



ELSEVIER

Contents lists available at ScienceDirect

Applied and Computational Harmonic Analysis

www.elsevier.com/locate/acha

A randomized algorithm for the decomposition of matrices [☆]Per-Gunnar Martinsson ^a, Vladimir Rokhlin ^b, Mark Tygert ^{c,*}^a Department of Applied Mathematics, University of Colorado at Boulder, Boulder, CO 80309, United States^b Departments of Computer Science, Mathematics, and Physics, Yale University, New Haven, CT 06511, United States^c Courant Institute of Mathematical Sciences, NYU, New York, NY 10012, United States

ARTICLE INFO

Article history:

Received 9 April 2008

Revised 23 August 2009

Accepted 13 February 2010

Available online 19 February 2010

Communicated by Charles K. Chui

Keywords:

Randomized

Algorithm

Low rank

Matrix

SVD

Lanczos

ABSTRACT

Given an $m \times n$ matrix A and a positive integer k , we describe a randomized procedure for the approximation of A with a matrix Z of rank k . The procedure relies on applying A^T to a collection of l random vectors, where l is an integer equal to or slightly greater than k ; the scheme is efficient whenever A and A^T can be applied rapidly to arbitrary vectors. The discrepancy between A and Z is of the same order as \sqrt{lm} times the $(k+1)$ st greatest singular value σ_{k+1} of A , with negligible probability of even moderately large deviations. The actual estimates derived in the paper are fairly complicated, but are simpler when $l-k$ is a fixed small nonnegative integer. For example, according to one of our estimates for $l-k=20$, the probability that the spectral norm $\|A-Z\|$ is greater than $10\sqrt{(k+20)m}\sigma_{k+1}$ is less than 10^{-17} . The paper contains a number of estimates for $\|A-Z\|$, including several that are stronger (but more detailed) than the preceding example; some of the estimates are effectively independent of m . Thus, given a matrix A of limited numerical rank, such that both A and A^T can be applied rapidly to arbitrary vectors, the scheme provides a simple, efficient means for constructing an accurate approximation to a singular value decomposition of A . Furthermore, the algorithm presented here operates reliably independently of the structure of the matrix A . The results are illustrated via several numerical examples.

© 2010 Elsevier Inc. All rights reserved.

1. Introduction

In many practical circumstances, it is desirable to approximate a matrix A with a sum of rank-1 matrices. Such an approximation of A often facilitates understanding of the properties of A . Moreover, if the approximation involves only a small number of rank-1 matrices, then the approximation also facilitates rapid calculations involving A .

There are at least two classical forms of such matrix approximations. One is an approximation to a singular value decomposition (SVD), which is known in the statistical literature as a principal component analysis (PCA). The other is an approximation obtained via subset selection; we will refer to the matrix representation obtained via subset selection as an interpolative decomposition. These two types of matrix approximations are defined as follows.

An approximation to an SVD of a real $m \times n$ matrix A consists of nonnegative real numbers $\sigma_1, \sigma_2, \dots, \sigma_{k-1}, \sigma_k$ known as singular values, orthonormal real $m \times 1$ column vectors $u^1, u^2, \dots, u^{k-1}, u^k$ known as left singular vectors, and orthonormal real $n \times 1$ column vectors $v^1, v^2, \dots, v^{k-1}, v^k$ known as right singular vectors, such that

[☆] Partially supported by the U.S. DoD under Air Force STTR Grant F49620-03-C-0031 and DARPA Grant HR0011-05-1-0002.

* Corresponding author.

E-mail addresses: per-gunnar.martinsson@colorado.edu (P.-G. Martinsson), tygert@aya.yale.edu (M. Tygert).

$$\left\| A - \sum_{j=1}^k u^j \sigma_j (v^j)^T \right\| \leq \delta, \quad (1)$$

where k , m , and n are positive integers with $k < m$ and $k < n$, δ is a positive real number specifying the precision of the approximation, and, for any matrix B , $\|B\|$ denotes the spectral (l^2 -operator) norm of B , that is, $\|B\|$ is the greatest singular value of B . An approximation to an SVD of A is often written in the equivalent form

$$\|A - U \Sigma V^T\| \leq \delta, \quad (2)$$

where U is a real $m \times k$ matrix whose columns are orthonormal, V is a real $n \times k$ matrix whose columns are orthonormal, and Σ is a real $k \times k$ matrix whose entries are all nonnegative and whose entries off of the main diagonal are zero. See, for example, [15] for a discussion of SVDs.

An interpolative decomposition of a real $m \times n$ matrix A consists of a real $m \times k$ matrix B whose columns constitute a subset of the columns of A , and a real $k \times n$ matrix P , such that

1. some subset of the columns of P makes up the $k \times k$ identity matrix,
2. no entry of P has an absolute value greater than 2, and
3. $A = BP$.

See, for example, [5,8,12,18], or Sections 4 and 5 of [3] for a discussion of interpolative decompositions.

Given an algorithm permitting the fast application of a numerically low-rank matrix A , and an algorithm permitting the fast application of A^T , the algorithm of the present paper provides a simple, efficient way for computing an accurate approximation to an SVD of A . Moreover, the algorithm provides a similar method for computing an accurate approximation to an interpolative decomposition of A under the same conditions.

Our scheme also provides an efficient, robust means for approximating the k greatest singular values and corresponding singular vectors of any matrix A for which a representation enabling the fast application of both A and A^T is available. The precision δ of the resulting approximation given by formula (2) is at most a reasonably small multiple of the $(k+1)$ st greatest singular value of A .

The algorithm of the present paper is a randomized one, and fails with a rather negligible probability. Examples of the probabilities involved can be found in (111) and (124) in Section 5 below.

Some potential applications of the algorithm include finding the eigenmodes of certain networks, mining digital documents for information via latent semantic analysis, simplifying the implementation of algorithms for fast matrix inversion that are based on the compression of blocks within matrices, and improving condition number estimation and subspace determination algorithms that are based on inverse iteration or other iterative methods.

Accelerating the approximation of matrices which can be applied rapidly has been a popular topic of recent papers. [17] and [1] design deterministic algorithms for the low-rank approximation of sparse matrices. Referring the reader to [16] for a detailed survey of the extensive literature on randomized algorithms, we observe that several recent articles introduce algorithms which, given any positive integer k , produce an approximation Z to the matrix A such that

$$\|A - Z\|_F \leq (1 + \varepsilon) \|A - Z_{\text{optimal}}\|_F, \quad (3)$$

where $\varepsilon \leq 1$ is an arbitrarily small positive real number, $\|A - Z\|_F$ is the Frobenius norm of $A - Z$, and Z_{optimal} is the rank- k matrix for which $\|A - Z_{\text{optimal}}\|_F$ is minimal, where $\|A - Z_{\text{optimal}}\|_F$ is the Frobenius norm of $A - Z_{\text{optimal}}$. [16] reports that the most efficient algorithms with proven bounds of the type (3) have worst-case costs of

$$C_{Dk/\varepsilon} \log_2 \left(\frac{1}{p} \right) \quad (4)$$

floating-point operations, where D is an (unspecified, but probably computable) universal constant, k is the rank of the approximation Z_{optimal} to A , ε is the positive real number in (3), p is the probability that the algorithm fails, and

$$C_{Dk/\varepsilon} = k \cdot C_A + (Dk/\varepsilon) \cdot C_{A^T} + E \cdot (k^2 m + (Dk/\varepsilon)^2 n), \quad (5)$$

where C_A is the cost of applying the real $m \times n$ matrix A to a real $n \times 1$ column vector, C_{A^T} is the cost of applying A^T to a real $m \times 1$ column vector, and E is a reasonably small positive real number.

The present paper describes an algorithm which, given any positive integer k , costs

$$C_l = k \cdot C_A + l \cdot C_{A^T} + E \cdot (k^2 m + l^2 n) \quad (6)$$

floating-point operations, where again C_A is the cost of applying the real $m \times n$ matrix A to a real $n \times 1$ column vector, C_{A^T} is the cost of applying A^T to a real $m \times 1$ column vector, E is a reasonably small positive real number, and l is a user-specified integer greater than or equal to k (for example, $l = k + 8$ or $l = k + 20$). The algorithm produces a rank- k approximation Z to the $m \times n$ matrix A such that

$$\|A - Z\| \leq 10\sqrt{lm}\|A - Z_{\text{optimal}}\| \tag{7}$$

with probability not less than $1 - q(l - k)$ (independent of the structure of the matrix A), where $\|A - Z\|$ is the spectral norm of $A - Z$, Z_{optimal} is the rank- k matrix for which the spectral norm $\|A - Z_{\text{optimal}}\|$ is minimal, and q is a complicated but rapidly decaying function – for example, $q(8) < 10^{-5}$ and $q(20) < 10^{-17}$. The present article also provides several bounds other than (7), some of which are effectively independent of m .

Thus, we focus on minimizing costs, while still attaining accuracy of roughly the order of that attained by the best possible rank- k approximation. It should also be pointed out that the algorithm of the present article is very similar to the algorithms of [14] and [16], though our analysis is quite different.

We do not analyze in detail the effects of round-off upon the algorithm of the present paper. However, most of the bounds that we discuss have finite-precision analogues. This is confirmed by both our preliminary analysis and our numerical experiments (some of which are described in Section 6 below). For simplicity, we discuss only real matrices; the analysis below extends easily to the complex case. The technical report version of the present article is [13].

The present paper has the following structure: Section 2 introduces the algorithm from a heuristic point of view. Section 3 collects together various known facts which later sections utilize. Section 4 provides the principal lemmas which Section 5 uses to construct algorithms. Section 5 describes the algorithm of the present paper, providing details about its accuracy and computational costs. Section 6 illustrates the algorithm via several numerical examples. Section 7 presents several conclusions about the algorithm.

2. Informal description of the algorithm

In this section, we provide a heuristic description of the algorithm for computing an approximation to an SVD. Section 5 provides a detailed description of the algorithm described intuitively in the present section, utilizing the tools described in Sections 3 and 4.

Suppose that k , m , and n are positive integers with $k < m$ and $k < n$, and A is a real $m \times n$ matrix. We will construct an approximation to an SVD of A such that

$$\|U\Sigma V^T - A\| \lesssim \sqrt{lm}\sigma_{k+1}, \tag{8}$$

where U is a real $m \times k$ matrix whose columns are orthonormal, V is a real $n \times k$ matrix whose columns are orthonormal, Σ is a diagonal real $k \times k$ matrix whose entries are all nonnegative, $\|U\Sigma V^T - A\|$ is the spectral (l^2 -operator) norm of $U\Sigma V^T - A$, σ_{k+1} is the $(k + 1)$ st greatest singular value of A , and l is a user-specified integer with $l > k$, such that $l < m$ and $l < n$ (for example, $l = k + 20$). To do so, we identify an orthonormal basis of most of the range of A^T , via the following two steps:

1. Using a random number generator, form a real $m \times l$ matrix G whose entries are i.i.d. Gaussian random variables of zero mean and unit variance, and compute the $n \times l$ product matrix

$$R = A^T G. \tag{9}$$

2. Using an SVD, form a real $n \times k$ matrix Q whose columns are orthonormal, such that there exists a real $k \times l$ matrix S for which

$$\|QS - R\| \leq \rho_{k+1}, \tag{10}$$

where ρ_{k+1} is the $(k + 1)$ st greatest singular value of R . (See Observation 3.6 for details concerning the construction of such a matrix Q .)

