randUTV: A Blocked Randomized Algorithm for Computing a Rank-Revealing UTV Factorization

P. G. MARTINSSON, University of Oxford
G. QUINTANA-ORTÍ, Universitat Jaume I
N. HEAVNER, University of Colorado at Boulder

A randomized algorithm for computing a so-called UTV factorization efficiently is presented. Given a matrix $\mathbf{A}$, the algorithm “randUTV” computes a factorization $\mathbf{A} = \mathbf{UTV}^*$, where $\mathbf{U}$ and $\mathbf{V}$ have orthonormal columns, and $\mathbf{T}$ is triangular (either upper or lower, whichever is preferred). The algorithm randUTV is developed primarily to be a fast and easily parallelized alternative to algorithms for computing the Singular Value Decomposition (SVD). randUTV provides accuracy very close to that of the SVD for problems such as low-rank approximation, solving ill-conditioned linear systems, and determining bases for various subspaces associated with the matrix. Moreover, randUTV produces highly accurate approximations to the singular values of $\mathbf{A}$. Unlike the SVD, the randomized algorithm proposed builds a UTV factorization in an incremental, single-stage, and noniterative way, making it possible to halt the factorization process once a specified tolerance has been met. Numerical experiments comparing the accuracy and speed of randUTV to the SVD are presented. Other experiments also demonstrate that in comparison to column-pivoted QR, which is another factorization that is often used as a relatively economic alternative to the SVD, randUTV compares favorably in terms of speed while providing far higher accuracy.

CCS Concepts: • **Mathematics of computing** → **Mathematical software performance**; **Computations on matrices**;

Additional Key Words and Phrases: Numerical linear algebra, rank-revealing matrix factorization, singular value decomposition, high performance, randomized methods

ACM Reference format:

P. G. Martinsson, G. Quintana-Ortí, and N. Heavner. 2019. randUTV: A Blocked Randomized Algorithm for Computing a Rank-Revealing UTV Factorization. *ACM Trans. Math. Softw.* 45, 1, Article 4 (March 2019), 26 pages.
<https://doi.org/10.1145/3242670>

The research reported was supported by DARPA, under contract N66001-13-1-4050, and by the NSF, under awards DMS-1407340 and DMS-1620472.

Authors' addresses: P.-G. Martinsson, Institute for Computational Engineering and Sciences, The University of Texas at Austin, C0200, POB 4.102, Austin, Texas 78712-1229, USA; email: pgm@ices.utexas.edu; G. Quintana-Ortí, Depto. de Ingeniería y Ciencia de Computadores, Universitat Jaume I, 12.071 Castellón, Spain; email: gquintan@icc.uji.es; N. Heavner, Department of Applied Mathematics, University of Colorado at Boulder, 526 UCB, Boulder, CO 80309-0526, USA; email: nathan.heavner@colorado.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.

0098-3500/2019/03-ART4 \$15.00

<https://doi.org/10.1145/3242670>

1 INTRODUCTION

1.1 Overview

Given an $m \times n$ matrix $\mathbf{A}$, the so-called UTV decomposition [37, p. 400] takes the form

$$\begin{array}{ccccc} \mathbf{A} & = & \mathbf{U} & \mathbf{T} & \mathbf{V}^*, \\ m \times n & & m \times m & m \times n & n \times n \end{array} \quad (1)$$

where $\mathbf{U}$ and $\mathbf{V}$ are unitary matrices, and $\mathbf{T}$ is a triangular matrix (either lower or upper triangular). The UTV decomposition can be viewed as a generalization of other standard factorizations, such as the Singular Value Decomposition (SVD) or the Column-Pivoted QR (CPQR) decomposition. (To be precise, the SVD is the special case where $\mathbf{T}$ is diagonal, and the CPQR is the special case where $\mathbf{V}$ is a permutation matrix.) The additional flexibility inherent in the UTV decomposition enables the design of efficient updating procedures (see [37, Ch. 5, Sec. 4] and [13, 31, 36]). In this article, we describe a randomized algorithm that we call *randUTV*, which exploits the additional flexibility provided by the UTV format to build a factorization algorithm that combines some of the most desirable properties of standard algorithms for computing the SVD and CPQR factorizations.

1.2 Where randUTV Fits in the Pantheon of Matrix Factorization Algorithms

The algorithm we describe is designed primarily to serve as a competitive alternative to the CPQR factorization for applications such as detecting the numerical rank of a matrix, solving ill-conditioned linear systems, and computing least-squares solutions to linear systems with rectangular or numerically rank-deficient coefficient matrices. The optimal factorization to use for many of these tasks is the SVD, but the CPQR remains popular because it is faster to compute and has a crucial advantage in that the factorization can be halted once a specified tolerance has been met.

We argue that *randUTV* outperforms CPQR in almost every regard. It does a much better job than CPQR of revealing the rank structure of the underlying matrix, which makes it a better tool for revealing the numerical rank, and for solving ill-conditioned linear systems. It is often faster to compute than the CPQR and parallelizes better as the number of cores is increased. Finally, it shares the key advantage of the CPQR that the factorization can be halted once a specified tolerance has been met.

The reason for the computational efficiency of *randUTV* is that it can (unlike the CPQR) be *blocked*, and can be implemented so that most flops are spent in BLAS3 operations rather than in BLAS2 operations. We will demonstrate that this leads to high performance on a single core and excellent performance as the number of cores is increased. We provide detailed numerical experiments in Section 6.1, but the key findings are summarized in Figure 1. We see that *randUTV* is much faster than state-of-the-art functions for computing the SVD, and that it is often faster than the CPQR, particularly as the number of cores increases. Since *randUTV* can be blocked, it is also very well suited for implementation on distributed memory machines, for matrices stored out-of-core, and so forth.

Remark 1 (randUTV as an Alternative to the SVD). The method we propose is sufficiently accurate and can sometimes be used as an economical method for computing an approximate SVD. The numerical experiments in Section 6.2 indicate that the matrix $\mathbf{T}$ often approximates the diagonal matrix $\mathbf{D}$ in the SVD to two or three digits of accuracy or more. This level of accuracy is often sufficient for purposes of data analysis, and we believe that *randUTV* provides an excellent alternative to the SVD for applications such as Principal Component Analysis (PCA). However, there of course remain many applications where computing the SVD using existing methodologies remains the superior choice, including situations where the actual singular values are required to high accuracy, where the singular values of the matrix decay very slowly, or simply situations where the matrix is small enough that computing the full SVD is very fast.

Fig. 1. Solid lines show the acceleration of randUTV with the parameter choice $q = 1$ (cf. Section 4.3) compared to Intel’s MKL SVD (dgesdd). Dashed lines show the analogous acceleration of CPQR over SVD.

Remark 2 (Where CPQR is Better than randUTV). Our numerical experiments indicate that for small matrices, our implementation of randUTV is at the current level of optimization slower than the best implementations of LAPACK’s dgeqp3 function [33] (see also Figure 1). Moreover, the CPQR provides important information about how to pick subsets of the columns of a matrix to use as a basis for its column space. This makes the CPQR an important component in algorithms for computing co-called interpolatory and CUR decompositions [25, 44].

1.3 A Randomized Algorithm for Computing the UTV Decomposition

The algorithm we propose is blocked for computational efficiency. For concreteness, let us assume that $m \geq n$, and that an *upper* triangular factor $\mathbf{T}$ is sought. With b a block size, randUTV proceeds through $\lceil n/b \rceil$ steps, where at the i ’th step the i ’th block of b columns of $\mathbf{A}$ is driven to upper triangular form, as illustrated later in Figure 2.

In the first iteration, randUTV uses a randomized subspace iteration inspired by Halko et al. [19] and Rokhlin et al. [34] to build two sets of b orthonormal vectors that approximately span the spaces spanned by the b dominant left and right singular vectors of $\mathbf{A}$, respectively. These basis vectors form the first b columns of two unitary “transformation matrices,” $\mathbf{U}^{(1)}$ and $\mathbf{V}^{(1)}$. We use these to build a new matrix

$$\mathbf{A}^{(1)} = (\mathbf{U}^{(1)})^* \mathbf{A} \mathbf{V}^{(1)}$$

that has a blocked structure as follows:

$$\mathbf{A}^{(1)} = \begin{bmatrix} \mathbf{A}_{11}^{(1)} & \mathbf{A}_{12}^{(1)} \\ \mathbf{0} & \mathbf{A}_{22}^{(1)} \end{bmatrix} = \begin{bmatrix} \text{diagonal} & \text{small} \\ \text{zero} & \text{large} \end{bmatrix}.$$

In other words, the top left $b \times b$ block is diagonal, and the bottom left block is zero. Ideally, we would want the top right block to also be zero, and this is what we would obtain if we used the exact left and right singular vectors in the transformation. The randomized sampling does not exactly identify the correct subspaces, but it does to high enough accuracy that the elements in the top right block have very small moduli. For purposes of low-rank approximation, such a format strikes an attractive balance between computational efficiency and close to optimal accuracy. We

will demonstrate that $\|\mathbf{A}_{22}^{(1)}\| \approx \inf\{\|\mathbf{A} - \mathbf{B}\| : \mathbf{B} \text{ has rank } b\}$, and that the diagonal entries of $\mathbf{A}_{11}^{(1)}$ form accurate approximations to the first b singular values of $\mathbf{A}$.

Once the first step has been completed, the second step applies the same procedure to the remaining block $\mathbf{A}_{22}^{(1)}$, which has size $(m - b) \times (n - b)$, and then continues in the same fashion to process all remaining blocks.

1.4 Relationship to Earlier Work

The UTV factorization was introduced and popularized by G. W. Stewart in a sequence of works, including references [36, 38–40] and the textbook chapter [37, p. 400]. Among these, one work [38] is of particular relevance, as it discusses explicitly how the UTV decomposition can be used for low-rank approximation and for finding approximations to the singular values of a matrix; in Section 6.2, we compare the accuracy of the method in Reference [38] to the accuracy of randUTV. Of relevance here is another work [30], which describes deterministic iterative methods for driving a triangular matrix toward diagonality. Another path toward better rank revelation and more accurate rank estimation is described in Reference [12]. A key advantage of the UTV factorization is the ease with which it can be updated (e.g., as discussed in References [3, 4, 31]). Implementation aspects are discussed in Reference [13].

The work presented here relies crucially on previous work on randomized subspace iteration for approximating the linear spaces spanned by groups of singular vectors, including References [18, 19, 28, 29, 34]. This prior body of literature focused on the task of computing *partial* factorizations of matrices. More recently, it was observed [11, 23, 24] that analogous techniques can be utilized to accelerate methods for computing *full* factorizations, such as the CPQR. The gain in speed is attained from *blocking* of the algorithms and the use of BLAS3 rather than by reducing the asymptotic flop count. Our work is also related to a randomized algorithm for computing a UTV decomposition described in Section 5 of Demmel et al. [8]; that algorithm relies on the observation that if an $n \times n$ matrix $\mathbf{V}$ is drawn from a Haar distribution, then the unpivoted QR factorization of $\mathbf{AV}$ is with high probability “rank revealing” in a certain sense. This algorithm is conceptually similar to randUTV for the special case where no oversampling is done, and no steps of power iteration are taken. This means that the algorithm of Demmel et al. [8] does not provide a path to improve on a factorization beyond what the most basic randomized sampling produces, which limits its usefulness in many practical applications. (Numerical results in Section 6.2 illustrate that the use of power iteration is often imperative.)

