arXiv:2002.07138v1 [math.NA] 17 Feb 2020

PASS-EFFICIENT RANDOMIZED LU ALGORITHMS FOR COMPUTING LOW-RANK MATRIX APPROXIMATION

BOLONG ZHANG* AND MICHAEL MASCAGNI†

Abstract. Low-rank matrix approximation is extremely useful in the analysis of data that arises in scientific computing, engineering applications, and data science. However, as data sizes grow, traditional low-rank matrix approximation methods, such as singular value decomposition (SVD) and column pivoting QR decomposition (CPQR), are either prohibitively expensive or cannot provide sufficiently accurate results. A solution is to use randomized low-rank matrix approximation methods such as randomized SVD , and randomized LU decomposition on extremely large data sets. In this paper, we focus on the randomized LU decomposition method. First, we employ a reorthogonalization procedure to perform the power iteration of the existing randomized LU algorithm to compensate for the rounding errors caused by the power method. Then we propose a novel randomized LU algorithm, called PowerLU, for the fixed low-rank approximation problem. PowerLU allows for an arbitrary number of passes of the input matrix, $v \geq 2$. Recall that the existing randomized LU decomposition only allows an even number of passes. We prove the theoretical relationship between PowerLU and the existing randomized LU. Numerical experiments show that our proposed PowerLU is generally faster than the existing randomized LU decomposition, while remaining accurate. We also propose a version of PowerLU, called PowerLU_FP, for the fixed precision low-rank matrix approximation problem. PowerLU_FP is based on an efficient blocked adaptive rank determination Algorithm 4.1 proposed in this paper. We present numerical experiments that show that PowerLU_FP can achieve almost the same accuracy and is faster than the randomized blocked QB algorithm by Martinsson and Voronin. We finally propose a single-pass algorithm based on LU factorization. Tests show that the accuracy of our single-pass algorithm is comparable with the existing single-pass algorithms.

Key words. randomized numerical linear algebra, low-rank matrix approximation, randomized SVD, randomized LU

AMS subject classifications. 65F99, 65C99

1. Introduction.

In research areas such as data mining, scientific computing, and many engineering applications, the matrices encountered are extremely large. In [35], the authors state that big data matrices are generally low-rank. Low-rank matrix approximation is a critical technique used to analyze data in many disciplines that have extremely large matrices [21]. Applications of randomized low-rank approximation including imaging processing [10], data mining [7], and machine learning [3] have already been widely explored. Mathematically low-rank approximation of a matrix can be defined as follows.

DEFINITION 1.1. *For a given matrix* $\mathbf{A} \in \mathbb{R}^{m\times n}$ *with rank* k*, we seek two low-rank matrices* $\mathbf{B} \in \mathbb{R}^{m\times k}$*, and* $\mathbf{C} \in \mathbb{R}^{k\times n}$ *with given rank* $k \ll min(m, n)$ *such that the following norm equation (1.1) holds:*

$$\|\mathbf{A} - \mathbf{BC}\|_N \ll 1. \tag{1.1}$$

We want to approximate $\mathbf{A}$ *so that* $\|\mathbf{A} - \mathbf{BC}\|_N$ *as small as possible. We can use any matrix norm, but it customary to use either the spectral norm,* $N = 2$*, or the Frobenius norm,* $N = F$ *in numerical linear algebra.*

The above definition (1.1) is called the fixed rank problem in low-rank approximation, where we know the rank, k, of the matrix in advance. However, in some cases, the rank of the matrix cannot be known in advance, and instead of giving us the rank, k,

*Department of Computer Science, Florida State University, FL, USA (bzhang@cs.fsu.edu).

†Department of Computer Science, Florida State University, FL and Applied and Computational Mathematics Division, NIST, Gaithersburg, MD, USA (mascagni@fsu.edu).

a real-valued tolerance ϵ is provided for the norm equation (1.1). In the other words, we want to find the minimum rank k for $\mathbf{B}$ and $\mathbf{C}$, such that equation (1.1) holds for the given tolerance, ϵ. This is another type of low-rank approximation problem, called fixed precision.

The singular value decomposition (SVD) is a common deterministic method to compute the low-rank matrix approximation. The famous theorem by Eckart and Young [9] states that the SVD can achieve optimal low-rank approximation results in both the spectral and Frobenius norms. Suppose the SVD decomposition for matrix $\mathbf{A} \in \mathbb{R}^{m\times n}, m \geq n$ is the following:

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \tag{1.2}$$

where $\mathbf{U} \in \mathbb{R}^{m\times n}$ has orthonormal columns, $\mathbf{V} \in \mathbb{R}^{n\times n}$ is an orthogonal matrix, and $\mathbf{\Sigma} \in \mathbb{R}^{n\times n}$ is a diagonal matrix whose diagonal entries are the singular values of $\mathbf{A}$ in descending order. Then the optimal solution to the fixed rank problem (1.1) for a given rank k is $\mathbf{A}_k = \mathbf{BC}$, where we set $\mathbf{B} = \mathbf{U}(:, 1:k)\mathbf{\Sigma}(1:k, 1:k)$ and $\mathbf{C} = \mathbf{V}(:, 1:k)^T$[1]. This gives us the explicit matrix norms of our rank k approximation errors as:

$$\|\mathbf{A} - \mathbf{A}_k\|_2 = \sigma_{k+1} \quad \text{and} \quad \|\mathbf{A} - \mathbf{A}_k\|_F = \sqrt{\sum_{i=k+1}^{n} \sigma_i^2}. \tag{1.3}$$

However, classical algorithms to calculate the SVD are extremely costly. So it is impractical and often impossible to perform the SVD on a sufficiently large matrix due to the prohibitive time complexity and current machine constraints. There exists an alternative method, which is obtained from column pivoting QR (CPQR). However, CPQR can not guarantee the accuracy of the low-rank approximation. If we allow the diagonal matrix $\mathbf{\Sigma}$ in equation (1.2) to be an upper triangular matrix, $\mathbf{R}$, we obtain the URV algorithm. If instead we choose a lower triangular matrix, $\mathbf{L}$, for the decomposition, we get the ULV algorithm sometimes called the QLP algorithm, [29, 30]. This is a middle ground, as URV is faster than SVD yet more accurate than CPQR. Besides URV, there exist other deterministic methods such as rank-revealing QR (RRQR) [14] and rank-revealing LU (RRLU) [23, 25] for low-rank approximation. Both RRQR and RRLU use permutations of the matrix columns and rows, which will finally produces a leading submatrix that captures most of the product of the singular values of the original matrix. RRQR and RRLU are expensive; however, they do not directly compute a low-rank approximation of the matrix [1].

Randomized low-rank approximation algorithms are a relatively recent development [15]. Compared with the deterministic methods, such as SVD, RRQR and RRLU, randomized methods are usually faster and maintain high accuracy. Generally, there are two classes of randomized low-rank approximation algorithms: sampling-based algorithms and random projection-based algorithms. Sampling algorithms use randomly selected columns or rows based on sampling probabilities derived from the original matrix. One then performs a deterministic algorithm, such as SVD, on the smaller subsampled problem [4, 5, 6, 8, 16, 20, 21]. In contrast to the sampling-based algorithms, the core idea behind the random projection-based methods is to project the high-dimensional space spanned by the columns of the matrix into a low-dimensional space where deterministic methods can be inexpensively applied. In this

[1] We use MATLAB matrix indexing notation here. For a given matrix $\mathbf{A}$, $\mathbf{A}(:, 1:k)$ and $\mathbf{A}(1:k, :)$ extract the first k columns and rows of $\mathbf{A}$, respectively.

paper, we mainly discuss the random projection-based methods. Formally, for a given matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, we multiply it with a random Gaussian matrix $\mathbf{\Omega} \in \mathbb{R}^{n\times l}$, where $l = k + q$, k is the target rank and q is the oversampling parameter. Then $\mathbf{Y} = \mathbf{A\Omega}$ will capture the most action (i.e. capture l largest singular values) of $\mathbf{A}$ and more importantly the dimension of the space spanned by the columns of $\mathbf{Y}$ is much smaller. Suppose $\mathbf{Q} \in \mathbb{R}^{m\times l}$ is the orthogonal basis for the approximate space $\mathbf{Y}$. Then we have the following random QB low-rank approximation

$$\mathbf{A} \approx \mathbf{Q}\mathbf{Q}^T\mathbf{A} = \mathbf{Q}\mathbf{B}. \tag{1.4}$$

Here $\mathbf{B} = \mathbf{Q}^T\mathbf{A} \in \mathbb{R}^{k\times n}$, which is a smaller matrix on which we could perform some standard factorization such as the SVD.

Algorithm 1.1 A Basic Randomized SVD Algorithm

Input: Matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, a target rank k, oversampling parameter q and $l = k+q$, $p \geq 0$.

Output: Orthogonal matrix $\mathbf{U} \in \mathbb{R}^{m\times n}$, $\mathbf{V} \in \mathbb{R}^{n\times n}$ and diagonal $\mathbf{\Sigma} \in \mathbb{R}^{n\times n}$ in an approximate rank-(l) SVD of $\mathbf{A}$, such that $\mathbf{A} \approx \mathbf{U\Sigma V}^T$

1: Draw a random $n \times l$ test matrix $\mathbf{\Omega}$;
2: Form the matrix product $\mathbf{Y} = (\mathbf{A}\mathbf{A}^T)^p\mathbf{A\Omega}$;
3: Perform QR decomposition to obtain orthonormal basis $\mathbf{Q} = qr(\mathbf{Y})$; // $\mathcal{C}_{qr}ml^2$
4: Form the $l \times n$ matrix $\mathbf{B} = \mathbf{Q}^T\mathbf{A}$; // $\mathcal{C}_{mm}mnl$
5: Form the SVD of the small matrix $\mathbf{B} : \mathbf{B} = \hat{\mathbf{U}}\mathbf{\Sigma V}^T$; // $\mathcal{C}_{svd}nl^2$
6: Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$; // $\mathcal{C}_{mm}ml^2$

For a matrix whose singular values decay slowly, The above procedure may not provide good results. However, the power iteration $(\mathbf{A}\mathbf{A}^T)^p\mathbf{A}$ can be applied to solve this problem, where p is the exponent of the power iteration. Suppose $\mathbf{A}$ has the SVD decomposition (1.2), then

$$(\mathbf{A}\mathbf{A}^T)^p\mathbf{A} = \mathbf{U}\mathbf{\Sigma}^{2p+1}\mathbf{V}^T. \tag{1.5}$$

Equation (1.5) implies that the singular values of the matrix $(\mathbf{A}\mathbf{A}^T)^p\mathbf{A}$ decay faster than the matrix $\mathbf{A}$. However, both of them have the same left and right eigenvectors. A basic randomized SVD algorithm (RandSVD) is shown here as Algorithm 1.1 [15]. However, to avoid the rounding error of float point arithmetic obtained from performing the power iteration, reorthogonalization is needed. A typical way to perform the power iteration reorthogonalization is shown in Algorithm 1.2 [15, 36], where a reduced QR decomposition is applied in each reorthogonalization. An optimized version of reorthogonalization was proposed in [19], where the authors used LU factorization to replace QR every time except the last iteration. To solve the fixed precision low-rank approximation problem, the adaptive version of the randomized QB algorithm has been proposed to find the rank incrementally in [15, 22] and was optimized in [38]. For the RandSVD Algorithm 1.1, we need to access the matrix an even number times. In some cases, the input matrix is enormous, and accessing the matrix is very expensive. Randomized SVD algorithms that minimize accesses of the matrix have been proposed in [2]. For the extremely large matrices, single-pass (single access) algorithms have also been studied recently. A general scheme for the single-pass algorithm was proposed in [31]. With the single-pass algorithm, the matrix can processed via its random Gaussian projection. In [15], the authors propose a single-pass algorithm, which uses two random Gaussian matrices to compress the original matrix

and finally solve a linear equation. However, the single-pass algorithms in [15, 31] are less accurate, to some extent. For a matrix in row or column major format, a more accurate algorithm has been proposed in [38].