Intuitively, the columns of Q in (10) constitute an orthonormal basis of most of the range of A^T .

Having identified a good approximation to the range of A^T , we perform some simple linear algebraic manipulations in order to obtain a good approximation to an SVD of A , via the following three steps:

3. Compute the $m \times k$ product matrix

$$T = A Q. \tag{11}$$

4. Form an SVD of T ,

$$T = U \Sigma W^T, \tag{12}$$

where U is a real $m \times k$ matrix whose columns are orthonormal, W is a real $k \times k$ matrix whose columns are orthonormal, and Σ is a real $k \times k$ matrix whose entries are all nonnegative and zero off of the main diagonal. (See, for example, Chapter 8 in [7] for details concerning the construction of such an SVD.)

5. Compute the $n \times k$ product matrix

$$V = QW. \quad (13)$$

We will prove in Section 5 that the matrices U , Σ , and V obtained via Steps 1–5 above satisfy (8). We encourage the reader to begin with Sections 5 and 6, referring back to the relevant portions of Sections 3 and 4 as they are referenced.

3. Preliminaries from linear algebra and the theory of probability

In this section, we summarize various facts about matrices. Section 3.1 discusses the approximation of arbitrary matrices. Section 3.2 discusses the singular values of arbitrary matrices. Section 3.3 discusses the singular values of certain random matrices.

In the present section and throughout the rest of the paper, we employ the following notation. In accordance with the standard practice, we will denote the base of the natural logarithm by e . We will denote an identity matrix by $\mathbf{1}$, and a matrix whose entries are all zero by $\mathbf{0}$. For any matrix A , we define the norm $\|A\|$ of A to be the spectral (l^2 -operator) norm of A , that is, $\|A\|$ is the greatest singular value of A . For any positive integer n , and real $n \times 1$ column vector $v \in \mathbb{R}^n$, we define the norm $\|v\|$ of v to be the root-sum-square (l^2 norm) of the entries of v , that is,

$$\|v\| = \sqrt{\sum_{k=1}^n (v_k)^2}, \quad (14)$$

where v_k is the k th entry of v . (Of course, the norm of v as viewed as a real $n \times 1$ matrix is equal to the norm of v as viewed as a real $n \times 1$ column vector.)

3.1. Approximation of general matrices

The following lemma states that, for any $m \times n$ matrix A whose rank is k , where k , m , and n are positive integers, there exist an $m \times k$ matrix B whose columns constitute a subset of the columns of A , and a $k \times n$ matrix P , such that

1. some subset of the columns of P makes up the $k \times k$ identity matrix,
2. P is not too large, and
3. $BP = A$.

Moreover, the lemma provides an analogous approximation BP to A when the exact rank of A is not k , but the $(k+1)$ st singular value of A is nevertheless small. The lemma is a reformulation of Theorem 3.2 in [12] and Theorem 3 in [5].

Lemma 3.1. *Suppose that m and n are positive integers, and A is a real $m \times n$ matrix.*

Then, for any positive integer k with $k \leq m$ and $k \leq n$, there exist a real $k \times n$ matrix P , and a real $m \times k$ matrix B whose columns constitute a subset of the columns of A , such that

1. some subset of the columns of P makes up the $k \times k$ identity matrix,
2. no entry of P has an absolute value greater than 1,
3. $\|P\| \leq \sqrt{k(n-k)+1}$,
4. the least (that is, the k th greatest) singular value of P is at least 1,
5. $BP = A$ when $k = m$ or $k = n$, and
6. $\|BP - A\| \leq \sqrt{k(n-k)+1}\sigma_{k+1}$ when $k < m$ and $k < n$, where σ_{k+1} is the $(k+1)$ st greatest singular value of A .

Remark 3.2. Properties 1, 2, 3, and 4 in Lemma 3.1 ensure that the interpolative decomposition BP of A is numerically stable. Also, Property 3 follows directly from Properties 1 and 2, and Property 4 follows directly from Property 1.

Observation 3.3. Existing algorithms for computing the matrices B and P in Lemma 3.1 are computationally expensive. We use an algorithm to produce B and P which satisfy somewhat weaker conditions than those in Lemma 3.1. We compute B and P such that

1. some subset of the columns of P makes up the $k \times k$ identity matrix,
2. no entry of P has an absolute value greater than 2,
3. $\|P\| \leq \sqrt{4k(n-k)+1}$,
4. the least (that is, the k th greatest) singular value of P is at least 1,
5. $BP = A$ when $k = m$ or $k = n$, and
6. $\|BP - A\| \leq \sqrt{4k(n-k)+1}\sigma_{k+1}$ when $k < m$ and $k < n$, where σ_{k+1} is the $(k+1)$ st greatest singular value of A .

For any positive real number ε , the algorithm can identify the least k such that $\|BP - A\| \approx \varepsilon$. Furthermore, there exists a real number C such that the algorithm computes both B and P using at most $Ckmn \log(n)$ floating-point operations. The algorithm is based upon the Cramer rule and the ability to obtain the minimal-norm (or at least roughly minimal-norm) solutions to linear algebraic systems of equations (see [5,11,12]).

Remark 3.4. For further discussion of interpolative decompositions, see, for example, [3,8–11,18].

The following classical lemma provides an approximation QS to an $n \times l$ matrix R via an $n \times k$ matrix Q whose columns are orthonormal, and a $k \times l$ matrix S . As remarked in Observation 3.6, the proof of this lemma provides a classic algorithm for computing Q and S , given R . We include the proof since we will be using this algorithm.

Lemma 3.5. *Suppose that k, l , and n are positive integers with $k < l$ and $l \leq n$, and R is a real $n \times l$ matrix.*

Then, there exist a real $n \times k$ matrix Q whose columns are orthonormal, and a real $k \times l$ matrix S , such that

$$\|QS - R\| \leq \rho_{k+1}, \tag{15}$$

where ρ_{k+1} is the $(k + 1)$ st greatest singular value of R .

Proof. We start by forming an SVD of R ,

$$R = U \Sigma V^T, \tag{16}$$

where U is a real $n \times l$ matrix whose columns are orthonormal, V is a real $l \times l$ matrix whose columns are orthonormal, and Σ is a real $l \times l$ matrix whose entries are nonnegative everywhere and zero off of the main diagonal, such that

$$\Sigma_{j,j} = \rho_j \tag{17}$$

for all $j = 1, 2, \dots, l - 1, l$, where $\Sigma_{j,j}$ is the entry in row j and column j of Σ , and ρ_j is the j th greatest singular value of R . We define Q to be the leftmost $n \times k$ block of U , and P to be the rightmost $n \times (l - k)$ block of U , so that

$$U = \left(Q \begin{array}{c} \vdots \\ P \end{array} \right). \tag{18}$$

We define S to be the uppermost $k \times l$ block of ΣV^T , and T to be the lowermost $(l - k) \times l$ block of ΣV^T , so that

$$\Sigma V^T = \begin{pmatrix} S \\ T \end{pmatrix}. \tag{19}$$

Combining (16), (17), (18), (19), and the fact that the columns of U are orthonormal, as are the columns of V , yields (15). \square

Observation 3.6. In order to compute the matrices Q and S in (15) from the matrix R , we can construct (16), and then form Q and S according to (18) and (19). (See, for example, Chapter 8 in [7] for details concerning the computation of the SVD.)

3.2. Singular values of general matrices

The following trivial technical lemma will be needed in Section 4.

Lemma 3.7. *Suppose that m and n are positive integers with $m \geq n$. Suppose further that A is a real $m \times n$ matrix such that $A^T A$ is invertible.*

Then,

$$\|(A^T A)^{-1} A^T\| = \frac{1}{\sigma_n}, \tag{20}$$

where σ_n is the least (that is, the n th greatest) singular value of A .

The following lemma provides what is known as the Courant–Fischer maximin characterization of singular values; Theorem 8.1.2 in [7] provides an equivalent formulation of (21).

Lemma 3.8. *Suppose that m and n are positive integers, and A is a real $m \times n$ matrix.*

Then, the k th greatest singular value σ_k of A is given by the formula

$$\sigma_k = \max_{S \subseteq \mathbb{R}^n: \dim S=k} \min_{v \in S: \|v\| \neq 0} \frac{\|Av\|}{\|v\|} \quad (21)$$

for all $k = 1, 2, \dots, \min(m, n) - 1, \min(m, n)$, where the maximum is taken over all k -dimensional subspaces of \mathbb{R}^n , and the minimum is taken over all vectors in S that have nonzero norms.

The following lemma states that the singular values of the product GA of matrices G and A are at most $\|G\|$ times greater than the corresponding singular values of A .

Lemma 3.9. *Suppose that l, m , and n are positive integers, A is a real $m \times n$ matrix, and G is a real $l \times m$ matrix.*

Then, the k th greatest singular value ρ_k of the product GA is at most a factor of $\|G\|$ times the k th greatest singular value σ_k of A , that is,

$$\rho_k \leq \|G\| \sigma_k \quad (22)$$

for all $k = 1, 2, \dots, \min(l, m, n) - 1, \min(l, m, n)$.

Proof. For any vector $v \in \mathbb{R}^n$ with $\|v\| \neq 0$,

$$\frac{\|GA v\|}{\|v\|} \leq \|G\| \frac{\|A v\|}{\|v\|}. \quad (23)$$

Combining (21) and (23) yields (22). \square

The following lemma states that the greatest singular value of a matrix A is at least as large as the greatest singular value of any rectangular block of entries in A ; the lemma is a straightforward consequence of the minimax properties of singular values (see, for example, Section 47 of Chapter 2 in [19]).

Lemma 3.10. *Suppose that k, l, m , and n are positive integers with $k \leq m$ and $l \leq n$. Suppose further that A is a real $m \times n$ matrix, and B is a $k \times l$ rectangular block of entries in A .*

Then, the greatest singular value of B is at most the greatest singular value of A .

The following lemma states that the singular values of an $(n - 1) \times n$ block of rows of an $n \times n$ matrix A interlace the singular values of A ; Corollary 8.6.3 in [7] provides an equivalent formulation of (24).

Lemma 3.11. *Suppose that n is a positive integer with $n > 1$, A is a real $n \times n$ matrix, and R is an $(n - 1) \times n$ rectangular block of entries in A .*

Then, the singular values $\sigma_1, \sigma_2, \dots, \sigma_{n-1}, \sigma_n$ of A and the singular values $\rho_1, \rho_2, \dots, \rho_{n-2}, \rho_{n-1}$ of R satisfy the inequalities

$$\sigma_1 \geq \rho_1 \geq \sigma_2 \geq \rho_2 \geq \dots \geq \sigma_{n-2} \geq \rho_{n-2} \geq \sigma_{n-1} \geq \rho_{n-1} \geq \sigma_n. \quad (24)$$

The following lemma states that if the norm of the difference of two matrices is small, then their corresponding singular values are close; Corollary 8.6.2 in [7] provides an equivalent formulation of (25).

Lemma 3.12. *Suppose that m and n are positive integers, and A and R are real $m \times n$ matrices.*

Then, the k th greatest singular value τ_k of the sum $A + R$ and the k th greatest singular value σ_k of A differ by at most $\|R\|$, that is,

$$|\tau_k - \sigma_k| \leq \|R\| \quad (25)$$

for all $k = 1, 2, \dots, \min(m, n) - 1, \min(m, n)$.

