

Randomized LU Decomposition

Gil Shabat^{1*} Yaniv Shmueli² Amir Averbuch²

¹School of Electrical Engineering, Tel Aviv University, Israel

²School of Computer Science, Tel Aviv University, Israel

October 31, 2013

Abstract

We present a fast randomized algorithm that computes a low rank LU decomposition. Our algorithm uses random projections type techniques to efficiently compute a low rank approximation of large matrices. The randomized LU algorithm can be parallelized and further accelerated by using sparse random matrices in its projection step. Several different error bounds are proven for the algorithm approximations. To prove these bounds, recent results from random matrix theory related to subgaussian matrices are used. As an application, we also show how the algorithm can be utilized to solve problems such as the rank-deficient least squares problem. Numerical examples, which illustrate the performance of the algorithm and compare it to other decomposition methods, are presented.

Keywords. LU decomposition, random matrices, randomized algorithms, sparse matrices.

1 Introduction

Matrix factorizations and low rank approximations play a major role in today's applications. In mathematics, matrix decompositions are used for low rank approximations that often reveal interesting properties of a matrix. Matrix decompositions are used for solving linear equations and for finding least squares solutions. In engineering, matrix decompositions are used in computer vision [13], machine learning [24], collaborative filtering and Big Data analytics [19]. As the size of the data grows exponentially, analysis of large datasets has gained an increasing importance. Such analysis can involve a factorization step of the input data given as a large sample-by-feature matrix or by a sample affinity matrix. Two main reasons for the difficulty in analyzing huge data structures are the high memory consumption and the computational complexity of the factorization step (typically $\mathcal{O}(n^3)$ operations, where n is the size of the matrix).

*Corresponding author: gil@eng.tau.ac.il

Recently, there is an on-going interest in applying mathematical tools from the field of random matrices to deal with these difficulties.

Several randomized algorithms for processing a large scale matrix factorizations have been developed, allowing to overcome the memory and performance challenges. Most of the randomized algorithms use random projections, which projects the matrix to a set of random vectors. Formally, given a matrix A of size $m \times n$ (suppose $m \geq n$) and a random matrix G of size $n \times k$, the product AG is computed to obtain a smaller matrix that potentially captures most of the data activities in A . In most applications, k is set to be much smaller than n in order to obtain a compact approximation for A that is more suitable for computational needs.

This can lead to an approximation error, for example, when the rank of G is smaller than the rank of A . Among the uses of fast randomized matrix decomposition algorithms, we find applications for tracking objects in videos [30], multiscale extensions for data [2] and detecting anomalies in network traffic for finding cyber attacks [8]. There are randomized versions for many different matrix factorizations [17] such as singular value decomposition (SVD), interpolative decomposition (ID) [7], pseudo-skeleton decomposition [16] (in which its randomized version is given by the CUR decomposition [12]), compressed sensing [11] and a randomized version for solving least squares problems [27].

In this paper, we present a randomized LU decomposition algorithm. Given an $m \times n$ matrix A , we seek a lower triangular $m \times k$ matrix L and an upper triangular $k \times n$ matrix U such that

$$\|LU - PAQ\|_2 = \mathcal{O}(\sigma_{k+1}), \quad (1.1)$$

where P and Q are orthogonal permutation matrices and σ_{k+1} is the $k+1$ largest singular value of A . Notice that L and U are not perfect square matrices as they contain a $k \times k$ triangular sub-matrix and some additional rows or columns. We prove several error bounds to approximate the $\|LU - PAQ\|_2$ error.

LU decomposition for sparse matrices can be done using fast methods that utilize sparse matrix multiplication (see [33] for example). Moreover, LU decomposition with full pivoting on sparse matrices is effective since it can generate large regions of zeros [9, 10, 29]. In this paper, we show how sparsity can be utilized to accelerate the computation of the randomized LU. We compare its performance to a sparse SVD algorithm performance that is based on Lanczos bidiagonalization [15] implemented in the PROPACK package [20].

Graphics Processing Units (GPU) are mostly used for computer games, graphics and visualization such as movies and 3D display. GPUs have powerful computation capabilities since they can do fast vectors and matrices computations and fast basic integer and floating arithmetics. For some computations, such as matrix multiplications and Fast Fourier Transform, GPUs can achieve up to 1,000 speedup over a general purpose CPU [18]. New video cards that are based on GPUs can achieve up to several Tera-Flops per second compared to a standard CPUs that are limited to approximately 10 Giga-Flops per second. Our randomized LU algorithm uses matrix multiplication and deterministic LU

decompositions that can be parallelized effectively. We show that by using a GPU we significantly accelerate the LU decomposition.

The paper is organized as follows: in Section 2, we overview related work on matrix decomposition and approximation using randomized methods. Section 3 reviews some mathematical facts that are needed for the development of the randomized LU. Section 4 presents the randomized LU algorithm and proves several error bounds on the approximation. We discuss the case of sparse matrices and also show how to solve rank deficient least squares problems using the randomized algorithm as an example. Section 5 presents numerical results on the approximation error, the computational complexity of the algorithm and compares it with other methods. We compare the performance of the algorithms using random matrices, images and large sparse matrices as their input data.

2 Related Work

Efficient matrix decomposition serves as a basis for many studies and algorithms that deal with data analysis applications. There is a variety of methods and algorithms that factorize a matrix into several matrices. Typically, the factorized terms have properties such as being triangular, orthogonal, diagonal, sparse or low-rank. When a matrix is factorized, it is possible to have a certain control on the desired approximation error.

SVD approximates an $m \times n$ matrix A by a low rank k matrix. It maintains the first k singular values while zeroing the others. Assume that the SVD of A is given by $U\Sigma V^T$ where Σ is a diagonal matrix whose elements are $\sigma_1, \dots, \sigma_{\min(m,n)}$. Then, the matrix $\tilde{A} = U\tilde{\Sigma}V^T$ with $\tilde{\sigma}_i = \sigma_i$ for $i = 1, \dots, k$ and $\tilde{\sigma}_i = 0$ for $i > k$ is its best rank k approximation under the spectral and Frobenius norms $\|A - \tilde{A}\|_2 = \sigma_{k+1}$ and $\|A - \tilde{A}\|_F^2 = \sum_{i=k+1}^{\min(m,n)} \sigma_i^2$, respectively.

Other rank revealing factorizations can be used to achieve low rank approximations. For example, both QR and LU factorizations have rank revealing versions such as RRQR [5] and RRLU [26], respectively. Rank revealing factorization uses permutation matrices on the columns and rows of A so that the factorized matrices structure have a strong rank portion and a rank deficient portion.

Other matrix factorization methods such as Interpolative Decomposition (ID) [7] and CUR decomposition [12], use columns and rows of the original matrix A in the factorization process. Such a property exposes the most important elements that construct A .

An ID factorization of order k of an $m \times n$ matrix A consists of an $m \times k$ matrix B whose columns consist of a subset of the columns of A , as well as a $k \times n$ matrix P , such that a subset of the columns of P becomes a $k \times k$ identity matrix and $A \approx BP$ such that $\|A - BP\| \lesssim \mathcal{O}(n, \sigma_{k+1}(A))$. Usually, k is chosen to be the numerical rank $k = \#\{j : \sigma_j(A) \geq \delta\sigma_1(A)\}$ of A up to a certain accuracy $\delta > 0$. This selection of k guarantees that the columns of B constitute a well conditioned basis to the range of A [7].

A CUR matrix k -approximation of a matrix A is constructed of matrices C

U and R where C is a subset of k columns from A , U is a small $k \times k$ matrix and R is a subset of k rows from A . If A is a matrix of samples where each column represents a feature, then CUR approximation is easier to interpret than SVD since its factorized terms are the most important samples and features from A . CUR approximation is less accurate than SVD but it also has a provable worst-case error bound.

All the above decompositions are impractical for processing large matrices since they require $\mathcal{O}(mn^2)$ operations (for $n < m$) and it is impossible to fit them into memory. This can be overcome by applying acceleration methods and memory optimization techniques and by trading accuracy for speed. Randomized algorithms play an important role since they offer an effective way to factorize matrices up to any precision. Randomized versions of SVD, QR and ID factorizations have been developed and studied [23]. Computing a low-rank matrix approximation using randomization methods can be done in two steps. First, a low-dimensional space, which captures most of the “energy” of A , is found. Then, the matrix A is projected into the retrieved subspace and the factorization of the projected matrix is computed [17]. Most random methods for matrix decomposition use a random Gaussian matrix as part of their randomization step. Such methods rely on known provable bounds for the singular values of the random matrix. Special random matrices such as sparse random matrices and subgaussian matrices were not used in such decomposition algorithms. In Section 4, we show how the usage of such random matrices can assist in providing efficient algorithms with improved error bounds in some cases.

3 Preliminaries

In this section, we review the RRLU decomposition and several singular values bounds on random matrices that will be used to prove the error bounds for the randomized LU algorithm. Throughout the paper, we use the following notation: for any matrix A , $\sigma_j(A)$ is the j th largest singular value and $\|A\|$ is the spectral norm (the largest singular value or l_2 operator norm). If x is a vector then $\|x\|$ is the standard l_2 (Euclidean) norm. A^\dagger denotes the pseudo-inverse of A .

