

FAST MONTE CARLO ALGORITHMS FOR MATRICES II: COMPUTING A LOW-RANK APPROXIMATION TO A MATRIX*

PETROS DRINEAS[†], RAVI KANNAN[‡], AND MICHAEL W. MAHONEY[§]

Abstract. In many applications, the data consist of (or may be naturally formulated as) an $m \times n$ matrix A . It is often of interest to find a low-rank approximation to A , i.e., an approximation D to the matrix A of rank not greater than a specified rank k , where k is much smaller than m and n . Methods such as the singular value decomposition (SVD) may be used to find an approximation to A which is the best in a well-defined sense. These methods require memory and time which are superlinear in m and n ; for many applications in which the data sets are very large this is prohibitive. Two simple and intuitive algorithms are presented which, when given an $m \times n$ matrix A , compute a description of a low-rank approximation D^* to A , and which are qualitatively faster than the SVD. Both algorithms have provable bounds for the error matrix $A - D^*$. For any matrix X , let $\|X\|_F$ and $\|X\|_2$ denote its Frobenius norm and its spectral norm, respectively. In the first algorithm, c columns of A are randomly chosen. If the $m \times c$ matrix C consists of those c columns of A (after appropriate rescaling), then it is shown that from $C^T C$ approximations to the top singular values and corresponding singular vectors may be computed. From the computed singular vectors a description D^* of the matrix A may be computed such that $\text{rank}(D^*) \leq k$ and such that

$$\|A - D^*\|_\xi^2 \leq \min_{D: \text{rank}(D) \leq k} \|A - D\|_\xi^2 + \text{poly}(k, 1/c) \|A\|_F^2$$

holds with high probability for both $\xi = 2, F$. 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 $O(cm + c^2)$ additional RAM. The second algorithm is similar except that it further approximates the matrix C by randomly sampling r rows of C to form a $r \times c$ matrix W . Thus, it has additional error, but it can be implemented in three passes over the matrix using only constant additional RAM. To achieve an additional error (beyond the best rank k approximation) that is at most $\epsilon \|A\|_F^2$, both algorithms take time which is polynomial in k , $1/\epsilon$, and $\log(1/\delta)$, 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 . Our bounds improve previously published results with respect to the rank parameter k for both the Frobenius and spectral norms. In addition, the proofs for the error bounds use a novel method that makes important use of matrix perturbation theory. The probability distribution over columns of A and the rescaling are crucial features of the algorithms which must be chosen judiciously.

Key words. randomized algorithms, Monte Carlo methods, massive data sets, singular value decomposition

AMS subject classification. 68W20

DOI. 10.1137/S0097539704442696

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 the singular value decomposition (SVD); in two related papers we consider matrix multiplication and a new method for computing a compressed approximate decomposition of a large matrix [13, 14]. Since such computations generally require

*Received by the editors April 5, 2004; accepted for publication (in revised form) November 17, 2005; published electronically May 26, 2006. The technical report version of this paper appeared as *Fast Monte Carlo Algorithms for Matrices II: Computing a Low-Rank Approximation to a Matrix*, by P. Drineas, R. Kannan, and M. W. Mahoney [15].

<http://www.siam.org/journals/sicomp/36-1/44269.html>

[†]Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY 12180 (drinep@cs.rpi.edu).

[‡]Department of Computer Science, Yale University, New Haven, CT 06520 (kannan@cs.yale.edu). This author was supported in part by a grant from the NSF.

[§]Department of Mathematics, Yale University, New Haven, CT 06520 (mahoney@cs.yale.edu).

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. Our algorithms will be allowed to read the matrices a few, e.g., one, 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 [13, 12].

In many applications, the data consist of (or may be naturally formulated as) an $m \times n$ matrix A which is either low-rank or is well approximated by a low-rank matrix [7, 9, 25, 3, 26, 28, 29, 24, 22]. In these application areas, e.g., latent semantic indexing, DNA microarray analysis, facial and object recognition, and web search models, the data may consist of m points in \mathbb{R}^n . Let $A \in \mathbb{R}^{m \times n}$ be the matrix with these points as rows. Two methods for dealing with such high-dimensional data are the SVD (and the related principal components analysis) and multidimensional scaling [18, 23]. Thus, it is often of interest to find a low-rank approximation to A , i.e., an approximation D , of rank no greater than a specified rank k , to the matrix A , where k is much smaller than m and n . For example, this rank reduction is used in many applications of linear algebra and statistics as well as in image processing, lossy data compression, text analysis, and cryptography [6]. The SVD may be used to find an approximation to A which is the best in a well-defined sense [18, 19], but it requires a superlinear (in m and n) polynomial time dependence that is prohibitive for many applications in which the data sets are very large. Another method that has attracted interest recently is the traditional “random projection” method where one projects the problem into a randomly chosen low-dimensional subspace [21, 30, 20]. This dimensional reduction requires performing an operation that amounts to premultiplying the given $m \times n$ matrix A by an $s \times m$ matrix which takes time dependent in a superlinear manner on $m + n$.

In this paper we present two simple and intuitive algorithms which, when given an $m \times n$ matrix A , compute a description of a low-rank approximation D^* to A , and which are qualitatively faster than the SVD. Both algorithms have provable bounds for the error matrix $A - D^*$. For any matrix X , let $\|X\|_F$ and $\|X\|_2$ denote its Frobenius norm and its spectral norm (as defined in section 3.1), respectively. In the first algorithm, the LINEARTIMESVD algorithm of section 4, c columns of A are randomly chosen. If the $m \times c$ matrix C consists of those c columns of A (after appropriate rescaling), then it is shown that from $C^T C$ approximations to the top singular values and corresponding singular vectors of A may be computed. From the computed singular vectors a description D^* of the matrix A may be computed such that $\text{rank}(D^*) \leq k$ and such that

$$(1) \quad \|A - D^*\|_\xi^2 \leq \min_{D: \text{rank}(D) \leq k} \|A - D\|_\xi^2 + \text{poly}(k, 1/c) \|A\|_F^2$$

holds with high probability for each of $\xi = 2, F$. 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 $O(cm + c^2)$ additional RAM. The second algorithm, the CONSTANTIMESVD algorithm of section 5, is similar except that it further approximates the matrix C by randomly sampling r rows of C to form an $r \times c$

TABLE 1
Summary of sampling complexity.

Additional error for:	LINEARTIMESVD	CONSTANTTIMESVD	REF. [16, 17]
$\ A - D^*\ _2^2$	$1/\epsilon^2$	$1/\epsilon^4$	k^4/ϵ^3
$\ A - D^*\ _F^2$	k/ϵ^2	k^2/ϵ^4	k^4/ϵ^3

matrix W . Thus, it has additional error but can be implemented in three passes over the matrix using additional RAM that is $O(cr)$, i.e., that is constant for constant c and r . To achieve an additional error that is at most $\epsilon \|A\|_F^2$, both algorithms take time which is polynomial in k , $1/\epsilon$, and $\log(1/\delta)$, where $\delta > 0$ is a failure probability; see Table 1 for a summary of the dependence of the sampling complexity on k and ϵ . The first algorithm takes time linear in $\max(m, n)$ and the other takes time independent of m and n . Our bounds improve previously published results with respect to the rank parameter k for both the Frobenius and the spectral norms. In addition, the proofs for the error bounds use a novel method that makes important use of matrix perturbation theory. The probability distribution over columns of A and the rescaling are crucial features of the algorithms which must be chosen judiciously.

It is worth emphasizing how this work fits into recent work on computing low-rank matrix approximations. In the original work of Frieze, Kannan, and Vempala [16] (see also [17]) it was shown that by working with a randomly chosen and constant-sized submatrix of A , one could obtain bounds of the form (1) for the Frobenius norm (and thus indirectly for the spectral norm). To achieve an additional error that is at most $\epsilon \|A\|_F^2$, the size of the submatrix was a constant with respect to m and n but depended polynomially on k and $1/\epsilon$; although the submatrix was constant-sized, its construction (in particular, the construction of the sampling probabilities) required space and thus time that was linear in $m + n$. In this work, we modify the algorithm of [16] so that both the construction of *and* the computation on the constant-sized submatrix requires only constant additional space and time; thus, it fits within the framework of the pass-efficient model of data-streaming computation [12, 13]. In addition, we provide a different proof of the main result of [16] for the Frobenius norm and improve the polynomial dependence on k . Our proof method is quite different than that of [16]; it relies heavily on the approximate matrix multiplication result of [13] and [11] and it uses the Hoffman–Wielandt inequality. In addition, we provide a proof of a direct and significantly improved bound with respect to the spectral norm. Since these results are technically quite complex, we also present the corresponding proofs for both norms in the linear additional space and time framework [12, 13]. These latter results have been presented in the context of clustering applications [10], but are included here for completeness and to provide motivation and clarity for the more complex constant time results. Table 1 provides a summary of our results, for both the linear and the constant time models, and shows the number of rows and columns to be sampled sufficient to ensure, with high probability, an additional error of $\epsilon \|A\|_F^2$ in (1); see section 6 for more discussion.

In other related work, Achlioptas and McSherry have also computed low-rank approximations using somewhat different sampling techniques [2, 1]. The primary focus of their work was in introducing methods to accelerate orthogonal iteration and Lanczos iteration, which are two commonly used methods for computing low-rank approximations to a matrix. Also included in [2, 1] is a comparison of their methods with those of [10, 12, 16] and thus with the results we present here. Our algorithms

and those of [16] and [2, 1] come with mathematically rigorous guarantees of the running time and of the quality of the approximation produced. As far as we know, so-called incremental SVD algorithms, which bring as much of the data as possible into memory, compute the SVD, and then update this SVD in an incremental fashion with the remaining data, do not come with such guarantees.

In section 2 several applications areas that deal with large matrices are discussed, and in section 3 we provide a review of relevant linear algebra, the pass-efficient model, and an approximate matrix multiplication result that will be used extensively. Then, in section 4 our linear additional space and time approximation algorithm, the LINEARTIMESVD algorithm, is presented and analyzed; in section 5 our constant additional space and time approximation algorithm, the CONSTANTTIMESVD algorithm, is presented and analyzed. Finally, in section 6 a discussion and conclusion are presented.

2. Some applications. There are numerous applications in which the data are well approximated by a low-rank matrix. In this section we discuss several such applications to provide motivation for our algorithms.