3.3. Singular values of random matrices

The following lemma provides a highly probable upper bound on the greatest singular value of a square matrix whose entries are independent, identically distributed (i.i.d.) Gaussian random variables of zero mean and unit variance; Formula 8.8 in [6] provides an equivalent formulation of the lemma.

Lemma 3.13. *Suppose that n is a positive integer, G is a real $n \times n$ matrix whose entries are i.i.d. Gaussian random variables of zero mean and unit variance, and γ is a positive real number, such that $\gamma > 1$ and*

$$1 - \frac{1}{4(\gamma^2 - 1)\sqrt{\pi n\gamma^2}} \left(\frac{2\gamma^2}{e\gamma^2 - 1} \right)^n \tag{26}$$

is nonnegative.

Then, the greatest singular value of G is at most $\sqrt{2n}\gamma$ with probability not less than the amount in (26).

Combining Lemmas 3.10 and 3.13 yields the following lemma, providing a highly probable upper bound on the greatest singular value of a rectangular matrix whose entries are i.i.d. Gaussian random variables of zero mean and unit variance.

Lemma 3.14. *Suppose that $l, m,$ and n are positive integers with $n \geq l$ and $n \geq m$. Suppose further that G is a real $l \times m$ matrix whose entries are i.i.d. Gaussian random variables of zero mean and unit variance, and γ is a positive real number, such that $\gamma > 1$ and (26) is nonnegative.*

Then, the greatest singular value of G is at most $\sqrt{2n}\gamma$ with probability not less than the amount in (26).

The following lemma provides a highly probable lower bound on the least singular value of a rectangular matrix whose entries are i.i.d. Gaussian random variables of zero mean and unit variance; Formula 2.5 in [4] and the proof of Lemma 4.1 in [4] together provide an equivalent formulation of Lemma 3.15.

Lemma 3.15. *Suppose that k and l are positive integers with $k \leq l$. Suppose further that G is a real $l \times k$ matrix whose entries are i.i.d. Gaussian random variables of zero mean and unit variance, and β is a positive real number, such that*

$$1 - \frac{1}{\sqrt{2\pi(l-k+1)}} \left(\frac{e}{(l-k+1)\beta} \right)^{l-k+1} \tag{27}$$

is nonnegative.

Then, the least (that is, the k th greatest) singular value of G is at least $1/(\sqrt{l}\beta)$ with probability not less than the amount in (27).

4. Mathematical apparatus

In this section, we describe the principal tools used in Section 5.

The following lemma states that the product BP of matrices B and P is a good approximation to a matrix A , provided that there exists a matrix G such that

1. the columns of B constitute a subset of the columns of A ,
2. $\|P\|$ is not too large,
3. GBP is a good approximation to GA , and
4. there exists a matrix F such that $\|F\|$ is not too large, and FGA is a good approximation to A .

Lemma 4.1. *Suppose that $k, l, m,$ and n are positive integers with $k \leq n$. Suppose further that A is a real $m \times n$ matrix, B is a real $m \times k$ matrix whose columns constitute a subset of the columns of A , P is a real $k \times n$ matrix, F is a real $m \times l$ matrix, and G is a real $l \times m$ matrix.*

Then,

$$\|BP - A\| \leq \|FGA - A\|(\|P\| + 1) + \|F\|\|GBP - GA\|. \tag{28}$$

Proof. We observe that

$$\|BP - A\| \leq \|BP - FGBP\| + \|FGBP - FGA\| + \|FGA - A\|, \tag{29}$$

$$\|BP - FGBP\| \leq \|B - FGB\|\|P\|, \tag{30}$$

and

$$\|FGBP - FGA\| \leq \|F\|\|GBP - GA\|. \tag{31}$$

Since the columns of B constitute a subset of the columns of A , it follows that the columns of $B - FGB$ constitute a subset of the columns of $A - FGA$, and therefore,

$$\|B - FGB\| \leq \|A - FGA\|. \tag{32}$$

Combining (29), (30), (31), and (32) yields (28). \square

Remark 4.2. Since the columns of B constitute a subset of the columns of A in Lemma 4.1, it follows that the columns of GB constitute a subset of the columns of GA . Conversely, whenever a matrix S is formed by gathering distinct columns of GA together into S , then clearly $S = GB$ for some matrix B whose columns constitute a subset of the columns of A .

The following lemma states that the product AQQ^T of matrices A , Q , and Q^T is a good approximation to a matrix A , provided that there exist matrices G and S such that

1. the columns of Q are orthonormal,
2. QS is a good approximation to $(GA)^T$, and
3. there exists a matrix F such that $\|F\|$ is not too large, and FGA is a good approximation to A .

Lemma 4.3. Suppose that k, l, m , and n are positive integers with $k \leq n$. Suppose further that A is a real $m \times n$ matrix, Q is a real $n \times k$ matrix whose columns are orthonormal, S is a real $k \times l$ matrix, F is a real $m \times l$ matrix, and G is a real $l \times m$ matrix.

Then,

$$\|AQQ^T - A\| \leq 2\|FGA - A\| + 2\|F\|\|QS - (GA)^T\|. \quad (33)$$

Proof. The proof is straightforward, but tedious, as follows.

We obtain from the triangle inequality that

$$\|AQQ^T - A\| \leq \|AQQ^T - FGAQQ^T\| + \|FGAQQ^T - FGA\| + \|FGA - A\|. \quad (34)$$

First, we provide a bound for $\|AQQ^T - FGAQQ^T\|$. Clearly,

$$\|AQQ^T - FGAQQ^T\| \leq \|A - FGA\|\|Q\|\|Q^T\|. \quad (35)$$

It follows from the fact that the columns of Q are orthonormal that

$$\|Q\| \leq 1 \quad (36)$$

and

$$\|Q^T\| \leq 1. \quad (37)$$

Combining (35), (36), and (37) yields

$$\|AQQ^T - FGAQQ^T\| \leq \|A - FGA\|. \quad (38)$$

Next, we provide a bound for $\|FGAQQ^T - FGA\|$. Clearly,

$$\|FGAQQ^T - FGA\| \leq \|F\|\|GAQQ^T - GA\|. \quad (39)$$

It follows from the triangle inequality that

$$\|GAQQ^T - GA\| \leq \|GAQQ^T - S^TQ^TQQ^T\| + \|S^TQ^TQQ^T - S^TQ^T\| + \|S^TQ^T - GA\|. \quad (40)$$

Furthermore,

$$\|GAQQ^T - S^TQ^TQQ^T\| \leq \|GA - S^TQ^T\|\|Q\|\|Q^T\|. \quad (41)$$

Combining (41), (36), and (37) yields

$$\|GAQQ^T - S^TQ^TQQ^T\| \leq \|GA - S^TQ^T\|. \quad (42)$$

Also, it follows from the fact that the columns of Q are orthonormal that

$$Q^TQ = \mathbf{1}. \quad (43)$$

It follows from (43) that

$$\|S^TQ^TQQ^T - S^TQ^T\| = 0. \quad (44)$$

Combining (40), (42), and (44) yields

$$\|GAQQ^T - GA\| \leq 2\|S^TQ^T - GA\|. \quad (45)$$

Combining (39) and (45) yields

$$\|FGAQ Q^T - FGA\| \leq 2\|F\| \|S^T Q^T - GA\|. \tag{46}$$

Combining (34), (38), and (46) yields (33). \square

The following lemma states that, for any matrix A , and matrix G whose entries are i.i.d. Gaussian random variables of zero mean and unit variance, with very high probability there exists a matrix F with a reasonably small norm, such that FGA is a good approximation to A .

Lemma 4.4. *Suppose that k, l, m , and n are positive integers with $k \leq l$, such that $l < m$ and $l < n$. Suppose further that A is a real $m \times n$ matrix, G is a real $l \times m$ matrix whose entries are i.i.d. Gaussian random variables of zero mean and unit variance, and β and γ are positive real numbers, such that $\gamma > 1$ and*

$$1 - \frac{1}{\sqrt{2\pi(l-k+1)}} \left(\frac{e}{(l-k+1)\beta} \right)^{l-k+1} - \frac{1}{4(\gamma^2-1)\sqrt{\pi m \gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^m \tag{47}$$

is nonnegative.

Then, there exists a real $m \times l$ matrix F such that

$$\|FGA - A\| \leq \sqrt{2lm\beta^2\gamma^2 + 1}\sigma_{k+1} \tag{48}$$

and

$$\|F\| \leq \sqrt{l}\beta \tag{49}$$

with probability not less than the amount in (47), where σ_{k+1} is the $(k+1)$ st greatest singular value of A .

Proof. We prove the existence of a matrix F satisfying (48) and (49) by constructing one.

We start by forming an SVD of A ,

$$A = U\Sigma V^T, \tag{50}$$

where U is a real unitary $m \times m$ matrix, Σ is a real $m \times n$ matrix whose entries are nonnegative everywhere and zero off of the main diagonal, and V is a real unitary $n \times n$ matrix, such that

$$\Sigma_{i,i} = \sigma_i \tag{51}$$

for all $i = 1, 2, \dots, \min(m, n) - 1, \min(m, n)$, where $\Sigma_{i,i}$ is the entry in row i and column i of Σ , and σ_i is the i th greatest singular value of A .

Next, we define auxiliary matrices H, R , and P . We define H to be the leftmost $l \times k$ block of the $l \times m$ matrix GU , and R to be the rightmost $l \times (m-k)$ block of GU , so that

$$GU = \left(H \parallel R \right). \tag{52}$$

Combining the fact that U is real and unitary, and the fact that the entries of G are i.i.d. Gaussian random variables of zero mean and unit variance, we see that the entries of H are also i.i.d. Gaussian random variables of zero mean and unit variance, as are the entries of R . We define $H^{(-1)}$ to be the real $k \times l$ matrix given by the formula

$$H^{(-1)} = (H^T H)^{-1} H^T. \tag{53}$$

We define P to be the $m \times l$ matrix whose uppermost $k \times l$ block is $H^{(-1)}$, and whose entries in the lowermost $(m-k) \times l$ block are zero, so that

$$P = \begin{pmatrix} H^{(-1)} \\ \mathbf{0} \end{pmatrix}. \tag{54}$$

Finally, we define F to be the $m \times l$ matrix given by

$$F = UP = U \begin{pmatrix} H^{(-1)} \\ \mathbf{0} \end{pmatrix}. \tag{55}$$

Combining (53), (20), the fact that the entries of H are i.i.d. Gaussian random variables of zero mean and unit variance, and Lemma 3.15 yields

$$\|H^{(-1)}\| \leq \sqrt{l}\beta \tag{56}$$

with probability not less than

$$1 - \frac{1}{\sqrt{2\pi(l-k+1)}} \left(\frac{e}{(l-k+1)\beta} \right)^{l-k+1}. \tag{57}$$

Combining (55), (56), and the fact that U is unitary yields (49).

We now show that F defined in (55) satisfies (48).

