

A randomized algorithm for principal component analysis

Vladimir Rokhlin, Arthur Szlam, and Mark Tygert

UCLA Computational and Applied Math. Technical Report 08-60
September 5, 2008

Abstract

Principal component analysis (PCA) requires the computation of a low-rank approximation to a matrix containing the data being analyzed. In many applications of PCA, the best possible accuracy of any rank-deficient approximation is at most a few digits (measured in the spectral norm, relative to the spectral norm of the matrix being approximated). In such circumstances, existing efficient algorithms do not guarantee good accuracy for the approximations they produce, unless one or both dimensions of the matrix being approximated are small. We describe an efficient algorithm for the low-rank approximation of matrices that produces accuracy very close to the best possible, for matrices of arbitrary sizes. We illustrate our theoretical results via several numerical examples.

1 Introduction

Principal component analysis (PCA) is among the most widely used techniques in statistics, data analysis, and data mining. PCA is the basis of many machine learning methods, including the latent semantic analysis of large databases of text and HTML documents described in [2]. Computationally, PCA amounts to the low-rank approximation of a matrix containing the data being analyzed. The present article describes an algorithm for the low-rank approximation of matrices, suitable for PCA. This paper demonstrates both theoretically and via numerical examples that the algorithm efficiently produces low-rank approximations whose accuracies are very close to the best possible.

The canonical construction of the best possible rank- k approximation to a real $m \times n$ matrix A uses the singular value decomposition (SVD) of A ,

$$A = U \Sigma V^T, \tag{1}$$

where U is a real unitary $m \times m$ matrix, V is a real unitary $n \times n$ matrix, and Σ is a real $m \times n$ matrix whose only nonzero entries appear in nonincreasing order on the diagonal and are nonnegative. The diagonal entries $\sigma_1, \sigma_2, \dots, \sigma_{m-1}, \sigma_m$ of Σ are known as the singular values of A . The best rank- k approximation to A is

$$A \approx \tilde{U} \tilde{\Sigma} \tilde{V}^T, \tag{2}$$

where \tilde{U} is the $m \times k$ matrix whose columns are the leftmost k columns of U , \tilde{V} is the $n \times k$ matrix whose columns are the leftmost k columns of V , and $\tilde{\Sigma}$ is the $k \times k$ matrix whose only nonzero entries appear in nonincreasing order on the diagonal and are the k greatest singular values of A . This approximation is “best” in the sense that the spectral norm $\|A - B\|$ of the difference between A and a rank- k matrix B is minimal for $B = \tilde{U} \tilde{\Sigma} \tilde{V}^T$. In fact,

$$\|A - \tilde{U} \tilde{\Sigma} \tilde{V}^T\| = \sigma_{k+1}, \quad (3)$$

where σ_{k+1} is the $(k + 1)^{\text{st}}$ greatest singular value of A . For more information about the SVD, see, for example, Chapter 8 in [5].

For definiteness, let us assume that $m \leq n$ and that A is an arbitrary (dense) real $m \times n$ matrix. To compute a rank- k approximation to A , one might form the matrices U , Σ , and V in (1), and then use them to construct \tilde{U} , $\tilde{\Sigma}$, and \tilde{V} in (2). However, even computing just Σ , the leftmost m columns of U , and the leftmost m columns of V requires at least $\mathcal{O}(nm^2)$ floating-point operations using any of the standard algorithms (see, for example, Chapter 8 in [5]). Alternatively, one might use pivoted QR -decomposition algorithms, which require $\mathcal{O}(nmk)$ flops and typically produce a rank- k approximation B to A such that

$$\|A - B\| \leq 10\sqrt{m} \sigma_{k+1}, \quad (4)$$

where $\|A - B\|$ is the spectral norm of $A - B$, and σ_{k+1} is the $(k + 1)^{\text{st}}$ greatest singular value of A (see, for example, Chapter 5 in [5]). Better yet, the algorithms of [6] require only about $\mathcal{O}(nmk)$ flops to produce a rank- k approximation guaranteed to satisfy (4).

Unfortunately, while the accuracy in (4) is sufficient for many applications of low-rank approximation, PCA often involves $m \geq 10,000$, and a “signal-to-noise ratio” $\sigma_1/\sigma_{k+1} \leq 100$, where $\sigma_1 = \|A\|$ is the greatest singular value of A , and σ_{k+1} is the $(k + 1)^{\text{st}}$ greatest. Moreover, the singular values $\leq \sigma_{k+1}$ often arise from noise in the process generating the data in A , making the singular values of A decay so slowly that $\sigma_m \geq \sigma_{k+1}/10$. When $m \geq 10,000$, $\sigma_1/\sigma_{k+1} \leq 100$, and $\sigma_m \geq \sigma_{k+1}/10$, the rank- k approximation B produced by a pivoted QR -decomposition algorithm typically satisfies $\|A - B\| \sim \|A\|$ — the “approximation” B is effectively unrelated to the matrix A being approximated! For large matrices whose “signal-to-noise ratio” σ_1/σ_{k+1} is less than 10,000, the \sqrt{m} factor in (4) may be unacceptable. Now, pivoted QR -decomposition algorithms are not the only algorithms which can compute a rank- k approximation using $\mathcal{O}(nmk)$ flops. However, other algorithms, such as those of [11] and [8], also yield accuracies involving factors of \sqrt{m} when the singular values σ_{k+1} , σ_{k+2} , σ_{k+3} , \dots of A decay slowly. (The decay is rather slow if, for example, $\sigma_{k+j} \sim j^\alpha \sigma_{k+1}$ for $j = 1, 2, 3, \dots$, with $-1/2 < \alpha \leq 0$.)

The algorithm described in the present paper produces a rank- k approximation B to A such that

$$\|A - B\| \leq C m^{1/(4i+2)} \sigma_{k+1} \quad (5)$$

with very high probability (typically $1 - 10^{-15}$, independent of A), where $\|A - B\|$ is the spectral norm of $A - B$, i is a nonnegative integer specified by the user, σ_{k+1} is the $(k + 1)^{\text{st}}$ greatest singular value of A , and C is a constant independent of A that theoretically may depend on the parameters of the algorithm. (Numerical evidence such as that in Section 5 suggests at the very least that $C < 10$; (83) and (70) in Section 4 provide more complicated

theoretical bounds on C .) The algorithm requires $\mathcal{O}(nmki)$ floating-point operations. In many applications of PCA, $i = 1$ or $i = 2$ is sufficient, and the algorithm then requires only $\mathcal{O}(nmk)$ flops. The algorithm provides the rank- k approximation B in the form of an SVD, outputting three matrices, \tilde{U} , $\tilde{\Sigma}$, and \tilde{V} , such that $B = \tilde{U} \tilde{\Sigma} \tilde{V}^T$, where the columns of \tilde{U} are orthonormal, the columns of \tilde{V} are orthonormal, and the entries of $\tilde{\Sigma}$ are all nonnegative and zero off of the diagonal.