3.1 Rank Revealing LU (RRLU)

The following theorem is adapted from [26]:

Theorem 3.1 ([26]). *Let A be an $m \times n$ matrix ($m \geq n$). Given an integer $1 \leq k < n$, then the following factorization holds*

$$PAQ = \begin{pmatrix} L_{11} & 0 \\ L_{21} & I_{n-k} \end{pmatrix} \begin{pmatrix} U_{11} & U_{12} \\ 0 & U_{22} \end{pmatrix}, \quad (3.1)$$

where L_{11} is a unit lower triangular and U_{11} is an upper triangular. Let $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$ be the singular values of A , then:

$$\sigma_k \geq \sigma_{\min}(L_{11}U_{11}) \geq \frac{\sigma_k}{k(n-k)+1}, \quad (3.2)$$

and

$$\sigma_{k+1} \leq \|U_{22}\| \leq (k(n-k) + 1)\sigma_{k+1}. \quad (3.3)$$

Based on Theorem 3.1, we have the following definition:

Definition 3.1 (RRLU Rank k Approximation denoted $RRLU_k$). *Given a RRLU decomposition (Theorem 3.1) with integer k of a matrix A as in Eq. 3.1 such that $PAQ = LU$. The RRLU rank k approximation is defined by taking the k columns from L and the k rows of U such that*

$$RRLU_k(PAQ) = \begin{pmatrix} L_{11} \\ L_{21} \end{pmatrix} (U_{11}U_{12}). \quad (3.4)$$

Lemma 3.2 (RRLU Approximation Error). *The error of the $RRLU_k$ approximation of A is*

$$\|PAQ - RRLU_k(PAQ)\| \leq (k(n-k) + 1)\sigma_{k+1}. \quad (3.5)$$

Proof. From Eqs. 3.1 and 3.4 we have

$$\begin{aligned} \|PAQ - RRLU_k(PAQ)\| &= \left\| \begin{pmatrix} L_{11} & 0 \\ L_{21} & I_{n-k} \end{pmatrix} \begin{pmatrix} U_{11} & U_{12} \\ 0 & U_{22} \end{pmatrix} - \begin{pmatrix} L_{11} \\ L_{21} \end{pmatrix} (U_{11}U_{12}) \right\| \\ &= \|U_{22}\| \leq (k(n-k) + 1)\sigma_{k+1}. \end{aligned} \quad (3.6)$$

The last inequality derives from Eq. 3.3. \square

Lemma 3.3 appears in [3], page 75:

Lemma 3.3 ([3]). *Let A and B be two matrices and let $\sigma_j(\cdot)$ denotes the j th singular value of a matrix. Then, $\sigma_j(AB) \leq \|A\|\sigma_j(B)$ and $\sigma_j(AB) \leq \|B\|\sigma_j(A)$.*

Lemma 3.4 was taken from [23] and it is an equivalent formulation for Eq. 8.8 in [14].

Lemma 3.4 ([23]). *Suppose that G is a real $n \times l$ matrix whose entries are i.i.d Gaussian random variables with zero mean and unit variance and let m be an integer such that $m \geq l$, $m \geq n$, $\gamma > 1$ and*

$$1 - \frac{1}{4(\gamma^2 - 1)\sqrt{\pi m \gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2 - 1}} \right)^m \quad (3.7)$$

is nonnegative. Then, $\|G\| \leq \sqrt{2m}\gamma$ with probability not less than the value in Eq. 3.7.

3.2 Sparse Random Matrices

Sparse matrices have a significant importance in many applications. As a popular example, computing the degrees of separation between two individuals using Facebook 2011 connection matrix requires to factorize a sparse matrix of size $720,000,000 \times 720,000,000$ with 69 billions connections. It means that only 1.33×10^{-5} percent of the matrix is non zero [1]. The advantage of using sparse matrices is its low memory consumption and the fact that some algorithms are more efficient when applied to sparse matrices.

Definition 3.2 (sparse Gaussian matrix). $A = (\xi_{ij})$ is a sparse Gaussian matrix where each entry is normally distributed with probability ρ and zero with probability $1 - \rho$. That is, ξ_{ij} is a random variable whose probability density function (PDF) is given by

$$p(x) = (1 - \rho)\delta(x) + \frac{\rho}{\sqrt{2\pi}\sigma^2} e^{-\frac{x^2}{2\sigma^2}}, \quad (3.8)$$

where $\delta(x)$ is the Dirac delta function and $0 < \rho \leq 1$. If A is a sparse matrix whose non-zero entries are Gaussian variables, then we refer to ρ as the density of the matrix.

Subgaussian random matrices processing is based on Definition 3.3:

Definition 3.3. A real valued random variable X is called b -subgaussian if for all $t > 0$ we have $\mathbb{E}e^{tX} \leq e^{b^2 t^2 / 2}$ where \mathbb{E} denotes expectation.

Suppose X is distributed as in Eq. 3.8 and \mathbb{E} is the expectation. One can easily verify that:

1. $\mathbb{E}X = 0$;
2. $\mathbb{E}X^2 = \rho\sigma^2$;
3. $\mathbb{E}|X|^3 = \frac{4\rho\sigma^3}{\sqrt{2\pi}}$;
4. X is subgaussian.

We review several facts adapted from [22] and [28] about random matrices whose entries are subgaussian. We focus on the case where A is a tall $m \times n$ matrix ($m > (1 + \frac{1}{\ln n})n$). Tall matrices can be used for low rank LU approximations, where the approximation rank is much smaller than the size of the matrix. Similar results can be found in [21] for almost square and square matrices.

Definition 3.4. For parameters $\mu \geq 1$, $a_1 > 0$, $a_2 > 0$, we define $\mathcal{A}(\mu, a_1, a_2, m, n)$ to be a set of all $m \times n$ ($m > n$) random matrices $A = (\xi_{ij})$ whose entries are i.i.d real valued centered random variables satisfying the following conditions:

1. Moments: $\mathbb{E}|\xi_{ij}|^3 \leq \mu^3$.
2. Norm: $\mathbb{P}(\|A\| > a_1\sqrt{m}) \leq e^{-a_2 m}$.

3. Variance: $\mathbb{E}\xi_{ij}^2 \geq 1$

where \mathbb{P} is the probability function.

It is shown in [22] that if A is subgaussian then $A \in \mathcal{A}$. In particular, a Gaussian matrix whose entries are zero with probability $1-\rho$ is also subgaussian. Hence, this model can also be used for sparse Gaussian matrices with density ρ . For simplicity, we work with random matrices with unit variance. In the case of sparse Gaussian matrices, we set:

$$\sigma^2 = \frac{1}{\rho}, \quad (3.9)$$

which leads to

$$\mu = \left(\frac{4}{\sqrt{2\pi\rho}} \right)^{\frac{1}{3}}. \quad (3.10)$$

The following theorems are taken from [22]:

Theorem 3.5 ([22]). *Every matrix A of size $m \times n$ ($m \geq n$) whose entries are subgaussian with $\mu \geq 1$ and $a_2 \geq 0$ satisfies:*

$$\mathbb{P}(\|A\| \geq a_1\sqrt{m}) \leq e^{-a_2m} \quad (3.11)$$

with $a_1 = 6\mu\sqrt{a_2 + 4}$.

Theorem 3.5 provides an upper bound for the largest singular value depending on the desired probability. Theorem 3.6 is used to bound from below the smallest singular value of sparse Gaussian matrices.

Theorem 3.6 ([22]). *Let $\mu \geq 1$, $a_1, a_2 > 0$. Let A be an $m \times n$ matrix with $m > (1 + \frac{1}{\ln n})n$. m can be written as $m = (1 + \delta)n$. Suppose the entries of A are independent centered random variables such that conditions 1, 2, 3 in Definition 3.4 hold. Then, there exist positive constants c_1 and c_2 such that:*

$$\mathbb{P}(\sigma_n(A) \leq c_1\sqrt{m}) \leq e^{-m} + e^{-c'm/(2\mu^6)} + e^{-a_2m} \leq e^{-c_2m}. \quad (3.12)$$

The exact values of constants c_1, c_2 and c'' are given by:

$$c_1 = \frac{b}{e^2 c_3} \left(\frac{b}{3e^2 c_3 a_1} \right)^{\frac{1}{3}}, \quad (3.13)$$

$$c'' = \frac{27}{2^{11}}. \quad (3.14)$$

Here, $c_3 = 4\sqrt{\frac{2}{\pi}} \left(\frac{2\mu^9}{a_1^3} + \sqrt{\pi} \right)$, $b = \min \left(\frac{1}{4}, \frac{c'}{5a_1\mu^3} \right)$ and $c' = \left(\frac{27}{2^{13}} \right)^{\frac{1}{2}}$. For the constant c_2 , we need a constant small enough to satisfy the inequality in Eq. 3.12 and set is as (for simplification):

$$c_2 = \min \left(1, \frac{c'}{(2\mu^6)}, a_2 \right) - \frac{\ln 3}{m} \quad (3.15)$$

4 Randomized LU

In this section, we present a randomized LU algorithm for computing the rank k LU approximation of a full matrix (Algorithm 4.1). In addition, we present a version that approximates a sparse matrix (Algorithm 4.2). Error bounds are proven for each algorithm.