2.1. Latent semantic indexing. Latent semantic indexing (LSI) is a general technique for analyzing a collection of documents which are assumed to be related [7, 9, 25]. Approaches to retrieving textual information from databases that depend on a lexical match between words in the query and words in the document can be inaccurate, both because often users want to retrieve information on the basis of conceptual content and because individual words do not in general provide reliable evidence about the conceptual topic of a document. LSI is an alternative matching method that attempts to overcome problems associated with lexical matching; it does so by assuming that there is some underlying or latent semantic structure that is partially obscured by variability in word choice and then using techniques such as SVD to remove the noise and estimate this latent structure.

Suppose that there are m documents and n terms which occur in the documents. Latent semantic structure analysis starts with a term-document matrix, e.g., a matrix $A \in \mathbb{R}^{m \times n}$, where A_{ij} is frequency of the j th term in the i th document. A topic is modeled as an n -vector of nonnegative reals summing to 1, where the j th component of a topic vector is interpreted as the frequency with which the j th term occurs in a discussion of the topic. By assumption, the number of topics that the documents are about is small relative to the number of unique terms n . It can be argued that, for a given k , finding a set of k topics which best describe the documents corresponds to keeping only the top k singular vectors of A ; most of the important underlying structure in the association of terms and documents will then be kept and most of the noise or variability in word usage will be removed.

2.2. DNA microarray data. DNA microarray technology has been used to study a variety of biological processes since it permits the monitoring of the expression levels of thousands of genes under a range of experimental conditions [3, 26, 28]. Depending on the particular technology, either the absolute or the relative expression levels of m genes, which for model organisms may constitute nearly the entire genome, are probed simultaneously by a single microarray. A series of n arrays probe genome-wide expression levels in n different samples, i.e., under n different experimental conditions. The data from microarray experiments may thus be represented as a matrix $A \in \mathbb{R}^{m \times n}$, where A_{ij} represents the expression level of gene i under experimental condition j . From this matrix, both the relative expression level of the

i th gene under every condition and also the relative expression level of every gene under the j th condition may be easily extracted.

This matrix is low-rank and thus a small number of eigengenes and corresponding eigenarrays (left and right singular vectors) are sufficient to capture most of the gene expression information. Removing the rest, which correspond to noise or experimental artifacts, enables meaningful comparison of the expression levels of different genes. When processing and modeling genome-wide expression data, the SVD and its low-rank approximation provides a framework such that the mathematical variables and operations suggest assigned biological meaning, e.g., in terms of cellular regulatory processes and cellular states, that may be hidden in the original data due to experimental noise or hidden dependencies. Expression data has been used for inference tasks such as to identify genes based on coexpression, predict regulatory elements, and reverse-engineer transcription networks, but this inference is difficult with noise or dependencies.

2.3. Eigenfaces and facial recognition. Applications of SVD and low-rank approximations in computer vision include pattern estimation, image compression and restoration, and facial and object recognition, where the concept of eigenfaces has been useful [29, 24].

The goal of facial recognition is to recognize a certain face given a database of photographs of human faces under variations in lighting conditions and pose view-points. A common approach is to represent the database as a matrix in which the rows of the matrix are the images represented as vectors. Thus, if there are m images, each of which is of size $n \times n$, the matrix $A \in \mathbb{R}^{m \times n^2}$ represents the database of images, where A_{ij} is the j th pixel value in the i th image. Typically, $m \ll n^2$ and, although many of the singular vectors are needed for very accurate reconstruction of an image, often only a few of the singular vectors are needed to extract the major appearance characteristics of an image. The right singular vectors of the matrix A are known as eigenfaces since they are the principal components or eigenvectors of the associated correlation matrix of the set of face images. The eigenfaces are computed and they are used to project the database of photographs to a lower-dimensional space that spans the significant variations among known facial images. Then, given a new image, it is projected to the same low-dimensional space, and its position is then compared to the images in the database.

2.4. Web search model. The problem of how to extract information from the network structure of a hyperlinked environment such as the World Wide Web was considered by Kleinberg [22]. This is of interest, for example, if one wants to find web pages that are relevant to a given query and one is using a keyword-based web search program, since there is no obvious endogenous measure of an authoritative page that would favor it under a text-based ranking system.

Starting with a set of pages returned by a text-based search engine, a document is defined to be an authority if many other documents returned by the search point to it, i.e., have a hypertext link to it. A document is defined to be a hub if it points to many other documents. More generally, suppose n documents are returned by the search engine. Then, a matrix $A \in \mathbb{R}^{m \times n}$ is defined, where A_{ij} is 1 or 0 depending on whether the i th document points to the j th document. Kleinberg attempts to find two n -vectors, x and y , where x_i is the hub weight of document i and y_j is the authority weight of document j . He then argues that it is desirable to find $\max_{|x|=|y|=1} x^T A y$, where $|\cdot|$ denotes the Euclidean length, since in maximizing x, y one expects the hub weights and authority weights to be mutually consistent. This is simply the problem of

finding the singular vectors of A . Since A is large, he judiciously chooses a submatrix of A and computes only the singular vectors of it. In the case when the key word has multiple meanings, not only the top but also some of the other singular vectors with large singular values are interesting. Thus, it is of interest to find the k largest singular vectors form some small k . This is the problem we are considering, and we also find the singular vectors of a submatrix, but a randomly chosen one.

3. Review of relevant background. This section contains a review of linear algebra that will be useful throughout the paper; for more detail, see [18, 19, 27, 8] and the references therein. This section also contains a review of the pass-efficient model of data-streaming computation (which provides a framework within which our SVD results may be viewed) and a matrix multiplication result that will be used extensively in our proofs; see [11, 12, 13] for more details.

3.1. Review of linear algebra. For a vector $x \in \mathbb{R}^n$ we let $x_i, i = 1, \dots, n$, denote the i th element of x and we let $|x| = (\sum_{i=1}^n |x_i|^2)^{1/2}$. 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; thus, if A_{ij} denotes the (i, j) th element of A , $A_{ij} = (A^{(j)})_i = (A_{(i)})_j$. The range of an $A \in \mathbb{R}^{m \times n}$ is

$$\text{range}(A) = \{y \in \mathbb{R}^m : y = Ax \text{ for some } x \in \mathbb{R}^n\} = \text{span}(A^{(1)}, \dots, A^{(n)}).$$

The rank of A , $\text{rank}(A)$, is the dimension of $\text{range}(A)$ and is equal to the number of linearly independent columns of A ; since this is equal to $\text{rank}(A^T)$ it also equals the number of linearly independent rows of A . The null space of A is

$$\text{null}(A) = \{x \in \mathbb{R}^n : Ax = 0\}.$$

For a matrix $A \in \mathbb{R}^{m \times n}$ we denote matrix norms by $\|A\|_\xi$, using subscripts to distinguish between various norms. Of particular interest will be the Frobenius norm, which is defined by

$$(2) \quad \|A\|_F = \sqrt{\sum_{i=1}^m \sum_{j=1}^n A_{ij}^2}.$$

If $\text{Tr}(A)$ is the matrix trace which is the sum of the diagonal elements of A , then $\|A\|_F^2 = \text{Tr}(A^T A) = \text{Tr}(AA^T)$. Also of interest is the spectral norm, which is defined by

$$(3) \quad \|A\|_2 = \sup_{x \in \mathbb{R}^n, x \neq 0} \frac{|Ax|}{|x|}.$$

Both of these norms are submultiplicative and unitarily invariant and they are related to each other as

$$\|A\|_2 \leq \|A\|_F \leq \sqrt{n} \|A\|_2.$$

Both of these norms provide a measure of the “size” of the matrix A . Note that if $A \in \mathbb{R}^{m \times n}$, then there exists an $x \in \mathbb{R}^n$ such that $|x| = 1$ and $A^T Ax = \|A\|_2^2 x$ and that if $\{x^1, x^2, \dots, x^n\}$ is any basis of \mathbb{R}^n and if $A \in \mathbb{R}^{m \times n}$, then $\|A\|_F^2 = \sum_{i=1}^n |Ax^i|^2$.

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

$$(4) \quad U^T A V = \Sigma = \mathbf{diag}(\sigma_1, \dots, \sigma_\rho),$$

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 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. The columns of U and V satisfy the relations $A v^i = \sigma_i u^i$ and $A^T u^i = \sigma_i v^i$. For symmetric matrices the left and right singular vectors are the same. The singular values of A are the nonnegative square roots of the eigenvalues of $A^T A$ and of $A A^T$; furthermore, the columns of U , i.e., the left singular vectors, are eigenvectors of $A A^T$ and the columns of V , i.e., the right singular vectors, are eigenvectors of $A^T A$.

The SVD can reveal important information about the structure of a matrix. If we define r by $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > \sigma_{r+1} = \dots = \sigma_\rho = 0$, then $\text{rank}(A) = r$, $\text{null}(A) = \text{span}(v^{r+1}, \dots, v^\rho)$, and $\text{range}(A) = \text{span}(u^1, \dots, u^r)$. If we let $U_r \in \mathbb{R}^{m \times r}$ denote the matrix consisting of the first r columns of U , $V_r \in \mathbb{R}^{r \times n}$ denote the matrix consisting of the first r columns of V , and $\Sigma_r \in \mathbb{R}^{r \times r}$ denote the principal $r \times r$ submatrix of Σ , then

$$(5) \quad A = U_r \Sigma_r V_r^T = \sum_{t=1}^r \sigma_t u^t v^{tT}.$$

Note that this decomposition property provides a canonical description of a matrix as a sum of r rank-one matrices of decreasing importance. If $k \leq r$ and we define

$$(6) \quad A_k = U_k \Sigma_k V_k^T = \sum_{t=1}^k \sigma_t u^t v^{tT},$$

then $A_k = U_k U_k^T A = (\sum_{t=1}^k u^t u^{tT}) A$ and $A_k = A V_k V_k^T = A (\sum_{t=1}^k v^t v^{tT})$, i.e., A_k is the projection of A onto the space spanned by the top k singular vectors of A . Furthermore, 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.,

$$(7) \quad \min_{D \in \mathbb{R}^{m \times n} : \text{rank}(D) \leq k} \|A - D\|_2 = \|A - A_k\|_2 = \sigma_{k+1}(A)$$

and

$$(8) \quad \min_{D \in \mathbb{R}^{m \times n} : \text{rank}(D) \leq k} \|A - D\|_F^2 = \|A - A_k\|_F^2 = \sum_{t=k+1}^r \sigma_t^2(A).$$

Thus, A_k constructed from the k largest singular triplets of A is the optimal rank- k approximation to A with respect to both $\|\cdot\|_F$ and $\|\cdot\|_2$. More generally, one can also show that $\|A\|_2 = \sigma_1$ and that $\|A\|_F^2 = \sum_{i=1}^r \sigma_i^2$.