1.5 Outline of the Article

Section 2 introduces notation and lists some relevant existing results that we need. Section 3 provides a high-level description of the proposed algorithm. Section 4 describes in detail how to drive one block of columns to upper triangular form, which forms one step of the blocked algorithm. Section 5 describes the whole multistep algorithm, discusses some implementation issues, and provides an estimate of the asymptotic flop count. Section 6 gives the results of several numerical experiments that illustrate the speed and accuracy of the proposed method. Section 7 describes availability of software.

2 PRELIMINARIES

This section introduces our notation and reviews some established techniques that will be needed. The material in Sections 2.1 through 2.4 is described in any standard text on numerical linear algebra (e.g., [14, 37, 41]). The material in Section 2.5 on randomized algorithms is described in further detail in the survey Halkin et al. [19] and the lecture notes in Martinsson [27].

2.1 Basic Notation

Throughout this article, we measure vectors in $\mathbb{R}^n$ using their Euclidean norm. The default norm for matrices will be the corresponding operator norm $\|\mathbf{A}\| = \sup\{\|\mathbf{A}\mathbf{x}\| : \|\mathbf{x}\| = 1\}$, although we will sometimes also use the Frobenius norm $\|\mathbf{A}\|_{\text{Fro}} = \left(\sum_{i,j} |\mathbf{A}(i,j)|^2\right)^{1/2}$. We use the notation of Golub and Van Loan [14] to specify submatrices: If $\mathbf{B}$ is an $m \times n$ matrix, and $I = [i_1, i_2, \dots, i_k]$ and $J = [j_1, j_2, \dots, j_\ell]$ are index vectors, then $\mathbf{B}(I, J)$ denotes the corresponding $k \times \ell$ submatrix. We let $\mathbf{B}(I, :)$ denote the matrix $\mathbf{B}(I, [1, 2, \dots, n])$, and define $\mathbf{B}(:, J)$ analogously. $\mathbf{I}_n$ denotes the $n \times n$ identity matrix, and $\mathbf{0}_{m,n}$ is the $m \times n$ zero matrix. The transpose of $\mathbf{B}$ is denoted by $\mathbf{B}^*$, and we say that an $m \times n$ matrix $\mathbf{U}$ is *orthonormal* if its columns are orthonormal so that $\mathbf{U}^* \mathbf{U} = \mathbf{I}_n$. A square orthonormal matrix is said to be *unitary*.

2.2 Singular Value Decomposition

Let $\mathbf{A}$ be an $m \times n$ matrix, and set $r = \min(m, n)$. Then, the SVD of $\mathbf{A}$ takes the form

$$\begin{array}{ccccc} \mathbf{A} & = & \mathbf{U} & \mathbf{D} & \mathbf{V}^*, \\ m \times n & & m \times r & r \times r & r \times n \end{array} \quad (2)$$

where matrices $\mathbf{U}$ and $\mathbf{V}$ are orthonormal, and $\mathbf{D}$ is diagonal. We have $\mathbf{U} = [\mathbf{u}_1 \ \mathbf{u}_2 \ \dots \ \mathbf{u}_r]$, $\mathbf{V} = [\mathbf{v}_1 \ \mathbf{v}_2 \ \dots \ \mathbf{v}_r]$, and $\mathbf{D} = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_r)$, where $\{\mathbf{u}_j\}_{j=1}^r$ and $\{\mathbf{v}_j\}_{j=1}^r$ are the left and right singular vectors of $\mathbf{A}$, respectively, and $\{\sigma_j\}_{j=1}^r$ are the singular values of $\mathbf{A}$. The singular values are ordered so that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r \geq 0$.

A principal advantage of the SVD is that it furnishes an explicit solution to the low-rank approximation problem. To be precise, for $k = 1, 2, \dots, r$, let us define the rank- k matrix

$$\mathbf{A}_k = \mathbf{U}(:, 1 : k) \mathbf{D}(1 : k, 1 : k) \mathbf{V}(:, 1 : k)^* = \sum_{j=1}^k \sigma_j \mathbf{u}_j \mathbf{v}_j^*.$$

Then, the Eckart-Young theorem asserts that

$$\|\mathbf{A} - \mathbf{A}_k\| = \inf\{\|\mathbf{A} - \mathbf{B}\| : \mathbf{B} \text{ has rank } k\}.$$

A disadvantage of the SVD is that the cost to compute the full SVD is $O(n^3)$ for a square matrix, and $O(mnr)$ for a rectangular matrix, with large prefactors. Moreover, standard techniques for computing the SVD are challenging to parallelize and cannot readily be modified to compute partial factorizations.

2.3 CPQR Decomposition

Let $\mathbf{A}$ be an $m \times n$ matrix, and set $r = \min(m, n)$. Then, the CPQR decomposition of $\mathbf{A}$ takes the form

$$\begin{array}{ccccc} \mathbf{A} & = & \mathbf{Q} & \mathbf{R} & \mathbf{P}^*, \\ m \times n & & m \times r & r \times n & n \times n \end{array} \quad (3)$$

where $\mathbf{Q}$ is orthonormal, $\mathbf{R}$ is upper triangular, and $\mathbf{P}$ is a permutation matrix. The permutation matrix $\mathbf{P}$ is typically chosen to ensure monotonic decay in magnitude of the diagonal entries of $\mathbf{R}$ so that $|\mathbf{R}(1, 1)| \geq |\mathbf{R}(2, 2)| \geq |\mathbf{R}(3, 3)| \geq \dots$. The factorization (3) is commonly computed using the *Householder QR algorithm* [14, Sec. 5.2], which is exceptionally stable.

An advantage of the CPQR is that it is computed via an incremental algorithm that can be halted to produce a partial factorization once any given tolerance has been met. A disadvantage is that it is not ideal for low-rank approximation. In the typical case, the error incurred is noticeably worse than what you get from the SVD but usually not disastrously so. However, there exist matrices

for which CPQR leads to very suboptimal approximation errors [21]. Specialized pivoting strategies have been developed that can in some circumstances improve the low-rank approximation property, resulting in so-called rank-revealing QR factorizations [7, 16].

The classical column-pivoted Householder QR factorization algorithm drives the matrix $\mathbf{A}$ to upper triangular form via a sequence of $r - 1$ rank-one updates and column pivotings. Due to the column pivoting performed after each update, it is difficult to *block*, making it hard to achieve high computational efficiency on modern processors [10]. It has recently been demonstrated that randomized methods can be used to resolve this difficulty [11, 23, 24]. However, we do not yet have rigorous theory backing up such randomized techniques, and they have not yet been incorporated into standard software packages.

2.4 Efficient Factorizations of Tall and Thin Matrices

The algorithms proposed in this article rely crucially on the fact that factorizations of “tall and thin” matrices can be computed efficiently. To be precise, suppose that we are given a matrix $\mathbf{B}$ of size $m \times b$, where $m \gg b$, and that we seek to compute an *unpivoted* QR factorization

$$\begin{array}{ccc} \mathbf{B} & = & \mathbf{Q} \quad \mathbf{R}, \\ m \times b & & m \times m \quad m \times b \end{array} \quad (4)$$

where $\mathbf{Q}$ is unitary and $\mathbf{R}$ is upper triangular. When the Householder QR factorization procedure is used to compute the factorization (4), the matrix $\mathbf{Q}$ is formed as a product of b so-called Householder reflectors [14, Sec. 5.2] and can be written in the form

$$\begin{array}{ccccc} \mathbf{Q} & = & \mathbf{I} & + & \mathbf{W} \quad \mathbf{Y}, \\ m \times m & & m \times m & & m \times b \quad b \times m \end{array} \quad (5)$$

for some matrices $\mathbf{W}$ and $\mathbf{Y}$ that can be readily computed given the b Householder vectors that are formed in the QR factorization procedure [6, 20, 35]. As a consequence, we need only $O(mb)$ elements to store $\mathbf{Q}$, and we can apply $\mathbf{Q}$ to an $m \times n$ matrix using $\sim 2mbn$ flops. In this article, the different typeface in $\mathbf{Q}$ is used as a reminder that this is a matrix that can be stored and applied efficiently, without being explicitly built.

Next, suppose that we seek to compute the SVD of the tall and thin matrix $\mathbf{B}$. This can be done efficiently in a two-step procedure, where the first step is to compute the unpivoted QR factorization (4). Let $\mathbf{R}_{\text{small}}$ denote the top $b \times b$ block of $\mathbf{R}$ so that

$$\mathbf{R} = \begin{bmatrix} \mathbf{R}_{\text{small}} \\ \mathbf{0}_{m-b,b} \end{bmatrix}. \quad (6)$$

Then, compute the SVD of $\mathbf{R}_{\text{small}}$ to obtain the factorization

$$\begin{array}{ccccc} \mathbf{R}_{\text{small}} & = & \mathbf{U}_{\text{small}} & \mathbf{D}_{\text{small}} & \mathbf{V}_{\text{small}}^*, \\ b \times b & & b \times b & b \times b & b \times b \end{array} \quad (7)$$

Combining (4), (6), and (7), we obtain the factorization

$$\begin{array}{ccccc} \mathbf{B} & = & \mathbf{Q} & \begin{bmatrix} \mathbf{U}_{\text{small}} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{m-b} \end{bmatrix} & \begin{bmatrix} \mathbf{D}_{\text{small}} \\ \mathbf{0}_{m-b,b} \end{bmatrix} & \mathbf{V}_{\text{small}}^*, \\ m \times b & & m \times m & m \times m & m \times b & b \times b \end{array} \quad (8)$$

which we recognize as an SVD of $\mathbf{B}$. Observe that the cost of computing this factorization is $O(mb^2)$, and that only $O(mb)$ stored elements are required, despite the fact that the matrix of left singular vectors is ostensibly an $m \times m$ dense matrix.

Remark 3. We mentioned in Section 2.3 that it is challenging to achieve high performance when implementing CPQR on modern processors. In contrast, *unpivoted* QR is highly efficient because this algorithm can readily be blocked (e.g., see Figures 4.1 and 4.2 in Martinsson et al. [24]).

2.5 Randomized Power Iterations

This section summarizes key results of references [19, 27, 34] on randomized algorithms for constructing sets of orthonormal vectors that approximately span the row or column spaces of a given matrix. To be precise, let $\mathbf{A}$ be an $m \times n$ matrix, let b be an integer such that $1 \leq b < \min(m, n)$, and suppose that we seek to find an $n \times b$ orthonormal matrix $\mathbf{Q}$ such that

$$\|\mathbf{A} - \mathbf{A}\mathbf{Q}\mathbf{Q}^*\| \approx \inf\{\|\mathbf{A} - \mathbf{B}\| : \mathbf{B} \text{ has rank } b\}.$$

Informally, the columns of $\mathbf{Q}$ should approximately span the same space as the dominant b right singular vectors of $\mathbf{A}$. This is a task that is very well suited for subspace iteration (see Demmel [9, Sec. 4.4.3] and Bathe [5]), particularly when the starting matrix is a Gaussian random matrix [19, 34]. With q a (small) integer denoting the number of steps taken in the power iteration, the following algorithm leads to close to optimal results:

- (1) Draw a Gaussian random matrix $\mathbf{G}$ of size $m \times b$.
- (2) Form a “sampling matrix” $\mathbf{Y}$ of size $n \times b$ via $\mathbf{Y} = (\mathbf{A}^* \mathbf{A})^q \mathbf{A}^* \mathbf{G}$.
- (3) Orthonormalize the columns of $\mathbf{Y}$ to form the matrix $\mathbf{Q}$.

