# Fast Approximation of Matrix Coherence and Statistical Leverage 

Petros Drineas<br>Malik Magdon-Ismail<br>DRINEP@CS.RPI.EDU<br>MAGDON@CS.RPI.EDU<br>Department of Computer Science<br>Rensselaer Polytechnic Institute<br>Troy, NY 12180

Michael W. Mahoney<br>MMAHONEY@CS.STANFORD.EDU<br>Department of Mathematics<br>Stanford University<br>Stanford, CA 94305

David P. Woodruff<br>DPWOODRU@US.IBM.COM<br>IBM Almaden Research Center<br>650 Harry Road<br>San Jose, CA 95120

Editor: Mehryar Mohri


#### Abstract

The statistical leverage scores of a matrix $A$ are the squared row-norms of the matrix containing its (top) left singular vectors and the coherence is the largest leverage score. These quantities are of interest in recently-popular problems such as matrix completion and Nyström-based low-rank matrix approximation as well as in large-scale statistical data analysis applications more generally; moreover, they are of interest since they define the key structural nonuniformity that must be dealt with in developing fast randomized matrix algorithms. Our main result is a randomized algorithm that takes as input an arbitrary $n \times d$ matrix $A$, with $n \gg d$, and that returns as output relative-error approximations to all $n$ of the statistical leverage scores. The proposed algorithm runs (under assumptions on the precise values of $n$ and $d$ ) in $O(n d \log n)$ time, as opposed to the $O\left(n d^{2}\right)$ time required by the naïve algorithm that involves computing an orthogonal basis for the range of $A$. Our analysis may be viewed in terms of computing a relative-error approximation to an underconstrained least-squares approximation problem, or, relatedly, it may be viewed as an application of Johnson-Lindenstrauss type ideas. Several practically-important extensions of our basic result are also described, including the approximation of so-called cross-leverage scores, the extension of these ideas to matrices with $n \approx d$, and the extension to streaming environments.


Keywords: matrix coherence, statistical leverage, randomized algorithm

## 1. Introduction

The concept of statistical leverage measures the extent to which the singular vectors of a matrix are correlated with the standard basis and as such it has found usefulness recently in large-scale data analysis and in the analysis of randomized matrix algorithms (Velleman and Welsch, 1981; Mahoney and Drineas, 2009; Drineas et al., 2008). A related notion is that of matrix coherence, which has been of interest in recently popular problems such as matrix completion and Nyström-based lowrank matrix approximation (Candes and Recht, 2008; Talwalkar and Rostamizadeh, 2010). Defined
more precisely below, the statistical leverage scores may be computed as the squared Euclidean norms of the rows of the matrix containing the top left singular vectors and the coherence of the matrix is the largest statistical leverage score. Statistical leverage scores have a long history in statistical data analysis, where they have been used for outlier detection in regression diagnostics (Hoaglin and Welsch, 1978; Chatterjee and Hadi, 1986). Statistical leverage scores have also proved crucial recently in the development of improved worst-case randomized matrix algorithms that are also amenable to high-quality numerical implementation and that are useful to domain scientists (Drineas et al., 2008; Mahoney and Drineas, 2009; Boutsidis et al., 2009; Drineas et al., 2006b; Sarlós, 2006; Drineas et al., 2010b); see Mahoney (2011) for a detailed discussion. The naïve and best previously existing algorithm to compute these scores would compute an orthogonal basis for the dominant part of the spectrum of $A$, for example, the basis provided by the Singular Value Decomposition (SVD) or a basis provided by a QR decomposition (Golub and Loan, 1996), and then use that basis to compute diagonal elements of the projection matrix onto the span of that basis.

We present a randomized algorithm to compute relative-error approximations to every statistical leverage score in time qualitatively faster than the time required to compute an orthogonal basis. For the case of an arbitrary $n \times d$ matrix $A$, with $n \gg d$, our main algorithm runs (under assumptions on the precise values of $n$ and $d$, see Theorem 2 for an exact statement) in $O\left(n d \log n / \varepsilon^{2}\right)$ time, as opposed to the $\Theta\left(n d^{2}\right)$ time required by the naïve algorithm. As a corollary, our algorithm provides a relative-error approximation to the coherence of an arbitrary matrix in the same time. In addition, several practically-important extensions of the basic idea underlying our main algorithm are also described in this paper.

### 1.1 Overview and Definitions

We start with the following definition of the statistical leverage scores of a matrix.
Definition 1 Given an arbitrary $n \times d$ matrix $A$, with $n>d$, let $U$ denote the $n \times d$ matrix consisting of the $d$ left singular vectors of $A$, and let $U_{(i)}$ denote the $i$-th row of the matrix $U$ as a row vector. Then, the statistical leverage scores of the rows of $A$ are given by

$$
\ell_{i}=\left\|U_{(i)}\right\|_{2}^{2}
$$

for $i \in\{1, \ldots, n\}$; the coherence $\gamma$ of the rows of $A$ is

$$
\gamma=\max _{i \in\{1, \ldots, n\}} \ell_{i},
$$

that is, it is the largest statistical leverage score of $A$; and the $(i, j)$-cross-leverage scores $c_{i j}$ are

$$
c_{i j}=\left\langle U_{(i)}, U_{(j)}\right\rangle,
$$

that is, they are the dot products between the $i^{\text {th }}$ row and the $j^{\text {th }}$ row of $U$.
Although we have defined these quantities in terms of a particular basis, they clearly do not depend on that particular basis, but only on the space spanned by that basis. To see this, let $P_{A}$ denote the projection matrix onto the span of the columns of $A$. Then,

$$
\begin{equation*}
\ell_{i}=\left\|U_{(i)}\right\|_{2}^{2}=\left(U U^{T}\right)_{i i}=\left(P_{A}\right)_{i i} . \tag{1}
\end{equation*}
$$

That is, the statistical leverage scores of a matrix $A$ are equal to the diagonal elements of the projection matrix onto the span of its columns. ${ }^{1}$ Similarly, the $(i, j)$-cross-leverage scores are equal to the off-diagonal elements of this projection matrix, that is,

$$
\begin{equation*}
c_{i j}=\left(P_{A}\right)_{i j}=\left\langle U_{(i)}, U_{(j)}\right\rangle . \tag{2}
\end{equation*}
$$

Clearly, $O\left(n d^{2}\right)$ time suffices to compute all the statistical leverage scores exactly: simply perform the SVD or compute a QR decomposition of $A$ in order to obtain any orthogonal basis for the range of $A$ and then compute the Euclidean norm of the rows of the resulting matrix. Thus, in this paper, we are interested in algorithms that run in $o\left(n d^{2}\right)$ time.

Several additional comments are worth making regarding this definition. First, since $\sum_{i=1}^{n} \ell_{i}=$ $\|U\|_{F}^{2}=d$, we can define a probability distribution over the rows of $A$ as $p_{i}=\ell_{i} / d$. As discussed below, these probabilities have played an important role in recent work on randomized matrix algorithms and an important algorithmic question is the degree to which they are uniform or nonuniform. ${ }^{2}$ Second, one could also define leverage scores for the columns of a "tall" matrix $A$, but clearly those are all equal to one unless $n<d$ or $A$ is rank-deficient. Third, and more generally, given a rank parameter $k$, one can define the statistical leverage scores relative to the best rank-k approximation to $A$ to be the $n$ diagonal elements of the projection matrix onto the span of $A_{k}$, the best rank- $k$ approximation to $A$.

### 1.2 Our Main Result

Our main result is a randomized algorithm for computing relative-error approximations to every statistical leverage score, as well as an additive-error approximation to all of the large cross-leverage scores, of an arbitrary $n \times d$ matrix, with $n \gg d$, in time qualitatively faster than the time required to compute an orthogonal basis for the range of that matrix. Our main algorithm for computing approximations to the statistical leverage scores (see Algorithm 1 in Section 3) will amount to constructing a "randomized sketch" of the input matrix and then computing the Euclidean norms of the rows of that sketch. This sketch can also be used to compute approximations to the large cross-leverage scores (see Algorithm 2 of Section 3).

The following theorem provides our main quality-of-approximation and running time result for Algorithm 1.

Theorem 2 Let A be a full-rank $n \times d$ matrix, with $n \gg d$; let $\varepsilon \in(0,1 / 2]$ be an error parameter; and recall the definition of the statistical leverage scores $\ell_{i}$ from Definition 1. Then, there exists a randomized algorithm (Algorithm 1 of Section 3 below) that returns values $\tilde{\ell}_{\text {, }}$, for all $i \in\{1, \ldots, n\}$, such that with probability at least 0.8 ,

$$
\left|\ell_{i}-\tilde{\ell}_{i}\right| \leq \varepsilon \ell_{i}
$$

1. In this paper, for simplicity of exposition, we consider the case that the matrix $A$ has rank equal to $d$, that is, has full column rank. Theoretically, the extension to rank-deficient matrices $A$ is straightforward-simply modify Definition 1 and thus Equations (1) and (2) to let $U$ be any orthogonal matrix (clearly, with fewer than $d$ columns) spanning the column space of $A$. From a numerical perspective, things are substantially more subtle, and we leave this for future work that considers numerical implementations of our algorithms.
2. Observe that if $U$ consists of $d$ columns from the identity, then the leverage scores are extremely nonuniform: $d$ of them are equal to one and the remainder are equal to zero. On the other hand, if $U$ consists of $d$ columns from a normalized Hadamard transform (see Section 2.3 for a definition), then the leverage scores are very uniform: all $n$ of them are equal to $d / n$.
holds for all $i \in\{1, \ldots, n\}$. Assuming $d \leq n \leq e^{d}$, the running time of the algorithm is

$$
O\left(n d \ln \left(d \varepsilon^{-1}\right)+n d \varepsilon^{-2} \ln n+d^{3} \varepsilon^{-2}(\ln n)\left(\ln \left(d \varepsilon^{-1}\right)\right)\right) .
$$

Algorithm 1 provides a relative-error approximation to all of the statistical leverage scores $\ell_{i}$ of $A$ and, assuming $d \ln d=o\left(\frac{n}{\ln n}\right), \ln n=o(d)$, and treating $\varepsilon$ as a constant, its running time is $o\left(n d^{2}\right)$, as desired. As a corollary, the largest leverage score (and thus the coherence) is approximated to relative-error in $o\left(n d^{2}\right)$ time.

The following theorem provides our main quality-of-approximation and running time result for Algorithm 2.

Theorem 3 Let A be a full-rank $n \times d$ matrix, with $n \gg d$; let $\varepsilon \in(0,1 / 2]$ be an error parameter; let $\kappa$ be a parameter; and recall the definition of the cross-leverage scores $c_{i j}$ from Definition 1. Then, there exists a randomized algorithm (Algorithm 2 of Section 3 below) that returns the pairs $\{(i, j)\}$ together with estimates $\left\{\tilde{c}_{i j}\right\}$ such that, with probability at least 0.8 ,
i. If $c_{i j}^{2} \geq \frac{d}{\kappa}+12 \varepsilon \ell_{i} \ell_{j}$, then $(i, j)$ is returned; if $(i, j)$ is returned, then $c_{i j}^{2} \geq \frac{d}{\kappa}-30 \varepsilon \ell_{i} \ell_{j}$.
ii. For all pairs $(i, j)$ that are returned, $\tilde{c}_{i j}^{2}-30 \varepsilon \ell_{i} \ell_{j} \leq c_{i j}^{2} \leq \tilde{c}_{i j}^{2}+12 \varepsilon \ell_{i} \ell_{j}$.

This algorithm runs in $O\left(\varepsilon^{-2} n \ln n+\varepsilon^{-3} \kappa d \ln ^{2} n\right)$ time.
Note that by setting $\kappa=n \ln n$, we can compute all the "large" cross-leverage scores, that is, those satisfying $c_{i j}^{2} \geq \frac{d}{n \ln n}$, to within additive-error in $O\left(n d \ln ^{3} n\right)$ time (treating $\varepsilon$ as a constant). If $\ln ^{3} n=$ $o(d)$ the overall running time is $o\left(n d^{2}\right)$, as desired.

### 1.3 Significance and Related Work

Our results are important for their applications to fast randomized matrix algorithms, as well as their applications in numerical linear algebra and large-scale data analysis more generally.

Significance in theoretical computer science. The statistical leverage scores define the key structural nonuniformity that must be dealt with (i.e., either rapidly approximated or rapidly uniformized at the preprocessing step) in developing fast randomized algorithms for matrix problems such as least-squares regression (Sarlós, 2006; Drineas et al., 2010b) and low-rank matrix approximation (Papadimitriou et al., 2000; Sarlós, 2006; Drineas et al., 2008; Mahoney and Drineas, 2009; Boutsidis et al., 2009). Roughly, the best random sampling algorithms use these scores (or the generalized leverage scores relative to the best rank- $k$ approximation to $A$ ) as an importance sampling distribution to sample with respect to. On the other hand, the best random projection algorithms rotate to a basis where these scores are approximately uniform and thus in which uniform sampling is appropriate. See Mahoney (2011) for a detailed discussion.

As an example, the CUR decomposition of Drineas et al. (2008) and Mahoney and Drineas (2009) essentially computes $p_{i}=\ell_{i} / k$, for all $i \in\{1, \ldots, n\}$ and for a rank parameter $k$, and it uses these as an importance sampling distribution. The computational bottleneck for these and related random sampling algorithms is the computation of the importance sampling probabilities. On the other hand, the computational bottleneck for random projection algorithms is the application of the random projection, which is sped up by using variants of the Fast Johnson-Lindenstrauss Transform (Ailon and Chazelle, 2006, 2009). By our main result, the leverage scores (and thus
these probabilities) can be approximated in time that depends on an application of a Fast JohnsonLindenstrauss Transform. In particular, the random sampling algorithms of Drineas et al. (2006b), Drineas et al. (2008) and Mahoney and Drineas (2009) for least-squares approximation and lowrank matrix approximation now run in time that is essentially the same as the best corresponding random projection algorithm for those problems (Sarlós, 2006).

Applications to numerical linear algebra. Recently, high-quality numerical implementations of variants of the basic randomized matrix algorithms have proven superior to traditional deterministic algorithms (Rokhlin and Tygert, 2008; Rokhlin et al., 2009; Avron et al., 2010). An important question raised by our main results is how these will compare with an implementation of our main algorithm. More generally, density functional theory (Bekas et al., 2007) and uncertainty quantification (Bekas et al., 2009) are two scientific computing areas where computing the diagonal elements of functions (such as a projection or inverse) of very large input matrices is common. For example, in the former case, "heuristic" methods based on using Chebychev polynomials have been used in numerical linear algebra to compute the diagonal elements of the projector (Bekas et al., 2007). Our main algorithm should have implications in both of these areas.

Applications in large-scale data analysis. The statistical leverage scores and the scores relative to the best rank- $k$ approximation to $A$ are equal to the diagonal elements of the so-called "hat matrix" (Hoaglin and Welsch, 1978; Chatterjee and Hadi, 1988). As such, they have a natural statistical interpretation in terms of the "leverage" or "influence" associated with each of the data points (Hoaglin and Welsch, 1978; Chatterjee and Hadi, 1986, 1988). In the context of regression problems, the $i^{\text {th }}$ leverage score quantifies the leverage or influence of the $i^{t h}$ constraint/row of $A$ on the solution of the overconstrained least squares optimization problem $\min _{x}\|A x-b\|_{2}$ and the $(i, j)$-th cross leverage score quantifies how much influence or leverage the $i^{\text {th }}$ data point has on the $j^{\text {th }}$ leastsquares fit (see Hoaglin and Welsch, 1978; Chatterjee and Hadi, 1986, 1988, for details). When applied to low-rank matrix approximation problems, the leverage score $\ell_{j}$ quantifies the amount of leverage or influence exerted by the $j^{t h}$ column of $A$ on its optimal low-rank approximation. Historically, these quantities have been widely-used for outlier identification in diagnostic regression analysis (Velleman and Welsch, 1981; Chatterjee et al., 2000).