The algorithm of the present paper is randomized, but succeeds with very high probability; for example, the bound (83) on its accuracy holds with probability greater than $1 - 10^{-15}$. The algorithm is similar to many recently discussed randomized algorithms for low-rank approximation, but produces approximations of higher accuracy when the singular values $\sigma_{k+1}, \sigma_{k+2}, \sigma_{k+3}, \dots$ of the matrix being approximated decay slowly; see, for example, [11] or [8]. The algorithm is a variant of that in [10]; [10] stimulated the authors' collaboration. The algorithm may be regarded as a generalization of the randomized power methods of [3] and [7], and in fact we use the latter to ascertain the approximations' accuracy rapidly and reliably.

The algorithm admits obvious “out-of-core” and parallel implementations (assuming that the user chooses the parameter i in (5) to be reasonably small). As with the algorithms of [7], [8], [9], [10], and [11], the core steps of the algorithm of the present paper involve the application of the matrix A being approximated and its transpose A^T to random vectors. The algorithm is more efficient when A and A^T can be applied rapidly to arbitrary vectors, such as when A is sparse.

Throughout the present paper, we use $\mathbf{1}$ to denote an identity matrix. We use $\mathbf{0}$ to denote a matrix whose entries are all zeros. For any matrix A , we use $\|A\|$ to denote the spectral norm of A , that is, $\|A\|$ is the greatest singular value of A . Furthermore, the entries of all matrices in the present paper are real valued, though the algorithm and analysis extend trivially to matrices whose entries are complex valued.

The present paper has the following structure: Section 2 collects together various known facts which later sections utilize. Section 3 provides the principal lemmas used in bounding the accuracy of the algorithm in Section 4. Section 4 describes the algorithm of the present paper. Section 5 illustrates the performance of the algorithm via several numerical examples.

2 Preliminaries

In this section, we summarize various facts about matrices and functions. Subsection 2.1 discusses the singular values of arbitrary matrices. Subsection 2.2 discusses the singular values of certain random matrices. Subsection 2.3 observes that a certain function is monotone.

2.1 Singular values of general matrices

The following trivial technical lemma will be needed in Section 3.

Lemma 2.1. *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 the least (that is, the n^{th} greatest) singular value σ_n of A is nonzero.*

Then,

$$\|(A^T A)^{-1} A^T\| = \frac{1}{\sigma_n}. \quad (6)$$

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 [12]).

Lemma 2.2. *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 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 2.4, the proof of this lemma provides an algorithm for computing Q and S , given R .

Lemma 2.3. *Suppose that k , l , and n are positive integers with $k < l \leq n$, 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

$$\|QS - R\| \leq \rho_{k+1}, \quad (7)$$

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

Proof. We start by forming an SVD of R ,

$$R = U \Sigma V^T, \quad (8)$$

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 diagonal $l \times l$ matrix, such that

$$\Sigma_{j,j} = \rho_j \quad (9)$$

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

$$U = \left(Q \parallel P \right). \quad (10)$$

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

$$\Sigma V^T = \begin{pmatrix} S \\ T \end{pmatrix}. \quad (11)$$

Combining (8), (9), (10), (11), and the fact that the columns of U are orthonormal, as are the columns of V , yields (7). \square

Observation 2.4. In order to compute the matrices Q and S in (7) from the matrix R , we can construct (8), and then form Q and S according to (10) and (11). (See, for example, Chapter 8 in [5] for details concerning the computation of the SVD.)

2.2 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 [4] provides an equivalent formulation of the lemma.

Lemma 2.5. *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*

$$1 - \frac{1}{4(\gamma^2 - 1)\sqrt{\pi n \gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2 - 1}} \right)^n \quad (12)$$

is nonnegative.

Then, the greatest singular value of G is at most $\sqrt{2n} \gamma$ with probability not less than the amount in (12).

Combining Lemmas 2.2 and 2.5 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 2.6. *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 (12) is nonnegative.*

Then, the greatest singular value of G is at most $\sqrt{2n} \gamma$ with probability not less than the amount in (12).

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 [1] and the proof of Lemma 4.1 in [1] together provide an equivalent formulation of Lemma 2.7.

Lemma 2.7. *Suppose that j and l are positive integers with $j \leq l$. Suppose further that G is a real $l \times j$ 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*

$$1 - \frac{1}{\sqrt{2\pi}(l-j+1)} \left(\frac{e}{(l-j+1)\beta} \right)^{l-j+1} \quad (13)$$

is nonnegative.

Then, the least (that is, the j^{th} greatest) singular value of G is at least $1/(\sqrt{l} \beta)$ with probability not less than the amount in (13).

2.3 A monotone function

The following technical lemma will be needed in Section 4.

Lemma 2.8. *Suppose that α is a nonnegative real number, and f is the function defined on $(0, \infty)$ via the formula*

$$f(x) = \frac{1}{\sqrt{2\pi x}} \left(\frac{e\alpha}{x} \right)^x. \quad (14)$$

Then, f decreases monotonically for $x > \alpha$.

Proof. The derivative of f is

$$f'(x) = f(x) \left(\ln \left(\frac{\alpha}{x} \right) - \frac{1}{2x} \right) \quad (15)$$

for any positive real number x . The right-hand side of (15) is negative when $x > \alpha$. \square

3 Mathematical apparatus

In this section, we provide lemmas to be used in Section 4 in bounding the accuracy of the algorithm of the present paper.

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 $(G(AA^T)^i A)^T$, and
3. there exists a matrix F such that $\|F\|$ is not too large, and $FG(AA^T)^i A$ is a good approximation to A .