Algorithm 4.1: Randomized LU Decomposition

Input: A matrix of size $m \times n$ to decompose; k desired rank; l number of columns to use.

Output: Matrices P, Q, L, U such that $\|PAQ = LU\| \leq \mathcal{O}(\sigma_{k+1}(A))$ where P and Q are orthogonal permutation matrices, L and U are the lower and upper triangular matrices, respectively.

- 1: Create a matrix G of size $n \times l$ whose entries are i.i.d. Gaussian random variables with zero mean and unit standard deviation.
 - 2: $Y \leftarrow AG$.
 - 3: Perform RRLU decomposition (Theorem 3.1) of Y such that $PYQ_y = L_yU_y$.
 - 4: Truncate L_y and U_y by choosing the first k columns and k rows, respectively: $L_y \leftarrow L_y(:, 1:k)$ and $U_y \leftarrow U_y(1:k, :)$
 - 5: $B \leftarrow L_y^\dagger PA$.
 - 6: Perform LU decomposition of B with column pivoting $BQ = L_bU_b$.
 - 7: $L \leftarrow L_yL_b$.
 - 8: $U \leftarrow U_b$.
-

Remark 4.1. The pseudo-inverse of L_y in step 5 can be computed by $L_y^\dagger = (L_y^T L_y)^{-1} L_y^T$. This can be done efficiently when it is computed on platforms such as GPUs that can multiply matrices efficiently. Usually, the inversion is done on a small matrix since in many cases $k \ll n$ and it can be done simply by the application of Gaussian elimination.

Remark 4.2. In practice, it is sufficient to perform step 3 in Algorithm 4.1 using standard LU decomposition with partial pivoting instead of applying RRLU. The cases where U grows exponentially are extremely rare (section 3.4.5 in [15] and [32]).

We now present our main error bound for Algorithm 4.1:

Theorem 4.3. Given a matrix A of size $m \times n$. Then, its randomized LU decomposition produced by Algorithm 4.1 with integers k and l ($l \geq k$) satisfies:

$$\|LU - PAQ\| \leq \left(2\sqrt{2nl\beta^2\gamma^2 + 1} + 2\sqrt{2nl}\beta\gamma(k(n-k) + 1) \right) \sigma_{k+1}(A), \quad (4.1)$$

with probability not less than

$$1 - \frac{1}{\sqrt{2\pi(l-k+1)}} \left(\frac{e}{(l-k+1)\beta} \right)^{l-k+1} - \frac{1}{4(\gamma^2-1)\sqrt{\pi n\gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^n, \quad (4.2)$$

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

The proof for Theorem 4.3 is given in Section 4.2. To show that the success probability ξ in Eq. 4.2 is sufficiently high, we present in Table 4.1 several calculated values of ξ . We omitted the value of n from Table 4.1, since it does not affect the value of ξ since the second term in Eq. 4.2 decays fast.

Table 4.1: Calculated values of the success probability ξ (Eq. 4.2).

$l - k$	β	γ	ξ
3	5	5	$1 - 6.8 \times 10^{-5}$
5	5	5	$1 - 9.0 \times 10^{-8}$
10	5	5	$1 - 5.2 \times 10^{-16}$
3	30	5	$1 - 5.2 \times 10^{-8}$
5	30	5	$1 - 1.9 \times 10^{-12}$
10	30	5	$1 - 1.4 \times 10^{-24}$
3	30	10	$1 - 5.2 \times 10^{-8}$
5	30	10	$1 - 1.9 \times 10^{-12}$
10	30	10	$1 - 1.4 \times 10^{-24}$

In Section 5, we show that in practice, Algorithm 4.1 produces comparable results to other well known randomized factorization methods such as randomized SVD and randomized ID.

4.1 Computational Complexity Analysis

To compute the number of floating points operations of Algorithm 4.1, we evaluate the complexity of each step:

1. Generating an $n \times l$ random matrix requires $\mathcal{O}(nl)$ operations.
2. Multiplying A by G to form Y requires $l\mathcal{C}_A$ operations, where \mathcal{C}_A is the complexity of applying A to a column vector $n \times 1$.
3. Computing partial pivoting LU for Y requires $\mathcal{O}(ml^2)$ operations.
4. Selecting the first k columns (we do not modify them) requires $\mathcal{O}(1)$ operations.
5. Computing the pseudo inverse of L_y requires $\mathcal{O}(k^2m + k^3 + k^2m)$ operations and multiplying it by A requires $k\mathcal{C}_{A^T}$ operations. Note that P is a permutation matrix that does not modify the rows of A .
6. Computing the partial pivoting LU for B requires $\mathcal{O}(k^2n)$ operations.
7. Computing L requires $\mathcal{O}(k^2m)$ operations.
8. Computing U requires $\mathcal{O}(1)$ operations.

By summing up the complexities of all the steps above, then Algorithm 4.1 necessitated

$$C_{RandLU} = lC_A + kC_{A^T} + \mathcal{O}(l^2m + k^3 + k^2n) \quad (4.3)$$

operations. Here, we used C_A (or C_{A^T}) as the complexity of applying A (or A^T) to a vector, since in some cases this can be done in an efficient way when for example A is an FFT matrix [25].

4.2 Bounds for the Randomized LU

In this section, we prove Theorem 4.3 and an additional complementary bound. This is done by finding a basis to a smaller matrix AG which is achieved in practice by using RRLU. The assumptions are that L is numerically stable so its pseudo-inverse can be computed accurately, that there exists a matrix U such that LU is a good approximation to AG and that there exists a matrix F such that $\|AGF - A\|$ is small. As for the numerical stability of L , it is always stable since it has a small condition number [31].

For the proof of Theorem 4.3, several theorems are needed and introduced. Lemma 4.4 states that a given basis L can form a good approximation to a matrix A by bounding $\|LL^\dagger A - A\|$.

Lemma 4.4. *Assume that A is an $m \times n$ matrix, L is an $m \times k$ matrix, G is an $n \times l$ matrix, U is a $k \times l$ matrix and F is $l \times n$ ($k \leq m$) matrix. Then,*

$$\|LL^\dagger A - A\| \leq 2\|AGF - A\| + 2\|F\|\|LU - AG\|. \quad (4.4)$$

Proof. By using the triangular inequality we get

$$\|LL^\dagger A - A\| \leq \|LL^\dagger A - LL^\dagger AGF\| + \|LL^\dagger AGF - AGF\| + \|AGF - A\|. \quad (4.5)$$

Clearly, the first term can be bounded by

$$\|LL^\dagger A - LL^\dagger AGF\| \leq \|LL^\dagger\|\|A - AGF\| \leq \|A - AGF\|. \quad (4.6)$$

The second term can be bounded by

$$\|LL^\dagger AGF - AGF\| \leq \|F\|\|LL^\dagger AG - AG\|. \quad (4.7)$$

Also,

$$\|LL^\dagger AG - AG\| \leq \|LL^\dagger AG - LL^\dagger LU\| + \|LL^\dagger LU - LU\| + \|LU - AG\|. \quad (4.8)$$

Since $L^\dagger L = I$, it follows that $\|LL^\dagger LU - LU\| = 0$. Also $\|LL^\dagger AG - LL^\dagger LU\| \leq \|AG - LU\|$, and therefore we can use it in Eq. 4.8 to obtain:

$$\|LL^\dagger AG - AG\| \leq 2\|LU - AG\|. \quad (4.9)$$

By substituting Eq. 4.9 in Eq. 4.7 we get

$$\|LL^\dagger AGF - AGF\| \leq 2\|F\|\|LU - AG\|. \quad (4.10)$$

By substituting Eqs. 4.6 and 4.10 in Eq. 4.5 we get

$$\|LL^\dagger A - A\| \leq 2\|AGF - A\| + 2\|F\|\|LU - AG\|. \quad (4.11)$$

□

Lemma 4.5 appears in [23]. It uses a lower bound for the least singular value of a Gaussian matrix with zero mean and unit variance. This bound can be found in [6].

Lemma 4.5 ([23]). *Assume that k, l, m and n are positive integers such that $k \leq l$, $l \leq m$ and $l \leq n$. Assume that A is a real $m \times n$ matrix, G is $n \times l$ whose entries are i.i.d Gaussian random variables of zero mean and unit variance, β and γ are real numbers, such that $\beta > 0$, $\gamma > 1$ and the quantity*

$$1 - \frac{1}{\sqrt{2\pi(l-k+1)}} \left(\frac{e}{(l-k+1)\beta} \right)^{l-k+1} - \frac{1}{4(\gamma^2-1)\sqrt{\pi n \gamma^2}} \left(\frac{2\gamma^2}{e^{\gamma^2-1}} \right)^n \quad (4.12)$$

is nonnegative. Then, there exists a real $l \times n$ matrix F such that

$$\|AGF - A\| \leq \sqrt{2nl\beta^2\gamma^2 + 1}\sigma_{k+1} \quad (4.13)$$

and

$$\|F\| \leq \sqrt{l}\beta \quad (4.14)$$

with probability not less than the value in Eq. 4.12.