More recently, these scores (usually the largest scores) often have an interpretation in terms of the data and processes generating the data that can be exploited. For example, depending on the setting, they can have an interpretation in terms of high-degree nodes in data graphs, very small clusters in noisy data, coherence of information, articulation points between clusters, the value of a customer in a network, space localization in sensor networks, etc. (Bonacich, 1987; Richardson and Domingos, 2002; Newman, 2005; Jonckheere et al., 2007; Mahoney, 2011). In genetics, dense matrices of size thousands by hundreds of thousands (a size scale at which even traditional deterministic QR algorithms fail to run) constructed from DNA Single Nucleotide Polymorphisms (SNP) data are increasingly common, and the statistical leverage scores can correlate strongly with other metrics of genetic interest (Paschou et al., 2007; Mahoney and Drineas, 2009; Drineas et al., 2010a; Paschou et al., 2010). Our main result will permit the computation of these scores and related quantities for significantly larger SNP data sets than has been possible previously (Paschou et al., 2007; Drineas et al., 2010a; Paschou et al., 2010; Georgiev and Mukherjee, 2011).

Remark. Lest there be any confusion, we should emphasize our main contributions. First, note that statistical leverage and matrix coherence are important concepts in statistics and machine learning. Second, recall that several random sampling algorithms for ubiquitous matrix problems such as least-squares approximation and low-rank matrix approximation use leverage scores in a
crucial manner; but until now these algorithms were $\Omega\left(T_{S V D}\right)$, where $T_{S V D}$ is the time required to compute a QR decomposition or a partial SVD of the input matrix. Third, note that, in some cases, $o\left(T_{S V D}\right)$ algorithms exist for these problems based on fast random projections. But recall that the existence of those projection algorithms in no way implies that it is easy or obvious how to compute the statistical leverage scores efficiently. Fourth, one implication of our main result is that those random sampling algorithms can now be performed just as efficiently as those random projection algorithms; thus, the solution for those matrix problems can now be obtained while preserving the identity of the rows. That is, these problems can now be solved just as efficiently by using actual rows, rather than the arbitrary linear combinations of rows that are returned by random projections. Fifth, we provide a generalization to "fat" matrices and to obtaining the cross-leverage scores. Sixth, we develop algorithms that can compute leverage scores and related statistics even in streaming environments.

### 1.4 Empirical Discussion of Our Algorithms

Although the main contribution of our paper is to provide a rigorous theoretical understanding of fast leverage score approximation, our paper does analyze the theoretical performance of what is meant to be a practical algorithm. Thus, one might wonder about the empirical performance of our algorithms-for example, whether hidden constants render the algorithms useless for data of realistic size. Not surprisingly, this depends heavily on the quality of the numerical implementation, whether one is interested in "tall" or more general matrices, etc. We will consider empirical and numerical aspects of these algorithms in forthcoming papers, for example, Gittens and Mahoney (2012). We will, however, provide here a brief summary of several numerical issues for the reader interested in these issues.

Empirically, the running time bottleneck for our main algorithm (Algorithm 1 of Section 3) applied to "tall" matrices is the application of the random projection $\Pi_{1}$. Thus, empirically the running time is similar to the running time of random projection based methods for computing approximations to the least-squares problem, which is also dominated by the application of the random projection. The state of the art here is the Blendenpik algorithm of Avron et al. (2010) and the LSRN algorithm of Meng et al. (2011). In their Blendenpik paper, Avron, Maymounkov, and Toledo showed that their high-quality numerical implementation of a Hadamard-based random projection (and associated least-squares computation) "beats LAPACK's ${ }^{3}$ direct dense least-squares solver by a large margin on essentially any dense tall matrix," and they concluded that their empirical results "suggest that random projection algorithms should be incorporated into future versions of LAPACK" (Avron et al., 2010). The LSRN algorithm of Meng, Saunders, and Mahoney improves Blendenpik in several respects, for example, providing better handling of sparsity and rank deficiency, but most notably the random projection underlying LSRN is particularly appropriate for solving large problems on clusters with high communication cost, for example, it has been shown to scale well on Amazon Elastic Cloud Compute clusters. Thus, our main algorithm should extend easily to these environments with the use of the random projection underlying LSRN. Moreover, for both Blendenpik and LSRN (when implemented with a Hadamard-based random projection), the hidden constants in the Hadamard-based random projection are so small that the random projection algorithm (and thus the empirical running time of our main algorithm for approximating leverage

[^0]scores) beats the traditional $O\left(n d^{2}\right)$ time algorithm for dense matrices as small as thousands of rows by hundreds of columns.

### 1.5 Outline

In Section 2, we will provide a brief review of relevant notation and concepts from linear algebra. Then, in Sections 3 and 4, we will present our main results: Section 3 will contain our main algorithm and Section 4 will contain the proof of our main theorem. Section 5 will then describe extensions of our main result to general "fat" matrices, that is, those with $n \approx d$. Section 6 will conclude by describing the relationship of our main result with another related estimator for the statistical leverage scores, an application of our main algorithm to the under-constrained least-squares approximation problem, and extensions of our main algorithm to streaming environments.

## 2. Preliminaries on Linear Algebra and Fast Random Projections

We will start with a review of basic linear algebra and notation, and then we will describe the Fast Johnson-Lindenstrauss Transform and the related Subsampled Randomized Hadamard Transform.

### 2.1 Basic Linear Algebra and Notation

Let $[n]$ denote the set of integers $\{1,2, \ldots, n\}$. For any matrix $A \in \mathbb{R}^{n \times d}$, let $A_{(i)}, i \in[n]$, denote the $i$-th row of $A$ as a row vector, and let $A^{(j)}, j \in[d]$ denote the $j$-th column of $A$ as a column vector. Let $\|A\|_{F}^{2}=\sum_{i=1}^{n} \sum_{j=1}^{d} A_{i j}^{2}$ denote the square of the Frobenius norm of $A$, and let $\|A\|_{2}=$ $\sup _{\|x\|_{2}=1}\|A x\|_{2}$ denote the spectral norm of $A$. Relatedly, for any vector $x \in \mathbb{R}^{n}$, its Euclidean norm (or $\ell_{2}$-norm) is the square root of the sum of the squares of its elements. The dot product between two vectors $x, y \in \mathbb{R}^{n}$ will be denoted $\langle x, y\rangle$, or alternatively as $x^{T} y$. Finally, let $e_{i} \in \mathbb{R}^{n}$, for all $i \in[n]$, denote the standard basis vectors for $\mathbb{R}^{n}$ and let $I_{n}$ denote the $n \times n$ identity matrix.

Let the rank of $A$ be $\rho \leq \min \{n, d\}$, in which case the "compact" or "thin" SVD of $A$ is denoted by $A=U \Sigma V^{T}$, where $U \in \mathbb{R}^{n \times \rho}, \Sigma \in \mathbb{R}^{\rho \times \rho}$, and $V \in \mathbb{R}^{d \times \rho}$. (For a general matrix $X$, we will write $X=U_{X} \Sigma_{X} V_{X}^{T}$.) Let $\sigma_{i}(A), i \in[\rho]$ denote the $i$-th singular value of $A$, and let $\sigma_{\max }(A)$ and $\sigma_{\min }(A)$ denote the maximum and minimum singular values of $A$, respectively. The Moore-Penrose pseudoinverse of $A$ is the $d \times n$ matrix defined by $A^{\dagger}=V \Sigma^{-1} U^{T}$ (Nashed, 1976). Finally, for any orthogonal matrix $U \in \mathbb{R}^{n \times \ell}$, let $U^{\perp} \in \mathbb{R}^{n \times(n-\ell)}$ denote an orthogonal matrix whose columns are an orthonormal basis spanning the subspace of $\mathbb{R}^{n}$ that is orthogonal to the subspace spanned by the columns of $U$ (i.e., the range of $U$ ). It is always possible to extend an orthogonal matrix $U$ to a full orthonormal basis of $\mathbb{R}^{n}$ as $\left[\begin{array}{ll}U & U^{\perp}\end{array}\right]$.

The SVD is important for a number of reasons (Golub and Loan, 1996). For example, the projection of the columns of $A$ onto the $k$ left singular vectors associated with the top $k$ singular values gives the best rank- $k$ approximation to $A$ in the spectral and Frobenius norms. Relatedly, the solution to the least-squares (LS) approximation problem is provided by the SVD: given an $n \times d$ matrix $A$ and an $n$-vector $b$, the LS problem is to compute the minimum $\ell_{2}$-norm vector $x$ such that $\|A x-b\|_{2}$ is minimized over all vectors $x \in \mathbb{R}^{d}$. This optimal vector is given by $x_{o p t}=$ $A^{\dagger} b$. We call a LS problem overconstrained (or overdetermined) if $n>d$ and underconstrained (or underdetermined) if $n<d$.

### 2.2 The Fast Johnson-Lindenstrauss Transform (FJLT)

Given $\varepsilon>0$ and a set of points $x_{1}, \ldots, x_{n}$ with $x_{i} \in \mathbb{R}^{d}$, a $\varepsilon$-Johnson-Lindenstrauss Transform ( $\varepsilon$-JLT), denoted $\Pi \in \mathbb{R}^{r \times d}$, is a projection of the points into $\mathbb{R}^{r}$ such that

$$
(1-\varepsilon)\left\|x_{i}\right\|_{2}^{2} \leq\left\|\Pi x_{i}\right\|_{2}^{2} \leq(1+\varepsilon)\left\|x_{i}\right\|_{2}^{2} .
$$

To construct an $\varepsilon$-JLT with high probability, simply choose every entry of $\Pi$ independently, equal to $\pm \sqrt{3 / r}$ with probability $1 / 6$ each and zero otherwise (with probability $2 / 3$ ) (Achlioptas, 2003). Let $\Pi_{J L T}$ be a matrix drawn from such a distribution over $r \times d$ matrices. ${ }^{4}$ Then, the following lemma holds.

Lemma 4 (Theorem 1.1 of Achlioptas (2003)) Let $x_{1}, \ldots, x_{n}$ be an arbitrary (but fixed) set of points, where $x_{i} \in \mathbb{R}^{d}$ and let $0<\varepsilon \leq 1 / 2$ be an accuracy parameter. If

$$
r \geq \frac{1}{\varepsilon^{2}}\left(12 \ln n+6 \ln \frac{1}{\delta}\right)
$$

then, with probability at least $1-\delta, \Pi_{J L T} \in \mathbb{R}^{r \times d}$ is an $\varepsilon$-JLT .
For our main results, we will also need a stronger requirement than the simple $\varepsilon$-JLT and so we will use a version of the Fast Johnson-Lindenstrauss Transform (FJLT), which was originally introduced in Ailon and Chazelle (2006) and Ailon and Chazelle (2009). Consider an orthogonal matrix $U \in$ $\mathbb{R}^{n \times d}$, viewed as $d$ vectors in $\mathbb{R}^{n}$. A FJLT projects the vectors from $\mathbb{R}^{n}$ to $\mathbb{R}^{r}$, while preserving the orthogonality of $U$; moreover, it does so very quickly. Specifically, given $\varepsilon>0, \Pi \in \mathbb{R}^{r \times n}$ is an $\varepsilon$-FJLT for $U$ if

- $\left\|I_{d}-U^{T} \Pi^{T} \Pi U\right\|_{2} \leq \varepsilon$, and
- given any $X \in \mathbb{R}^{n \times d}$, the matrix product $\Pi X$ can be computed in $O(n d \ln r)$ time.

The next lemma follows from the definition of an $\varepsilon$-FJLT, and its proof can be found in Drineas et al. (2006b) and Drineas et al. (2010b).

Lemma 5 Let $A$ be any matrix in $\mathbb{R}^{n \times d}$ with $n \gg d$ and $\operatorname{rank}(A)=d$. Let the $S V D$ of $A$ be $A=$ $U \Sigma V^{T}$, let $\Pi$ be an $\varepsilon$-FJLT for $U$ (with $0<\varepsilon \leq 1 / 2$ ) and let $\Psi=\Pi U=U_{\Psi} \Sigma_{\Psi} V_{\Psi}^{T}$. Then, all the following hold:

$$
\begin{align*}
\operatorname{rank}(\Pi A) & =\operatorname{rank}(\Pi U)=\operatorname{rank}(U)=\operatorname{rank}(A)=d, \\
\left\|I-\Sigma_{\Psi}^{-2}\right\|_{2} & \leq \varepsilon /(1-\varepsilon), \text { and }  \tag{3}\\
(\Pi A)^{\dagger} & =V \Sigma^{-1}(\Pi U)^{\dagger} . \tag{4}
\end{align*}
$$

[^1]
### 2.3 The Subsampled Randomized Hadamard Transform (SRHT)

One can use a Randomized Hadamard Transform (RHT) to construct, with high probability, an $\varepsilon$-FJLT. Our main algorithm will use this efficient construction in a crucial way. ${ }^{5}$ Recall that the (unnormalized) $n \times n$ matrix of the Hadamard transform $\hat{H}_{n}$ is defined recursively by

$$
\hat{H}_{2 n}=\left[\begin{array}{cc}
\hat{H}_{n} & \hat{H}_{n} \\
\hat{H}_{n} & -\hat{H}_{n}
\end{array}\right],
$$

with $\hat{H}_{1}=1$. The $n \times n$ normalized matrix of the Hadamard transform is equal to

$$
H_{n}=\hat{H}_{n} / \sqrt{n}
$$

From now on, for simplicity and without loss of generality, we assume that $n$ is a power of 2 and we will suppress $n$ and just write $H$. (Variants of this basic construction that relax this assumption and that are more appropriate for numerical implementation have been described and evaluated in Avron et al. (2010).) Let $D \in \mathbb{R}^{n \times n}$ be a random diagonal matrix with independent diagonal entries $D_{i i}=+1$ with probability $1 / 2$ and $D_{i i}=-1$ with probability $1 / 2$. The product $H D$ is a RHT and it has three useful properties. First, when applied to a vector, it "spreads out" its energy. Second, computing the product $H D x$ for any vector $x \in \mathbb{R}^{n}$ takes $O\left(n \log _{2} n\right)$ time. Third, if we only need to access $r$ elements in the transformed vector, then those $r$ elements can be computed in $O\left(n \log _{2} r\right)$ time (Ailon and Liberty, 2008). The Subsampled Randomized Hadamard Transform (SRHT) randomly samples (according to the uniform distribution) a set of $r$ rows of a RHT.

Using the sampling matrix formalism described previously (Drineas et al., 2006a,b, 2008, 2010b), we will represent the operation of randomly sampling $r$ rows of an $n \times d$ matrix $A$ using an $r \times n$ linear sampling operator $S^{T}$. Let the matrix $\Pi_{F J L T}=S^{T} H D$ be generated using the SRHT. ${ }^{6}$ The most important property about the distribution $\Pi_{F J L T}$ is that if $r$ is large enough, then, with high probability, $\Pi_{F J L T}$ generates an $\varepsilon$-FJLT. We summarize this discussion in the following lemma (which is essentially a combination of Lemmas 3 and 4 from Drineas et al. (2010b), restated to fit our notation).
Lemma 6 Let $\Pi_{F J L T} \in \mathbb{R}^{r \times n}$ be generated using the SRHT as described above and let $U \in \mathbb{R}^{n \times d}$ ( $n \gg d$ ) be an (arbitrary but fixed) orthogonal matrix. If

$$
r \geq \frac{14^{2} d \ln (40 n d)}{\varepsilon^{2}} \ln \left(\frac{30^{2} d \ln (40 n d)}{\varepsilon^{2}}\right)
$$

then, with probability at least $0.9, \Pi_{F J L T}$ is an $\varepsilon$-FJLT for $U$.