Algorithm 1.2 Power Iteration Reorthogonalization for Algorithm 1.1

Input: Given a $\mathbf{A} \in \mathbb{R}^{m\times n}$, $l < min(m,n)$ and $p \geq 0$.
Output: Orthonormal basis $\mathbf{Q} \in \mathbb{R}^{m\times l}$ for input matrix $\mathbf{A}$

1: Generate random Gaussian matrix $\mathbf{\Omega}$ with size $n \times l$.
2: $[\mathbf{Q}, \sim] = qr(\mathbf{A\Omega})$ // $\mathcal{C}_{qr}mnl + \mathcal{C}_{qr}ml^2$
3: **for** $i = 1 : 1 : p$ **do**
4: $\quad[\mathbf{Q}, \sim] = qr(\mathbf{A^T Q})$ // $\mathcal{C}_{qr}mnl + \mathcal{C}_{qr}nl^2$
5: $\quad[\mathbf{Q}, \sim] = qr(\mathbf{AQ})$ // $\mathcal{C}_{qr}mnl + \mathcal{C}_{qr}ml^2$
6: **end for**

A randomized LU factorization algorithm (RandLU) for low-rank matrix approximation based on the RandSVD was first proposed in [28]. The authors give two reasons for their motivation for the RandLU decomposition as: (a) compared with SVD, LU decomposition is usually faster and also very efficient for sparse matrices with computation time related to the nonzero elements; (b) LU decomposition can be fully parallelized, which makes it very efficient in computation on modern hardware such as the GPU [28]. The overall algorithmic flow of RandLU is shown in the Algorithm 2.1. In [28], the authors showed the efficiency and accuracy of the algorithm [2]. However, costly matrix computations such as the pseudoinverse are needed in RandLU. In [28], the authors demonstrated the speed superiority of the RandLU Algorithm 2.1 over RandSVD with relatively small matrix size. Our tests show that RandLU can lose its advantage over RandSVD with larger matrices. Besides, RandLU can only solve the fixed-rank low-rank approximation problem. For the fixed precision low-rank approximation problem, RandLU will not work. Also, RandLU is not a single-pass algorithm.

In this paper, we propose a pass-efficient randomized LU algorithm, called PowerLU, to solve the fixed-rank low-rank approximation problem. Compared with RandLU, PowerLU is usually faster and also maintains high accuracy. Moreover, PowerLU can access the matrix, $\mathbf{A}$, an arbitrary number of times. Next we propose a version of PowerLU, called PowerLU_FP, for the fixed precision low-rank approximation problem. PowerLU_FP is based on an efficient blocked adaptive rank determination Algorithm 4.1 proposed in this paper. Lastly, we propose a single-pass algorithm for low-rank approximation based on LU decomposition under the assumption that the matrix is stored in column-major or row-major order. Experiments show that our proposed single-pass algorithm can achieve better results compared with the single-pass algorithm in [15].

The rest of this paper is organized as follows. In §2, we define some basic notation and review the existing RandLU algorithm. We then further discuss the RandSVD and RandLU algorithms' commonality. Then we describe our proposed PowerLU algorithm and it's analysis in §3. In §4, we explore a version of PowerLU algorithm, called PowerLU_FP, for the fixed precision low-rank approximation problem. In §5, the single-pass algorithm is discussed. In §6, we present the numerical results from numerous experiments to demonstrate the efficiency of our proposed algorithms.

[2]Code available at [27]

2. Technical Preliminaries and Related Work. In this section, we first introduce some linear algebra basics needed in our paper. Then we briefly review the RandLU decomposition in [28]. To avoid the round-off error caused by computing the power iteration in RandLU, we employ a practical reorthogonalization procedure. Finally, we give a discussion of the RandSVD and RandLU methods. Throughout the paper, we use the following notation: for any matrix $\mathbf{A}$, $\|\mathbf{A}\|$ denotes the spectral norm by default, which is the largest singular value of $\mathbf{A}$. We use $\|\mathbf{A}\|_F$ for the Frobenius norm, which is $(\sum_{i,j}|a_{ij}|^2)^{1/2}$. σ_k signifies the kth largest singular value of $\mathbf{A}$. To describe algorithms in this paper, we use MATLAB notation in our pseudocode, where "$lu(\cdot)$" and "$qr(\cdot)$" to denote the MATLAB builtin LU and QR factorizations.

2.1. Linear Algebra Basics. In this section, we will briefly review the definitions and properties of the orthogonal projection and the pseudoinverse [12].

Orthogonal Projection. An orthogonal projection is a linear transformation from a vector space to itself. For a matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, we use $Range(\mathbf{A})$ to denote the space spanned by the columns of $\mathbf{A}$. Suppose $\mathbf{A}$ has full column rank, then we denote the orthogonal projection of $Range(\mathbf{A})$ as $\mathbf{P_A}$, which is defined as follows:

$$\mathbf{P_A} = \mathbf{A}(\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T. \tag{2.1}$$

It is easy to prove the following proprieties for $\mathbf{P_A}$:

- $\mathbf{P_A^2} = \mathbf{P_A}$.
- $Range(\mathbf{A}) = Range(\mathbf{P_A})$.
- $\mathbf{P_A} = \mathbf{A}\mathbf{A}^T$, if $\mathbf{A}$ has orthonormal columns.

Pseudoinverse. In linear algebra, the pseudoinverse of a matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$ is the generalization of the inverse for non-square matrices. It is customary to denote the pseudoinverse of $\mathbf{A} \in \mathbb{R}^{m\times n}$ with a dagger, $\mathbf{A}^\dagger$. If $m \geq n$ and $\mathbf{A} \in \mathbb{R}^{m\times n}$ has full column rank, then the pseudoinverse of $\mathbf{A}^{m\times n}$ can be computed as

$$\mathbf{A}^\dagger = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T, \tag{2.2}$$

where $\mathbf{A}^\dagger \in \mathbb{R}^{n\times m}$. $\mathbf{A}^\dagger \in \mathbb{R}^{n\times m}$ has the following properties:

- $(\mathbf{A}^\dagger)^T = (\mathbf{A}^T)^\dagger$.
- If $\mathbf{A} \in \mathbb{R}^{m\times n}$ has full column rank, then $\mathbf{P}_A = \mathbf{A}\mathbf{A}^\dagger$ is an orthogonal projection.

2.2. Randomized LU Decomposition. LU factorization is a fundamental problem in numerical linear algebra, and plays an important role in solving systems of linear equations. Compared with the QR and SVD decompositions, LU is generally faster. In [28], the authors developed RandLU decomposition for the fixed rank low-rank approximation problem. RandLU decomposition can be applied to problems such as rank deficient least square, image reconstruction, and dictionary construction [28, 26].

The RandLU algorithm is shown in Algorithm 2.1, which produces matrices $\mathbf{P}, \mathbf{Q}, \mathbf{L}, \mathbf{U}$ satisfying Theorem 2.3, where $\mathbf{P}, \mathbf{Q}$ are permutation matrices, $\mathbf{L}$ is a lower triangular matrix and $\mathbf{U}$ is an upper triangular matrix when $p = 0$, where p is the exponent of the power iteration in step 2. When $p > 0$, power iteration makes the RandLU algorithm more accurate for a matrix with slowly decaying singular values. However, if we compute step 2 in floating-point arithmetic, rounding errors will overwhelm the small singular values compared to the spectral norm $\|\mathbf{A}\|$ [15, 38]. To solve

this problem, reorthogonalization with reduced QR when performing the power iteration was used [15]. In [19], an accelerated power iteration was proposed by replacing the QR with partial pivoting LU except for the last iteration. However, for randomized LU, there is no need for an orthonormal basis at the end, so one need not do the QR factorization at the end. In this paper, we propose a power iteration reorthogonalization variant, which is shown as Algorithm 2.2, for the RandLU Algorithm 2.1. In practice, we replace steps 1-3 in RandLU Algorithm 2.1 with Algorithm 2.2. In §6, we show that with reorthogonalization RandLU can achieve better accuracy in practice.

Remark 2.1. In Algorithm 2.1, we use partial pivoting LU instead of applying RRLU [25] to $\mathbf{Y}$. The authors pointed out that partial pivoting LU works well for most cases, and they also use partial pivoting LU in their own implementation [27]. In Algorithm 2.1, we need to multiply $\mathbf{PA}$ by the pseudoinverse of $\mathbf{L}_y$. We use (2.2) to compute $\mathbf{L}_y^\dagger$.

Remark 2.2. Power iteration has been applied to RandSVD, RandLU and our proposed PowerLU Algorithms. Power iteration reorthogonalization can remove rounding error caused by floating point arithmetic. However, in theory, reorthogonalization algorithms cannot improve the accuracy of randomized algorithms if exact arithmetic is used. The accuracy of the randomized algorithms is related the number of passes of the matrix, which in turn speeds the decay of the singular values.

THEOREM 2.3 (**Error Bound for RandLU**). *For a given matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, then, its randomized LU produced by Algorithm 2.1 with integers k , $p = 0$ and $\ell(\ell \geq k)$ satisfies:*

$$\|\mathbf{LU} - \mathbf{PAQ}\| \leq \left(2\sqrt{2n\ell\beta^2\gamma^2 + 1} + 2\sqrt{2n\ell}\beta\gamma(k(n-k)+1)\right)\sigma_{k+1}(\mathbf{A}), \tag{2.3}$$

with probability greater than or equal to

$$\Theta = 1 - \frac{1}{\sqrt{2\pi(\ell-k+1)}}\left(\frac{e}{(\ell-k+1)\beta}\right)^{\ell-k+1} - \frac{1}{4(\gamma^2-1)\sqrt{\pi n\gamma^2}}\left(\frac{2\gamma^2}{e^{\gamma^2-1}}\right)^{n}, \tag{2.4}$$

where $\beta > 0$ and $\gamma > 1$.

This is a result from [28], and by probability we mean the measure on the random Gaussian matrix used in their algorithm. We believe that we can prove similar bounds for our proposed algorithms in this paper based on our Theorem 3.4.

To describe the time complexity for the algorithms listed in this paper, we use the same notation as in [22]: let $\mathcal{C}_{mm}, \mathcal{C}_{lu}, \mathcal{C}_{qr}$, $\mathcal{C}_{svd}$ and $\mathcal{C}_{inv}$ denote the scaling constants for the cost of executing a matrix-matrix multiplication, partial pivoting LU factorization, a full QR factorization, an SVD factorization, and a matrix inversion respectively. Specifically, we have:

- Multiplying two matrices of size $m \times n$ and $n \times r$ costs $\mathcal{C}_{mm}mnr$
- Performing a partial pivoting LU factorization of a matrix of size $m \times n$, with $m \geq n$, costs $\mathcal{C}_{lu}mn^2$
- Performing a full QR factorization of a matrix of size $m \times n$, with $m \geq n$, costs $\mathcal{C}_{qr}mn^2$
- Performing an SVD factorization of a matrix of size $m \times n$, with $m \geq n$, costs $\mathcal{C}_{svd}mn^2$
- Inverting a matrix of size $n \times n$, costs $\mathcal{C}_{inv}n^3$

Algorithm 2.1 Randomized LU Decomposition

Input: Matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, desired rank k, $l \geq k$ number of columns to use, and $p \geq 0$.

Output: Matrix: $\mathbf{P} \in \mathbb{R}^{m \times n}, \mathbf{Q} \in \mathbb{R}^{n \times n}, \mathbf{L} \in \mathbb{R}^{m \times n}, \mathbf{U} \in \mathbb{R}^{n \times n}$ such that $\mathbf{PAQ} \approx \mathbf{LU}$, where $\mathbf{P}, \mathbf{Q}$ are orthogonal permutation matrices, $\mathbf{L}$ and $\mathbf{U}$ are lower and upper triangular matrices, respectively.

1: Generate a randomized Gaussian matrix $\mathbf{\Omega}$ with size $n \times l$;
2: $\mathbf{Y} = \mathbf{A}(\mathbf{A}^T\mathbf{A})^p\mathbf{\Omega}$;
3: $[\mathbf{L}_y, \mathbf{U}_y, \mathbf{P}] = lu(\mathbf{Y})$; // $\mathcal{C}_{lu}ml^2$
4: Truncate $\mathbf{L}_y$ and $\mathbf{U}_y$ by choosing the first k columns and first k rows, respectively, such that $\mathbf{L}_y = \mathbf{L}_y(:, 1:k)$ and $\mathbf{U}_y = \mathbf{U}_y(1:k, :)$;
5: $\mathbf{B} = \mathbf{L}_y^{\dagger}\mathbf{PA}$; // $\mathcal{C}_{mm}(mnk + 2mk^2) + \mathcal{C}_{inv}k^3$
6: Apply LU decomposition to $\mathbf{B}$ with column pivoting $\mathbf{BQ} = \mathbf{L}_b\mathbf{U}_b$;
7: $\mathbf{L} = \mathbf{L}_y\mathbf{L}_b$; // $\mathcal{C}_{mm}mk^2$
8: $\mathbf{U} = \mathbf{U}_b$;

Algorithm 2.2 Power Iteration Reorthogonalization for Algorithm 2.1

Input: Given a $\mathbf{A} \in \mathbb{R}^{m \times n}$, $l < min(m, n)$ and $p \geq 0$.

Output: $\mathbf{P} \in \mathbb{R}^{m \times m}$, $\mathbf{L} \in \mathbb{R}^{m \times l}$ and $\mathbf{U} \in \mathbb{R}^{l \times l}$ such that $\mathbf{P}^T\mathbf{LU} = \mathbf{A}(\mathbf{A}^T\mathbf{A})^p\mathbf{\Omega}$, where $\mathbf{\Omega}$ is the random Gaussian matrix generated in this algorithm.

