

Fast Monte Carlo Algorithms for Matrices III: Computing a Compressed Approximate Matrix Decomposition *

Petros Drineas [†] Ravi Kannan ^{‡§} Michael W. Mahoney [¶]

Technical Report, YALEU/DCS/TR-1271, February 2004.

Abstract

In many applications, the data consist of (or may be naturally formulated as) an $m \times n$ matrix A which may be stored on disk but which is too large to be read into RAM or to practically perform superlinear polynomial time computations on it. Two algorithms are presented which, when given an $m \times n$ matrix A , compute approximations to A which are the product of three smaller matrices, C , U , and R , each of which may be computed rapidly. Let $A' = CUR$ be the computed approximate decomposition; both algorithms have provable bounds for the error matrix $A - A'$. In the first algorithm, $c = O(1)$ columns of A and $r = O(1)$ rows of A are randomly chosen. If the $m \times c$ matrix C consists of those c columns of A (after appropriate rescaling) and the $r \times n$ matrix R consists of those r rows of A (also after appropriate rescaling) then the $c \times r$ matrix U may be calculated from C and R . For any matrix X , let $\|X\|_F$ and $\|X\|_2$ denote its Frobenius norm and its spectral norm, respectively. It is proven that

$$\|A - A'\|_\xi \leq \min_{D: \text{rank}(D) \leq k} \|A - D\|_\xi + \text{poly}(k, 1/c) \|A\|_F$$

holds in expectation and with high probability for both $\xi = 2, F$ and for all $k = 1, \dots, \text{rank}(A)$; thus by appropriate choice of k

$$\|A - A'\|_2 \leq \epsilon \|A\|_F$$

also holds in expectation and with high probability. This algorithm may be implemented without storing the matrix A in Random Access Memory (RAM), provided it can make two passes over the matrix stored in external memory and use $O(m + n)$ additional RAM memory. The second algorithm is similar except that it approximates the matrix C by randomly sampling $O(1)$ rows of C . Thus, it has additional error but it can be implemented in three passes over the matrix using only constant additional RAM memory. To achieve an additional error (beyond the best rank k approximation) that is at most $\epsilon \|A\|_F$, both algorithms take time which is a low-degree polynomial in k , $1/\epsilon$, and $1/\delta$, where $\delta > 0$ is a failure probability; the first takes time linear in $\max(m, n)$ and the second takes time independent of m and n . The proofs for the error bounds make important use of matrix perturbation theory and previous work on approximating matrix multiplication and computing low-rank approximations to a matrix. The probability distribution over columns and rows and the rescaling are crucial features of the algorithms and must be chosen judiciously.

* A preliminary version of this paper appeared as “Pass efficient algorithms for approximating large matrices” by P. Drineas and R. Kannan in the *Proceedings of the 14th ACM-SIAM Symposium on Discrete Algorithms*, 2003, pp. 223–232.

[†]Department of Computer Science, Rensselaer Polytechnic Institute, Troy, New York 12180, drinep@cs.rpi.edu

[‡]Department of Computer Science, Yale University, New Haven, Connecticut, USA 06520, kannan@cs.yale.edu

[§]Supported in part by a grant from the NSF.

[¶]Department of Mathematics, Yale University, New Haven, Connecticut, USA 06520, mahoney@cs.yale.edu

1 Introduction

We are interested in developing and analyzing fast Monte Carlo algorithms for performing useful computations on large matrices. In this paper we consider a new method for computing a compressed approximate decomposition of a large matrix; in two related papers we consider matrix multiplication and the Singular Value Decomposition (SVD) [10, 11]. Since such computations generally require time which is superlinear in the number of nonzero elements of the matrix, we expect our algorithms to be useful in many applications where data sets are modeled by matrices and are extremely large. In all these cases, we assume that the input matrices are prohibitively large to store in Random Access Memory (RAM) and thus that only external memory storage is possible. Thus, our algorithms will be allowed to read the matrices a few, e.g., one or two or three, times and keep a small randomly-chosen and rapidly-computable “sketch” of the matrices in RAM; computations will then be performed on this “sketch”. We will work within the framework of the Pass-Efficient computational model, in which the scarce computational resources are the number of passes over the data, the additional RAM space required, and the additional time required [10, 8].

In many applications, an $m \times n$ matrix A is stored on disk and is too large to be read into RAM or to practically perform superlinear polynomial time computations on it; thus, one may be interested in a succinctly-described, easily-computed $m \times n$ matrix A' that is an approximation to A . Let c and r be positive, usually constant, integers that we choose, and for any matrix X let $\|X\|_F$ and $\|X\|_2$ denote its Frobenius and spectral norm (as defined in Section 2.1), respectively. We present two algorithms that compute an approximation A' to the matrix A that has the following properties:

1. $A' = CUR$, where C is an $m \times c$ matrix consisting of c randomly picked columns of A , R is an $r \times n$ matrix consisting of r randomly picked rows of A and U is a $c \times r$ matrix computed from C, R ,
2. C, U , and R can be defined after making a small constant number of passes through the whole matrix A from disk,
3. U can be constructed using additional RAM space and time that is $O(m + n)$ (for the LINEARTIMECUR algorithm) or is $O(1)$ (for the CONSTANTTIMECUR algorithm),
4. for every $\epsilon > 0$ and every k such that $1 \leq k \leq \text{rank}(A)$ we can choose c and r such that, with high probability, A' satisfies

$$\|A - A'\|_2 \leq \min_{D: \text{rank}(D) \leq k} \|A - D\|_2 + \epsilon \|A\|_F,$$

and thus we can choose c and r such that $\|A - A'\|_2 \leq \epsilon \|A\|_F$, and

5. for every $\epsilon > 0$ and every k such that $1 \leq k \leq \text{rank}(A)$ we can choose c and r such that, with high probability, A' satisfies

$$\|A - A'\|_F \leq \min_{D: \text{rank}(D) \leq k} \|A - D\|_F + \epsilon \|A\|_F.$$

In the first algorithm, the LINEARTIMECUR algorithm of Section 3, $c = O(1)$ columns of A and $r = O(1)$ rows of A are randomly sampled. If the $m \times c$ matrix C consists of those c columns of A (after appropriate rescaling) and the $r \times n$ matrix R consists of those R rows

addl. error for:	LINEARTIMECUR	CONSTANTTIMECUR
$\ A - A'\ _2$	$c, r = \frac{\eta^2}{\epsilon^4}, \frac{k}{\delta^2 \epsilon^2}$	$c, w, r = \frac{\eta^2}{\epsilon^8}, \frac{\eta^2}{\epsilon^8}, \frac{k}{\delta^2 \epsilon^2}$
$\ A - A'\ _F$	$c, r = \frac{k\eta^2}{\epsilon^4}, \frac{k}{\delta^2 \epsilon^2}$	$c, w, r = \frac{k^2\eta^2}{\epsilon^8}, \frac{k^2\eta^2}{\epsilon^8}, \frac{k}{\delta^2 \epsilon^2}$

Figure 1: Summary of sampling complexity

of A (also after appropriate rescaling) then from C and R the $c \times r$ matrix U is computed. This algorithm may be implemented without storing the matrix A in RAM, provided it can make two passes over the matrix stored in external memory and use $O(m + n)$ additional RAM memory. The second algorithm, the CONSTANTTIMECUR algorithm of Section 4, is similar except that it approximates the matrix C by randomly sampling a constant number w of rows of C . Thus, our second algorithm only requires constant additional RAM memory but it uses a third pass over the data and has additional error. To achieve an additional error (beyond the best rank k approximation) that is at most $\epsilon \|A\|_F$, both algorithms take time which is a low-degree polynomial in k , $1/\epsilon$, and $1/\delta$, where $\delta > 0$ is a failure probability; the first algorithm takes time linear in $\max(m, n)$ and the second takes time independent of m and n . See Figure 1 for a summary of the dependence of the sampling complexity on k and ϵ , δ , and $\eta = 1 + \sqrt{8 \log(1/\delta)}$.

The proofs for the error bounds make important use of linear algebra and matrix perturbation theory; see [14, 15, 19, 5] for an overview of these topics. In particular, the proofs use the approximate matrix multiplication results of [10] and the approximate singular value decomposition results of [11]. As with those previous works, the probability distribution over the columns, the probability distribution over the rows, and the respective rescaling are crucial features of the algorithms which must be chosen judiciously. In addition, as a by-product of the CUR decomposition, we can estimate the top k singular values of A .

Our CUR approximations may be viewed as a “dimension reduction” technique. Two common techniques for dimension reduction are the SVD and Multidimensional Scaling; see [14, 18]. Another method that has attracted renewed interest recently is the traditional “Random Projection” method where one projects the problem into a randomly chosen low-dimensional subspace [17, 20, 16]. These methods do not share (1) and (2) and thus are not suited for very large problems. Our algorithms achieve (1) and (2) at the cost of the $\epsilon \|A\|_F$ error. In addition, our CUR approximations may be viewed as an approximate decomposition of a matrix $A \approx CUR$; as with other decompositions, the CUR decomposition both reveals information about the structure of the matrix and allows computations to be performed more efficiently. For example, in applications of the LINEARTIMECUR algorithm, A' could be stored in RAM since it can be stored in $O(m + n)$ space (instead of $O(mn)$ space); in addition, A' could then be operated upon by, e.g., applying $A' = CUR$ to a query vector $x \in \mathbb{R}^n$, which is an operation that can be performed in $O(m + n)$ space (instead of $O(mn)$ space if A is used).