We define S to be the leftmost uppermost $k \times k$ block of Σ , and T to be the rightmost lowermost $(m-k) \times (n-k)$ block of Σ , so that

$$\Sigma = \begin{pmatrix} S & \mathbf{0} \\ \mathbf{0} & T \end{pmatrix}. \tag{58}$$

Combining (50), (52), and (55) yields

$$FGA - A = U \left(\begin{pmatrix} H^{(-1)} \\ \mathbf{0} \end{pmatrix} (H \parallel R) - \mathbf{1} \right) \Sigma V^T. \tag{59}$$

Combining (53) and (58) yields

$$\left(\begin{pmatrix} H^{(-1)} \\ \mathbf{0} \end{pmatrix} (H \parallel R) - \mathbf{1} \right) \Sigma = \begin{pmatrix} \mathbf{0} \parallel H^{(-1)}RT \\ \mathbf{0} \parallel -T \end{pmatrix}. \tag{60}$$

Furthermore,

$$\left\| \begin{pmatrix} \mathbf{0} \parallel H^{(-1)}RT \\ \mathbf{0} \parallel -T \end{pmatrix} \right\|^2 \leq \|H^{(-1)}RT\|^2 + \|T\|^2. \tag{61}$$

Moreover,

$$\|H^{(-1)}RT\| \leq \|H^{(-1)}\| \|R\| \|T\|. \tag{62}$$

Combining (58) and (51) yields

$$\|T\| \leq \sigma_{k+1}. \tag{63}$$

Combining (59), (60), (61), (62), (63), and the fact that U and V are unitary yields

$$\|FGA - A\| \leq \sqrt{\|H^{(-1)}\|^2 \|R\|^2 + 1} \sigma_{k+1}. \tag{64}$$

Combining Lemma 3.14 and the fact that the entries of R are i.i.d. Gaussian random variables of zero mean and unit variance shows that

$$\|R\| \leq \sqrt{2m}\gamma \tag{65}$$

with probability not less than

$$1 - \frac{1}{4(\gamma^2 - 1)\sqrt{\pi m\gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^m. \tag{66}$$

Combining (64), (56), and (65) yields (48). \square

The following lemma is very similar to Lemma 4.4. Lemma 4.5 is tighter than Lemma 4.4 when the singular values of the matrix A decay sufficiently fast, and the numbers j and l in the lemma are both much less than m .

Lemma 4.5. *Suppose that $j, k, l, m,$ and n are positive integers with $k \leq l$, such that $k + j < m$ and $k + j < n$, as well as $l < m$ and $l < n$. Suppose further that A is a real $m \times n$ matrix, G is a real $l \times m$ matrix whose entries are i.i.d. Gaussian random variables of zero mean and unit variance, and β and γ are positive real numbers, such that $\gamma > 1$ and*

$$\Phi = 1 - \frac{1}{\sqrt{2\pi(l-k+1)}} \left(\frac{e}{(l-k+1)\beta} \right)^{l-k+1} - \frac{1}{4(\gamma^2 - 1)\sqrt{\pi \max(m-k-j, l)\gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^{\max(m-k-j, l)} - \frac{1}{4(\gamma^2 - 1)\sqrt{\pi \max(j, l)\gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^{\max(j, l)} \tag{67}$$

is nonnegative.

Then, there exists a real $m \times l$ matrix F such that

$$\|FGA - A\| \leq \sqrt{2l \max(j, l) \beta^2 \gamma^2 + 1} \sigma_{k+1} + \sqrt{2l \max(m - k - j, l) \beta^2 \gamma^2 + 1} \sigma_{k+j+1} \tag{68}$$

and

$$\|F\| \leq \sqrt{l} \beta \tag{69}$$

with probability not less than the amount in (67), where σ_{k+1} is the $(k + 1)$ st greatest singular value of A , and σ_{k+j+1} is the $(k + j + 1)$ st greatest singular value of A .

Proof. We prove the existence of a matrix F satisfying (68) and (69) by constructing one.

We start by forming an SVD of A ,

$$A = U \Sigma V^T, \tag{70}$$

where U is a real unitary $m \times m$ matrix, Σ is a real $m \times n$ matrix whose entries are nonnegative everywhere and zero off of the main diagonal, and V is a real unitary $n \times n$ matrix, such that

$$\Sigma_{i,i} = \sigma_i \tag{71}$$

for all $i = 1, 2, \dots, \min(m, n) - 1, \min(m, n)$, where $\Sigma_{i,i}$ is the entry in row i and column i of Σ , and σ_i is the i th greatest singular value of A .

Next, we define auxiliary matrices H , R , Γ , and P . We define H to be the leftmost $l \times k$ block of the $l \times m$ matrix GU , R to be the $l \times j$ block of GU whose first column is the $(k + 1)$ st column of GU , and Γ to be the rightmost $l \times (m - j - k)$ block of GU , so that

$$GU = (H \parallel R \parallel \Gamma). \tag{72}$$

Combining the fact that U is real and unitary, and the fact that the entries of G are i.i.d. Gaussian random variables of zero mean and unit variance, we see that the entries of H are also i.i.d. Gaussian random variables of zero mean and unit variance, as are the entries of R , and as are the entries of Γ . We define $H^{(-1)}$ to be the real $k \times l$ matrix given by the formula

$$H^{(-1)} = (H^T H)^{-1} H^T. \tag{73}$$

We define P to be the real $m \times l$ matrix whose uppermost $k \times l$ block is $H^{(-1)}$, whose entries are zero in the $j \times l$ block whose first row is the $(k + 1)$ st row of P , and whose entries in the lowermost $(m - k - j) \times l$ block are zero, so that

$$P = \begin{pmatrix} H^{(-1)} \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix}. \tag{74}$$

Finally, we define F to be the $m \times l$ matrix given by

$$F = UP = U \begin{pmatrix} H^{(-1)} \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix}. \tag{75}$$

Combining (73), (20), the fact that the entries of H are i.i.d. Gaussian random variables of zero mean and unit variance, and Lemma 3.15 yields

$$\|H^{(-1)}\| \leq \sqrt{l} \beta \tag{76}$$

with probability not less than

$$1 - \frac{1}{\sqrt{2\pi(l - k + 1)}} \left(\frac{e}{(l - k + 1)\beta} \right)^{l - k + 1}. \tag{77}$$

Combining (75), (76), and the fact that U is unitary yields (69).

We now show that F defined in (75) satisfies (68).

We define S to be the leftmost uppermost $k \times k$ block of Σ , T to be the $j \times j$ block of Σ whose leftmost uppermost entry is the entry in the $(k + 1)$ st row and $(k + 1)$ st column of Σ , and Θ to be the rightmost lowermost $(m - k - j) \times (n - k - j)$ block of Σ , so that

$$\Sigma = \begin{pmatrix} S & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & T & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \Theta \end{pmatrix}. \quad (78)$$

Combining (70), (72), and (75) yields

$$FGA - A = U \left(\begin{pmatrix} H^{(-1)} \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix} (H \parallel R \parallel \Gamma) - \mathbf{1} \right) \Sigma V^T. \quad (79)$$

Combining (73) and (78) yields

$$\left(\begin{pmatrix} H^{(-1)} \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix} (H \parallel R \parallel \Gamma) - \mathbf{1} \right) \Sigma = \begin{pmatrix} \mathbf{0} & H^{(-1)}RT & H^{(-1)}\Gamma\Theta \\ \mathbf{0} & -T & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & -\Theta \end{pmatrix}. \quad (80)$$

Furthermore,

$$\left\| \begin{pmatrix} \mathbf{0} & H^{(-1)}RT & H^{(-1)}\Gamma\Theta \\ \mathbf{0} & -T & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & -\Theta \end{pmatrix} \right\|^2 \leq \|H^{(-1)}RT\|^2 + \|H^{(-1)}\Gamma\Theta\|^2 + \|T\|^2 + \|\Theta\|^2. \quad (81)$$

Moreover,

$$\|H^{(-1)}RT\| \leq \|H^{(-1)}\| \|R\| \|T\| \quad (82)$$

and

$$\|H^{(-1)}\Gamma\Theta\| \leq \|H^{(-1)}\| \|\Gamma\| \|\Theta\|. \quad (83)$$

Combining (78) and (71) yields

$$\|T\| \leq \sigma_{k+1} \quad (84)$$

and

$$\|\Theta\| \leq \sigma_{k+j+1}. \quad (85)$$

Combining (79), (80), (81), (82), (83), (84), (85), and the fact that U and V are unitary yields

$$\|FGA - A\|^2 \leq (\|H^{(-1)}\|^2 \|R\|^2 + 1)(\sigma_{k+1})^2 + (\|H^{(-1)}\|^2 \|\Gamma\|^2 + 1)(\sigma_{k+j+1})^2. \quad (86)$$

Combining Lemma 3.14 and the fact that the entries of R are i.i.d. Gaussian random variables of zero mean and unit variance, as are the entries of Γ , yields

$$\|R\| \leq \sqrt{2 \max(j, l)} \gamma \quad (87)$$

and

$$\|\Gamma\| \leq \sqrt{2 \max(m-k-j, l)} \gamma, \quad (88)$$

with probability not less than

$$1 - \frac{1}{4(\gamma^2 - 1)\sqrt{\pi} \max(m-k-j, l)\gamma^2} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^{\max(m-k-j, l)} - \frac{1}{4(\gamma^2 - 1)\sqrt{\pi} \max(j, l)\gamma^2} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^{\max(j, l)}. \quad (89)$$

Combining (86), (76), (87), and (88) shows that

$$\|FGA - A\|^2 \leq (2l \max(j, l)\beta^2\gamma^2 + 1)(\sigma_{k+1})^2 + (2l \max(m-k-j, l)\beta^2\gamma^2 + 1)(\sigma_{k+j+1})^2 \quad (90)$$

with probability not less than the amount in (67). Combining (90) and the fact that

$$\sqrt{x+y} \leq \sqrt{x} + \sqrt{y} \quad (91)$$

for any nonnegative real numbers x and y yields (68). \square

Given an $m \times n$ matrix A , and a matrix G whose entries are i.i.d. Gaussian random variables of zero mean and unit variance, the following lemma provides a highly probable upper bound on the singular values of the product GA in terms of the singular values of A ; the lemma is most useful when the singular values of A decay sufficiently fast, and the numbers j and l in the lemma are both much less than m .

Lemma 4.6. *Suppose that j, k, l, m , and n are positive integers with $k < l$, such that $k + j < m$ and $k + j < n$. Suppose further that A is a real $m \times n$ matrix, G is a real $l \times m$ matrix whose entries are i.i.d. Gaussian random variables of zero mean and unit variance, and γ is a positive real number, such that $\gamma > 1$ and*

$$\begin{aligned} \Psi = & 1 - \frac{1}{4(\gamma^2 - 1)\sqrt{\pi \max(m - k - j, l)\gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2 - 1}} \right)^{\max(m - k - j, l)} \\ & - \frac{1}{4(\gamma^2 - 1)\sqrt{\pi \max(k + j, l)\gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2 - 1}} \right)^{\max(k + j, l)} \end{aligned} \tag{92}$$

is nonnegative.

Then,

$$\rho_{k+1} \leq \sqrt{2 \max(k + j, l)\gamma} \sigma_{k+1} + \sqrt{2 \max(m - k - j, l)\gamma} \sigma_{k+j+1} \tag{93}$$

with probability not less than Ψ defined in (92), where ρ_{k+1} is the $(k + 1)$ st greatest singular value of GA , σ_{k+1} is the $(k + 1)$ st greatest singular value of A , and σ_{k+j+1} is the $(k + j + 1)$ st greatest singular value of A .