1: Generate random Gaussian matrix $\mathbf{\Omega}$ with size $n \times l$;
2: $[\mathbf{L}, \mathbf{U}, \mathbf{P}] = lu(\mathbf{A\Omega})$; // $\mathcal{C}_{mm}mnl + \mathcal{C}_{lu}ml^2$
3: **if** $p > 0$ **then**
4: $\quad \mathbf{L} = \mathbf{P}^T\mathbf{L}$;
5: $\quad$ **for** $i = 1:1:p$ **do**
6: $\quad\quad [\mathbf{L}, \sim] = lu(\mathbf{A}^T\mathbf{L})$; // $\mathcal{C}_{mm}mnl + \mathcal{C}_{lu}nl^2$
7: $\quad\quad$ **if** $i == p$ **then**
8: $\quad\quad\quad [\mathbf{L}, \mathbf{U}, \mathbf{P}] = lu(\mathbf{AL}))$; // $\mathcal{C}_{mm}mnl + \mathcal{C}_{lu}ml^2$
9: $\quad\quad$ **else**
10: $\quad\quad\quad [\mathbf{L}, \sim] = lu(\mathbf{AL})$; // $\mathcal{C}_{mm}mnl + \mathcal{C}_{lu}ml^2$
11: $\quad\quad$ **end if**
12: $\quad$ **end for**
13: **end if**

Here, we compare the time complexity for RandLU and RandSVD. Reorthogonalization of the power iteration will be employed to reduce the rounding error. So for both RandLU and RandSVD, we analyze their time complexity together with their corresponding power iteration reorthogonalization variants. It is easy to see the time complexity for the RandLU Algorithm 2.1, where the power iteration is computed by Algorithm 2.2, is given by the formula (2.5) with $p \geq 0$ as following:

$$\begin{aligned} \mathcal{C}_{RandLU} \sim \mathcal{C}_{mm}mnl + \mathcal{C}_{lu}ml^2 + p \cdot \left(2\mathcal{C}_{mm}mnl + \mathcal{C}_{lu}ml^2 + \mathcal{C}_{lu}nl^2\right) + \\ \mathcal{C}_{mm}(mnk + 2mk^2) + \mathcal{C}_{inv}k^3 + \mathcal{C}_{mm}mk^2. \end{aligned} \tag{2.5}$$

According to Algorithm 1.1, the time complexity for RandSVD Algorithm 1.1, where the power iteration is computed by Algorithm 1.2 [15] is as follows:

$$\begin{aligned} \mathcal{C}_{RandSVD} \sim \mathcal{C}_{mm}mnl + \mathcal{C}_{qr}ml^2 + p \cdot \left(2\mathcal{C}_{mm}mnl + \mathcal{C}_{qr}ml^2 + \mathcal{C}_{qr}nl^2\right) + \\ \mathcal{C}_{mm}mnl + \mathcal{C}_{svd}nl^2 + \mathcal{C}_{mm}ml^2. \end{aligned} \tag{2.6}$$

It is hard to tell which of the two is larger, (2.5) or (2.6). When $p = 0$,the authors [28] showed that RandLU is faster than RandSVD with fixed matrix size $m = n = 3000$, and various target rank values. Our own tests with the same matrix size confirmed this. However, our tests show that RandLU will lose this advantage over RandSVD when we use larger matrices. Details can be found in §6.

So far, we have introduced RandLU factorization shown in [28], and we employed the power iteration reorthogonalization for RandLU as shown in Algorithm 2.2. RandLU can solve the fixed rank low-rank approximation problem for a given rank. However, in some cases, it is impossible to obtain the rank of the matrix in advance, so now we will explore algorithms for the fixed precision low-rank approximation problem based on LU factorization.

2.3. Discussion. In this section, we develop the intuition to help understand our proposed randomized LU algorithm. We previously discussed the RandSVD and RandLU algorithms. And we will use our understanding of RandSVD Algorithm 1.1 to examine RandLU Algorithm 2.1.

We can organize the RandSVD Algorithm 1.1 into two stages as follows:

1. Randomized Stage: Find an orthonormal basis $\mathbf{Q} \in \mathbb{R}^{m\times l}$ of the column space of $\mathbf{A} \in \mathbb{R}^{m\times n}$ via random projection. Then we approximate $\mathbf{A} \approx \mathbf{P_Q A} = \mathbf{Q}\mathbf{Q}^T\mathbf{A}$.
2. Deterministic SVD Stage: Let $\mathbf{B} = \mathbf{Q}^T\mathbf{A}$, then perform a deterministic SVD on $\mathbf{B}$, which is much smaller than $\mathbf{A}$.

RandLU is organized similarly. First, RandLU produces a lower triangular matrix, $\mathbf{L}_y$. This matrix is used with its psuedoinverse to create an orthogonal projection (a permuted) of matrix $\mathbf{A}$. This is analogous to the RandSVD randomization stage:

$$\mathbf{PA} \approx \mathbf{L}_y\mathbf{L}_y^\dagger\mathbf{PA}. \tag{2.7}$$

Then partial pivoting LU is performed on $\mathbf{B} = \mathbf{L}_y^\dagger\mathbf{PA}$, where $\mathbf{B}$ is much smaller than $\mathbf{A}$. This is the analog of the deterministic stage in RandSVD.

However, it is costly to compute equation (2.7) in RandLU, since the pseudoinverse of a lower triangular matrix is needed. In Algorithm 2.1, step 2, power iteration is performed as $\mathbf{A}(\mathbf{A}^T\mathbf{A})^p\mathbf{\Omega}$. As we discussed, to eliminate rounding errors in power iteration, a reorthogonalization procedure is necessary. As Algorithm 3.2 shows, we can obtain an orthonormal basis $\mathbf{V} \in \mathbb{R}^{n\times l}$ for $\mathbf{A}^T$ by performing reorthogonalization on $(\mathbf{A}^T\mathbf{A})^p\mathbf{\Omega}$ when $p \geq 1$. Then we get the approximation

$$\mathbf{A} \approx \mathbf{A}\mathbf{V}\mathbf{V}^T. \tag{2.8}$$

In this paper, we present a randomized LU algorithm, called PowerLU, which is based on the approximation (2.8). We named our proposed algorithm PowerLU because the power iteration exponent $p \geq 1$ is required by our algorithm. The basic idea of PowerLU is that we perform partial pivoting LU on $\mathbf{AV}$ to get lower and upper triangular matrices $\mathbf{L}_1$, $\mathbf{U}_1$. Let $\mathbf{B} = \mathbf{U}_1\mathbf{V}^T$, then we perform another partial pivoting LU on $\mathbf{B}^T$. PowerLU does not need to compute a pseudoinverse and so is generally faster than RandLU. Another advantage of using (2.8) is that if $\mathbf{V}$ has sufficiently long columns, we can adaptively determine the rank, $k \leq l$, of $\mathbf{A}$ quite easily. We will discuss the details in the upcoming sections.

3. PowerLU: A Method for the Fixed Rank Low-Rank Matrix Approximation. In this section, we discuss a new randomized LU factorization, called

PowerLU, for the fixed rank low-rank matrix approximation problem (1.1). Our proposed method is based on (2.8). It is generally faster than the RandLU and RandSVD without loss of accuracy. In some cases, accessing the matrix $\mathbf{A}$ can be expensive. RandLU and RandSVD using power iteration access the matrix an even number of times. Our proposed PowerLU allows an arbitrary number of matrix accesses, which can reduce computation time when an odd number of accesses is sufficient.

3.1. PowerLU: A Low-Rank Matrix Approximation Method. In this section, we will present our proposed PowerLU method in detail.

The PowerLU algorithm is shown in Algorithm 3.1. Given a matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, suppose we want to find a low-rank approximation for $\mathbf{A}$ with target rank k. First we generate a random Gaussian matrix $\mathbf{\Omega} \in \mathbb{R}^{n\times l}$ where $l \geq k$. In practice we choose $l = k + q$, where $q = 5$ or 10 is the oversampling size. This is needed to ensure that the final rank is at least k with high probability.

In next step, to deal with a matrix with singular values decaying slowly, we compute an orthonormal basis $\mathbf{V} \in \mathbb{R}^{n\times l}$ by using QR decomposition on

$$(\mathbf{A}^T\mathbf{A})^p\mathbf{\Omega} = \mathbf{VZ}. \tag{3.1}$$

So $\mathbf{A}$ can be approximated by

$$\mathbf{A} \approx \mathbf{AVV}^T. \tag{3.2}$$

Then we form the partial pivoting LU decomposition of $\mathbf{AV}$

$$\mathbf{L_1U_1} = \mathbf{PAV}, \tag{3.3}$$

where $\mathbf{P}$ is a permutation matrix. Since $\mathbf{L_1U_1}$ derive from $\mathbf{AV}$, we can obtain an approximation for $\mathbf{PA}$ as

$$\mathbf{PA} \approx \mathbf{PAVV}^T = \mathbf{L_1U_1V}^T. \tag{3.4}$$

Let $\mathbf{B} = \mathbf{U}_1\mathbf{V}^T$, and we use standard partial pivoting LU on $\mathbf{B}^T$ to obtain

$$\mathbf{QB}^T = \mathbf{L}_2\mathbf{U}_2. \tag{3.5}$$

Combining the above equations (3.4) and (3.5), we obtain

$$\mathbf{PAQ} \approx \mathbf{LU}, \tag{3.6}$$

where $\mathbf{L} = \mathbf{L}_1\mathbf{U}_2^T$ and $\mathbf{U} = \mathbf{L}_2^T$.

In (3.5), we form $\mathbf{B}$ through matrix-matrix multiplication. Compared with RandLU, PowerLU is does not compute a pseudoinverse, and so is usually much faster. In step 2 of the PowerLU Algorithm 3.1, we need to perform the power iteration. As we discussed in §2.2, we might lose some accuracy if we calculate it without reorthogonalization. In practice we use Algorithm 3.2 to replace steps 1-2 in Algorithm 3.1. In Algorithm 3.1, we need to obtain the orthonormal basis $\mathbf{V}$. Therefore, Algorithm 3.2 for PowerLU is different with Algorithm 2.2.

3.2. Accuracy of Algorithm 3.1. The accuracy of the PowerLU Algorithm 3.1 is the approximation error obtained from (3.2). Theorem 4.1 in [2] gives us an error bound for (3.2), which we state here as Theorem 3.1.

Algorithm 3.1 The PowerLU Algorithm

Input: Given $\mathbf{A} \in \mathbb{R}^{m\times n}$, desired rank k, oversampling parameter q and $l = k + q$, and $p \geq 1$.

Output: Matrix: $\mathbf{P}, \mathbf{Q}, \mathbf{L}, \mathbf{U}$ such that $\mathbf{PAQ} \approx \mathbf{LU}$, where $\mathbf{P}, \mathbf{Q}$ are orthogonal permutation matrices, $\mathbf{L} \in \mathbb{R}^{m\times k}$ and $\mathbf{U} \in \mathbb{R}^{k\times n}$ are the lower and upper triangular matrix, respectively.

1: Generate a randomized Gaussian matrix $\mathbf{\Omega}$ with size $n \times l$;
2: Compute $(\mathbf{A}^T\mathbf{A})^p\mathbf{\Omega}$ and its orthonormal basis $\mathbf{V}$;
3: $\mathbf{Y} = \mathbf{AV}(:, 1 : k)$; // $\mathcal{C}_{mm}mnk$
4: $[\mathbf{L}_1, \mathbf{U}_1, \mathbf{P}] = lu(\mathbf{Y})$; // $\mathcal{C}_{lu}mk^2$
5: $\mathbf{B} = \mathbf{U}_1\mathbf{V}(:, 1 : k)^T$; // $\mathcal{C}_{mm}nk^2$
6: $[\mathbf{L}_2, \mathbf{U}_2, \mathbf{Q}] = lu(\mathbf{B}^T)$; // $\mathcal{C}_{lu}nk^2$
7: $\mathbf{L} = \mathbf{L}_1\mathbf{U}_2^T$; // $\mathcal{C}_{mm}mk^2$
8: $\mathbf{U} = \mathbf{L}_2^T$;

THEOREM 3.1. *Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ with nonnegative singular values $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_{min(m,n)}$, and let $k \geq 2$ be the target rank, $q \geq 2$ be an oversampling parameter, with $k + q \leq min(m, n)$. Draw a Gaussian random matrix $\mathbf{\Omega} \in \mathbb{R}^{n\times(k+q)}$ and set $\mathbf{Y} = (\mathbf{A}^T\mathbf{A})^p\mathbf{\Omega}$ for $p \geq 1$. Let $\mathbf{V} \in \mathbb{R}^{n\times(k+q)}$ be an orthonormal matrix which forms a basis for the range of $\mathbf{Y}$. Then*

(3.7)
$$\mathbb{E}[\|\mathbf{A} - \mathbf{AVV}^T\|] \leq \left[\left(1 + \sqrt{\frac{k}{q-1}}\sigma_{k+1}^{2p}\right)\sigma_{k+1}^{2p} + \frac{e\sqrt{k+q}}{q}\left(\sum_{j>k}\sigma_j^{4p}\right)^{1/2}\right]^{1/(2p)}.$$

Using Theorem 3.1, PowerLU can achieve high accuracy results for a matrix with rapidly decaying singular values. Otherwise, we can improve the accuracy by employing a larger value of exponent, p, in the power iteration.