Our CUR approximations have been used in applications such as the reconstruction of a matrix given a sample of the matrix in a recommendation systems context and for “similarity query” problems which are widely used in areas such as Information Retrieval [12]. In this application, after A has been preprocessed, one gets “query” vectors x and must find the similarity of x to each row of A . Here, the similarity of two vectors is defined to be their dot product or their normalized dot product; our technique can handle both. See [6] for related discussion. Note that the measure $\|A\|_2$ is a worst case measure and that this is more useful in many contexts than an average case measure like $\|A\|_F$, since the relevant query x often comes from a small dimensional subspace and is not random. See also [2, 1] for a nice discussion of these issues. Our CUR approximations have also been used in theoretical applications such as designing and

analyzing approximation algorithms for the Max-Cut problem [9]. In these applications, the use of the constant additional space and time framework is essential.

In other related work, Achlioptas and McSherry have also computed succinctly-described matrix approximations using somewhat different sampling techniques [2, 1]. Also included in [2, 1] is a comparison of their methods with those of [7, 8, 13] and thus with the results we present here. When compared with our LINEARTIMECUR algorithm, they achieve the same results for the Frobenius norm bound and slightly better results (with respect to $1/\epsilon$) for the spectral norm bound [1]; in their work, however, there is no analogue of our CONSTANTTIMECUR algorithm.

After this introduction, we provide in Section 2 a review of linear algebra, the Pass-Efficient model, and of several previous results from [10, 11] that are used in this paper. In Section 3 we describe and analyze the LINEARTIMECUR algorithm which computes an approximate CUR decomposition of a matrix A using linear (in m and n) additional space and time, and in Section 4 we describe and analyze the CONSTANTTIMECUR algorithm which computes a description of an approximate CUR decomposition of a matrix A using only constant additional space and time. Finally, in Section 5 we provide a discussion and conclusion.

2 Review of Relevant Background

2.1 Review of Linear Algebra

This section contains a review of linear algebra that will be useful throughout the paper; for more details, see [14, 15, 19, 5] and references therein.

For a vector $x \in \mathbb{R}^n$ we let $|x| = (\sum_{i=1}^n |x_i|^2)^{1/2}$ denote its Euclidean length. For a matrix $A \in \mathbb{R}^{m \times n}$ we let $A^{(j)}$, $j = 1, \dots, n$, denote the j -th column of A as a column vector and $A_{(i)}$, $i = 1, \dots, m$, denote the i -th row of A as a row vector. We denote matrix norms by $\|A\|_\xi$, using subscripts to distinguish between various norms. Of particular interest will be the Frobenius norm, the square of which is $\|A\|_F^2 = \sum_{i=1}^m \sum_{j=1}^n A_{ij}^2$, and the spectral norm, which is defined by $\|A\|_2 = \sup_{x \in \mathbb{R}^n, x \neq 0} \frac{|Ax|}{|x|}$. These norms are related to each other as: $\|A\|_2 \leq \|A\|_F \leq \sqrt{n} \|A\|_2$. If $A \in \mathbb{R}^{m \times n}$, then there exist orthogonal matrices $U = [u^1 u^2 \dots u^m] \in \mathbb{R}^{m \times m}$ and $V = [v^1 v^2 \dots v^n] \in \mathbb{R}^{n \times n}$ where $\{u^t\}_{t=1}^m \in \mathbb{R}^m$ and $\{v^t\}_{t=1}^n \in \mathbb{R}^n$ are such that

$$U^T A V = \Sigma = \mathbf{diag}(\sigma_1, \dots, \sigma_\rho),$$

where $\Sigma \in \mathbb{R}^{m \times n}$, $\rho = \min\{m, n\}$ and $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_\rho \geq 0$. Equivalently, $A = U \Sigma V^T$. The three matrices U , V , and Σ constitute the Singular Value Decomposition (SVD) of A . The σ_i are the singular values of A and the vectors u^i , v^i are the i -th left and the i -th right singular vectors, respectively. If $k \leq r = \text{rank}(A)$ and we define $A_k = U_k \Sigma_k V_k^T = \sum_{t=1}^k \sigma_t u^t v^{tT}$ then the distance (as measured by both $\|\cdot\|_2$ and $\|\cdot\|_F$) between A and any rank k approximation to A is minimized by A_k , i.e.,

$$\min_{D \in \mathbb{R}^{m \times n}: \text{rank}(D) \leq k} \|A - D\|_2 = \|A - A_k\|_2 = \sigma_{k+1}(A), \text{ and} \quad (6)$$

$$\min_{D \in \mathbb{R}^{m \times n}: \text{rank}(D) \leq k} \|A - D\|_F^2 = \|A - A_k\|_F^2 = \sum_{t=k+1}^r \sigma_t^2(A). \quad (7)$$

2.2 Review of the Pass-Efficient Model

The Pass-Efficient model of data-streaming computation is a computational model that is motivated by the observation that in modern computers the amount of disk storage, i.e., sequential

access memory, has increased very rapidly, while RAM and computing speeds have increased at a substantially slower pace [10, 8]. In the Pass-Efficient model the three scarce computational resources are number of passes over the data and the additional RAM space and additional time required by the algorithm. The data are assumed to be stored on a disk, to consist of elements whose size is bounded by a constant, and to be presented to an algorithm on a read-only tape. See [10] for more details.

2.3 Review of Approximate Matrix Multiplication

The BASICMATRIXMULTIPLICATION algorithm to approximate the product of two matrices is presented and analyzed in [10]. When this algorithm is given as input two matrices, $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{n \times p}$, a probability distribution $\{p_i\}_{i=1}^n$, and a number $c \leq n$, it returns as output two matrices, C and R , such that $CR \approx AB$; $C \in \mathbb{R}^{m \times c}$ is a matrix whose columns are c randomly chosen columns of A (suitably rescaled) and $R \in \mathbb{R}^{c \times p}$ is a matrix whose rows are the c corresponding rows of B (also suitably rescaled). An important aspect of this algorithm is the probability distribution $\{p_i\}_{i=1}^n$ used to choose column-row pairs. Although one could always use a uniform distribution, superior results are obtained if the probabilities are chosen judiciously. In particular, a set of sampling probabilities $\{p_i\}_{i=1}^n$ are the *optimal probabilities* (with respect to approximating the product AB) if they are of the form (8); for an explanation and discussion, see [10]. In [10] we prove the following.

Theorem 1 *Suppose $A \in \mathbb{R}^{m \times n}$, $B \in \mathbb{R}^{n \times p}$, $c \in \mathbb{Z}^+$ such that $1 \leq c \leq n$, and $\{p_i\}_{i=1}^n$ are such that $p_i \geq 0$ and $\sum_{i=1}^n p_i = 1$. Construct C and R with the BASICMATRIXMULTIPLICATION algorithm of [10] and let CR be an approximation to AB . If the probabilities $\{p_i\}_{i=1}^n$ are such that*

$$p_k = \frac{|A^{(k)}| |B_{(k)}|}{\sum_{k'=1}^n |A^{(k')}| |B_{(k')}|}, \quad (8)$$

then

$$\mathbf{E} [\|AB - CR\|_F] \leq \frac{1}{\sqrt{c}} \|A\|_F \|B\|_F. \quad (9)$$

If, in addition, we let $\delta \in (0, 1)$ and $\eta = 1 + \sqrt{8 \log(1/\delta)}$, then with probability at least $1 - \delta$

$$\|AB - CR\|_F \leq \frac{\eta}{\sqrt{c}} \|A\|_F \|B\|_F. \quad (10)$$

Furthermore, if the probabilities $\{p_i\}_{i=1}^n$ are such that

$$p_k = \frac{|B_{(k)}|^2}{\|B\|_F^2}, \quad (11)$$

then

$$\mathbf{E} [\|AB - CR\|_F] \leq \frac{1}{\sqrt{c}} \|A\|_F \|B\|_F. \quad (12)$$

Note that with probabilities of the form (11), we do not get a bound of the form $\|AB - CR\|_F \leq \frac{1}{\sqrt{c}} \|A\|_F \|B\|_F$ by sampling $O(\log(1/\delta))$ columns without making additional, awkward assumptions on the input matrices; see [10]. Of course, by Markov's inequality we can (and will) obtain such a bound by sampling $O(1/\delta)$ columns.

2.4 Review of Approximate Singular Value Decomposition

The LINEARTIMESVD algorithm is presented in [11]. It is an algorithm which, when given a matrix $A \in \mathbb{R}^{m \times n}$, uses $O(m + n)$ additional space and time to compute an approximation to the top k singular values and the corresponding left singular vectors of A by randomly choosing c columns of A and rescaling each appropriately to construct a matrix $C \in \mathbb{R}^{m \times c}$, computing the top k singular values and corresponding right singular vectors of C , and using them to construct a matrix $H_k \in \mathbb{R}^{m \times k}$ consisting of approximations to the top k left singular vectors of A . In [11] we prove the following.

Theorem 2 *Suppose $A \in \mathbb{R}^{m \times n}$ and let H_k be constructed from the LINEARTIMESVD algorithm of [11]. Then,*

$$\|A - H_k H_k^T A\|_F^2 \leq \|A - A_k\|_F^2 + 2\sqrt{k} \|AA^T - CC^T\|_F \quad (13)$$

$$\|A - H_k H_k^T A\|_2^2 \leq \|A - A_k\|_2^2 + 2 \|AA^T - CC^T\|_2. \quad (14)$$

The CONSTANTTIMESVD algorithm is also presented in [11]. It is an algorithm which, when given a matrix $A \in \mathbb{R}^{m \times n}$, uses constant additional space and time to compute a description of an approximation to the top k singular values and the corresponding left singular vectors of A . It does so in a manner similar to that of the LINEARTIMESVD algorithm except that it performs a second level of sampling in order to estimate (rather than compute exactly) the top k singular values and corresponding singular vectors of C . The γ in the following theorem is a parameter of the CONSTANTTIMESVD algorithm of [11] that is related to the second level of sampling; it also appears in our CONSTANTTIMECUR algorithm, and is thus discussed in Section 4. In [11] we prove the following.