Lemma 4.6 rephrases Lemma 4.5 by utilizing the bounds that appear in [22]. The proof is close to the argumentation that appear in the proof of Lemma 4.5.

Lemma 4.6. *Let A be a real $m \times n$ ($m \geq n$) matrix. Let G be a real $n \times l$ matrix whose entries are Gaussian with zero mean and unit variance. Let k and l be integers such that $l < m$, $l < n$ and $l > (1 + \frac{1}{\ln k})k$. Let F be a real matrix of size $l \times n$. We define a_1, a_2, c_1, c_2 as in Theorem 3.6. Then,*

$$\|AGF - A\| \leq \sqrt{\frac{a_1^2 n}{c_1^2 l} + 1}\sigma_{k+1}(A), \quad (4.15)$$

and

$$\|F\| \leq \frac{1}{c_1 \sqrt{l}} \quad (4.16)$$

with probability not less than $1 - e^{-c_2 l} - e^{-a_2 n}$.

Proof. We begin by forming the SVD of A

$$A = U\Sigma V^T, \quad (4.17)$$

where U is orthogonal $m \times m$ matrix, Σ is $m \times n$ diagonal matrix with non-negative entries and V is orthogonal matrix $n \times n$. Given V^T and G , suppose that

$$V^T G = \begin{pmatrix} H \\ R \end{pmatrix}, \quad (4.18)$$

where H is $k \times l$ and R is $(n-k) \times l$. Since G is a Gaussian i.i.d. matrix and V is an orthogonal matrix, $V^T G$ is also a Gaussian i.i.d. matrix. Therefore, H is a Gaussian i.i.d. matrix. Let us define $F = PV^T$, where P is of size $l \times n$ such that

$$P = \begin{pmatrix} H^\dagger & 0 \end{pmatrix}.$$

Therefore,

$$F = \begin{pmatrix} H^\dagger & 0 \end{pmatrix} V^T. \quad (4.19)$$

Computing $\|F\|$ using Theorem 3.6 gives:

$$\|F\| = \|PV^T\| = \|H^\dagger\| = \|H^T(HH^T)^{-1}\| = \frac{1}{\sigma_k(H)} \leq \frac{1}{c_1\sqrt{l}} \quad (4.20)$$

with probability not less than $1 - e^{-c_2 l}$. Now we can bound $\|AGF - A\|$. By using Eqs. 4.17, 4.18 and 4.19 we get

$$AGF - A = U\Sigma \left(\begin{pmatrix} H \\ R \end{pmatrix} \begin{pmatrix} H^\dagger & 0 \end{pmatrix} - I \right) V^T. \quad (4.21)$$

We define S to be the upper-left $k \times k$ block of Σ and T to be the lower-right $(n-k) \times (n-k)$ block. Then,

$$\Sigma \left(\begin{pmatrix} H \\ R \end{pmatrix} \begin{pmatrix} H^\dagger & 0 \end{pmatrix} - I \right) = \begin{pmatrix} S & 0 \\ 0 & T \end{pmatrix} \begin{pmatrix} 0 & 0 \\ RH^\dagger & -I \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ TRH^\dagger & -T \end{pmatrix}.$$

The norm of the last term can be rewritten as:

$$\left\| \begin{pmatrix} 0 & 0 \\ TRH^\dagger & -T \end{pmatrix} \right\|^2 \leq \|TRH^\dagger\|^2 + \|T\|^2. \quad (4.22)$$

Therefore, by using Eqs. 4.21, 4.22 and the fact that $\|T\| = \sigma_{k+1}(A)$, we get

$$\|AGF - A\| \leq \sqrt{\|TRH^\dagger\|^2 + \|T\|^2} \leq \sqrt{\|H^\dagger\|^2 \|R\|^2 + 1} \sigma_{k+1}(A). \quad (4.23)$$

Also we know that

$$\|R\| \leq \|V^T G\| = \|G\| \leq a_1 \sqrt{n}$$

with probability not less than $1 - e^{-a_2 n}$. Combining Eq. 4.23 with the fact that $\|H^\dagger\| \leq \frac{1}{c_1\sqrt{l}}$ and $\|R\| \leq a_1\sqrt{n}$ gives:

$$\|AGF - A\| \leq \sigma_{k+1}(A) \sqrt{\frac{a_1^2 n}{c_1^2 l} + 1}. \quad (4.24)$$

□

Remark 4.7. In contrast to Lemma 4.5 where $\|AGF - A\| = \mathcal{O}(\sqrt{nl})$, Lemma 4.6 provides the bound $\|AGF - A\| = \mathcal{O}(\sqrt{\frac{n}{l}})$ that is better for large values of l .

Remark 4.8. *The condition $l > \left(1 + \frac{1}{\ln k}\right)k$ is satisfied without dramatic increase of the computational complexity of Algorithm 4.1. However, there are bounds for the case where H is almost square ($l \approx k$) and square ($l = k$). These bounds are given in [21].*

We now prove Theorem 4.3:

Proof of Theorem 4.3. The error is given by the expression $\|LU - PAQ\|$ where L, U, P and Q are the output of Algorithm 4.1 whose inputs were a matrix A , integers k and l . From Steps 7 and 8 in Algorithm 4.1 we have

$$\|LU - PAQ\| = \|L_y L_b U_b - PAQ\|. \quad (4.25)$$

Here, L_y is the $m \times k$ matrix from step 4. By using the fact that $BQ = L_b U_b = L_y^\dagger PAQ$, we get

$$\|LU - PAQ\| = \|L_y L_b U_b - PAQ\| = \|L_y L_y^\dagger PAQ - PAQ\|. \quad (4.26)$$

Applying Lemma 4.4 gives that

$$\begin{aligned} \|LU - PAQ\| &= \|L_y L_y^\dagger PAQ - PAQ\| \\ &\leq 2\|PAQ\tilde{G}F - PAQ\| + 2\|F\| \|L_y U_y - PAQ\tilde{G}\|. \end{aligned} \quad (4.27)$$

Here, U_y is the $k \times n$ matrix from step 4 in Algorithm 4.1. This holds for any random matrix \tilde{G} whose entries are i.i.d.. In particular, there exists a matrix G such that $Q\tilde{G} = GQ_y$, where G is also a random Gaussian i.i.d. matrix. This is in fact \tilde{G} after row and columns permutations. Therefore, the last term can be reformulated as $\|L_y U_y - PAQ\tilde{G}\| = \|L_y U_y - PAGQ_y\|$ where G is the random matrix from Algorithm 4.1. Applying Lemmas 3.2 and 3.3 to $\|L_y U_y - PAQ\tilde{G}\|$ gives

$$\begin{aligned} \|L_y U_y - PAQ\tilde{G}\| &= \|L_y U_y - PAGQ_y\| \\ &\leq (k(n-k) + 1)\sigma_{k+1}(AG) \\ &\leq (k(n-k) + 1)\|G\|\sigma_{k+1}(A). \end{aligned} \quad (4.28)$$

Lemma 4.5 gives that $\|PAQ\tilde{G}F - PAQ\| \leq \sqrt{2nl\beta^2\gamma^2 + 1}\sigma_{k+1}(A)$ and $\|F\| \leq \sqrt{l}\beta$. By combining Lemmas 4.5 and 3.4 we get

$$\|LU - PAQ\| \leq \left(2\sqrt{2nl\beta^2\gamma^2 + 1} + 2\sqrt{2nl}\beta\gamma(k(n-k) + 1)\right)\sigma_{k+1}(A) \quad (4.29)$$

which completes the proof. \square

Remark 4.9. *The error in Theorem 4.3 may appear large, especially for the case where $k \approx \frac{n}{2}$ and n is large. Yet, we performed extensive numerical experiments showing that the actual error was much smaller when using Gaussian elimination with partial pivoting. Note that the error can be decreased by increasing k . This can be applicable in certain applications. Numerical illustrations appear in section 5.*

We now present an additional error bound, which relies on the work in [22].

Theorem 4.10. *Given a matrix A of size $m \times n$, integers k and l such that $l > \left(1 + \frac{1}{\ln k}\right)k$ and $a_2 > 0$. Applying Algorithm 4.1 with A, k and l as its input parameters, produces a randomized LU decomposition that satisfies*

$$\|LU - PAQ\| \leq \left(2\sqrt{\frac{a_1^2 n}{c_1^2 l}} + 1 + \frac{2a_1\sqrt{n}}{c_1\sqrt{l}}(k(n-k) + 1)\right) \sigma_{k+1}(A), \quad (4.30)$$

with probability not less than $1 - e^{-a_2 n} - e^{-c_2 l}$. The value of c_1 is given in Eq. 3.13 and the value of c_2 is given in Eq. 3.15. Both values depend on a_2 .