From the perturbation theory of matrices it is known that the size of the difference between two matrices can be used to bound the difference between the singular value spectrum of the two matrices [27, 8]. In particular, if $A, E \in \mathbb{R}^{m \times n}$, $m \geq n$, then

$$(9) \quad \max_{t: 1 \leq t \leq n} |\sigma_t(A + E) - \sigma_t(A)| \leq \|E\|_2$$

and

$$(10) \quad \sum_{k=1}^n (\sigma_k(A + E) - \sigma_k(A))^2 \leq \|E\|_F^2.$$

The latter inequality is known as the Hoffman–Wielandt inequality.

3.2. 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 [13, 12]. 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 [13] for more details.

3.3. Review of matrix multiplication. The BASICMATRIXMULTIPLICATION algorithm to approximate the product of two matrices is presented and analyzed in [13]. 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 *nearly optimal probabilities* if they are of the form (11) and are the *optimal probabilities* (with respect to approximating the product AB) if they are of the form (11) with $\beta = 1$. In [13] we prove the following theorem.

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 $p_i \geq 0$, $\sum_{i=1}^n p_i = 1$ and such that for some positive constant $\beta \leq 1$*

$$(11) \quad p_k \geq \frac{\beta |A^{(k)}| |B_{(k)}|}{\sum_{k'=1}^n |A^{(k')}| |B_{(k')}|}.$$

Construct C and R with the BASICMATRIXMULTIPLICATION algorithm of [13] and let CR be an approximation to AB . Then

$$(12) \quad \mathbf{E}[\|AB - CR\|_F^2] \leq \frac{1}{\beta c} \|A\|_F^2 \|B\|_F^2.$$

Furthermore, let $\delta \in (0, 1)$ and $\eta = 1 + \sqrt{(8/\beta) \log(1/\delta)}$. Then, with probability at least $1 - \delta$,

$$(13) \quad \|AB - CR\|_F^2 \leq \frac{\eta^2}{\beta c} \|A\|_F^2 \|B\|_F^2.$$

In [13] it is shown that after one pass over the matrices nearly optimal probabilities can be constructed. In the present paper, we will be particularly interested in the case that $B = A^T$. In this case, using the SELECT algorithm of [13] random samples can be drawn according to nearly optimal probabilities using $O(1)$ additional space and time.

LINEARTIMESVD Algorithm.

Input: $A \in \mathbb{R}^{m \times n}$, $c, k \in \mathbb{Z}^+$ such that $1 \leq k \leq c \leq n$, $\{p_i\}_{i=1}^n$ such that $p_i \geq 0$ and $\sum_{i=1}^n p_i = 1$.

Output: $H_k \in \mathbb{R}^{m \times k}$ and $\sigma_t(C), t = 1, \dots, k$.

1. For $t = 1$ to c ,
 - (a) Pick $i_t \in 1, \dots, n$ with $\Pr[i_t = \alpha] = p_\alpha, \alpha = 1, \dots, n$.
 - (b) Set $C^{(t)} = A^{(i_t)} / \sqrt{cp_{i_t}}$.
2. Compute $C^T C$ and its SVD; say $C^T C = \sum_{t=1}^c \sigma_t^2(C) y^t y^{tT}$.
3. Compute $h^t = C y^t / \sigma_t(C)$ for $t = 1, \dots, k$.
4. Return H_k , where $H_k^{(t)} = h^t$, and $\sigma_t(C), t = 1, \dots, k$.

FIG. 1. The LINEARTIMESVD algorithm.




FIG. 2. Diagram for the LINEARTIMESVD algorithm.

4. Linear time SVD approximation algorithm.

4.1. The algorithm. Given a matrix $A \in \mathbb{R}^{m \times n}$ we wish to approximate its top k singular values and the corresponding singular vectors in a constant number of passes through the data and $O(cm + c^2)$ additional space and $O(c^2m + c^3)$ additional time. The strategy behind the LINEARTIMESVD algorithm is to pick c columns of the matrix A , rescale each by an appropriate factor to form a matrix $C \in \mathbb{R}^{m \times c}$, and then compute the singular values and corresponding left singular vectors of the matrix C , which will be approximations to the singular values and left singular vectors of A , in a sense we make precise later. These are calculated by performing an SVD of the matrix $C^T C$ to compute the right singular vectors of C and from them calculating the left singular vectors of C .

The LINEARTIMESVD algorithm is described in Figure 1; it takes as input a matrix A and returns as output an approximation to the top k left singular values and the corresponding singular vectors. Note that by construction the SVD of C is

$C = H\Sigma_C Y^T$. A diagram illustrating the action of the LINEARTIMESVD algorithm is presented in Figure 2. The transformation represented by the matrix A is shown along with its SVD, and the transformation represented by the matrix C is also shown along with its SVD. It will be shown that if the probabilities $\{p_i\}_{i=1}^n$ are chosen judiciously, then the left singular vectors of C are with high probability approximations to the left singular vectors of A .

In section 4.2 we discuss running-time issues, and in section 4.3 we will prove the correctness of the algorithm.

4.2. Analysis of the implementation and running time. Assuming that nearly optimal sampling probabilities (as defined in section 3.3) are used, then in the LINEARTIMESVD algorithm the sampling probabilities p_k can be used to select columns to be sampled in one pass and $O(c)$ additional space and time using the SELECT algorithm of [13]. 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 $O(mc)$. Given $C \in \mathbb{R}^{m \times c}$, computing $C^T C$ requires $O(mc)$ additional space and $O(mc^2)$ additional time, and computing the SVD of $C^T C$ requires $O(c^3)$ additional time. Then computing H_k requires k matrix-vector multiplications for a total of $O(mck)$ additional space and time. Thus, overall $O(cm + c^2)$ additional space and $O(c^2 m + c^3)$ additional time are required by the LINEARTIMESVD algorithm. Note that the “description” of the solution that is computable in the allotted additional space and time is the explicit approximation to the top k singular values and corresponding left singular vectors.

4.3. Analysis of the sampling step. Approximating A by $A_k = U_k U_k^T A$ incurs an error equal to $\|A - A_k\|_F^2 = \sum_{t=k+1}^r \sigma_t^2(A)$ and $\|A - A_k\|_2 = \sigma_{k+1}(A)$, since A_k is the “optimal” rank- k approximation to A with respect to both $\|\cdot\|_F$ and $\|\cdot\|_2$. We will show that in addition to this error the matrix $H_k H_k^T A$ has an error that depends on $\|AA^T - CC^T\|_F$. Then, using the results of Theorem 1, we will show that this additional error depends on $\|A\|_F^2$. We first consider obtaining a bound with respect to the Frobenius norm.

THEOREM 2. *Suppose $A \in \mathbb{R}^{m \times n}$ and let H_k be constructed from the LINEAR-TIMESVD algorithm. Then*

$$\|A - H_k H_k^T A\|_F^2 \leq \|A - A_k\|_F^2 + 2\sqrt{k} \|AA^T - CC^T\|_F.$$

Proof. Recall that for matrices X and Y , $\|X\|_F^2 = \mathbf{Tr}(X^T X)$, $\mathbf{Tr}(X + Y) = \mathbf{Tr}(X) + \mathbf{Tr}(Y)$, and also that $H_k^T H_k = I_k$. Thus, we may express $\|A - H_k H_k^T A\|_F^2$ as

$$\begin{aligned} \|A - H_k H_k^T A\|_F^2 &= \mathbf{Tr}((A - H_k H_k^T A)^T (A - H_k H_k^T A)) \\ &= \mathbf{Tr}(A^T A - 2A^T H_k H_k^T A + A^T H_k H_k^T H_k H_k^T A) \\ &= \mathbf{Tr}(A^T A) - \mathbf{Tr}(A^T H_k H_k^T A) \\ (14) \qquad &= \|A\|_F^2 - \|A^T H_k\|_F^2. \end{aligned}$$

We may relate $\|A^T H_k\|_F^2$ and $\sum_{t=1}^k \sigma_t^2(C)$ by the following:

$$\begin{aligned}
\left| \|A^T H_k\|_F^2 - \sum_{t=1}^k \sigma_t^2(C) \right| &\leq \sqrt{k} \left(\sum_{t=1}^k \left(|A^T h^t|^2 - \sigma_t^2(C) \right)^2 \right)^{1/2} \\
&= \sqrt{k} \left(\sum_{t=1}^k \left(|A^T h^t|^2 - |C^T h^t|^2 \right)^2 \right)^{1/2} \\
&= \sqrt{k} \left(\sum_{t=1}^k \left(h^{tT} (AA^T - CC^T) h^t \right)^2 \right)^{1/2} \\
(15) \qquad &\leq \sqrt{k} \|AA^T - CC^T\|_F.
\end{aligned}$$

The first inequality follows by applying the Cauchy–Schwarz inequality; the last inequality follows by writing AA^T and CC^T with respect to a basis containing $\{h^t\}_{t=1}^k$. By again applying the Cauchy–Schwarz inequality, noting that $\sigma_t^2(X) = \sigma_t(XX^T)$ for a matrix X , and applying the Hoffman–Wielandt inequality (10), we may also relate $\sum_{k=1}^k \sigma_t^2(C)$ and $\sum_{k=1}^k \sigma_t^2(A)$ by the following:

$$\begin{aligned}
\left| \sum_{t=1}^k \sigma_t^2(C) - \sum_{t=1}^k \sigma_t^2(A) \right| &\leq \sqrt{k} \left(\sum_{t=1}^k \left(\sigma_t^2(C) - \sigma_t^2(A) \right)^2 \right)^{1/2} \\
&= \sqrt{k} \left(\sum_{t=1}^k \left(\sigma_t(CC^T) - \sigma_t(AA^T) \right)^2 \right)^{1/2} \\
&\leq \sqrt{k} \left(\sum_{t=1}^m \left(\sigma_t(CC^T) - \sigma_t(AA^T) \right)^2 \right)^{1/2} \\
(16) \qquad &\leq \sqrt{k} \|CC^T - AA^T\|_F.
\end{aligned}$$

Combining the results of (15) and (16) allows us to relate $\|A^T H_k\|_F^2$ and $\sum_{t=1}^k \sigma_t^2(A)$ by the following:

$$(17) \qquad \left| \|A^T H_k\|_F^2 - \sum_{t=1}^k \sigma_t^2(A) \right| \leq 2\sqrt{k} \|AA^T - CC^T\|_F.$$

Combining (17) with (14) yields the theorem. \square

We next prove a similar result for the spectral norm; note that the factor \sqrt{k} is not present.

