R be the rows of A corresponding to C_1 . Then*

$$\|A - AR^\dagger R\| \leq \sqrt{1 + \|C_2 C_1^\dagger\|_2^2} \cdot \|A - BB^\dagger A\|.$$

THEOREM 4.6. *If Assumption 3.9 is satisfied and we set*

$$\ell_0 = \alpha \mu (k_1 + k_2), \quad \ell_1^{\text{sRRQR, col}} = k_1 + k_2, \quad \ell_1^{\text{new, col}} = \alpha \mu (k_1 + k_3), \quad \ell_1^{\text{sRRQR, row}} = k, \\ \ell_1^{\text{new, row}} = 0, \quad \ell_1^{\text{sRRQR, row}} = \ell_1^{\text{sRRQR, col}} + \ell_1^{\text{new, col}},$$

for some $\alpha > 0$, then the index sets I_1 and J_1 selected after one full iteration of Algorithm 4.1 satisfy

$$\|A - A(:, J_1)A(:, J_1)^\dagger AA(I_1, :)^{\dagger} A(I_1, :)\| \\ \leq \mathcal{O} \left(\varepsilon m \sqrt{nk} \cdot \frac{\sigma_1(XZY^T)}{\sigma_k(XZY^T)} \cdot (\|X^\dagger\| + \|Y^\dagger\|) \right)$$

with a probability of at least

$$1 - k \left(\frac{\beta \ell}{n} + \frac{\beta \ell}{m} + c_1^\alpha + c_2^\alpha \right).$$

Proof. Using (3.10), we can split the error into two terms; the first one is directly bounded by Theorem 3.10, and we can use Theorem 4.5 to bound the second term $\|A - AA(I_1, :)^{\dagger}A(I_1, :)\|$. More specifically, we apply Theorem 4.5 with $B = A(:, J_1)$; note that, thanks to the sRRQR properties, we have

$$\|A(I_1^c, J_1)A(I_1, J_1)^{-1}\| \leq \sqrt{\ell(n-\ell)}\|A(I_1^c, J_1)A(I_1, J_1)^{-1}\|_{\max} \leq \sqrt{\ell(n-\ell)},$$

where we denoted by $I_1^c \in \{1, \dots, n\}^{n-|I_1|}$ the index set corresponding to all the rows not selected by I_1 . It follows that

$$\|A - AA(I_1, :)^{\dagger}A(I_1, :)\| \leq \sqrt{1 + \ell(n-\ell)}\|A - A(:, J_1)A(:, J_1)^{\dagger}A\|$$

and the quantity $\|A - A(:, J_1)A(:, J_1)^{\dagger}A\|$ can be bounded, again, using Theorem 3.10. \square

5. Numerical examples. In this section, we explore numerically the behavior of Algorithms 2.1 and 4.1 on some examples. The numerical experiments have been performed in MATLAB, and the implementation of the sRRQR algorithm from [14] is taken from [28], with parameter $\eta = 1.1$.

Example 5.1. Let us consider the bivariate function

$$f(x, y) = \frac{5 \sin(3x)}{5y - 4} + 2e^{x/2} \cos(10y) + \frac{20y}{4x - 1}.$$

We construct the matrix A as the discretization of $f(x, y)$ on an equispaced grid of size 1000×1000 in the square $[0, 1]^2$, to which we add a random matrix E with norm $\|E\| = 10^{-5}$. The matrix A is of rank approximately 3, and some of the first three singular vectors have high coherence. A visual representation of A and of its first three right and left singular vectors is in Figure 2. We run Algorithm 4.1 with $\ell_0 = 6$ and all other ℓ -parameters set to 3. After one full iteration, we are able to find a good subset of rows and columns. The quality does not improve with more iterations; see Figure 3 (left), where we also compare this with the low-rank approximation error given by the index sets returned by Algorithm 2.1 with $\ell_0 = 6, \ell_a = \ell_b = 3$. While the bounds we obtain in this paper (Theorem 3.10 in the previous section and Theorem 4.6 below)