Proof. We start by forming an SVD of A ,

$$A = U \Sigma V^T, \tag{94}$$

where U is a real unitary $m \times m$ matrix, Σ is a real $m \times n$ matrix whose entries are nonnegative everywhere and zero off of the main diagonal, and V is a real unitary $n \times n$ matrix, such that

$$\Sigma_{i,i} = \sigma_i \tag{95}$$

for all $i = 1, 2, \dots, \min(m, n) - 1, \min(m, n)$, where $\Sigma_{i,i}$ is the entry in row i and column i of Σ , and σ_i is the i th greatest singular value of A .

Combining (94) and the fact that V is unitary yields that GA has the same singular values as $GU\Sigma$.

Next, we define auxiliary matrices H and R . We define H to be the leftmost $l \times (k + j)$ block of the $l \times m$ matrix GU , and R to be the rightmost $l \times (m - k - j)$ block of GU , so that

$$GU = \left(H \begin{smallmatrix} \vdots \\ \mathbf{0} \end{smallmatrix} R \right). \tag{96}$$

Combining the fact that U is real and unitary, and the fact that the entries of G are i.i.d. Gaussian random variables of zero mean and unit variance, we see that the entries of H are also i.i.d. Gaussian random variables of zero mean and unit variance, as are the entries of R .

Combining (96) and the fact that GA has the same singular values as $GU\Sigma$ yields that GA has the same singular values as $\left(H \begin{smallmatrix} \vdots \\ \mathbf{0} \end{smallmatrix} \right) \Sigma + \left(\mathbf{0} \begin{smallmatrix} \vdots \\ R \end{smallmatrix} \right) \Sigma$.

It follows from (95) that

$$\left\| \left(\mathbf{0} \begin{smallmatrix} \vdots \\ R \end{smallmatrix} \right) \Sigma \right\| \leq \|R\| \sigma_{k+j+1}. \tag{97}$$

Combining (25) and (97) yields

$$\rho_{k+1} \leq \tau_{k+1} + \|R\| \sigma_{k+j+1}, \tag{98}$$

where ρ_{k+1} is the $(k + 1)$ st greatest singular value of $\left(H \begin{smallmatrix} \vdots \\ \mathbf{0} \end{smallmatrix} \right) \Sigma + \left(\mathbf{0} \begin{smallmatrix} \vdots \\ R \end{smallmatrix} \right) \Sigma$, and τ_{k+1} is the $(k + 1)$ st greatest singular value of $\left(H \begin{smallmatrix} \vdots \\ \mathbf{0} \end{smallmatrix} \right) \Sigma$; ρ_{k+1} is also the $(k + 1)$ st greatest singular value of GA , since GA has the same singular values as $\left(H \begin{smallmatrix} \vdots \\ \mathbf{0} \end{smallmatrix} \right) \Sigma + \left(\mathbf{0} \begin{smallmatrix} \vdots \\ R \end{smallmatrix} \right) \Sigma$.

Furthermore,

$$\left\| \left(H \begin{smallmatrix} \vdots \\ \mathbf{0} \end{smallmatrix} \right) \right\| \leq \|H\|. \tag{99}$$

Combining (22), (99), and (95) yields

$$\tau_{k+1} \leq \|H\| \sigma_{k+1}. \tag{100}$$

Combining Lemma 3.14 and the fact that the entries of H are i.i.d. Gaussian random variables of zero mean and unit variance, as are the entries of R , shows that

$$\|H\| \leq \sqrt{2 \max(k+j, l)} \gamma \quad (101)$$

and

$$\|R\| \leq \sqrt{2 \max(m-k-j, l)} \gamma \quad (102)$$

with probability not less than the amount in (92).

Combining (98), (100), (101), and (102) yields (93). \square

5. Detailed description of the algorithm

In this section, we describe the algorithm of the present paper. In Section 5.1, we discuss approximations to interpolative decompositions. In Section 5.2, we discuss approximations to SVDs. In Section 5.3, we tabulate the computational costs of various parts of the algorithm. In Section 5.4, we describe Table 1, providing numerical bounds on the probability that the randomized algorithm fails to meet its accuracy estimates. We remind the reader that we denote the spectral norm of any matrix A by $\|A\|$.

5.1. Interpolative decomposition

Suppose that k , m , and n are positive integers with $k < m$ and $k < n$, and A is a real $m \times n$ matrix. In this subsection, we will collect together k appropriately chosen columns of A into a real $m \times k$ matrix B , and construct a real $k \times n$ matrix P , such that

$$\|P\| \leq \sqrt{4k(n-k) + 1} \quad (103)$$

and

$$\|BP - A\| \lesssim \sqrt{klmn} \sigma_{k+1}, \quad (104)$$

where σ_{k+1} is the $(k+1)$ st greatest singular value of A , and l is a user-specified integer with $l > k$, such that $l < m$ and $l < n$ (for example, $l = k + 20$). To do so, we make the following three steps:

1. Using a random number generator, form a real $l \times m$ matrix G whose entries are i.i.d. Gaussian random variables of zero mean and unit variance, and compute the $l \times n$ product matrix

$$R = GA. \quad (105)$$

2. Using Observation 3.3, form a real $l \times k$ matrix S whose columns constitute a subset of the columns of R , and a real $k \times n$ matrix P satisfying (103), such that

$$\|SP - R\| \leq \sqrt{4k(n-k) + 1} \rho_{k+1}, \quad (106)$$

where ρ_{k+1} is the $(k+1)$ st greatest singular value of R .

3. Due to Step 2, the columns of S constitute a subset of the columns of R . In other words, there exists a finite sequence $i_1, i_2, \dots, i_{k-1}, i_k$ of integers such that, for any $j = 1, 2, \dots, k-1, k$, the j th column of S is the i_j th column of R . Collect the corresponding columns of A into a real $m \times k$ matrix B , so that, for any $j = 1, 2, \dots, k-1, k$, the j th column of B is the i_j th column of A .

It is easy to see that the matrices B and P satisfy (103) and (104). Indeed, Step 2 above guarantees (103) by construction. Moreover, combining (105) and Remark 4.2 yields

$$S = GB. \quad (107)$$

Combining (106), (105), and (107) yields

$$\|GBP - GA\| \leq \sqrt{4k(n-k) + 1} \rho_{k+1}, \quad (108)$$

where ρ_{k+1} is the $(k+1)$ st greatest singular value of R . Suppose that β and γ are positive real numbers such that $\gamma > 1$ and

$$\chi = 1 - \frac{1}{\sqrt{2\pi(l-k+1)}} \left(\frac{e}{(l-k+1)\beta} \right)^{l-k+1} - \frac{1}{2(\gamma^2-1)\sqrt{\pi m \gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^m \quad (109)$$

is nonnegative. Then, combining (28), (48), (49), (103), (108), (22), (105), and Lemma 3.14 yields

$$\|BP - A\| \leq (\sqrt{2lm\beta^2\gamma^2 + 1}(\sqrt{4k(n-k) + 1} + 1) + \beta\gamma\sqrt{2lm}\sqrt{4k(n-k) + 1})\sigma_{k+1} \quad (110)$$

with probability not less than χ defined in (109), where σ_{k+1} is the $(k + 1)$ st greatest singular value of A . The bound (110) is a precise version of (104). For example, choosing $\beta = 3/4$, $\gamma^2 = 5$, and $l = k + 20$, and combining (110) and (109), we obtain

$$\|BP - A\| \leq 10\sqrt{k(k + 20)mn}\sigma_{k+1} \quad (111)$$

with probability not less than $1 - 10^{-17}$. Table 1 contains similar results obtained by taking other values for $l - k$, β , and γ .

Observation 5.1. When the singular values of A decay sufficiently fast, and l is much less than m , the factors $\sqrt{2lm\beta^2\gamma^2 + 1}$ and $\sqrt{2lm}$ in (110) are much larger than necessary. Indeed, suppose that j is a positive integer with $k + j < m$ and $k + j < n$, and β and γ are positive real numbers, such that $\gamma > 1$ and $\Phi + \Psi > 1$, where Φ is defined in (67), and Ψ is defined in (92). Then, combining (28), (68), (69), (103), (108), (93), and (105) yields

$$\|BP - A\| \leq \xi\sigma_{k+1} + \eta\sigma_{k+j+1} \quad (112)$$

with probability not less than $\Phi + \Psi - 1$, where Φ is defined in (67), Ψ is defined in (92), σ_{k+1} is the $(k + 1)$ st greatest singular value of A , and σ_{k+j+1} is the $(k + j + 1)$ st greatest singular value of A , and where

$$\xi = \sqrt{2l \max(j, l)\beta^2\gamma^2 + 1}(\sqrt{4k(n-k) + 1} + 1) + \beta\gamma\sqrt{2l \max(k + j, l)}\sqrt{4k(n-k) + 1} \quad (113)$$

and

$$\eta = \sqrt{2l \max(m - k - j, l)\beta^2\gamma^2 + 1}(\sqrt{4k(n-k) + 1} + 1) + \beta\gamma\sqrt{2l \max(m - k - j, l)}\sqrt{4k(n-k) + 1}. \quad (114)$$

When j , k , and l are all much less than m , clearly ξ is much less than η . Moreover, if $\eta\sigma_{k+j+1}$ is less than $\xi\sigma_{k+1}$, then (112) yields

$$\|BP - A\| \leq 2\xi\sigma_{k+1} \quad (115)$$

with probability not less than $\Phi + \Psi - 1$, where Φ is defined in (67), Ψ is defined in (92), σ_{k+1} is the $(k + 1)$ st greatest singular value of A , and ξ is defined in (113). Please note that the right-hand side of (115) is independent of m .

Remark 5.2. If we choose $l = k$ in the algorithm of the present subsection (instead of choosing $l > k$), then we must replace (106) with the formula

$$\|SP - R\| = 0. \quad (116)$$

All other aspects of the algorithm stay the same in the case that $l = k$. In particular, (110) and (112) hold in the case that $l = k$, too.

5.2. Singular value decomposition

Suppose that k , m , and n are positive integers with $k < m$ and $k < n$, and A is a real $m \times n$ matrix. In this subsection, we will construct an approximation to an SVD of A such that

$$\|U\Sigma V^T - A\| \lesssim \sqrt{lm}\sigma_{k+1}, \quad (117)$$

where U is a real $m \times k$ matrix whose columns are orthonormal, V is a real $n \times k$ matrix whose columns are orthonormal, Σ is a diagonal real $k \times k$ matrix whose entries are all nonnegative, σ_{k+1} is the $(k + 1)$ st greatest singular value of A , and l is a user-specified integer with $l > k$, such that $l < m$ and $l < n$ (for example, $l = k + 20$). To do so, we make the following five steps:

1. Using a random number generator, form a real $l \times m$ matrix G whose entries are i.i.d. Gaussian random variables of zero mean and unit variance, and compute the $l \times n$ product matrix

$$R = GA. \quad (118)$$

2. Using an SVD, form a real $n \times k$ matrix Q whose columns are orthonormal, such that there exists a real $k \times l$ matrix S for which

$$\|QS - R^T\| \leq \rho_{k+1}, \quad (119)$$

where ρ_{k+1} is the $(k+1)$ st greatest singular value of R . (See Observation 3.6 for details concerning the construction of such a matrix Q .)