Observe that in step (2), the matrix $\mathbf{Y}$ is computed via alternating application of $\mathbf{A}^*$ and $\mathbf{A}$ to a tall thin matrix with b columns. In some situations, orthonormalization is required between applications to avoid loss of accuracy due to floating point arithmetic. In Halko et al. [19] and Rokhlin et al. [34], it is demonstrated that by using a Gaussian random matrix as the starting point, it is often sufficient to take just a few steps of power iteration, say $q = 1$ or 2 , or even $q = 0$.

Remark 4 (Oversampling). In the analysis of power iteration with a Gaussian random matrix as the starting point, it is common to draw a few “extra” samples. In other words, one picks a small integer p representing the amount of oversampling done, say $p = 5$ or $p = 10$, and starts with a Gaussian matrix of size $m \times (b + p)$. This results in an orthonormal matrix $\mathbf{Q}$ of size $n \times (b + p)$. Then, with probability almost 1, the error $\|\mathbf{A} - \mathbf{A}\mathbf{Q}\mathbf{Q}^*\|$ is close to the minimal error in rank- b approximation in both spectral and Frobenius norm [19, Sec. 10]. When no oversampling is used, one risks losing some accuracy in the last couple of modes of the SVD. However, our experience shows that in the context of the present article, this loss is hardly noticeable.

Remark 5 (RSVD). The randomized range finder described in this section is simply the first stage in the two-stage Randomized SVD (RSVD) procedure for computing an approximate rank- b SVD of a given matrix $\mathbf{A}$ of size $m \times n$. To wit, suppose that we have used the randomized range finder to build an orthonormal matrix $\mathbf{Q}$ of size $n \times b$ such that $\mathbf{A} = \mathbf{A}\mathbf{Q}\mathbf{Q}^* + \mathbf{E}$, for some error matrix $\mathbf{E}$. Then, we can compute an approximate factorization

$$\begin{array}{ccccc} \mathbf{A} & = & \mathbf{U} & \mathbf{D} & \mathbf{V}^* + \mathbf{E}, \\ m \times n & & m \times b & b \times b & b \times n \quad m \times n \end{array} \quad (9)$$

where $\mathbf{U}$ and $\mathbf{V}$ are orthonormal, and $\mathbf{D}$ is diagonal, via the following steps (which together form the “second stage” of RSVD). First, set $\mathbf{B} = \mathbf{A}\mathbf{Q}$ so that $\mathbf{A}\mathbf{Q}\mathbf{Q}^* = \mathbf{B}\mathbf{Q}^*$. Second, compute a full SVD of the small matrix $\mathbf{B}$ so that

$$\begin{array}{ccccc} \mathbf{B} & = & \mathbf{U} & \mathbf{D} & \hat{\mathbf{V}}^*, \\ m \times b & & m \times b & b \times b & b \times b \end{array} \quad (10)$$

Fig. 2. Cartoon illustrating the process by which a given matrix $\mathbf{A}$ is driven to block upper triangular form, with most of the mass concentrated into the diagonal entries. With $\mathbf{A}^{(0)} = \mathbf{A}$, we form a sequence of matrices $\mathbf{A}^{(i)} = (\mathbf{U}^{(i)})^* \mathbf{A}^{(i-1)} \mathbf{V}^{(i)}$ by applying unitary matrices $\mathbf{U}^{(i)}$ and $\mathbf{V}^{(i)}$ from the left and the right. Each $\mathbf{U}^{(i)}$ and $\mathbf{V}^{(i)}$ consists predominantly of a product of b Householder reflectors, where b is the block size. Black blocks represent nonzero entries. Grey blocks represent entries that are not necessarily zero but are small in magnitude.

Third, set $\mathbf{V} = \mathbf{Q}\hat{\mathbf{V}}$. Observe that these last three steps are exact up to floating point arithmetic, so the error in (9) is exactly the same as the error in the range finder alone: $\mathbf{E} = \mathbf{A} - \mathbf{A}\mathbf{Q}\mathbf{Q}^* = \mathbf{A} - \mathbf{U}\mathbf{D}\mathbf{V}^*$.

3 AN OVERVIEW OF THE RANDOMIZED UTV FACTORIZATION ALGORITHM

This section describes at a high level the overall template of the algorithm randUTV that given an $m \times n$ matrix $\mathbf{A}$ computes its UTV decomposition (1). For simplicity, we assume that $m \geq n$, that an upper triangular middle factor $\mathbf{T}$ is sought, and that we work with a block size b that evenly divides n so that the matrix $\mathbf{A}$ can be partitioned into s blocks of b columns each; in other words, we assume that $n = sb$. The algorithm randUTV iterates over s steps, where at the i 'th step the i 'th block of b columns is driven to upper triangular form via the application of unitary transformations from the left and the right. A cartoon of the process is given later in Figure 2.

To be slightly more precise, we build at the i 'th step unitary matrices $\mathbf{U}^{(i)}$ and $\mathbf{V}^{(i)}$ of sizes $m \times m$ and $n \times n$, respectively, such that

$$\mathbf{U} = \mathbf{U}^{(1)} \mathbf{U}^{(2)} \dots \mathbf{U}^{(s)}, \quad \text{and} \quad \mathbf{V} = \mathbf{V}^{(1)} \mathbf{V}^{(2)} \dots \mathbf{V}^{(s)}.$$

Using these matrices, we drive $\mathbf{A}$ toward upper triangular form through a sequence of transformations

$$\begin{aligned}
\mathbf{A}^{(0)} &= \mathbf{A}, \\
\mathbf{A}^{(i)} &= \left(\mathbf{U}^{(i)}\right)^* \mathbf{A}^{(i-1)} \mathbf{V}^{(i)}, \quad i = 1, 2, 3, \dots, s, \\
\mathbf{T} &= \mathbf{A}^{(s)}.
\end{aligned}$$

The objective at step i is to transform the i 'th diagonal block to diagonal form, to zero out all blocks beneath the i 'th diagonal block, and to make all blocks to the right of the i 'th diagonal block as small in magnitude as possible.

Each matrix $\mathbf{U}^{(i)}$ and $\mathbf{V}^{(i)}$ consists predominantly of a product of b Householder reflectors. To be precise, each such matrix is a product of b Householder reflectors but with the i 'th block of b columns right-multiplied by a small $b \times b$ unitary matrix.

The next two sections provide additional details. Section 4 describes exactly how to build the transformation matrices $\mathbf{U}^{(1)}$ and $\mathbf{V}^{(1)}$ that are needed in the first step of the iteration. Then, Section 5 shows how to apply the techniques described in Section 4 repeatedly to build the full factorization.

4 A RANDOMIZED ALGORITHM FOR FINDING A PAIR OF TRANSFORMATION MATRICES FOR THE FIRST STEP

4.1 Objectives for the Construction

In this section, we describe a randomized algorithm for finding unitary matrices $\mathbf{U}$ and $\mathbf{V}$ that execute the first step of the process outlined in Section 3 and illustrated in Figure 2. To avoid notational clutter, we simplify the notation slightly, describing what we consider a basic single step of the process. Given an $m \times n$ matrix $\mathbf{A}$ and a block size b , we seek two orthonormal matrices $\mathbf{U}$ and $\mathbf{V}$ of sizes $m \times m$ and $n \times n$, respectively, such that the matrix

$$\mathbf{T} = \mathbf{U}^* \mathbf{A} \mathbf{V} \quad (11)$$

has a diagonal leading $b \times b$ block, and the entries beneath this block are all zeroed out. In Section 3, we referred to the matrices $\mathbf{U}$ and $\mathbf{V}$ as $\mathbf{U}^{(1)}$ and $\mathbf{V}^{(1)}$, respectively, and $\mathbf{T}$ as $\mathbf{A}^{(1)}$.

To make the discussion precise, let us partition $\mathbf{U}$ and $\mathbf{V}$ so that

$$\mathbf{U} = [\mathbf{U}_1 \mid \mathbf{U}_2], \quad \text{and} \quad \mathbf{V} = [\mathbf{V}_1 \mid \mathbf{V}_2], \quad (12)$$

where $\mathbf{U}_1$ and $\mathbf{V}_1$ each contain b columns. Then, set $\mathbf{T}_{ij} = \mathbf{U}_i^* \mathbf{A} \mathbf{V}_j$ for $i, j = 1, 2$ so that

$$\mathbf{U}^* \mathbf{A} \mathbf{V} = \begin{bmatrix} \mathbf{T}_{11} & \mathbf{T}_{12} \\ \mathbf{T}_{21} & \mathbf{T}_{22} \end{bmatrix}. \quad (13)$$

Our objective is now to build matrices $\mathbf{U}$ and $\mathbf{V}$ that accomplish the following:

- (i) $\mathbf{T}_{11}$ is diagonal, with entries that closely approximate the leading b singular values of $\mathbf{A}$.
- (ii) $\mathbf{T}_{21} = \mathbf{0}$.
- (iii) $\mathbf{T}_{12}$ has small magnitude.
- (iv) The norm of $\mathbf{T}_{22}$ should be close to optimally small so that $\|\mathbf{T}_{22}\| \approx \inf\{\|\mathbf{A} - \mathbf{C}\| : \mathbf{C} \text{ has rank } b\}$.

The purpose of condition (iv) is to minimize the error in the rank- b approximation to $\mathbf{A}$ (cf. Section 5.3).

4.2 A Theoretically Ideal Choice of Transformation Matrices

Suppose that we could somehow find two matrices $\mathbf{U}$ and $\mathbf{V}$ whose first b columns *exactly* span the subspaces spanned by the dominant left and right singular vectors, respectively. Finding such matrices is of course computationally hard, but *if* we could build them, then we would get the optimal result that

- (ii) $\mathbf{T}_{21} = \mathbf{0}$.
- (iii) $\mathbf{T}_{12} = \mathbf{0}$.
- (iv) $\|\mathbf{T}_{22}\| = \sigma_{b+1} = \min\{\|\mathbf{A} - \mathbf{C}\| : \text{matrix } \mathbf{C} \text{ has rank } b\}$.

Enforcing condition (i) is then very easy, as the dominant $b \times b$ block is now disconnected from the rest of the matrix. Simply executing a full SVD of this small $b \times b$ block and then updating $\mathbf{U}$ and $\mathbf{V}$ will do the job.

4.3 A Randomized Technique for Approximating the Span of the Dominant Singular Vectors

Inspired by the observation in Section 4.2 that a theoretically ideal right transformation matrix $\mathbf{V}$ is a matrix whose b first columns span the space spanned by the dominant b right singular vectors of $\mathbf{A}$, we use the randomized power iteration described in Section 2.5 to execute this task. We let q denote a small integer specifying how many steps of power iteration we take. Typically, $q = 0, 1, 2$ are good choices. Then, take the following steps. First, draw a Gaussian random matrix $\mathbf{G}$ of size $m \times b$. Second, compute a sampling matrix $\mathbf{Y} = (\mathbf{A}^* \mathbf{A})^q \mathbf{A}^* \mathbf{G}$ of size $n \times b$. Third, perform an unpivoted Householder QR factorization of $\mathbf{Y}$ so that

$$\begin{array}{ccc} \mathbf{Y} & = & \mathbf{V} \quad \mathbf{R} \\ n \times b & & n \times n \quad n \times b \end{array}$$

Observe that $\mathbf{V}$ will be a product of b Householder reflectors, and that the first b columns of $\mathbf{V}$ will form an orthonormal basis for the space spanned by the columns of $\mathbf{Y}$. In consequence, the first b columns of $\mathbf{V}$ form an orthonormal basis for a space that approximately spans the space spanned by the b dominant right singular vectors of $\mathbf{A}$. (The font used for $\mathbf{V}$ is a reminder that it is a product of Householder reflectors, which is exploited when it is stored and operated on (cf. Section 2.4).)