Theorem 3 *Suppose $A \in \mathbb{R}^{m \times n}$, let a description of \tilde{H}_ℓ be constructed from the CONSTANTTIMESVD algorithm of [11] by sampling c columns of A with probabilities $\{p_i\}_{i=1}^n$ and w rows of C with probabilities $\{q_j\}_{j=1}^m$ where $p_i = |A^{(i)}|^2 / \|A\|_F^2$ and $q_j = |C_{(j)}|^2 / \|C\|_F^2$. Let $\eta = 1 + \sqrt{8 \log(2/\delta)}$ and $\epsilon > 0$.*

If a Frobenius norm bound is desired, and hence the CONSTANTTIMESVD algorithm is run with $\gamma = \epsilon/100k$, then by choosing $c = \Omega(k^2 \eta^2 / \epsilon^4)$ columns of A and $w = \Omega(k^2 \eta^2 / \epsilon^4)$ rows of C we have that with probability at least $1 - \delta$,

$$\|A - \tilde{H}_\ell \tilde{H}_\ell^T A\|_F^2 \leq \|A - A_k\|_F^2 + \epsilon \|A\|_F^2. \quad (15)$$

If a spectral norm bound is desired, and hence the CONSTANTTIMESVD algorithm is run with $\gamma = \epsilon/100$, then by choosing $c = \Omega(\eta^2 / \epsilon^4)$ columns of A and $w = \Omega(\eta^2 / \epsilon^4)$ rows of C we have that with probability at least $1 - \delta$,

$$\|A - \tilde{H}_\ell \tilde{H}_\ell^T A\|_2^2 \leq \|A - A_k\|_2^2 + \epsilon \|A\|_F^2. \quad (16)$$

3 The Linear Time CUR Decomposition

In this section we describe and analyze the LINEARTIMECUR algorithm which computes an approximate CUR decomposition of a matrix $A \in \mathbb{R}^{m \times n}$ using linear (in m and n) additional space and time. In Section 4 we describe and analyze the CONSTANTTIMECUR algorithm which computes a description of an approximate CUR decomposition of a matrix A using only constant

LINEARTIMECUR Algorithm

Input: $A \in \mathbb{R}^{m \times n}$, $r, c, k \in \mathbb{Z}^+$ s.t. $1 \leq r \leq m$, $1 \leq c \leq n$, and $1 \leq k \leq \min(r, c)$, $\{p_i\}_{i=1}^m$ s.t. $p_i \geq 0$ and $\sum_{i=1}^m p_i = 1$, and $\{q_j\}_{j=1}^n$ s.t. $q_j \geq 0$ and $\sum_{j=1}^n q_j = 1$.

Output: $C \in \mathbb{R}^{m \times c}$, $U \in \mathbb{R}^{c \times r}$, and $R \in \mathbb{R}^{r \times n}$.

- For $t = 1$ to c ,
 - Pick $j_t \in \{1, \dots, n\}$ with $\Pr[j_t = \alpha] = q_\alpha$, $\alpha = 1, \dots, n$.
 - Set $C^{(t)} = A^{(j_t)} / \sqrt{c q_{j_t}}$.
- Compute $C^T C$ and its singular value decomposition; say $C^T C = \sum_{t=1}^c \sigma_t^2(C) y^t y^{tT}$.
- If $\sigma_k(C) = 0$ then let $k = \max\{k' : \sigma_{k'}(C) \neq 0\}$.
- For $t = 1$ to r ,
 - Pick $i_t \in \{1, \dots, m\}$ with $\Pr[i_t = \alpha] = p_\alpha$, $\alpha = 1, \dots, m$.
 - Set $R_{(t)} = A_{(i_t)} / \sqrt{r p_{i_t}}$.
 - Set $\Psi_{(t)} = C_{(i_t)} / \sqrt{r p_{i_t}}$.
- Let $\Phi = \sum_{t=1}^k \frac{1}{\sigma_t^2(C)} y^t y^{tT}$ and let $U = \Phi \Psi^T$.
- Return C , U , and R .

Figure 2: The LINEARTIMECUR Algorithm

additional space and time. Both algorithms will make extensive use of the corresponding results from [11] for approximating the SVD of a matrix as well as results from [10] on approximating the product of two matrices. As with the SVD algorithms, the CONSTANTTIMECUR algorithm has a similar flavor to the LINEARTIMECUR algorithm, but is technically more complex due to the second level of sampling required. Thus, in this section we provide extensive discussion of the LINEARTIMECUR algorithm and the motivation and intuition behind it, and in Section 4 we highlight the differences between the linear additional time framework and the constant additional time framework.

3.1 The Algorithm

Given a matrix $A \in \mathbb{R}^{m \times n}$, we wish to compute a succinctly-described, easily-computed matrix A' that is decomposable as $A' = CUR \in \mathbb{R}^{m \times n}$ and that satisfies (1)–(5) of Section 1. The LINEARTIMECUR algorithm, which is presented in Figure 2, accomplishes this by first forming a matrix $C \in \mathbb{R}^{m \times c}$ by rescaling a randomly chosen subset of c columns of A ; the columns are chosen in c independent identical trials where in each trial the α -th column of A is chosen with probability q_α and if the α -th column is chosen it is rescaled by $1/\sqrt{c q_\alpha}$ before inclusion in C . The algorithm then forms a matrix $R \in \mathbb{R}^{r \times n}$ by rescaling a randomly chosen subset of r rows of A ; the rows are chosen in r independent identical trials where in each trial the α -th row of A is chosen with the probability p_α and if the α -th row is chosen it is rescaled by $1/\sqrt{r p_\alpha}$ before inclusion in R . Using the same randomly chosen rows to construct R from A the algorithm also constructs a




Figure 3: Diagram for the LINEAR TIME CUR Algorithm

matrix Ψ from C in an identical manner. Thus, $\Psi \in \mathbb{R}^{r \times c}$ and $\Psi_{ij} = A_{i_{t_1} j_{t_2}} / \sqrt{c r p_{i_{t_1}} q_{j_{t_2}}}$, where i_{t_1} is the element of $\{1, \dots, m\}$ selected in the t_1 -th row sampling trial and j_{t_2} is the element of $\{1, \dots, n\}$ selected in the t_2 -th column sampling trial.

The following sampling matrix formalism provides a convenient representation of our ideas, and will be used extensively in this section and the next. (See [10] for another use of this sampling matrix formalism.) Let us define the column sampling matrix $S_C \in \mathbb{R}^{n \times c}$ to be the zero-one matrix where $(S_C)_{ij} = 1$ if the i -th column of A is chosen in the j -th independent random trial and $S_{ij} = 0$ otherwise; let us also define the associated rescaling matrix $D_C \in \mathbb{R}^{c \times c}$ to be the diagonal matrix with $(D_C)_{tt} = 1/\sqrt{c p_{i_t}}$, where i_t is the element of $\{1, \dots, n\}$ chosen in the t -th sampling trial. Let us similarly define $S_R \in \mathbb{R}^{r \times m}$ and $D_R \in \mathbb{R}^{r \times r}$ to be the row sampling matrix and associated diagonal rescaling matrix, respectively. In this notation,

$$C = A S_C D_C \text{ and } R = D_R S_R A, \quad (17)$$

where $S_C D_C$ postmultiplies (and thus samples and rescales columns of) A to form C , and where $D_R S_R$ premultiplies (and thus samples and rescales rows of) A to form R . Thus, in this notation,

$$\Psi = D_R S_R C = D_R S_R A S_C D_C. \quad (18)$$

Given C , the LINEAR TIME CUR algorithm computes the top k singular values, $\sigma_t^2(C)$, $t = 1, \dots, k$, and the corresponding singular vectors, y^t , $t = 1, \dots, k$, of $C^T C$. Note that these are also the (squares of the) singular values and the corresponding right singular vectors of C . Using these quantities, a matrix $\Phi \in \mathbb{R}^{c \times c}$ may be defined as

$$\Phi = \sum_{t=1}^k \frac{1}{\sigma_t^2(C)} y^t y^{t^T}, \quad (19)$$

from which $U \in \mathbb{R}^{c \times r}$ is constructed as $U = \Phi \Psi^T$. We could, of course, have defined a matrix Φ (and thus constructed a matrix U) from the singular vectors and singular values of $R R^T$ in a manner analogous to that described above; in that case, the roles of the row sampling and column sampling would be reversed relative to the discussion below.

Figure 3 presents a diagram illustrating the action of the LINEAR TIME CUR algorithm. The matrix A is shown as operating between the high dimensional spaces \mathbb{R}^n and \mathbb{R}^m . In addition, the matrix C is shown as operating between \mathbb{R}^c and \mathbb{R}^m and the matrix R is shown as operating between \mathbb{R}^n and \mathbb{R}^r . Intuitively, one may think of \mathbb{R}^c and \mathbb{R}^r as being the most significant parts of \mathbb{R}^n and \mathbb{R}^m , respectively, in terms of the action of A . Indeed, this will be the case when the sampling probabilities $\{p_i\}_{i=1}^m$ and $\{q_j\}_{j=1}^n$ satisfy certain conditions, as stated, e.g., in Theorem 4. The diagram also illustrates that $C = A S_C D_C$, $R = D_R S_R A$, and the matrix U which can be seen to be $U = \Phi \Psi^T = \Phi C^T (D_R S_R)^T$.