Lemma 3.1. *Suppose that i , k , l , m , and n are positive integers with $k \leq l \leq m \leq n$. 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,

$$\|AQQ^T - A\| \leq 2\|FG(AA^T)^i A - A\| + 2\|F\| \|QS - (G(AA^T)^i A)^T\|. \quad (16)$$

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

To simplify notation, we define

$$B = (AA^T)^i A. \quad (17)$$

We obtain from the triangle inequality that

$$\|AQQ^T - A\| \leq \|AQQ^T - FGBQQ^T\| + \|FGBQQ^T - FGB\| + \|FGB - A\|. \quad (18)$$

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

$$\|AQQ^T - FGBQQ^T\| \leq \|A - FGB\| \|Q\| \|Q^T\|. \quad (19)$$

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

$$\|Q\| \leq 1 \quad (20)$$

and

$$\|Q^T\| \leq 1. \quad (21)$$

Combining (19), (20), and (21) yields

$$\|AQQ^T - FGBQQ^T\| \leq \|A - FGB\|. \quad (22)$$

Next, we provide a bound for $\|FGBQQ^T - FGB\|$. Clearly,

$$\|FGBQQ^T - FGB\| \leq \|F\| \|GBQQ^T - GB\|. \quad (23)$$

It follows from the triangle inequality that

$$\|GBQQ^T - GB\| \leq \|GBQQ^T - S^T Q^T Q Q^T\| + \|S^T Q^T Q Q^T - S^T Q^T\| + \|S^T Q^T - GB\|. \quad (24)$$

Furthermore,

$$\|GBQQ^T - S^T Q^T Q Q^T\| \leq \|GB - S^T Q^T\| \|Q\| \|Q^T\|. \quad (25)$$

Combining (25), (20), and (21) yields

$$\|GBQQ^T - S^T Q^T Q Q^T\| \leq \|GB - S^T Q^T\|. \quad (26)$$

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

$$Q^T Q = \mathbf{1}. \quad (27)$$

It follows from (27) that

$$\|S^T Q^T Q Q^T - S^T Q^T\| = 0. \quad (28)$$

Combining (24), (26), and (28) yields

$$\|GBQQ^T - GB\| \leq 2 \|S^T Q^T - GB\|. \quad (29)$$

Combining (23) and (29) yields

$$\|FGBQQ^T - FGB\| \leq 2 \|F\| \|S^T Q^T - GB\|. \quad (30)$$

Combining (18), (22), (30), and (17) yields (16). \square

The following lemma states that, for any positive integer i , 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 $FG(AA^T)^i A$ is a good approximation to A . This lemma is similar to Lemma 15 of [9].

Lemma 3.2. *Suppose that $i, j, k, l, m,$ and n are positive integers with $j < k < l < m \leq n$. 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 the j^{th} greatest singular value σ_j of A is positive, $\gamma > 1$, and*

$$\begin{aligned} \Phi = 1 - \frac{1}{\sqrt{2\pi(l-j+1)}} \left(\frac{e}{(l-j+1)\beta} \right)^{l-j+1} \\ - \frac{1}{4(\gamma^2-1)\sqrt{\pi \max(m-k, l)} \gamma^2} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^{\max(m-k, l)} \\ - \frac{1}{4(\gamma^2-1)\sqrt{\pi l} \gamma^2} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^l \end{aligned} \quad (31)$$

is nonnegative.

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

$$\begin{aligned} \|FG(AA^T)^i A - A\| \leq \sqrt{2l^2 \beta^2 \gamma^2 + 1} \sigma_{j+1} \\ + \sqrt{2l \max(m-k, l) \beta^2 \gamma^2 \left(\frac{\sigma_{k+1}}{\sigma_j} \right)^{4i} + 1} \sigma_{k+1} \end{aligned} \quad (32)$$

and

$$\|F\| \leq \frac{\sqrt{l} \beta}{(\sigma_j)^{2i}} \quad (33)$$

with probability not less than Φ defined in (31), where σ_j is the j^{th} greatest singular value of A , σ_{j+1} is the $(j+1)^{\text{st}}$ greatest singular value of A , and σ_{k+1} is the $(k+1)^{\text{st}}$ greatest singular value of A .

Proof. We prove the existence of a matrix F satisfying (32) and (33) by constructing one.

We start by forming an SVD of A ,

$$A = U \Sigma V^T, \quad (34)$$

where U is a real unitary $m \times m$ matrix, Σ is a real diagonal $m \times m$ matrix, and V is a real $n \times m$ matrix whose columns are orthonormal, such that

$$\Sigma_{p,p} = \sigma_p \quad (35)$$

for all $p = 1, 2, \dots, m-1, m$, where $\Sigma_{p,p}$ is the entry in row p and column p of Σ , and σ_p is the p^{th} greatest singular value of A .

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

$$GU = \left(H \mid R \mid \Gamma \right). \quad (36)$$

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 $j \times l$ matrix given by the formula

$$H^{(-1)} = (H^T H)^{-1} H^T. \quad (37)$$

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

$$\Sigma = \begin{pmatrix} S & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & T & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \Theta \end{pmatrix}. \quad (38)$$

We define P to be the real $m \times l$ matrix whose uppermost $j \times l$ block is the product $S^{-2i} H^{(-1)}$, whose entries are zero in the $(k-j) \times l$ block whose first row is the $(j+1)^{\text{st}}$ row of P , and whose entries in the lowermost $(m-k) \times l$ block are zero, so that

$$P = \begin{pmatrix} S^{-2i} H^{(-1)} \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix}. \quad (39)$$

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

$$F = U P = U \begin{pmatrix} S^{-2i} H^{(-1)} \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix}. \quad (40)$$

Combining (37), (6), the fact that the entries of H are i.i.d. Gaussian random variables of zero mean and unit variance, and Lemma 2.7 yields

$$\|H^{(-1)}\| \leq \sqrt{l} \beta \quad (41)$$

with probability not less than

$$1 - \frac{1}{\sqrt{2\pi(l-j+1)}} \left(\frac{e}{(l-j+1)\beta} \right)^{l-j+1}. \quad (42)$$

Combining (40), (41), (38), (35), the fact that Σ is zero off of its main diagonal, and the fact that U is unitary yields (33).

We now show that F defined in (40) satisfies (32).

Combining (34), (36), and (40) yields

$$F G (A A^T)^i A - A = U \left(\begin{pmatrix} S^{-2i} H^{(-1)} \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix} (H \parallel R \parallel \Gamma) \Sigma^{2i} - \mathbf{1} \right) \Sigma V^T. \quad (43)$$

Combining (37) and (38) yields

$$\begin{aligned} & \left(\left(\begin{array}{c|c} S^{-2i} H^{(-1)} & \\ \hline \mathbf{0} & \\ \hline \mathbf{0} & \end{array} \right) (H \mid R \mid \Gamma) \Sigma^{2i} - \mathbf{1} \right) \Sigma \\ &= \left(\begin{array}{c|cc} \mathbf{0} & S^{-2i} H^{(-1)} R T^{2i+1} & S^{-2i} H^{(-1)} \Gamma \Theta^{2i+1} \\ \hline \mathbf{0} & -T & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{0} & -\Theta \end{array} \right). \end{aligned} \quad (44)$$