4.4 Construction of the Left Transformation Matrix

The process for finding the left transformation matrix $\mathbf{U}$ is deterministic and will exactly transform the first b columns of $\mathbf{A}\mathbf{V}$ into a diagonal matrix. Observe that with $\mathbf{V}$, the unitary matrix constructed using the procedure in Section 4.3, we have the identity

$$\mathbf{A} = \mathbf{A}\mathbf{V}\mathbf{V}^* = [\mathbf{A}\mathbf{V}_1 \mid \mathbf{A}\mathbf{V}_2] \mathbf{V}^*, \quad (14)$$

where the partitioning $\mathbf{V} = [\mathbf{V}_1 \mid \mathbf{V}_2]$ is such that $\mathbf{V}_1$ holds the first b columns of $\mathbf{V}$. We now perform a full SVD on the matrix $\mathbf{A}\mathbf{V}_1$, which is of size $m \times b$ so that

$$\begin{array}{ccc} \mathbf{A}\mathbf{V}_1 & = & \mathbf{U} \quad \mathbf{D} \quad \mathbf{V}_{\text{small}}^* \\ m \times b & & m \times m \quad m \times b \quad b \times b \end{array} \quad (15)$$

Inserting (15) into (14), we obtain the identity

$$\mathbf{A} = [\mathbf{U}\mathbf{D}\mathbf{V}_{\text{small}}^* \mid \mathbf{A}\mathbf{V}_2] \mathbf{V}^*. \quad (16)$$

Factor out $\mathbf{U}$ in (16) to the left to get

$$\mathbf{A} = \mathbf{U} [\mathbf{D}\mathbf{V}_{\text{small}}^* \mid \mathbf{U}^* \mathbf{A}\mathbf{V}_2] \mathbf{V}^*.$$

Finally, factor out $\mathbf{V}_{\text{small}}$ to the right to yield the factorization

$$\mathbf{A} = \mathbf{U} \underbrace{[\mathbf{D} \mid \mathbf{U}^* \mathbf{A}\mathbf{V}_2]}_{=\mathbf{T}} \mathbf{V}^*, \quad \text{with } \mathbf{V} = \mathbf{V} \begin{bmatrix} \mathbf{V}_{\text{small}} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{n-b} \end{bmatrix}. \quad (17)$$

Equation (17) is the factorization $\mathbf{A} = \mathbf{UTV}^*$ that we seek, with

$$\mathbf{T} = \begin{bmatrix} \mathbf{T}_{11} & \mathbf{T}_{12} \\ \mathbf{T}_{21} & \mathbf{T}_{22} \end{bmatrix} = \begin{bmatrix} \mathbf{D}(1:b, 1:b) & \mathbf{U}_1^* \mathbf{A} \mathbf{V}_2 \\ \mathbf{0} & \mathbf{U}_2^* \mathbf{A} \mathbf{V}_2 \end{bmatrix}.$$

Remark 6. As we saw in (17), the right transformation matrix $\mathbf{V}$ consists of a product $\mathbb{V}$ of b Householder reflectors, with the first b columns rotated by a small unitary matrix $\mathbf{V}_{\text{small}}$. We will next demonstrate that the left transformation matrix $\mathbf{U}$ can be built in such a way that it can be written in an entirely analogous form. Simply observe that the matrix $\mathbf{A} \mathbf{V}_1$ in (15) is a tall thin matrix so that the SVD in (15) can efficiently be computed by first performing an unpivoted QR factorization of $\mathbf{A} \mathbf{V}_1$ to yield a factorization

$$\mathbf{A} \mathbf{V}_1 = \begin{matrix} & \mathbf{U} & \\ m \times b & m \times m & m \times b \end{matrix} \begin{bmatrix} \mathbf{R}_{11} \\ \mathbf{0} \end{bmatrix},$$

where $\mathbf{R}_{11}$ is of size $b \times b$, and $\mathbf{U}$ is a product of b Householder reflectors (cf. Section 2.4). Then, perform an SVD of $\mathbf{R}_{11}$ to obtain

$$\mathbf{R}_{11} = \begin{matrix} \mathbf{U}_{\text{small}} & \mathbf{D}_{\text{small}} & \mathbf{V}_{\text{small}}^* \\ b \times b & b \times b & b \times b \end{matrix}.$$

The factorization (15) then becomes

$$\mathbf{A} \mathbf{V}_1 = \begin{matrix} & \mathbf{U} & \\ m \times b & m \times m & m \times b \end{matrix} \begin{bmatrix} \mathbf{D}_{\text{small}} \\ \mathbf{0} \end{bmatrix} \begin{matrix} \mathbf{V}_{\text{small}}^* \\ b \times b \end{matrix}, \quad \text{with } \mathbf{U} = \mathbf{U} \begin{bmatrix} \mathbf{U}_{\text{small}} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}. \quad (18)$$

We see that the expression for $\mathbf{U}$ in (18) is exactly analogous to the expression for $\mathbf{V}$ in (17).

4.5 Summary of the Construction of Transformation Matrices

Even though the derivation of the transformation matrices got slightly long, the final algorithm is simple. It can be written down precisely with just a few lines of Matlab code, as shown in the right panel in Figure 3 (the single-step process described in this section is the subroutine stepUTV).

5 THE ALGORITHM RANDUTV

In this section, we describe the algorithm randUTV that given an $m \times n$ matrix $\mathbf{A}$ computes a UTV factorization of the form (1). For concreteness, we assume that $m \geq n$, and that we seek to build an *upper* triangular middle matrix $\mathbf{T}$. The modifications required for the other cases are straightforward. Section 5.1 describes the most basic version of the scheme, and Section 5.2 describes a computationally efficient version. Section 5.3 describes the connection between the presented algorithm and the randomized SVD (RSVD) algorithm for computing a partial factorization. Finally, Section 5.4 provides a calculation of the asymptotic flop count of the randUTV algorithm.

5.1 A Simplistic Algorithm

The algorithm randUTV is obtained by applying the single-step algorithm described in Section 4 repeatedly, to drive $\mathbf{A}$ to upper triangular form one block of b columns at a time. We recall that a cartoon of the process is shown in Figure 2. At the start of the process, we create three arrays that hold the output matrices $\mathbf{T}$, $\mathbf{U}$, and $\mathbf{V}$, and initialize them by setting

$$\mathbf{T} = \mathbf{A}, \quad \mathbf{U} = \mathbf{I}_m, \quad \mathbf{V} = \mathbf{I}_n.$$

In the first step of the iteration, we use the single-step technique described in Section 4 to create two unitary “left and right transformation matrices” $\mathbf{U}^{(1)}$ and $\mathbf{V}^{(1)}$ and then update $\mathbf{T}$, $\mathbf{U}$, and $\mathbf{V}$

<pre> function [U,T,V] = randUTV(A,b,q) T = A; U = eye(size(A,1)); V = eye(size(A,2)); for i = 1:ceil(size(A,2)/b) I1 = 1:(b*(i-1)); I2 = (b*(i-1)+1):size(A,1); J2 = (b*(i-1)+1):size(A,2); if (length(J2) > b) [UU,TT,VV] = stepUTV(T(I2,J2),b,q); else [UU,TT,VV] = svd(T(I2,J2)); end U(:,I2) = U(:,I2)*UU; V(:,J2) = V(:,J2)*VV; T(I2,J2) = TT; T(I1,J2) = T(I1,J2)*VV; end return </pre>	<pre> function [U,T,V] = stepUTV(A,b,q) G = randn(size(A,1),b); Y = A'*G; for i = 1:q Y = A'*(A*Y); end [V,~] = qr(Y); [U,D,W] = svd(A*V(:,1:b)); T = [D,U'*A*... V(:,(b+1):end)]; V(:,1:b) = V(:,1:b)*W; return </pre>
---	---

Fig. 3. Matlab code for the algorithm randUTV that given an $m \times n$ matrix $\mathbf{A}$ computes its UTV factorization $\mathbf{A} = \mathbf{UTV}^*$ (cf. (1)). The input parameters b and q reflect the block size and the number of steps of power iteration, respectively. This code is simplistic in that products of Householder reflectors are stored simply as dense matrices, making the overall complexity $O(n^4)$; it also assumes that $m \geq n$. An efficient implementation is described in Figure 4.

accordingly:

$$\mathbf{T} \leftarrow (\mathbf{U}^{(1)})^* \mathbf{T} \mathbf{V}^{(1)}, \quad \mathbf{U} \leftarrow \mathbf{U} \mathbf{U}^{(1)}, \quad \mathbf{V} \leftarrow \mathbf{V} \mathbf{V}^{(1)}.$$

This leaves us with a matrix $\mathbf{T}$, whose b leading columns are upper triangular (like matrix $\mathbf{A}^{(1)}$ in Figure 2). For the second step, we build transformation matrices $\mathbf{U}^{(2)}$ and $\mathbf{V}^{(2)}$ by applying the single-step algorithm described in Section 4 to the remainder matrix $\mathbf{T}((b+1):m, (b+1):n)$ and then updating $\mathbf{T}$, $\mathbf{U}$, and $\mathbf{V}$ accordingly.

The process then continues to drive one block of b columns at a time to upper triangular form. With $s = \lceil n/b \rceil$ denoting the total number of steps, we find that after $s-1$ steps, all that remains to process is the bottom right block of $\mathbf{T}$ (cf. matrix $\mathbf{A}^{(2)}$ in Figure 2). This block consists of b columns if n is a multiple of the block size, and is otherwise even thinner. For this last block, we obtain the final left and right transformation matrices $\mathbf{U}^{(s)}$ and $\mathbf{V}^{(s)}$ by computing a full SVD of the remaining matrix $\mathbf{T}(((s-1)b+1):m, ((s-1)b+1):n)$, and updating the matrices $\mathbf{T}$, $\mathbf{U}$, and $\mathbf{V}$ accordingly. (In the cartoon in Figure 2, we have $s=3$, and the matrices $\mathbf{U}^{(3)}$ and $\mathbf{V}^{(3)}$ are built by computing a full SVD of a dense matrix of size $2s \times s$.)

We call the algorithm described in this section randUTV. It can be coded using just a few lines of Matlab code, as illustrated in Figure 3. In this simplistic version, all unitary matrices are represented as dense matrices, which makes the overall complexity $O(n^4)$ for an $n \times n$ matrix.

5.2 A Computationally Efficient Version

In this section, we describe how the basic version of randUTV, as given in Figure 3, can be turned into a highly efficient procedure via three simple modifications. The resulting algorithm is summarized in Figure 4. We note that the two versions of randUTV described in Figures 3 and 4 are

mathematically equivalent; if they were to be executed in exact arithmetic, their outputs would be identical.

The first modification is that all operations on the matrix $\mathbf{T}$ and on the unitary matrices $\mathbf{U}$ and $\mathbf{V}$ should be carried out “in place” to not unnecessarily move any data. To be precise, using the notation in Section 4, we generate at each iteration four unitary matrices $\mathbf{U}$, $\mathbf{V}$, $\mathbf{U}_{\text{small}}$, and $\mathbf{V}_{\text{small}}$. As soon as such a matrix is generated, it is immediately applied to $\mathbf{T}$ and then used to update either $\mathbf{U}$ or $\mathbf{V}$.

Second, we exploit that the two “large” unitary transforms $\mathbf{U}$ and $\mathbf{V}$ both consist of products of b Householder reflectors. We generate them by computing unpivoted Householder QR factorizations of tall and thin matrices, using a subroutine that outputs simply the b Householder vectors. Then, $\mathbf{U}$ and $\mathbf{V}$ can both be stored and applied efficiently, as described in Section 2.4.