THEOREM 3. *Suppose $A \in \mathbb{R}^{m \times n}$ and let H_k be constructed from the LINEAR-TIMESVD algorithm. Then*

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

Proof. Let $\mathcal{H}_k = \text{range}(H_k) = \text{span}(h^1, \dots, h^k)$ and let \mathcal{H}_{m-k} be the orthogonal complement of \mathcal{H}_k . Let $x \in \mathbb{R}^m$ and let $x = \alpha y + \beta z$, where $y \in \mathcal{H}_k$, $z \in \mathcal{H}_{m-k}$, and

$\alpha^2 + \beta^2 = 1$; then

$$\begin{aligned} \|A - H_k H_k^T A\|_2 &= \max_{x \in \mathbb{R}^m, |x|=1} |x^T (A - H_k H_k^T A)| \\ &= \max_{y \in \mathcal{H}_k, |y|=1, z \in \mathcal{H}_{m-k}, |z|=1, \alpha^2 + \beta^2 = 1} |(\alpha y^T + \beta z^T)(A - H_k H_k^T A)| \\ (18) \quad &\leq \max_{y \in \mathcal{H}_k, |y|=1} |y^T (A - H_k H_k^T A)| + \max_{z \in \mathcal{H}_{m-k}, |z|=1} |z^T (A - H_k H_k^T A)| \end{aligned}$$

$$(19) \quad = \max_{z \in \mathcal{H}_{m-k}, |z|=1} |z^T A|.$$

Inequality (18) follows since $\alpha, \beta \leq 1$ and (19) follows since $y \in \mathcal{H}_k$ and $z \in \mathcal{H}_{m-k}$. We next bound (19):

$$(20) \quad |z^T A|^2 = z^T C C^T z + z^T (A A^T - C C^T) z$$

$$(21) \quad \leq \sigma_{k+1}^2(C) + \|A A^T - C C^T\|_2$$

$$(22) \quad \leq \sigma_{k+1}^2(A) + 2\|A A^T - C C^T\|_2$$

$$(22) \quad = \|A - A_k\|_2^2 + 2\|A A^T - C C^T\|_2.$$

Inequality (20) follows since $\max_{z \in \mathcal{H}_{m-k}} |z^T C|$ occurs when z is the $(k+1)$ st left singular vector, i.e., the maximum possible in the \mathcal{H}_{m-k} subspace. Inequality (21) follows since $\sigma_{k+1}^2(C) = \sigma_{k+1}(C C^T)$ and since by (9) we have that $\sigma_{k+1}^2(C) \leq \sigma_{k+1}(A A^T) + \|A A^T - C C^T\|_2$; (22) follows since $\|A - A_k\|_2 = \sigma_{k+1}(A)$. The theorem then follows by combining (19) and (22). \square

Theorems 2 and 3 hold regardless of the sampling probabilities $\{p_i\}_{i=1}^n$. Since $\|A - A_k\|_\xi$, $\xi = 2, F$, is a property of the matrix A , the choice of sampling probabilities enters into the error of $\|A - H_k H_k^T A\|_\xi^2$ only through the term involving the additional error beyond the optimal rank- k approximation, i.e., the term $\|A A^T - C C^T\|_\xi$. Although the additional error in Theorem 3 depends on $\|A A^T - C C^T\|_2$, we note that $\|A A^T - C C^T\|_2 \leq \|A A^T - C C^T\|_F$ and will use a bound for the latter quantity to bound the former in the following. Note that the prefactor of the additional error is $2\sqrt{k}$ for $\|\cdot\|_F^2$, while that for $\|\cdot\|_2^2$ is only 2.

In the following theorem we specialize the sampling probabilities to be those that are nearly optimal; by choosing enough columns, the error in the approximation of the SVD can be made arbitrarily small.

THEOREM 4. *Suppose $A \in \mathbb{R}^{m \times n}$; let H_k be constructed from the LINEAR-TIMESVD algorithm by sampling c columns of A with probabilities $\{p_i\}_{i=1}^n$ such that $p_i \geq \beta |A^{(i)}|^2 / \|A\|_F^2$ for some positive $\beta \leq 1$, and let $\eta = 1 + \sqrt{(8/\beta) \log(1/\delta)}$. Let $\epsilon > 0$. If $c \geq 4k/\beta\epsilon^2$, then*

$$(23) \quad \mathbf{E}[\|A - H_k H_k^T A\|_F^2] \leq \|A - A_k\|_F^2 + \epsilon \|A\|_F^2,$$

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

$$(24) \quad \|A - H_k H_k^T A\|_F^2 \leq \|A - A_k\|_F^2 + \epsilon \|A\|_F^2.$$

In addition, if $c \geq 4/\beta\epsilon^2$, then

$$(25) \quad \mathbf{E}[\|A - H_k H_k^T A\|_2^2] \leq \|A - A_k\|_2^2 + \epsilon \|A\|_F^2,$$

and if $c \geq 4\eta^2/\beta\epsilon^2$, then with probability at least $1 - \delta$,

$$(26) \quad \|A - H_k H_k^T A\|_2^2 \leq \|A - A_k\|_2^2 + \epsilon \|A\|_F^2.$$

Proof. By combining Theorems 2 and 3 with Theorem 1 we have that

$$(27) \quad \mathbf{E}[\|A - H_k H_k^T A\|_F^2] \leq \|A - A_k\|_F^2 + \left(\frac{4k}{\beta c}\right)^{1/2} \|A\|_F^2,$$

$$(28) \quad \mathbf{E}[\|A - H_k H_k^T A\|_2^2] \leq \|A - A_k\|_2^2 + \left(\frac{4}{\beta c}\right)^{1/2} \|A\|_F^2,$$

and that with probability at least $1 - \delta$,

$$(29) \quad \|A - H_k H_k^T A\|_F^2 \leq \|A - A_k\|_F^2 + \left(\frac{4\eta^2 k}{\beta c}\right)^{1/2} \|A\|_F^2,$$

$$(30) \quad \|A - H_k H_k^T A\|_2^2 \leq \|A - A_k\|_2^2 + \left(\frac{4\eta^2}{\beta c}\right)^{1/2} \|A\|_F^2.$$

The theorem follows by using the appropriate value of c . \square

Note that alternatively one could sample rows instead of columns of a matrix; in this case, a modified version of the LINEARTIMESVD algorithm leads to results analogous to Theorems 2 through 4.

5. Constant time SVD approximation algorithm.

5.1. The algorithm. Given a matrix $A \in \mathbb{R}^{m \times n}$ we now wish to approximate its top k singular values and the corresponding singular vectors in a constant number of passes through the data and additional space and time that are $O(1)$, independent of m and n . The strategy behind the CONSTANTTIMESVD algorithm is to pick c columns of the matrix A , rescale each by an appropriate factor to form a matrix $C \in \mathbb{R}^{m \times c}$, and then compute approximations to the singular values and left singular vectors of the matrix C , which will then be approximations to the singular values and left singular vectors of A . In the LINEARTIMESVD algorithm of section 4, the left singular vectors of the matrix C are computed exactly; as the analysis of section 4.2 showed, this computation takes additional space and time that is linear in $m + n$ (assuming that c is constant). With the CONSTANTTIMESVD algorithm, in order to use only a constant $O(1)$ additional space and time, sampling is performed again, drawing rows of C to construct a matrix $W \in \mathbb{R}^{w \times c}$. The SVD of $W^T W$ is then computed; let $W^T W = Z \Sigma_W Z^T = Z \Sigma_W^2 Z^T$. The singular values and corresponding singular vectors so obtained are with high probability approximations to the singular values and singular vectors of $C^T C$ and thus to the singular values and right singular vectors of C . Note that this is simply using the LINEARTIMESVD algorithm to approximate the right singular vectors of C by randomly sampling rows of C .

The CONSTANTTIMESVD algorithm is described in Figure 3; it takes as input a matrix A and returns as output a “description” of an approximation to the top k left singular values and the corresponding singular vectors. This “description” of the approximations to the left singular vectors of A may, at the expense of one additional pass and linear additional space and time, be converted into an explicit

CONSTANTTIMESVD Algorithm.

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

Output: $\sigma_t(W)$, $t = 1, \dots, \ell$ and a “description” of $\tilde{H}_\ell \in \mathbb{R}^{m \times \ell}$.

1. For $t = 1$ to c ,
 - (a) Pick $i_t \in 1, \dots, n$ with $\Pr[i_t = \alpha] = p_\alpha$, $\alpha = 1, \dots, n$, and save $\{(i_t, p_{j_t}) : t = 1, \dots, c\}$.
 - (b) Set $C^{(t)} = A^{(i_t)} / \sqrt{c p_{i_t}}$. (Note that C is not explicitly constructed in RAM.)
2. Choose $\{q_j\}_{j=1}^m$ such that $q_j = |C_{(j)}|^2 / \|C\|_F^2$.
3. For $t = 1$ to w ,
 - (a) Pick $j_t \in 1, \dots, m$ with $\Pr[j_t = \alpha] = q_\alpha$, $\alpha = 1, \dots, m$.
 - (b) Set $W_{(t)} = C_{(j_t)} / \sqrt{w q_{j_t}}$.
4. Compute $W^T W$ and its SVD. Say $W^T W = \sum_{t=1}^c \sigma_t^2(W) z^t z^{tT}$.
5. If a $\|\cdot\|_F$ bound is desired, set $\gamma = \epsilon/100k$,
Else if a $\|\cdot\|_2$ bound is desired, set $\gamma = \epsilon/100$.
6. Let $\ell = \min\{k, \max\{t : \sigma_t^2(W) \geq \gamma \|W\|_F^2\}\}$.
7. Return singular values $\{\sigma_t(W)\}_{t=1}^\ell$ and their corresponding singular vectors $\{z^t\}_{t=1}^\ell$.

FIG. 3. The CONSTANTTIMESVD algorithm.

approximation to the left singular vectors of A by using $C = \tilde{H} \Sigma_W Z^T$ to compute \tilde{H} , whose columns are approximations of the left singular vectors of C . Note that γ in the CONSTANTTIMESVD algorithm is introduced to bound small singular values of C that may be perturbed by the second level of sampling; as indicated, the particular value of γ that is chosen depends on the norm bound which is desired. Note also that the probabilities $\{q_j\}_{j=1}^m$ used in the algorithm are optimal (in the sense of section 3.3), as will be the probabilities $\{p_i\}_{i=1}^n$ which will enter into Theorem 5.