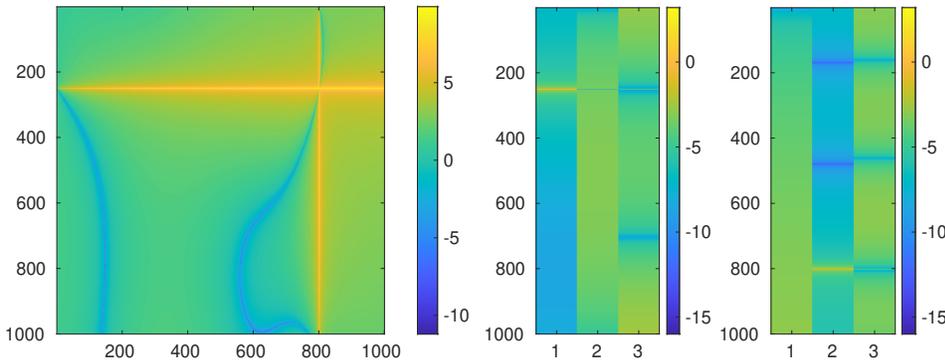


FIG. 2. *Logarithm of the magnitude of the entries of the matrix A from Example 5.1 (left) and of its first three left and right singular vectors (middle and right). The first left singular vector and the second right singular vector are approximately sparse, while all the other ones are incoherent.*

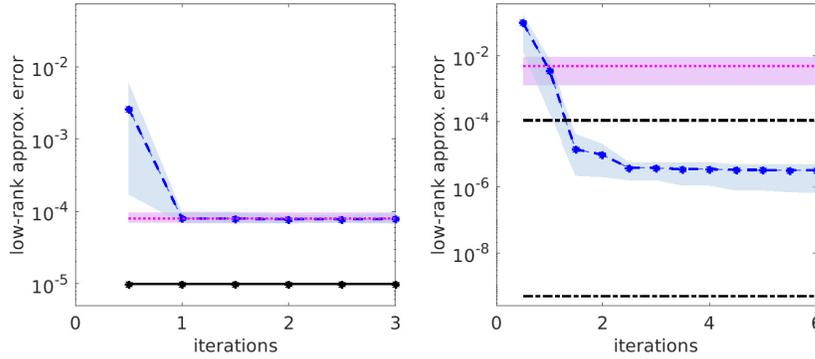


FIG. 3. The blue part of the plot indicates the low-rank approximation error, namely, the quantity $\|A - A(:, J)A(:, J)^\dagger AA(I, :)\|$, where I and J are the index sets returned by Algorithm 4.1 for the matrix from Example 5.1 (left) and the one from Example 5.2 (right). The x -axis denotes the number of iterations of Algorithm 4.1. We run Algorithm 4.1 100 times; the blue dashed line represents the average error, and the light blue areas represent the 90% confidence region. The dotted magenta line is the approximation error given by the index sets returned by Algorithm 2.1 (it is constant because the method is not iterative), and the pink area is the 90% confidence region. Note that, while for Example 5.1 the two algorithms give similar errors after the first step, a few iterations of Algorithm 4.1 help improve the error in the case of Example 5.2. The values of the ℓ -parameters are $\ell_0 = 6, \ell_a = \ell_b = 3$ (left plot) and $\ell_0 = 16, \ell_a = \ell_b = 8$ (right plot) for Algorithm 2.1. For Algorithm 4.1, in the left plot we let $\ell_0 = 6$ and all other ℓ -parameters are set to 3, and in the right plot we let $\ell_0 = 16$ and all other ℓ -parameters are set to 8. For comparison, the black line in the left plot is the error of the best rank-3 approximation of the matrix, and the two black lines in the right plot are the errors of the best rank-8 and rank-16 approximations, respectively. (Color available online.)

involve an “amplification” factor on the error ε that depends on the size of the matrix A , in this example we observe that the error is roughly 10 times larger than the best rank-3 approximation error in the spectral norm, which is $\approx 10^{-5}$.