The third and final modification pertains to the situation where the input matrix $\mathbf{A}$ is nonsquare. In this case, the full SVD that is computed in the last step involves a rectangular matrix. When $\mathbf{A}$ is tall ($m > n$), we find at this step that J_3 is empty, so the matrix to be processed is $\mathbf{T}([I_2, I_3], J_2)$. When computing the SVD of this matrix, we use the efficient version described in Remark 6, which outputs a factorization in which $\mathbf{U}_{\text{small}}$ consists in part of a product of Householder reflectors. (In this situation, $\mathbf{U}_{\text{small}}$ is in fact not necessarily “small,” but it can be stored and applied efficiently.)

Remark 7 (The Case $m < n$). In a situation where the matrix has fewer rows than columns but we still seek an upper triangular matrix $\mathbf{T}$, randUTV proceeds exactly as described for the first $s - 1$ steps. In the final step, we now find that I_3 is empty, but J_3 is not, and so we need to compute the SVD of the “fat” matrix $\mathbf{T}(I_2, [J_2, J_3])$. We do this in a manner entirely analogous to how we handle a “tall” matrix, by first performing an unpivoted QR factorization of the rows of $\mathbf{T}(I_2, [J_2, J_3])$. In this situation, it is the matrix of right singular vectors at the last step that consists in part of a product of Householder reflectors.

Remark 8 (The Cases $m \ll n$ or $m \gg n$). The algorithm shown in Figure 4 is very efficient when m and n are of the same order of magnitude. However, if $m \ll n$ (or vice versa), then this algorithm could be accelerated by computing first an unpivoted QR factorization $\mathbf{A} = \mathbf{QR}$ and then applying the algorithm in Figure 4 to $\mathbf{R}$. This is standard practice for the computation of other factorizations, such as the SVD (cf. Section 2.4).

5.3 Connection Between RSVD and randUTV

The proposed algorithm randUTV is directly inspired by the RSVD algorithm described in Remark 5 (as originally described in References [22, 28, 34] and later elaborated in References [19, 29]). In this section, we explore this connection in more detail and demonstrate that the low-rank approximation error that results from the single-step UTV-factorization described in Section 4 is identical to the error produced by the RSVD (with a twist). This means that the detailed error analysis that is available for the RSVD (e.g., see References [15, 19, 45]) immediately applies to the procedure described here. To be precise, we present the following theorem.

THEOREM 1. *Let $\mathbf{A}$ be an $m \times n$ matrix, let b be an integer denoting step size such that $1 \leq b < \min(m, n)$, and let q denote a non-negative integer. Let $\mathbf{G}$ denote the $m \times b$ Gaussian matrix drawn in Section 4.3, and let $\mathbf{U}$, $\mathbf{T}$, and $\mathbf{V}$ be the factors in the factorization $\mathbf{A} = \mathbf{UTV}^*$ built in Sections 4.3 and 4.4, partitioned as in (12) and (13). We have the following:*

- (a) *Let $\mathbf{Y} = (\mathbf{A}^* \mathbf{A})^q \mathbf{A}^* \mathbf{G}$ denote a sampling matrix, and let $\mathbf{Q}$ denote an $n \times b$ orthonormal matrix whose columns form a basis for the column space of $\mathbf{Y}$. Then, the error $\|\mathbf{A} - \mathbf{A} \mathbf{Q} \mathbf{Q}^*\|$ precisely equals the error incurred by the RSVD with q steps of power iteration, as analyzed*

```

[U, T, V] = randUTV(A, b, q)
% Initialize output variables:
T = A; U = Im; V = In;
for i = 1 : min(⌈m/b⌉, ⌈n/b⌉) do
    % Create partitions 1 : m = [I1, I2, I3] and 1 : n = [J1, J2, J3] so that (I2, J2) points to the “active” block
    % that is to be diagonalized:
    I1 = 1 : (b(i - 1)); I2 = (b(i - 1) + 1) : min(bi, m); I3 = (bi + 1) : m;
    J1 = 1 : (b(i - 1)); J2 = (b(i - 1) + 1) : min(bi, n); J3 = (bi + 1) : n;
    if (I3 and J3 are both nonempty) then
        % Generate the sampling matrix Y whose columns approximately span the space spanned by the b
        % dominant right singular vectors of the matrix X = T([I2, I3], [J2, J3]). We do this via randomized
        % sampling, setting Y = (X*X)qX*G, where G is a Gaussian random matrix with b columns.
        G = randn(m - b(i - 1), b)
        Y = T([I2, I3], [J2, J3])*G
        for j = 1 : q do
            Y = T([I2, I3], [J2, J3])*(T([I2, I3], [J2, J3])Y).
        end for
        % Build a unitary matrix V whose first b columns form an orthonormal basis for the columns of Y. Then,
        % apply the transformations. (We exploit that V is a product of b Householder reflectors, cf. Section 2.4.)
        [V, ~] = qr(Y)
        T(:, [J2, J3]) = T(:, [J2, J3])V
        V(:, [J2, J3]) = V(:, [J2, J3])V
        % Build a unitary matrix U whose first b columns form an orthonormal basis for the columns of
        % T([I2, I3], J2). Then, apply the transformations. (We exploit that U is a product of b Householder
        % reflectors, cf. Section 2.4.)
        [U, R] = qr(T([I2, I3], J2))
        U(:, [I2, I3]) = U(:, [I2, I3])U
        T([I2, I3], J3) = U*T([I2, I3], J3)
        T(I3, J2) = 0
        % Perform the local SVD that diagonalizes the active diagonal block. Then, apply the transformations.
        [Usmall, Dsmall, Vsmall] = svd(R(1 : b, 1 : b))
        T(I2, J2) = Dsmall
        T(I2, J3) = Usmall*T(I2, J3)
        U(:, I2) = U(:, I2)Usmall
        T(I1, J2) = T(I1, J2)Vsmall
        V(:, J2) = V(:, J2)Vsmall
    else
        % Perform the local SVD that diagonalizes the last diagonal block. Then, apply the transformations. If
        % either I3 or J3 is long, this should be done economically, cf. Section 5.2.
        [Usmall, Dsmall, Vsmall] = svd(T([I2, I3], [J2, J3]))
        U(:, [I2, I3]) = U(:, [I2, I3])Usmall
        V(:, [J2, J3]) = V(:, [J2, J3])Vsmall
        T([I2, I3], [J2, J3]) = Dsmall
        T(I1, [J2, J3]) = T(I1, [J2, J3])Vsmall
    end if
end for
return

```

Fig. 4. The algorithm randUTV that given an $m \times n$ matrix $\mathbf{A}$ computes the UTV factorization $\mathbf{A} = \mathbf{U}\mathbf{T}\mathbf{V}^*$ (cf. (1)). The input parameters b and q reflect the block size and the number of steps of power iteration, respectively.

in Halo et al. [19, Sec. 10]. It holds that

$$\|\mathbf{A} - \mathbf{A}\mathbf{Q}\mathbf{Q}^*\| = \|\mathbf{A} - \mathbf{U}_1\mathbf{T}_{11}\mathbf{V}_1^*\| = \left\| \begin{bmatrix} \mathbf{T}_{12} \\ \mathbf{T}_{22} \end{bmatrix} \right\|. \quad (19)$$

- (b) Let $\mathbf{Z} = \mathbf{A}\mathbf{Y} = (\mathbf{A}\mathbf{A}^*)^{q+1}\mathbf{G}$ denote a sampling matrix, and let $\mathbf{W}$ denote an $m \times b$ orthonormal matrix whose columns form a basis for the column space of $\mathbf{Z}$. If the rank of $\mathbf{A}$ is at least b , then

$$\|\mathbf{A} - \mathbf{W}\mathbf{W}^*\mathbf{A}\| = \|\mathbf{A} - \mathbf{U}_1(\mathbf{T}_{11}\mathbf{V}_1^* + \mathbf{T}_{12}\mathbf{V}_2^*)\| = \|\mathbf{T}_{22}\|. \quad (20)$$

We observe that the term $\|\mathbf{A} - \mathbf{W}\mathbf{W}^*\mathbf{A}\|$ that arises in part (b) can informally be said to be the error resulting from RSVD with “ $q + 1/2$ ” steps of power iteration. This conforms with what one might have optimistically hoped for, given that the RSVD involves $2q + 2$ applications of either $\mathbf{A}$ or $\mathbf{A}^*$ to thin matrices with b columns, and randUTV involves $2q + 3$ such operations at each step ($2q + 1$ applications in building $\mathbf{Y}$, and then the computations of $\mathbf{A}\mathbf{V}$ and $\mathbf{U}^*\mathbf{A}$, which are in practice applications of $\mathbf{A}$ to thin matrices due to the identity (5)).

PROOF. The proofs for the two parts rest on the fact that $\mathbf{A}$ can be decomposed as follows:

$$\mathbf{A} = \mathbf{U}\mathbf{T}\mathbf{V}^* = \begin{bmatrix} \mathbf{U}_1 & \mathbf{U}_2 \end{bmatrix} \begin{bmatrix} \mathbf{T}_{11} & \mathbf{T}_{12} \\ \mathbf{0} & \mathbf{T}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{V}_1^* \\ \mathbf{V}_2^* \end{bmatrix} = \mathbf{U}_1\mathbf{T}_{11}\mathbf{V}_1^* + \mathbf{U}_1\mathbf{T}_{12}\mathbf{V}_2^* + \mathbf{U}_2\mathbf{T}_{22}\mathbf{V}_2^*, \quad (21)$$

where $\mathbf{U}_1$ and $\mathbf{V}_1$ have b columns each, and $\mathbf{T}_{11}$ is of size $b \times b$. With the identity (21) in hand, the claims in (a) follow once we have established that the orthogonal projection $\mathbf{Q}\mathbf{Q}^*$ equals the projection $\mathbf{V}_1\mathbf{V}_1^*$. The claims in (b) follow once we establish that $\mathbf{W}\mathbf{W}^* = \mathbf{U}_1\mathbf{U}_1^*$. Thus, the proofs for (a) and (b) proceed as follows:

- (a) Observe that the matrix $\mathbf{Q}$ is by construction identical to the matrix $\mathbb{V}_1$ built in Sections 4.3 and 4.4. Since $\mathbb{V}_1 = \mathbf{V}_1\mathbf{V}_{\text{small}}^*$, we find that $\mathbf{Q}\mathbf{Q}^* = \mathbb{V}_1\mathbb{V}_1^* = (\mathbf{V}_1\mathbf{V}_{\text{small}}^*)(\mathbf{V}_1\mathbf{V}_{\text{small}}^*)^* = \mathbf{V}_1(\mathbf{V}_{\text{small}}^*\mathbf{V}_{\text{small}})\mathbf{V}_1^*$. Since $\mathbf{V}_{\text{small}}^*\mathbf{V}_{\text{small}} = \mathbf{I}_b$, it follows that $\mathbf{Q}\mathbf{Q}^* = \mathbf{V}_1\mathbf{V}_1^*$. Then,

$$\mathbf{A}\mathbf{Q}\mathbf{Q}^* = \mathbf{A}\mathbf{V}_1\mathbf{V}_1^* = \{\text{Use (21) and that } \mathbf{V}_2^*\mathbf{V}_1 = \mathbf{0} \text{ and } \mathbf{V}_1^*\mathbf{V}_1 = \mathbf{I}_b\} = \mathbf{U}_1\mathbf{T}_{11}\mathbf{V}_1^*. \quad (22)$$

The first identity in (19) follows immediately from (22). The second identity holds since (21) implies that $\mathbf{A} - \mathbf{U}_1\mathbf{T}_{11}\mathbf{V}_1^* = \mathbf{U}_1\mathbf{T}_{12}\mathbf{V}_2^* + \mathbf{U}_2\mathbf{T}_{22}\mathbf{V}_2^* = \mathbf{U} \begin{bmatrix} \mathbf{T}_{12} \\ \mathbf{T}_{22} \end{bmatrix} \mathbf{V}_2^*$, with $\mathbf{U}$ unitary and $\mathbf{V}_2$ orthonormal.