Proof. By using steps 5,6,7 and 8 in Algorithm 4.1, we get that

$$\|LU - PAQ\| = \|L_y L_y^\dagger PAQ - PAQ\|. \quad (4.31)$$

Then, from Lemma 4.4,

$$\|L_y L_y^\dagger PAQ - PAQ\| \leq 2\|PAQ\tilde{G}F - PAQ\| + 2\|F\|\|L_y U_y - PAQ\tilde{G}\|. \quad (4.32)$$

From Lemma 4.6 we get that

$$\|PAQ\tilde{G}F - PAQ\| \leq \sqrt{\frac{a_1^2 n}{c_1^2 l}} + 1\sigma_{k+1}(A). \quad (4.33)$$

Using the same arguments given in Theorem 4.3, we get

$$\|L_y U_y - PAQ\tilde{G}\| = \|L_y U_y - PAGQ_y\| \leq (k(n-k) + 1)\|G\|\sigma_{k+1}(A) \quad (4.34)$$

where G is the matrix used in Algorithm 4.1 Step 1. Combining Eqs. 4.32, 4.33, 4.34 and since $\|F\| \leq \frac{1}{c_1\sqrt{l}}$, $\|G\| \leq a_1\sqrt{n}$ (see Lemma 4.6 and Theorem 3.5, respectively), we get that

$$\|LU - PAQ\| \leq 2\sqrt{\frac{a_1^2 n}{c_1^2 l}} + 1\sigma_{k+1}(A) + \frac{2a_1\sqrt{n}}{c_1\sqrt{l}}(k(n-k) + 1)\sigma_{k+1}(A). \quad (4.35)$$

Here, $\mu = \left(\frac{4}{\sqrt{2\pi}}\right)^{\frac{1}{3}}$, a_1 is given by Theorem 3.5 and c_1 is given in Eq. 3.13. \square

4.3 Randomized LU for Sparse Matrices

Assume that A is a sparse matrix. We want to compute its approximated LU factorization by applying Algorithm 4.1 to A . If the random matrix G , which is used in step 1 in Algorithm 4.1, is sparse, then AG computation can be done by using a sparse matrix multiplication. If this product is also sparse, then its LU decomposition can be done by using sparse LU which is more efficient since it depends only on the non-zero elements. Moreover, the result matrices L and U

are also sparse. This observation can help to accelerate even more the randomized LU on large sparse matrices. In this section, we derive the randomized LU approximation error bound using the tools introduced in Section 3.2. As can be seen from Theorem 4.3, the error depends on both the largest and the smallest singular values of the random matrix G which is now a sparse random matrix. Therefore, we shall compute the error bounds for the singular values of G when G is a sparse random matrix with density ρ . We present the Randomized LU algorithm for sparse matrices in Algorithm 4.2, and its error bound in Theorem 4.11.

Algorithm 4.2: Randomized LU Decomposition for Sparse Matrices

Input: A sparse matrix of size $m \times n$ to decompose; k desired rank; l number of columns to use; ρ - random matrix density.

Output: Matrices P, Q, L and U such that $\|PAQ = LU\| \leq \mathcal{O}(\sigma_{k+1}(A))$ where P and Q are orthogonal permutation matrices and L and U are sparse lower and upper triangular matrices, respectively.

- 1: Apply Algorithm 4.1 where G is a sparse random Gaussian matrix, with density ρ and its non-zero entries are $g_{ij} \sim \mathcal{N}(0, \frac{1}{\rho})$.

Algorithm 4.2 finds a rank k LU decomposition of A since it is based on Algorithm 4.1. However, following the proof of the error bound (Theorem 4.10) for a sparse Gaussian G is not applicable since it relies on the fact that GU is also a Gaussian i.i.d (U is an orthogonal matrix), which is not the case when G is a sparse Gaussian matrix. Therefore, we present another error bound (Theorem 4.11) for Algorithm 4.2.

Theorem 4.11. *Given a matrix A , integers k and l such that $l > (1 + \frac{1}{\ln k})k$, $k \approx \frac{l}{2}$, $a_2 > 0$ and $0 < \rho < 1$. By applying Algorithm 4.2 with A, k, l, ρ as its input parameters, we get a randomized LU decomposition that satisfies*

$$\|LU - PAQ\| \leq \left(2\sqrt{\frac{a_1^2 n}{c_1^2 l}} + 1 + \frac{2a_1 \sqrt{n}}{c_1 \sqrt{l}} (k(n-k) + 1) \right) \sigma_{k+1}(A), \quad (4.36)$$

with probability not less than $1 - e^{-a_2 n} - e^{-c_2 l}$. The values of c_1 and c_2 are given in Eqs. 3.13 and 3.15 respectively, which depend on ρ .

To prove the error bound for Algorithm 4.2, we present Conjecture 4.12 that bounds the k th singular value of a sparse subgaussian matrix multiplied by an orthogonal matrix.

Conjecture 4.12. *Let G be an $n \times l$ sparse Gaussian matrix with density $\rho \ll 1$. Assume that l and k are integers such that $k \approx \frac{l}{2}$. Q is an $n \times n$ orthogonal matrix. Define G_1 as a matrix with the k top rows of G and B_1 as a matrix with the k top rows of QG . Then, for n and l sufficiently large, $\sigma_k(G_1) \leq \sigma_k(B_1)$ with high probability.*

To verify that Conjecture 4.12 holds experimentally, we calculated the failure probability $\mathbb{P}(\sigma_k(G_1) > \sigma_k(B_1))$ by computing the k th singular value of B_1 and G_1 10,000 times for different values of n, l, k and ρ . The results are presented in Table 4.2.

Table 4.2: Probability of the correctness of Conjecture 4.12.

n	l	k	ρ	Average $\sigma_k(G_1)$	$\mathbb{P}(\sigma_k(G_1) > \sigma_k(B_1))$
3000	200	100	0.03	0.611	0
3000	300	200	0.03	1.324	0
3000	400	200	0.02	2.602	0
3000	700	400	0.02	4.165	0
3000	900	400	0.03	9.215	0
4000	300	100	0.03	3.663	0
4000	500	300	0.02	3.356	0
4000	700	400	0.03	5.718	0
4000	700	300	0.01	0.923	0
4000	700	300	0.02	6.578	0
4000	700	300	0.03	6.738	0

Conjecture 4.12 is used to prove Lemma 4.13, which is similar to Lemma 4.6.

Lemma 4.13. *Let A be a real $m \times n$ ($m \geq n$) matrix. Let G be a real $n \times l$ sparse Gaussian matrix with density $\rho \ll 1$. Assume that the entries of G are i.i.d. with zero mean and unit variance. Let k and l be integers such that $l < m$, $l < n$, $l > (1 + \frac{1}{\ln k})k$ and $k \approx \frac{l}{2}$. Let F be a real matrix of size $l \times n$. We define a_1, a_2, c_1, c_2 as in Theorem 3.6. Then*

$$\|AGF - A\| \leq \sqrt{\frac{a_1^2 n}{c_1^2 l} + 1} \sigma_{k+1}(A), \quad (4.37)$$

and

$$\|F\| \leq \frac{1}{c_1 \sqrt{l}} \quad (4.38)$$

with probability not less than $1 - e^{-c_2 l} - e^{-a_2 n}$.

Proof. The proof is almost identical to the proof of Lemma 4.6. Here, in order to prove Eq. 4.20 when G is a sparse Gaussian matrix, we use Conjecture 4.12 and take $Q = V^T$. Therefore, $\sigma_k(H) \geq \sigma_k(G_1)$ where G_1 is the upper $k \times l$ block of G . By combining Conjecture 4.12 and Theorem 3.6 we get

$$\|F\| = \frac{1}{\sigma_k(H)} \leq \frac{1}{\sigma_k(G_1)} \leq \frac{1}{c_1 \sqrt{l}} \quad (4.39)$$

where c_1 depends on $\mu = \left(\frac{4}{\sqrt{2\pi\rho}}\right)^{\frac{1}{3}}$. □

Proof of Theorem 4.11. The proof is almost identical to the proof of Theorem 4.10, except that we use Lemma 4.13 instead of Lemma 4.6 that is applicable to sparse Gaussian random matrices. The bound is the same, except that a_1 , c_1 and c_2 depend on $\mu = \frac{4}{\sqrt{2\pi\rho}}$. \square

Remark 4.14. *In practice, step 3 in Algorithm 4.2 can be done using standard LU decomposition with either partial or full pivoting. By using full pivoting we can sparsify L even further.*

4.4 Rank Deficient Least Squares

In this section, we present an application for the randomized LU and show how it can be used to solve efficiently the Rank Deficient Least Squares (RDLS) problem.

Assume that A is an $m \times n$ matrix ($m \geq n$) with $\text{rank}(A) = k$, $k < n$ and let b be a column vector of size $m \times 1$. We want to minimize $\|Ax - b\|$. Because A is a rank deficient matrix, then the problem has an infinite number of solutions since if x is a minimizer and $z \in \text{null}(A)$, then $x + z$ is also a minimizer (i.e. a valid solution). We now show that the complexity of the solution depends on the rank of A and that the problem is equivalent to solving the following two problems: a full rank Least Square (LS) problem of size $m \times k$, and a simplified undetermined linear system of equations that requires a matrix inversion of size $k \times k$.