A diagram illustrating the action of the CONSTANTTIMESVD algorithm is presented in Figure 4. The transformation represented by the matrix A is represented along with its SVD, and the transformation represented by the matrix C is also shown (but note that its SVD is not shown). The transformation represented by the matrix W , which is constructed from C with the second level of sampling, is also shown along with its SVD. In addition, approximations to the right singular vectors of C and to the left singular vectors of C calculated from $C = \tilde{H} \Sigma_W Z^T$ are shown.

In section 5.2 we will show that this algorithm takes $O(1)$, i.e., a constant with respect to m and n , additional space and time, assuming that c and w are constant. In section 5.3 we will state Theorem 5, which will establish the correctness of the algorithm; this theorem is the main result of this section and is the analogue of Theorem 4. Finally, in section 5.4 we will prove Theorem 5.

5.2. Analysis of the implementation and running time. Assuming that optimal sampling probabilities (as defined in section 3.3) are used, then in the CONSTANTTIMESVD algorithm the sampling probabilities p_k can be used to select columns to be sampled in one pass and $O(c)$ additional space and time using the SELECT algorithm of [13]. Given the columns of A to be sampled, we do not explicitly construct the




FIG. 4. Diagram for the CONSTANTTIMESVD algorithm.

matrix C but instead perform a second level of sampling and select w rows of C with probabilities $\{q_i\}_{i=1}^m$ (as described in the CONSTANTTIMESVD algorithm) in order to construct the matrix W . We do this by performing a second pass and using $O(w)$ additional space and time, again using the SELECT algorithm. Then in a third pass we explicitly construct W ; this requires additional space and time that is $O(cw)$. Then, given W , computing $W^T W$ requires $O(cw)$ additional space and $O(c^2w)$ additional time, and computing the SVD of $W^T W$ requires $O(c^3)$ additional time. The singular values and corresponding singular vectors thus computed can then be returned as the “description” of the solution. The total additional time for the CONSTANTTIMESVD algorithm is then $O(c^3 + cw^2)$; this is a constant if c and w are assumed to be a constant. To explicitly compute \tilde{H}_k would require k matrix-vector multiplications which would require another pass over the data and $O(mck)$ additional space and time.

5.3. Statement of Theorem 5. This subsection and the next provide an analysis of the CONSTANTTIMESVD algorithm similar to the analysis of the LINEAR-TIMESVD algorithm found in section 4.3. Recall that in section 4 we were interested in bounding $\|A - H_k H_k^T A\|_\xi^2$, where $\xi = F, 2$. In that case, $H_k^T H_k = I_k$, $H_k H_k^T$ was an orthonormal projection, and $H_k H_k^T A$ was our rank at most k approximation to A . In the constant time model, we do not have access to H_k but instead to \tilde{H}_ℓ , where the columns of \tilde{H}_ℓ , i.e., $\tilde{h}^t = Cz^t/\sigma_t(W), t = 1, \dots, \ell$, do not form an orthonormal set. However, by Lemma 2 of section 5.4.1, if C and W are constructed by sampling with optimal probabilities, then with high probability the columns of \tilde{H}_ℓ are approximately orthonormal, $\tilde{H}_\ell^T \tilde{H}_\ell \approx I_\ell$, and $\tilde{H}_\ell \tilde{H}_\ell^T = \sum_{t=1}^\ell \tilde{h}^t \tilde{h}^{t^T}$ is approximately an orthonormal projection. Applying this to A , we will get our low-rank approximation. Note that in dealing with this nonorthogonality the original proof of [16] contained a small error which was corrected in the journal version [17].

In this section and the next we use the following notation. Recall that the SVD of $W^T W \in \mathbb{R}^{c \times c}$ is

$$(31) \quad W^T W = \sum_{t=1}^c \sigma_t^2(W) z^t z^{tT} = Z \Sigma_W^2 Z^T,$$

where $Z \in \mathbb{R}^{c \times c}$. 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 $W^T W$. Then

$$(32) \quad \tilde{H}_\ell = C Z_{1, \ell} T,$$

where $T \in \mathbb{R}^{\ell \times \ell}$ is the diagonal matrix with elements $T_{tt} = 1/\sigma_t(W)$. In addition, let the SVD of \tilde{H}_ℓ be

$$(33) \quad \tilde{H}_\ell = B_\ell \Sigma_{\tilde{H}_\ell} D_\ell^T,$$

and let us define the matrix $\Delta \in \mathbb{R}^{\ell \times \ell}$ to be

$$(34) \quad \Delta = T Z_{1, \ell}^T (C^T C - W^T W) Z_{1, \ell} T.$$

We will see that Δ is a measure of the degree to which the columns of \tilde{H}_ℓ are not orthonormal.

Theorem 5 is the constant time analogue of Theorem 4 and is the main result of this section. Note that since the results from sampling at the second step, i.e., sampling from the matrix C to form the matrix W , depend on the samples chosen in the first sampling step, we do not state the following results in expectation, but instead state them with high probability.

THEOREM 5. *Suppose $A \in \mathbb{R}^{m \times n}$; let a description of \tilde{H}_ℓ be constructed from the CONSTANTTIMESVD algorithm 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 $1 - \delta$,

$$(35) \quad \|A - \tilde{H}_\ell \tilde{H}_\ell^T A\|_F^2 \leq \|A - A_k\|_F^2 + \epsilon \|A\|_F^2.$$

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 $1 - \delta$,

$$(36) \quad \|A - \tilde{H}_\ell \tilde{H}_\ell^T A\|_2^2 \leq \|A - A_k\|_2^2 + \epsilon \|A\|_F^2.$$

Proof. See section 5.4 for the proof. \square

Recall that in section 4 we first proved Theorems 2 and 3, which provided a bound on $\|A - H_k H_k^T A\|_F^2$ and $\|A - H_k H_k^T A\|_2^2$, respectively, for arbitrary probabilities, and then we proved Theorem 4 for the nearly optimal probabilities. Although a similar presentation strategy could be adopted in this section, in the interests of simplicity (due to the technically more complicated proofs in the constant time model) we instead immediately restrict ourselves in Theorem 5 to the case of optimal sampling probabilities and defer the proofs of the supporting lemmas to section 5.4.

5.4. Proof of Theorem 5. In this section, we prove Theorem 5. We start in section 5.4.1 with several lemmas that are common to both the Frobenius and spectral norms. Then in section 5.4.2 we provide the proof of (35). Finally, in section 5.4.3 we provide the proof of (36).

5.4.1. General lemmas. In this section, we prove four lemmas that are used in the proofs of both the Frobenius and spectral norm results.

First, we relate $\|A - \tilde{H}_\ell \tilde{H}_\ell^T A\|_\xi^2$, for $\xi = 2, F$, to $\|A - B_\ell B_\ell^T A\|_\xi^2$ plus an error term; we do so since the columns of B_ℓ are orthonormal, which will allow us to bound $\|A - B_\ell B_\ell^T A\|_\xi^2$ using arguments similar to those used to bound $\|A - H_k H_k^T A\|_\xi^2$ in Theorems 2 and 3.

LEMMA 1. *For $\xi = 2, F$ and for any $\epsilon > 0$,*

$$\|A - \tilde{H}_\ell \tilde{H}_\ell^T A\|_\xi^2 \leq \left(1 + \frac{\epsilon}{100}\right) \|A - B_\ell B_\ell^T A\|_\xi^2 + \left(1 + \frac{100}{\epsilon}\right) \|B_\ell B_\ell^T - \tilde{H}_\ell \tilde{H}_\ell^T\|_\xi^2 \|A\|_\xi^2.$$

Proof. By subadditivity and submultiplicativity,

$$\|A - \tilde{H}_\ell \tilde{H}_\ell^T A\|_\xi^2 \leq (\|A - B_\ell B_\ell^T A\|_\xi + \|B_\ell B_\ell^T - \tilde{H}_\ell \tilde{H}_\ell^T\|_\xi \|A\|_\xi)^2.$$

The lemma follows since $(\alpha + \beta)^2 \leq (1 + \epsilon)\alpha^2 + (1 + 1/\epsilon)\beta^2$ for all $\epsilon \geq 0$. \square

Second, although the vectors $\tilde{h}^t = Cz^t/\sigma_t(W)$, $t = 1, \dots, \ell$, do not in general form an orthonormal set, one would expect from their construction that if the matrix $W^T W$ is close to the matrix $C^T C$, then with high probability they will be approximately orthonormal. Lemma 2 establishes that Δ , defined in (34), characterizes how far \tilde{H}_ℓ is from having orthonormal columns and shows that the error introduced due to this nonorthonormality is bounded by a simple function of γ and the error introduced at the second level of sampling.

LEMMA 2. *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 $\tilde{H}_\ell = CZ_{1,\ell}T$ and that $T^T Z_{1,\ell}^T W^T W Z_{1,\ell} T = I_\ell$, so that

$$(37) \quad \|\tilde{H}_\ell^T \tilde{H}_\ell - I_\ell\|_\xi = \|T^T Z_{1,\ell}^T C^T C Z_{1,\ell} T - T^T Z_{1,\ell}^T W^T W Z_{1,\ell} T\|_\xi$$

$$(38) \quad = \|T^T Z_{1,\ell}^T (C^T C - W^T W) Z_{1,\ell} T\|_\xi.$$

Using the submultiplicativity properties of the 2-norm, and in particular

$$(39) \quad \|AB\|_\xi \leq \|A\|_2 \|B\|_\xi,$$

$$(40) \quad \|AB\|_\xi \leq \|A\|_\xi \|B\|_2,$$

for both $\xi = 2, F$, we get

$$(41) \quad \|\tilde{H}_\ell^T \tilde{H}_\ell - I_\ell\|_\xi \leq \|T^T Z_{1,\ell}^T\|_2 \|C^T C - W^T W\|_\xi \|Z_{1,\ell} T\|_2$$

$$(42) \quad \leq \|T\|_2^2 \|C^T C - W^T W\|_\xi$$

$$(43) \quad \leq \max_{t=1,\dots,\ell} (1/\sigma_t^2(W)) \|C^T C - W^T W\|_\xi,$$

since $\|Z_{1,\ell}\|_2 = 1$. The lemma follows since $\sigma_t^2(W) \geq \gamma \|W\|_F^2$ for all $t = 1, \dots, \ell$ by the definition of ℓ . \square

Third, we consider the second term in Lemma 1, $\|B_\ell B_\ell^T - \tilde{H}_\ell \tilde{H}_\ell^T\|_\xi^2$ and show that it can be related to $\|\Delta\|_\xi$.