Figure 4: Another diagram for the LINEARTIMECUR Algorithm

and then we can bound the two terms separately. The first term in (23) can be bounded using the SVD results of [11] if the column sampling probabilities satisfy certain conditions; in particular we will require that they are the optimal probabilities. Since $H_k^T H_k = I_k$ Lemma 2 states that

$$\begin{aligned} \|H_k H_k^T A - CUR\|_F &= \|H_k^T A - H_k^T (D_R S_R)^T D_R S_R A\|_F \\ &= \|H_k^T A - \widetilde{H}_k^T \widetilde{A}\|_F. \end{aligned} \quad (24)$$

Thus, the second term in (23) can be bounded by the matrix multiplication results of [10]; it will follow that if the sampling probabilities $\{p_i\}_{i=1}^m$ satisfy certain conditions then $\widetilde{H}_k^T \widetilde{A} \approx H_k^T A$ in the sense that the error in $\|H_k^T A - \widetilde{H}_k^T \widetilde{A}\|_F$ can be bounded. Note that since the optimal probabilities depend on both H^T and A , and since we do not have access to H^T , our probabilities will not be optimal; nevertheless, although we will not obtain bounds with very high probability, as in [10] and [11], we will be able to apply Markov's inequality and thus achieve the bounds we desire.

A diagram illustrating the method (just described) that will be used to prove the correctness of the CUR algorithm is presented in Figure 4. In this figure, the locations of \mathbb{R}^c and \mathbb{R}^r and thus the directions of R , S_R , C , S_C , and U have been switched relative to their location in Figure 3. This presentation has several advantages: first, the SVD of C and the SVD of A can both be presented in the same figure as the CUR decomposition of A ; second, one can see that bounding $\|A - H_k H_k^T A\|_\xi$ well in terms of $\|AA^T - CC^T\|_\xi$ depends on the probabilities used to sample the columns of A ; and third, one can also see that bounding $\|H_k^T A - \widetilde{H}_k^T \widetilde{A}\|_F$ well depends on the probabilities used to sample the (columns of H_k^T and the corresponding) rows of A . See the corresponding figures in [10] and [11] for a comparison.

3.2 Analysis of the Implementation and Running Time

In the LINEARTIMECUR algorithm the sampling probabilities $\{p_i\}_{i=1}^m$ and $\{q_j\}_{j=1}^n$ (if they are chosen to be of the form used in Theorem 4 and Theorem 5) can be computed in one pass and $O(c+r)$ additional space and time using the SELECT algorithm of [10]. Given the elements to be sampled, the matrix C can then be constructed in one additional pass; this requires additional space and time that is $O(mc)$. Similarly, the matrix R can then be constructed in the same pass

using additional space and time that is $O(nr)$. Given $C \in \mathbb{R}^{m \times c}$ computing $C^T C$ requires $O(mc^2)$ additional space and time and computing the SVD of $C^T C$ requires $O(c^3)$ additional space and time. The matrix Ψ can be computed in the same second pass by sampling the same r rows of C that were used to construct R from A ; this requires additional space and time that is $O(cr)$. The matrix Φ can be explicitly computed using $O(c^2 k)$ additional space and time and then the matrix $U = \Phi \Psi^T$ can be computed using $O(c^2 r)$ additional space and time. Thus, since c , r , and k are assumed to be a constant, overall $O(m+n)$ additional space and time are required by the LINEARTIMECUR algorithm, and (1)–(3) of Section 1 are satisfied. Note that the “description” of the solution that is computable in the allotted additional space and time is the explicit matrices C , U , and R .

3.3 Analysis of the Sampling Step

Before stating and proving the main theorem of this section, we will first prove two useful lemmas. Lemma 1 will establish (22) and Lemma 2 will establish (24).

Lemma 1

$$CUR = H_k \widetilde{H}_k^T \widetilde{A}$$

Proof: Note that the singular value decomposition of C is $C = \sum_{t=1}^c \sigma_t(C) h^t y^{tT}$, that the matrix $\Psi = D_R S_R C$, and that $U = \Phi \Psi^T$, where Φ is given by (19). Thus, we have that

$$\begin{aligned} CUR &= C \left(\sum_{t=1}^k \frac{1}{\sigma_t^2(C)} y^t y^{tT} \right) C^T (D_R S_R)^T R \\ &= \left(\sum_{t_1=1}^k \sigma_{t_1}(C) h^{t_1} y^{t_1T} \right) \left(\sum_{t_2=1}^k \frac{1}{\sigma_{t_2}^2(C)} y^{t_2} y^{t_2T} \right) \left(\sum_{t_3=1}^k \sigma_{t_3}(C) y^{t_3} h^{t_3T} \right) (D_R S_R)^T R \\ &= \left(\sum_{t=1}^k h^t h^{tT} \right) (D_R S_R)^T R. \end{aligned}$$

The lemma follows since $\sum_{t=1}^k h^t h^{tT} = H_k H_k^T$, since $R = D_R S_R A$, and from the definitions (21). \diamond

Lemma 2

$$\|H_k H_k^T A - CUR\|_F = \|H_k^T A - \widetilde{H}_k^T \widetilde{A}\|_F$$

Proof: From Lemma 1 we have that $CUR = H_k \widetilde{H}_k^T \widetilde{A}$. Let us define the matrix $\Omega \in \mathbb{R}^{k \times n}$ as

$$\Omega = H_k^T A - H_k^T (D_R S_R)^T D_R S_R A = H_k^T A - \widetilde{H}_k^T \widetilde{A}.$$

In addition, note that

$$\|H_k H_k^T A - H_k \widetilde{H}_k^T \widetilde{A}\|_F^2 = \|H_k \Omega\|_F^2 = \mathbf{Tr}(\Omega^T H_k^T H_k \Omega).$$

The lemma follows since $H_k^T H_k = I_k$ and since $\mathbf{Tr}(\Omega^T \Omega) = \|\Omega\|_F^2$. \diamond

Here is our main theorem regarding the LINEARTIMECUR algorithm described in Section 3.1. Note that in this theorem we restrict ourselves to sampling probabilities that are optimal in the sense of Section 2.3.

Theorem 4 Suppose $A \in \mathbb{R}^{m \times n}$ and let C , U , and R be constructed from the LINEARTIMECUR algorithm by sampling c columns of A with probabilities $\{q_j\}_{j=1}^n$ and r rows of A with probabilities $\{p_i\}_{i=1}^m$. Assume that $p_i = |A_{(i)}|^2 / \|A\|_F^2$ and $q_j = |A^{(j)}|^2 / \|A\|_F^2$. Then

$$\mathbf{E} [\|A - CUR\|_F] \leq \|A - A_k\|_F + \left(\left(\frac{4k}{c} \right)^{1/4} + \left(\frac{k}{r} \right)^{1/2} \right) \|A\|_F \quad (25)$$

$$\mathbf{E} [\|A - CUR\|_2] \leq \|A - A_k\|_2 + \left(\left(\frac{4}{c} \right)^{1/4} + \left(\frac{k}{r} \right)^{1/2} \right) \|A\|_F. \quad (26)$$

In addition, if we let $\eta_c = 1 + \sqrt{8 \log(1/\delta_c)}$ and let $\delta = \delta_r + \delta_c$, then with probability at least $1 - \delta$

$$\|A - CUR\|_F \leq \|A - A_k\|_F + \left(\left(\frac{4k\eta_c^2}{c} \right)^{1/4} + \left(\frac{k}{\delta_r^2 r} \right)^{1/2} \right) \|A\|_F \quad (27)$$

$$\|A - CUR\|_2 \leq \|A - A_k\|_2 + \left(\left(\frac{4\eta_c^2}{c} \right)^{1/4} + \left(\frac{k}{\delta_r^2 r} \right)^{1/2} \right) \|A\|_F. \quad (28)$$

Proof: By the submultiplicativity of $\|\cdot\|_F$ and $\|\cdot\|_2$ we have that

$$\|A - CUR\|_\xi \leq \|A - H_k H_k^T A\|_\xi + \|H_k H_k^T A - CUR\|_\xi \quad (29)$$

for both $\xi = 2, F$; thus by Lemma 1 and Lemma 2 we have that

$$\|A - CUR\|_\xi \leq \|A - H_k H_k^T A\|_\xi + \|H_k^T A - H_k^T (D_R S_R)^T D_R S_R A\|_F \quad (30)$$

$$= \|A - H_k H_k^T A\|_\xi + \|H_k^T A - \widetilde{H}_k^T \widetilde{A}\|_F \quad (31)$$

for both $\xi = 2, F$, where (31) follows from the definitions (21). Then, note that the column sampling satisfies the requirements for the LINEARTIMESVD algorithm of [11]. Thus, by Theorem 2 it follows from (30) that

$$\|A - CUR\|_F \leq \|A - A_k\|_F + (4k)^{1/4} \|AA^T - CC^T\|_F^{1/2} + \|H_k^T A - \widetilde{H}_k^T \widetilde{A}\|_F \quad (32)$$

$$\|A - CUR\|_2 \leq \|A - A_k\|_2 + \sqrt{2} \|AA^T - CC^T\|_F^{1/2} + \|H_k^T A - \widetilde{H}_k^T \widetilde{A}\|_F. \quad (33)$$

Note that from the LINEARTIMECUR algorithm the column sampling probabilities are of the form (8) with $B = A^T$; thus, they are optimal and $\mathbf{E} [\|AA^T - CC^T\|_F] \leq \frac{1}{\sqrt{c}} \|A\|_F^2$. In addition, although the row sampling probabilities are not optimal, they are of the form (11); thus, since $\|H_k^T\|_F = \sqrt{k}$ we have that

$$\mathbf{E} \left[\|H_k^T A - \widetilde{H}_k^T \widetilde{A}\|_F \right] \leq \sqrt{\frac{k}{r}} \|A\|_F. \quad (34)$$

Thus, by taking expectations of (32) and (33), by using Jensen's inequality and Theorem 1, (25) and (26) follow.