## 3. Our Main Algorithmic Results

In this section, we will describe our main results for computing relative-error approximations to every statistical leverage score (see Algorithm 1) as well as additive-error approximations to all of the large cross-leverage scores (see Algorithm 2) of an arbitrary matrix $A \in \mathbb{R}^{n \times d}$, with $n \gg d$. Both algorithms make use of a "randomized sketch" of $A$ of the form $A\left(\Pi_{1} A\right)^{\dagger} \Pi_{2}$, where $\Pi_{1}$ is an $\varepsilon$-FJLT and $\Pi_{2}$ is an $\varepsilon$-JLT. We start with a high-level description of the basic ideas.

[^2]
### 3.1 Outline of Our Basic Approach

Recall that our first goal is to approximate, for all $i \in[n]$, the quantities

$$
\begin{equation*}
\ell_{i}=\left\|U_{(i)}\right\|_{2}^{2}=\left\|e_{i}^{T} U\right\|_{2}^{2} \tag{5}
\end{equation*}
$$

where $e_{i}$ is a standard basis vector. The hard part of computing the scores $\ell_{i}$ according to Equation (5) is computing an orthogonal matrix $U$ spanning the range of $A$, which takes $O\left(n d^{2}\right)$ time. Since $U U^{T}=A A^{\dagger}$, it follows that

$$
\begin{equation*}
\ell_{i}=\left\|e_{i}^{T} U U^{T}\right\|_{2}^{2}=\left\|e_{i}^{T} A A^{\dagger}\right\|_{2}^{2}=\left\|\left(A A^{\dagger}\right)_{(i)}\right\|_{2}^{2} \tag{6}
\end{equation*}
$$

where the first equality follows from the orthogonality of (the columns of) $U$. The hard part of computing the scores $\ell_{i}$ according to Equation (6) is two-fold: first, computing the pseudoinverse; and second, performing the matrix-matrix multiplication of $A$ and $A^{\dagger}$. Both of these procedures take $O\left(n d^{2}\right)$ time. As we will see, we can get around both of these bottlenecks by the judicious application of random projections to Equation (6).

To get around the bottleneck of $O\left(n d^{2}\right)$ time due to computing $A^{\dagger}$ in Equation (6), we will compute the pseudoinverse of a "smaller" matrix that approximates $A$. A necessary condition for such a smaller matrix is that it preserves rank. So, naïve ideas such as uniformly sampling $r_{1} \ll n$ rows from $A$ and computing the pseudoinverse of this sampled matrix will not work well for an arbitrary $A$. For example, this idea will fail (with high probability) to return a meaningful approximation for matrices consisting of $n-1$ identical rows and a single row with a nonzero component in the direction perpendicular to that the identical rows; finding that "outlying" row is crucial to obtaining a relative-error approximation. This is where the SRHT enters, since it preserves important structures of $A$, in particular its rank, by first rotating $A$ to a random basis and then uniformly sampling rows from the rotated matrix (see Drineas et al., 2010b for more details). More formally, recall that the SVD of $A$ is $U \Sigma V^{T}$ and let $\Pi_{1} \in \mathbb{R}^{r_{1} \times n}$ be an $\varepsilon$-FJLT for $U$ (using, for example, the SRHT of Lemma 6 with the appropriate choice for $r_{1}$ ). Then, one could approximate the $\ell_{i}$ 's of Equation (6) by

$$
\begin{equation*}
\hat{\ell}_{i}=\left\|e_{i}^{T} A\left(\Pi_{1} A\right)^{\dagger}\right\|_{2}^{2}, \tag{7}
\end{equation*}
$$

where we approximated the $n \times d$ matrix $A$ by the $r_{1} \times d$ matrix $\Pi_{1} A$. Computing $A\left(\Pi_{1} A\right)^{\dagger}$ in this way takes $O\left(n d r_{1}\right)$ time, which is not efficient because $r_{1}>d$ (from Lemma 6).

To get around this bottleneck, recall that we only need the Euclidean norms of the rows of the matrix $A\left(\Pi_{1} A\right)^{\dagger} \in \mathbb{R}^{n \times r_{1}}$. Thus, we can further reduce the dimensionality of this matrix by using an $\varepsilon$-JLT to reduce the dimension $r_{1}=\Omega(d)$ to $r_{2}=O(\ln n)$. Specifically, let $\Pi_{2}^{T} \in \mathbb{R}^{r_{2} \times r_{1}}$ be an $\varepsilon$-JLT for the rows of $A\left(\Pi_{1} A\right)^{\dagger}$ (viewed as $n$ vectors in $\mathbb{R}^{r_{1}}$ ) and consider the matrix $\Omega=A\left(\Pi_{1} A\right)^{\dagger} \Pi_{2}$. This $n \times r_{2}$ matrix $\Omega$ may be viewed as our "randomized sketch" of the rows of $A A^{\dagger}$. Then, we can compute and return

$$
\tilde{\ell}_{i}=\left\|e_{i}^{T} A\left(\Pi_{1} A\right)^{\dagger} \Pi_{2}\right\|_{2}^{2},
$$

for each $i \in[n]$, which is essentially what Algorithm 1 does. Not surprisingly, the sketch $A\left(\Pi_{1} A\right)^{\dagger} \Pi_{2}$ can be used in other ways: for example, by considering the dot product between two different rows of this randomized sketching matrix (and some additional manipulations) Algorithm 2 approximates the large cross-leverage scores of $A$.

Input: $A \in \mathbb{R}^{n \times d}$ (with $\operatorname{SVD} A=U \Sigma V^{T}$ ), error parameter $\varepsilon \in(0,1 / 2]$.
Output: $\tilde{\ell}_{i}, i \in[n]$.

1. Let $\Pi_{1} \in \mathbb{R}^{r_{1} \times n}$ be an $\varepsilon$-FJLT for $U$, using Lemma 6 with

$$
r_{1}=\Omega\left(\frac{d \ln n}{\varepsilon^{2}} \ln \left(\frac{d \ln n}{\varepsilon^{2}}\right)\right) .
$$

2. Compute $\Pi_{1} A \in \mathbb{R}^{r_{1} \times d}$ and its SVD, $\Pi_{1} A=U_{\Pi_{1} A} \Sigma_{\Pi_{1} A} V_{\Pi_{1} A}^{T}$. Let $R^{-1}=V_{\Pi_{1} A} \Sigma_{\Pi_{1} A}^{-1} \in \mathbb{R}^{d \times d}$.
(Alternatively, $R$ could be computed by a $Q R$ factorization of $\Pi_{1} A$.)
3. View the normalized rows of $A R^{-1} \in \mathbb{R}^{n \times d}$ as $n$ vectors in $\mathbb{R}^{d}$, and construct $\Pi_{2} \in \mathbb{R}^{d \times r_{2}}$ to be an $\varepsilon$-JLT for $n^{2}$ vectors (the aforementioned $n$ vectors and their $n^{2}-n$ pairwise sums), using Lemma 4 with

$$
r_{2}=O\left(\varepsilon^{-2} \ln n\right) .
$$

4. Construct the matrix product $\Omega=A R^{-1} \Pi_{2}$.
5. For all $i \in[n]$ compute and return $\tilde{\ell}_{i}=\left\|\Omega_{(i)}\right\|_{2}^{2}$.

Algorithm 1: Approximating the (diagonal) statistical leverage scores $\ell_{i}$.

### 3.2 Approximating All the Statistical Leverage Scores

Our first main result is Algorithm 1, which takes as input an $n \times d$ matrix $A$ and an error parameter $\varepsilon \in(0,1 / 2]$, and returns as output numbers $\tilde{\ell}_{i}, i \in[n]$. Although the basic idea to approximate $\left\|\left(A A^{\dagger}\right)_{(i)}\right\|^{2}$ was described in the previous section, we can improve the efficiency of our approach by avoiding the full sketch of the pseudoinverse. In particular, let $\hat{A}=\Pi_{1} A$ and let its SVD be $\hat{A}=U_{\hat{A}} \Sigma_{\hat{A}} V_{\hat{A}}^{T}$. Let $R^{-1}=V_{\hat{A}} \Sigma_{\hat{A}}^{-1}$ and note that $R^{-1} \in \mathbb{R}^{d \times d}$ is an orthogonalizer for $\hat{A}$ since $U_{\hat{A}}=$ $\hat{A} R^{-1}$ is an orthogonal matrix. ${ }^{7}$ In addition, note that $A R^{-1}$ is approximately orthogonal. Thus, we can compute $A R^{-1}$ and use it as an approximate orthogonal basis for $A$ and then compute $\hat{\ell}_{i}$ as the squared row-norms of $A R^{-1}$. The next lemma states that this is exactly what our main algorithm does; even more, we could get the same estimates by using any "orthogonalizer" of $\Pi_{1} A$.

Lemma 7 Let $R^{-1}$ be such that $Q=\Pi_{1} A R^{-1}$ is an orthogonal matrix with $\operatorname{rank}(Q)=\operatorname{rank}\left(\Pi_{1} A\right)$. Then, $\left\|\left(A R^{-1}\right)_{(i)}\right\|_{2}^{2}=\hat{\ell}_{i}$.
7. This preprocessing is reminiscent of how Rokhlin and Tygert (2008) and Avron et al. (2010) preprocessed the input to provide numerical implementations of the fast relative-error algorithm (Drineas et al., 2010b) for approximate LS approximation. From this perspective, Algorithm 1 can be viewed as specifying a particular basis $Q$, that is, as choosing $Q$ to be the left singular vectors of $\Pi_{1} A$.

Proof Since $\hat{A}=\Pi_{1} A$ has rank $d$ (by Lemma 5) and $R^{-1}$ preserves this rank, $R^{-1}$ is a $d \times d$ invertible matrix. Using $\hat{A}=Q R$ and properties of the pseudoinverse, we get $(\hat{A})^{\dagger}=R^{-1} Q^{T}$. Thus,

$$
\hat{\ell}_{i}=\left\|\left(A\left(\Pi_{1} A\right)^{\dagger}\right)_{(i)}\right\|_{2}^{2}=\left\|\left(A R^{-1} Q^{T}\right)_{(i)}\right\|_{2}^{2}=\left\|\left(A R^{-1}\right)_{(i)} Q^{T}\right\|_{2}^{2}=\left\|\left(A R^{-1}\right)_{(i)}\right\|_{2}^{2} .
$$

This lemma says that the $\hat{\ell}_{i}$ of Equation (7) can be computed with any QR decomposition, rather than with the SVD; but note that one would still have to post-multiply by $\Pi_{2}$, as in Algorithm 1, in order to compute "quickly" the approximations of the leverage scores.

### 3.3 Approximating the Large Cross-leverage Scores

By combining Lemmas 9 and 10 (in Section 4.2 below) with the triangle inequality, one immediately obtains the following lemma.

Lemma 8 Let $\Omega$ be either the sketching matrix constructed by Algorithm 1, that is, $\Omega=A R^{-1} \Pi_{2}$, or $\Omega=A\left(\Pi_{1} A\right)^{\dagger} \Pi_{2}$ as described in Section 3.1. Then, the pairwise dot-products of the rows of $\Omega$ are additive-error approximations to the leverage scores and cross-leverage scores:

$$
\left|\left\langle U_{(i)}, U_{(j)}\right\rangle-\left\langle\Omega_{(i)}, \Omega_{(j)}\right\rangle\right| \leq \frac{3 \varepsilon}{1-\varepsilon}\left\|U_{(i)}\right\|_{2}\left\|U_{(j)}\right\|_{2}
$$

That is, if one were interested in obtaining an approximation to all the cross-leverage scores to within additive error (and thus the diagonal statistical leverage scores to relative-error), then the algorithm which first computes $\Omega$ followed by all the pairwise inner products achieves this in time $T(\Omega)+O\left(n^{2} r_{2}\right)$, where $T(\Omega)$ is the time to compute $\Omega$ from Section 3.2 and $r_{2}=O\left(\varepsilon^{-2} \ln n\right) .{ }^{8}$ The challenge is to avoid the $n^{2}$ computational complexity and this can be done if one is interested only in the large cross-leverage scores.

Our second main result is provided by Algorithms 2 and 3. Algorithm 2 takes as input an $n \times d$ matrix $A$, a parameter $\kappa>1$, and an error parameter $\varepsilon \in(0,1 / 2]$, and returns as output a subset of $[n] \times[n]$ and estimates $\tilde{c}_{i j}$ satisfying Theorem 3. The first step of the algorithm is to compute the matrix $\Omega=A R^{-1} \Pi_{2}$ constructed by Algorithm 1. Then, Algorithm 2 uses Algorithm 3 as a subroutine to compute "heavy hitter" pairs of rows from a matrix.

## 4. Proofs of Our Main Theorems

We will start with a sketch of the proofs of Theorems 2 and 3, and then we will provide the details.

### 4.1 Sketch of the Proof of Theorems 2 and 3

We will start by providing a sketch of the proof of Theorems 2 and 3 . A detailed proof is provided in the next two subsections. In our analysis, we will condition on the events that $\Pi_{1} \in \mathbb{R}^{r_{1} \times n}$ is an $\varepsilon$-FJLT for $U$ and $\Pi_{2} \in \mathbb{R}^{r_{1} \times r_{2}}$ is an $\varepsilon$-JLT for $n^{2}$ points in $\mathbb{R}^{r_{1}}$. Note that by setting $\delta=0.1$ in Lemma 4 , both events hold with probability at least 0.8 , which is equal to the success probability of Theorems 2 and 3. The algorithm estimates $\tilde{\ell}_{i}=\left\|\tilde{u}_{i}\right\|_{2}^{2}$, where $\tilde{u}_{i}=e_{i}^{T} A\left(\Pi_{1} A\right)^{\dagger} \Pi_{2}$. First, observe
8. The exact algorithm which computes a basis first and then the pairwise inner products requires $O\left(n d^{2}+n^{2} d\right)$ time. Thus, by using the sketch, we can already improve on this running time by a factor of $d / \ln n$.

Input: $A \in \mathbb{R}^{n \times d}$ and parameters $\kappa>1, \varepsilon \in(0,1 / 2]$.
Output: The set $\mathcal{H}$ consisting of pairs $(i, j)$ together with estimates $\tilde{c}_{i j}$ satisfying Theorem 3.

1. Compute the $n \times r_{2}$ matrix $\Omega=A R^{-1} \Pi_{2}$ from Algorithm 1 .
2. Use Algorithm 3 with inputs $\Omega$ and $\kappa^{\prime}=\kappa(1+30 d \varepsilon)$ to obtain the set $\mathcal{H}$ containing all the $\kappa^{\prime}$-heavy pairs of $\Omega$.
3. Return the pairs in $\mathcal{H}$ as the $\kappa$-heavy pairs of $A$.