3. Compute the $m \times k$ product matrix

$$T = AQ. \quad (120)$$

4. Form an SVD of T ,

$$T = U\Sigma W^T, \quad (121)$$

where U is a real $m \times k$ matrix whose columns are orthonormal, W is a real $k \times k$ matrix whose columns are orthonormal, and Σ is a real $k \times k$ matrix whose entries are all nonnegative and zero off of the main diagonal. (See, for example, Chapter 8 in [7] for details concerning the construction of such an SVD.)

5. Compute the $n \times k$ product matrix

$$V = QW. \quad (122)$$

It is easy to see that the matrices U , Σ , and V satisfy (117). Indeed, suppose that β and γ are positive real numbers such that $\gamma > 1$ and χ defined in (109) is nonnegative. Then, combining (33), (48), (49), (120), (121), (122), (119), (22), (118), and Lemma 3.14 yields

$$\|U\Sigma V^T - A\| \leq (2\sqrt{2lm\beta^2\gamma^2 + 1} + 2\sqrt{2lm}\beta\gamma)\sigma_{k+1} \quad (123)$$

with probability not less than χ defined in (109), where σ_{k+1} is the $(k+1)$ st greatest singular value of A . The bound (123) is a precise version of (117). For example, choosing $\beta = 3/4$, $\gamma^2 = 5$, and $l = k + 20$, and combining (123) and (109), we obtain

$$\|U\Sigma V^T - A\| \leq 10\sqrt{(k+20)m}\sigma_{k+1} \quad (124)$$

with probability not less than $1 - 10^{-17}$. Table 1 contains similar results obtained by taking other values for $l - k$, β , and γ .

Observation 5.3. When the singular values of A decay sufficiently fast, and l is much less than m , the factors $\sqrt{2lm\beta^2\gamma^2 + 1}$ and $\sqrt{2lm}$ in (123) are much larger than necessary. Indeed, suppose that j is a positive integer with $k + j < m$ and $k + j < n$, and β and γ are positive real numbers, such that $\gamma > 1$ and $\Phi + \Psi > 1$, where Φ is defined in (67), and Ψ is defined in (92). Then, combining (33), (68), (69), (120), (121), (122), (119), (93), and (118) yields

$$\|U\Sigma V^T - A\| \leq \xi\sigma_{k+1} + \eta\sigma_{k+j+1} \quad (125)$$

with probability not less than $\Phi + \Psi - 1$, where Φ is defined in (67), Ψ is defined in (92), σ_{k+1} is the $(k+1)$ st greatest singular value of A , and σ_{k+j+1} is the $(k+j+1)$ st greatest singular value of A , and where

$$\xi = 2\sqrt{2l\max(j, l)\beta^2\gamma^2 + 1} + 2\sqrt{2l\max(k+j, l)\beta\gamma} \quad (126)$$

and

$$\eta = 2\sqrt{2l\max(m-k-j, l)\beta^2\gamma^2 + 1} + 2\sqrt{2l\max(m-k-j, l)\beta\gamma}. \quad (127)$$

When j , k , and l are all much less than m , clearly ξ is much less than η . Moreover, if $\eta\sigma_{k+j+1}$ is less than $\xi\sigma_{k+1}$, then (125) yields

$$\|U\Sigma V^T - A\| \leq 2\xi\sigma_{k+1} \quad (128)$$

with probability not less than $\Phi + \Psi - 1$, where Φ is defined in (67), Ψ is defined in (92), σ_{k+1} is the $(k+1)$ st greatest singular value of A , and ξ is defined in (126). Please note that the right-hand side of (128) is independent of m .

Remark 5.4. In a certain sense, the factor $10\sqrt{(k+20)m}$ in (124) is not too large. Indeed, if $k \leq 80$ and $m \leq 10,000,000,000$, then $10\sqrt{(k+20)m}\sigma_{k+1} \leq 10^7\sigma_{k+1}$. In many practical circumstances, we can make $10^7\sigma_{k+1}$ arbitrarily small by choosing k appropriately.

Remark 5.5. If we choose $l = k$ in the algorithm of the present subsection (instead of choosing $l > k$), then we must replace (119) with the formula

$$\|QS - R^T\| = 0. \tag{129}$$

All other aspects of the algorithm stay the same in the case that $l = k$. In particular, (123) and (125) hold in the case that $l = k$, too.

5.3. CPU time requirements

In this subsection, we tabulate the numbers of floating-point operations required by the algorithms described in Sections 5.1 and 5.2, as applied once to a matrix A .

5.3.1. Interpolative decomposition

The algorithm of Section 5.1 incurs the following costs in order to compute an approximation to an interpolative decomposition of A :

1. Forming R in (105) requires applying A^T to l vectors.
2. Computing S and P in (106) or (116) costs $\mathcal{O}(lkn \log(n))$.
3. Forming B in (107) requires applying A to k vectors, where each vector has a single entry of 1 and $n - 1$ entries of 0.

Summing up the costs in Steps 1–3 above, we conclude that the algorithm of Section 5.1 costs

$$C_{ID} = k \cdot C_A + l \cdot C_{A^T} + \mathcal{O}(lkn \log(n)), \tag{130}$$

where C_A is the cost of applying A to a real $n \times 1$ column vector, and C_{A^T} is the cost of applying A^T to a real $m \times 1$ column vector.

5.3.2. Singular value decomposition

The algorithm of Section 5.2 incurs the following costs in order to compute an approximation to an SVD of A :

1. Forming R in (118) requires applying A^T to l vectors.
2. Computing Q in (119) or (129) costs $\mathcal{O}(l^2n)$.
3. Forming T in (120) requires applying A to k vectors.
4. Computing the SVD (121) of T costs $\mathcal{O}(k^2m)$.
5. Forming V in (122) costs $\mathcal{O}(k^2n)$.

Summing up the costs in Steps 1–5 above, we conclude that the algorithm of Section 5.2 costs

$$C_{SVD} = k \cdot C_A + l \cdot C_{A^T} + \mathcal{O}(k^2m + l^2n), \tag{131}$$

where C_A is the cost of applying A to a real $n \times 1$ column vector, and C_{A^T} is the cost of applying A^T to a real $m \times 1$ column vector.

Remark 5.6. We observe that the algorithm of the present paper only requires applying A to k vectors and A^T to l vectors; it does not require explicit access to the individual entries of A . This consideration can be important when the entries of A are not available explicitly, but instead A and A^T are available solely in the form of procedures for their applications to arbitrary vectors. Often such procedures for applying A and A^T cost much less than the standard procedure for applying a dense matrix to a vector.

Remark 5.7. Without any further analysis, we can observe that the cost in (130) of applying the algorithm of Section 5.1 to any $m \times n$ matrix A is in principle less than the $\mathcal{O}(kmn \log(n))$ cost of using the algorithm discussed in Observation 3.3 directly on A , provided that l is sufficiently less than $k \log(n)$, and that both m and n are much greater than both k and l .

Remark 5.8. When “ QR ” decompositions are used as in [5] to compute the matrices S and P in (106) and (116), the cost of the algorithm of Section 5.1 is usually less than the cost of the algorithm of Section 5.2.

5.4. Description of Table 1

Tables 1.1–1.6 provide an upper bound $\Pi_{l-k, \beta, \gamma}$ on the probability that

$$\|U\Sigma V^T - A\| > \zeta \sqrt{lm} \sigma_{k+1}, \tag{132}$$

Table 1

For description of the table see Section 5.4.

$l-k$	β	γ^2	ζ	$\Pi_{l-k,\beta,\gamma}$
0	10	3	10^2	.15E-00
0	88	4	10^3	.18E-01
0	790	5	10^4	.21E-02
0	6600	7	10^5	.18E-03
0	58000	9	10^6	.19E-04

(1.1)

$l-k$	β	γ^2	ζ	$\Pi_{l-k,\beta,\gamma}$
2	7.9	5	10^2	.37E-03
2	66	7	10^3	.61E-06
2	589	9	10^4	.85E-09
2	5300	11	10^5	.12E-11
2	51000	12	10^6	.14E-14

(1.3)

$l-k$	β	γ^2	ζ	$\Pi_{l-k,\beta,\gamma}$
0	10	3	100	.15E-00
1	7.9	5	100	.85E-02
2	7.9	5	100	.37E-03
4	7.2	6	100	.44E-06
8	6.6	7	100	.12E-12
16	6.6	7	100	.55E-28
32	6.2	8	100	.83E-63

(1.5)

$l-k$	β	γ^2	ζ	$\Pi_{l-k,\beta,\gamma}$
1	7.9	5	10^2	.85E-02
1	72	6	10^3	.11E-03
1	620	8	10^4	.15E-05
1	5500	10	10^5	.17E-07
1	53000	11	10^6	.19E-09

(1.2)

$l-k$	β	γ^2	ζ	$\Pi_{l-k,\beta,\gamma}$
2	1.02	3	10	.18E-00
4	0.88	4	10	.16E-01
8	0.88	4	10	.10E-04
16	0.79	5	10	.16E-12
32	0.72	6	10	.59E-32

(1.4)

$l-k$	β	γ^2	ζ	$\Pi_{l-k,\beta,\gamma}$
0	10	3	1000	.18E-01
1	72	6	1000	.11E-03
2	66	7	1000	.61E-06
4	62	8	1000	.93E-11
8	58	9	1000	.37E-21
16	55	10	1000	.73E-44
32	53	11	1000	.14E-93

(1.6)

where U , Σ , and V are the matrices in the approximation to an SVD of the $m \times n$ matrix A in (123). In (132), k and l are any positive integers with $k \leq l$, such that $l < m$ and $l < n$, σ_{k+1} is the $(k+1)$ st greatest singular value of A , and ζ takes on the values specified by the penultimate columns of the tables. The quantity $\Pi_{l-k,\beta,\gamma}$ is defined by the formula

$$\Pi_{l-k,\beta,\gamma} = \frac{1}{\sqrt{2\pi(l-k+1)}} \left(\frac{e}{(l-k+1)\beta} \right)^{l-k+1} + \frac{1}{2(\gamma^2-1)\sqrt{\pi(l-k+2)\gamma^2}} \left(\frac{2\gamma^2}{e\gamma^2-1} \right)^{l-k+2}. \tag{133}$$

It follows from the fact that $m > l > l-k$ that $\Pi_{l-k,\beta,\gamma}$ provides an upper bound on $1 - \chi$ defined in (109). The quantities $l-k$, β , and γ take on the values specified by the first, second, and third columns of the tables. The quantity ζ is specified by β and γ via (123). Please note that $\Pi_{l-k,\beta,\gamma}$ depends only on $l-k$, β , and γ , and provides an upper bound that is otherwise independent of k , m , n , and A ; (124) provides a similar bound. When the singular values of A decay sufficiently fast, and l is much less than m , the factor of \sqrt{m} in the right-hand side of (132) is larger than necessary; see Observation 5.3 above.

Remark 5.9. Due to (110), the quantity $\Pi_{l-k,\beta,\gamma}$ defined in (133) also provides an upper bound on the probability that

$$\|BP - A\| > \zeta \sqrt{klmn} \sigma_{k+1}, \tag{134}$$

where B and P are the matrices in the approximation to an interpolative decomposition of the $m \times n$ matrix A in (110).