Remark 3.2. In the PowerLU Algorithm 3.1, we set $p \geq 1$. Since if $p = 0$, Algorithm 3.2 cannot output a orthonormal basis of $\mathbf{A}$.

Algorithm 3.2 Power Iteration Reorthogonalization for Algorithm 3.1

Input: Given a matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, $l < min(m, n)$ and $p \geq 1$.

Output: Orthonormal basis $\mathbf{V} \in \mathbb{R}^{m\times l}$ of the matrix $(\mathbf{AA}^T)^p\mathbf{\Omega}$

1: Generate a random Gaussian matrix $\mathbf{\Omega}$ with size $n \times l$;
2: **for** $i = 1 : 1 : p$ **do**
3: $[\mathbf{\Omega}, \sim] = lu(\mathbf{A\Omega})$; // $\mathcal{C}_{mm}mnl + \mathcal{C}_{lu}ml^2$
4: **if** $i == p$ **then**
5: $[\mathbf{V}, \ ] = qr(\mathbf{A}^T\mathbf{\Omega}))$; // $\mathcal{C}_{mm}mnl + \mathcal{C}_{qr}nl^2$
6: **else**
7: $[\mathbf{\Omega}, \sim] = lu(\mathbf{A}^T\mathbf{\Omega})$; // $\mathcal{C}_{mm}mnl + \mathcal{C}_{lu}nl^2$
8: **end if**
9: **end for**

3.3. Relationship of PowerLU to RandLU. So far we have discussed both RandLU and PowerLU. In practice, reorthogonalization of the power iteration is used

in both algorithms. The time complexity for the PowerLU Algorithm 3.1, where the power iteration is computed by Algorithm 3.2 is

$$\begin{aligned}\mathcal{C}_{PowerLU} &\sim (p-1)\cdot\left[\left(2\mathcal{C}_{mm}mnl + \mathcal{C}_{lu}ml^2 + \mathcal{C}_{lu}nl^2\right)\right] + 2\mathcal{C}_{mm}mnl \\ &\quad + \mathcal{C}_{lu}ml^2 + \mathcal{C}_{qr}nl^2 + \mathcal{C}_{mm}(mnk + nk^2 + mk^2) + \mathcal{C}_{lu}(nk^2 + mk^2).\end{aligned} \tag{3.8}$$

So the difference in time complexity between RandLU and PowerLU is

$$\begin{aligned}\mathcal{C}_{RandLU} - \mathcal{C}_{PowerLU} &= \mathcal{C}_{mm}mnl + \mathcal{C}_{lu}ml^2 + (\mathcal{C}_{lu} - \mathcal{C}_{qr})nl^2 \\ &\quad + \mathcal{C}_{mm}(2mk^2 - nk^2) + \mathcal{C}_{inv}k^3 - \mathcal{C}_{lu}(nk^2 + mk^2).\end{aligned} \tag{3.9}$$

Remark 3.3. For RandLU and PowerLU, we assume the input matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, where $m \geq n$. Since $l \approx k$ equation (3.9) will most likely be positive, which means that our proposed PowerLU algorithm is less costly than RandLU. Our numerical tests confirm this for dense matrices.

Theorem 3.4 below states the relation between the PowerLU and RandLU error bounds in the manner shown in Theorem 2.3.

THEOREM 3.4. *Let $\mathbf{A}$ be an $m \times n$ matrix, let l be a positive integer such that $l < min(m, n)$, let p be a positive integer, and let $\mathbf{\Omega}$ be a random Gaussian matrix of size $n \times l$. Let $\mathbf{P}_p\mathbf{A}\mathbf{Q}_p \approx \mathbf{L}_p\mathbf{U}_p$ be the factorization resulting from the PowerLU Algorithm 3.1. Let $\mathbf{P}_r\mathbf{A}\mathbf{Q}_r \approx \mathbf{L}_r\mathbf{U}_r$ be the factorization resulting from the RandLU Algorithm 2.1. Suppose the rank of the matrix is at least l and both algorithms are performed in exact arithmetic. Then we have:*

$$Range(\mathbf{P}_r^T\mathbf{L}_r) = Range(\mathbf{P}_p^T\mathbf{L}_p). \tag{3.10}$$

Proof. First, we notice from Algorithm 3.1 that we have the following unpivoted QR decomposition

$$(\mathbf{A}^T\mathbf{A})^p\mathbf{\Omega} = \mathbf{V}\mathbf{Z}, \tag{3.11}$$

where $\mathbf{V}$ is the orthogonal basis for $(\mathbf{A}^T\mathbf{A})^p\mathbf{\Omega}$. Then we have

$$\mathbf{Y} = \mathbf{A}(\mathbf{A}^T\mathbf{A})^p\mathbf{\Omega} = \mathbf{A}\mathbf{V}\mathbf{Z} = \mathbf{P}_p^T\mathbf{L}_1\mathbf{U}_1\mathbf{Z}. \tag{3.12}$$

From the RandLU Algorithm 2.1, we have the following:

$$\mathbf{P}_r\mathbf{Y} = \mathbf{L}_y\mathbf{U}_y, \tag{3.13}$$

which means $\mathbf{U}_y$ is invertable, and so

$$\mathbf{L}_y = \mathbf{P}_r\mathbf{Y}\mathbf{U}_y^{-1}. \tag{3.14}$$

Then we replace $\mathbf{Y}$ with the equation (3.12):

$$\mathbf{L}_y = \mathbf{P}_r\mathbf{P}_p^T\mathbf{L}_1\mathbf{U}_1\mathbf{Z}\mathbf{U}_y^{-1}. \tag{3.15}$$

Since the matrices $\mathbf{U}_1, \mathbf{Z}, \mathbf{U}_y$ are non-singular matrices with size $l \times l$, we now have

$$Range(\mathbf{P}_r^T\mathbf{L}_y) = Range(\mathbf{P}_p^T\mathbf{L}_1). \tag{3.16}$$

The RandLU Algorithm 2.1 implies we have $\mathbf{L}_r = \mathbf{L}_y\mathbf{L}_b$, and similarly the PowerLU Algorithm 3.1 implies $\mathbf{L}_p = \mathbf{L}_1\mathbf{U}_2^T$. Since both $\mathbf{L}_b, \mathbf{U}_2$ are non-singular, we arrive at

$$Range(\mathbf{P}_r^T\mathbf{L}_r) = Range(\mathbf{P}_p^T\mathbf{L}_p). \tag{3.17}$$

□

Theorem 3.4 established the relationship between RandLU and PowerLU. We obtain the following equation (3.18) by using Theorem 3.4. So all the theoretical analysis for randomized LU [28] can be applied to the PowerLU decomposition:

$$\mathbf{P}_r^T \mathbf{L}_r \mathbf{L}_r^\dagger \mathbf{P}_r \mathbf{A} = \mathbf{P}_p^T \mathbf{L}_p \mathbf{L}_p^\dagger \mathbf{P}_p \mathbf{A}. \tag{3.18}$$

We observe that the RandLU Algorithm 2.1 accesses the matrix $\mathbf{A}$ a total of $2p+2$ times. In much of the literature, accessing $\mathbf{A}$ is often referred to as a pass of $\mathbf{A}$. However, only $2p+1$ passes of $\mathbf{A}$ are needed in PowerLU. Therefore, for the same value p, PowerLU will have one less pass than RandLU and will therefore be less accurate. In the rest of the section, we will extend the PowerLU Algorithm 3.1 so that works with any number of passes of the matrix $\mathbf{A}$.

3.4. Generalized PowerLU Decomposition. As we discussed in the last section, PowerLU has an odd number passes of the matrix $\mathbf{A}$. For some cases, accessing the matrix will be expensive. However, to increase the accuracy of the computation, we have to increment by 2 passes each time due to multiplication by $\mathbf{A}^T\mathbf{A}$. In [2], the authors proposed the generalized randomized SVD algorithm, which allows any number of passes $v \geq 2$ by using generalized randomized subspace iteration. A similar algorithm was proposed in [11]. Here we use v to denote the number of passes of $\mathbf{A}$. In this section, we will introduce the generalized PowerLU algorithm which can have any number passes $v \geq 2$ of $\mathbf{A}$.

The idea for the generalized PowerLU is very simple. Let $p = \lfloor \frac{v-1}{2} \rfloor$ when v is odd in Algorithm 3.1. We modify steps 1-2 of Algorithm 3.1 when $v \geq 2$ is even as follows:

1. Generate a random Gaussian matrix $\mathbf{\Omega} \in \mathbb{R}^{m\times l}$;
2. Compute $(\mathbf{A}^T\mathbf{A})^p\mathbf{A}^T\mathbf{\Omega}$ and its orthonormal basis $\mathbf{V}$;

We describe the power iteration reorthogonalization procedure for the generalized PowerLU in Algorithm 3.3. Here we need only replace steps 1-2 in Algorithm 3.1 by Algorithm 3.3, which allows an arbitrary number of passes $v \geq 2$ of the matrix $\mathbf{A}$.

When v is odd, then the error bound for generalized PowerLU can be estimated by Theorem 3.1 if we set $p = \lfloor \frac{v-1}{2} \rfloor$. When v is even, the error bound for the generalized PowerLU is given by Theorem 3.5 [2].

THEOREM 3.5. *Let $\mathbf{A} \in \mathbb{R}^{m\times n}$ with nonnegative singular values $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_{min(m,n)}$. Let $k \geq 2$ be the target rank, $q \geq 2$ be an oversampling parameter, with $k+q \leq min(m,n)$. Draw a random Gaussian matrix $\mathbf{\Omega} \in \mathbb{R}^{n\times(k+q)}$ and set $\mathbf{Y} = (\mathbf{A}^T\mathbf{A})^p\mathbf{A}^T\mathbf{\Omega}$ for $p \geq 1$. Let $\mathbf{V} \in \mathbb{R}^{n\times(k+q)}$ be an orthonormal matrix which forms the basis for the range of $\mathbf{Y}$. Then*

$$\begin{aligned} &\mathbb{E}[\|\mathbf{A} - \mathbf{A}\mathbf{V}\mathbf{V}^T\|] = \mathbb{E}[\|\mathbf{A}^T - \mathbf{V}\mathbf{V}^T\mathbf{A}^T\|] \\ &\leq \left[\left(1 + \sqrt{\frac{k}{q-1}\sigma_{k+1}^{2p+1}}\right)\sigma_{k+1}^{2p+1} + \frac{e\sqrt{k+q}}{q}\left(\sum_{j>k}\sigma_j^{2(2p+1)}\right)^{1/2}\right]^{1/(2p+1)}. \end{aligned} \tag{3.19}$$

Its easy to see that if the power iteration is computed by Algorithm 3.3, the cost for the generalized PowerLU is the same as PowerLU when v is odd, When v is even, an extra factorization (LU or QR) and an extra matrix-matrix multiplication will be required. Subsequently PowerLU will refer to the generalized PowerLU decomposition just presented. Numerical experiments to show the efficiency of our proposed algorithm are presented in §6. In the next section, we will discuss how to create a version of PowerLU to solve the fixed precision problem.

Algorithm 3.3 Generalized Power Iteration Reorthogonalization for Algorithm 3.1

Input: Given a matrix $\mathbf{A}\ in\mathbb{R}^{m\times n}$, $l < min(m,n)$ and $v \geq 2$.
Output: Orthonormal basis $\mathbf{V} \in \mathbb{R}^{n\times l}$ for matrix $(\mathbf{A}^T\mathbf{A})^{\lfloor\frac{v-1}{2}\rfloor}\mathbf{A}^T\mathbf{\Omega}$ or $(\mathbf{A}^T\mathbf{A})^{\lfloor\frac{v-1}{2}\rfloor}\mathbf{\Omega}$.

```
1:  if v even then
2:     Generate a random Gaussian matrix Ω with size m × l;
3:     if v > 2 then
4:        [V, ~] = lu(A^T Ω);                    // C_mm mnl + C_lu nl^2
5:     else
6:        [V, ~] = qr(A^T Ω);                    // C_mm mnl + C_qr nl^2
7:     end if
8:  else
9:     Generate a random Gaussian matrix V with size n × l;
10: end if
11: for i = 1 : ⌊(v-1)/2⌋ do
12:    [V, ~] = lu(AV);                          // C_mm mnl + C_lu ml^2
13:    if i == ⌊(v-1)/2⌋ then
14:       [V, ~] = qr(A^T V);                    // C_mm mnl + C_qr nl^2
15:    else
16:       [V, ~] = lu(A^T V);                    // C_mm mnl + C_lu nl^2
17:    end if
18: end for
```

4. PowerLU_FP: A Method for Fixed Precision Low-Rank Matrix Approximation. We now introduce a version of the PowerLU algorithm used to solve the fixed precision low-rank matrix approximation problem. PowerLU_FP is based on an efficient blocked adaptive rank determination algorithm proposed in this section. Our proposed blocked adaptive rank determination algorithm is inspired by the blocked random QB algorithm in [22]. In [22], the authors try to build an orthonormal basis $\mathbf{Q}$ of $\mathbf{A}\in\mathbb{R}^{m\times n}$ incrementally so that for a given tolerance, ϵ,

$$\|\mathbf{A}-\mathbf{Q}\mathbf{Q}^T\mathbf{A}\|_F \leq \epsilon. \tag{4.1}$$

In [38], the authors present an improved version of the blocked random QB algorithm, which uses an inexpensive error calculation as the stopping criterion. They compute a sufficiently large orthonormal basis, $\mathbf{V}$, for $\mathbf{A}^T$ when the power iteration exponent $p \geq 1$. However, their algorithm also tries to build $\mathbf{Q}$ such that (4.1) holds. They fail to use the properties of (2.8) to determine the rank of $\mathbf{A}$.