- (b) We will first prove that with probability 1, the two $m \times b$ matrices $\mathbf{A}\mathbb{V}_1$ and $\mathbf{Z}$ have the same column spaces. To this end, note that since $\mathbb{V}_1$ is obtained by performing an unpivoted QR factorization of the matrix $\mathbf{Y}$ defined in (a), we know that $\mathbb{V}_1\mathbf{R} = \mathbf{Y}$ for some $b \times b$ upper triangular matrix $\mathbf{R}$. The assumption that $\mathbf{A}$ has rank at least b implies that $\mathbf{R}$ is invertible with probability 1. Consequently, $\mathbf{A}\mathbb{V}_1 = \mathbf{A}\mathbf{Y}\mathbf{R}^{-1} = \mathbf{Z}\mathbf{R}^{-1}$, since $\mathbf{Z} = \mathbf{A}\mathbf{Y}$. Since right multiplication by an invertible matrix does not change the column space of a matrix, the claim follows.

Since $\mathbf{A}\mathbb{V}_1$ and $\mathbf{Z}$ have the same column spaces, it follows from the definition of $\mathbf{U}$ that $\mathbf{U}_1\mathbf{U}_1^* = \mathbf{W}\mathbf{W}^*$. Since $\mathbf{U}_1 = \mathbf{U}_1\mathbf{U}_{\text{small}}$, where $\mathbf{U}_{\text{small}}$ is unitary, we see that $\mathbf{U}_1\mathbf{U}_1^* = \mathbf{U}_1\mathbf{U}_1^*$. Consequently,

$$\mathbf{W}\mathbf{W}^*\mathbf{A} = \mathbf{U}_1\mathbf{U}_1^*\mathbf{A} = \{\text{Use (21).}\} = \mathbf{U}_1\mathbf{T}_{11}\mathbf{V}_1^* + \mathbf{U}_1\mathbf{T}_{12}\mathbf{V}_2^*,$$

which establishes the first identity in (20). The second identity holds since (21) implies that $\mathbf{A} - \mathbf{U}_1(\mathbf{T}_{11}\mathbf{V}_1^* + \mathbf{T}_{12}\mathbf{V}_2^*) = \mathbf{U}_2\mathbf{T}_{22}\mathbf{V}_2^*$, with $\mathbf{U}_2$ and $\mathbf{V}_2$ orthonormal. $\square$

Remark 9 (Oversampling). We recall that the accuracy of randUTV depends on how well the space $\text{col}(\mathbb{V}_1)$ aligns with the space spanned by the b dominant right singular vectors of $\mathbf{A}$. If these two spaces were to match exactly, then the truncated UTV factorization would achieve perfectly

optimal accuracy. One way to improve the alignment is to increase the power parameter q . A second way to make the two spaces align better is to use *oversampling*, as described in Remark 4. With p an oversampling parameter (say $p = 5$ or $p = 10$), we would draw a Gaussian random matrix $\mathbf{G}$ of size $m \times (b + p)$ and then compute an “extended” sampling matrix $\mathbf{Y}' = (\mathbf{A}^* \mathbf{A})^q \mathbf{A}^* \mathbf{G}$ of size $n \times (b + p)$. The $n \times b$ sampling matrix $\mathbf{Y}$ that we would use to compute $\mathbb{V}$ would then be formed by the b dominant left singular vectors of $\mathbf{Y}'$ (see the first figure in the appendix). Oversampling in this fashion does improve the accuracy (see the appendix), but in our experience the additional computational cost is not worth it. Incorporating oversampling would also introduce an additional tuning parameter p , which is in many ways undesirable.

5.4 Theoretical Cost of randUTV

Now we analyze the theoretical cost of the implementation of randUTV, and we compare it to those of CPQR and SVD.

The theoretical cost of the CPQR factorization and the unpivoted QR factorization of an $m \times n$ matrix is $2mn^2 - 2n^3/3$ flops, when no orthonormal matrices are required. Although the theoretical cost of both the pivoted and the unpivoted QR is the same, other factors should be considered, the most important one being the quality of flops. In modern architectures, flops performed inside BLAS-3 operations can be about 5 to 10 times faster than flops performed inside BLAS-1 and BLAS-2 operations since BLAS-3 is CPU bound, whereas BLAS-1 and BLAS-2 are memory bound. Hence, the unpivoted QR factorization in high-performance libraries such as in Anderson et al. [2] can be much faster because most of the flops are performed inside BLAS-3 operations, whereas only half of the flops in the best implementations of CPQR (dgeqp3) [33] in Anderson et al. [2] are performed inside BLAS-3 operations. In addition, this low performance of CPQR can even be smaller because of the appearance of catastrophic cancellations during the computations. The appearance of just one catastrophic cancellation will stop the building of a block Householder reflector before all of it has been built. This sudden stop forces the algorithm to work on smaller block sizes, which are suboptimal, and hence performances are even lower.

The SVD usually comprises two steps: the reduction to bidiagonal form and then the reduction from bidiagonal to diagonal form. The first step is a direct step, whereas the second step is an iterative one. The theoretical cost of the reduction to bidiagonal form of an $m \times n$ matrix is $4mn^2 - 4n^3/3$ flops, when no singular vectors are needed. If $m \gg n$, the cost can be reduced to $2mn^2 + 2n^3$ flops by performing first a QR factorization. The cost of the reduction to diagonal form depends on the number of iterations, which is unknown *a priori*, but it is usually small when no singular vectors are built. On the one hand, in the bidiagonalization, a large share of the flops are performed inside the not-so-efficient BLAS-1 and BLAS-2 operations. Therefore, no high performances are obtained in the reduction to bidiagonal form. On the other hand, the reduction to diagonal form uses just BLAS-1 operations. These operations are memory bound, and additionally they cannot be efficiently parallelized within BLAS, which might reduce performances on multicore machines. In conclusion, usual implementations of the SVD will render low performances on both single-core and multicore architectures.

The theoretical cost of the randUTV factorization of an $m \times n$ matrix is $(5 + 2q)mn^2 - (3 + 2q)n^3/3$ flops, when no orthonormal matrices are required, and q steps of power iteration are applied. If $q = 0$, the theoretical cost of randUTV is three times as high as the theoretical cost of CPQR; if $q = 1$, it is four times as high; and if $q = 2$, it is five times as high. Although randUTV seems computationally more expensive than CPQR, the quality of flops should be considered. The share of BLAS-1 and BLAS-2 flops in randUTV is very small: BLAS-1 and BLAS-2 flops are only employed inside the CPQR factorization of the sampling matrix $\mathbf{Y}$, the QR factorization of the current column block, and the SVD of the diagonal block. As these operations only apply to blocks of di-

mensions $n \times b$, $m \times b$, and $b \times b$, respectively, the total amount of these types of flops is negligible, and therefore most of the flops performed in randUTV are BLAS-3 flops. Hence, the algorithm for computing the randUTV will be much faster than what the theoretical cost predicts. In conclusion, this heavy use of BLAS-3 operations will render good performances on single-core architectures, multicore architectures, GPUs, and distributed-memory architectures.

6 NUMERICAL RESULTS

6.1 Computational Speed

In this section, we investigate the speed of the proposed algorithm randUTV and compare it to the speeds of highly optimized methods for computing the SVD and the CPQR factorization.

All experiments reported in this article were performed on an Intel Xeon E5-2670 processor at 2.6GHz, with 16 cores and 64GB of RAM. Other details of interest include that the OS used was Linux (Version 2.6.32-431.el6.x86_64), and the code was compiled with gcc (Version 4.4.7). Main routines for computing the SVD (dgesvd and dgesdd) and the CPQR (dgeqp3) were taken from Intel's MKL library (Version 11.1.3), as this library usually delivers much higher performances than Netlib's LAPACK codes. Our implementations were coded with libflame [42, 43] (Release 11104).

Each of the three algorithms we tested (randUTV, SVD, CPQR) was applied to double-precision real matrices of size $n \times n$. We report the following times:

T_{svd}	The time in seconds for the LAPACK functions dgesvd and dgesdd from Intel's MKL.
T_{cpqr}	The time in seconds for the LAPACK function dgeqp3 from Intel's MKL.
T_{randUTV}	The time in seconds for our implementation of randUTV.

For the purpose of a fair comparison, the three implementations were linked to the BLAS library from Intel's MKL. In all cases, we used an algorithmic block size of $b = 64$. Although likely not optimal for all problem sizes, this block size yields near best performance, and regardless, it allows us to easily compare and contrast the performance of the different implementations.

Table 1 shows the measured computational times when executed on 1, 8, and 16 cores, respectively. In these experiments, all orthonormal matrices (**U** and **V** for SVD and UTV, and **Q** for CPQR) are explicitly formed. This slightly favors CPQR because only one orthonormal matrix is required. The corresponding numbers obtained when orthonormal matrices are not built are given in the appendix. To better illustrate the relative performance of the various techniques, in Figure 5 we plot the computational times measured divided by n^3 . Since all techniques under consideration have asymptotic complexity $O(n^3)$ when applied to an $n \times n$ matrix, these graphs better reveal the computational efficiency. (We plot time divided by n^3 rather than the more commonly reported "normalized Gigaflows," as the algorithms we compare have different scaling factors multiplying the dominant n^3 -term in the asymptotic flop count.) Figure 5 also shows the timings we measured when the orthonormal matrices were not formed.

The results in Figure 5 lead us to make several observations. First, the algorithm randUTV is decisively faster than the SVD in almost all cases (the exceptions involve the situation where no unitary matrices are sought and the input matrix is small). Second, comparing the speeds of CPQR and randUTV, when both methods are executed on a single core, CPQR is similar when no orthonormal matrices are built or slightly faster when orthonormal matrices are built (because in this case randUTV must build two matrices). Third, as the matrix size grows, and as the number of cores increases, randUTV gains an edge on CPQR in terms of speed.

6.2 Errors

In this section, we describe the results of the numerical experiments that were conducted to investigate how good randUTV is at computing an accurate rank- k approximation to a given matrix.