LEMMA 3. For $\xi = 2, F$

$$\|B_\ell B_\ell^T - \tilde{H}_\ell \tilde{H}_\ell^T\|_\xi = \|\Delta\|_\xi.$$

Proof. Since $\tilde{H}_\ell = B_\ell \Sigma_{\tilde{H}_\ell} D_\ell^T$, we have

$$\begin{aligned} \|B_\ell B_\ell^T - \tilde{H}_\ell \tilde{H}_\ell^T\|_\xi &= \|B_\ell (I_\ell - \Sigma_{\tilde{H}_\ell}^2) B_\ell^T\|_\xi \\ &= \|I_\ell - \Sigma_{\tilde{H}_\ell}^2\|_\xi \\ &= \|D_\ell (I_\ell - \Sigma_{\tilde{H}_\ell}^2) D_\ell^T\|_\xi \\ &= \|I_\ell - \tilde{H}_\ell^T \tilde{H}_\ell\|_\xi. \quad \square \end{aligned}$$

Fourth, Lemma 4 considers the special case in which the probabilities $\{p_i\}_{i=1}^n$ that are entered into the CONSTANTTIMESVD algorithm are optimal, as is the case for Theorem 5.

LEMMA 4. Let $A \in \mathbb{R}^{m \times n}$ and let \tilde{H}_ℓ be constructed from the CONSTANTTIMESVD algorithm 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 $p_i = \mathbf{Pr}[i_t = i] = |A^{(i)}|^2 / \|A\|_F^2$ and $q_j = \mathbf{Pr}[j_t = j] = |C_{(j)}|^2 / \|C\|_F^2$. Then

$$\|W\|_F = \|C\|_F = \|A\|_F.$$

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

5.4.2. Lemmas for the Frobenius norm proof. In this section we prove (35). We do this by first proving lemmas sufficient to bound $\|A - B_\ell B_\ell^T A\|_F^2$; when this is combined with the lemmas of section 5.4.1 we obtain a bound on $\|A - \tilde{H}_\ell \tilde{H}_\ell^T A\|_F^2$. The bound on $\|A - B_\ell B_\ell^T A\|_F^2$ depends on the error for the optimal rank- k approximation to A , i.e., $\|A - A_k\|_F^2$, and additional errors that depend on the quality of the sampling approximations, i.e., on $\|AA^T - CC^T\|_F$ and $\|C^T C - W^T W\|_F$. This will be the analogue of Theorem 2 applied to the constant additional space and time model. The result and associated proof will have a similar structure to that of Theorem 2, but will be more complicated due to the nonorthonormality of the vectors $\tilde{h}^t, t = 1, \dots, \ell$, and will involve additional error terms since two levels of approximation are involved.

We now prove several lemmas which will provide a bound for the first term in Lemma 1 when applied to the Frobenius norm. We first rewrite the $\|A - B_\ell B_\ell^T A\|_F^2$ term from Lemma 1. Note that Lemma 5 is the constant time analogue of (14).

LEMMA 5.

$$\|A - B_\ell B_\ell^T A\|_F^2 = \|A\|_F^2 - \|B_\ell^T A\|_F^2.$$

Proof.

$$\begin{aligned} \|A - B_\ell B_\ell^T A\|_F^2 &= \mathbf{Tr}((A - B_\ell B_\ell^T A)^T (A - B_\ell B_\ell^T A)) \\ &= \mathbf{Tr}(A^T A - A^T B_\ell B_\ell^T A). \quad \square \end{aligned}$$

Next, we want to provide a lower bound for $\|B_\ell^T A\|_F^2$ in terms of the singular values of W . We do so in several steps. First, we relate $\|B_\ell^T A\|_F^2$ to $\|\tilde{H}_\ell^T A\|_F^2$. We note that the assumption $\|\Delta\|_F < 1$ is made since in Theorem 5 optimal probabilities are used and sufficiently many columns and rows are drawn; if this assumption is dropped, then bounds of the form in Theorem 5 may be obtained with slightly worse sampling complexity.

LEMMA 6. *If $\|\Delta\|_F < 1$, then*

$$\|B_\ell^T A\|_F^2 \geq (1 - \|\Delta\|_F) \|\tilde{H}_\ell^T A\|_F^2.$$

Proof. Since $\tilde{H}_\ell = B_\ell \Sigma_{\tilde{H}_\ell} D_\ell^T$

$$(44) \quad \|\tilde{H}_\ell^T A\|_F^2 = \|\Sigma_{\tilde{H}_\ell} B_\ell^T A\|_F^2 \leq \|\Sigma_{\tilde{H}_\ell}\|_2^2 \|B_\ell^T A\|_F^2 = \|\tilde{H}_\ell^T \tilde{H}_\ell\|_2^2 \|B_\ell^T A\|_F^2,$$

using (39). From the triangle inequality

$$(45) \quad \|\tilde{H}_\ell^T \tilde{H}_\ell\|_2 \geq |\|I_\ell\|_2 - \|\tilde{H}_\ell^T \tilde{H}_\ell - I_\ell\|_2| = |1 - \|\Delta\|_2|.$$

The lemma follows since $\|\Delta\|_2 \leq \|\Delta\|_F < 1$ and by observing that $1 + x \leq 1/(1 - x)$ for all $x \leq 1$. \square

Second, we relate $\|\tilde{H}_\ell^T A\|_F^2$ to $\|\tilde{H}_\ell^T C\|_F^2$.

LEMMA 7.

$$\|\tilde{H}_\ell^T A\|_F^2 \geq \|\tilde{H}_\ell^T C\|_F^2 - (k + \sqrt{k} \|\Delta\|_F) \|AA^T - CC^T\|_F.$$

Proof. Since $\|\tilde{H}_\ell^T A\|_F^2 = \mathbf{Tr}(\tilde{H}_\ell^T AA^T \tilde{H}_\ell^T)$, we have that

$$\begin{aligned} \|\tilde{H}_\ell^T A\|_F^2 &= \mathbf{Tr}(\tilde{H}_\ell^T CC^T \tilde{H}_\ell^T) + \mathbf{Tr}(\tilde{H}_\ell^T (AA^T - CC^T) \tilde{H}_\ell^T) \\ &\geq \|\tilde{H}_\ell^T C\|_F^2 - \|AA^T - CC^T\|_2 \|\tilde{H}_\ell\|_F^2, \end{aligned}$$

where the inequality follows since

$$\begin{aligned} \left| \mathbf{Tr}(\tilde{H}_\ell^T (AA^T - CC^T) \tilde{H}_\ell^T) \right| &\leq \sum_t \left| (\tilde{H}_\ell^T)_{(t)} (AA^T - CC^T) (\tilde{H}_\ell)_{(t)} \right| \\ &\leq \|AA^T - CC^T\|_2 \|\tilde{H}_\ell\|_F^2. \end{aligned}$$

The lemma follows since $\|\cdot\|_2 \leq \|\cdot\|_F$ and since

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

Third, we relate $\|\tilde{H}_\ell^T C\|_F^2$ to $\sum_{t=1}^\ell \sigma_t^2(W)$.

LEMMA 8.

$$\left\| \tilde{H}_\ell^T C \right\|_F^2 \geq \sum_{t=1}^\ell \sigma_t^2(W) - \frac{2}{\sqrt{\gamma}} \|C^T C - W^T W\|_F.$$

Proof. Since $\|\tilde{H}_\ell^T C\|_F^2 = \|C^T \tilde{H}_\ell\|_F^2 = \|C^T C Z_{1,\ell} T\|_F^2$, we have

$$\begin{aligned} \left\| \tilde{H}_\ell^T C \right\|_F^2 &\geq \left(\|W^T W Z_{1,\ell} T\|_F - \|(C^T C - W^T W) Z_{1,\ell} T\|_F \right)^2 \\ &\geq \left(\left(\sum_{t=1}^\ell \sigma_t^2(W) \right)^{1/2} - \frac{1}{\sqrt{\gamma} \|W\|_F} \|(C^T C - W^T W)\|_F \right)^2, \end{aligned}$$

where the second inequality uses that $\|XZ\|_F \leq \|X\|_F$ for any matrix X if the matrix Z has orthonormal columns. By multiplying out the right-hand side and ignoring terms that reinforce the inequality, the lemma follows since $(\sum_{t=1}^\ell \sigma_t^2(W))^{1/2} / \|W\|_F \leq 1$. \square

By combining Lemmas 6, 7, and 8, we have our desired bound on $\|B_\ell^T A\|_F^2$ in terms of the singular values of W . Finally, we use matrix perturbation theory to relate $\sum_{t=1}^\ell \sigma_t^2(W)$ to $\sum_{t=1}^k \sigma_t^2(A)$.

LEMMA 9.

$$\sum_{t=1}^\ell \sigma_t^2(W) \geq \sum_{t=1}^k \sigma_t^2(A) - \sqrt{k} \|AA^T - CC^T\|_F - \sqrt{k} \|C^T C - W^T W\|_F - (k-\ell)\gamma \|W\|_F^2.$$

Proof. Recalling the Hoffman–Wielandt inequality, we see that

$$\begin{aligned} \left| \sum_{t=1}^k (\sigma_t^2(C) - \sigma_t^2(A)) \right| &\leq \sqrt{k} \left(\sum_{t=1}^k (\sigma_t^2(C) - \sigma_t^2(A))^2 \right)^{1/2} \\ &\leq \sqrt{k} \left(\sum_{t=1}^k (\sigma_t(CC^T) - \sigma_t(AA^T))^2 \right)^{1/2} \\ (46) \quad &\leq \sqrt{k} \|AA^T - CC^T\|_F, \end{aligned}$$

and, similarly, that

$$\begin{aligned} \left| \sum_{t=1}^k (\sigma_t^2(W) - \sigma_t^2(C)) \right| &\leq \sqrt{k} \left(\sum_{t=1}^k (\sigma_t^2(W) - \sigma_t^2(C))^2 \right)^{1/2} \\ &\leq \sqrt{k} \left(\sum_{t=1}^k (\sigma_t(WW^T) - \sigma_t(CC^T))^2 \right)^{1/2} \\ (47) \quad &\leq \sqrt{k} \|C^T C - W^T W\|_F. \end{aligned}$$

By combining (46) and (47) we see that

$$(48) \quad \left| \sum_{t=1}^k \sigma_t^2(W) - \sum_{t=1}^k \sigma_t^2(A) \right| \leq \sqrt{k} \|AA^T - CC^T\|_F + \sqrt{k} \|C^T C - W^T W\|_F.$$