To establish (27) and (28) first let the events $\mathcal{E}_\xi, \xi = c, r$ be defined as follows:

$$\begin{aligned} \mathcal{E}_c : \quad & \|AA^T - CC^T\|_F \leq \frac{\eta_c}{\sqrt{c}} \|A\|_F^2 \\ \mathcal{E}_r : \quad & \|H_k^T A - \widetilde{H}_k^T \widetilde{A}\|_F \leq \frac{1}{\delta_r} \sqrt{\frac{k}{r}} \|A\|_F. \end{aligned}$$

Thus, from Theorem 1 we have that $\Pr[\mathcal{E}_c] \geq 1 - \delta_c$. By applying Markov's inequality to $\left\|H_k^T A - \widetilde{H}_k^T \widetilde{A}\right\|_F$ and using (34) we see that

$$\Pr \left[\left\|H_k^T A - \widetilde{H}_k^T \widetilde{A}\right\|_F \geq \frac{1}{\delta_r} \sqrt{\frac{k}{r}} \|A\|_F \right] \leq \delta_r$$

and thus that $\Pr[\mathcal{E}_r] \geq 1 - \delta_r$. The theorem then follows from (32) and (33) by considering the event $\mathcal{E}_c \cap \mathcal{E}_r$. \diamond

Note that in the proof of Theorem 4 (and similarly in that of Theorem 6 in Section 4) $\|A - CUR\|_\xi$ is bounded by bounding each of the terms $\|A - H_k H_k^T A\|_\xi$ and $\|H_k H_k^T A - CUR\|_\xi$ independently. In order to bound $\|A - H_k H_k^T A\|_\xi$, the SVD results for arbitrary probabilities from [11] are used, and then it is noted that the column sampling probabilities used in the LINEARTIMECUR algorithm are optimal for bounding $\|AA^T - CC^T\|_F$. Then, independently, $\|H_k H_k^T A - CUR\|_\xi$ is bounded by using the matrix multiplication results of [10]. Since the probabilities that are used for the row sampling are not optimal with respect to bounding $\left\|H_k^T A - \widetilde{H}_k^T \widetilde{A}\right\|_F$, we do not obtain that bound with high probability. Due to the use of Markov's inequality, we must sample a number of rows that is $O(1/\delta)$ whereas we only need to sample a number of columns that is $O(\log(1/\delta))$.

As a corollary of Theorem 4 we have the following theorem. In this theorem, in addition to using sampling probabilities that are optimal in the sense of Section 2.3, we choose sufficiently many columns and rows to ensure that the additional error is less than $\epsilon' \|A\|_F$.

Theorem 5 *Suppose $A \in \mathbb{R}^{m \times n}$ and let C , U , and R be constructed from the LINEARTIMECUR algorithm by sampling c columns of A with probabilities $\{q_j\}_{j=1}^n$ and r rows of A with probabilities $\{p_i\}_{i=1}^m$. Assume that $p_i = |A_{(i)}|^2 / \|A\|_F^2$ and $q_j = |A^{(j)}|^2 / \|A\|_F^2$ and let $\epsilon, \epsilon' > 0$ with $\epsilon = \epsilon'/2$. If $c \geq 4k/\epsilon^4$ and $r \geq k/\epsilon^2$, then*

$$\mathbf{E} [\|A - CUR\|_F] \leq \|A - A_k\|_F + \epsilon' \|A\|_F, \quad (35)$$

and if $c \geq 4/\epsilon^4$ and $r \geq k/\epsilon^2$, then

$$\mathbf{E} [\|A - CUR\|_2] \leq \|A - A_k\|_2 + \epsilon' \|A\|_F. \quad (36)$$

In addition, if we let $\eta_c = 1 + \sqrt{8 \log(1/\delta_c)}$ and let $\delta = \delta_r + \delta_c$, and if $c \geq 4k\eta_c^2/\epsilon^4$ and $r \geq k/\delta_r^2\epsilon^2$, then with probability at least $1 - \delta$

$$\|A - CUR\|_F \leq \|A - A_k\|_F + \epsilon' \|A\|_F, \quad (37)$$

and if $c \geq 4\eta_c^2/\epsilon^4$ and $r \geq k/\delta_r^2\epsilon^2$, then with probability at least $1 - \delta$

$$\|A - CUR\|_2 \leq \|A - A_k\|_2 + \epsilon' \|A\|_F. \quad (38)$$

The results of Theorem 4 and Theorem 5 for both the Frobenius norm and the spectral norm hold for all k and are of particular interest when A is well approximated by a matrix of low rank since then one may choose $k = O(1)$ and obtain a good approximation. In addition, since $\|A - A_t\|_2 \leq \|A\|_F / \sqrt{t}$, for all $t = 1, 2, \dots, r$, the bounds with respect to the spectral norm have the following interesting property: from (36) we can see that

$$\mathbf{E} [\|A - CUR\|_2] \leq \left(1/\sqrt{k} + \epsilon'\right) \|A\|_F,$$

and similarly for (38). Thus, under the assumptions of Theorem 5 if we choose $k = 1/\epsilon'^2$ and let $\epsilon'' = 2\epsilon'$ then we have that

$$\mathbf{E} [\|A - CUR\|_2] \leq \epsilon'' \|A\|_F, \quad (39)$$

and that

$$\|A - CUR\|_2 \leq \epsilon'' \|A\|_F \quad (40)$$

holds with probability at least $1 - \delta$.

4 The Constant Time $C\tilde{U}R$ Decomposition

4.1 The Algorithm

The CONSTANTTIMECUR algorithm is very similar in spirit to the LINEARTIMECUR algorithm; thus, we only highlight its main features with an emphasis on similarities and differences between the two algorithms. Given a matrix $A \in \mathbb{R}^{m \times n}$, we wish to compute a succinctly-described, easily-computed matrix A' that is decomposable as $A' = C\tilde{U}R \in \mathbb{R}^{m \times n}$ and that satisfies (1)–(5) of Section 1, where the additional RAM space and time to compute \tilde{U} is $O(1)$. The CONSTANTTIMECUR algorithm, which is presented in Figure 5, accomplishes this by forming a matrix $C \in \mathbb{R}^{m \times c}$ by rescaling a randomly chosen subset of c columns of A , forming a matrix $R \in \mathbb{R}^{r \times n}$ by rescaling a randomly chosen subset of r rows of A , and forming a matrix $\Psi \in \mathbb{R}^{r \times c}$ from C by choosing the same randomly chosen rows used to construct R from A and rescaling appropriately. Given C , the CONSTANTTIMECUR algorithm randomly chooses and rescales w rows of C to form a matrix $W \in \mathbb{R}^{w \times c}$ and then computes the top ℓ singular values, $\sigma_t^2(W)$, $t = 1, \dots, \ell$, and the corresponding singular vectors, z^t , $t = 1, \dots, \ell$, of $W^T W$. Note that these are also approximations to the (squares of the) singular values and the corresponding right singular vectors of C . Using these quantities, a matrix $\tilde{\Phi} \in \mathbb{R}^{c \times c}$ may be defined as

$$\tilde{\Phi} = \sum_{t=1}^{\ell} \frac{1}{\sigma_t^2(W)} z^t z^{tT}, \quad (41)$$

from which $\tilde{U} \in \mathbb{R}^{c \times r}$ is constructed as $\tilde{U} = \tilde{\Phi} \Psi^T$. Note that in the constant additional space and time framework the actual matrices C and R are not explicitly computed; instead the constant-sized matrix \tilde{U} is computed and only a constant number of bits are stored to specify which columns and rows of A are kept (along with their associated rescaling factors) in the construction of C and R , respectively. Thus, the length of the succinct representation of A is a constant.

Figure 3 of Section 3 provides a diagram illustrating the action of LINEARTIMECUR algorithm, but the diagram and associated discussion are also relevant for the CONSTANTTIMECUR algorithm. Figure 6, in this section, also provides a diagram illustrating the action of the CONSTANTTIMECUR algorithm, and is the analogue for the constant time $C\tilde{U}R$ of Figure 4; Figure 6 also illustrates that the matrix $Y (= V_C)$ (of Figure 4) consisting of the top k right singular vectors of C is not exactly computed, but is instead approximated by the matrix $Z (= V_W)$, where Z is a matrix whose columns $Z^{(t)} = z^t$ consist of the right singular vectors of W . Thus, the matrix H_k consisting of the left singular vectors of C is not exactly computed but is only approximated by \tilde{H}_ℓ , where $\tilde{H}_\ell^{(t)} = \tilde{h}^t = Cz^t/\sigma_t(W)$, for $t = 1, \dots, \ell$. Since by construction it is still the case that $R = D_R S_R A$ (and also that $C = A S_C D_C$ and $\Psi = D_R S_R C$), where the sampling matrices and the diagonal rescaling matrices are defined as in Section 3.1, it follows from Lemma 3 that

$$C\tilde{U}R = \tilde{H}_\ell \tilde{H}_\ell^T (D_R S_R)^T D_R S_R A. \quad (42)$$

CONSTANTTIMECUR Algorithm

Input: $A \in \mathbb{R}^{m \times n}$, $r, c, k \in \mathbb{Z}^+$ s.t. $1 \leq r \leq m$, $1 \leq c \leq n$, and $1 \leq k \leq \min(r, c)$, $\{p_i\}_{i=1}^m$ s.t. $p_i \geq 0$ and $\sum_{i=1}^m p_i = 1$, and $\{q_j\}_{j=1}^n$ s.t. $q_j \geq 0$ and $\sum_{j=1}^n q_j = 1$.

Output: $\tilde{U} \in \mathbb{R}^{c \times r}$ and a “description” of $C \in \mathbb{R}^{m \times c}$ and $R \in \mathbb{R}^{r \times n}$.