Table 1. Computational Times of Different Factorizations When Executed on One Core (Top), 8 Cores (Middle), and 16 Cores (Bottom)

Computational times when executed on a single core						
n	T_{svd}		T_{cpqr}	T_{randUTV}		
	dgesvd	dgesdd		$q = 0$	$q = 1$	$q = 2$
500	1.93e-01	1.24e-01	3.78e-02	8.43e-02	9.37e-02	1.03e-01
1000	1.22e+00	7.82e-01	2.62e-01	5.08e-01	5.74e-01	6.45e-01
2000	9.04e+00	5.88e+00	2.18e+00	3.58e+00	4.09e+00	4.61e+00
3000	2.98e+01	2.00e+01	7.59e+00	1.14e+01	1.31e+01	1.50e+01
4000	7.02e+01	4.70e+01	1.80e+01	2.71e+01	3.13e+01	3.56e+01
5000	1.35e+02	9.04e+01	3.48e+01	5.07e+01	5.90e+01	6.74e+01
6000	2.29e+02	1.54e+02	5.95e+01	8.74e+01	1.01e+02	1.16e+02
8000	5.47e+02	3.68e+02	1.43e+02	2.10e+02	2.45e+02	2.81e+02
10000	1.03e+03	6.98e+02	2.74e+02	3.95e+02	4.61e+02	5.26e+02

Computational times when executed on 8 cores						
n	T_{svd}		T_{cpqr}	T_{randUTV}		
	dgesvd	dgesdd		$q = 0$	$q = 1$	$q = 2$
500	1.12e-01	7.70e-02	2.32e-02	5.36e-02	5.64e-02	5.92e-02
1000	5.27e-01	3.23e-01	1.02e-01	2.18e-01	2.35e-01	2.51e-01
2000	2.63e+00	1.57e+00	4.92e-01	1.14e+00	1.24e+00	1.33e+00
3000	8.23e+00	5.04e+00	2.36e+00	2.92e+00	3.25e+00	3.56e+00
4000	1.94e+01	1.34e+01	6.45e+00	5.92e+00	6.57e+00	7.28e+00
5000	3.79e+01	2.71e+01	1.28e+01	1.01e+01	1.13e+01	1.25e+01
6000	6.46e+01	4.82e+01	2.24e+01	1.62e+01	1.84e+01	2.05e+01
8000	1.55e+02	1.21e+02	5.37e+01	3.70e+01	4.20e+01	4.70e+01
10000	2.90e+02	2.29e+02	1.01e+02	6.48e+01	7.43e+01	8.39e+01

Computational times when executed on 16 cores						
n	T_{svd}		T_{cpqr}	T_{randUTV}		
	dgesvd	dgesdd		$q = 0$	$q = 1$	$q = 2$
500	1.08e-01	7.74e-02	2.32e-02	5.83e-02	6.14e-02	6.44e-02
1000	4.69e-01	3.03e-01	9.06e-02	2.24e-01	2.38e-01	2.52e-01
2000	2.26e+00	1.33e+00	4.09e-01	9.51e-01	1.02e+00	1.09e+00
3000	7.05e+00	4.34e+00	2.16e+00	2.34e+00	2.57e+00	2.81e+00
4000	1.73e+01	1.21e+01	6.08e+00	4.62e+00	5.10e+00	5.57e+00
5000	3.34e+01	2.48e+01	1.22e+01	7.54e+00	8.31e+00	9.09e+00
6000	5.86e+01	4.47e+01	2.15e+01	1.24e+01	1.38e+01	1.52e+01
8000	1.38e+02	1.12e+02	5.15e+01	2.72e+01	3.02e+01	3.32e+01
10000	2.62e+02	2.14e+02	9.70e+01	4.57e+01	5.12e+01	5.68e+01

All orthonormal matrices are built explicitly.

Specifically, we compared how well partial factorizations reveal the numerical ranks of four different test matrices:

- *Matrix 1 (Fast Decay)*: This is an $n \times n$ matrix of the form $\mathbf{A} = \mathbf{U}\mathbf{D}\mathbf{V}^*$, where $\mathbf{U}$ and $\mathbf{V}$ are randomly drawn matrices with orthonormal columns (obtained by performing an unpivoted

Fig. 5. Computational cost (the lower, the better performances) of randUTV compared to the cost of the LAPACK drivers dgesvd and dgesdd (SVD), and dgeqp3 (CPQR). The algorithms were applied to double-precision real matrices of size $n \times n$. The measured computational time divided by n^3 is plotted against n (e.g., the red lines are T_{svd}/n^3). The three rows of plots correspond to executing on 1, 8, and 16 cores, respectively.

QR factorization on a random Gaussian matrix), and where $\mathbf{D}$ is diagonal with entries given by $d_j = \beta^{(j-1)/(n-1)}$ with $\beta = 10^{-5}$.

- **Matrix 2 (S-Shaped Decay):** This matrix is built in the same manner as Matrix 1, but now the diagonal entries of $\mathbf{D}$ are chosen to first hover around 1, then decay rapidly, and then level out at 10^{-2} , as shown later in Figure 7 (black line).
- **Matrix 3 (Gap):** This matrix is built in the same manner as Matrix 1, but now there is a sudden drop in magnitudes of the singular values so that $\sigma_{151} = 0.1 \sigma_{150}$. Specifically, $\mathbf{D}(j, j) = 1/j$ for $j \leq 150$ and $\mathbf{D}(j, j) = 0.1/j$ for $j > 150$.
- **Matrix 4 (BIE):** This matrix is the result of discretizing a Boundary Integral Equation (BIE) defined on a smooth closed curve in the plane. To be precise, we discretized the so-called single-layer operator associated with the Laplace equation using a sixth-order quadrature rule designed by Alpert [1]. This operator is well known to be ill conditioned, which neces-

Fig. 6. Rank- k approximation errors for the matrix “Fast Decay” (see Section 6.2) of size 4000×4000 . The block size was $b = 100$. Left: Absolute errors in spectral norm. The black line (circles) marks the theoretically minimal errors. Right: Relative errors, as defined by (27).

sitates the use of a rank-revealing factorization to solve the corresponding linear system in as stable a manner as possible.

For each test matrix, we computed the error

$$e_k = \|A - A_k\|, \quad (23)$$

where A_k is the rank- k approximation resulting from either of the three techniques discussed in this article:

$$\text{SVD:} \quad A_k = A_k^{\text{optimal}} = U(:, 1:k) D(1:k, 1:k) V(:, 1:k)^*, \quad (24)$$

$$\text{CPQR:} \quad A_k = Q(:, 1:k) R(1:k, :) P^*, \quad (25)$$

$$\text{randUTV:} \quad A_k = U(:, 1:k) T(1:k, :) V^*. \quad (26)$$

For randUTV, we ran the experiment with zero, one, and two steps of power iteration ($q = 0, 1, 2$). In addition to the direct errors defined by (23), we also calculated the relative errors, as defined via

$$e_k^{\text{relative}} = 100\% \times \frac{\|A - A_k\|}{\|A - A_k^{\text{optimal}}\|}. \quad (27)$$

The results are shown in Figures 6 through 9. The figures also report the errors resulting from a UTV factorization that G. W. Stewart proposed in reference [38], precisely for purposes of low-rank approximation and estimation of singular values. To be precise, Stewart builds a factorization $A = ULV^*$, where U and V are orthonormal, and L is lower triangular. The procedure is to first compute a CPQR factorization of A so that $A = Q_1 R_1 P_1^*$. Then, compute a CPQR of the transpose of the upper triangular matrix R_1 so that $R_1^* = Q_2 L^* P_2^*$. Finally, set $U = Q_1 P_2$ and $V = P_1 Q_2$. The rank- k approximation is then

$$\text{QLP:} \quad A_k = U L(:, 1:k) V(:, 1:k)^*. \quad (28)$$

Fig. 7. Rank- k approximation errors for the matrix “S-shaped decay” (see Section 6.2) of size 4000×4000 . The block size was $b = 100$. Left: Absolute errors in spectral norm. The black line (circles) marks the theoretically minimal errors. Right: Relative errors, as defined by (27).

Fig. 8. Rank- k approximation errors for the matrix “Gap” (see Section 6.2) of size 4000×4000 . The block size was $b = 100$. Left: Absolute errors in spectral norm. The black line (circles) marks the theoretically minimal errors. (Observe that we zoomed in on the area around the gap.) Right: Relative errors, as defined by (27).

Based on the errors shown in Figures 6 through 9, we make several empirical observations. First, randUTV is much better than CPQR at computing low-rank approximations. Even when no power iteration ($q = 0$) is used, errors from randUTV are substantially smaller. When one or two steps of power iteration are taken ($q = 1$ or $q = 2$), the errors become close to optimal in all cases studied. Second, for the matrix with a gap in its singular values (cf. Figure 8), randUTV performs

Fig. 9. Rank- k approximation errors for the matrix “BIE” (see Section 6.2) of size 4000×4000 . The block size was $b = 100$. Left: Absolute errors in spectral norm. The black line (circles) marks the theoretically minimal errors. Right: Relative errors, as defined by (27).

remarkably well in that both σ_{150} and σ_{151} are approximated to high accuracy. Third, the relative errors resulting from randUTV are consistently small, and much more reliably small than those resulting from CPQR. Fourth, comparing the QLP factorization of Stewart [38] to randUTV, we see that Stewart’s algorithm results in errors that are similar to those resulting from randUTV with $q = 0$. As soon as the power parameter is increased, randUTV tends to perform better. (We observe that in terms of speed, Stewart’s QLP algorithm relies on two CPQR factorizations, which makes it much slower than randUTV.)

All error results shown in this section refer to errors measured in the spectral norm. When errors are measured in the Frobenius norm, randUTV performs even better, as shown in the figures in the appendix. The effects of including oversampling in the algorithm, as described in Remark 9, are illustrated in numerical experiments given in the appendix. These experiments show that oversampling does improve the error, but also that the improvement is almost imperceptible. For most applications, oversampling is, in our experience, not worth the additional effort.

Remark 10 (Is randUTV Rank-revealing?). Although we do not have rigorous proofs that randUTV is rank revealing, the numerical experiments in this section strongly indicate that it is, in the sense that a truncated UTV factorization approximates a given matrix to nearly optimal accuracy (particularly when a small number of steps of power iteration is used). We did not formally investigate whether randUTV is “strongly rank revealing” in the sense of Gu and Eisenstat [16], but the numerical experiments certainly indicate that it does very well in this regard.

6.3 Concentration of Mass to the Diagonal

In this section, we investigate our claim that randUTV produces a matrix $\mathbf{T}$ whose diagonal entries are close approximations to the singular values of the given matrix. (We recall that in the factorization $\mathbf{A} = \mathbf{UTV}^*$, the matrix $\mathbf{T}$ is upper triangular but with the entries above the diagonal very small in magnitude, which forces the diagonal entries to approach the corresponding singular values.)

Fig. 10. The results shown here illustrate how well randUTV approximates the singular values of four different matrices of size 4000×4000 . For reference, we also include the diagonal values of the middle factors in a plain CPQR and in Stewart’s QLP factorization [38]. The left column shows the diagonal entries themselves. The right column shows relative errors so that, for example, the magenta line (crosses) for CPQR plots $100\% \times |\mathbf{R}(i, i) - \sigma_i|/\sigma_i$ versus i .

Figure 10 shows the results for the four different test matrices described in Section 6.2, again for matrices of size 4000×4000 . For reference, we also show the diagonal entries of the “R-factor” in a CPQR and the diagonal entries of the “L-factor” in Stewart’s QLP factorization [38]. We see that both Stewart’s and our algorithm result in far better results than plain CPQR. randUTV roughly matches the accuracy of the QLP when $q = 0$ and does better as q is increased to one or two. (We recall that randUTV is much faster than the QLP method.)

7 AVAILABILITY OF CODE

Implementations of the discussed algorithm are available under 3-clause (modified) BSD license from <https://github.com/flame/randutv>.

This repository includes two different implementations: one to be used with the LAPACK library [2], and the other one to be used with the libFLAME [17, 32] library.

8 CONCLUDING REMARKS AND FUTURE WORK

We have described a randomized algorithm that we call *randUTV* for computing an economical alternative to the SVD of a matrix. The new method is much faster than classical algorithms for computing the SVD (cf. Section 6.1), and it provides accuracy that is very close to optimal for tasks such as low-rank approximation or subspace identification (cf. Section 6.2). The new algorithm is blocked and executes well on modern communication-constrained hardware. It is incremental and can be used to compute partial factorizations.

Compared to the CPQR factorization, which is commonly used as an economical alternative to the SVD for low-rank approximation, *randUTV* is similar or better in terms of speed in modern architectures and much more accurate. Compared to methods designed specifically for low-rank approximation, such as the strongly rank-revealing QR factorizations [16] and the QLP factorization [38], *randUTV* provides similar or better accuracy and is much faster.