Example 5.2. While our analysis in section 4.1 is restricted to the first full iteration of Algorithm 4.1, it is worth noting that, in many cases, the low-rank approximation error continues to decrease if more iterations are performed. As an example of this behavior, let us consider the matrix $A \in \mathbb{R}^{1000 \times 1000}$ with entries $a_{ij} = \frac{1}{i+j^2+1}$, and let us run Algorithm 4.1 with $\ell_0 = 16$ and all other ℓ -parameters set to 8. The result is shown in Figure 3 (right), where we also compare this with the low-rank approximation error given by the index sets returned by Algorithm 2.1 with $\ell_0 = 16, \ell_a = \ell_b = 8$. The matrix A has rapidly decaying singular values. The low-rank approximations obtained by Algorithms 2.1 and 4.1 have rank 16. For illustration purposes, in Figure 3, we compare the errors of Algorithm 2.1 and 4.1 with the error of the best rank-8 and rank-16 approximations of A . However, it is not completely fair to compare their error to the best rank-16 approximation of A since our theory results predict the need for an oversampling factor (the ℓ -parameter is chosen to be mildly larger than the target rank).

Example 5.3 (comparison between Algorithm 2.1 and Algorithm 4.1). We construct two matrices of size $n = m = 128$ that satisfy Assumption 3.9 and compare the quality of the low-rank approximation given by Algorithm 2.1 with the one given by one or three full iterations of Algorithm 4.1. Let us define

$$U_1 = \sqrt{\frac{1}{n}} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{bmatrix} \otimes e_{n/8}, \quad U_2 = \sqrt{\frac{1}{n}} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \\ -1 & -1 & -1 & -1 \\ -1 & 1 & -1 & 1 \\ -1 & -1 & 1 & 1 \\ -1 & 1 & 1 & -1 \end{bmatrix} \otimes e_{n/8},$$

where $e_{n/8}$ denotes a vector of all ones of length $n/8$ and \otimes denotes the Kronecker product. The matrix $[U_1 \ U_2]$ has orthonormal columns and has coherence 1, the smallest possible coherence. Let us define $U_3 = \begin{bmatrix} \mathbb{I}_4 \\ 0 \end{bmatrix} \in \mathbb{R}^{n \times 4}$, where \mathbb{I}_4 denotes the 4×4 identity matrix. Then, we let $\tilde{A} = [U_1 \ U_2 \ U_3] [U_1 \ U_3 \ U_2]^T$ and this matrix satisfies Assumptions 3.7 with $k_1 = k_2 = k_3 = 4$. We construct a matrix A by adding a random perturbation of norm $\varepsilon = 10^{-5}$ to \tilde{A} . To construct the second matrix, let us consider $V_1, V_2 \in \mathbb{R}^{n \times 4}$ such that $[V_1 \ V_2] = \sqrt{\frac{8}{n}} \mathbb{I}_8 \otimes e_{n/8}$. Note that the coherence in this case is 8. The matrix $\tilde{B} = [V_1 \ V_2 \ U_3] [V_1 \ U_3 \ V_2]^T$ satisfies Assumption 3.7 and we take B to be an additive ε -perturbation of \tilde{B} . Note that A and B satisfy Assumption 3.9, and they are approximately of rank 12. In Figure 4, we compare the low-rank approximation error for different values of the parameter ℓ between 12 and 56, of

- Algorithm 2.1;
- one full iteration of Algorithm 4.1;
- three full iterations of Algorithm 4.1.

In both algorithms, for a given value of ℓ we set $\ell_0 = \ell$ and all other parameters are set to $\ell/2$. Some remarks are in order.