Algorithm 2: Approximating the large (off-diagonal) cross-leverage scores $c_{i j}$.

```
Input: \(X \in \mathbb{R}^{n \times r}\) with rows \(x_{1}, \ldots, x_{n}\) and a parameter \(\kappa>1\).
Output: \(\mathcal{H}=\left\{(i, j), \tilde{c}_{i j}\right\}\) containing all heavy (unordered) pairs. The pair \((i, j), \tilde{c}_{i j} \in \mathcal{H}\)
if and only if \(\tilde{c}_{i j}^{2}=\left\langle x_{i}, x_{j}\right\rangle^{2} \geq\left\|X^{T} X\right\|_{F}^{2} / \kappa\).
    Compute the norms \(\left\|x_{i}\right\|_{2}\) and sort the rows according to norm, so that
    \(\left\|x_{1}\right\|_{2} \leq \cdots \leq\left\|x_{n}\right\|_{2}\).
    \(\mathcal{H} \leftarrow\left\} ; z_{1} \leftarrow n ; z_{2} \leftarrow 1\right.\).
    while \(z_{2} \leq z_{1}\) do
        while \(\left\|x_{z_{1}}\right\|_{2}^{2}\left\|x_{z_{2}}\right\|_{2}^{2}<\left\|X^{T} X\right\|_{F}^{2} / \kappa\) do
            \(z_{2} \leftarrow z_{2}+1\).
            if \(z_{2}>z_{1}\) then
                return \(\mathcal{H}\).
            end if
        end while
        for each pair \((i, j)\) where \(i=z_{1}\) and \(j \in\left\{z_{2}, z_{2}+1, \ldots, z_{1}\right\}\) do
            \(\tilde{c}_{i j}^{2}=\left\langle x_{i}, x_{j}\right\rangle^{2}\).
            if \(\tilde{c}_{i j}^{2} \geq\left\|X^{T} X\right\|_{F}^{2} / \kappa\) then
                add \((i, j)\) and \(\tilde{c}_{i j}\) to \(\mathcal{H}\).
            end if
            \(z_{1} \leftarrow z_{1}-1\).
        end for
    end while
    return \(\mathcal{H}\).
```

Algorithm 3: Computing heavy pairs of a matrix.
that the sole purpose of $\Pi_{2}$ is to improve the running time while preserving pairwise inner products;
this is achieved because $\Pi_{2}$ is an $\varepsilon$-JLT for $n^{2}$ points. So, the results will follow if

$$
e_{i}^{T} A\left(\Pi_{1} A\right)^{\dagger}\left(\left(\Pi_{1} A\right)^{\dagger}\right)^{T} A^{T} e_{j} \approx e_{i}^{T} U U^{T} e_{j}
$$

and $\left(\Pi_{1} A\right)^{\dagger}$ can be computed efficiently. Since $\Pi_{1}$ is an $\varepsilon$-FJLT for $U$, where $A=U \Sigma V^{T},\left(\Pi_{1} A\right)^{\dagger}$ can be computed in $O\left(n d \ln r_{1}+r_{1} d^{2}\right)$ time. By Lemma 5, $\left(\Pi_{1} A\right)^{\dagger}=V \Sigma^{-1}\left(\Pi_{1} U\right)^{\dagger}$, and so

$$
e_{i}^{T} A\left(\Pi_{1} A\right)^{\dagger}\left(\left(\Pi_{1} A\right)^{\dagger}\right)^{T} A^{T} e_{j}=e_{i}^{T} U\left(\Pi_{1} U\right)^{\dagger}\left(\Pi_{1} U\right)^{\dagger^{T}} U^{T} e_{j} .
$$

Since $\Pi_{1}$ is an $\varepsilon$-FJLT for $U$, it follows that $\left(\Pi_{1} U\right)^{\dagger}\left(\Pi_{1} U\right)^{\dagger^{T}} \approx I_{d}$, that is, that $\Pi_{1} U$ is approximately orthogonal. Theorem 2 follows from this basic idea. However, in order to prove Theorem 3, having a sketch which preserves inner products alone is not sufficient. We also need a fast algorithm to identify the large inner products and to relate these to the actual cross-leverage scores. Indeed, it is possible to efficiently find pairs of rows in a general matrix with large inner products. Combining this with the fact that the inner products are preserved, we obtain Theorem 3.

### 4.2 Proof of Theorem 2

We condition all our analysis on the events that $\Pi_{1} \in \mathbb{R}^{r_{1} \times n}$ is an $\varepsilon$-FJLT for $U$ and $\Pi_{2} \in \mathbb{R}^{r_{1} \times r_{2}}$ is an $\varepsilon$-JLT for $n^{2}$ points in $\mathbb{R}^{r_{1}}$. Define

$$
\begin{aligned}
\hat{u}_{i} & =e_{i}^{T} A\left(\Pi_{1} A\right)^{\dagger}, \text { and } \\
\tilde{u}_{i} & =e_{i}^{T} A\left(\Pi_{1} A\right)^{\dagger} \Pi_{2} .
\end{aligned}
$$

Then, $\hat{\ell}_{i}=\left\|\hat{u}_{i}\right\|_{2}^{2}$ and $\tilde{\ell}_{i}=\left\|\tilde{u}_{i}\right\|_{2}^{2}$. The proof will follow from the following two lemmas.
Lemma 9 For $i, j \in[n]$,

$$
\left|\left\langle U_{(i)}, U_{(j)}\right\rangle-\left\langle\hat{u}_{i}, \hat{u}_{j}\right\rangle\right| \leq \frac{\varepsilon}{1-\varepsilon}\left\|U_{(i)}\right\|_{2}\left\|U_{(j)}\right\|_{2} .
$$

Lemma 10 For $i, j \in[n]$,

$$
\left|\left\langle\hat{u}_{i}, \hat{u}_{j}\right\rangle-\left\langle\tilde{u}_{i}, \tilde{u}_{j}\right\rangle\right| \leq 2 \varepsilon\left\|\hat{u}_{i}\right\|_{2}\left\|\hat{u}_{j}\right\|_{2} .
$$

Lemma 9 states that $\left\langle\hat{u}_{i}, \hat{u}_{j}\right\rangle$ is an additive error approximation to all the cross-leverage scores $(i \neq j)$ and a relative error approximation for the diagonals $(i=j)$. Similarly, Lemma 10 shows that these cross-leverage scores are preserved by $\Pi_{2}$. Indeed, with $i=j$, from Lemma 9 we have $\left|\hat{\ell}_{i}-\ell_{i}\right| \leq \frac{\varepsilon}{1-\varepsilon} \ell_{i}$, and from Lemma 10 we have $\left|\hat{\ell}_{i}-\tilde{\ell}_{i}\right| \leq 2 \hat{\ell} \hat{\ell}_{i}$. Using the triangle inequality and $\varepsilon \leq 1 / 2$ :

$$
\left|\ell_{i}-\tilde{\ell}_{i}\right|=\left|\ell_{i}-\hat{\ell}_{i}+\hat{\ell}_{i}-\tilde{\ell}_{i}\right| \leq\left|\ell_{i}-\hat{\ell}_{i}\right|+\left|\hat{\ell}_{i}-\tilde{\ell}_{i}\right| \leq\left(\frac{\varepsilon}{1-\varepsilon}+2 \varepsilon\right) \ell_{i} \leq 4 \varepsilon \ell_{i} \text {. }
$$

The theorem follows after rescaling $\varepsilon$.

Proof of Lemma 9 Let $A=U \Sigma V^{T}$. Using this SVD of $A$ and Equation (4) in Lemma 5,

$$
\left\langle\hat{u}_{i}, \hat{u}_{j}\right\rangle=e_{i}^{T} U \Sigma V^{T} V \Sigma^{-1}\left(\Pi_{1} U\right)^{\dagger}\left(\Pi_{1} U\right)^{\dagger^{T}} \Sigma^{-1} V^{T} V \Sigma U^{T} e_{j}=e_{i}^{T} U\left(\Pi_{1} U\right)^{\dagger}\left(\Pi_{1} U\right)^{\dagger T} U^{T} e_{j}
$$

By performing standard manipulations, we can now bound $\left|\left\langle U_{(i)}, U_{(j)}\right\rangle-\left\langle\hat{u}_{i}, \hat{u}_{j}\right\rangle\right|$ :

$$
\begin{aligned}
\left|\left\langle U_{(i)}, U_{(j)}\right\rangle-\left\langle\hat{u}_{i}, \hat{u}_{j}\right\rangle\right| & =e_{i}^{T} U U^{T} e_{j}-e_{i}^{T} U\left(\Pi_{1} U\right)^{\dagger}\left(\Pi_{1} U\right)^{\dagger T} U^{T} e_{j} \\
& =e_{i}^{T} U\left(I_{d}-\left(\Pi_{1} U\right)^{\dagger}\left(\Pi_{1} U\right)^{\dagger T}\right) U^{T} e_{j} \\
& \leq\left\|I_{d}-\left(\Pi_{1} U\right)^{\dagger}\left(\Pi_{1} U\right)^{\dagger T}\right\|_{2}\left\|U_{(i)}\right\|_{2}\left\|U_{(j)}\right\|_{2} .
\end{aligned}
$$

Let the SVD of $\Psi=\Pi_{1} U$ be $\Psi=U_{\Psi} \Sigma_{\Psi} V_{\Psi}^{T}$, where $V_{\Psi}$ is a full rotation in $d$ dimensions (because $\operatorname{rank}(A)=\operatorname{rank}\left(\Pi_{1} U\right)$ ). Then, $\Psi^{\dagger} \Psi^{\dagger^{T}}=V_{\Psi} \Sigma_{\Psi}^{-2} V_{\Psi}^{T}$. Thus,

$$
\begin{aligned}
\left|\left\langle U_{(i)}, U_{(j)}\right\rangle-\left\langle\hat{u}_{i}, \hat{u}_{j}\right\rangle\right| & \leq\left\|I_{d}-V_{\Psi} \Sigma_{\Psi}^{-2} V_{\Psi}^{T}\right\|_{2}\left\|U_{(i)}\right\|_{2}\left\|U_{(j)}\right\|_{2} \\
& =\left\|V_{\Psi} V_{\Psi}^{T}-V_{\Psi} \Sigma_{\Psi}^{-2} V_{\Psi}^{T}\right\|_{2}\left\|U_{(i)}\right\|_{2}\left\|U_{(j)}\right\|_{2} \\
& =\left\|I_{d}-\Sigma_{\Psi}^{-2}\right\|_{2}\left\|U_{(i)}\right\|_{2}\left\|U_{(j)}\right\|_{2}
\end{aligned}
$$

where we used the fact that $V_{\Psi} V_{\Psi}^{T}=V_{\Psi}^{T} V_{\Psi}=I_{d}$ and the unitary invariance of the spectral norm. Finally, using Equation (3) of Lemma 5 the result follows.
Proof of Lemma 10. Since $\Pi_{2}$ is an $\varepsilon$-JLT for $n^{2}$ vectors, it preserves the norms of an arbitrary (but fixed) collection of $n^{2}$ vectors. Let $x_{i}=\hat{u}_{i} /\left\|\hat{u}_{i}\right\|_{2}$. Consider the following $n^{2}$ vectors:

$$
\begin{array}{cl}
x_{i} & \text { for } i \in[n], \text { and } \\
x_{i}+x_{j} & \text { for } i, j \in[n], i \neq j
\end{array}
$$

By the $\varepsilon$-JLT property of $\Pi_{2}$ and the fact that $\left\|x_{i}\right\|_{2}=1$,

$$
\begin{array}{cl}
1-\varepsilon \leq\left\|x_{i} \Pi_{2}\right\|_{2}^{2} \leq 1+\varepsilon & \text { for } i \in[n], \text { and } \\
(1-\varepsilon)\left\|x_{i}+x_{j}\right\|_{2}^{2} \leq\left\|x_{i} \Pi_{2}+x_{j} \Pi_{2}\right\|_{2}^{2} \leq(1+\varepsilon)\left\|x_{i}+x_{j}\right\|_{2}^{2} & \text { for } i, j \in[n], i \neq j \tag{9}
\end{array}
$$

Combining Equations (8) and (9) after expanding the squares using the identity $\|a+b\|^{2}=\|a\|^{2}+$ $\|b\|^{2}+2\langle a, b\rangle$, substituting $\left\|x_{i}\right\|=1$, and after some algebra, we obtain

$$
\left\langle x_{i}, x_{j}\right\rangle-2 \varepsilon \leq\left\langle x_{i} \Pi_{2}, x_{j} \Pi_{2}\right\rangle \leq\left\langle x_{i}, x_{j}\right\rangle+2 \varepsilon
$$

To conclude the proof, multiply throughout by $\left\|\hat{u}_{i}\right\|\left\|\hat{u}_{j}\right\|$ and use the homogeneity of the inner product, together with the linearity of $\Pi_{2}$, to obtain:

$$
\left\langle\hat{u}_{i}, \hat{u}_{j}\right\rangle-2 \varepsilon\left\|\hat{u}_{i}\right\|\left\|\hat{u}_{j}\right\| \leq\left\langle\hat{u}_{i} \Pi_{2}, \hat{u}_{j} \Pi_{2}\right\rangle \leq\left\langle\hat{u}_{i}, \hat{u}_{j}\right\rangle+2 \varepsilon\left\|\hat{u}_{i}\right\|\left\|\hat{u}_{j}\right\|
$$

Running Times. By Lemma 7, we can use $V_{\Pi_{1} A} \Sigma_{\Pi_{1} A}^{-1}$ instead of $\left(\Pi_{1} A\right)^{\dagger}$ and obtain the same estimates. Since $\Pi_{1}$ is an $\varepsilon$-FJLT, the product $\Pi_{1} A$ can be computed in $O\left(n d \ln r_{1}\right)$ while its SVD takes an additional $O\left(r_{1} d^{2}\right)$ time to return $V_{\Pi_{1} A} \Sigma_{\Pi_{1} A}^{-1} \in \mathbb{R}^{d \times d}$. Since $\Pi_{2} \in \mathbb{R}^{d \times r_{2}}$, we obtain $V_{\Pi_{1} A} \Sigma_{\Pi_{1} A}^{-1} \Pi_{2} \in \mathbb{R}^{d \times r_{2}}$ in an additional $O\left(r_{2} d^{2}\right)$ time. Finally, premultiplying by $A$ takes $O\left(n d r_{2}\right)$
time, and computing and returning the squared row-norms of $\Omega=A V_{\Pi_{1} A} \Sigma_{\Pi_{1} A}^{-1} \Pi_{2} \in \mathbb{R}^{n \times r_{2}}$ takes $O\left(n r_{2}\right)$ time. So, the total running time is the sum of all these operations, which is

$$
O\left(n d \ln r_{1}+n d r_{2}+r_{1} d^{2}+r_{2} d^{2}\right) .
$$

Recall that for our implementations of the $\varepsilon$-JLTs and $\varepsilon$-FJLTs, we have $\delta=0.1$, and we have $r_{1}=O\left(\varepsilon^{-2} d(\ln n)\left(\ln \left(\varepsilon^{-2} d \ln n\right)\right)\right)$ and $r_{2}=O\left(\varepsilon^{-2} \ln n\right)$. It follows that the asymptotic running time is

$$
O\left(n d \ln \left(d \varepsilon^{-1}\right)+n d \varepsilon^{-2} \ln n+d^{3} \varepsilon^{-2}(\ln n)\left(\ln \left(d \varepsilon^{-1}\right)\right)\right) .
$$

To simplify, suppose that $d \leq n \leq e^{d}$ and treat $\varepsilon$ as a constant. Then, the asymptotic running time is

$$
O\left(n d \ln n+d^{3}(\ln n)(\ln d)\right)
$$

### 4.3 Proof of Theorem 3

We first construct an algorithm to estimate the large inner products among the rows of an arbitrary matrix $X \in \mathbb{R}^{n \times r}$ with $n>r$. This general algorithm will be applied to the matrix $\Omega=A V_{\Pi_{1} A} \Sigma_{\Pi_{1} A}^{-1} \Pi_{2}$. Let $x_{1}, \ldots, x_{n}$ denote the rows of $X$; for a given $\mathrm{\kappa}>1$, the pair $(i, j)$ is heavy if

$$
\left\langle x_{i}, x_{j}\right\rangle^{2} \geq \frac{1}{\kappa}\left\|X^{T} X\right\|_{F}^{2}
$$

By the Cauchy-Schwarz inequality, this implies that

$$
\begin{equation*}
\left\|x_{i}\right\|_{2}^{2}\left\|x_{j}\right\|_{2}^{2} \geq \frac{1}{\kappa}\left\|X^{T} X\right\|_{F}^{2}, \tag{10}
\end{equation*}
$$

so it suffices to find all the pairs $(i, j)$ for which Equation (10) holds. We will call such pairs normheavy. Let $s$ be the number of norm-heavy pairs satisfying Equation (10). We first bound the number of such pairs.