Furthermore,

$$\begin{aligned} & \left\| \left(\begin{array}{c|cc} \mathbf{0} & S^{-2i} H^{(-1)} R T^{2i+1} & S^{-2i} H^{(-1)} \Gamma \Theta^{2i+1} \\ \hline \mathbf{0} & -T & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{0} & -\Theta \end{array} \right) \right\|^2 \\ & \leq \|S^{-2i} H^{(-1)} R T^{2i+1}\|^2 + \|S^{-2i} H^{(-1)} \Gamma \Theta^{2i+1}\|^2 + \|T\|^2 + \|\Theta\|^2. \end{aligned} \quad (45)$$

Moreover,

$$\|S^{-2i} H^{(-1)} R T^{2i+1}\| \leq \|S^{-1}\|^{2i} \|H^{(-1)}\| \|R\| \|T\|^{2i+1} \quad (46)$$

and

$$\|S^{-2i} H^{(-1)} \Gamma \Theta^{2i+1}\| \leq \|S^{-1}\|^{2i} \|H^{(-1)}\| \|\Gamma\| \|\Theta\|^{2i+1}. \quad (47)$$

Combining (38) and (35) yields

$$\|S^{-1}\| \leq \frac{1}{\sigma_j}, \quad (48)$$

$$\|T\| \leq \sigma_{j+1}, \quad (49)$$

and

$$\|\Theta\| \leq \sigma_{k+1}. \quad (50)$$

Combining (43), (44), (45), (46), (47), (48), (49), (50), and the fact that the columns of U are orthonormal, as are the columns of V , yields

$$\begin{aligned} \|FG(AA^T)^i A - A\|^2 & \leq \left(\|H^{(-1)}\|^2 \|R\|^2 \left(\frac{\sigma_{j+1}}{\sigma_j} \right)^{4i} + 1 \right) (\sigma_{j+1})^2 \\ & \quad + \left(\|H^{(-1)}\|^2 \|\Gamma\|^2 \left(\frac{\sigma_{k+1}}{\sigma_j} \right)^{4i} + 1 \right) (\sigma_{k+1})^2. \end{aligned} \quad (51)$$

Combining Lemma 2.6 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

$$\|R\| \leq \sqrt{2l} \gamma \quad (52)$$

and

$$\|\Gamma\| \leq \sqrt{2 \max(m-k, l)} \gamma, \quad (53)$$

with probability not less than

$$1 - \frac{1}{4(\gamma^2 - 1) \sqrt{\pi \max(m-k, l) \gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^{\max(m-k, l)} - \frac{1}{4(\gamma^2 - 1) \sqrt{\pi l \gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^l. \quad (54)$$

Combining (51), (41), (52), and (53) yields that

$$\begin{aligned} \|F G (A A^T)^i A - A\|^2 &\leq \left(2l^2 \beta^2 \gamma^2 \left(\frac{\sigma_{j+1}}{\sigma_j} \right)^{4i} + 1 \right) (\sigma_{j+1})^2 \\ &\quad + \left(2l \max(m-k, l) \beta^2 \gamma^2 \left(\frac{\sigma_{k+1}}{\sigma_j} \right)^{4i} + 1 \right) (\sigma_{k+1})^2 \end{aligned} \quad (55)$$

with probability not less than Φ defined in (31). Combining (55), the fact that $\sigma_{j+1} \leq \sigma_j$, and the fact that

$$\sqrt{x+y} \leq \sqrt{x} + \sqrt{y} \quad (56)$$

for any nonnegative real numbers x and y yields (32). \square

Given a 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 . This lemma is reproduced from [9], where it appears as Lemma 16.

Lemma 3.3. *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*

$$\begin{aligned} \Xi = 1 - \frac{1}{4(\gamma^2 - 1) \sqrt{\pi \max(m-k-j, l) \gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^{\max(m-k-j, l)} \\ - \frac{1}{4(\gamma^2 - 1) \sqrt{\pi \max(k+j, l) \gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^{\max(k+j, l)} \end{aligned} \quad (57)$$

is nonnegative.

Then,

$$\rho_{k+1} \leq \sqrt{2 \max(k+j, l)} \gamma \sigma_{k+1} + \sqrt{2 \max(m-k-j, l)} \gamma \sigma_{k+j+1} \quad (58)$$

with probability not less than Ξ defined in (57), 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 .

The following corollary follows immediately from the preceding lemma, by replacing the matrix A with $(A A^T)^i A$, the integer k with j , and the integer j with $k-j$.

Corollary 3.4. *Suppose i, j, k, l, m , and n are positive integers with $j < k < l < m \leq n$. 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*

$$\Psi = 1 - \frac{1}{4(\gamma^2 - 1) \sqrt{\pi \max(m-k, l) \gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^{\max(m-k, l)} - \frac{1}{4(\gamma^2 - 1) \sqrt{\pi l \gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^l \quad (59)$$

is nonnegative.

Then,

$$\rho_{j+1} \leq \sqrt{2l} \gamma (\sigma_{j+1})^{2i+1} + \sqrt{2 \max(m-k, l)} \gamma (\sigma_{k+1})^{2i+1} \quad (60)$$

with probability not less than Ψ defined in (59), where ρ_{j+1} is the $(j+1)^{\text{st}}$ greatest singular value of $G(AA^T)^i A$, σ_{j+1} is the $(j+1)^{\text{st}}$ greatest singular value of A , and σ_{k+1} is the $(k+1)^{\text{st}}$ greatest singular value of A .

4 Description of the algorithm

In this section, we describe the algorithm of the present paper, providing details about its accuracy and computational costs. Subsection 4.1 describes the basic algorithm. Subsection 4.2 tabulates the computational costs of the algorithm. Subsection 4.3 describes a complementary algorithm. Subsection 4.4 describes a computationally more expensive variant that is somewhat more accurate and tolerant to roundoff.

4.1 The algorithm

Suppose that i, k, m , and n are positive integers with $2k < m \leq n$, and A is a real $m \times n$ matrix. In this subsection, we will construct an approximation to an SVD of A such that

$$\|A - U \Sigma V^T\| \leq C m^{1/(4i+2)} \sigma_{k+1} \quad (61)$$

with very high probability, 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 real diagonal $k \times k$ matrix whose entries are all nonnegative, σ_{k+1} is the $(k+1)^{\text{st}}$ greatest singular value of A , and C is a constant independent of A that depends on the parameters of the algorithm. (Section 5 will give an empirical indication of the size of C , and (83) will give one of our best theoretical estimates to date.)

Intuitively, we could apply A^T to several random vectors, in order to identify the part of its range corresponding to the larger singular values. To enhance the decay of the singular values of A , we apply $A^T(AA^T)^i$ instead. Once we have identified most of the range of A^T , we perform several linear-algebraic manipulations in order to recover an approximation to A .