- Both algorithms need fewer samples to recover a good low-rank approximation for the matrix A , which has lower coherence, as is to be expected.
- The bounds from Theorem 3.10 and Theorem 4.6 seem to be quite pessimistic, at least for the given examples. In particular, for larger values of ℓ

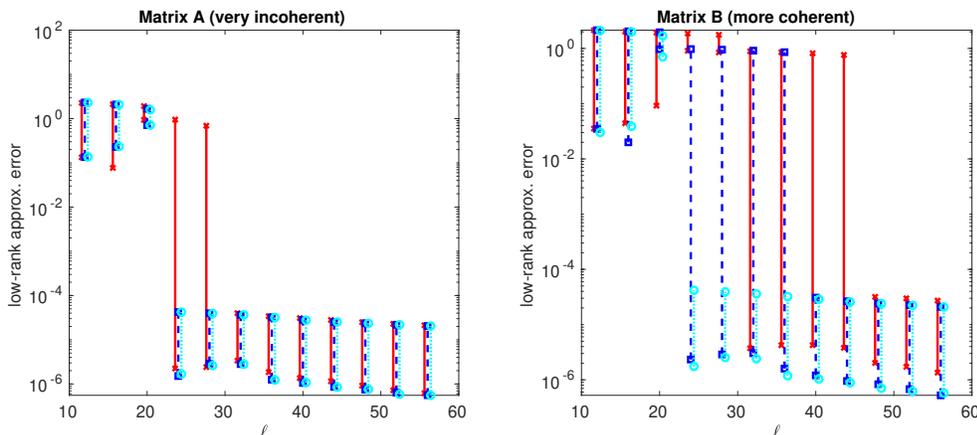


FIG. 4. Comparison of Algorithm 2.1 (red continuous line), one full iteration of Algorithm 4.1 (blue dashed line), and three full iterations of Algorithm 4.1 (cyan dotted line) for the matrices from Example 5.3. Each vertical line corresponds to one algorithm and one parameter ℓ . For each algorithm and each value of $\ell = 12, 16, \dots, 56$, we have performed 1,000 independent runs of each algorithm; the bottom of each line represents the 10th percentile, and the top of each line represents the 90th percentile of the obtained low-rank approximation errors. (Color available online.)

(e.g., $\ell \geq 48$) the error of the low-rank approximation is comparable to $\varepsilon = 10^{-5}$ with high probability.

- Although Theorem 4.6 features an extra factor $\sqrt{n\ell}$ compared to Theorem 3.10, in this example one full iteration of Algorithm 4.1 is always at least as good as Algorithm 2.1.
- For the matrix B , we notice an improvement when performing more iterations of Algorithm 4.1; this may be due to the fact that we are sampling more new rows/columns at each iteration.

6. Discussion. In this work, we have shown that, with high probability, a uniform sampling strategy combined with sRRQR (Algorithm 2.1) selects rows and columns that give a good CUR approximation for matrices satisfying Assumption 3.9. Our error bounds are informative when the perturbation E to the low-rank matrix XZY^T is small enough that $\|E\| \ll m\sqrt{k/\ell}$ and the ratio $\sigma_1(XZY^T)/\sigma_k(XZY^T)$ is not too large. Some dependence of the error on the size of the matrix is inevitable in the spectral norm when working with row and column-based low-rank approximations; see (1.2) and the results in [4]. While Algorithm 2.1 is reminiscent of [4, Algorithm 3], a key difference is that Algorithm 2.1 requires some additional random rows and columns after the sRRQR step (lines 3 and 6), which are necessary to obtain good low-rank approximation for a larger class of matrices than the one considered in [4], that is, matrices that admit a decomposition into factors with suitable incoherence and/or sparsity assumptions up to a small additive error.

While it is difficult, in general, to check whether a matrix A admits a decomposition satisfying Assumption 3.9, the objective of this paper is to shed some light on the excellent practical performance of simple sublinear-time algorithms for column and row subset selection. Note that, in Assumption 3.9, it is easier to think of XZY^T as the singular value decomposition of A or its best rank- k approximation, but in fact, we do not require X and Y to have orthonormal columns, as long as they are well-conditioned. This flexibility allows us to apply our bounds to a larger class of matrices.

In section 4, we introduced Algorithm 4.1, directly inspired by [17, Algorithm 2.2], and we provided an error bound for the CUR approximation constructed with the indices returned after one iteration. Note that, with respect to Theorem 3.10 on the column subset selection error, the constant in the bound of Theorem 4.6 has been multiplied approximately by $\sqrt{n\ell}$; this increase in the error is a result of Theorem 4.5. However, in all the examples that we have considered, the row and column indices returned by Algorithm 4.1 were at least as good as the ones returned by Algorithm 2.1 with a corresponding choice of the parameters. Since iterative refinement has been shown to work for similar settings [13, 27], an interesting direction for future research is the analysis of Algorithm 4.1 beyond the first iteration.