- For $t = 1$ to c ,
 - Pick $j_t \in \{1, \dots, n\}$ with $\Pr[j_t = \alpha] = q_\alpha$, and save $\{(j_t, q_{j_t}) : t = 1, \dots, c\}$.
 - Set $C^{(t)} = A^{(j_t)} / \sqrt{c q_{j_t}}$. (Note that C is not explicitly constructed in RAM.)
- Choose $\{\pi_i\}_{i=1}^m$ s.t. $\pi_i = |C_{(i)}|^2 / \|C\|_F^2$.
- For $t = 1$ to w ,
 - Pick $i_t \in 1, \dots, m$ with $\Pr[i_t = \alpha] = \pi_\alpha$, $\alpha = 1, \dots, m$.
 - Set $W_{(t)} = C_{(i_t)} / \sqrt{w \pi_{i_t}}$.
- Compute $W^T W$ and its singular value decomposition; say $W^T W = \sum_{t=1}^c \sigma_t^2(W) z^t z^{tT}$.
- If a $\|\cdot\|_F$ bound is desired, set $\gamma = \epsilon/100k$,
Else if a $\|\cdot\|_2$ bound is desired, set $\gamma = \epsilon/100$.
- Let $\ell = \min\{k, \max\{t : \sigma_t^2(W) \geq \gamma \|W\|_F^2\}\}$.
- Keep singular values $\{\sigma_t(W)\}_{t=1}^\ell$ and their corresponding singular vectors $\{z^t\}_{t=1}^\ell$.
- For $t = 1$ to r ,
 - Pick $i_t \in \{1, \dots, m\}$ with $\Pr[i_t = \alpha] = p_\alpha$, and save $\{(i_t, p_{i_t}) : t = 1, \dots, r\}$.
 - Set $R_{(t)} = A_{(i_t)} / \sqrt{r p_{i_t}}$. (Note that R is not explicitly constructed in RAM.)
 - Set $\Psi_{(t)} = C_{(i_t)} / \sqrt{r p_{i_t}}$.
- Let $\tilde{\Phi} = \sum_{t=1}^\ell \frac{1}{\sigma_t^2(W)} z^t z^{tT}$ and let $\tilde{U} = \tilde{\Phi} \Psi^T$.
- Return \tilde{U} , c column labels $\{(j_t, q_{j_t}) : t = 1, \dots, c\}$, and r row labels $\{(i_t, p_{i_t}) : t = 1, \dots, r\}$.

Figure 5: The CONSTANTTIMECUR Algorithm




Figure 6: Diagram for the CONSTANTTIMECUR Algorithm

In the linear time case we had that $H_k^T H_k = I_k$ since the columns of H_k were k of the left singular vectors of C . This allowed us to prove Lemma 2, as described in Section 3.1. In the constant additional time setting, the columns of \tilde{H}_ℓ are only approximations to k of the left singular vectors of C , but we still have that $\tilde{H}_\ell^T \tilde{H}_\ell \approx I_\ell$. To quantify this, define $Z_{\alpha,\beta} \in \mathbb{R}^{c \times (\beta - \alpha + 1)}$ to be the matrix whose columns are the α -th through the β -th singular vectors of $W^T W$ and $T \in \mathbb{R}^{\ell \times \ell}$ to be the diagonal matrix with elements $T_{tt} = 1/\sigma_t(W)$. If we define the matrix $\Delta \in \mathbb{R}^{\ell \times \ell}$ to be

$$\Delta = T Z_{1,\ell}^T (C^T C - W^T W) Z_{1,\ell} T, \quad (43)$$

then Lemma 5 will establish that

$$\left\| \tilde{H}_\ell \tilde{H}_\ell^T A - C \tilde{U} R \right\|_F \leq \left(1 + \|\Delta\|_F^{1/2} \right) \left\| \tilde{H}_\ell^T A - \tilde{H}_\ell^T (D_R S_R)^T D_R S_R A \right\|_F, \quad (44)$$

and Lemma 4 establishes that $\tilde{H}_\ell^T \tilde{H}_\ell = I_\ell + \Delta$. Thus, as in the linear time case, we can split

$$\left\| A - C \tilde{U} R \right\|_\xi \leq \left\| A - \tilde{H}_\ell \tilde{H}_\ell^T A \right\|_\xi + \left\| \tilde{H}_\ell \tilde{H}_\ell^T A - C \tilde{U} R \right\|_\xi$$

for $\xi = 2, F$ and then bound the two terms separately. The first term can be bounded by the constant time SVD results of [11], the second term can be bounded by the matrix multiplication results of [10], and bounding the overall error depends on both results. Due to the two levels of sampling, the additional error will be larger than in the linear additional time framework, but it can be made arbitrarily small by choosing a constant number of rows and columns.

4.2 Analysis of the Implementation and Running Time

In the CONSTANTTIMECUR algorithm the sampling probabilities $\{p_i\}_{i=1}^m$ and $\{q_j\}_{j=1}^n$ (if they are chosen to be of the form used in Theorem 6) can be computed in one pass and $O(c+r)$ additional space and time using the SELECT algorithm of [10]. Given the columns of A to be sampled,

we do not explicitly construct the matrix C but instead perform a second level of sampling and select w rows of C with probabilities $\{\pi_i\}_{i=1}^m$ in order to construct the matrix W ; this requires a second pass and $O(w)$ additional space and time. Then, in a third pass we explicitly construct W ; this requires additional space and time that is $O(cw)$. Similarly, a description of the matrix R can then be constructed in the same third pass using additional space and time that is $O(r)$. Then, given W , computing $W^T W$ requires $O(cw^2)$ additional space and time and computing the SVD of $W^T W$ requires $O(c^3)$ additional space and time. The matrix Ψ can be computed in the same third pass by sampling the same r rows of C that were used to construct R from A ; this requires additional space and time that is $O(cr)$. The matrix $\tilde{\Phi}$ can be explicitly computed using $O(c^2 k)$ additional space and time and then the matrix $U = \tilde{\Phi} \Psi^T$ can be computed using $O(c^2 r)$ additional space and time. Thus, since c , r , and k are assumed to be constants, overall $O(1)$ additional space and time are required by the CONSTANTTIMECUR algorithm, and (1)–(3) of Section 1 are satisfied. Note that the “description” of the solution that is computable in the allotted additional space and time is the matrix \tilde{U} and the labels i_1, \dots, i_r and j_1, \dots, j_c indicating the rows chosen to construct C and R as well as the corresponding probabilities $\{p_{i_t}\}_{t=1}^r$ and $\{q_{j_t}\}_{t=1}^c$; we note that we need to know p_i only for the sampled rows i and q_j only for the sampled columns j .

4.3 Analysis of the Sampling Step

Before stating and proving the main theorem of this section, we will first prove several useful lemmas. First, in Lemma 3 we will establish (42).

Lemma 3

$$C\tilde{U}R = \tilde{H}_\ell \tilde{H}_\ell^T (D_R S_R)^T D_R S_R A$$

Proof: Since $\tilde{h}^t = Cz^t/\sigma_t(W)$ for $t = 1, \dots, \ell$, we have that $C = \sum_{t=1}^{\ell} \sigma_t(W) \tilde{h}^t z^{tT}$. Thus, we have that

$$\begin{aligned} C\tilde{U}R &= C \left(\sum_{t=1}^{\ell} \frac{1}{\sigma_t^2(W)} z^t z^{tT} \right) C^T (D_R S_R)^T R \\ &= \left(\sum_{t=1}^{\ell} \tilde{h}^t \tilde{h}^{tT} \right) (D_R S_R)^T R. \end{aligned}$$

The lemma follows since $\sum_{t=1}^{\ell} \tilde{h}^t \tilde{h}^{tT} = \tilde{H}_\ell \tilde{H}_\ell^T$ and since $R = D_R S_R A$. ◊

Next, Lemma 4 will characterize, in terms of Δ , the degree to which the columns of \tilde{H}_ℓ are not orthonormal. Note that it appeared in [11].

Lemma 4 *When written in the basis with respect to Z :*

$$\tilde{H}_\ell^T \tilde{H}_\ell = I_\ell + \Delta.$$

Furthermore, for $\xi = 2, F$,

$$\|\Delta\|_\xi \leq \frac{1}{\gamma \|W\|_F^2} \|C^T C - W^T W\|_\xi.$$

Proof: Recall that $\tilde{h}^t = Cz^t/\sigma_t(W)$ is the t -th column of \tilde{H}_ℓ and that $|z^{tT}W^TWz^t| = \sigma_t^2(W)\delta_{tt}$, where $\delta_{tt'}$ is the Kronecker delta function. First note that

$$\tilde{h}^{tT}\tilde{h}^t = \frac{z^{tT}C^TCz^t}{\sigma_t^2(W)} = \frac{z^{tT}W^TWz^t}{\sigma_t^2(W)} + \frac{z^{tT}(C^TC - W^TW)z^t}{\sigma_t^2(W)} = 1 + \Delta_{tt}.$$

Next note that if $t \neq t'$ then

$$\tilde{h}^{tT}\tilde{h}^{t'} = \frac{z^{tT}C^TCz^{t'}}{\sigma_t(W)\sigma_{t'}(W)} = \frac{z^{tT}W^TWz^{t'}}{\sigma_t(W)\sigma_{t'}(W)} + \frac{z^{tT}(C^TC - W^TW)z^{t'}}{\sigma_t(W)\sigma_{t'}(W)} = \Delta_{tt'}$$

Finally, since $|\Delta_{tt'}| \leq \frac{1}{\gamma\|W\|_F^2} \left| \left(Z_{1,\ell}^T(C^TC - W^TW)Z_{1,\ell} \right)_{tt'} \right|$ for every t and t' , it follows that

$$\|\Delta\|_F \leq \frac{1}{\gamma\|W\|_F^2} \|C^TC - W^TW\|_F.$$

Similarly, by submultiplicativity,

$$\|\Delta\|_2 \leq \|T\|_2^2 \|Z_{1,\ell}\|_2^2 \|C^TC - W^TW\|_2 \leq \frac{1}{\gamma\|W\|_F^2} \|C^TC - W^TW\|_2.$$

The lemma then follows. ◇

Next, in Lemma 5 we will establish (44).