More precisely, we choose an integer $l > k$ such that $l \leq m - k$ (for example, $l = k + 12$), and make the following five steps:

- 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

$$R = G(AA^T)^i A. \quad (62)$$

- 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

$$\|QS - R^T\| \leq \rho_{k+1}, \quad (63)$$

where ρ_{k+1} is the $(k+1)^{\text{st}}$ greatest singular value of R . (See Observation 2.4 for details concerning the construction of such a matrix Q .)

- Compute the $m \times k$ product matrix

$$T = AQ. \quad (64)$$

- Form an SVD of T ,

$$T = U\Sigma W^T, \quad (65)$$

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 diagonal $k \times k$ matrix whose entries are all nonnegative. (See, for example, Chapter 8 in [5] for details concerning the construction of such an SVD.)

- Compute the $n \times k$ product matrix

$$V = QW. \quad (66)$$

The following theorem states precisely that the matrices U , Σ , and V satisfy (61). See (83) for a more compact (but less general) formulation.

Theorem 4.1. *Suppose that i , k , l , m , and n are positive integers with $k < l \leq m - k$ and $m \leq n$, and A is a real $m \times n$ matrix. Suppose further that β and γ are positive real numbers such that $\gamma > 1$,*

$$(l - k + 1)\beta \geq 1, \quad (67)$$

$$2l^2\gamma^2\beta^2 \geq 1, \quad (68)$$

and

$$\begin{aligned} \Pi = 1 - \frac{1}{2(\gamma^2 - 1)\sqrt{\pi(m-k)\gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2-1}}\right)^{m-k} - \frac{1}{2(\gamma^2 - 1)\sqrt{\pi l \gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2-1}}\right)^l \\ - \frac{1}{\sqrt{2\pi(l-k+1)}} \left(\frac{e}{(l-k+1)\beta}\right)^{l-k+1} \end{aligned} \quad (69)$$

is nonnegative. Suppose in addition that U , Σ , and V are the matrices produced via the five-step algorithm of the present subsection, given above.

Then,

$$\|A - U \Sigma V^T\| \leq 16 \gamma \beta l \left(\frac{m-k}{l} \right)^{1/(4i+2)} \sigma_{k+1} \quad (70)$$

with probability not less than Π , where Π is defined in (69), and σ_{k+1} is the $(k+1)^{\text{st}}$ greatest singular value of A .

Proof. It is sufficient to prove that

$$\|A Q Q^T - A\| \leq 16 \gamma \beta l \left(\frac{m-k}{l} \right)^{1/(4i+2)} \sigma_{k+1} \quad (71)$$

with probability Π , where Q is the matrix from (63), since combining (71), (64), (65), and (66) yields (70). We now prove (71).

First, we consider the case when

$$\|A\| \leq \left(\frac{m-k}{l} \right)^{1/(4i+2)} \sigma_{k+1}. \quad (72)$$

Clearly,

$$\|A Q Q^T - A\| \leq \|A\| \|Q\| \|Q^T\| + \|A\|. \quad (73)$$

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

$$\|Q\| \leq 1 \quad (74)$$

and

$$\|Q^T\| \leq 1. \quad (75)$$

Combining (73), (74), (75), (72), and (68) yields (71), completing the proof for the case when (72) holds.

For the remainder of the proof, we consider the case when

$$\|A\| > \left(\frac{m-k}{l} \right)^{1/(4i+2)} \sigma_{k+1}. \quad (76)$$

To prove (71), we will use (16), that is,

$$\|A Q Q^T - A\| \leq 2 \|F G (A A^T)^i A - A\| + 2 \|F\| \|Q S - (G (A A^T)^i A)^T\| \quad (77)$$

for any real $m \times l$ matrix F , where G is from (62), and Q and S are from (63). We now choose an appropriate matrix F .

First, we define j to be the positive integer such that

$$\sigma_{j+1} \leq \left(\frac{m-k}{l} \right)^{1/(4i+2)} \sigma_{k+1} < \sigma_j, \quad (78)$$

where σ_j is the j^{th} greatest singular value of A , and σ_{j+1} is the $(j+1)^{\text{st}}$ greatest (such an integer j exists due to (76) and the supposition of the theorem that $l \leq m-k$). We then

use the matrix F from (32) and (33) associated with this integer j , so that (as stated in (32) and (33))

$$\begin{aligned} \|F G (A A^T)^i A - A\| &\leq \sqrt{2l^2 \beta^2 \gamma^2 + 1} \sigma_{j+1} \\ &\quad + \sqrt{2l \max(m-k, l) \beta^2 \gamma^2 \left(\frac{\sigma_{k+1}}{\sigma_j}\right)^{4i} + 1} \sigma_{k+1} \end{aligned} \quad (79)$$

and

$$\|F\| \leq \frac{\sqrt{l} \beta}{(\sigma_j)^{2i}} \quad (80)$$

with probability not less than Φ defined in (31). Formula (79) bounds the first term in the right-hand side of (77).

To bound the second term in the right-hand side of (77), we observe that $j \leq k$, due to (78) and the supposition of the theorem that $l \leq m - k$. Combining (63), (62), (60), and the fact that $j \leq k$ yields that

$$\|Q S - (G (A A^T)^i A)^T\| \leq \sqrt{2l} \gamma (\sigma_{j+1})^{2i+1} + \sqrt{2 \max(m-k, l)} \gamma (\sigma_{k+1})^{2i+1} \quad (81)$$

with probability not less than Ψ defined in (59). Combining (80) and (81) yields that

$$\|F\| \|Q S - (G (A A^T)^i A)^T\| \leq \sqrt{2l^2 \gamma^2 \beta^2} \sigma_{j+1} + \sqrt{2l \max(m-k, l) \gamma^2 \beta^2 \left(\frac{\sigma_{k+1}}{\sigma_j}\right)^{4i}} \sigma_{k+1} \quad (82)$$

with probability not less than Π defined in (69). The combination of Lemma 2.8, (67), and the fact that $j \leq k$ justifies the use of k (rather than the j used in (31) for Φ) in the last term in the right-hand side of (69).

Combining (77), (79), (82), (78), (68), and the supposition of the theorem that $l \leq m - k$ yields (71), completing the proof. \square

Remark 4.2. Choosing $l = k + 12$, $\beta = 2.57$, and $\gamma = 2.43$ in (69) and (70) yields that

$$\|A - U \Sigma V^T\| \leq 100l \left(\frac{m-k}{l}\right)^{1/(4i+2)} \sigma_{k+1} \quad (83)$$

with probability greater than $1 - 10^{-15}$, where σ_{k+1} is the $(k+1)^{\text{st}}$ greatest singular value of A . Numerical experiments (some of which are reported in Section 5) indicate that the factor $100l$ in the right-hand side of (83) is much greater than necessary.