In rest of this section, we will partition the orthonormal basis of $\mathbf{A}^T$ into blocks, and propose a new blocked adaptive rank determination Algorithm 4.1. We organize PowerLU_FP to solve the fixed precision problem based on Algorithm 4.1.

4.1. Blocked Adaptive Rank Determination for the Fixed Precision Problem. We will describe our blocked adaptive rank determination algorithm. The partition is inspired by the blocked random QB algorithm in [22].

For a given matrix $\mathbf{A}\in\mathbb{R}^{m\times n}$ and target rank k, we compute an orthonormal basis, $\mathbf{V}\in\mathbb{R}^{n\times l}$, for $\mathbf{A}^T$, by Algorithm 3.3 with sufficient large l, where we set $l = k = sb$. Here we introduce b as the block size, and s as the number of blocks. Then, we can partition the matrix $\mathbf{V}$ into blocks $\{\mathbf{V}_j\}_{j=1}^{s}$, each of size $n\times b$,

$$\mathbf{V} = [\mathbf{V}_1, \mathbf{V}_2, \cdots, \mathbf{V}_s]. \tag{4.2}$$

To determine the rank adaptively, we should build a block version of the method to compute an approximate error

$$\|\mathbf{A} - \mathbf{A}\mathbf{V}\mathbf{V}^T\|_F. \tag{4.3}$$

Initially, we set

$$\mathbf{A}^{(0)} = \mathbf{A}, \tag{4.4}$$

and iterate for $i = 1, 2, \cdots, s$ as follows

$$\mathbf{A}^{(i)} = \mathbf{A}^{(i-1)} - \mathbf{A}^{(i-1)}\mathbf{V}_i\mathbf{V}_i^T. \tag{4.5}$$

Equations (4.4)-(4.5) show us how to compute the error (4.3) via blocks of $\mathbf{V}$.

Next, we need to verify

$$\|\mathbf{A}^{(s)}\|_F = \|\mathbf{A} - \mathbf{A}\mathbf{V}\mathbf{V}^T\|_F. \tag{4.6}$$

We state Theorem 4.1 to demonstrate (4.6).

THEOREM 4.1. *Let $\mathbf{A} \in \mathbb{R}^{m\times n}$, b is the block size, and s is the number of blocks. Suppose that the rank of $\mathbf{A}$ is at least sb. Let $\mathbf{V}$ be the orthonormal basis for $\mathbf{A}^T$, and we partition it as in (4.2), with each $\mathbf{V}_i$ of size $n \times b$. Let $\{\mathbf{A}^{(j)}\}_{j=1}^{i}$ be defined by (4.4)-(4.5). We set:*

$$\mathbf{P}_i = \sum_{j=1}^{i} \mathbf{V}_j\mathbf{V}_j^T. \tag{4.7}$$

Then for every $i = 1, 2, \cdots s$, we have

(a) The $\mathbf{P}_i$ are orthogonal projections.

(b) $\mathbf{A}^{(i)} = \mathbf{A}(\mathbf{I} - \mathbf{P}_i)$.

Proof. The proof is as follows.

1. To prove (a), we need to prove that

$$\mathbf{P}_i^2 = \mathbf{P}_i. \tag{4.8}$$

Since $\mathbf{V}$ has orthonormal columns, $\mathbf{V}_j^T\mathbf{V}_k = \delta_{jk}\mathbf{I}$. Therefore,

$$\mathbf{P}_i^2 = (\sum_{j=1}^{i} \mathbf{V}_j\mathbf{V}_j^T)^2 = \mathbf{P}_i. \tag{4.9}$$

So (a) is proved.

2. We prove (b) by induction on i. Suppose (b) holds for $i - 1$, then

$$\mathbf{A}^{(i-1)} = \mathbf{A}(\mathbf{I} - \mathbf{P}_{i-1}). \tag{4.10}$$

We can obtain $\mathbf{A}^{(i)}$ by (4.5)

$$\begin{aligned}
\mathbf{A}^{(i)} &= \mathbf{A}^{(i-1)} - \mathbf{A}^{(i-1)}\mathbf{V}_i\mathbf{V}_i^T \\
&= \mathbf{A}^{(i-1)}(\mathbf{I} - \mathbf{V}_i\mathbf{V}_i^T) \\
&= \mathbf{A}(\mathbf{I} - \mathbf{P}_{i-1})(\mathbf{I} - \mathbf{V}_i\mathbf{V}_i^T) \\
&= \mathbf{A}(\mathbf{I} - (\mathbf{V}_i\mathbf{V}_i^T + \mathbf{P}_{i-1})) = \mathbf{A}(\mathbf{I} - \mathbf{P}_i).
\end{aligned} \tag{4.11}$$

The last stop follows from (4.7), and so (b) is proved. ∎

Using Theorem 4.1, we verify (4.6) as follows:

$$\|\mathbf{A}^{(s)}\|_F = \|\mathbf{A} - \mathbf{A}\mathbf{P}_s\|_F = \|\mathbf{A} - \mathbf{A}\sum_{j=1}^{s}\mathbf{V}_j\mathbf{V}_j^T\|_F = \|\mathbf{A} - \mathbf{A}\mathbf{V}\mathbf{V}^T\|_F. \tag{4.12}$$

Now we can solve the fixed precision low-rank approximation for a given tolerance, ϵ. We now use the remainder of $\mathbf{A}$, $\mathbf{A}^{(i)}$, to exit the loop when the tolerance is satisfied. We stop the computation when $\|\mathbf{A}^{(i)}\|_F^2 < \epsilon^2\|\mathbf{A}\|_F^2$. We can use steps (4.4) - (4.5) to determine the rank of $\mathbf{A}$. However, we need to update $\mathbf{A}^{(i)}$. This is costly, especially when $\mathbf{A}$ is a sparse matrix (fill-in will occur when updating the matrix) [38]. In [38], the authors proposed an error indicator to avoid updating $\mathbf{A}^{(i)}$. In this section, we will introduce a similar error indicator, which can estimate the error quickly without updating.

THEOREM 4.2. *Let $\mathbf{A}$ be an $m \times n$ matrix, and $\mathbf{V}$ is an $n \times b$ orthogonal column matrix (i.e., $\mathbf{V}^T\mathbf{V} = \mathbf{I}_b$). Suppose $b < n$ and $\mathbf{B} = \mathbf{A}\mathbf{V}$. Then*

$$\|\mathbf{A} - \mathbf{B}\mathbf{V}^T\|_F^2 = \|\mathbf{A}\|_F^2 - \|\mathbf{B}\|_F^2. \tag{4.13}$$

Proof. Our proof is similar to that of Theorem 1 in [38]. By the properties of the Frobenius norm, for any matrix $\mathbf{M}$, we have

$$\|\mathbf{M}\|_F^2 = tr(\mathbf{M}\mathbf{M}^T), \tag{4.14}$$

where $tr(\cdot)$ is the trace of the matrix. Then we compute

$$\begin{aligned}(\mathbf{A} - \mathbf{B}\mathbf{V}^T)(\mathbf{A} - \mathbf{B}\mathbf{V}^T)^T &= (\mathbf{A} - \mathbf{B}\mathbf{V}^T)(\mathbf{A}^T - \mathbf{V}\mathbf{B}^T)\\ &= \mathbf{A}\mathbf{A}^T - \mathbf{A}\mathbf{V}\mathbf{B}^T - \mathbf{B}\mathbf{V}^T\mathbf{A}^T + \mathbf{B}\mathbf{V}^T\mathbf{V}\mathbf{B}^T\\ &= \mathbf{A}\mathbf{A}^T - \mathbf{B}\mathbf{B}^T - \mathbf{B}\mathbf{B}^T + \mathbf{B}\mathbf{B}^T\\ &= \mathbf{A}\mathbf{A}^T - \mathbf{B}\mathbf{B}^T.\end{aligned} \tag{4.15}$$

Since the trace is a linear operator, taking the trace of the above proves (4.13). ∎

THEOREM 4.3. *After the ith iteration of the loop (4.4) - (4.5), the error, $E = \|\mathbf{A}^{(i)}\|_F^2$, is given by*

$$E = \|\mathbf{A}\|_F^2 - \sum_{j=1}^{i}\|\mathbf{A}\mathbf{V}_j\|_F^2. \tag{4.16}$$

Proof. The proof is as follows.

$$\begin{aligned}E = \|\mathbf{A}^{(i)}\|_F^2 &= \|\mathbf{A}^{(i-1)} - \mathbf{A}^{(i-1)}\mathbf{V}_i\mathbf{V}_i^T\|_F^2\\ &= \|\mathbf{A}^{(i-1)}\|_F^2 - \|\mathbf{A}^{(i-1)}\mathbf{V}_i\|_F^2\\ &= \|\mathbf{A}^{(i-2)} - \mathbf{A}^{(i-2)}\mathbf{V}_{i-1}\mathbf{V}_{i-1}^T\|_F^2 - \|\mathbf{A}^{(i-1)}\mathbf{V}_i\|_F^2\\ &= \|\mathbf{A}^{(i-2)}\|_F^2 - \|\mathbf{A}^{(i-2)}\mathbf{V}_{i-1}\|_F^2 - \|\mathbf{A}^{(i-1)}\mathbf{V}_i\|_F^2\\ &= \cdots\\ &= \|\mathbf{A}\|_F^2 - \sum_{j=1}^{i}\|\mathbf{A}^{(j-1)}\mathbf{V}_j\|_F^2.\end{aligned} \tag{4.17}$$

In (4.17), we need to prove that $\|\mathbf{A}^{(j-1)}\mathbf{V}_j\|_F^2 = \|\mathbf{A}\mathbf{V}_j\|_F^2$ when $j \leq i$. We have

$$\begin{aligned} \|\mathbf{A}^{(j-1)}\mathbf{V}_j\|_F^2 &= \|(\mathbf{A}^{(j-2)} - \mathbf{A}^{(j-2)}\mathbf{V}_{j-1}\mathbf{V}_{j-1}^T)\mathbf{V}_j\|_F^2 \\ &= \|\mathbf{A}^{(j-2)}\mathbf{V}_j - \mathbf{A}^{(j-2)}\mathbf{V}_{j-1}\mathbf{V}_{j-1}^T\mathbf{V}_j\|_F^2 \end{aligned} \tag{4.18}$$

In equation (4.18), $\mathbf{A}^{(j-2)}\mathbf{V}_{j-1}\mathbf{V}_{j-1}^T\mathbf{V}_j = \mathbf{0}$ since $\mathbf{V}$ has orthonormal columns, and $\mathbf{V}_{j-1}^T\mathbf{V}_j = \mathbf{0}$. Then

$$\|\mathbf{A}^{(j-1)}\mathbf{V}_j\|_F^2 = \|\mathbf{A}^{(j-2)}\mathbf{V}_j\|_F^2 = \cdots = \|\mathbf{A}\mathbf{V}_j\|_F^2. \tag{4.19}$$ □

Theorem 4.3 is an error indicator without the need to update the remainder of $\mathbf{A}$. The efficient blocked adaptive rank determination algorithm without updating is shown in Algorithm 4.1.

Algorithm 4.1 Effective Blocked Adaptive Rank Determination Algorithm

Input: Given $\mathbf{A} \in \mathbb{R}^{m\times n}$, desired accuracy tolerance ϵ, sufficient large $l < min(m,n)$, block size b and $v \geq 2$.

Output: rank k, $\mathbf{V}(:, 1:k)$and $\mathbf{G}(:, 1:k)$.