Since $\sigma_t^2(W) < \gamma \|W\|_F^2$ for all $t = \ell + 1, \dots, k$ we have that $\sum_{t=\ell+1}^k \sigma_t^2(W) \leq (k - \ell)\gamma \|W\|_F^2$. Combining this with (48) allows us to relate $\sum_{t=1}^\ell \sigma_t^2(W)$ and $\sum_{t=1}^k \sigma_t^2(A)$, thus establishing the lemma. \square

Now we combine these results in order to prove (35). Let $E_{AA^T} = AA^T - CC^T$ and $E_{C^T C} = C^T C - W^T W$. First, we establish a lower bound on $\|B_\ell^T A\|_F^2$. By combining Lemmas 6 and 7 and dropping terms that reinforce the inequality, we have that

$$\|B_\ell^T A\|_F^2 \geq \|\tilde{H}_\ell^T C\|_F^2 - \|\Delta\|_F \|\tilde{H}_\ell^T C\|_F^2 - (k + \sqrt{k} \|\Delta\|_F) \|E_{AA^T}\|_F.$$

By combining this with Lemmas 8 and 9 and dropping terms that reinforce the inequality, we have that

$$(49) \quad \begin{aligned} \|B_\ell^T A\|_F^2 &\geq \sum_{t=1}^k \sigma_t^2(A) - (k + \sqrt{k}) \|E_{AA^T}\|_F - \left(\sqrt{k} + \frac{2}{\sqrt{\gamma}}\right) \|E_{C^T C}\|_F \\ &\quad - \|\Delta\|_F \sum_{t=1}^k \sigma_t^2(A) - \sqrt{k} \|\Delta\|_F \|E_{AA^T}\|_F - (k - \ell)\gamma \|W\|_F^2. \end{aligned}$$

From Lemma 5 this immediately leads to the upper bound on $\|A - B_\ell B_\ell^T A\|_F^2$,

$$(50) \quad \begin{aligned} \|A - B_\ell B_\ell^T A\|_F^2 &\leq \|A - A_k\|_F^2 + (k + \sqrt{k}) \|E_{AA^T}\|_F + \left(\sqrt{k} + \frac{2}{\sqrt{\gamma}}\right) \|E_{C^T C}\|_F \\ &\quad + \|\Delta\|_F \sum_{t=1}^k \sigma_t^2(A) + \sqrt{k} \|\Delta\|_F \|E_{AA^T}\|_F + (k - \ell)\gamma \|W\|_F^2. \end{aligned}$$

From Lemmas 1 and 3,

$$(51) \quad \|A - \tilde{H}_\ell \tilde{H}_\ell^T A\|_F^2 \leq \left(1 + \frac{\epsilon}{100}\right) \|A - B_\ell B_\ell^T A\|_F^2 + \left(1 + \frac{100}{\epsilon}\right) \|\Delta\|_F^2 \|A\|_F^2.$$

Recall that $\gamma = \epsilon/100k$, that $\sum_{t=1}^k \sigma_t^2(A) \leq \|A\|_F^2$, that $\|\Delta\|_F \leq \|E_{C^T C}\|_F / \gamma \|W\|_F^2$ by Lemma 2, and that $\|W\|_F = \|C\|_F = \|A\|_F$ by Lemma 4; (35) then follows by combining (50) and (51), using the sampling probabilities indicated in the statement of the theorem, and by choosing $c, w = \Omega(k^2 \eta^2 / \epsilon^4)$.

5.4.3. Lemmas for the spectral norm proof. In this section we prove (36). We do this by first proving lemmas sufficient to bound $\|A - B_\ell B_\ell^T A\|_2^2$; when this is combined with the lemmas of section 5.4.1, we obtain a bound on $\|A - \tilde{H}_\ell \tilde{H}_\ell^T A\|_2^2$. The bound on $\|A - B_\ell B_\ell^T A\|_2^2$ depends on the error for the optimal rank- k approximation to A , i.e., $\|A - A_k\|_2^2$, and additional errors that depend on the quality of the sampling approximations, i.e., on $\|AA^T - CC^T\|_2$ and $\|C^T C - W^T W\|_2$. This will be the analogue of Theorem 3 applied to the constant additional space and time model. The result and associated proof will have a similar structure to that of Theorem 3, but will be more complicated due to the nonorthonormality of the vectors $\tilde{h}^t, t = 1, \dots, \ell$, and will involve additional error terms since two levels of approximation are involved.

We now prove three lemmas which will provide a bound for the first term in Lemma 1 when applied to the spectral norm. We first rewrite the $\|A - B_\ell B_\ell^T A\|_2^2$ term from Lemma 1.

LEMMA 10.

$$\|A - B_\ell B_\ell^T A\|_2^2 \leq \|Z_{k+1,c}^T C^T C Z_{k+1,c}\|_2 + \|Z_{\ell+1,k}^T C^T C Z_{\ell+1,k}\|_2 + \|AA^T - CC^T\|_2.$$

Proof. In order to bound $\|A - B_\ell B_\ell^T A\|_2$ we will project onto the subspace spanned by B_ℓ and its orthogonal complement in a manner analogous to that used in the proof of Theorem 3. Let $\mathcal{B}_\ell = \text{range}(B_\ell)$ and let $\mathcal{B}_{m-\ell}$ be the orthogonal complement of \mathcal{B}_ℓ . Let $x = \alpha y + \beta z$, where $y \in \mathcal{B}_\ell$, $z \in \mathcal{B}_{m-\ell}$, and $\alpha^2 + \beta^2 = 1$. Then

$$\begin{aligned} \|A - B_\ell B_\ell^T A\|_2 &= \max_{x \in \mathbb{R}^m, |x|=1} |x^T (A - B_\ell B_\ell^T A)| \\ &= \max_{y \in \mathcal{B}_\ell, |y|=1, z \in \mathcal{B}_{m-\ell}, |z|=1, \alpha^2 + \beta^2 = 1} |(\alpha y^T + \beta z^T)(A - B_\ell B_\ell^T A)| \\ (52) \quad &\leq \max_{y \in \mathcal{B}_\ell, |y|=1} |y^T (A - B_\ell B_\ell^T A)| + \max_{z \in \mathcal{B}_{m-\ell}, |z|=1} |z^T (A - B_\ell B_\ell^T A)| \\ (53) \quad &= \max_{z \in \mathcal{B}_{m-\ell}, |z|=1} |z^T A|. \end{aligned}$$

Inequality (52) follows since $\alpha, \beta \leq 1$ and (53) follows since $y \in \mathcal{B}_\ell$ and $z \in \mathcal{B}_{m-\ell}$. To bound (53), let $z \in \mathcal{B}_{m-\ell}$, $|z| = 1$; then

$$\begin{aligned} |z^T A|^2 &= z^T (AA^T) z \\ &= z^T (CC^T) z + z^T (AA^T - CC^T) z \\ (54) \quad &= z^T (CC^T - CZ_{1,k} Z_{1,k}^T C^T) z + z^T (CZ_{1,k} Z_{1,k}^T C^T) z + z^T (AA^T - CC^T) z \end{aligned}$$

$$(55) \quad = z^T (CZ_{k+1,c} Z_{k+1,c}^T C^T) z + z^T (CZ_{\ell+1,k} Z_{\ell+1,k}^T C^T) z + z^T (AA^T - CC^T) z.$$

Equation (55) follows since $I_c = ZZ^T = Z_{1,k} Z_{1,k}^T + Z_{k+1,c} Z_{k+1,c}^T$ and, since

$$CZ_{1,\ell} Z_{1,\ell}^T C^T = \sum_{t=1}^{\ell} C z^t z^{t^T} C^T = \sum_{t=1}^{\ell} \sigma_t^2(W) \tilde{h}^t \tilde{h}^{t^T},$$

implies that

$$(56) \quad z^T CZ_{1,\ell} Z_{1,\ell}^T C^T z = 0$$

for $z \in \mathcal{B}_{m-\ell}$. Thus, by combining (53) and (55)

$$\|A - B_\ell B_\ell^T A\|_2^2 \leq \|CZ_{k+1,c} Z_{k+1,c}^T C^T\|_2 + \|CZ_{\ell+1,k} Z_{\ell+1,k}^T C^T\|_2 + \|AA^T - CC^T\|_2.$$

The lemma follows since $\|X^T X\|_2 = \|X X^T\|_2$ for any matrix X . \square

We next bound the $\|Z_{k+1,c}^T C^T C Z_{k+1,c}\|_2$ term from Lemma 10; note that matrix perturbation theory is used in (59).

LEMMA 11.

$$\|Z_{k+1,c}^T C^T C Z_{k+1,c}\|_2 \leq \|A - A_k\|_2^2 + \|AA^T - CC^T\|_2 + 2\|C^T C - W^T W\|_2.$$

Proof. First note that

$$(57) \quad \|Z_{k+1,c}^T C^T C Z_{k+1,c}\|_2 \leq \|Z_{k+1,c}^T W^T W Z_{k+1,c}\|_2 + \|Z_{k+1,c}^T (C^T C - W^T W) Z_{k+1,c}\|_2.$$

Since

$$\begin{aligned} \|Z_{k+1,c}^T (C^T C - W^T W) Z_{k+1,c}\|_2 &\leq \|C^T C - W^T W\|_2 \|Z_{k+1,c}\|_2^2 \\ &= \|C^T C - W^T W\|_2 \end{aligned}$$

and

$$\|Z_{k+1,c}^T W^T W Z_{k+1,c}\|_2 = \sigma_{k+1}^2(W),$$

it follows from (57) that

$$(58) \quad \|Z_{k+1,c}^T C^T C Z_{k+1,c}\|_2 \leq \sigma_{k+1}^2(W) + \|C^T C - W^T W\|_2.$$

By a double application of (9), we see that

$$(59) \quad \sigma_{k+1}^2(W) \leq \sigma_{k+1}^2(A) + \|AA^T - CC^T\|_2 + \|C^T C - W^T W\|_2.$$

The lemma follows by combining (58) and (59) since $\|A - A_k\|_2 = \sigma_{k+1}(A)$. \square

Finally, we bound the $\|Z_{\ell+1,k}^T C^T C Z_{\ell+1,k}\|_2$ term from Lemma 10; note that if $\ell = k$, it is unnecessary.