Our results do not cover all the matrices for which *there is hope*, in the sense of section 2.3. For example, a scenario that is not covered by the current theory and is left for future work consists of matrices that have some incoherence in the blocks X_1, X_2, Y_1, Y_3 but do not satisfy any sparseness assumption.

Acknowledgment. We thank the referees for providing helpful feedback.

REFERENCES

- [1] M. BEBENDORF, *Approximation of boundary element matrices*, Numer. Math., 86 (2000), pp. 565–589, <https://doi.org/10.1007/PL00005410>.
- [2] H. CAI, K. HAMM, L. HUANG, AND D. NEEDELL, *Robust CUR decomposition: Theory and imaging applications*, SIAM J. Imaging Sci., 14 (2021), pp. 1472–1503, <https://doi.org/10.1137/20M1388322>.

- [3] H. CAI, L. HUANG, P. LI, AND D. NEEDELL, *Matrix completion with cross-concentrated sampling: Bridging uniform sampling and CUR sampling*, IEEE Trans. Pattern Anal., 45 (2023), pp. 10100–10113, <https://doi.org/10.1109/TPAMI.2023.3261185>.
- [4] J. CHIU AND L. DEMANET, *Sublinear randomized algorithms for skeleton decompositions*, SIAM J. Matrix Anal. Appl., 34 (2013), pp. 1361–1383, <https://doi.org/10.1137/110852310>.
- [5] M. B. COHEN, C. MUSCO, AND C. MUSCO, *Input sparsity time low-rank approximation via ridge leverage score sampling*, in Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, 2017, pp. 1758–1777, <https://doi.org/10.1137/1.9781611974782.115>.
- [6] A. CORTINOVIS AND D. KRESSNER, *Low-rank approximation in the Frobenius norm by column and row subset selection*, SIAM J. Matrix Anal. Appl., 41 (2020), pp. 1651–1673, <https://doi.org/10.1137/19M1281848>.
- [7] A. CORTINOVIS, D. KRESSNER, AND S. MASSEI, *On maximum volume submatrices and cross approximation for symmetric semidefinite and diagonally dominant matrices*, Linear Algebra Appl., 593 (2020), pp. 251–268, <https://doi.org/10.1016/j.laa.2020.02.010>.
- [8] A. DESHPANDE, L. RADEMACHER, S. VEMPALA, AND G. WANG, *Matrix approximation and projective clustering via volume sampling*, Theory Comput., 2 (2006), pp. 225–247, <https://doi.org/10.4086/toc.2006.v002a012>.
- [9] Y. DONG AND P.-G. MARTINSSON, *Simpler is better: A comparative study of randomized pivoting algorithms for CUR and interpolative decompositions*, Adv. Comput. Math., 49 (2023), 66, <https://doi.org/10.1007/s10444-023-10061-z>.
- [10] P. DRINEAS, M. W. MAHONEY, AND S. MUTHUKRISHNAN, *Relative-error CUR matrix decompositions*, SIAM J. Matrix Anal. Appl., 30 (2008), pp. 844–881, <https://doi.org/10.1137/07070471X>.
- [11] C. ECKART AND G. YOUNG, *The approximation of one matrix by another of lower rank*, Psychometrika, 1 (1936), pp. 211–218, <https://doi.org/10.1007/BF02288367>.
- [12] A. FRIEZE, R. KANNAN, AND S. VEMPALA, *Fast Monte-Carlo algorithms for finding low-rank approximations*, J. ACM, 51 (2004), pp. 1025–1041, <https://doi.org/10.1145/1039488.1039494>.
- [13] S. A. GOREINOV, I. V. OSELEDETS, D. V. SAVOSTYANOV, E. E. TYRTYSHNIKOV, AND N. L. ZAMARASHKIN, *How to find a good submatrix*, in Matrix Methods: Theory, Algorithms and Applications, World Scientific, Hackensack, NJ, 2010, pp. 247–256, <https://doi.org/10.1142/9789812836021>.