1: $\mathbf{L} = []$, $\mathbf{U} = []$, $\mathbf{A}^{(0)} = \mathbf{A}$, $E = \|A\|_F^2$ and $acc = \epsilon^2 E$;
2: Passing matrix $\mathbf{A}$, l, and v into Algorithm 3.3 to get the orthonormal basis $\mathbf{V}$;
3: $\mathbf{G} = \mathbf{A}\mathbf{V}$; // $\mathcal{C}_{mm}mnl$
4: **for** $i = 1, 2, 3, \cdots$ **do**
5: Let $t_1 = (i-1)b + 1$ and $t_2 = ib$;
6: $E = E - \|\mathbf{G}(:, t_1 : t_2)\|_F^2$;
7: **if** $E \leq acc$ **then**
8: STOP;
9: **end if**
10: **end for**

Remark 4.4. Algorithm 4.1 is very efficient with the most costly step being 2, which is actually Algorithm 3.3. If $\mathbf{V} \in \mathbb{R}^{n\times l}$ obtained from step 2 has failed to achieve the given tolerance, then we need to regenerate the random Gaussian matrix and rerun the algorithm to obtain additional rank information from the remainder matrix, $\mathbf{A} - \mathbf{A}\mathbf{V}\mathbf{V}^T$. Although extra computation might be needed, this rank determination method is still efficient, since it requires fewer passes of $\mathbf{A}$ or it's remainder compared with the blocked random QB algorithm [22].

4.2. PowerLU_FP: A Method for Fixed Precision Low-Rank Matrix Approximation. So far we discussed a new blocked adaptive rank determination Algorithm 4.1. It produces a randomized LU algorithm suitable for the fixed precision low-rank approximation problem. We call it PowerLU_FP; FP for fixed precision. PowerLU_FP is shown in Algorithm 4.2.

Remark 4.5. The time complexity of PowerLU_FP is related to l. If $l \approx k$, then the time complexity of PowerLU_FP is almost the same as PowerLU. Thus we can make PowerLU_FP more efficient with a suitable l.

5. Single-Pass Randomized LU for the Low-Rank Matrix Approximation. In this section, we will discuss the single-pass algorithm. In [15], the authors present a single-pass randomized QB algorithm, which will compress the input matrix

Algorithm 4.2 PowerLU_FP: PowerLU for the Fixed Precision Problem

Input: Given $\mathbf{A} \in \mathbb{R}^{m\times n}$, desired accuracy tolerance ϵ, sufficiently large $l < min(m,n)$, block size b and $v \geq 2$.

Output: rank k, $\mathbf{P}, \mathbf{Q}, \mathbf{L}, \mathbf{U}$ such that $\mathbf{PAQ} \approx \mathbf{LU}$, where $\mathbf{P}, \mathbf{Q}$ are orthogonal permutation matrices, $\mathbf{L} \in \mathbb{R}^{m\times k}$ and $\mathbf{U} \in \mathbb{R}^{k\times n}$ are the lower and upper triangular matrix, respectively..

1: Use Algorithm 4.1 to take $\mathbf{A}$, ϵ, b, l, and v to obtain k, $\mathbf{V}$, and $\mathbf{G}$;
2: $[\mathbf{L}_1, \mathbf{U}_1, \mathbf{P}] = lu(\mathbf{G})$; // $\mathcal{C}_{lu}mk^2$
3: $\mathbf{B} = \mathbf{U}_1\mathbf{V}^T$; // $\mathcal{C}_{mm}nk^2$
4: $[\mathbf{L}_2, \mathbf{U}_2, \mathbf{Q}] = lu(\mathbf{B}^T)$; // $\mathcal{C}_{lu}nk^2$
5: $\mathbf{L} = \mathbf{L}_1\mathbf{U}_2^T$; // $\mathcal{C}_{mm}mk^2$
6: $\mathbf{U} = \mathbf{L}_2^T$;

$\mathbf{A}$ by using two randomized Gaussian matrices. This algorithm needs to solve a system of linear equations, but is less accurate in practice [15, 37]. A more general single-pass algorithm can be found in [31]. In [37, 38], the authors proposed a single-pass scheme for principal component analysis, which is quite accurate when the matrix is stored in column-major or row-major format. Inspired by this last single-pass algorithms, we propose the single-pass randomized LU for the low-rank matrix approximation. Our proposed single-pass algorithm requires that the matrix is stored in column-major format. For a matrix in row-major format, a slight modification allows the algorithm to work. The proposed algorithm is shown in the Algorithm 5.1.

Given $\mathbf{A} \in \mathbb{R}^{m\times n}$, suppose we want to find it's low-rank approximation with target rank k. First we generate a random Gaussian matrix $\mathbf{\Omega} \in \mathbb{R}^{n\times k}$; then we compute

$$\mathbf{G} = \mathbf{A^T\Omega}, \tag{5.1}$$

and

$$\mathbf{H} = \mathbf{AG}. \tag{5.2}$$

Then we perform partial pivoting LU on $\mathbf{H}$ to define $\mathbf{L}_1$ and $\mathbf{U}_1$ with the permutation matrix, $\mathbf{P}$, encoding the pivots

$$\mathbf{PH} = \mathbf{PAG} = \mathbf{L}_1\mathbf{U}_1. \tag{5.3}$$

Taking the transpose gives us

$$\mathbf{G}^T\mathbf{A}^T\mathbf{P} = \mathbf{U}_1^T\mathbf{L}_1^T. \tag{5.4}$$

Now multiply both sides by $\mathbf{G}^{T\dagger}$

$$\mathbf{A}^T\mathbf{P} = \mathbf{G}^{T\dagger}\mathbf{U}_1^T\mathbf{L}_1^T. \tag{5.5}$$

We now perform partial pivoting LU on $\mathbf{G}^{T\dagger}\mathbf{U}_1^T$ to obtain $\mathbf{L}_2$ and $\mathbf{U}_2$ with the permutation matrix, $\mathbf{Q}$, encoding the pivots

$$\mathbf{QG}^{\dagger T}\mathbf{U}_1^T = \mathbf{L}_2\mathbf{U}_2. \tag{5.6}$$

We have now computed the desired lower and upper triangular matrices as

$$\mathbf{L} = \mathbf{L}_1\mathbf{U}_2^T, \tag{5.7}$$

and

$$\mathbf{U} = \mathbf{L}_2^T. \tag{5.8}$$

Remark 5.1. In practice, if a matrix is stored in column-major format, we compute step (5.1) and (5.2) with one pass over the matrix $\mathbf{A}$ as follows [38]:

Let $\mathbf{A}(:, i)$ be the ith column of $\mathbf{A}$, then the ith row of $\mathbf{G}(i, :)$ is computed by

$$\mathbf{G}(i, :) = \mathbf{A}(:, i)^T \mathbf{\Omega}. \tag{5.9}$$

By using $\mathbf{A}(:, i)$ and $\mathbf{G}(i, :)$, we compute $\mathbf{H}_i$ as an outer product

$$\mathbf{H}_i = \mathbf{A}(:, i)\mathbf{G}(i, :). \tag{5.10}$$

We only to access the ith column of $\mathbf{A}$ one time to compute (5.9) and (5.10). Finally we obtain $\mathbf{G}$ and assemble $\mathbf{H} = \sum_i \mathbf{H}_i$ by accessing $\mathbf{A}$ once per successive column access.

Algorithm 5.1 Single-Pass Randomized LU Factorization

Input: Given $\mathbf{A} \in \mathbb{R}^{m \times n}$, desired rank k.

Output: Matrix: $\mathbf{P}, \mathbf{Q}, \mathbf{L}, \mathbf{U}$ such that $\mathbf{PAQ} \approx \mathbf{LU}$, where $\mathbf{P}, \mathbf{Q}$ are orthogonal permutation matrices, $\mathbf{L} \in \mathbb{R}^{m \times k}$ and $\mathbf{U} \in \mathbb{R}^{k \times n}$ are lower and upper triangular matrices.

1: Generate a randomized Gaussian matrix $\mathbf{\Omega}$ with size $m \times k$;
2: Compute $\mathbf{G} = \mathbf{A}^T \mathbf{\Omega}$;
3: Compute $\mathbf{H} = \mathbf{AG}$;
4: $[\mathbf{L}_1, \mathbf{U}_1, \mathbf{P}] = lu(\mathbf{H})$;
5: $[\mathbf{L}_2, \mathbf{U}_2, \mathbf{Q}] = lu(\mathbf{G}^{\dagger T} \mathbf{U}_1^T)$;
6: $\mathbf{L} = \mathbf{L}_1 \mathbf{U}_2^T$;
7: $\mathbf{U} = \mathbf{L}^T$;

In Algorithm 5.1, we do not use the oversampling parameter for the random Gaussian matrix. If oversampling is used, we need to do the truncation to get the desired results for expected rank k. Experiments show that our proposed algorithm can achieve results almost as accurate as the method in [37], and produce better results than the single-pass algorithm in [15].

6. Numerical Results. In this section, we will present some numerical experiments to show the efficiency and accuracy of our proposed methods. The experiments were carried out on a machine with an Intel(R) Xeon(R) E5-2603 v4 @ 1.70GHz 6 core CPU with 16 GB of RAM. Our code was implemented in MATLAB. Since these algorithms involve randomness, all the results shown in the this section are the average of 20 random calculations. The numerical test settings are similar to those in [38]. The algorithms used in the this section are the following:

- PowerLU: The generalized PowerLU, where the power iteration is computed by the Algorithm 3.3.
- PowerLU_FP: Algorithm 4.2.
- SinglePass: Algorithm 5.1.
- RandSVD: RandSVD Algorithm 1.1 , where the power iteration is computed by the Algorithm 1.2 .

- RandLU: RandLU Algorithm 2.1 , where the power iteration is computed by the Algorithm 2.2.
- RandLU_Original: RandLU Algorithm 2.1, where the power iteration is computed directly without reorthogonalization.
- RangeFinder: Algorithm 4.2 in [15].
- RandQB_b: Blocked random QB Algorithm in [22].
- RandQB_FP: Algorithm 4 in [38]. Note that here FP means Few Passes in the RandQB_FP algorithm as opposed to our meaning of fixed precision.
- SinglePass2011: Single-Pass Algorithm in [15].

In this paper, v denotes the number of passes over the matrix $\mathbf{A}$ for the algorithms: PowerLU and PowerLU_FP, and p denotes the power iteration exponent for the algorithms: RandLU, RandSVD, RandQB_b, and RandQB_FP. However, we have the following relation between v and p:

$$v = 2p + 2. \tag{6.1}$$

Therefore, for simplicity, we only use p. The number of passes over the matrix will be computed by equation (6.1) for the PowerLU and PowerLU_FP algorithms.

6.1. Comparison of Execution Time. In this section, we will examine the execution time of our proposed algorithms. To test our algorithms, we generate random square matrices with variety of sizes ranging from 2000 to 32000.

In the first experiment, we set $l = 200$ for all matrices regardless of their size. All the experiments are done both with and without power iteration. The results are shown in the Fig. 1. For RandQB_b and RandLU results are not available when $n = 32000$ due to machine memory constraints. Fig. 1 on the left is the result without the power iteration. We can see that PowerLU achieves up to 1.80X speedup over RandLU, and has almost the same performance as RandSVD. The results with the power iteration, where $p = 1$, are shown on the right. The gap between RandLU and RandSVD is smaller since more time is spent by both on matrix-matrix multiplication. PowerLU gets a 1.44X speedup over RandLU. We notice that when $p = 0$, RandLU is actually the original randomized LU algorithm, without power iteration. However, our tests show that RandLU is slower than RandSVD, but our PowerLU algorithm provides results comparable with RandSVD.

For algorithms used in the fixed precision problem, we set the block size $b = 20$. PowerLU_FP obtains up to a 4.83X speedup without power iteration, and up to a 3.41X speedup when with power iteration parameter $p = 1$ compared with RandQB_b. From Fig. 1, we see our proposed PowerLU_FP has a 1.16X speedup over RandQB_FP when $n = 2000$ without power iteration. When power iteration is used, PowerLU_FP obtains up to a 1.21X speedup when $n = 2000$ over RandQB_FP. We can see that as the matrix size increases, the advantage over RandQB_FP becomes smaller. RandQB_FP is currently the state-of-art algorithm for the fixed precision low-rank approximation using QR factorization, and it produces an orthonormal basis for $\mathbf{A}$ satisfying (4.1). However, RandQB_FP does not actually produce a factorization, and there are extra steps needed to produce one, such as the SVD. Our PowerLU_FP does produce a low rank LU factorization.