4.2 Computational costs

In this subsection, we tabulate the number of floating-point operations required by the algorithm described in Subsection 4.1 as applied once to a matrix A .

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

1. Forming R in (62) requires applying A to il column vectors, and A^T to $(i+1)l$ column vectors.
2. Computing Q in (63) costs $\mathcal{O}(l^2 n)$.
3. Forming T in (64) requires applying A to k column vectors.
4. Computing the SVD (65) of T costs $\mathcal{O}(k^2 m)$.
5. Forming V in (66) costs $\mathcal{O}(k^2 n)$.

Summing up the costs in Steps 1–5 above, we conclude that the algorithm of Subsection 4.1 costs

$$C_{\text{PCA}} = (il + k) \cdot C_A + (il + l) \cdot C_{A^T} + \mathcal{O}(k^2 m + l^2 n) \quad (84)$$

floating-point operations, 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 4.3. We observe that the algorithm only requires applying A to $il + k$ vectors and A^T to $il + l$ 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.

4.3 A modified algorithm

In this subsection, we describe a simple modification of the algorithm described in Subsection 4.1. Again, suppose that i, k, l, m , and n are positive integers with $k < l \leq m - k$ and $m \leq n$, and A is a real $m \times n$ matrix. Then, the following five-step algorithm constructs an approximation to an SVD of A^T such that

$$\|A^T - U \Sigma V^T\| \leq C m^{1/(4i)} \sigma_{k+1} \quad (85)$$

with very high probability, where U is a real $n \times k$ matrix whose columns are orthonormal, V is a real $m \times k$ matrix whose columns are orthonormal, Σ is a real diagonal $k \times k$ matrix whose entries are all nonnegative, σ_{k+1} is the $(k+1)^{\text{st}}$ greatest singular value of A , and C is a constant independent of A that depends on the parameters of the algorithm:

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 $l \times m$ product matrix

$$R = G(AA^T)^i. \quad (86)$$

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

$$\|QS - R^T\| \leq \rho_{k+1}, \quad (87)$$

where ρ_{k+1} is the $(k+1)^{\text{st}}$ greatest singular value of R . (See Observation 2.4 for details concerning the construction of such a matrix Q .)

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

$$T = A^T Q. \quad (88)$$

4. Form an SVD of T ,

$$T = U \Sigma W^T, \quad (89)$$

where U is a real $n \times k$ matrix whose columns are orthonormal, W is a real $k \times k$ matrix whose columns are orthonormal, and Σ is a real diagonal $k \times k$ matrix whose entries are all nonnegative. (See, for example, Chapter 8 in [5] for details concerning the construction of such an SVD.)

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

$$V = Q W. \quad (90)$$

Clearly, (85) is similar to (61), as (86) is similar to (62).

4.4 Blanczos

In this subsection, we describe a modification of the algorithm of Subsection 4.1, enhancing the accuracy at some additional computational expense. Suppose that i, k, l, m , and n are positive integers with $k < l$ and $(i + 1)l \leq m - k$, and A is a real $m \times n$ matrix, such that $m \leq n$. Then, the following five-step algorithm constructs an approximation $U \Sigma V^T$ to an SVD of A :

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 $l \times n$ matrices $R^{(0)}, R^{(1)}, \dots, R^{(i-1)}, R^{(i)}$ defined via the formulae

$$R^{(0)} = G A, \quad (91)$$

$$R^{(1)} = R^{(0)} A^T A, \quad (92)$$

$$R^{(2)} = R^{(1)} A^T A, \quad (93)$$

\vdots

$$R^{(i-1)} = R^{(i-2)} A^T A, \quad (94)$$

$$R^{(i)} = R^{(i-1)} A^T A. \quad (95)$$

Form the $((i + 1)l) \times n$ matrix

$$R = \begin{pmatrix} R^{(0)} \\ R^{(1)} \\ \vdots \\ R^{(i-1)} \\ R^{(i)} \end{pmatrix}. \quad (96)$$

2. Using a pivoted QR -decomposition algorithm, form a real $n \times ((i+1)l)$ matrix Q whose columns are orthonormal, such that there exists a real $((i+1)l) \times ((i+1)l)$ matrix S for which

$$R^T = QS. \quad (97)$$

(See, for example, Chapter 5 in [5] for details concerning the construction of such a matrix Q .)

3. Compute the $m \times ((i+1)l)$ product matrix

$$T = AQ. \quad (98)$$

4. Form an SVD of T ,

$$T = U \Sigma W^T, \quad (99)$$

where U is a real $m \times ((i+1)l)$ matrix whose columns are orthonormal, W is a real $((i+1)l) \times ((i+1)l)$ matrix whose columns are orthonormal, and Σ is a real diagonal $((i+1)l) \times ((i+1)l)$ matrix whose entries are all nonnegative. (See, for example, Chapter 8 in [5] for details concerning the construction of such an SVD.)

5. Compute the $n \times ((i+1)l)$ product matrix

$$V = QW. \quad (100)$$

It is not difficult to see that the matrices U , Σ , and V produced by the algorithm of the present subsection satisfy the same upper bounds (70) and (83) as the matrices produced by the algorithm of Subsection 4.1. If desired, one may produce a similarly accurate rank- k approximation by arranging U , Σ , and V such that the diagonal entries of Σ appear in nonincreasing order, and then discarding all but the leftmost k columns of U and all but the leftmost k columns of V , and retaining only the leftmost uppermost $k \times k$ block of Σ . We will refer to the algorithm of the present subsection as “blanczos,” due to its similarity with the block Lanczos method (see, for example, Subsection 9.2.6 in [5] for a description of the block Lanczos method).

5 Numerical results

In this section, we illustrate the performance of the algorithm of the present paper via several numerical examples.

We use the algorithm to construct a rank- k approximation, with $k = 10$, to the $m \times (2m)$ matrix A defined via its singular value decomposition

$$A = U^{(A)} \Sigma^{(A)} (V^{(A)})^T, \quad (101)$$

where $U^{(A)}$ is an $m \times m$ Hadamard matrix (a unitary matrix whose entries are all $\pm 1/\sqrt{m}$), $V^{(A)}$ is a $(2m) \times (2m)$ Hadamard matrix, and $\Sigma^{(A)}$ is an $m \times (2m)$ matrix whose entries are

zero off of the main diagonal, and whose diagonal entries are defined in terms of the $(k+1)^{\text{st}}$ singular value σ_{k+1} via the formulae

$$\Sigma_{j,j}^{(A)} = \sigma_j = (\sigma_{k+1})^{\lfloor j/2 \rfloor / 5} \quad (102)$$

for $j = 1, 2, \dots, 9, 10$, where $\lfloor j/2 \rfloor$ is the greatest integer less than or equal to $j/2$, and

$$\Sigma_{j,j}^{(A)} = \sigma_j = \sigma_{k+1} \cdot \frac{m-j}{m-11} \quad (103)$$

for $j = 11, 12, \dots, m-1, m$. Thus, $\sigma_1 = 1$ and $\sigma_k = \sigma_{k+1}$ (recall that $k = 10$). We always choose $\sigma_{k+1} < 1$, so that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{m-1} \geq \sigma_m$.