6. Numerical results

In this section, we describe the results of five numerical tests of the algorithm of the present paper. Table 2 summarizes the numerical output of the examples described in the present section.

Tables 2.1–2.5 display the results of applying the algorithm of the present paper to a real $n \times n$ matrix A , for the indicated values of n . The matrix A is defined at the end of the present section. The numbers k and l are those from Section 5; k is the rank of the approximations to A , and l is the number of rows in the matrix G whose entries are i.i.d. Gaussian random variables of zero mean and unit variance (the algorithm uses the product GA). The number t_{ID} is the time in seconds required to compute the approximation to an interpolative decomposition. The number t_{SVD} is the time in seconds required to compute the approximation to an SVD of A . (Please note that our implementation is optimized for accuracy and for analyzing the numerical properties of the algorithm, and is probably not very efficient.) The number σ_{k+1} is that from the definition of A below; furthermore, σ_{k+1} appears in the bounds (110) and (123) on the errors of the approximations. The number δ_{ID} is the norm of the difference between A and the approximation BP to an interpolative decomposition of A , that is,

$$\delta_{ID} = \|BP - A\|, \tag{135}$$

Table 2

For description of the table see Section 6.

k	l	n	t_{ID}	t_{SVD}	σ_{k+1}	δ_{ID}	δ_{SVD}
10	10	10^2	.34E-03	.81E-03	.20E-15	.14E-13	.83E-14
10	10	10^3	.32E-02	.70E-02	.20E-15	.89E-14	.58E-14
10	10	10^4	.34E-01	.78E-01	.20E-15	.78E-14	.14E-13
10	10	10^5	.57E-00	.11E+01	.20E-15	.27E-13	.22E-13
10	10	10^6	.58E+01	.13E+02	.20E-15	.68E-13	.93E-13

(2.1)

k	l	n	t_{ID}	t_{SVD}	σ_{k+1}	δ_{ID}	δ_{SVD}
10	10	10^2	.31E-03	.87E-03	.10E-07	.42E-06	.28E-06
10	10	10^3	.31E-02	.71E-02	.10E-07	.38E-06	.27E-06
10	10	10^4	.34E-01	.78E-01	.10E-07	.22E-06	.18E-06
10	10	10^5	.55E-00	.11E+01	.10E-07	.83E-07	.73E-07
10	10	10^6	.56E+01	.12E+02	.10E-07	.15E-06	.14E-06

(2.2)

k	l	n	t_{ID}	t_{SVD}	σ_{k+1}	δ_{ID}	δ_{SVD}
30	30	10^5	.41E+01	.88E+01	.10E-08	.62E-07	.33E-07
30	31	10^5	.41E+01	.90E+01	.10E-08	.36E-07	.29E-07
30	32	10^5	.42E+01	.92E+01	.10E-08	.23E-07	.13E-07
30	34	10^5	.44E+01	.97E+01	.10E-08	.19E-07	.11E-07
30	38	10^5	.47E+01	.11E+02	.10E-08	.19E-07	.11E-07
30	46	10^5	.53E+01	.13E+02	.10E-08	.86E-08	.41E-08

(2.3)

k	l	n	t_{ID}	t_{SVD}	σ_{k+1}	δ_{ID}	δ_{SVD}
2	2	$4 \cdot 10^2$.13E-03	.28E-03	.10E-07	.12E-05	.17E-06
2	2	$4 \cdot 10^3$.13E-02	.28E-02	.10E-07	.43E-05	.12E-06
2	2	$4 \cdot 10^4$.13E-01	.29E-01	.10E-07	.10E-04	.26E-06
2	2	$4 \cdot 10^5$.16E-00	.36E-00	.10E-07	.28E-04	.98E-07
2	2	$4 \cdot 10^6$.16E+01	.35E+01	.10E-07	.93E-04	.39E-06

(2.4)

k	l	n	t_{ID}	t_{SVD}	σ_{k+1}	δ_{ID}	δ_{SVD}
10	10	10^2	.12E-03	.80E-03	.10E-06	.14E-05	.53E-06
10	10	10^3	.12E-02	.46E-02	.10E-06	.41E-05	.18E-05
10	10	10^4	.14E-01	.48E-01	.10E-06	.83E-05	.34E-05
10	10	10^5	.20E-00	.65E-00	.10E-06	.25E-04	.11E-04
10	10	10^6	.20E+01	.74E+01	.10E-06	.11E-03	.34E-04

(2.5)

where the matrices B and P are those from (110). The number δ_{SVD} is the norm of the difference between A and the approximation $U \Sigma V^T$ to an SVD of A , that is,

$$\delta_{SVD} = \|U \Sigma V^T - A\|, \tag{136}$$

where the matrices U , Σ , and V are those from (123).

All numbers displayed in Table 2 are the maximum values obtained from three independent realizations of the random variables involved.

The values of δ_{ID} and δ_{SVD} displayed in Tables 2.2, 2.3, 2.4, and 2.5 are those obtained via the power method for estimating the norm of a matrix, after the estimates stabilized to three significant figures. The values of δ_{ID} and δ_{SVD} displayed in Table 2.1 are those obtained after 100 iterations of the power method. The estimates of δ_{ID} and δ_{SVD} summarized in Table 2.1 did not stabilize to three significant figures after 100 (or any other number of) iterations, undoubtedly due to round-off.

We performed all computations using IEEE standard double-precision variables, whose mantissas have approximately one bit of precision less than 16 digits (so that the relative precision of the variables is approximately .2E-15). We ran all computations on one core of a 1.86 GHz Intel Centrino Core Duo microprocessor with 2 MB of L2 cache and 1 GB of RAM. We compiled the Fortran 77 code using the Lahey/Fujitsu Linux Express v6.2 compiler, with the optimization flag `--o2` enabled.

In our implementation, we computed SVDs using 2-sided plane (Jacobi/Givens) rotations (see, for example, Chapter 8 in [7]). We used an algorithm based upon pivoted “QR” decompositions to compute the matrices S and P in (106) and (116) (see, for example, Chapter 5 in [7] for a description of “QR” decompositions, and [5] for further details regarding our particular implementation). Please note that because we used the algorithm of [5] to compute the matrices S and P in (106) and (116), the cost of computing the interpolative decomposition was of the same order as the cost of computing the SVD, as mentioned in Remark 5.8.

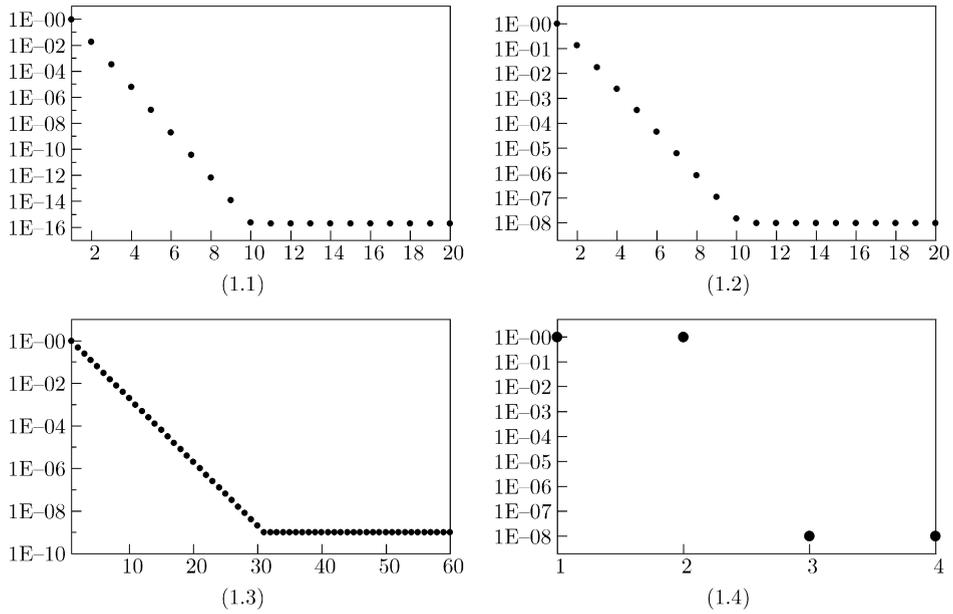


Fig. 1. For description of the figure see Section 6.

In Examples 1, 2, and 3, we use a pseudorandom number generator to construct real $n \times 1$ vectors $\mu^1, \mu^2, \dots, \mu^{j-1}, \mu^j$, and $v^1, v^2, \dots, v^{j-1}, v^j$, such that their entries are a realization of i.i.d. Gaussian random variables of zero mean and unit variance, with $j = 20$ in Examples 1 and 2, and $j = 60$ in Example 3. We orthonormalize $\mu^1, \mu^2, \dots, \mu^{j-1}, \mu^j$ via the Gram–Schmidt process with reorthogonalization (see, for example, [2]) to obtain real $n \times 1$ vectors $u^1, u^2, \dots, u^{j-1}, u^j$, and do the same with $v^1, v^2, \dots, v^{j-1}, v^j$ to obtain real $n \times 1$ vectors $v^1, v^2, \dots, v^{j-1}, v^j$. We denote by $\sigma_1, \sigma_2, \dots, \sigma_{j-1}, \sigma_j$ the positive real numbers displayed in Fig. 1 (we use the numbers in Fig. 1.1 for Example 1, the numbers in Fig. 1.2 for Example 2, and those in Fig. 1.3 for Example 3). We define A to be the $n \times n$ matrix given by the formula

$$A = \sum_{i=1}^j u^i \sigma_i (v^i)^T. \tag{137}$$

Clearly, the rank of A is j . Since $u^1, u^2, \dots, u^{j-1}, u^j$ are orthonormal, as are $v^1, v^2, \dots, v^{j-1}, v^j$, the i th singular value of A is σ_i , for all $i = 1, 2, \dots, j - 1, j$. Table 2 displays the results of applying the algorithm of the present paper to A , for various values of n (Table 2.1 displays the results for Example 1, Table 2.2 displays the results for Example 2, and Table 2.3 displays those for Example 3).