The solution can be derived by the application of Algorithm 4.1 to A to get

$$\|Ax - b\| = \|P^T LUQ^T x - b\| = \|LUQ^T x - Pb\|, \quad (4.40)$$

where L is an $m \times k$ matrix, U is an $k \times n$ matrix and both L and U are of rank k . Let $y = UQ^T x$ and $c = Pb$. Then, the problem can be reformulated as

$$\min \|Ly - c\|. \quad (4.41)$$

Note that L is a full rank matrix and the problem to be solved becomes a standard full rank LS problem. The solution is given by $y = L^\dagger c$. Next, we solve

$$Uz = y, \quad (4.42)$$

where $z = Q^T x$. Since U is $k \times n$ matrix, Eq. 4.42 it is an under-determined system. Assume that $U = [U_1 \ U_2]$ and $z = [z_1 \ z_2]^T$, where U_1 is a $k \times k$ matrix, z_1 is a $k \times 1$ vector and z_2 is a $(n - k) \times 1$ vector. Then, the solution is given by setting any value to z_2 and solving

$$U_1 z_1 = y - U_2 z_2. \quad (4.43)$$

For simplicity, we choose $z_2 = 0$. Therefore, we get that $z_1 = U_1^{-1} y$. The final solution is given by $x = Qz$. This procedure is summarized in Algorithm 4.3

that finds the solution to the deficient least square problem using Algorithm 4.1.

Algorithm 4.3: Solving Rank Deficient Least Squares with Randomized LU

Input: A matrix of size $m \times n$ with rank k ; l integer such that $l \geq k$; b vector of size $m \times 1$.

Output: x a solution that minimizes $\|Ax - b\|$.

- 1: Apply Algorithm 4.1 to A with parameters k and l and save the returned matrices P, Q, L, U, L^\dagger .
 - 2: $y \leftarrow L^\dagger P b$.
 - 3: $z_1 \leftarrow U_1^{-1} y$.
 - 4: $z \leftarrow \begin{pmatrix} z_1 \\ z_2 \end{pmatrix}$, where z_2 is an $n - k$ zero vector.
 - 5: $x \leftarrow Q z$.
-

The complexity of Algorithm 4.3 is equal to the randomized LU complexity (Algorithm 4.1) with an additional inversion of the matrix U_1 in Step 3, which is of size $k \times k$. Note that the solution given by Algorithm 4.3 is sparse in the sense that x contains at most k non-zero entries.

5 Numerical Results

In order to evaluate Algorithm 4.1, we present the numerical results by comparing the low rank approximation of several randomized algorithms with their execution times. We tested the algorithms and compared between them by applying them to random and sparse matrices and to images. All the results were computed using the standard MATLAB libraries on a machine with two Intel Xeon CPUs X5560 2.8GHz and with an nVidia GPU GTX TITAN card.

5.1 Error Rate and Computational Time Comparisons

The performance of the randomized LU (Algorithm 4.1) was tested and compared to a randomized SVD and to a randomized ID (see [17, 23]). The tests compare the normalized (relative) error of the low rank approximation obtained by the examined methods and also measure the computational time of each method. If A is the original matrix and \hat{A} is a low rank approximation of A , then the relative approximation error is given by:

$$\text{err} = \frac{\|A - \hat{A}\|}{\|A\|}. \quad (5.1)$$

First, we compare between the low rank approximation achieved by randomized SVD, randomized ID and randomized LU for different ranks k . Throughout the experiments, we chose $l = k + 3$ and the test matrix was a random matrix of size 2000×2000 with exponentially decaying singular values. The computation was done in a single precision. The results are presented in Fig. 5.1. The

experiment shows that the error of the randomized ID is significantly larger than the error from both randomized SVD and randomized LU (Algorithm 4.1), which are almost identical. Next, we compare between the execution time of these algorithms. The results are presented in Fig. 5.2. The results show that the execution time of the randomized LU is significantly faster (about $\times 30$ faster) than the execution time of the randomized SVD and randomized ID. LU can be parallelized very efficiently (see [15] section 3.6). Therefore, the execution time of computing a randomized LU of a matrix of size 2000×2000 was measured on an nVidia GTX TITAN GPU device and it is shown in Fig. 5.3. The execution time on the GPU was up to $\times 15$ faster than running it on an eight cores CPU. Thus, the algorithm scales well. For larger matrices (n and k are large) the difference between the performances on CPU and on GPU are more significant.

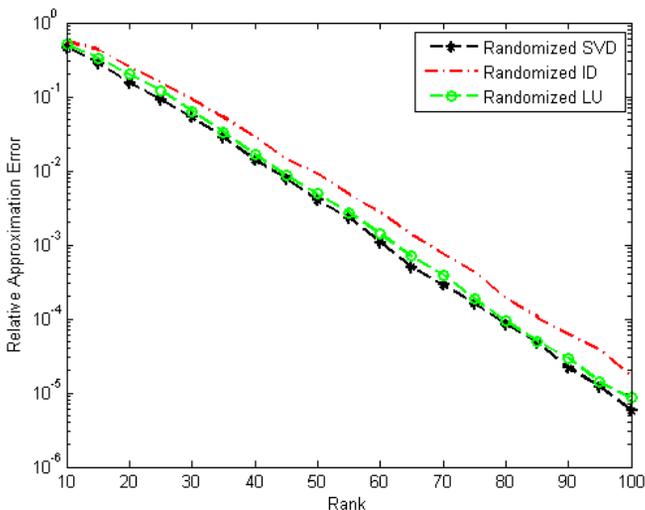


Figure 5.1: Comparison between the low rank approximation error of different algorithms: Randomized SVD, Randomized ID and Randomized LU.

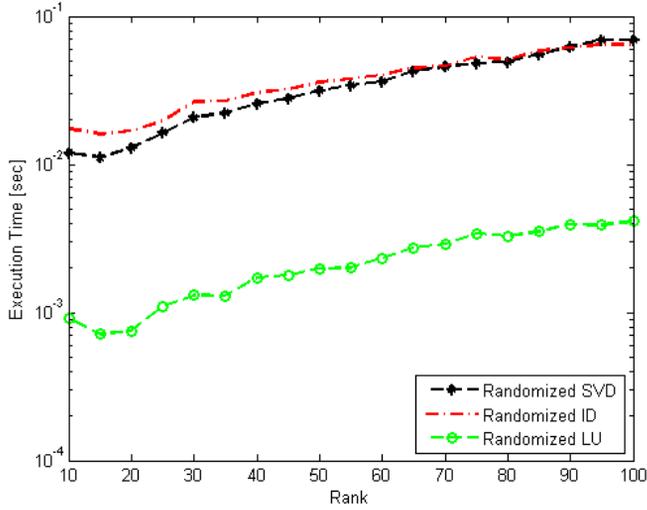


Figure 5.2: Comparison between the execution times of the same algorithms as in Fig. 5.1 running on a CPU.

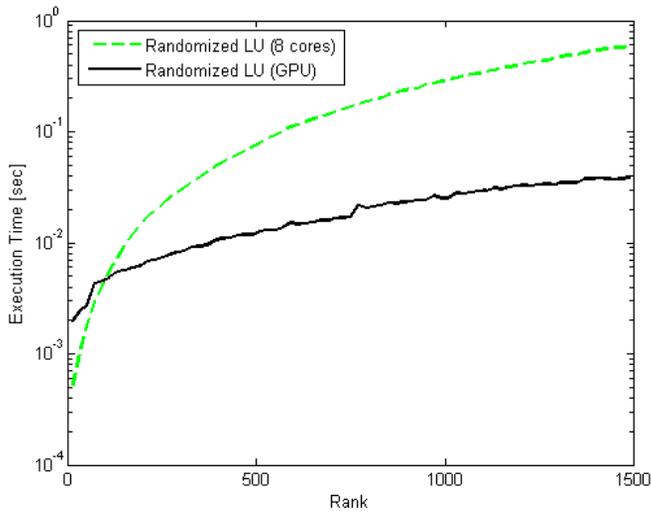


Figure 5.3: Comparison between the execution times from running Algorithm 4.1 on different computational platforms: CPU with 8 cores and GPU.

5.2 Image Matrix Factorization

We tested Algorithm 4.1 by applying it to images represented by a matrix. The factorization error and the execution time were compared with the randomized SVD and the randomized ID. We also added the SVD error computation and execution time as benchmark computed using Lanczos method with the PROPACK package. The image size was 2124×7225 pixels and it has 256 gray levels. The used parameters were $k = 200, l = 203$. The approximation quality (error) was measured in PSNR defined by:

$$\text{PSNR} = 20 \log_{10} \frac{\max_A \sqrt{N}}{\|A - \hat{A}\|_F} \quad (5.2)$$

where A is the original image, \hat{A} is the approximated image, \max_A is the maximal pixel value of A and N the total numbers of pixels. $\|\cdot\|_F$ stands for the Frobenius norm. The higher the PSNR is the better the approximation is.

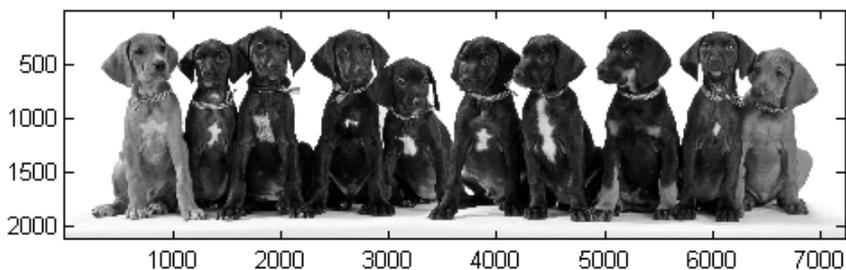


Figure 5.4: The original input image that was factorized by the randomized LU, randomized ID and randomized SVD algorithms.