- [14] M. GU AND S. C. EISENSTAT, *Efficient algorithms for computing a strong rank-revealing QR factorization*, SIAM J. Sci. Comput., 17 (1996), pp. 848–869, <https://doi.org/10.1137/0917055>.
- [15] K. HAMM, M. MESKINI, AND H. CAI, *Riemannian CUR decompositions for robust principal component analysis*, in Topological, Algebraic and Geometric Learning Workshops 2022, PMLR, 2022, pp. 152–160.
- [16] R. A. HORN AND C. R. JOHNSON, *Matrix Analysis*, 2nd ed., Cambridge University Press, Cambridge, 2013.
- [17] Y. LI, H. YANG, E. R. MARTIN, K. L. HO, AND L. YING, *Butterfly factorization*, Multiscale Model. Simul., 13 (2015), pp. 714–732, <https://doi.org/10.1137/15M1007173>.
- [18] T.-T. LU AND S.-H. SHIOU, *Inverses of 2×2 block matrices*, Comput. Math. Appl., 43 (2002), pp. 119–129, [https://doi.org/10.1016/S0898-1221\(01\)00278-4](https://doi.org/10.1016/S0898-1221(01)00278-4).
- [19] L. MIRSKY, *Symmetric gauge functions and unitarily invariant norms*, Quart. J. Math. Oxford Ser. (2), 11 (1960), pp. 50–59, <https://doi.org/10.1093/qmath/11.1.50>.
- [20] A. K. SAIBABA, *HOID: Higher order interpolatory decomposition for tensors based on Tucker representation*, SIAM J. Matrix Anal. Appl., 37 (2016), pp. 1223–1249, <https://doi.org/10.1137/15M1048628>.
- [21] D. C. SORENSEN AND M. EMBREE, *A DEIM induced CUR factorization*, SIAM J. Sci. Comput., 38 (2016), pp. A1454–A1482, <https://doi.org/10.1137/140978430>.
- [22] G. W. STEWART, *Four algorithms for the efficient computation of truncated pivoted QR approximations to a sparse matrix*, Numer. Math., 83 (1999), pp. 313–323, <https://doi.org/10.1007/s002110050451>.
- [23] A. TALWALKAR AND A. ROSTAMIZADEH, *Matrix coherence and the Nyström method*, in Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence, 2010, pp. 572–579.
- [24] J. A. TROPP, *Improved analysis of the subsampled randomized Hadamard transform*, Adv. Adapt. Data Anal., 3 (2011), pp. 115–126, <https://doi.org/10.1142/S1793536911000787>.
- [25] M. UDELL AND A. TOWNSEND, *Why are big data matrices approximately low rank?*, SIAM J. Math. Data Sci., 1 (2019), pp. 144–160, <https://doi.org/10.1137/18M1183480>.

- [26] S. VORONIN AND P.-G. MARTINSSON, *Efficient algorithms for CUR and interpolative matrix decompositions*, Adv. Comput. Math., 43 (2017), pp. 495–516, <https://doi.org/10.1007/s10444-016-9494-8>.
- [27] J. XIA, *Making the Nyström method highly accurate for low-rank approximations*, SIAM J. Sci. Comput., 46 (2024), pp. A1076–A1101, <https://doi.org/10.1137/23M1585039>.
- [28] X. XING, *Strong Rank Revealing QR Decomposition*, MATLAB Central File Exchange, <https://www.mathworks.com/matlabcentral/fileexchange/69139-strong-rank-revealing-qr-decomposition>, 2024. Retrieved February 13, 2024.
- [29] G. ZHANG, H. LI, AND Y. WEI, *CPQR-based randomized algorithms for generalized CUR decompositions*, Comput. Appl. Math., 43 (2024), 132, 28, <https://doi.org/10.1007/s40314-024-02642-5>.