Example 4 is designed to illustrate that factors of the order of $\sqrt{4k(n-k)+1}$ are necessary in bounds such as (110) and (112). This example is identical to Examples 1, 2, and 3, using the same matrix A defined in (137), but with n assumed to be divisible by 8, with $j = 4$, and using the numbers $\sigma_1, \sigma_2, \sigma_3$, and σ_4 displayed in Fig. 1.4 ($\sigma_1 = 1, \sigma_2 = 1, \sigma_3 = .1E-7$, and $\sigma_4 = .1E-7$), and with the following vectors u^1, u^2, u^3, u^4 , and v^1, v^2, v^3, v^4 :

$$(u^1)^T = \frac{1}{\sqrt{n}}(1 \ 1 \ \dots \ 1 \ 1), \tag{138}$$

$$(u^2)^T = \frac{1}{\sqrt{n}}(1 \ -1 \ 1 \ -1 \ \dots \ 1 \ -1 \ 1 \ -1), \tag{139}$$

$$(u^3)^T = \frac{1}{\sqrt{n}}(1 \ 1 \ -1 \ -1 \ 1 \ 1 \ -1 \ -1 \ \dots \ 1 \ 1 \ -1 \ -1 \ 1 \ 1 \ -1 \ -1), \tag{140}$$

$$(u^4)^T = \frac{1}{\sqrt{n}}(1 \ 1 \ 1 \ 1 \ -1 \ -1 \ -1 \ -1 \ \dots \ 1 \ 1 \ 1 \ 1 \ -1 \ -1 \ -1 \ -1), \tag{141}$$

$$(v^1)^T = \frac{1}{\sqrt{n-1}}(1 \ 1 \ \dots \ 1 \ 1 \ 1 \ 0), \tag{142}$$

$$(v^2)^T = (0 \ 0 \ \dots \ 0 \ 0 \ 0 \ 1), \tag{143}$$

$$(v^3)^T = \frac{1}{\sqrt{n-2}}(1 \ -1 \ 1 \ -1 \ \dots \ 1 \ -1 \ 1 \ -1 \ 0 \ 0), \tag{144}$$

$$(v^4)^T = \frac{1}{\sqrt{2}}(1 \ 0 \ -1 \ 0 \ 0 \ 0 \ \dots \ 0 \ 0), \tag{145}$$

that is,

1. a. every entry of u^1 is $1/\sqrt{n}$,
 b. entries $1, 2, \dots, n-2, n-1$ of v^1 are $1/\sqrt{n-1}$, and entry n of v^1 is 0,
2. a. every even entry of u^2 is $-1/\sqrt{n}$, and every odd entry of u^2 is $1/\sqrt{n}$,
 b. entries $1, 2, \dots, n-2, n-1$ of v^2 are 0, and entry n of v^2 is 1,
3. a. the first pair of entries of u^3 is $1/\sqrt{n}$, the second pair of entries is $-1/\sqrt{n}$, the third pair of entries is $1/\sqrt{n}$, the fourth pair of entries is $-1/\sqrt{n}$, and so on (with each successive pair of entries alternating sign),
 b. every even entry of v^3 except for entry n is $-1/\sqrt{n-2}$, every odd entry of v^3 except for entry $n-1$ is $1/\sqrt{n-2}$, and entries $n-1$ and n of v^3 are 0,
4. a. the first quadruplet of entries of u^4 is $1/\sqrt{n}$, the second quadruplet of entries is $-1/\sqrt{n}$, the third quadruplet of entries is $1/\sqrt{n}$, the fourth quadruplet of entries is $-1/\sqrt{n}$, and so on (with each successive quadruplet of entries alternating sign), and
 b. entry 1 of v^4 is $1/\sqrt{2}$, entry 2 of v^4 is 0, entry 3 is $-1/\sqrt{2}$, and entries $4, 5, \dots, n-1, n$ are 0.

For this example (Example 4),

$$u^1\sigma_1(v^1)^T + u^2\sigma_2(v^2)^T = \frac{1}{\sqrt{n(n-1)}} \begin{pmatrix} 1 & 1 & \dots & 1 & 1 & \sqrt{n-1} \\ 1 & 1 & \dots & 1 & 1 & -\sqrt{n-1} \\ 1 & 1 & \dots & 1 & 1 & \sqrt{n-1} \\ 1 & 1 & \dots & 1 & 1 & -\sqrt{n-1} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 1 & 1 & \dots & 1 & 1 & \sqrt{n-1} \\ 1 & 1 & \dots & 1 & 1 & -\sqrt{n-1} \end{pmatrix}. \tag{146}$$

Clearly, u^1, u^2, u^3 , and u^4 are orthonormal, as are v^1, v^2, v^3 , and v^4 . Therefore, the i th singular value of A is σ_i , for all $i = 1, 2, 3, 4$. Table 2.4 displays the results of applying the algorithm of the present paper to A , for various values of n .

Example 5 is designed to illustrate that factors of the order of \sqrt{m} are necessary in bounds such as (110) and (123). In this example, we define A to be the $n \times n$ matrix given by the formula

$$A = uv^T + \sigma \mathbf{1}, \tag{147}$$

where $\sigma = .1E-6$, and u and v are the real $n \times 1$ column vectors given by the formulae

$$u^T = (1 \ 0 \ 0 \ \dots \ 0 \ 0) \tag{148}$$

and

$$v^T = \frac{1}{\sqrt{n}}(1 \ 1 \ \dots \ 1 \ 1). \tag{149}$$

It follows from (24) that the k th greatest singular value σ_k of A is equal to $.1E-6$ for $k = 2, 3, \dots, n-2, n-1$. Table 2.5 displays the results of applying the algorithm of the present paper to A , for various values of n .

Remark 6.1. All numerical data that we have examined – including the data displayed in Table 2, as well as the data from further experiments – appear to satisfy the bounds (110), (112), (123), and (125). The loss of precision displayed in Table 2.1 as n increases is probably largely due to round-off (compare Table 2.2), whereas the loss of precision in δ_{ID} displayed in Table 2.4 as n increases suggests that any bounds such as (110) and (112) must contain factors of the order of $\sqrt{4k(n-k)+1}$. The loss of precision displayed in Table 2.5 as n increases suggests that any bounds such as (110) and (123) must contain factors of the order of \sqrt{m} . In contrast, in many practical situations the bounds mentioned in Observations 5.1 and 5.3 are effectively independent of m .

7. Conclusions

We have presented a randomized algorithm for the computation of an approximation to an interpolative decomposition and to an SVD of a matrix A for which A and A^T can be applied rapidly to arbitrary vectors. Given any positive integer k , the algorithm constructs a rank- k approximation whose accuracy measured in the spectral norm is of roughly the same order as

the accuracy of the best possible rank- k approximation. The algorithm has a rather negligible probability of failure (10^{-17} is typical), and operates reliably independently of the structure of A (unlike the classical Lanczos and power methods for computing an approximation to A). Indeed, for most interesting parameter settings, failure of the algorithm is too improbable to detect — running the algorithm the 10^{17} times usually required for failure is not really reasonable.

If the matrix A to be approximated has no degenerate or nearly degenerate singular values, then the number of times our scheme applies the matrices A and A^T to vectors is similar to (though often somewhat less than) the number of times required by the classical Lanczos method (for a description of the Lanczos method, see, for example, Chapter 9 in [7]). In the presence of degenerate singular values, the algorithm of the present paper tends to be more efficient for the low-rank approximation of matrices, even when compared to the block Lanczos method (for a description of the block Lanczos method, see, for example, Section 9.2.6 in [7]). Furthermore, in the parallel computing environment, the Lanczos method has well-known difficulties because it is iterative; our approach should not encounter such difficulties.

We are investigating several straightforward extensions of the scheme, and will report on these at a later date:

1. In the present article, the rank k of the approximation to be constructed and the number l of test vectors are fixed. In practice, we would like to adjust k and l during the course of the algorithm in order to guarantee that the approximation attains a prescribed accuracy, preferably using as small a number l as possible.
2. The present article constructs approximations to interpolative decompositions and to singular value decompositions. We have constructed a similar algorithm for approximating the Schur decomposition (see, for example, Theorem 7.1.3 in [7] for a description of the Schur decomposition).
3. When the matrix to be approximated is self-adjoint and nonnegative definite, Step 3 of the algorithm of Section 5.2 is unnecessary and may be replaced with less expensive manipulations.

Acknowledgments

We would like to thank R.R. Coifman, R. Kannan, and G.C. Papanicolaou for useful discussions.

References

- [1] M.W. Berry, S.A. Pulatova, G.W. Stewart, Algorithm 844: Computing sparse reduced-rank approximations to sparse matrices, *ACM Trans. Math. Software* 31 (2) (2005) 252–269.
- [2] Å. Björck, Numerics of Gram–Schmidt orthogonalization, *Linear Algebra Appl.* 197–198 (1994) 297–316.
- [3] T.F. Chan, P.C. Hansen, Some applications of the rank-revealing QR factorization, *SIAM J. Sci. Stat. Comput.* 13 (3) (1992) 727–741.
- [4] Z. Chen, J.J. Dongarra, Condition numbers of Gaussian random matrices, *SIAM J. Matrix Anal. Appl.* 27 (3) (2005) 603–620.
- [5] H. Cheng, Z. Gimbutas, P.-G. Martinsson, V. Rokhlin, On the compression of low rank matrices, *SIAM J. Sci. Comput.* 26 (4) (2005) 1389–1404.
- [6] H.H. Goldstine, J. von Neumann, Numerical inverting of matrices of high order, II, *Proc. Amer. Math. Soc.* 2 (1951) 188–202.
- [7] G.H. Golub, C.F. Van Loan, *Matrix Computations*, 3rd ed., Johns Hopkins University Press, Baltimore, Maryland, 1996.
- [8] S.A. Goreinov, E.E. Tyrtyshnikov, The maximal-volume concept in approximation by low-rank matrices, in: V. Olshevsky (Ed.), *Structured Matrices in Mathematics, Computer Science, and Engineering I: Proceedings of an AMS–IMS–SIAM Joint Summer Research Conference*, University of Colorado, Boulder, June 27–July 1, 1999, in: *Contemp. Math.*, vol. 280, AMS Publications, Providence, RI, 2001, pp. 47–51.
- [9] S.A. Goreinov, E.E. Tyrtyshnikov, N.L. Zamarashkin, Pseudo-skeleton approximations by matrices of maximal volume, *Math. Notes* 62 (4) (1997) 515–519.
- [10] S.A. Goreinov, E.E. Tyrtyshnikov, N.L. Zamarashkin, A theory of pseudoskeleton approximations, *Linear Algebra Appl.* 261 (1–3) (1997) 1–21.
- [11] M. Gu, S.C. Eisenstat, Efficient algorithms for computing a strong rank-revealing QR factorization, *SIAM J. Sci. Comput.* 17 (4) (1996) 848–869.
- [12] P.-G. Martinsson, V. Rokhlin, M. Tygert, On interpolation and integration in finite-dimensional spaces of bounded functions, *Commun. Appl. Math. Comput. Sci.* 1 (2006) 133–142.
- [13] P.-G. Martinsson, V. Rokhlin, M. Tygert, A randomized algorithm for the approximation of matrices, *Tech. Rep.*, 1361, Yale University Department of Computer Science, June 2006.
- [14] C.H. Papadimitriou, P. Raghavan, H. Tamaki, S. Vempala, Latent semantic indexing: A probabilistic analysis, *J. Comput. System Sci.* 61 (2000) 217–235.
- [15] W. Press, S. Teukolsky, W. Vetterling, B. Flannery, *Numerical Recipes*, 2nd ed., Cambridge University Press, Cambridge, UK, 1992.
- [16] T. Sarlós, Improved approximation algorithms for large matrices via random projections, in: *Proceedings of FOCS 2006, the 47th Annual IEEE Symposium on Foundations of Computer Science*, 2006, pp. 143–152, a revised version is available at <http://www.ilab.sztaki.hu/~stamas/publications/rp-long.pdf>.
- [17] G.W. Stewart, Four algorithms for the efficient computation of truncated pivoted QR approximations to a sparse matrix, *Numer. Math.* 83 (2) (1999) 313–323.
- [18] E.E. Tyrtyshnikov, Incomplete cross approximation in the mosaic-skeleton method, *Computing* 64 (4) (2000) 367–380.
- [19] J.H. Wilkinson, *The Algebraic Eigenvalue Problem*, Oxford University Press, Oxford, 1965.