Lemma 11 Using the above notation, $s \leq \kappa r$.
Proof Observe that

$$
\sum_{i, j=1}^{n}\left\|x_{i}\right\|_{2}^{2}\left\|x_{j}\right\|_{2}^{2}=\left(\sum_{i=1}^{n}\left\|x_{i}\right\|_{2}^{2}\right)^{2}=\|X\|_{F}^{4}=\left(\sum_{i=1}^{r} \sigma_{i}^{2}\right)^{2}
$$

where $\sigma_{1}, \ldots, \sigma_{r}$ are the singular values of $X$. To conclude, by the definition of a heavy pair,

$$
\sum_{i, j}\left\|x_{i}\right\|_{2}^{2}\left\|x_{j}\right\|_{2}^{2} \geq \frac{s}{\kappa}\left\|X^{T} X\right\|_{F}^{2}=\frac{s}{\kappa} \sum_{i=1}^{r} \sigma_{i}^{4} \geq \frac{s}{\kappa r}\left(\sum_{i=1}^{r} \sigma_{i}^{2}\right)^{2},
$$

where the last inequality follows by Cauchy-Schwarz.
Algorithm 3 starts by computing the norms $\left\|x_{i}\right\|_{2}^{2}$ for all $i \in[n]$ and sorting them (in $O(n r+n \ln n)$ time) so that we can assume that $\left\|x_{1}\right\|_{2} \leq \cdots \leq\left\|x_{n}\right\|_{2}$. Then, we initialize the set of norm-heavy pairs to $\mathcal{H}=\{ \}$ and we also initialize two pointers $z_{1}=n$ and $z_{2}=1$. The basic loop in the algorithm checks if $z_{2}>z_{1}$ and stops if that is the case. Otherwise, we increment $z_{2}$ to the first pair $\left(z_{1}, z_{2}\right)$ that
is norm-heavy. If none of pairs are norm heavy (i.e., $z_{2}>z_{1}$ occurs), then we stop and output $\mathcal{H}$; otherwise, we add $\left(z_{1}, z_{2}\right),\left(z_{1}, z_{2}+1\right), \ldots,\left(z_{1}, z_{1}\right)$ to $\mathcal{H}$. This basic loop computes all pairs $\left(z_{1}, i\right)$ with $i \leq z_{1}$ that are norm-heavy. Next, we decrease $z_{1}$ by one and if $z_{1}<z_{2}$ we stop and output $\mathcal{H}$; otherwise, we repeat the basic loop. Note that in the basic loop $z_{2}$ is always incremented. This occurs whenever the pair $\left(z_{1}, z_{2}\right)$ is not norm-heavy. Since $z_{2}$ can be incremented at most $n$ times, the number of times we check whether a pair is norm-heavy and fail is at most $n$. Every successful check results in the addition of at least one norm-heavy pair into $\mathcal{H}$ and thus the number of times we check if a pair is norm heavy (a constant-time operation) is at most $n+s$. The number of pair additions into $\mathcal{H}$ is exactly $s$ and thus the total running time is $O(n r+n \ln n+s)$. Finally, we must check each norm-heavy pair to verify whether or not it is actually heavy by computing $s$ inner products vectors in $\mathbb{R}^{r}$; this can be done in $O(s r)$ time. Using $s \leq \kappa r$ we get the following lemma.
Lemma 12 Algorithm 3 returns $\mathcal{H}$ including all the heavy pairs of $X$ in $O\left(n r+\kappa r^{2}+n \ln n\right)$ time.
To complete the proof, we apply Algorithm 3 with $\Omega=A V_{\Pi_{1} A} \Sigma_{\Pi_{1} A}^{-1} \Pi_{2} \in \mathbb{R}^{n \times r_{2}}$, where $r_{2}=O\left(\varepsilon^{-2} \ln n\right)$. Let $\tilde{u}_{1}, \ldots, \tilde{u}_{n}$ denote the rows of $\Omega$ and recall that $A=U \Sigma V^{T}$. Let $u_{1}, \ldots, u_{n}$ denote the rows of $U$; then, from Lemma 8,

$$
\begin{equation*}
\left\langle u_{i}, u_{j}\right\rangle-\frac{3 \varepsilon}{1-\varepsilon}\left\|u_{i}\right\|\left\|u_{j}\right\| \leq\left\langle\tilde{u}_{i}, \tilde{u}_{j}\right\rangle \leq\left\langle u_{i}, u_{j}\right\rangle+\frac{3 \varepsilon}{1-\varepsilon}\left\|u_{i}\right\|\left\|u_{j}\right\| . \tag{11}
\end{equation*}
$$

Given $\varepsilon, \kappa$, assume that for the pair of vectors $u_{i}$ and $u_{j}$

$$
\left\langle u_{i}, u_{j}\right\rangle^{2} \geq \frac{1}{\kappa}\left\|U^{T} U\right\|_{F}^{2}+12 \varepsilon\left\|u_{i}\right\|^{2}\left\|u_{j}\right\|^{2}=\frac{d}{\kappa}+12 \varepsilon\left\|u_{i}\right\|^{2}\left\|u_{j}\right\|^{2},
$$

where the last equality follows from $\left\|U^{T} U\right\|_{F}^{2}=\left\|I_{d}\right\|_{F}^{2}=d$. By Equation (11), after squaring and using $\varepsilon<0.5$,

$$
\begin{equation*}
\left\langle u_{i}, u_{j}\right\rangle^{2}-12 \varepsilon\left\|u_{i}\right\|^{2} \varepsilon\left\|u_{j}\right\|^{2} \leq\left\langle\tilde{u}_{i}, \tilde{u}_{j}\right\rangle^{2} \leq\left\langle u_{i}, u_{j}\right\rangle^{2}+30 \varepsilon\left\|u_{i}\right\|^{2}\left\|u_{j}\right\|^{2} . \tag{12}
\end{equation*}
$$

Thus, $\left\langle\tilde{u}_{i}, \tilde{u}_{j}\right\rangle^{2} \geq d / \kappa$ and summing Equation (12) over all $i, j$ we get $\left\|\Omega^{T} \Omega\right\|_{F}^{2} \leq d+30 \varepsilon d^{2}$, or, equivalently,

$$
d \geq \frac{\left\|\Omega^{T} \Omega\right\|_{F}^{2}}{1+30 d \varepsilon} .
$$

We conclude that

$$
\begin{equation*}
\left\langle u_{i}, u_{j}\right\rangle^{2} \geq \frac{d}{\kappa}+12 \varepsilon\left\|u_{i}\right\|^{2}\left\|u_{j}\right\|^{2} \Longrightarrow\left\langle\tilde{u}_{i}, \tilde{u}_{j}\right\rangle^{2} \geq \frac{d}{\kappa} \geq \frac{\left\|\Omega^{T} \Omega\right\|_{F}^{2}}{\kappa(1+30 d \varepsilon)} . \tag{13}
\end{equation*}
$$

By construction, Algorithm 3 is invoked with $\kappa^{\prime}=\kappa\left\|\Omega^{T} \Omega\right\|_{F}^{2} / d$ and thus it finds all pairs with $\left\langle\tilde{u}_{i}, \tilde{u}_{j}\right\rangle^{2} \geq\left\|\Omega^{T} \Omega\right\|_{F}^{2} / \kappa^{\prime}=d / \kappa$. This set contains all pairs for which

$$
\left\langle u_{i}, u_{j}\right\rangle^{2} \geq \frac{d}{\kappa}+12 \varepsilon\left\|u_{i}\right\|^{2}\left\|u_{j}\right\|^{2}
$$

Further, since every pair returned satisfies $\left\langle\tilde{u}_{i}, \tilde{u}_{j}\right\rangle^{2} \geq d / \kappa$, by Equation (12), $c_{i j} \geq d / \kappa-30 \varepsilon \ell_{i} \ell_{j}$. This proves the first claim of the Theorem; the second claim follows analogously from Equation (12).

Using Lemma 12, the running time of our approach is $O\left(n r_{2}+\kappa^{\prime} r_{2}^{2}+n \ln n\right)$. Since $r_{2}=$ $O\left(\varepsilon^{-2} \ln n\right)$, and, by Equation (13), $\kappa^{\prime}=\kappa\left\|\Omega^{T} \Omega\right\|_{F}^{2} / d \leq \kappa(1+30 d \varepsilon)$, the overall running time is $O\left(\varepsilon^{-2} n \ln n+\varepsilon^{-3} \kappa d \ln ^{2} n\right)$.

## 5. Extending Our Algorithm to General Matrices

In this section, we will describe an important extension of our main result, namely the computation of the statistical leverage scores relative to the best rank- $k$ approximation to a general matrix $A$. More specifically, we consider the estimation of leverage scores for the case of general "fat" matrices, namely input matrices $A \in \mathbb{R}^{n \times d}$, where both $n$ and $d$ are large, for example, when $d=n$ or $d=\Theta(n)$. Clearly, the leverage scores of any full rank $n \times n$ matrix are exactly uniform. The problem becomes interesting if one specifies a rank parameter $k \ll \min \{n, d\}$. This may arise when the numerical rank of $A$ is small (e.g., in some scientific computing applications, more than $99 \%$ of the spectral norm of $A$ may be captured by some $k \ll \min \{n, d\}$ directions), or, more generally, when one is interested in some low rank approximation to $A$ (e.g., in some data analysis applications, a reasonable fraction or even the majority of the Frobenius norm of $A$ may be captured by some $k \ll \min \{n, d\}$ directions, where $k$ is determined by some exogenously-specified model selection criterion). Thus, assume that in addition to a general $n \times d$ matrix $A$, a rank parameter $k<\min \{n, d\}$ is specified. In this case, we wish to obtain the statistical leverage scores $\ell_{i}=\left\|\left(U_{k}\right)_{(i)}\right\|_{2}^{2}$ for $A_{k}=U_{k} \Sigma_{k} V_{k}^{T}$, the best rank-k approximation to $A$. Equivalently, we seek the normalized leverage scores

$$
\begin{equation*}
p_{i}=\frac{\ell_{i}}{k} . \tag{14}
\end{equation*}
$$

Note that $\sum_{i=1}^{n} p_{i}=1$ since $\sum_{i=1}^{n} \ell_{i}=\left\|U_{k}\right\|_{F}^{2}=k$.
Unfortunately, as stated, this is an ill-posed problem. Indeed, consider the degenerate case when $A=I_{n}$ (i.e., the $n \times n$ identity matrix). In this case, $U_{k}$ is not unique and the leverage scores are not well-defined. Moreover, for the obvious $\binom{n}{k}$ equivalent choices for $U_{k}$, the leverage scores defined according to any one of these choices do not provide a relative error approximation to the leverage scores defined according to any other choices. More generally, removing this trivial degeneracy does not help. Consider the matrix

$$
A=\left(\begin{array}{cc}
I_{k} & 0 \\
0 & (1-\gamma) I_{n-k}
\end{array}\right) \in \mathbb{R}^{n \times n} .
$$

In this example, the leverage scores for $A_{k}$ are well defined. However, as $\gamma \rightarrow 0$, it is not possible to distinguish between the top- $k$ singular space and its complement. This example suggests that it should be possible to obtain some result conditioning on the spectral gap at the $k^{t h}$ singular value. For example, one might assume that $\sigma_{k}^{2}-\sigma_{k+1}^{2} \geq \gamma>0$, in which case the parameter $\gamma$ would play an important role in the ability to solve this problem. Any algorithm which cannot distinguish the singular values with an error less than $\gamma$ will confuse the $k$-th and $(k+1)$-th singular vectors and consequently will fail to get an accurate approximation to the leverage scores for $A_{k}$.

In the following, we take a more natural approach which leads to a clean problem formulation. To do so, recall that the leverage scores and the related normalized leverage scores of Equation (14) are used to approximate the matrix in some way, for example, we might be seeking a low-rank approximation to the matrix with respect to the spectral (Drineas et al., 2008) or the Frobenius (Boutsidis et al., 2009) norm, or we might be seeking useful features or data points in downstream data analysis applications (Paschou et al., 2007; Mahoney and Drineas, 2009), or we might be seeking to develop high-quality numerical implementations of low-rank matrix approximation algorithms (Halko et al., 2011), etc. In all these cases, we only care that the estimated leverage scores are a good approximation to the leverage scores of some "good" low-rank approximation to $A$. The
following definition captures the notion of a set of rank- $k$ matrices that are good approximations to A.

Definition 13 Given $A \in \mathbb{R}^{n \times d}$ and a rank parameter $k \ll \min \{n, d\}$, let $A_{k}$ be the best rank-k approximation to $A$. Define the set $\mathcal{S}_{\varepsilon}$ of rank-k matrices that are good approximations to $A$ as follows (for $\xi=2, F$ ):

$$
\mathcal{S}_{\varepsilon}=\left\{X \in \mathbb{R}^{n \times d}: \operatorname{rank}(X)=k \text { and }\|A-X\|_{\xi} \leq(1+\varepsilon)\left\|A-A_{k}\right\|_{\xi}\right\} .
$$

We are now ready to define our approximations to the normalized leverage scores of any matrix $A \in \mathbb{R}^{n \times d}$ given a rank parameter $k \ll \min \{n, d\}$. Instead of seeking to approximate the $p_{i}$ of Equation (14) (a problem that is ill-posed as discussed above), we will be satisfied if we can approximate the normalized leverage scores of some matrix $X \in \mathcal{S}_{\mathcal{\varepsilon}}$. This is an interesting relaxation of the task at hand: all matrices $X$ that are sufficiently close to $A_{k}$ are essentially equivalent, since they can be used instead of $A_{k}$ in applications.

Definition 14 Given $A \in \mathbb{R}^{n \times d}$ and a rank parameter $k \ll \min \{n, d\}$, let $\mathcal{S}_{\varepsilon}$ be the set of matrices of Definition 13. We call the numbers $\hat{p}_{i}$ (for all $i \in[n]$ ) $\beta$-approximations to the normalized leverage scores of $A_{k}$ (the best rank-k approximation to $A$ ) if, for some matrix $X \in \mathcal{S}_{\mathcal{E}}$,

$$
\hat{p}_{i} \geq \frac{\beta\left\|\left(U_{X}\right)_{(i)}\right\|_{2}^{2}}{k} \quad \text { and } \quad \sum_{i=1}^{n} \hat{p}_{i}=1 .
$$

Here $U_{X} \in \mathbb{R}^{n \times k}$ is the matrix of the left singular vectors of $X$.
Thus, we will seek algorithms whose output is a set of numbers, with the requirement that those numbers are good approximations to the normalized leverage scores of some matrix $X \in \mathcal{S}_{\varepsilon}$ (instead of $A_{k}$ ). This removes the ill-posedness of the original problem. Next, we will give two examples of algorithms that compute such $\beta$-approximations to the normalized leverage scores of a general matrix $A$ with a rank parameter $k$ for two popular norms, the spectral norm and the Frobenius norm. ${ }^{9}$

### 5.1 Leverage Scores for Spectral Norm Approximators

Algorithm 4 approximates the statistical leverage scores of a general matrix $A$ with rank parameter $k$ in the spectral norm case. It takes as inputs a matrix $A \in \mathrm{R}^{n \times d}$ with $\operatorname{rank}(A)=\rho$ and a rank parameter $k \ll \rho$, and outputs a set of numbers $\hat{p}_{i}$ for all $i \in[n]$, namely our approximations to the normalized leverage scores of $A$ with rank parameter $k$.