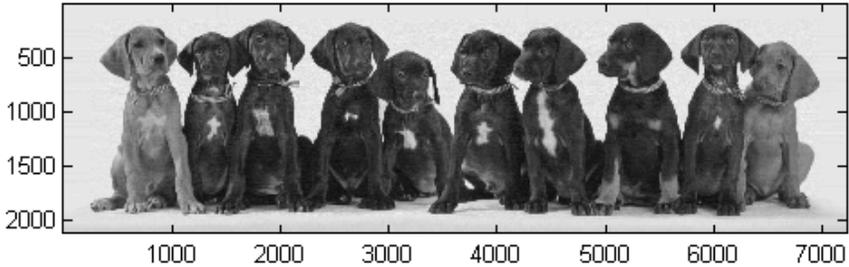


Figure 5.5: The reconstructed image from the randomized LU factorization with $k = 200$ and $l = 203$.

Figures 5.4 and 5.5 show the original and the reconstructed images, respectively. The image approximation quality (measured in PSNR) is shown in Fig. 5.6 where the PSNR of Algorithm 4.1 is higher than the randomized ID and almost identical to the randomized SVD. Also, the PSNR is close to the result achieved by Lanczos SVD which is the best possible rank k approximation. The execution time of each algorithm is shown in Fig. 5.7. All the computations were done using double precision. Here, the randomized LU is significantly faster than all the other compared methods making it applicable for real time applications.

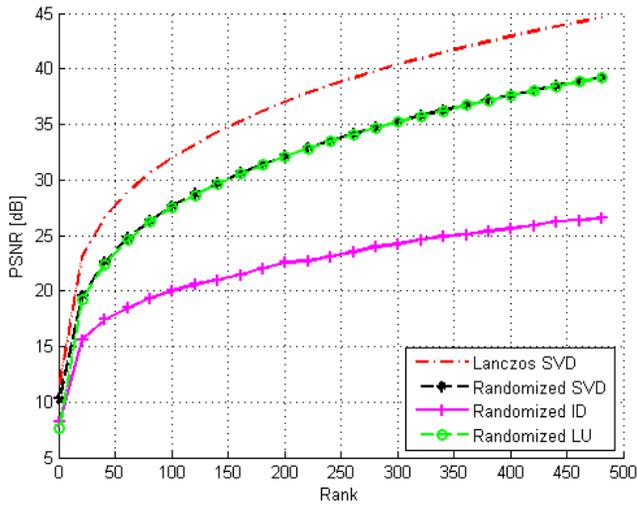


Figure 5.6: Comparison between the PSNR values from image reconstruction application using randomized LU, randomized ID, randomized SVD and Lanczos SVD algorithms.

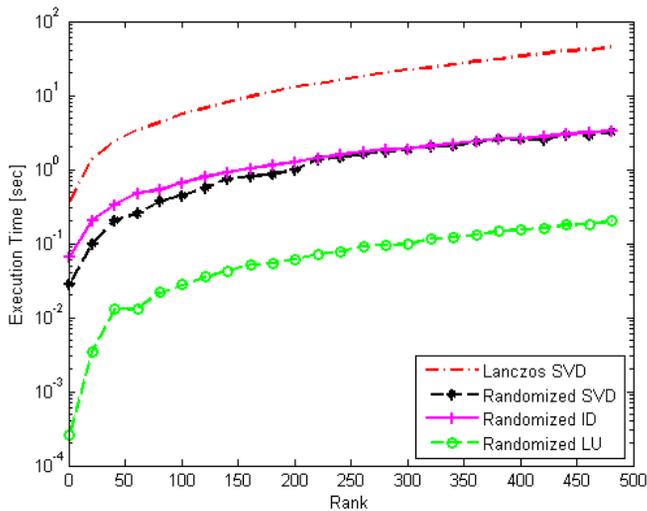


Figure 5.7: Comparison between the execution time of the randomized LU, randomized ID, randomized SVD and Lanczos SVD algorithms.

5.3 Sparse Matrix Factorization

In order to test Algorithm 4.1 on sparse matrices, we used a binary sparse matrix of size $862,664 \times 862,664$ with 19,235,140 non zero elements ($\rho = 2.58 \times 10^{-5}$) that contains the results of crawling the .eu domain. The eu-2005 matrix was generated and studied in [4]. Each edge in the eu-2005 graph represents a link between two websites. The approximation error of each algorithm applied to the eu-2005 matrix is shown in Fig. 5.8, and the execution time is shown in Fig. 5.9.

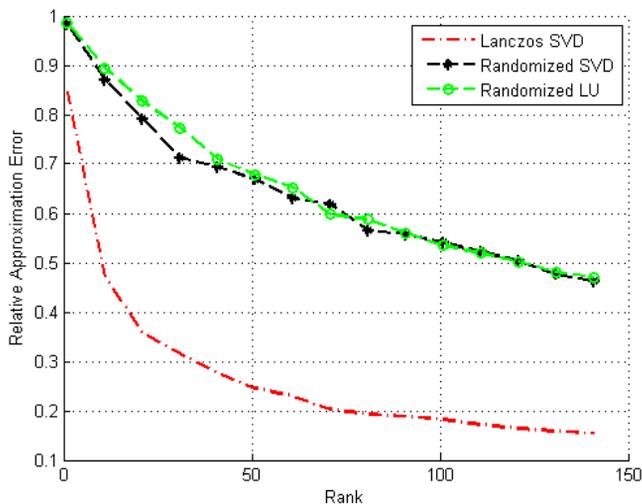


Figure 5.8: Comparison between the approximation error of the randomized LU, randomized ID and randomized SVD algorithms, executed on the sparse matrix eu-2005.

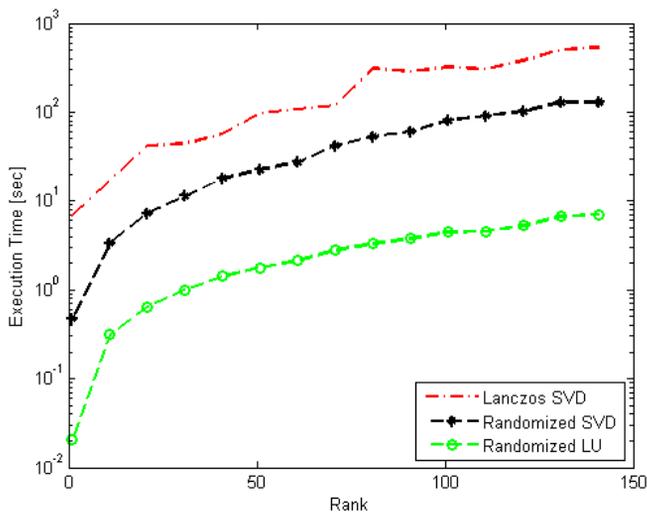


Figure 5.9: Comparison between the execution time of the randomized LU, randomized ID and randomized SVD algorithms, executed on the sparse matrix eu-2005.

To test the factorization for large values of k , we applied Algorithms 4.1 and 4.2 to a $100,000 \times 100,000$ random generated sparse matrix A with density $\rho = 10^{-6}$. Applying Algorithm 4.1 to this matrix using $k = 10,000$ requires a large memory size and CPU resources. This resulted in a very long execution time. A similar experiment was conducted by applying Algorithm 4.2 to the same matrix A with k ranging from 13,000 to 25,000. The random matrix G , selected in Step 1 of Algorithm 4.2, was a random generated sparse matrix with two possible densities $\rho = 10^{-4}$ and $\rho = 3 \times 10^{-4}$. The approximation error for each density-based execution is shown in Fig. 5.10 and the execution time is shown in Fig. 5.11. The graphs represent an averaging from successive application of the algorithm.

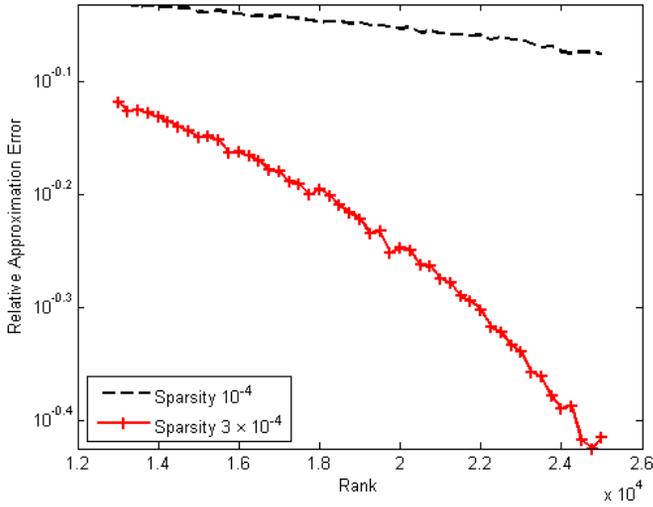


Figure 5.10: Approximation error from the application of Algorithm 4.2 to the matrix A with two different densities.