Lemma 5

$$\left\| \tilde{H}_\ell \tilde{H}_\ell^T A - C\tilde{U}R \right\|_F \leq \left(1 + \|\Delta\|_F^{1/2} \right) \left\| \tilde{H}_\ell^T A - \tilde{H}_\ell^T (D_R S_R)^T D_R S_R A \right\|_F$$

Proof: From Lemma 3 we have that $C\tilde{U}R = \tilde{H}_\ell \tilde{H}_\ell^T (D_R S_R)^T D_R S_R A$. Let us define the matrix $\Omega \in \mathbb{R}^{\ell \times n}$ as

$$\Omega = \tilde{H}_\ell^T A - \tilde{H}_\ell^T (D_R S_R)^T D_R S_R A.$$

Thus, since $\|XX^T\|_F = \|X\|_F^2$ for a matrix X , we have that

$$\begin{aligned} \left\| \tilde{H}_\ell \Omega \right\|_F^2 &= \mathbf{Tr} \left(\Omega^T \tilde{H}_\ell^T \tilde{H}_\ell \Omega \right) \\ &= \mathbf{Tr} \left(\Omega^T (I_\ell + \Delta) \Omega \right) \end{aligned} \tag{45}$$

$$\begin{aligned} &= \|\Omega\|_F^2 + \mathbf{Tr} \left(\Omega^T \Delta \Omega \right) \\ &\leq \|\Omega\|_F^2 + \|\Delta\|_2 \|\Omega\|_F^2, \end{aligned} \tag{46}$$

where (45) follows from Lemma 4 and (46) follows since $|\mathbf{Tr}(\Omega^T \Delta \Omega)| \leq \|\Delta\|_2 \mathbf{Tr}(\Omega^T \Omega)$. The lemma then follows. ◇

Finally, in Lemma 6 we show that $\|W\|_F = \|C\|_F = \|A\|_F$ when optimal probabilities are used. It also appeared in [11].

Lemma 6 *Suppose $A \in \mathbb{R}^{m \times n}$ and run the CONSTANTTIMECUR algorithm by sampling c columns of A with probabilities $\{q_j\}_{j=1}^n$ (and then sampling w rows of C with probabilities $\{\pi_i\}_{i=1}^m$ to construct W) and r rows of A with probabilities $\{p_i\}_{i=1}^m$. Assume that $p_i = |A_{(i)}|^2 / \|A\|_F^2$ and $q_j = |A^{(j)}|^2 / \|A\|_F^2$. Then, $\|W\|_F = \|C\|_F = \|A\|_F$.*

Proof: If $p_i = |A^{(i)}|^2 / \|A\|_F^2$ then we have that $\|C\|_F^2 = \sum_{t=1}^c |C^{(t)}|^2 = \sum_{t=1}^c \frac{|A^{(i_t)}|^2}{cp_{i_t}} = \|A\|_F^2$. Similarly, if $q_j = |C_{(j)}|^2 / \|C\|_F^2$ then we have that $\|W\|_F^2 = \sum_{t=1}^w |W_{(t)}|^2 = \sum_{t=1}^w \frac{|C_{(i_t)}|^2}{wq_{i_t}} = \|C\|_F^2$. The lemma follows. \diamond

Here is our main theorem regarding the CONSTANTTIMECUR algorithm described in Section 4.1. It is the constant time analogue of Theorem 5. Note that in this theorem we restrict ourselves to sampling probabilities that are optimal in the sense of Section 2.3, and to choosing sufficiently many columns and rows to ensure that the additional error is less than $\epsilon \|A\|_F$.

Theorem 6 *Suppose $A \in \mathbb{R}^{m \times n}$ and let C, \tilde{U} , and R be constructed from the CONSTANTTIMECUR algorithm by sampling c columns of A with probabilities $\{q_j\}_{j=1}^n$ (and then sampling w rows of C with probabilities $\{\pi_i\}_{i=1}^m$ to construct W) and r rows of A with probabilities $\{p_i\}_{i=1}^m$. Assume that $p_i = |A_{(i)}|^2 / \|A\|_F^2$ and $q_j = |A^{(j)}|^2 / \|A\|_F^2$. Let $\eta = 1 + \sqrt{8 \log(3/\delta)}$ and $\epsilon > 0$.*

If a Frobenius norm bound is desired, and hence the CONSTANTTIMESVD algorithm is run with $\gamma = \epsilon/100k$, then if we let $c = \Omega(k^2 \eta^2 / \epsilon^8)$, $w = \Omega(k^2 \eta^2 / \epsilon^8)$, and $r = \Omega(k/\delta^2 \epsilon^2)$, then with probability at least $1 - \delta$

$$\left\| A - C\tilde{U}R \right\|_F \leq \|A - A_k\|_F + \epsilon \|A\|_F. \quad (47)$$

If a spectral norm bound is desired, and hence the CONSTANTTIMESVD algorithm is run with $\gamma = \epsilon/100$, then if we let $c = \Omega(\eta^2 / \epsilon^8)$, $w = \Omega(\eta^2 / \epsilon^8)$, and $r = \Omega(k/\delta^2 \epsilon^2)$, then with probability at least $1 - \delta$

$$\left\| A - C\tilde{U}R \right\|_2 \leq \|A - A_k\|_2 + \epsilon \|A\|_F. \quad (48)$$

Proof: Let us define the events:

$$\mathcal{E}_c : \left\| AA^T - CC^T \right\|_F \leq \frac{\eta}{\sqrt{c}} \|A\|_F^2 \quad (49)$$

$$\mathcal{E}_w : \left\| C^T C - W^T W \right\|_F \leq \frac{\eta}{\sqrt{w}} \|A\|_F^2 \quad (50)$$

$$\mathcal{E}_r : \left\| \tilde{H}_\ell^T A - \tilde{H}_\ell^T (D_R S_R)^T D_R S_R A \right\|_F \leq \frac{3}{\delta \sqrt{r}} \left\| \tilde{H}_\ell \right\|_F \|A\|_F \quad (51)$$

Under the assumptions of this theorem, event \mathcal{E}_c holds with probability greater than $1 - \delta/3$ by Theorem 1, and similarly for event \mathcal{E}_w . Next, we claim that \mathcal{E}_r holds with probability greater than $1 - \delta/3$. To prove this it suffices to prove that

$$\mathbf{E} \left[\left\| \tilde{H}_\ell^T A - \tilde{H}_\ell^T (D_R S_R)^T D_R S_R A \right\|_F \right] \leq \frac{1}{\sqrt{r}} \left\| \tilde{H}_\ell \right\|_F \|A\|_F, \quad (52)$$

since the claim that $\Pr[\mathcal{E}_r] \geq 1 - \delta/3$ follows immediately from (52) by Markov's inequality; but (52) follows from Theorem 1 since the probabilities $\{p_i\}_{i=1}^m$ used to sample the columns of \tilde{H}_ℓ and the corresponding rows of A are of the form (11). Thus, under the assumptions of the theorem,

$$\Pr[\mathcal{E}_\xi] \geq 1 - \delta/3 \text{ for } \xi = c, w, r.$$

Next, from Lemma 4 and the Cauchy-Schwartz inequality, it follows that

$$\left\| \tilde{H}_\ell \right\|_F^2 = \sum_{t=1}^{\ell} \left| \tilde{h}^{t^T} \tilde{h}^t \right| = \sum_{t=1}^{\ell} 1 + \Delta_{tt} \leq k + \sqrt{k} \|\Delta\|_F. \quad (53)$$

Since $\sqrt{1+x} \leq 1 + \sqrt{x}$ for $x \geq 0$ it follows from (53) and (51) that under the event \mathcal{E}_r we have

$$\left\| \tilde{H}_\ell^T A - \tilde{H}_\ell^T (D_R S_R)^T D_R S_R A \right\|_F \leq \frac{3}{\delta \sqrt{r}} \left(\sqrt{k} + k^{1/4} \|\Delta\|_F^{1/2} \right) \|A\|_F. \quad (54)$$

By combining (54) with Lemma 5 we have that

$$\begin{aligned} \left\| \tilde{H}_\ell \tilde{H}_\ell^T A - C \tilde{U} R \right\|_F &\leq \frac{3}{\delta \sqrt{r}} \left(1 + \|\Delta\|_F^{1/2} \right) \left(\sqrt{k} + k^{1/4} \|\Delta\|_F^{1/2} \right) \|A\|_F \\ &\leq \left(\frac{9k}{\delta^2 r} \right)^{1/2} \left(1 + \|\Delta\|_F^{1/2} \right)^2 \|A\|_F \end{aligned} \quad (55)$$

$$\leq \left(\frac{9k}{\delta^2 r} \right)^{1/2} \left(1 + 3 \|\Delta\|_F^{1/2} \right) \|A\|_F \quad (56)$$

$$\leq \left(\frac{9k}{\delta^2 r} \right)^{1/2} \left(1 + \frac{3 \|C^T C - W^T W\|_F^{1/2}}{\sqrt{\gamma} \|W\|_F} \right) \|A\|_F, \quad (57)$$

where (55) follows since $k^{1/4} \leq \sqrt{k}$, (56) follows by multiplying out terms and since $\|\Delta\|_F \leq 1$ under the assumptions of the theorem, and (57) follows from the bound on $\|\Delta\|_F$ in Lemma 4.