Figure 1 plots the singular values $\sigma_1, \sigma_2, \dots, \sigma_{m-1}, \sigma_m$ of A with $m = 512$ and $\sigma_{k+1} = .10\%$; these parameters correspond to the first row of numbers in Table 1 and the first row of numbers in Table 2.

Table 1 reports the results of applying the algorithm of Subsection 4.1 to matrices of various sizes, with $i = 1$. Table 2 reports the results of applying the algorithm of Subsection 4.1 to matrices of various sizes, with $i = 0$.

Table 3 reports the results of applying the algorithms of Subsections 4.1 and 4.3 with varying numbers of iterations i . Rows in the table where i is enclosed in parentheses correspond to the algorithm of Subsection 4.3; rows where i is not enclosed in parentheses correspond to the algorithm of Subsection 4.1.

Table 4 reports the results of applying the algorithm of Subsection 4.1 to matrices whose best rank- k approximations have varying accuracies. Table 5 reports the results of applying the blanczos algorithm of Subsection 4.4 to matrices whose best rank- k approximations have varying accuracies.

The headings of the tables have the following meanings:

- m is the number of rows in A , the matrix being approximated.
- n is the number of columns in A , the matrix being approximated.
- i is the integer parameter used in the algorithms of Subsections 4.1, 4.3, and 4.4. Rows in the tables where i is enclosed in parentheses correspond to the algorithm of Subsection 4.3; rows where i is not enclosed in parentheses correspond to either the algorithm of Subsection 4.1 or that of Subsection 4.4.
- t is the time in seconds required by the algorithm to create an approximation and compute its accuracy δ .
- σ_{k+1} is the $(k+1)^{\text{st}}$ greatest singular value of A , the matrix being approximated; σ_{k+1} is also the accuracy of the best possible rank- k approximation to A .
- δ is the accuracy of the approximation $U \Sigma V^T$ to A constructed by the algorithm, that is,

$$\delta = \|A - U \Sigma V^T\|. \quad (104)$$

The values for t are the average values over 3 independent randomized trials of the algorithm. The values for δ are the worst (maximum) values encountered in 3 independent randomized trials of the algorithm. The values for δ in each trial are those produced by 20 iterations of the power method applied to $A - U \Sigma V^T$, started with a vector whose entries are i.i.d. centered Gaussian random variables. The theorems of [3] and [7] guarantee that this power method produces accurate results with overwhelmingly high probability.

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. We implemented a fast Walsh-Hadamard transform to apply rapidly the Hadamard matrices $U^{(A)}$ and $V^{(A)}$ in (101). We used the LAPACK 3.1.1 divide-and-conquer SVD routine `dgesdd` to compute all full SVDs. For the parameter l , we set $l = 12 (= k + 2)$ for all of the examples reported here.

The experiments reported here and our further tests point to the following observations:

1. The accuracies in Table 1 are superior to those in Table 2; the algorithm performs much better with $i > 0$.
2. The accuracies in Tables 1–3 appear to be proportional to $m^{1/(4i+2)} \sigma_{k+1}$ for the algorithm of Subsection 4.1, and to be proportional to $m^{1/(4i)} \sigma_{k+1}$ for the algorithm of Subsection 4.3, in accordance with (61) and (85). The numerical results reported here, as well as our further experiments, indicate that the theoretical bound (70) on the accuracy should remain valid with a greatly reduced constant in the right-hand side, independent of the matrix A being approximated. See item 4 below for a discussion of Tables 4 and 5.
3. The timings in the tables are consistent with (84), as we could (and did) apply the Hadamard matrices $U^{(A)}$ and $V^{(A)}$ in (101) to vectors via fast Walsh-Hadamard transforms at a cost of $\mathcal{O}(m \log(m))$ floating-point operations per matrix-vector multiplication.
4. The accuracies in Table 5 are superior to those in Table 4, particularly when σ_{k+1} is very small. Understandably, the algorithm of Subsection 4.1 would seem to break down when $(\sigma_{k+1})^{2i+1}$ is less than the machine precision, unlike the blanczos algorithm of Subsection 4.4. When $(\sigma_{k+1})^{2i+1}$ is much less than the machine precision, the accuracy of blanczos in the presence of roundoff is similar to that of the algorithm of Subsection 4.1 run with a reduced i . When $(\sigma_{k+1})^{2i+1}$ is much greater than the machine precision, the accuracy of blanczos is similar to that of the algorithm of Subsection 4.1 run with i being the same as in the blanczos algorithm. Since the blanczos algorithm of Subsection 4.4 is so tolerant of roundoff, we suspect that the blanczos algorithm is a better general-purpose black-box tool for the computation of principal component analyses, despite its somewhat higher cost as compared with the algorithms of Subsections 4.1 and 4.3.

Acknowledgements

We would like to thank Ming Gu for suggesting combining the Lanczos method with randomized methods for the low-rank approximation of matrices. We would also like to thank R. Raphael Coifman and Yoel Shkolnisky for many discussions. This work was supported in part by DARPA/AFOSR Grant FA9550-07-1-0541. A.S. wishes to thank the NSF for generous support under DMS-0811203.

Appendix

(The appendix appears in this technical report, but will not appear in the published version.)

A slight modification of the algorithm of Subsection 4.1 seems to admit a tighter bound on its accuracy, though the derivation of this bound is currently incomplete. The modified algorithm has the following six steps (using the same notation as in Subsection 4.1):

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

$$R = G(AA^T)^i A. \quad (105)$$

2. Using a pivoted QR -decomposition algorithm, form a real $n \times l$ matrix Q whose columns are orthonormal, such that there exists a real $l \times l$ matrix S for which

$$R^T = QS. \quad (106)$$

(See, for example, Chapter 5 in [5] for details concerning the construction of such a matrix Q .)

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

$$T = AQ. \quad (107)$$

4. Form an SVD of T ,

$$T = \tilde{U} \tilde{\Sigma} W^T, \quad (108)$$

where \tilde{U} is a real $m \times l$ matrix whose columns are orthonormal, W is a real $l \times l$ matrix whose columns are orthonormal, and $\tilde{\Sigma}$ is a real diagonal $l \times l$ matrix whose only nonzero entries are nonnegative and appear in nonincreasing order on the diagonal. (See, for example, Chapter 8 in [5] for details concerning the construction of such an SVD.)