LEMMA 12.

$$\|Z_{\ell+1,k}^T C^T C Z_{\ell+1,k}\|_2 \leq \|C^T C - W^T W\|_2 + \gamma \|W\|_F^2.$$

Proof. First note that

$$(60) \quad \|Z_{\ell+1,k}^T C^T C Z_{\ell+1,k}\|_2 \leq \|Z_{\ell+1,k}^T (C^T C - W^T W) Z_{\ell+1,k}\|_2 + \|Z_{\ell+1,k}^T W^T W Z_{\ell+1,k}\|_2.$$

Since

$$\begin{aligned} \|Z_{\ell+1,k}^T (C^T C - W^T W) Z_{\ell+1,k}\|_2 &\leq \|C^T C - W^T W\|_2 \|Z_{\ell+1,k}\|_2^2 \\ &= \|C^T C - W^T W\|_2 \end{aligned}$$

and

$$\|Z_{\ell+1,k}^T W^T W Z_{\ell+1,k}\|_2 = \sigma_{\ell+1}^2(W),$$

it follows from (60) that

$$(61) \quad \|Z_{\ell+1,k}^T C^T C Z_{\ell+1,k}\|_2 \leq \|C^T C - W^T W\|_2 + \sigma_{\ell+1}^2(W).$$

The lemma follows since $\sigma_t^2(W) < \gamma \|W\|_F^2$ for all $t = \ell + 1, \dots, k$. \square

Now we combine these results in order to prove (36). Recall that $E_{AA^T} = AA^T - CC^T$ and $E_{C^TC} = C^TC - W^TW$. By combining Lemmas 10, 11, and 12, we have that

$$(62) \quad \|A - B_\ell B_\ell^T A\|_2^2 \leq \|A - A_k\|_2^2 + 2\|E_{AA^T}\|_2 + 3\|E_{C^TC}\|_2 + \gamma\|W\|_F^2.$$

From Lemmas 1 and 3,

$$(63) \quad \|A - \tilde{H}_\ell \tilde{H}_\ell^T A\|_2^2 \leq \left(1 + \frac{\epsilon}{100}\right) \|A - B_\ell B_\ell^T A\|_2^2 + \left(1 + \frac{100}{\epsilon}\right) \|\Delta\|_2^2 \|A\|_2^2.$$

Recall that $\gamma = \epsilon/100$, that $\|\cdot\|_2 \leq \|\cdot\|_F$, and that $\|\Delta\|_2 \leq \|E_{C^TC}\|_2 / \gamma \|W\|_F^2$ by Lemma 2; (36) follows by combining (62) and (63), using the sampling probabilities indicated in the statement of the theorem, and by choosing $c, w = \Omega(\eta^2/\epsilon^4)$.

6. Discussion and conclusion. We have presented two algorithms to compute approximations to the SVD of a matrix $A \in \mathbb{R}^{m \times n}$ which do not require that A be stored in RAM, but for which the additional space and time required (in addition to a constant number of passes over the matrix) is either linear in $m + n$ or is a constant independent of m and n ; we have also proven error bounds for both algorithms with respect to both the Frobenius and spectral norms. Table 1 in section 1 presents a summary of the dependence of the sampling complexity on k and ϵ . With the LINEARTIMESVD algorithm, the additional error (beyond the optimal rank- k approximation) in the spectral norm bound can be made less than $\epsilon \|A\|_F^2$ by sampling $\Theta(1/\epsilon^2)$ columns, and the additional error in the Frobenius norm can be made less than $\epsilon \|A\|_F^2$ by sampling $\Theta(k/\epsilon^2)$ columns. Likewise, with the CONSTANTTIMESVD algorithm, the additional error in the spectral norm can be made less than $\epsilon \|A\|_F^2$ by sampling $\Theta(1/\epsilon^4)$ columns and rows, and the additional error in the Frobenius norm can be made less than $\epsilon \|A\|_F^2$ by sampling $\Theta(k^2/\epsilon^4)$ columns and rows. The results of [16] require $\Theta(k^4/\epsilon^3)$ columns and rows for the Frobenius (and thus the spectral) norm bound.

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 of error [4, 5]. In [5] these methods have been applied to the low-rank matrix approximation problem (defined as approximating the SVD with respect to the Frobenius norm) and to the matrix reconstruction problem. It is shown that any sampling algorithm that with high probability finds a good low-rank approximation requires $\Omega(m + n)$ queries. In addition, it is shown that even if the algorithm is given the exact weight distribution over the columns of a matrix, it will still require $\Omega(k/\epsilon^2)$ column queries to approximate A . Thus, the LINEARTIMESVD algorithm (see also the original [10]) is optimal with respect to Frobenius norm bounds for the rank parameter k and the CONSTANTTIMESVD algorithm (see also the original [16]) is optimal with respect to Frobenius norm bounds up to polynomial factors.

Acknowledgments. We would like to thank the following individuals for comments and fruitful discussions: Dimitris Achlioptas, Alan Frieze, Mauro Maggioni, Frank McSherry, and Santosh Vempala. We would also like to thank Dimitris Achlioptas and Frank McSherry for providing us with a preprint of [1], i.e., the journal version of [2], which provides a useful comparison of their results with ours. We would like to

thank the National Science Foundation for partial support of this work. Finally, we would like to thank an anonymous reviewer for carefully reading the paper and making numerous useful suggestions; in particular, the reviewer provided elegant, short proofs for Lemmas 2 and 6.

REFERENCES

- [1] D. ACHLIOPTAS AND F. MCSHERRY, *Fast computation of low rank matrix approximations*, J. ACM, to appear.
- [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, 2001, pp. 611–618.
- [3] O. ALTER, P. O. BROWN, AND D. BOTSTEIN, *Singular value decomposition for genome-wide expression data processing and modeling*, Proc. Natl. Acad. Sci. USA, 97 (2000), pp. 10101–10106.
- [4] Z. BAR-YOSSEF, *The Complexity of Massive Data Set Computations*, Ph.D. thesis, University of California, Berkeley, 2002.
- [5] Z. BAR-YOSSEF, *Sampling lower bounds via information theory*, in Proceedings of the 35th Annual ACM Symposium on Theory of Computing, 2003, pp. 335–344.
- [6] M. W. BERRY, Z. DRMAČ, AND E. R. JESSUP, *Matrices, vector spaces, and information retrieval*, SIAM Rev., 41 (1999), pp. 335–362.
- [7] M. W. BERRY, S. T. DUMAIS, AND G. W. O'BRIAN, *Using linear algebra for intelligent information retrieval*, SIAM Rev., 37 (1995), pp. 573–595.
- [8] R. BHATIA, *Matrix Analysis*, Springer-Verlag, New York, 1997.
- [9] S. T. DEERWESTER, S. T. DUMAIS, G. W. FURNAS, T. K. LANDAUER, AND R. HARSHMAN, *Indexing by latent semantic analysis*, J. Amer. Soc. Inform. Sci., 41 (1990), pp. 391–407.
- [10] 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, SIAM, Philadelphia, 1999, pp. 291–299.
- [11] P. DRINEAS AND R. KANNAN, *Fast Monte-Carlo algorithms for approximate matrix multiplication*, in Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science, 2001, pp. 452–459.
- [12] P. DRINEAS AND R. KANNAN, *Pass efficient algorithms for approximating large matrices*, in Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, 2003, pp. 223–232.
- [13] P. DRINEAS, R. KANNAN, AND M. W. MAHONEY, *Fast Monte Carlo algorithms for matrices I: Approximating matrix multiplication*, SIAM J. Comput., 36 (2006), pp. 132–157.
- [14] P. DRINEAS, R. KANNAN, AND M. W. MAHONEY, *Fast Monte Carlo algorithms for matrices III: Computing a compressed approximate matrix decomposition*, SIAM J. Comput., 36 (2006), pp. 184–206.
- [15] P. DRINEAS, R. KANNAN, AND M. W. MAHONEY, *Fast Monte Carlo Algorithms for Matrices II: Computing a Low-Rank Approximation to a Matrix*, Tech. Report YALEU/DCS/TR-1270, Department of Computer Science, Yale University, New Haven, CT, 2004.
- [16] 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, 1998, pp. 370–378.
- [17] A. FRIEZE, R. KANNAN, AND S. VEMPALA, *Fast Monte-Carlo algorithms for finding low-rank approximations*, J. ACM, 51 (2004), pp. 1025–1041.
- [18] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, Johns Hopkins University Press, Baltimore, MD, 1989.
- [19] R. A. HORN AND C. R. JOHNSON, *Matrix Analysis*, Cambridge University Press, New York, 1985.
- [20] P. INDYK, *Stable distributions, pseudorandom generators, embeddings and data stream computation*, in Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science, 2000, pp. 189–197.
- [21] J. KLEINBERG, *Two algorithms for nearest-neighbor search in high dimensions*, in Proceedings of the 29th Annual ACM Symposium on Theory of Computing, 1997, pp. 599–608.
- [22] J. KLEINBERG, *Authoritative sources in a hyperlinked environment*, in Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms, 1998, pp. 668–677.
- [23] K. V. MARDIA, J. T. KENT, AND J. M. BIBBY, *Multivariate Analysis*, Academic Press, London,

- 1979.
- [24] H. MURASE AND S. K. NAYAR, *Visual learning and recognition of 3-d objects from appearance*, *Internat. J. Comput. Vision*, 14 (1995), pp. 5–24.
 - [25] C. H. PAPADIMITRIOU, P. RAGHAVAN, H. TAMAKI, AND S. VEMPALA, *Latent semantic indexing: A probabilistic analysis*, in *Proceedings of the 17th ACM Symposium on Principles of Database Systems*, 1998, pp. 159–168.
 - [26] S. RAYCHAUDHURI, J. M. STUART, AND R. B. ALTMAN, *Principal components analysis to summarize microarray experiments: Application to sporulation time series*, in *Proceedings of the Pacific Symposium on Biocomputing 2000*, 2000, pp. 455–466.
 - [27] G. W. STEWART AND J. G. SUN, *Matrix Perturbation Theory*, Academic Press, New York, 1990.
 - [28] O. TROYANSKAYA, M. CANTOR, G. SHERLOCK, P. BROWN, T. HASTIE, R. TIBSHIRANI, D. BOSTEIN, AND R. B. ALTMAN, *Missing value estimation methods for DNA microarrays*, *Bioinformatics*, 17 (2001), pp. 520–525.
 - [29] M. TURK AND A. PENTLAND, *Eigenfaces for recognition*, *J. Cognitive Neurosci.*, 3 (1991), pp. 71–96.
 - [30] S. VEMPALA, *Random projection: A new approach to VLSI layout*, in *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science*, 1998, pp. 389–395.