Let us first consider establishing the Frobenius norm bound of (47). Recall that in this case we have set $\gamma = \epsilon/100k$. We will use the submultiplicativity of $\|\cdot\|_F$,

$$\left\| A - C \tilde{U} R \right\|_F \leq \left\| A - \tilde{H}_\ell \tilde{H}_\ell^T A \right\|_F + \left\| \tilde{H}_\ell \tilde{H}_\ell^T A - C \tilde{U} R \right\|_F, \quad (58)$$

and will bound each term separately. First, using the probabilities $\{q_j\}_{j=1}^n$ and the values of $c, w = \Omega(k^2 \eta^2 / \epsilon^8)$, then under the event $\mathcal{E}_c \cap \mathcal{E}_w$ we have that

$$\left\| A - \tilde{H}_\ell \tilde{H}_\ell^T A \right\|_F \leq \|A - A_k\|_F + \frac{\epsilon}{2} \|A\|_F, \quad (59)$$

by Theorem 3; see also [11]. In addition, using the sampling probabilities $\{p_i\}_{i=1}^m$ and values of $r = \Omega(k/\delta^2 \epsilon^2)$ and $w = \Omega(k^2 \eta^2 / \epsilon^8)$, and noting Lemma 6, it follows from (57) that under the event $\mathcal{E}_r \cap \mathcal{E}_w$:

$$\left\| \tilde{H}_\ell \tilde{H}_\ell^T A - C \tilde{U} R \right\|_F \leq \frac{\epsilon}{2} \|A\|_F. \quad (60)$$

Thus, under the event $\mathcal{E}_c \cap \mathcal{E}_w \cap \mathcal{E}_r$, which has probability at least $1 - \delta$, by combining (59) and (60) we see that (47) follows.

Let us next consider establishing the spectral norm bound of (48). Recall that in this case we have set $\gamma = \epsilon/100$ and that

$$\left\| A - C \tilde{U} R \right\|_2 \leq \left\| A - \tilde{H}_\ell \tilde{H}_\ell^T A \right\|_2 + \left\| \tilde{H}_\ell \tilde{H}_\ell^T A - C \tilde{U} R \right\|_F \quad (61)$$

by the submultiplicativity of $\|\cdot\|_2$ and since $\|\cdot\|_2 \leq \|\cdot\|_F$. First, using the probabilities $\{q_j\}_{j=1}^n$ and the values of $c, w = \Omega(\eta^2 / \epsilon^8)$, then under the event $\mathcal{E}_c \cap \mathcal{E}_w$ we have that

$$\left\| A - \tilde{H}_\ell \tilde{H}_\ell^T A \right\|_2 \leq \|A - A_k\|_2 + \frac{\epsilon}{2} \|A\|_F, \quad (62)$$

by Theorem 3; see also [11]. In addition, using the sampling probabilities $\{p_i\}_{i=1}^m$ and values of $r = \Omega(k/\delta^2 \epsilon^2)$ and $w = \Omega(\eta^2 / \epsilon^8)$, and noting Lemma 6, it follows from (57) that under the event $\mathcal{E}_r \cap \mathcal{E}_w$:

$$\left\| \tilde{H}_\ell \tilde{H}_\ell^T A - C \tilde{U} R \right\|_F \leq \frac{\epsilon}{2} \|A\|_F. \quad (63)$$

Thus, under the event $\mathcal{E}_c \cap \mathcal{E}_w \cap \mathcal{E}_r$, which has probability at least $1 - \delta$, by combining (62) and (63) we see that (48) follows. \diamond

As in the linear additional time framework, the results of Theorem 6 hold for all k and are of particular interest when A is well approximated by a matrix of low rank since then one may choose $k = O(1)$ and obtain a good approximation. In addition, it follows from (48) that

$$\|A - C\tilde{U}R\|_2 \leq (1/\sqrt{k} + \epsilon') \|A\|_F.$$

Thus, under the assumptions of Theorem 6 if we choose $k = 1/\epsilon'^2$ and let $\epsilon'' = 2\epsilon'$ then we have that

$$\|A - CUR\|_2 \leq \epsilon'' \|A\|_F \tag{64}$$

holds with probability at least $1 - \delta$.

For future reference [9], we note the following lemma.

Lemma 7

$$\|\tilde{U}\|_2 \leq \frac{O(1)}{\gamma \|A\|_F}$$

Proof: First note that $\|\tilde{U}\|_2 = \|\tilde{\Phi}\Psi^T\|_2 \leq \|\tilde{\Phi}\|_2 \|\Psi^T\|_2$. Since $\tilde{\Phi} = \sum_{t=1}^{\ell} \frac{1}{\sigma_t^2(W)} z^t z^{tT}$, we see that $\|\tilde{\Phi}\|_2 = \frac{1}{\sigma_t^2(W)} \leq \frac{1}{\gamma \|W\|_F^2}$. The lemma follows since, when using the probabilities of Theorem 6 we have that $\|\Psi^T\|_F \leq O(1) \|A\|_F$ and that $1/\|W\|_F^2 \leq O(1)/\|A\|_F^2$. \diamond

5 Discussion and Conclusion

To put the CUR decomposition in context, it will be useful to contrast it with the SVD. The SVD of A expresses A as $A = \sum_{t=1}^{\rho} \sigma_t(A) u^t v^{tT}$. Keeping the first k terms of this expansion, i.e., keeping $A_k = \sum_{t=1}^k \sigma_t(A) u^t v^{tT} = U_k \Sigma_k V_k^T$, gives us the “optimal” rank k approximation to A with respect to both the spectral norm and the Frobenius norm [14, 15]. Thus, computing the SVD gives us a good succinct approximation, since it only takes space $O(k(m+n))$ to write down U_k, Σ_k, V_k . However, the computational problem of finding the SVD cannot be carried out in a small constant number of passes. Our theorems say that weaker bounds, which are similar in spirit, may be achieved by CUR . Note, however, that although the SVD may be thought of as a rotation followed by a rescaling followed by a rotation, the CUR decomposition is quite different; both C and R perform an action like A and thus U must involve a pseudoinverse-like operation. Note also that the last upper bound, i.e., (5), is much smaller than $\|A\|_F$ when A has a good low-rank approximation. This is indeed the case for matrices occurring in many contexts, such as matrices for which Principal Component Analysis is used.

Recent work has focused on developing new techniques for proving lower bounds on the number of queries a sampling algorithm is required to perform in order to approximate a given function accurately with a low probability or error [3, 4]. In [4] these methods have been applied to the low-rank matrix approximation problem and to the matrix reconstruction problem. In the latter problem, the input is a matrix $A \in \mathbb{R}^{m \times n}$ and the goal is to find a matrix B that is close to A . In [4] it is shown that finding a B such that $\|A - B\|_F \leq \epsilon \|A\|_F$ requires $\Omega(mn)$ queries and that finding a B such that $\|A - B\|_2 \leq \epsilon \|A\|_F$ requires $\Omega(m+n)$ queries. Thus, our algorithm is optimal for constant ϵ .

Acknowledgments We would like to thank the National Science Foundation for partial support of this work.

References

- [1] D. Achlioptas and F. McSherry. Fast computation of low rank matrix approximations. *submitted*.
- [2] D. Achlioptas and F. McSherry. Fast computation of low rank matrix approximations. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 611–618, 2001.
- [3] Z. Bar-Yossef. *The Complexity of Massive Data Set Computations*. PhD thesis, University of California, Berkeley, 2002.
- [4] Z. Bar-Yossef. Sampling lower bounds via information theory. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 335–344, 2003.
- [5] R. Bhatia. *Matrix Analysis*. Springer-Verlag, New York, 1997.
- [6] E. Cohen and D.D. Lewis. Approximating matrix multiplication for pattern recognition tasks. *Journal of Algorithms*, 30(2):211–252, 1999.
- [7] P. Drineas, A. Frieze, R. Kannan, S. Vempala, and V. Vinay. Clustering in large graphs and matrices. In *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 291–299, 1999.
- [8] P. Drineas and R. Kannan. Pass efficient algorithms for approximating large matrices. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 223–232, 2003.
- [9] P. Drineas, R. Kannan, and M.W. Mahoney. Sampling sub-problems of heterogeneous Max-Cut problems and approximation algorithms. *manuscript*.
- [10] P. Drineas, R. Kannan, and M.W. Mahoney. Fast Monte Carlo algorithms for matrices I: Approximating matrix multiplication. Technical Report YALEU/DCS/TR-1269, Yale University Department of Computer Science, New Haven, CT, February 2004.
- [11] P. Drineas, R. Kannan, and M.W. Mahoney. Fast Monte Carlo algorithms for matrices II: Computing a low-rank approximation to a matrix. Technical Report YALEU/DCS/TR-1270, Yale University Department of Computer Science, New Haven, CT, February 2004.
- [12] P. Drineas, I. Kerenidis, and P. Raghavan. Competitive recommendation systems. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 82–90, 2002.
- [13] A. Frieze, R. Kannan, and S. Vempala. Fast Monte-Carlo algorithms for finding low-rank approximations. In *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science*, pages 370–378, 1998.
- [14] G.H. Golub and C.F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, 1989.
- [15] R.A. Horn and C.R. Johnson. *Matrix Analysis*. Cambridge University Press, New York, 1985.

- [16] P. Indyk. Stable distributions, pseudorandom generators, embeddings and data stream computation. In *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science*, pages 189–197, 2000.
- [17] J. Kleinberg. Two algorithms for nearest-neighbor search in high dimensions. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pages 599–608, 1997.
- [18] K.V. Mardia, J.T. Kent, and J.M. Bibby. *Multivariate Analysis*. Academic Press, London, 1979.
- [19] G.W. Stewart and J.G. Sun. *Matrix Perturbation Theory*. Academic Press, New York, 1990.
- [20] S. Vempala. Random projection: A new approach to VLSI layout. In *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science*, pages 389–395, 1998.