In the second experiment, we test the efficiency of our proposed algorithms on sparse matrices. We generate our test matrices using the MATLAB command *sprand* with density 0.3%. The rest of the parameters are set to the values used above. The results are shown in Fig. 2. Again, results for RandLU are not available when $n = 32000$. As before, the left and right sides of the Fig. 2 are the results without and

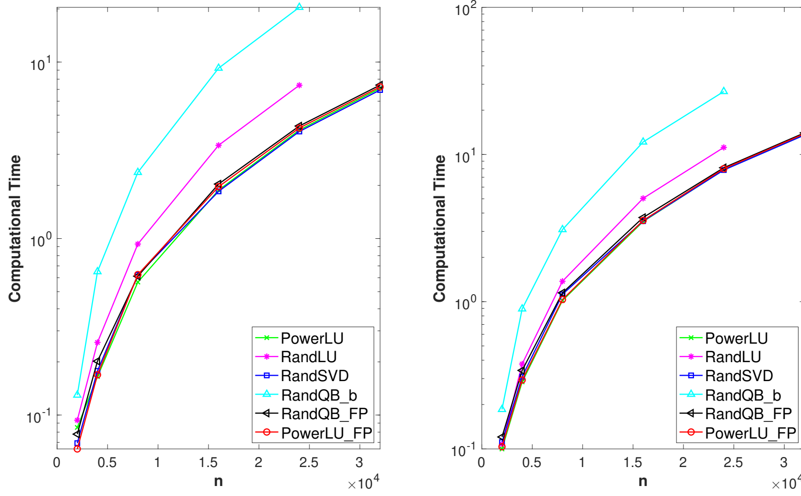

Fig. 1: Runtime of the algorithms with fixed $l = 200$ for dense matrices, without power iteration (left), and with power iteration (right).

with power iteration respectively. When $n < 20000$, RandLU achieves the best results. This is because a sparse matrix, matrix-matrix multiplication is faster than for a dense matrix, and because PowerLU must execute an extra QR decomposition to obtain the orthonormal basis.However, PowerLU performs slightly better than RandSVD. When $n \geq 20000$, we see that RandLU, PowerLU and RandSVD have comparable running times. PowerLU_FP can has up to an 11.36X speedup over RandQB_b, and is generally faster than RandQB_FP for sparse matrices.

In the third experiment, we fixed the size of the matrix and varied of the rank, k, from 100 to 1000. The results are shown in Fig. 3. PowerLU_FP shows up to a 4.21X speedup over randQB_b without power iteration. However, the speedup decreases a little when we set $p = 1$ in the power iteration, since more time will be spent on matrix-matrix multiplication. From Fig. 3, we can see that our PowerLU_FP generally outperforms RandQB_FP for different values of the rank, k, especially when k is large. In terms of the fixed rank algorithms, PowerLU always outperforms RandLU and RandSVD. However, RandLU is actually slower than RandSVD with a relatively small value of k when compared with RandSVD.

6.2. Comparison of Accuracy. To test the accuracy of our methods, we randomly generate three types of the matrices based on the singular value decay speed as follows [13, 38]:

- *Matrix Type 1:(Slow decay)*: This is a matrix $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$, where $\mathbf{U}$ and $\mathbf{V}$ are random matrices with orthonormal columns, and $\mathbf{\Sigma}$ is a rectangular, diagonal matrix with entries $\mathbf{\Sigma}(k,k) = \frac{1}{k^2}$.
- *Matrix Type 2:(Fast decay)*: $\mathbf{\Sigma}(k,k) = e^{-k/7}$.
- *Matrix Type 3:(S-Shaped decay)*: $\mathbf{\Sigma}(k,k) = 0.0001 + (1 + e^{k-30})^{-1}$.

Suppose $\mathbf{A}$ is the original matrix and $\mathbf{A}_k$ is the approximated matrix obtained by using the randomized algorithm. Then we can compute the relative Frobenius error

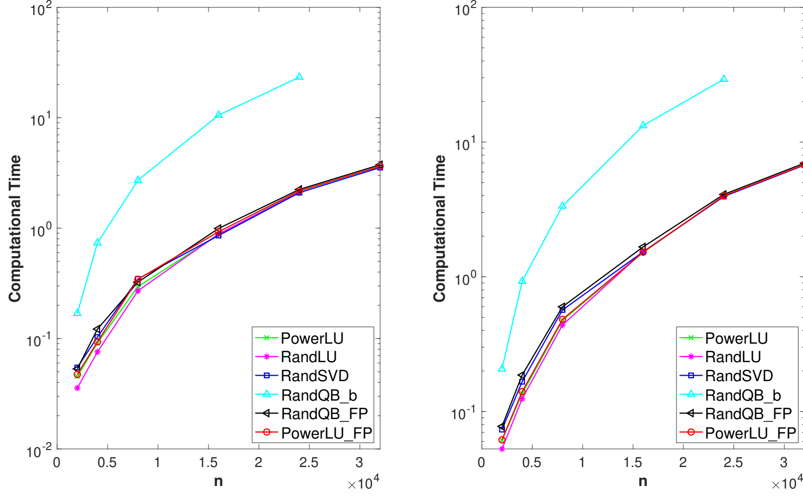

Fig. 2: Runtime of the algorithms with fixed $l = 200$ for sparse matrices, without power iteration (left), and with power iteration (right).

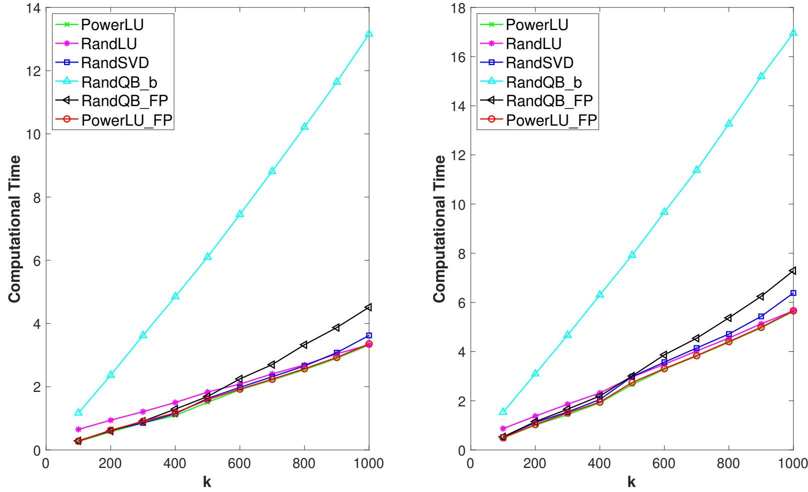

Fig. 3: Runtime of the algorithms with fixed matrix size $n = 8000$, without power iteration (left), and with power iteration (right).

as following:

$$\frac{\|\mathbf{A} - \mathbf{A}_k\|_F}{\|\mathbf{A}\|_F}. \tag{6.2}$$

For each matrix type, we generate a square matrix of size 2000, for which we compare the errors of our proposed algorithms with truncated SVD, RandSVD, and

RandLU for different values of l. We will also set different power iteration parameter values, p. For PowerLU and PowerLU_FP, we need to use equation (6.1) to compute the actually number of matrix passes.As shown in the Figs. 4, 5 and 6, our proposed PowerLU and PowerLU_FP have accuracy comparable to RandSVD with power iteration. For our proposed algorithms, when $p = 1$, they achieves almost the same accuracy as when $p = 2$. RandLU_Original, without reorthogonalization, loses some accuracy as l grows. For our proposed algorithm, any number of passes $v \geq 2$ of $\mathbf{A}$ is possible as opposed an even number of passes. For example, If PowerLU uses 3 passes of $\mathbf{A}$ to achieve the required accuracy, RandSVD and RandLU would need 4 passes since these two algorithms can only accept even number of passes.

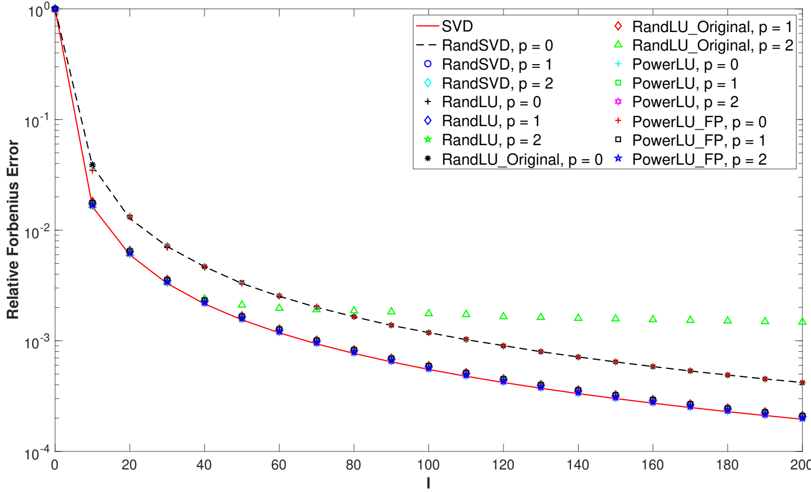

Fig. 4: Errors for Matrix Type 1 (Slow decay)

6.3. Performance of the Single-Pass Algorithm. In this section, we measure the accuracy of our proposed LU-based single-pass low-rank approximation Algorithm 5.1. The results are shown in Fig. 7. We can see that our proposed algorithm has almost the same accuracy as RandQB_FP [38], and both of them are better than SinglePass11[15]. We only show results for Matrix Types 1 and 2.

6.4. Results of Fixed Precision Approximation. In this section, we will present the results of our proposed algorithms used to solve the fixed precision low-rank matrix approximation problem. The Eckart and Young Theorem [9] implies that the optimal solution can be achieved via truncated SVD. For the optimal solution, we could compute the SVD and then calculate the error as $(\sum_{i=k+1}^{min(m,n)} \sigma_i^2)^{1/2}$, where k is the rank that satisfies the given tolerance. Our proposed method, PowerLU_FP, will be compared with truncated SVD, RandQB_FP and RangeFinder. For both PowerLU_FP and RandQB_FP, we set $l = 50b$ sufficiently large, where b is the block size.PowerLU_FP can only generate a matrix of rank that is a multiple of the block size b. To produce the rank more precisely, we can replace steps 7-9 in Algorithm 4.1 by the following steps 7 - 12 .

First, we generate the three types of square matrices as in §6.2, with size 8000.

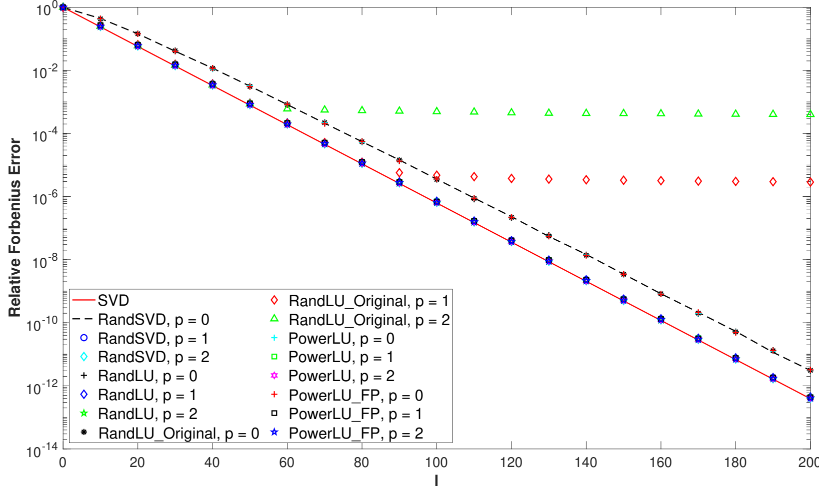

Fig. 5: Errors for Matrix Type 2 (Fast decay)

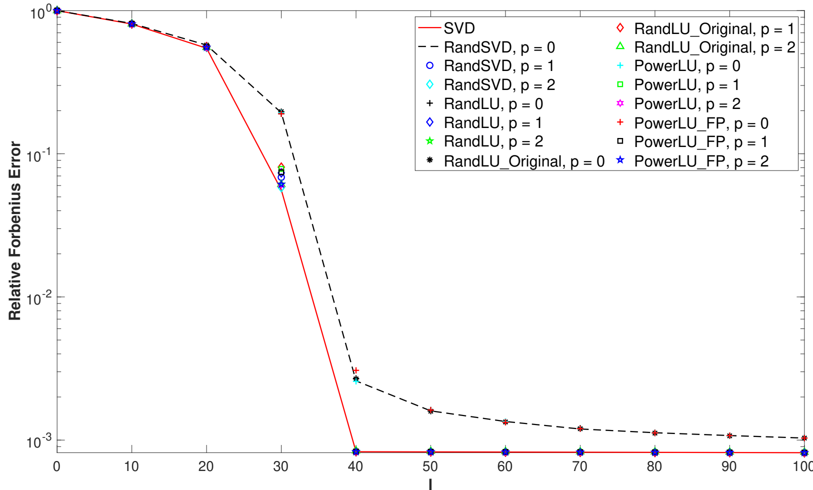

Fig. 6: Errors for Matrix Type3 (S-shaped decay)

For PowerLU_FP and RandQB_FP, we set $p = 1$ (corresponding to $v = 4$ for PowerLU_FP), and we set $b = 10$ for both algorithms except for the last case where we set $b = 40$ which is needed to produce a sufficiently large l. We test our algorithm with differentϵ values. The results are listed in Table 1. We can see that our proposed PowerLU_FP produces a rank close to the truncated SVD but with a speedup greater than 10X.