The next lemma argues that there exists a matrix $X \in \mathbb{R}^{n \times d}$ of rank $k$ that is sufficiently close to $A$ (in particular, it is a member of $\mathcal{S}_{\varepsilon}$ with constant probability) and, additionally, can be written as $X=$ $B Y$, where $Y \in \mathbb{R}^{2 k \times d}$ is a matrix of rank $k$. A version of this lemma was essentially proven in Halko et al. (2011), but see also Rokhlin et al. (2009) for computational details; we will use the version of the lemma that appeared in Boutsidis et al. (2011b). (See also the conference version (Boutsidis et al., 2011a), but in the remainder we refer to the technical report version (Boutsidis et al., 2011b) for consistency of numbering.) Note that for our purposes in this section, the computation of $Y$ is not relevant and we defer the reader to Halko et al. (2011) and Boutsidis et al. (2011b) for details.

[^3]Input: $A \in \mathbb{R}^{n \times d}$ with $\operatorname{rank}(A)=\rho$ and a rank parameter $k \ll \rho$
Output: $\hat{p}_{i}, i \in[n]$

1. Construct $\Pi \in \mathbb{R}^{d \times 2 k}$ with entries drawn in i.i.d. trials from the normal distribution $\mathcal{N}(0,1)$.
2. Compute $B=\left(A A^{T}\right)^{q} A \Pi \in \mathbb{R}^{n \times 2 k}$, with $q$ as in Equation (16).
3. Approximately compute the statistical leverage scores of the "tall" matrix $B$ by calling Algorithm 1 with inputs $B$ and $\varepsilon$; let $\hat{\ell}_{i}$ (for all $i \in[n]$ ) be the outputs of Algorithm 1.
4. Return

$$
\begin{equation*}
\hat{p}_{i}=\frac{\hat{\ell}_{i}}{\sum_{j=1}^{n} \hat{\ell}_{j}} \tag{15}
\end{equation*}
$$

for all $i \in[n]$.

Algorithm 4: Approximating the statistical leverage scores of a general matrix $A$ (spectral norm case).

Lemma 15 (Spectral Sketch) Given $A \in \mathbb{R}^{n \times d}$ of rank $\rho$, a rank parameter $k$ such that $2 \leq k<\rho$, and an error parameter $\varepsilon$ such that $0<\varepsilon<1$, let $\Pi \in \mathbb{R}^{d \times 2 k}$ be a standard Gaussian matrix (with entries selected in i.i.d. trials from $\mathfrak{N}(0,1)$ ). If $B=\left(A A^{T}\right)^{q} A \Pi$, where

$$
\begin{equation*}
q \geq\left\lceil\frac{\ln \left(1+\sqrt{\frac{k}{k-1}}+e \sqrt{\frac{2}{k}} \sqrt{\min \{n, d\}-k}\right)}{2 \ln (1+\varepsilon / 10)-1 / 2}\right] \tag{16}
\end{equation*}
$$

then there exists a matrix $X \in \mathbb{R}^{n \times d}$ of rank $k$ satisfying $X=B Y$ (with $Y \in \mathbb{R}^{2 k \times d}$ ) such that

$$
\mathbf{E}\left[\|A-X\|_{2}\right] \leq\left(1+\frac{\varepsilon}{10}\right)\left\|A-A_{k}\right\|_{2}
$$

The matrix $B$ can be computed in $O(n d k q)$ time.
This version of the above lemma is proven in Boutsidis et al. (2011b). ${ }^{10}$ Now, since $X$ has rank $k$, it follows that $\|A-X\|_{2} \geq\left\|A-A_{k}\right\|_{2}$ and thus we can consider the non-negative random variable $\|A-X\|_{2}-\left\|A-A_{k}\right\|_{2}$ and apply Markov’s inequality to get that

$$
\|A-X\|_{2}-\left\|A-A_{k}\right\|_{2} \leq \varepsilon\left\|A-A_{k}\right\|_{2}
$$

holds with probability at least 0.9 . Thus, $X \in \mathcal{S}_{\varepsilon}$ with probability at least 0.9 .
10. More specifically, the proof may be found in Lemma 32 and in particular in Equation (14) in Section A.2; note that for our purposes here we replaced $\varepsilon / \sqrt{2}$ by $\varepsilon / 10$ after adjusting $q$ accordingly.

The next step of the proposed algorithm is to approximately compute the leverage scores of $B \in \mathbb{R}^{n \times 2 k}$ via Algorithm 1. Under the assumptions of Theorem 2, this step runs in $O\left(n k \varepsilon^{-2} \ln n\right)$ time. Let $U_{X} \in \mathbb{R}^{n \times k}$ be the matrix containing the left singular vectors of the matrix $X$ of Lemma 15 . Then, since $X=B Y$ by Lemma 15, it follows that

$$
U_{B}=\left[\begin{array}{ll}
U_{X} & U_{R}
\end{array}\right]
$$

is a basis for the subspace spanned by the columns of $B$. Here $U_{R} \in \mathbb{R}^{n \times k}$ is an orthogonal matrix whose columns are perpendicular to the columns of $U_{X}$. Now consider the approximate leverage scores $\hat{\ell}_{i}$ computed by Algorithm 1 and note that (by Theorem 2),

$$
\left|\hat{\ell}_{i}-\left\|\left(U_{B}\right)_{(i)}\right\|_{2}^{2}\right| \leq \varepsilon\left\|\left(U_{B}\right)_{(i)}\right\|_{2}^{2}
$$

holds with probability at least 0.8 for all $i \in[n]$. It follows that

$$
\sum_{j=1}^{n} \hat{\ell}_{j} \leq(1+\varepsilon) \sum_{j=1}^{n}\left\|\left(U_{B}\right)_{(j)}\right\|_{2}^{2}=(1+\varepsilon) \sum_{j=1}^{n}\left\|U_{B}\right\|_{F}^{2}=2(1+\varepsilon) k
$$

Finally,

$$
\begin{aligned}
\hat{p}_{i}=\frac{\hat{\ell}_{i}}{\sum_{j=1}^{n} \hat{\ell}_{j}} & \geq(1-\varepsilon) \frac{\left\|\left(U_{B}\right)_{(i)}\right\|_{2}^{2}}{\sum_{j=1}^{n} \hat{\ell}_{j}} \\
& \geq(1-\varepsilon) \frac{\left\|\left(U_{X}\right)_{(i)}\right\|_{2}^{2}+\left\|\left(U_{R}\right)_{(i)}\right\|_{2}^{2}}{\sum_{j=1}^{n} \hat{\ell}_{j}} \\
& \geq \frac{1-\varepsilon}{2} \frac{\left\|\left(U_{X}\right)_{(i)}\right\|_{2}^{2}}{\sum_{j=1}^{n} \hat{\ell}_{j}} \\
& \geq \frac{1-\varepsilon}{2(1+\varepsilon)} \frac{\left\|\left(U_{X}\right)_{(i)}\right\|_{2}^{2}}{k} .
\end{aligned}
$$

Clearly, $\left\|\left(U_{X}\right)_{(i)}\right\|_{2}^{2} / k$ are the normalized leverage scores of the matrix $X$. Recall that $X \in \mathcal{S}_{\varepsilon}$ with probability at least 0.9 and use Definition 14 to conclude that the scores $\hat{p}_{i}$ of Equation (15) are $\left(\frac{1-\varepsilon}{2(1+\varepsilon)}\right)$-approximations to the normalized leverage scores of $A$ with rank parameter $k$. The following Theorem summarizes the above discussion:

Theorem 16 Given $A \in \mathbb{R}^{n \times d}$, a rank parameter $k$, and an accuracy parameter $\varepsilon$, Algorithm 4 computes a set of normalized leverage scores $\hat{p}_{i}$ that are $\left(\frac{1-\varepsilon}{2(1+\varepsilon)}\right)$-approximations to the normalized leverage scores of $A$ with rank parameter $k$ with probability at least 0.7 . The proposed algorithm runs in

$$
O\left(n d k \frac{\ln (\min \{n, d\})}{\ln (1+\varepsilon)}+n k \varepsilon^{-2} \ln n\right)
$$

time.

Input: $A \in \mathbb{R}^{n \times d}$ with $\operatorname{rank}(A)=\rho$ and a rank parameter $k \ll \rho$
Output: $\hat{p}_{i}, i \in[n]$

1. Let $r$ be as in Equation (18) and construct $\Pi \in \mathbb{R}^{d \times r}$ whose entries are drawn in i.i.d. trials from the normal distribution $\mathcal{N}(0,1)$.
2. Compute $B=A \Pi \in \mathbb{R}^{n \times r}$.
3. Compute a matrix $Q \in \mathbb{R}^{n \times r}$ whose columns form an orthonormal basis for the column space of $B$.
4. Compute the matrix $Q^{T} A \in \mathbb{R}^{r \times d}$ and its left singular vectors $U_{Q^{T} A} \in \mathbb{R}^{r \times d}$.
5. Let $U_{Q^{T} A, k} \in \mathbb{R}^{r \times k}$ denote the top $k$ left singular vectors of the matrix $Q^{T} A$ (the first $k$ columns of $U_{Q^{T} A}$ ) and compute, for all $i \in[n]$,

$$
\begin{equation*}
\hat{\ell}_{i}=\left\|\left(Q U_{Q^{T} A, k}\right)_{(i)}\right\|_{2}^{2} . \tag{17}
\end{equation*}
$$

6. Return $\hat{p}_{i}=\hat{\ell}_{i} / k$ for all $i \in[n]$.

Algorithm 5: Approximating the statistical leverage scores of a general matrix $A$ (Frobenius norm case).

### 5.2 Leverage Scores for Frobenius Norm Approximators

Algorithm 5 approximates the statistical leverage scores of a general matrix $A$ with rank parameter $k$ in the Frobenius norm case. It takes as inputs a matrix $A \in \mathrm{R}^{n \times d}$ with $\operatorname{rank}(A)=\rho$ and a rank parameter $k \ll \rho$, and outputs a set of numbers $\hat{p}_{i}$ for all $i \in[n]$, namely our approximations to the normalized leverage scores of $A$ with rank parameter $k$. It is worth noting that $\sum_{i=1}^{n} \hat{\ell}_{i}=\left\|Q U_{Q^{T} A, k}\right\|_{F}^{2}=\left\|U_{Q^{T} A, k}\right\|_{F}^{2}=k$ and thus the $\hat{p}_{i}$ sum up to one. The next lemma argues that there exists a matrix $X \in \mathbb{R}^{n \times d}$ of rank $k$ that is sufficiently close to $A$ (in particular, it is a member of $\mathcal{S}_{\varepsilon}$ with constant probability). Unlike the previous section (the spectral norm case), we will now be able to provide a closed-form formula for this matrix $X$ and, more importantly, the normalized leverage scores of $X$ will be exactly equal to the $\hat{p}_{i}$ returned by our algorithm. Thus, in the parlance of Definition 14, we will get a 1-approximation to the normalized leverage scores of $A$ with rank parameter $k$.

Lemma 17 (Frobenius Sketch) Given $A \in \mathbb{R}^{n \times d}$ of rank $\rho$, a rank parameter $k$ such that $2 \leq k<\rho$, and an error parameter $\varepsilon$ such that $0<\varepsilon<1$, let $\Pi \in \mathbb{R}^{d \times r}$ be a standard Gaussian matrix (with entries selected in i.i.d. trials from $\mathcal{N}(0,1)$ ) with

$$
\begin{equation*}
r \geq k+\left\lceil\frac{10 k}{\varepsilon}+1\right\rceil \tag{18}
\end{equation*}
$$

Let $B=A \Pi$ and let $X$ be as in Equation (19). Then,

$$
\mathbf{E}\left[\|A-X\|_{F}^{2}\right] \leq\left(1+\frac{\varepsilon}{10}\right)\left\|A-A_{k}\right\|_{F}^{2}
$$

The matrix B can be computed in $O\left(n d k \varepsilon^{-1}\right)$ time.
Let

$$
\begin{equation*}
X=Q\left(Q^{T} A\right)_{k} \in \mathbb{R}^{n \times d} \tag{19}
\end{equation*}
$$

where $\left(Q^{T} A\right)_{k}$ is the best rank- $k$ approximation to the matrix $Q^{T} A$; from standard linear algebra, $\left(Q^{T} A\right)_{k}=U_{Q^{T} A, k} U_{Q^{T} A, k}^{T} Q^{T} A$. Then, the above lemma is proven in Boutsidis et al. (2011b). ${ }^{11}$ Now, since $X$ has rank $k$, it follows that $\|A-X\|_{F}^{2} \geq\left\|A-A_{k}\right\|_{F}^{2}$ and thus we can consider the non-negative random variable $\|A-X\|_{F}^{2}-\left\|A-A_{k}\right\|_{F}^{2}$ and apply Markov's inequality to get that

$$
\|A-X\|_{F}^{2}-\left\|A-A_{k}\right\|_{F}^{2} \leq \varepsilon\left\|A-A_{k}\right\|_{F}^{2}
$$

holds with probability at least 0.9 . Rearranging terms and taking square roots of both sides implies that

$$
\|A-X\|_{F} \leq \sqrt{1+\varepsilon}\left\|A-A_{k}\right\|_{F} \leq(1+\varepsilon)\left\|A-A_{k}\right\|_{F} .
$$

Thus, $X \in \mathcal{S}_{\varepsilon}$ with probability at least 0.9 . To conclude our proof, recall that $Q$ is an orthonormal basis for the columns of $B$. From Equation (19),

$$
X=Q\left(Q^{T} A\right)_{k}=Q U_{Q^{T} A, k} U_{Q^{T} A, k}^{T} Q^{T} A=Q U_{Q^{T} A, k} \Sigma_{Q^{T} A, k} V_{Q^{T} A, k}^{T} .
$$

In the above, $\Sigma_{Q^{T} A, k} \in \mathbb{R}^{k \times k}$ is the diagonal matrix containing the top $k$ singular values of $Q^{T} A$ and $V_{Q^{T} A, k}^{T} \in \mathbb{R}^{k \times d}$ is the matrix whose rows are the top $k$ right singular vectors of $Q^{T} A$. Thus, the left singular vectors of the matrix $X$ are exactly equal to the columns of the orthogonal matrix $Q U_{Q^{T} A, k}$; it now follows that the $\hat{\ell}_{i}$ of Equation (17) are the leverage scores of the matrix $X$ and, finally, that the $\hat{p}_{i}$ returned by the proposed algorithm are the normalized leverage scores of the matrix $X$.

We briefly discuss the running time of the proposed algorithm. First, we can compute $B$ in $O(n d r)$ time. Then, the computation of $Q$ takes $O\left(n r^{2}\right)$ time. The computation of $Q^{T} A$ takes $O(n d r)$ time and the computation of $U_{Q^{T} A}$ takes $O\left(d r^{2}\right)$ time. Thus, the total time is equal to $O\left(n d r+(n+d) r^{2}\right)$. The following Theorem summarizes the above discussion.
Theorem 18 Given $A \in \mathbb{R}^{n \times d}$, a rank parameter $k$, and an accuracy parameter $\varepsilon$, Algorithm 5 computes a set of normalized leverage scores $\hat{p}_{i}$ that are 1-approximations to the normalized leverage scores of $A$ with rank parameter $k$ with probability at least 0.7 . The proposed algorithm runs in $O\left(n d k \varepsilon^{-1}+(n+d) k^{2} \varepsilon^{-2}\right)$ time.