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

$$\tilde{V} = QW. \quad (109)$$

6. Extract the leftmost $m \times k$ block U of \tilde{U} , the leftmost $n \times k$ block V of \tilde{V} , and the leftmost uppermost $k \times k$ block Σ of $\tilde{\Sigma}$.

If a certain conjecture (described below) is true, then the matrices U , Σ , and V produced by the above six-step algorithm satisfy

$$\|A - U \Sigma V^T\| \leq 8 \sqrt{2 \alpha^2 l \left(\frac{m - k}{l} \right)^{1/(2i+1)} + 1} \sigma_{k+1} \quad (110)$$

with very high probability (the probability depends on $l - k$ and the positive real number α from the conjecture described below), where σ_{k+1} is the $(k + 1)^{\text{st}}$ greatest singular value of A . Please note that the leading factor in the square root in (110) is not $2 \alpha^2 l^2$, but $2 \alpha^2 l$.

A proof of (110) would require the following.

Conjecture. *Suppose that j , l , and m are positive integers with $j < l \leq m$, H is an $l \times j$ matrix whose entries are i.i.d. Gaussian random variables of zero mean and unit variance, and Γ is an $l \times m$ matrix whose entries are i.i.d. Gaussian random variables of zero mean and unit variance. Suppose in addition that α is a positive real number.*

Then,

$$\|H^{(-1)} \Gamma\| \leq \sqrt{m} \alpha \quad (111)$$

with very high probability (the probability depends on $l - j$ and α), where

$$H^{(-1)} = (H^T H)^{-1} H^T. \quad (112)$$

Numerical experiments and heuristic reasoning leave little doubt that this conjecture is true, but our best rigorous estimates to date require an additional factor of \sqrt{l} in the right-hand side of (111).

m	n	i	t	σ_{k+1}	δ
512	1024	1	.13E-01	.10%	.11%
2048	4096	1	.56E-01	.10%	.13%
8192	16384	1	.25E-00	.10%	.18%
32768	65536	1	.12E+01	.10%	.24%
131072	262144	1	.75E+01	.10%	.37%
524288	1048576	1	.36E+02	.10%	.39%

Table 1: Algorithm of Subsection 4.1

m	n	i	t	σ_{k+1}	δ
512	1024	0	.14E-01	.10%	1.2%
2048	4096	0	.47E-01	.10%	2.7%
8192	16384	0	.22E-00	.10%	3.9%
32768	65536	0	.10E+01	.10%	5.3%
131072	262144	0	.60E+01	.10%	11%
524288	1048576	0	.29E+02	.10%	22%

Table 2: Algorithm of Subsection 4.1

m	n	i	t	σ_{k+1}	δ
524288	1048576	0	.29E+02	1.0%	86%
524288	1048576	(1)	.31E+02	1.0%	9.1%
524288	1048576	1	.36E+02	1.0%	3.7%
524288	1048576	(2)	.38E+02	1.0%	2.5%
524288	1048576	2	.43E+02	1.0%	2.2%
524288	1048576	(3)	.45E+02	1.0%	1.5%
524288	1048576	3	.49E+02	1.0%	1.0%

Table 3: Algorithms of Subsections 4.1 and 4.3 (parentheses around i designate Subsection 4.3)

m	n	i	t	σ_{k+1}	δ
262144	524288	1	.17E+02	.10E-02	.39E-02
262144	524288	1	.17E+02	.10E-04	.10E-03
262144	524288	1	.17E+02	.10E-06	.25E-05
262144	524288	1	.17E+02	.10E-08	.90E-06
262144	524288	1	.17E+02	.10E-10	.55E-07
262144	524288	1	.17E+02	.10E-12	.51E-08
262144	524288	1	.17E+02	.10E-14	.10E-05

Table 4: Algorithm of Subsection 4.1

m	n	i	t	σ_{k+1}	δ
262144	524288	1	.31E+02	.10E-02	.35E-02
262144	524288	1	.31E+02	.10E-04	.15E-04
262144	524288	1	.31E+02	.10E-06	.24E-05
262144	524288	1	.31E+02	.10E-08	.11E-06
262144	524288	1	.31E+02	.10E-10	.19E-08
262144	524288	1	.31E+02	.10E-12	.25E-10
262144	524288	1	.31E+02	.10E-14	.53E-11

Table 5: Algorithm of Subsection 4.4




Figure 1: Singular values with $m = 512$, $n = 1024$, and $\sigma_{k+1} = .10\%$

References

- [1] Z. CHEN AND J. J. DONGARRA, *Condition numbers of Gaussian random matrices*, SIAM J. Matrix Anal. Appl., 27 (2005), pp. 603–620.
- [2] S. DEERWESTER, S. T. DUMAIS, G. W. FURNAS, T. K. LANDAUER, AND R. HARSHMAN, *Indexing by latent semantic analysis*, J. Amer. Soc. Inform. Sci., 41 (1990), pp. 391–407.
- [3] DIXON JD, *Estimating extremal eigenvalues and condition numbers of matrices*, SIAM J Numer Anal, 20 (1983), pp. 812–814.
- [4] H. H. GOLDSTINE AND J. VON NEUMANN, *Numerical inverting of matrices of high order, II*, Amer. Math. Soc. Proc., 2 (1951), pp. 188–202.
- [5] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, Johns Hopkins University Press, Baltimore, Maryland, third ed., 1996.
- [6] 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.
- [7] J. KUCZYŃSKI AND H. WOŹNIAKOWSKI, *Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start*, SIAM J Matrix Anal Appl, 13 (1992), pp. 1094–1122.
- [8] E. LIBERTY, F. WOOLFE, P.-G. MARTINSSON, V. ROKHLIN, AND M. TYGERT, *Randomized algorithms for the low-rank approximation of matrices*, Proc. Natl. Acad. Sci. USA, 104 (2007), pp. 20167–20172.
- [9] P.-G. MARTINSSON, V. ROKHLIN, AND M. TYGERT, *A randomized algorithm for the approximation of matrices*, Tech. Rep. 1361, Yale University Department of Computer Science, June 2006.
- [10] S. ROWEIS, *EM algorithms for PCA and SPCA*, in Proceedings of the 1997 Conference on Advances in Neural Information Processing Systems 10, MIT Press, 1998, pp. 626–632.
- [11] T. SARLÓS, *Improved approximation algorithms for large matrices via random projections*, in Proceedings FOCS 2006, New York, October 2006, IEEE Press, pp. 143–152.
- [12] J. H. WILKINSON, *The Algebraic Eigenvalue Problem*, Oxford University Press, Oxford, 1965.