We next test our algorithm with a real data, which is the standard image [38]

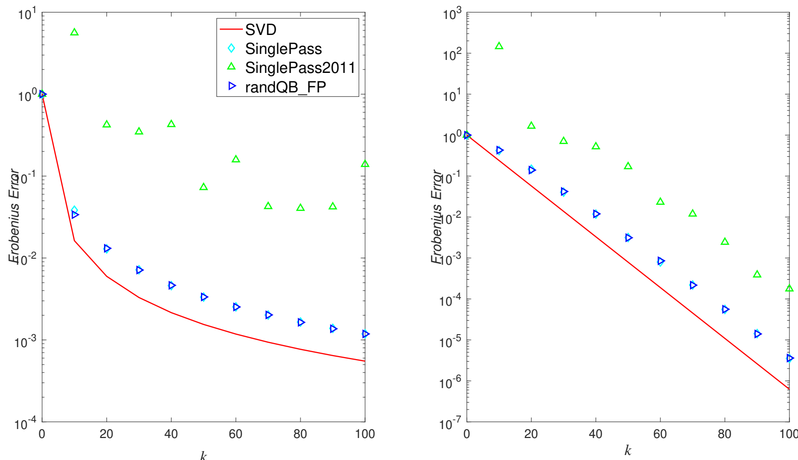

Fig. 7: Approximation errors of the single-pass algorithms and truncated SVD for Matrix Type 1 (left) and Matrix Type 2 (right).

7: **for** $j = 1, 2, \cdots, b$ **do**
8: $\quad E = E - \|\mathbf{G}(:, t1 + j)\|_F^2$;
9: $\quad$ **if** $E < err$ **then**
10: $\quad\quad$ break;
11: $\quad$ **end if**
12: **end for**

---

shown in Fig. 8 represented by a 9504×4752 matrix. We can see from Table 2, when we set $p = 2$ (corresponding $v = 6$ for PowerLU_FP), the rank is reduced significantly. With the same exponent, the results are almost the same with PowerLU_FP despite a different block size. Our proposed algorithm nearly achieves a 7X size reduction from the original, which is comparable to RandQB_FP. The compressed image obtained from PowerLU_FP is shown in Fig. 9.

Table 1: Results for fixed precision problems.

Matrix Type	ϵ	PowerLU_FP		RandQB_FP		Truncated SVD		Range Finder	
		Rank	Time(s)	Rank	Time(s)	Rank	Time(s)	Rank	Time(s)
1	1e-2	15	2.511	16	2.515	15	99.83	116	0.4271
	1e-4	328	2.620	328	2.847	313		2101	23.118
2	1e-4	66	2.515	66	2.462	65	86.06	99	0.3670
	1e-5	82	2.505	82	2.468	81		115	0.4079
3	1e-2	32	2.461	33	2.423	32	86.88	3635	61.477
	1.5e-3	1588	11.24	1588	13.47	1587		7925	272.62

Table 2: Results for fixed precision with real (image) data.

Data	ϵ	Parameters	PowerLU_FP		RandQB_FP		Truncated SVD		Range Finder	
			Rank	Time(s)	Rank	Time(s)	Rank	Time(s)	Rank	Time(s)
Image	0.1	$P=1, b=10$	472	1.967	471	2.415	426	40.62	2892	45.99
		$P=1, b=20$	471	3.753	472	4.008				
		$P=2, b=10$	443	2.749	443	3.310				
		$P=2, b=20$	443	5.388	443	6.122				

Fig. 8: The original image

7. Conclusions and Future Work. In this manuscript, we discussed two novel randomized LU algorithms: PowerLU and PowerLU_FP to solve both the fixed-rank and the fixed precision low-rank approximation problems. A randomized LU (RandLU) Algorithm 2.1 was first proposed in [28]. Our tests show that compared with RandSVD, RandLU is slower for large matrices. Moreover, round-off errors in power iteration will extinguish the small singular values of the input matrix. For the fixed precision problem, RandLU will not work. In this paper, we introduced a reorthogonalization Algorithm 2.2 for the RandLU Algorithm 2.1. With our reorthogonalization procedure, RandLU can eliminate this round-off issue. Compared with RandLU, our new randomized LU Algorithm 3.1, called PowerLU, is generally faster and accurate. PowerLU is based on the orthonormal basis of the input matrix $\mathbf{A}^T$ and the approximation (2.8). PowerLU allows an arbitrary number of passes $v \geq 2$ of the matrix $\mathbf{A}$ via the generalized power iteration reorthogonalization Algorithm 4.2.

Fig. 9: The compressed image obtained by PowerLU_FP with 10% error

However, RandLU only allows even number of passes of $\mathbf{A}$.

To solve the fixed precision problem, we proposed an efficient blocked adaptive rank determination Algorithm 4.1, which uses an efficient error indicator (4.16) without the need for a matrix update. It can determine a rank close to the optimal rank produced by SVD. Then we proposed PowerLU_FP Algorithm 4.2, which is based on Algorithm 4.1. This variant is faster than the randomized QB algorithm in [22]. We also proved the correctness of our proposed algorithms. For the problem where accessing the matrix is expensive, we proposed a single-pass LU-based algorithm, which requires that matrix be stored in column-major format. Tests establish the accuracy of our proposed single-pass Algorithm 5.1.

We plan to build on this work in several ways. First, we plan to implement the algorithms in a high-level computer language such as C/C++ for the sake of efficiency and portability. A longer term goal is to create a distributable piece of mathematical software that can form the basis of a library for randomized linear algebra. We believe that such software will be of interest to may different communities, including big data and data science. We also plan to implement our algorithms on GPUs. We are particularly enthusiastic to explore this with our blocked algorithm. One more potential research direction is the application of the proposed algorithms on image and video processing since these our single-pass version is a stream algorithm. Finally, we plan on considering higher-order tensor decomposition including CANDECOMP/PARAFAC (CP) decomposition [17, 18], Tucker decomposition[18, 32, 33, 34] and Tensor-Train decomposition [24].

REFERENCES

[1] D. Anderson and M. Gu, *An efficient, sparsity-preserving, online algorithm for low-rank approximation*, in Proceedings of the 34th International Conference on Machine Learning-Volume 70, JMLR. org, 2017, pp. 156–165.

[2] E. K. Bjarkason, *Pass-efficient randomized algorithms for low-rank matrix approximation using any number of views*, SIAM Journal on Scientific Computing, 41 (2019), pp. A2355–A2383.
[3] R. Dai, L. Li, and W. Yu, *Fast training and model compression of gated rnns via singular value decomposition*, in 2018 International Joint Conference on Neural Networks (IJCNN), IEEE, 2018, pp. 1–7.
[4] P. Drineas, R. Kannan, and M. W. Mahoney, *Fast monte carlo algorithms for matrices i: Approximating matrix multiplication*, SIAM Journal on Computing, 36 (2006), pp. 132–157.
[5] P. Drineas, R. Kannan, and M. W. Mahoney, *Fast monte carlo algorithms for matrices ii: Computing a low-rank approximation to a matrix*, SIAM Journal on computing, 36 (2006), pp. 158–183.
[6] P. Drineas, R. Kannan, and M. W. Mahoney, *Fast monte carlo algorithms for matrices iii: Computing a compressed approximate matrix decomposition*, SIAM Journal on Computing, 36 (2006), pp. 184–206.
[7] P. Drineas and M. W. Mahoney, *RandNLA: randomized numerical linear algebra*, Communications of the ACM, 59 (2016), pp. 80–90.
[8] P. Drineas and M. W. Mahoney, *Lectures on randomized numerical linear algebra*, arXiv preprint arXiv:1712.08880, (2017).
[9] C. Eckart and G. Young, *The approximation of one matrix by another of lower rank*, Psychometrika, 1 (1936), pp. 211–218.
[10] M. Elad and M. Aharon, *Image denoising via sparse and redundant representations over learned dictionaries*, IEEE Transactions on Image processing, 15 (2006), pp. 3736–3745.
[11] X. Feng, Y. Xie, M. Song, W. Yu, and J. Tang, *Fast randomized PCA for sparse data*, arXiv preprint arXiv:1810.06825, (2018).
[12] G. H. Golub and C. F. Van Loan, *Matrix computations*, vol. 3, JHU Press, 2012.
[13] A. Gopal and P.-G. Martinsson, *The PowerURV algorithm for computing rank-revealing full factorizations*, arXiv preprint arXiv:1812.06007, (2018).
[14] M. Gu and S. C. Eisenstat, *Efficient algorithms for computing a strong rank-revealing QR factorization*, SIAM Journal on Scientific Computing, 17 (1996), pp. 848–869.
[15] N. Halko, P.-G. Martinsson, and J. A. Tropp, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM review, 53 (2011), pp. 217–288.
[16] R. Kannan and S. Vempala, *Randomized algorithms in numerical linear algebra*, Acta Numerica, 26 (2017), pp. 95–135.
[17] H. A. Kiers, *Towards a standardized notation and terminology in multiway analysis*, Journal of Chemometrics: A Journal of the Chemometrics Society, 14 (2000), pp. 105–122.
[18] T. G. Kolda and B. W. Bader, *Tensor decompositions and applications*, SIAM review, 51 (2009), pp. 455–500.
[19] H. Li, G. C. Linderman, A. Szlam, K. P. Stanton, Y. Kluger, and M. Tygert, *Algorithm 971: An implementation of a randomized algorithm for principal component analysis*, ACM Transactions on Mathematical Software (TOMS), 43 (2017), p. 28.
[20] M. W. Mahoney and P. Drineas, *Cur matrix decompositions for improved data analysis*, Proceedings of the National Academy of Sciences, (2009), pp. pnas–0803205106.
[21] M. W. Mahoney et al., *Randomized algorithms for matrices and data*, Foundations and Trends® in Machine Learning, 3 (2011), pp. 123–224.
[22] P.-G. Martinsson and S. Voronin, *A randomized blocked algorithm for efficiently computing rank-revealing factorizations of matrices*, SIAM Journal on Scientific Computing, 38 (2016), pp. S485–S507.
[23] L. Miranian and M. Gu, *Strong rank revealing LU factorizations*, Linear algebra and its applications, 367 (2003), pp. 1–16.
[24] I. V. Oseledets, *Tensor-train decomposition*, SIAM Journal on Scientific Computing, 33 (2011), pp. 2295–2317.
[25] C.-T. Pan, *On the existence and computation of rank-revealing LU factorizations*, Linear Algebra and its Applications, 316 (2000), pp. 199–222.
[26] A. Rotbart, G. Shabat, Y. Shmueli, and A. Averbuch, *Randomized LU decomposition: An algorithm for dictionaries construction*, arXiv preprint arXiv:1502.04824, (2015).
[27] G. Shabat, *Randomized LU matrix factorization Code*. Link. Accessed: 2019-10-12.
[28] G. Shabat, Y. Shmueli, Y. Aizenbud, and A. Averbuch, *Randomized LU decomposition*, Applied and Computational Harmonic Analysis, 44 (2018), pp. 246–272.
[29] G. Stewart, *The QLP approximation to the singular value decomposition*, SIAM Journal on Scientific Computing, 20 (1999), pp. 1336–1348.

[30] G. W. Stewart, *Updating a rank-revealing ULV decomposition*, SIAM Journal on Matrix Analysis and Applications, 14 (1993), pp. 494–499.
[31] J. A. Tropp, A. Yurtsever, M. Udell, and V. Cevher, *Practical sketching algorithms for low-rank matrix approximation*, SIAM Journal on Matrix Analysis and Applications, 38 (2017), pp. 1454–1485.
[32] L. R. Tucker, *Implications of factor analysis of three-way matrices for measurement of change*, Problems in measuring change, 15 (1963), pp. 122–137.
[33] L. R. Tucker, *Some mathematical notes on three-mode factor analysis*, Psychometrika, 31 (1966), pp. 279–311.
[34] L. R. Tucker et al., *The extension of factor analysis to three-dimensional matrices*, Contributions to mathematical psychology, 110119 (1964).
[35] M. Udell and A. Townsend, *Why are big data matrices approximately low rank?*, SIAM Journal on Mathematics of Data Science, 1 (2019), pp. 144–160.
[36] S. Voronin and P.-G. Martinsson, *RSVDPACK: An implementation of randomized algorithms for computing the singular value, interpolative, and cur decompositions of matrices on multi-core and gpu architectures*, arXiv preprint arXiv:1502.05366, (2015).
[37] W. Yu, Y. Gu, J. Li, S. Liu, and Y. Li, *Single-pass PCA of large high-dimensional data*, arXiv preprint arXiv:1704.07669, (2017).
[38] W. Yu, Y. Gu, and Y. Li, *Efficient randomized algorithms for the fixed-precision low-rank matrix approximation*, SIAM Journal on Matrix Analysis and Applications, 39 (2018), pp. 1339–1359.