## 6. Discussion

We will conclude with a discussion of our main results in a broader context: understanding the relationship between our main algorithm and a related estimator for the statistical leverage scores; applying our main algorithm to solve under-constrained least squares problems; and implementing variants of the basic algorithm in streaming environments.
11. More specifically, the proof may be found in Lemma 33 in Section A.3; note that for our purposes here we set $p=\left\lceil\frac{10 k}{\varepsilon}+1\right\rceil$.

### 6.1 A Related Estimator for the Leverage Scores

Magdon-Ismail (2010) presented the following algorithm to estimate the statistical leverage scores: given as input an $n \times d$ matrix $A$, with $n \gg d$, the algorithm proceeds as follows.

- Compute $\Pi A$, where the $O\left(\frac{n \ln d}{\ln ^{2} n}\right) \times n$ matrix $\Pi$ is a SRHT or another FJLT.
- Compute $X=(\Pi A)^{\dagger} \Pi$.
- For $t=1, \ldots, n$, compute the estimate $\tilde{w}_{t}=A_{(t)}^{T} X^{(t)}$ and set $w_{t}=\max \left\{\frac{d \ln ^{2} n}{4 n}, \tilde{w}_{t}\right\}$.
- Return the quantities $\tilde{p}_{i}=w_{i} / \sum_{i^{\prime}=1}^{n} w_{i^{\prime}}$, for $i \in[n]$.

Magdon-Ismail (2010) argued that the output $\tilde{p}_{i}$ achieves an $O\left(\ln ^{2} n\right)$ approximation to all of the (normalized) statistical leverage scores of $A$ in roughly $O\left(n d^{2} / \ln n\right)$ time. (To our knowledge, prior to our work here, this is the only known estimator that obtains any nontrivial provable approximation to the leverage scores of a matrix in $o\left(n d^{2}\right)$ time.) To see the relationship between this estimator and our main result, recall that

$$
\ell_{i}=e_{i}^{T} U U^{T} e_{i}=e_{i}^{T} A A^{\dagger} e_{i}=x_{i}^{T} y_{i},
$$

where the vector $x_{i}^{T}=e_{i}^{T} A$ is cheap to compute and the vector $y_{i}=A^{\dagger} e_{i}$ is expensive to compute. The above algorithm effectively approximates $y_{i}=A^{\dagger} e_{i}$ via a random projection as $\tilde{y}_{i}=(\Pi A)^{\dagger} \Pi e_{i}$, where $\Pi$ is a SRHT or another FJLT. Since the estimates $x_{i}^{T} \tilde{y}_{i}$ are not necessarily positive, a truncation at the negative tail, followed by a renormalization step, must be performed in order to arrive at the final estimator returned by the algorithm. This truncation-renormalization step has the effect of inflating the estimates of the small leverage scores by an $O\left(\ln ^{2} n\right)$ factor. By way of comparison, Algorithm 1 essentially computes a sketch of $A A^{\dagger}$ of the form $A(\Pi A)^{\dagger} \Pi^{T}$ that maintains positivity for each of the row norm estimates.

Although both Algorithm 1 and the algorithm of this subsection estimate $A A^{\dagger}$ by a matrix of the form $A(\Pi A)^{\dagger} \Pi^{T}$, there are notable differences. The algorithm of this subsection does not actually compute or approximate $A A^{T}$ directly; instead, it separates the matrix into two parts and computes the dot product between $e_{i}^{T} A$ and $(\Pi A)^{\dagger} \Pi e_{i}$. Positivity is sacrificed and this leads to some complications in the estimator; however, the truncation step is interesting, since, despite the fact that the estimates are "biased" (in a manner somewhat akin to what is obtained with "thresholding" or "regularization" procedures), we still obtain provable approximation guarantees. The algorithm of this subsection is simpler (since it uses an application of only one random projection), albeit at the cost of weaker theoretical guarantees and a worse running time than our main algorithm. A direction of considerable practical interest is to evaluate empirically the performance of these two estimators, either for estimating all the leverage scores or (more interestingly) for estimating the largest leverage scores for data matrices for which the leverage scores are quite nonuniform.

### 6.2 An Application to Under-constrained Least-squares Problems

Consider the following under-constrained least-squares problem:

$$
\begin{equation*}
\min _{x \in \mathbb{R}^{d}}\|A x-b\|_{2}, \tag{20}
\end{equation*}
$$

where $A \in \mathbb{R}^{n \times d}$ has much fewer rows than columns, that is, $n \ll d$. It is well-known that we can solve this problem exactly in $O\left(n^{2} d\right)$ time and that the minimal $\ell_{2}$-norm solution is given by $x_{\text {opt }}=A^{\dagger} b$. For simplicity, let's assume that the input matrix $A$ has full $\operatorname{rank}$ (i.e., $\operatorname{rank}(A)=n$ ) and thus $\left\|A x_{\text {opt }}-b\right\|_{2}=0$.

In this section, we will argue that Algorithm 6 computes a simple, accurate estimator $\tilde{x}_{\text {opt }}$ for $x_{\text {opt }}$. In words, Algorithm 6 samples a small number of columns from $A$ (note that the columns of $A$ correspond to variables in our under-constrained problem) and uses the sampled columns to compute $\tilde{x}_{\text {opt }}$. However, in order to determine which columns will be included in the sample, the algorithm will make use of the statistical leverage scores of the matrix $A^{T}$; more specifically, columns (and thus variables) will be chosen with probability proportional to the corresponding statistical leverage score. We will state Algorithm 6 assuming that these probabilities are parts of the input; the following theorem is our main quality-of-approximation result for Algorithm 6.

Theorem 19 Let $A \in \mathbb{R}^{n \times d}$ be a full-rank matrix with $n \ll d$; let $\varepsilon \in(0,0.5]$ be an accuracy parameter; let $\delta \in(0,1)$ be a failure probability; and let $x_{\text {opt }}=A^{\dagger} b$ be the minimal $\ell_{2}$-norm solution to the least-squares problem of Equation (20). Let $p_{i} \geq 0, i \in[d]$, be a set of probabilities satisfying $\sum_{i=1}^{d} p_{i}=1$ and

$$
\begin{equation*}
p_{i} \geq \frac{\beta\left\|V_{(i)}\right\|_{2}^{2}}{n} \tag{21}
\end{equation*}
$$

for some constant $\beta \in(0,1]$. (Here $V \in \mathbb{R}^{d \times n}$ is the matrix of the right singular vectors of $A$.) If $\tilde{x}_{\text {opt }}$ is computed via Algorithm 6 then, with probability at least $1-\delta$,

$$
\left\|x_{\text {opt }}-\tilde{x}_{\text {opt }}\right\|_{2} \leq 2 \varepsilon\left\|x_{\text {opt }}\right\|_{2} .
$$

Algorithm 6 runs in $O\left(n^{3} \varepsilon^{-2} \beta^{-1} \ln (n / \varepsilon \beta \delta)+n d\right)$ time.
Proof Let the singular value decomposition of the full-rank matrix $A$ be $A=U \Sigma V^{T}$, with $U \in \mathbb{R}^{n \times n}$, $\Sigma \in \mathbb{R}^{n \times n}$, and $V \in \mathbb{R}^{d \times n}$; note that all the diagonal entries of $\Sigma$ are strictly positive since $A$ has full rank. We can now apply Theorem 4 of Section 6.1 of Drineas et al. (2010b) to get ${ }^{12}$ that

$$
\left\|I_{n}-V^{T} S S^{T} V\right\|_{2}=\left\|V^{T} V-V^{T} S S^{T} V\right\|_{2} \leq \varepsilon
$$

for our choice of $r$ with probability at least $1-\delta$. Note that $V^{T} S \in \mathbb{R}^{n \times r}$ (with $r \geq n$ ) and let $\sigma_{i}\left(V^{T} S\right)$ denote the singular values of $V^{T} S$ for all $i \in[n]$; the above inequality implies that for all $i \in[n]$

$$
\left|1-\sigma_{i}^{2}\left(V^{T} S\right)\right| \leq\left\|I_{n}-V^{T} S S^{T} V\right\|_{2} \leq \varepsilon \leq 0.5 .
$$

Thus, all the singular values of $V^{T} S$ are strictly positive and hence $V^{T} S$ has full rank equal to $n$. Also, using $\varepsilon \leq 0.5$,

$$
\begin{equation*}
\left|1-\sigma_{i}^{-2}\left(V^{T} S\right)\right| \leq \frac{\varepsilon}{1-\varepsilon} \leq 2 \varepsilon \tag{22}
\end{equation*}
$$

12. We apply Theorem 4 of Section 6.1 of Drineas et al. (2010b) with $A=V^{T}$ and note that $\left\|V^{T}\right\|_{F}^{2}=n \geq 1,\left\|V^{T}\right\|_{2}=1$, and $\left(V^{T}\right)^{(i)}=V_{(i)}$.

We are now ready to prove our theorem:

$$
\begin{aligned}
\left\|x_{\text {opt }}-\tilde{x}_{\text {opt }}\right\|_{2} & =\left\|A^{T}(A S)^{\dagger T}(A S)^{\dagger} b-A^{\dagger} b\right\|_{2} \\
& =\left\|V \Sigma U^{T}\left(U \Sigma V^{T} S\right)^{\dagger T}\left(U \Sigma V^{T} S\right)^{\dagger} b-V \Sigma^{-1} U^{T} b\right\|_{2} \\
& =\left\|\Sigma U^{T} U \Sigma^{-1}\left(V^{T} S\right)^{\dagger T}\left(V^{T} S\right)^{\dagger} \Sigma^{-1} U^{T} b-\Sigma^{-1} U^{T} b\right\|_{2} \\
& =\left\|\left(V^{T} S\right)^{\dagger T}\left(V^{T} S\right)^{\dagger} \Sigma^{-1} U^{T} b-\Sigma^{-1} U^{T} b\right\|_{2} .
\end{aligned}
$$

In the above derivations we substituted the SVD of $A$, dropped terms that do not change unitarily invariant norms, and used the fact that $V^{T} S$ and $\Sigma$ have full rank in order to simplify the pseudoinverse. Now let $\left(V^{T} S\right)^{\dagger T}\left(V^{T} S\right)^{\dagger}=I_{n}+E$ and note that Equation (22) and the fact that $V^{T} S$ has full rank imply

$$
\|E\|_{2}=\left\|I_{n}-\left(V^{T} S\right)^{\dagger T}\left(V^{T} S\right)^{\dagger}\right\|_{2}=\max _{i \in[n]}\left|1-\sigma_{i}^{-2}\left(V^{T} S\right)\right| \leq 2 \varepsilon .
$$

Thus, we conclude our proof by observing that

$$
\begin{aligned}
\left\|x_{\text {opt }}-\tilde{x}_{\text {opt }}\right\|_{2} & =\left\|\left(I_{n}+E\right) \Sigma^{-1} U^{T} b-\Sigma^{-1} U^{T} b\right\|_{2} \\
& =\left\|E \Sigma^{-1} U^{T} b\right\|_{2} \\
& \leq\|E\|_{2}\left\|\Sigma^{-1} U^{T} b\right\|_{2} \\
& \leq 2 \varepsilon\left\|x_{\text {opt }}\right\|_{2} .
\end{aligned}
$$

In the above we used the fact that $\left\|x_{\text {opt }}\right\|_{2}=\left\|A^{\dagger} b\right\|_{2}=\left\|V \Sigma^{-1} U^{T} b\right\|_{2}=\left\|\Sigma^{-1} U^{T} b\right\|_{2}$. The running time of the algorithm follows by observing that $A S$ is an $n \times r$ matrix and thus computing its pseudoinverse takes $O\left(n^{2} r\right)$ time; computing $x_{\text {opt }}$ takes an additional $O(n r+d n)$ time.

We conclude the section with a few remarks. First, assuming that $\varepsilon, \beta$, and $\delta$ are constants and $n \ln n=o(d)$, it immediately follows that Algorithm 6 runs in $o\left(n^{2} d\right)$ time. It should be clear that we can use Theorem 2 and the related Algorithm 1 to approximate the statistical leverage scores, thus bypassing the need to exactly compute them. Second, instead of approximating the statistical leverage scores needed in Algorithm 6, we could use the randomized Hadamard transform (essentially post-multiply $A$ by a randomized Hadamard transform to make all statistical leverage scores uniform). The resulting algorithm could be theoretically analyzed following the lines of Drineas et al. (2010b). It would be interesting to evaluate experimentally the performance of the two approaches in real data.

### 6.3 Extension to Streaming Environments

In this section, we consider the estimation of the leverage scores and of related statistics when the input data set is so large that an appropriate way to view the data is as a data stream (Muthukrishnan, 2005). In this context, one is interested in computing statistics of the data stream while making one pass (or occasionally a few additional passes) over the data from external storage and using only a small amount of additional space. For an $n \times d$ matrix $A$, with $n \gg d$, small additional space means that the space complexity only depends logarithmically on the high dimension $n$ and polynomially on the low dimension $d$. When we discuss bits of space, we assume that the entries of $A$ can be

Input: $A \in \mathbb{R}^{n \times d}, b \in \mathbb{R}^{n}$, error parameter $\varepsilon \in(0, .5]$, failure probability $\delta$, and a set of probabilities $p_{i}$ (for all $i \in[d]$ ) summing up to one and satisfying Equation (21).

Output: $\tilde{x}_{o p t} \in \mathbb{R}^{d}$.

1. Let $r=\frac{96 n}{\beta \varepsilon^{2}} \ln \left(\frac{96 n}{\beta \varepsilon^{2} \sqrt{\delta}}\right)$.
2. Let $S \in \mathbb{R}^{d \times r}$ be an all-zeros matrix.
3. For $t=1, \ldots, r$ do

- Pick $i_{t} \in[d]$ such that $\operatorname{Pr}\left(i_{t}=i\right)=p_{i}$.
- $S_{i_{t} t}=1 / \sqrt{r p_{i_{t}}}$.

4. Return $\tilde{x}_{o p t}=A^{T}(A S)^{\dagger T}(A S)^{\dagger} b$.

Algorithm 6: Approximately solving under-constrained least squares problems.
discretized to $O(\log n)$ bit integers, though all of our results can be generalized to arbitrary word sizes. The general strategy behind our algorithms is as follows.

- As the data streams by, compute $T A$, for an appropriate problem-dependent linear sketching matrix $T$, and also compute $\Pi A$, for a random projection matrix $\Pi .{ }^{13}$
- After the first pass over the data, compute the matrix $R^{-1}$, as described in Algorithm 1, corresponding to $\Pi A$ (or compute the pseudoinverse of $\Pi A$ or the $R$ matrix from any other QR decomposition of $A$ ).
- Compute $T A R^{-1} \Pi_{2}$, for a random projection matrix $\Pi_{2}$, such as the one used by Algorithm 1 .

With the procedure outlined above, the matrix $T$ is effectively applied to the rows of $A R^{-1} \Pi_{2}$, that is, to the sketch of $A$ that has rows with Euclidean norms approximately equal to the row norms of $U$, and pairwise inner products approximately equal to those in $U$. Thus statistics related to $U$ can be extracted.