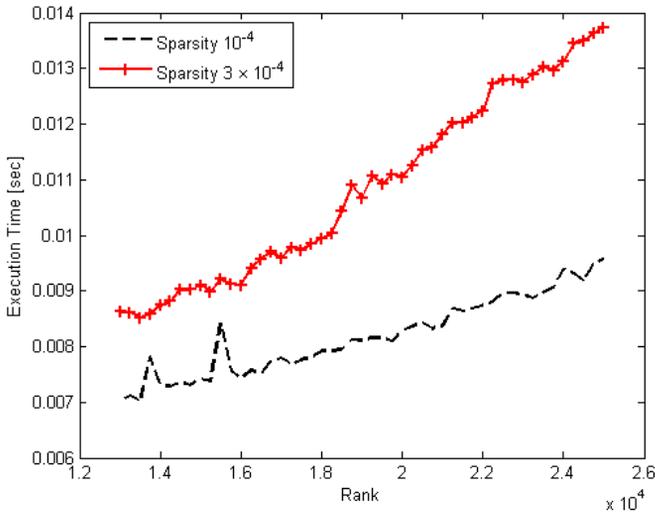


Figure 5.11: Execution time from the application of Algorithm 4.2 to the matrix A with two different densities.

The approximation error in Fig. 5.10 depends on the sparsity ρ of the random matrix G . The error increases as G became sparser. Figure 5.11 shows the dependency between execution time and sparsity. A sparser matrix reduces

the algorithm execution time. This trade-off enables the user to choose between speed and accuracy by selecting a proper ρ value.

6 Conclusion

In this work, we presented a randomized algorithm for computing a low rank LU decomposition. Given an integer k , the algorithm finds an LU decomposition such that both L and U are of rank k with negligible failure probability. We constructed error bounds for the approximation of the input matrix and proved that they are proportional the $k + 1$ singular value. We also compared the performance of the algorithm with regard to the error rate and to the computational time. We compared the results to randomized SVD, randomized ID and to Lanczos method of SVD running on sparse matrices. We also showed that our algorithm can be parallelized since it consists mostly of matrix multiplication and pivoted LU. The results on GPU show that it was possible to accelerate the computational time significantly even by using only the standard MATLAB libraries.

Acknowledgment

This research was partially supported by the Israel Science Foundation (Grant No. 1041/10), by the Israeli Ministry of Science & Technology (Grant No. 3-909), by US - Israel Binational Science Foundation (BSF 2012282) and by a Fellowship from Jyväskylä University. The authors would like to thank Yoel Shkolnisky for the helpful discussions.

References

- [1] L. BACKSTROM, P. BOLDI, M. ROSA, J. UGANDER, AND S. VIGNA, *Four degrees of separation*, in Proceedings of the 3rd Annual ACM Web Science Conference, ACM, 2012, pp. 33–42.
- [2] A. BERMANIS, A. AVERBUCH, AND R. COIFMAN, *Multiscale data sampling and function extension*, Applied and Computational Harmonic Analysis, 34 (2013), pp. 15–29.
- [3] R. BHATIA, *Matrix analysis*, vol. 169, Springer, 1997.
- [4] P. BOLDI, B. CODENOTTI, M. SANTINI, AND S. VIGNA, *Ubicrawler: A scalable fully distributed web crawler*, Software: Practice & Experience, 34 (2004), pp. 711–726.
- [5] T. F. CHAN, *Rank revealing QR factorizations*, Linear Algebra and Its Applications, 88 (1987), pp. 67–82.

- [6] Z. CHEN AND J. J. DONGARRA, *Condition numbers of gaussian random matrices*, SIAM Journal on Matrix Analysis and Applications, 27 (2005), pp. 603–620.
- [7] H. CHENG, Z. GIMBUTAS, P. MARTINSSON, AND V. ROKHLIN, *On the compression of low rank matrices*, SIAM Journal on Scientific Computing, 26 (2005), pp. 1389–1404.
- [8] G. DAVID, *Anomaly Detection and Classification via Diffusion Processes in Hyper-Networks*, PhD thesis, School of Computer Science, Tel Aviv University, March 2009.
- [9] T. A. DAVIS AND I. S. DUFF, *An unsymmetric-pattern multifrontal method for sparse lu factorization*, SIAM Journal on Matrix Analysis and Applications, 18 (1997), pp. 140–158.
- [10] J. W. DEMMEL, S. C. EISENSTAT, J. R. GILBERT, X. S. LI, AND J. W. LIU, *A supernodal approach to sparse partial pivoting*, SIAM Journal on Matrix Analysis and Applications, 20 (1999), pp. 720–755.
- [11] D. L. DONOHO, *Compressed sensing*, Information Theory, IEEE Transactions on, 52 (2006), pp. 1289–1306.
- [12] P. DRINEAS, M. W. MAHONEY, AND S. MUTHUKRISHNAN, *Relative-error cur matrix decompositions*, SIAM Journal on Matrix Analysis and Applications, 30 (2008), pp. 844–881.
- [13] M. ELAD AND M. AHARON, *Image denoising via sparse and redundant representations over learned dictionaries*, Image Processing, IEEE Transactions on, 15 (2006), pp. 3736–3745.
- [14] H. H. GOLDSTINE AND J. VON NEUMANN, *Numerical inverting of matrices of high order. ii*, Proceedings of the American Mathematical Society, 2 (1951), pp. 188–202.
- [15] G. H. GOLUB AND C. F. VAN LOAN, *Matrix computations*, vol. 4, John Hopkins University Press, 2012.
- [16] S. A. GOREINOV, E. E. TYRTYSHNIKOV, AND N. L. ZAMARASHKIN, *A theory of pseudoskeleton approximations*, Linear Algebra and its Applications, 261 (1997), pp. 1–21.
- [17] 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.
- [18] D. KIRK, *Nvidia cuda software and gpu parallel computing architecture*, in ISMM, vol. 7, 2007, pp. 103–104.
- [19] Y. KOREN, R. BELL, AND C. VOLINSKY, *Matrix factorization techniques for recommender systems*, Computer, 42 (2009), pp. 30–37.

- [20] R. LARSEN, *Lanczos bidiagonalization with partial reorthogonalization*, Tech. Report DAIMI PB-357, Department of Computer Science, Aarhus University, 1998.
- [21] A. LITVAK AND O. RIVASPLATA, *Smallest singular value of sparse random matrices*, *Stud. Math.*, 212 (2010), pp. 195–218.
- [22] A. E. LITVAK, A. PAJOR, M. RUDELSON, AND N. TOMCZAK-JAEGERMANN, *Smallest singular value of random matrices and geometry of random polytopes*, *Advances in Mathematics*, 195 (2005), pp. 491–523.
- [23] P. MARTINSSON, V. ROKHLIN, AND M. TYGERT, *A randomized algorithm for the decomposition of matrices*, *Applied and Computational Harmonic Analysis*, 30 (2011), pp. 47–68.
- [24] R. MAZUMDER, T. HASTIE, AND R. TIBSHIRANI, *Spectral regularization algorithms for learning large incomplete matrices*, *The Journal of Machine Learning Research*, 99 (2010), pp. 2287–2322.
- [25] A. V. OPPENHEIM, R. W. SCHAFER, J. R. BUCK, ET AL., *Discrete-time signal processing*, vol. 5, Prentice Hall Upper Saddle River, 1999.
- [26] C.-T. PAN, *On the existence and computation of rank-revealing LU factorizations*, *Linear Algebra and its Applications*, 316 (2000), pp. 199–222.
- [27] V. ROKHLIN AND M. TYGERT, *A fast randomized algorithm for overdetermined linear least-squares regression*, *Proceedings of the National Academy of Sciences*, 105 (2008), pp. 13212–13217.
- [28] M. RUDELSON AND R. VERSHYNIN, *Smallest singular value of a random rectangular matrix*, *Communications on Pure and Applied Mathematics*, 62 (2009), pp. 1707–1739.
- [29] O. SCHENK, K. GÄRTNER, AND W. FICHTNER, *Efficient sparse lu factorization with left-right looking strategy on shared memory multiprocessors*, *BIT Numerical Mathematics*, 40 (2000), pp. 158–176.
- [30] Y. SHMUELI, G. SHABAT, A. BERMANIS, AND A. AVERBUCH, *Accelerating particle filter using multiscale methods*, in *Electrical & Electronics Engineers in Israel (IEEEI)*, 2012 IEEE 27th Convention of, IEEE, 2012, pp. 1–4.
- [31] G. STEWART, *The triangular matrices of gaussian elimination and related decompositions*, Tech. Report TR-3533, Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD, 1995.
- [32] L. N. TREFETHEN AND R. S. SCHREIBER, *Average-case stability of gaussian elimination*, *SIAM Journal on Matrix Analysis and Applications*, 11 (1990), pp. 335–360.

- [33] R. YUSTER AND U. ZWICK, *Fast sparse matrix multiplication*, ACM Transactions on Algorithms (TALG), 1 (2005), pp. 2–13.