The new method can also be viewed as an alternative to existing randomized methods for low-rank approximation, as described in References [19, 22, 29]. These methods excel when the numerical rank k of a matrix is much smaller than the matrix dimensions m and n . Moreover, the methods of References [19, 22, 29] work best when the user has some rough idea of what k is in advance (although these requirements were relaxed substantially in Martinsson and Voronin [26]). In contrast, *randUTV* provides high speed for any rank and does not need any *a priori* information about the numerical rank.

In this article, we focused on the case of multicore CPUs with shared memory. We expect that the relative advantages of *randUTV* will be even more pronounced in more severely communication-constrained environments, such as GPUs, distributed memory parallel computers, or factorizations of matrices stored out of core. Work on variations of the method modified for such environments is currently under way.

ACKNOWLEDGMENTS

We thank Francisco D. Igual-Peña from the Universidad Complutense de Madrid (Spain) for granting access to the computing platform where some of the experiments were conducted.

REFERENCES

- [1] Bradley K. Alpert. 1999. Hybrid Gauss-trapezoidal quadrature rules. *SIAM Journal on Scientific Computing* 20, 5 (1999), 1551–1584.
- [2] E. Anderson, Z. Bai, C. Bischof, L. S. Blackford, J. Demmel, Jack J. Dongarra, et al. 1999. *LAPACK Users' Guide* (3rd ed.). SIAM, Philadelphia, PA.
- [3] Jesse L. Barlow. 2003. Modification and maintenance of ULV decompositions. In *Applied Mathematics and Scientific Computing*, Z. Drmac, V. Hari, L. Sopta, Z. Tutek, and K. Veselic (Eds.). Springer US, Boston, MA, 31–62. DOI: https://doi.org/10.1007/978-1-4757-4532-0_2
- [4] Jesse L. Barlow, Peter A. Yoon, and Hongyuan Zha. 1996. An algorithm and a stability theory for downdating the ULV decomposition. *BIT Numerical Mathematics* 36, 1 (1996), 14–40.
- [5] Klaus-Jürgen Bathe. 1971. *Solution Methods for Large Generalized Eigenvalue Problems in Structural Engineering*. National Technical Information Service, US Department of Commerce, Alexandria, VA.
- [6] Christian Bischof and Charles Van Loan. 1987. The WY representation for products of householder matrices. *SIAM Journal on Scientific and Statistical Computing* 8, 1 (Jan. 1987), S2–13.
- [7] Tony F. Chan. 1987. Rank revealing QR factorizations. *Linear Algebra and Its Applications* 88 (1987), 67–82.
- [8] James Demmel, Ioana Dumitriu, and Olga Holtz. 2007. Fast linear algebra is stable. *Numerische Mathematik* 108, 1 (2007), 59–91.
- [9] James W. Demmel. 1997. *Applied Numerical Linear Algebra*. SIAM, Philadelphia, PA.

- [10] Jack J. Dongarra, Iain S. Duff, Danny C. Sorensen, and Henk A. van der Vorst. 1991. *Solving Linear Systems on Vector and Shared Memory Computers*. SIAM, Philadelphia, PA.
- [11] Jed A. Duersch and Ming Gu. 2017. Randomized QR with column pivoting. *SIAM Journal on Scientific Computing* 39, 4 (2017), C263–291. DOI: <https://doi.org/10.1137/15M1044680>
- [12] Ricardo D. Fierro and Per Christian Hansen. 1997. Low-rank revealing UTV decompositions. *Numerical Algorithms* 15, 1 (1997), 37–55. DOI: <https://doi.org/10.1023/A:1019254318361>
- [13] Ricardo D. Fierro, Per Christian Hansen, and Peter Søren Kirk Hansen. 1999. UTV tools: Matlab templates for rank-revealing UTV decompositions. *Numerical Algorithms* 20, 2–3 (1999), 165–194.
- [14] Gene H. Golub and Charles F. Van Loan. 1996. *Matrix Computations* (3rd ed.). Johns Hopkins University Press, Baltimore, MD.
- [15] M. Gu. 2015. Subspace iteration randomization and singular value problems. *SIAM Journal on Scientific Computing* 37, 3 (2015), A1139–1173. DOI: <https://doi.org/10.1137/130938700>
- [16] Ming Gu and Stanley C. Eisenstat. 1996. Efficient algorithms for computing a strong rank-revealing QR factorization. *SIAM Journal on Scientific Computing* 17, 4 (1996), 848–869.
- [17] John A. Gunnels, Fred G. Gustavson, Greg M. Henry, and Robert A. van de Geijn. 2001. FLAME: Formal linear algebra methods environment. *ACM Transactions on Mathematical Software* 27, 4 (2001), 422–455. <http://doi.acm.org/10.1145/504210.504213>
- [18] Nathan Halko, Per-Gunnar Martinsson, Yoel Shkolnisky, and Mark Tygert. 2011. An algorithm for the principal component analysis of large data sets. *SIAM Journal on Scientific Computing* 33, 5 (2011), 2580–2594. DOI: <https://doi.org/10.1137/100804139>
- [19] Nathan Halko, Per-Gunnar Martinsson, and Joel A. Tropp. 2011. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review* 53, 2 (2011), 217–288.
- [20] Thierry Joffrain, Tze Meng Low, Enrique S. Quintana-Ortí, Robert van de Geijn, and Field G. Van Zee. 2006. Accumulating householder transformations, revisited. *ACM Transactions on Mathematical Software* 32, 2 (June 2006), 169–179. DOI: <https://doi.org/10.1145/1141885.1141886>
- [21] William Kahan. 1966. Numerical linear algebra. *Canadian Mathematical Bulletin* 9, 6 (1966), 757–801.
- [22] Edo Liberty, Franco Woolfe, Per-Gunnar Martinsson, Vladimir Rokhlin, and Mark Tygert. 2007. Randomized algorithms for the low-rank approximation of matrices. *Proceedings of the National Academy of Sciences of the United States of America* 104, 51 (2007), 20167–20172. DOI: <https://doi.org/10.1073/pnas.0709640104>
- [23] P. G. Martinsson. 2015. Blocked rank-revealing QR factorizations: How randomized sampling can be used to avoid single-vector pivoting. arXiv:1505.08115.
- [24] P. G. Martinsson, G. Quintana-Ortí, N. Heavner, and R. van de Geijn. 2017. Householder QR factorization with randomization for column pivoting (HQRPP). arXiv:1512.02671.
- [25] P. G. Martinsson, V. Rokhlin, and M. Tygert. 2006. On interpolation and integration in finite-dimensional spaces of bounded functions. *Communications in Applied Mathematics and Computational Science* 1, 1 (2006), 133–142.
- [26] P. G. Martinsson and S. Voronin. 2016. A randomized blocked algorithm for efficiently computing rank-revealing factorizations of matrices. *SIAM Journal on Scientific Computing* 38, 5 (2016), S485–507. DOI: <https://doi.org/10.1137/15M1026080>
- [27] Per-Gunnar Martinsson. 2016. Randomized methods for matrix computations and analysis of high dimensional data. *The Mathematics of Data*, M. W. Mahoney, J. C. Duchi, and A. C. Gilbert (Eds.). American Mathematical Society, 187–231.
- [28] Per-Gunnar Martinsson, Vladimir Rokhlin, and Mark Tygert. 2006. *A Randomized Algorithm for the Approximation of Matrices*. Technical Report YALEU/DCS/RR-1361. Computer Science Department, Yale University, New Haven, CT.
- [29] Per-Gunnar Martinsson, Vladimir Rokhlin, and Mark Tygert. 2011. A randomized algorithm for the decomposition of matrices. *Applied and Computational Harmonic Analysis* 30, 1 (2011), 47–68. DOI: <https://doi.org/10.1016/j.acha.2010.02.003>
- [30] R. Mathias and G. W. Stewart. 1993. A block {QR} algorithm and the singular value decomposition. *Linear Algebra Applications* 182 (1993), 91–100. DOI: [https://doi.org/10.1016/0024-3795\(93\)90493-8](https://doi.org/10.1016/0024-3795(93)90493-8)
- [31] Haesun Park and Lars Eldén. 1995. Downdating the rank-revealing URV decomposition. *SIAM Journal on Matrix Analysis and Applications* 16, 1 (1995), 138–155. DOI: <https://doi.org/10.1137/S0895479892241913>
- [32] Enrique S. Quintana, Gregorio Quintana, Xiaobai Sun, and Robert van de Geijn. 2001. A note on parallel matrix inversion. *SIAM Journal on Scientific Computing* 22, 5 (2001), 1762–1771.
- [33] Gregorio Quintana-Ortí, Xiaobai Sun, and Christian H. Bischof. 1998. A BLAS-3 version of the QR factorization with column pivoting. *SIAM Journal on Scientific Computing* 19, 5 (1998), 1486–1494. DOI: <https://doi.org/10.1137/S1064827595296732>
- [34] Vladimir Rokhlin, Arthur Szlam, and Mark Tygert. 2009. A randomized algorithm for principal component analysis. *SIAM Journal on Matrix Analysis and Applications* 31, 3 (2009), 1100–1124.

- [35] Robert Schreiber and Charles Van Loan. 1989. A storage-efficient WY representation for products of householder transformations. *SIAM Journal on Scientific and Statistical Computing* 10, 1 (Jan. 1989), 53–57.
- [36] G. W. Stewart. 1994. UTV decompositions. In *Numerical Analysis 1993: Proceedings of the 15th Dundee Conference, June–July 1993*, D. F. Griffiths and G. A. Watson (Eds.). Pitman Research Notes in Mathematics Series. Longman, 225.
- [37] G. W. Stewart. 1998. *Matrix Algorithms Volume 1: Basic Decompositions*. SIAM, Philadelphia, PA.
- [38] G. W. Stewart. 1999. The QLP approximation to the singular value decomposition. *SIAM Journal on Scientific Computing* 20, 4 (1999), 1336–1348.
- [39] G. W. Stewart. 1992. An updating algorithm for subspace tracking. *IEEE Transactions on Signal Processing* 40, 6 (June 1992), 1535–1541. DOI : <https://doi.org/10.1109/78.139256>
- [40] G. W. Stewart. 1993. Updating a rank-revealing ULV decomposition. *SIAM Journal on Matrix Analysis and Applications* 14, 2 (1993), 494–499. DOI : <https://doi.org/10.1137/0614034>
- [41] Lloyd N. Trefethen and David Bau III. 1997. *Numerical Linear Algebra*. Vol. 50. SIAM, Philadelphia, PA.
- [42] Field G. Van Zee. 2012. libflame: *The Complete Reference*. Available at <https://www.lulu.com>.
- [43] Field G. Van Zee, Ernie Chan, Robert van de Geijn, Enrique S. Quintana-Ortí, and Gregorio Quintana-Ortí. 2009. The libflame library for dense matrix computations. *IEEE Computation in Science and Engineering* 11, 6 (2009), 56–62.
- [44] S. Voronin and P. G. Martinsson. 2016. Efficient algorithms for cur and interpolative matrix decompositions. *Advances in Computational Mathematics* 43, 3 (2016), 495–516. DOI : <https://doi.org/10.1007/s10444-016-9494-8>
- [45] Rafi Witten and Emmanuel Candes. 2015. Randomized algorithms for low-rank matrix factorizations: Sharp performance bounds. *Algorithmica* 72, 1 (2015), 264–281.

Received March 2017; revised March 2018; accepted July 2018