Large Leverage Scores. Given any $n \times d$ matrix $A$ in a streaming setting, it is known how to find the indices of all rows $A_{(i)}$ of $A$ for which $\left\|A_{(i)}\right\|_{2}^{2} \geq \tau\|A\|_{F}^{2}$, for a parameter $\tau$, and in addition it is known how to compute a $(1+\varepsilon)$-approximation to $\left\|A_{(i)}\right\|_{2}^{2}$ for these large rows. The basic idea is to use the notion of $\ell_{2}$-sampling on matrix $A$, namely, to sample random entries $A_{i j}$ with probability $A_{i j}^{2} /\|A\|_{F}^{2}$. A single entry can be sampled from this distribution in a single pass using $O\left(\varepsilon^{-2} \log ^{3}(n d)\right)$ bits of space (Monemizadeh and Woodruff, 2010; Andoni et al., 2010). More precisely, these references demonstrate that there is a distribution over $O\left(d \varepsilon^{-2} \log ^{3}(n d)\right) \times n$ matrices

[^4]$T$ for which for any fixed matrix $A \in \mathbb{R}^{n \times d}$, there is a procedure which given $T A$, outputs a sample $(i, j) \in[n] \times[d]$ with probability $(1 \pm \varepsilon) \frac{A_{i, j}^{2}}{\|A\|_{F}^{2}} \pm n^{-O(1)}$. Technically, these references concern sampling from vectors rather than matrices, so $T(A)$ is a linear operator which treats $A$ as a length- $n d$ vector and applies the algorithm of Monemizadeh and Woodruff (2010) and Andoni et al. (2010). However, by simply increasing the number of rows in $T$ by a factor of the small dimension $d$, we can assume $T$ is left matrix multiplication. By considering the marginal along [n], the probability that $i=a$, for any $a \in[n]$, is
$$
(1 \pm \varepsilon) \frac{\left\|U_{(a)}\right\|_{2}^{2}}{\|U\|_{F}^{2}} \pm(n d)^{-O(1)} .
$$

By the coupon collector problem, running $O\left(\tau^{-1} \log \tau^{-1}\right)$ independent copies is enough to find a set containing all rows $A_{(i)}$ for which $\left\|A_{(i)}\right\|_{2}^{2} \geq \tau\|A\|_{F}^{2}$, and no rows $A_{(i)}$ for which $\left\|A_{(i)}\right\|_{2}^{2}<\frac{\tau}{2}\|A\|_{F}^{2}$ with probability at least 0.99 .

When applied to our setting, we can apply a random projection matrix $\Pi$ and a linear sketching matrix $T$ which has $O\left(d \tau^{-1} \varepsilon^{-2} \log ^{3}(n) \log \tau^{-1}\right)$ rows in the following manner. First, $T A$ and $\Pi A$ are computed in the first pass over the data; then, at the end of the first pass, we compute $R^{-1}$; and finally, we compute $T A R^{-1} \Pi_{2}$, for a random projection matrix $\Pi_{2}$. This procedure effectively applies the matrix $T$ to the rows of $A R^{-1} \Pi_{2}$, which have norms equal to the row norms of $U$, up to a factor of $1+\varepsilon$. The multiplication at the end by $\Pi_{2}$ serves only to speed up the time for processing $T A R^{-1}$. Thus, by the results of Monemizadeh and Woodruff (2010) and Andoni et al. (2010), we can find all the leverage scores $\left\|U_{(i)}\right\|_{2}^{2}$ that are of magnitude at least $\tau\|U\|_{F}^{2}$ in small space and a single pass over the data. By increasing the space by a factor of $O\left(\varepsilon^{-2} \log n\right)$, we can also use the $\ell_{2}$-samples to estimate the norms $\left\|U_{(i)}\right\|_{2}^{2}$ for the row indices $i$ that we find.

Entropy. Given a distribution $\rho$, a statistic of $\rho$ of interest is the entropy of this distribution, where the entropy is defined as $H(\rho)=\sum_{i} \rho(i) \log _{2}(1 / \rho(i))$. This statistic can be approximated in a streaming setting. Indeed, it is known that estimating $H(\rho)$ up to an additive $\varepsilon$ can be reduced to $(1+\tilde{\varepsilon})$-approximation of the $\ell_{p}$-norm of the vector $(\rho(1), \ldots, \rho(n))$, for $O(\log 1 / \varepsilon)$ different $p \in(0,1)$ (Harvey et al., 2008). Here $\tilde{\varepsilon}=\varepsilon /\left(\log ^{3} 1 / \varepsilon \cdot \log n\right)$. When applied to our setting, the distribution of interest is $\rho(i)=\frac{1}{d}\left\|U_{(i)}\right\|_{2}^{2}$. To compute the entropy of this distribution, there exist sketching matrices $T$ for providing $(1+\varepsilon)$-approximations to the quantity $F_{p}\left(F_{2}\right)$ of an $n \times d$ matrix $A$, where $F_{p}\left(F_{2}\right)$ is defined as $\sum_{i=1}^{n}\left\|A_{(i)}\right\|_{2}^{2 p}$, using $O\left(\varepsilon^{-4} \log ^{2} n \log 1 / \varepsilon\right)$ bits of space (see Theorem 1 of Ganguly et al. (2008)). Thus, to compute the entropy of the leverage score distribution, we can do the following. First, maintain $T A$ and $\Pi A$ in the first pass over the data, where $T$ is a sketching matrix for $F_{p}\left(F_{2}\right), p \in(0,1)$. At the end of the first pass, compute $R^{-1}$; and finally, compute $T A R^{-1} \Pi_{2}$, which effectively applies the $F_{p}\left(F_{2}\right)$-estimation matrix $T$ to the rows of the matrix $A R^{-1} \Pi_{2}$. Therefore, by the results of Harvey et al. (2008) and Ganguly et al. (2008), we can compute an estimate $\phi$ which is within an additive $\varepsilon$ of $H(\rho)$ using $O\left(d \varepsilon^{-4} \log ^{6} n \log ^{14} 1 / \varepsilon\right)$ bits of space and a single pass. We note that it is also possible to estimate $H(\rho)$ up to a multiplicative $1+\varepsilon$ factor using small, but more, space; see, for example, Harvey et al. (2008).

Sampling Row Identities. Another natural problem is that of obtaining samples of rows of A proportional to their leverage score importance sampling probabilities. To do so, we use $\ell_{2}$ sampling (Monemizadeh and Woodruff, 2010; Andoni et al., 2010) as used above for finding the large leverage scores. First, compute $T A$ and $\Pi A$ in the first pass over the data stream; then, compute $R^{-1}$; and finally, compute $T A R^{-1}$. Thus, by applying the procedures of Andoni et al. (2010) a total of $s$ times independently, we obtain $s$ samples $i_{1}, \ldots, i_{s}$, with replacement, of rows of $A$ proportional
to $\left\|U_{\left(i_{1}\right)}\right\|_{2}^{2}, \ldots,\left\|U_{\left(i_{s}\right)}\right\|_{2}^{2}$, that is, to their leverage score. The algorithm requires $O\left(s d \varepsilon^{-2} \log ^{4} n\right)$ bits of space and runs in a single pass. To obtain more than just the row identities $i_{1}, \ldots, i_{s}$, for example, to obtain the actual samples, one can read off these rows from $A$ in a second pass over the matrix.

## References

D. Achlioptas. Database-friendly random projections: Johnson-lindenstrauss with binary coins. Journal of Computer and System Sciences, 66(4):671-687, 2003.
N. Ailon and B. Chazelle. Approximate nearest neighbors and the fast Johnson-Lindenstrauss transform. In Proceedings of the 38th Annual ACM Symposium on Theory of Computing, pages 557563, 2006.
N. Ailon and B. Chazelle. The fast Johnson-Lindenstrauss transform and approximate nearest neighbors. SIAM Journal on Computing, 39(1):302-322, 2009.
N. Ailon and E. Liberty. Fast dimension reduction using Rademacher series on dual BCH codes. In Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1-9, 2008.
A. Andoni, R. Krauthgamer, and K. Onak. Streaming algorithms from precision sampling. Technical report, 2010. Preprint: arXiv:1011.1263.
H. Avron, P. Maymounkov, and S. Toledo. Blendenpik: Supercharging LAPACK's least-squares solver. SIAM Journal on Scientific Computing, 32:1217-1236, 2010.
C. Bekas, E. Kokiopoulou, and Y. Saad. An estimator for the diagonal of a matrix. Applied Numerical Mathematics, 57:1214-1229, 2007.
C. Bekas, A. Curioni, and I. Fedulova. Low cost high performance uncertainty quantification. In Proceedings of the 2nd Workshop on High Performance Computational Finance, page Article No.: 8, 2009.
P. Bonacich. Power and centrality: A family of measures. The American Journal of Sociology, 92 (5):1170-1182, 1987.
C. Boutsidis, M.W. Mahoney, and P. Drineas. An improved approximation algorithm for the column subset selection problem. In Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms, pages 968-977, 2009.
C. Boutsidis, P. Drineas, and M. Magdon-Ismail. Near-optimal column-based matrix reconstruction. In Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science, pages 305-314, 2011a.
C. Boutsidis, P. Drineas, and M. Magdon-Ismail. Near-optimal column-based matrix reconstruction. Technical report, 2011b. Preprint: arXiv:1103.0995.
E.J. Candes and B. Recht. Exact matrix completion via convex optimization. Technical report, 2008. Preprint: arXiv:0805.4471.
S. Chatterjee and A.S. Hadi. Influential observations, high leverage points, and outliers in linear regression. Statistical Science, 1(3):379-393, 1986.
S. Chatterjee and A.S. Hadi. Sensitivity Analysis in Linear Regression. John Wiley \& Sons, New York, 1988.
S. Chatterjee, A.S. Hadi, and B. Price. Regression Analysis by Example. John Wiley \& Sons, New York, 2000.
A. Dasgupta, R. Kumar, and T. Sarlós. A sparse Johnson-Lindenstrauss transform. In Proceedings of the 42nd Annual ACM Symposium on Theory of Computing, pages 341-350, 2010.
P. Drineas, R. Kannan, and M.W. Mahoney. Fast Monte Carlo algorithms for matrices I: Approximating matrix multiplication. SIAM Journal on Computing, 36:132-157, 2006 a.
P. Drineas, M.W. Mahoney, and S. Muthukrishnan. Sampling algorithms for $\ell_{2}$ regression and applications. In Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1127-1136, 2006b.
P. Drineas, M.W. Mahoney, and S. Muthukrishnan. Relative-error CUR matrix decompositions. SIAM Journal on Matrix Analysis and Applications, 30:844-881, 2008.
P. Drineas, J. Lewis, and P. Paschou. Inferring geographic coordinates of origin for Europeans using small panels of ancestry informative markers. PLoS ONE, 5(8): 11892, 2010a.
P. Drineas, M.W. Mahoney, S. Muthukrishnan, and T. Sarlós. Faster least squares approximation. Numerische Mathematik, 117(2):219-249, 2010b.
S. Ganguly, M. Bansal, and S. Dube. Estimating hybrid frequency moments of data streams. In Proceedings of the 2nd Annual International Workshop on Frontiers in Algorithmics, pages 5566, 2008.
S. Georgiev and S. Mukherjee, 2011. Unpublished results. (2011).
A. Gittens and M. W. Mahoney, 2012. In preparation. (2012).
G.H. Golub and C.F. Van Loan. Matrix Computations. Johns Hopkins University Press, Baltimore, 1996.
N. Halko, P.-G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. SIAM Review, 53(2):217-288, 2011.
N.J.A. Harvey, J. Nelson, and K. Onak. Sketching and streaming entropy via approximation theory. In Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science, pages 489-498, 2008.
D.C. Hoaglin and R.E. Welsch. The hat matrix in regression and ANOVA. The American Statistician, 32(1):17-22, 1978.
E. A. Jonckheere, M. Lou, J. Hespanha, and P. Barooah. Effective resistance of Gromov-hyperbolic graphs: Application to asymptotic sensor network problems. In Proceedings of the 46th IEEE Conference on Decision and Control, pages 1453-1458, 2007.
M. Magdon-Ismail. Row sampling for matrix algorithms via a non-commutative Bernstein bound. Technical report, 2010. Preprint: arXiv:1008.0587.
M. W. Mahoney. Randomized Algorithms for Matrices and Data. Foundations and trends in machine learning. NOW Publishers, Boston, 2011. Also available at: arXiv:1104.5557.
M.W. Mahoney and P. Drineas. CUR matrix decompositions for improved data analysis. Proc. Natl. Acad. Sci. USA, 106:697-702, 2009.
X. Meng, M. A. Saunders, and M. W. Mahoney. LSRN: A parallel iterative solver for strongly overor under-determined systems. Technical report, 2011. Preprint: arXiv:1109.5981.
M. Monemizadeh and D. P. Woodruff. 1-pass relative-error $l_{p}$-sampling with applications. In Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1143-1160, 2010.
S. Muthukrishnan. Data Streams: Algorithms and Applications. Foundations and Trends in Theoretical Computer Science. Now Publishers Inc, Boston, 2005.
M.Z. Nashed, editor. Generalized Inverses and Applications. Academic Press, New York, 1976.
M.E.J. Newman. A measure of betweenness centrality based on random walks. Social Networks, 27:39-54, 2005.
C.H. Papadimitriou, P. Raghavan, H. Tamaki, and S. Vempala. Latent semantic indexing: a probabilistic analysis. Journal of Computer and System Sciences, 61(2):217-235, 2000.
P. Paschou, E. Ziv, E.G. Burchard, S. Choudhry, W. Rodriguez-Cintron, M.W. Mahoney, and P. Drineas. PCA-correlated SNPs for structure identification in worldwide human populations. PLoS Genetics, 3:1672-1686, 2007.
P. Paschou, J. Lewis, A. Javed, and P. Drineas. Ancestry informative markers for finescale individual assignment to worldwide populations. Journal of Medical Genetics, page doi:10.1136/jmg.2010.078212, 2010.
M. Richardson and P. Domingos. Mining knowledge-sharing sites for viral marketing. In Proceedings of the 8th Annual ACM SIGKDD Conference, pages 61-70, 2002.
V. Rokhlin and M. Tygert. A fast randomized algorithm for overdetermined linear least-squares regression. Proc. Natl. Acad. Sci. USA, 105(36):13212-13217, 2008.
V. Rokhlin, A. Szlam, and M. Tygert. A randomized algorithm for principal component analysis. SIAM Journal on Matrix Analysis and Applications, 31(3):1100-1124, 2009.
T. Sarlós. Improved approximation algorithms for large matrices via random projections. In Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science, pages 143152, 2006.
A. Talwalkar and A. Rostamizadeh. Matrix coherence and the Nyström method. In Proceedings of the 26th Conference in Uncertainty in Artificial Intelligence, 2010.
P.F. Velleman and R.E. Welsch. Efficient computing of regression diagnostics. The American Statistician, 35(4):234-242, 1981.


[^0]:    3. LAPACK (short for Linear Algebra PACKage) is a high-quality and widely-used software library of numerical routines for solving a wide range of numerical linear algebra problems.
[^1]:    4. When no confusion can arise, we will use $\Pi_{J L T}$ to refer to this distribution over matrices as well as to a specific matrix drawn from this distribution.
[^2]:    5. Note that the RHT has also been crucial in the development of $o\left(n d^{2}\right)$ randomized algorithms for the general overconstrained LS problem (Drineas et al., 2010b) and its variants have been used to provide high-quality numerical implementations of such randomized algorithms (Rokhlin and Tygert, 2008; Avron et al., 2010).
    6. Again, when no confusion can arise, we will use $\Pi_{F J L T}$ to denote a specific SRHT or the distribution on matrices implied by the randomized process for constructing an SRHT.
[^3]:    9. Note that we will not compute $S_{\varepsilon}$, but our algorithms will compute a matrix in that set. Moreover, that matrix can be used for high-quality low-rank matrix approximation. See the comments in Section 1.4 for more details.
[^4]:    13. In the offline setting, one would use an SRHT or another FJLT, while in the streaming setting one could use either of the following. If the stream is such that one sees each entire column of $A$ at once, then one could do an FJLT on the column. Alternatively, if one see updates to the individual entries of $A$ in an arbitrary order, then one could apply any sketching matrix, such as those of Achlioptas (2003) or of Dasgupta et al. (2010).
