

Efficient Frequent Directions Algorithm for Sparse Matrices

Mina Ghashami
University of Utah
Utah, USA
ghashami@cs.utah.edu

Edo Liberty
Yahoo Research!
New York, USA
edo@yahoo-inc.com

Jeff M. Phillips^{*}
University of Utah
Utah, USA
jeffp@cs.utah.edu

ABSTRACT

This paper describes Sparse Frequent Directions, a variant of Frequent Directions for sketching sparse matrices. It resembles the original algorithm in many ways: both receive the rows of an input matrix $A^{n \times d}$ one by one in the streaming setting and compute a small sketch $B \in \mathbb{R}^{\ell \times d}$. Both share the same strong (provably optimal) asymptotic guarantees with respect to the space-accuracy tradeoff in the streaming setting. However, unlike Frequent Directions which runs in $O(nd\ell)$ time regardless of the sparsity of the input matrix A , Sparse Frequent Directions runs in $\tilde{O}(\text{nnz}(A)\ell + n\ell^2)$ time. Our analysis loosens the dependence on computing the Singular Value Decomposition (SVD) as a black box within the Frequent Directions algorithm. Our bounds require recent results on the properties of fast approximate SVD computations. Finally, we empirically demonstrate that these asymptotic improvements are practical and significant on real and synthetic data.

Keywords

Frequent Directions, Sparse Matrix, Matrix Sketching

1. INTRODUCTION

It is very common to represent data in the form of a matrix. For example, in text analysis under the bag-of-words model, a large corpus of documents can be represented as a matrix whose rows refer to the documents and columns correspond to words. A non-zero in the matrix corresponds to a word appearing in the corresponding document. Similarly, in recommendation systems [13], preferences of users are represented as a matrix with rows corresponding to users and columns corresponding to items. Non-zero entries correspond to user ratings or actions.

A large set of data analytic tasks rely on obtaining a low-rank approximation of the data matrix. These include clus-

^{*}Thanks to support by NSF CCF-1350888, IIS-1251019, ACI-1443046, and CNS-1514520.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '16, August 13 - 17, 2016, San Francisco, CA, USA

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4232-2/16/08...\$15.00

DOI: <http://dx.doi.org/10.1145/2939672.2939800>

tering, dimension reduction, principal component analysis (PCA), signal denoising, etc. Such approximations can be computed using the Singular Value Decompositions (SVD). For an $n \times d$ matrix A ($d \leq n$) computing the SVD requires $O(nd^2)$ time and $O(nd)$ space in memory on a single machine.

In many scenarios, however, data matrices are extremely large and computing their SVD exactly is infeasible. Efficient approximate solutions exist for distributed setting or when data access otherwise is limited. In the row streaming model, the matrix rows are presented to the algorithm one by one in an arbitrary order. The algorithm is tasked with processing the stream in one pass while being severely restricted in its memory footprint. At the end of the stream, the algorithm must provide a sketch matrix B which is a good approximation of A even though it is significantly more compact. This is called matrix sketching.

Matrix sketching methods are designed to be parallelizable, space and time efficient, and easily updatable. Computing the sketch on each machine and then combining the sketches together should be as good as sketching the combined data from all the different machines. The streaming model is especially attractive since a sketch can be obtained and maintained as the data is being collected. Therefore, eliminating the need for data storage altogether.

Often matrices, as above, are sparse; most of their entries are zero. The work of [9] argues that typical term-document matrices are sparse; documents contain no more than 5% of all words. On wikipedia, most words appear on only a small constant number of pages. Similarly, in recommendation systems, in average a user rates or interacts with a small fraction of the available items: less than 6% in some user-movies recommendation tasks [1] and much fewer in physical purchases or online advertising. As such, most of these datasets are stored as sparse matrices.

There exist several techniques for producing low rank approximations of sparse matrices whose running time is $O(\text{nnz}(A) \text{poly}(k, 1/\epsilon))$ for some error parameter $\epsilon \in (0, 1)$. Here $\text{nnz}(A)$ denotes the number of non-zeros in the matrix A . Examples include the power method [19], random projection techniques [35], projection-hashing [6], and instances of column selection techniques [11].

However, for a recent and popular technique FREQUENT-DIRECTIONS (best paper of KDD 2013 [24]), there is no known way to take advantage of the sparsity of the input matrix. While it is deterministic and its space-error bounds are known to be optimal for dense matrices in the row-update model [17], it runs in $O(nd\ell)$ time to produce a sketch of

size $\ell \times d$. In particular, it maintains a sketch with ℓ rows and updates it iteratively over a stream, periodically invoking a full SVD which requires $O(d\ell^2)$ time. Reliance on exact SVD computations seems to be the main hurdle in reducing the runtime to depend on $O(\text{nnz}(A))$. This paper shows a version of FREQUENTDIRECTIONS whose runtime depends on $O(\text{nnz}(A))$. This requires a new understanding and a more careful analysis of FREQUENTDIRECTIONS. It also takes advantage of block power methods (also known as Subspace Iteration, Simultaneous Iteration, or Orthogonal Iteration) that run in time proportional to $\text{nnz}(A)$ but incur small approximation error [29].

1.1 Linear Algebra Notations

Throughout the paper we identify an $n \times d$ matrix A with a set of n rows $[a_1; a_2; \dots; a_n]$ where each a_i is a vector in \mathbb{R}^d . The notation a_i stands for the i th row of the matrix A . By $[A; a]$ we mean the row vector a appended to the matrix A as its last row. Similarly, $[A; B]$ stands for stacking two matrices A and B vertically. The matrices I_n and $0^{n \times d}$ denote the n -dimensional identity matrix and the full zero matrix of dimension $n \times d$ respectively. The notation $\mathcal{N}(0, 1)^{d \times \ell}$ denotes the distribution over $d \times \ell$ matrices whose entries are drawn independently from the normal distribution $\mathcal{N}(0, 1)$. For a vector x the notation $\|\cdot\|$ refers to the Euclidian norm $\|x\| = (\sum_i x_i^2)^{1/2}$. The Frobenius norm of a matrix A is defined as $\|A\|_F = \sqrt{\sum_{i=1}^n \|a_i\|^2}$, and the operator (or spectral) norm of it is $\|A\|_2 = \sup_{x \neq 0} \|Ax\|/\|x\|$.

The notation $\text{nnz}(A)$ refers to the number of non-zeros in A , and $\rho = \text{nnz}(A)/(nd)$ denotes relative density of A . The Singular Value Decomposition of a matrix $A \in \mathbb{R}^{m \times d}$ for $m \leq d$ is denoted by $[U, \Sigma, V] = \text{SVD}(A)$. It guarantees that $A = U\Sigma V^T$, $U^T U = I_m$, $V^T V = I_m$, $U \in \mathbb{R}^{m \times m}$, $V \in \mathbb{R}^{d \times m}$, and $\Sigma \in \mathbb{R}^{m \times m}$ is a non-negative diagonal matrix such that $\Sigma_{i,i} = \sigma_i$ and $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_m \geq 0$. It is convenient to denote by U_k , and V_k the matrices containing the first k columns of U and V and by $\Sigma_k \in \mathbb{R}^{k \times k}$ the top left $k \times k$ block of Σ . The matrix $A_k = U_k \Sigma_k V_k^T$ is the best rank k approximation of A in the sense that $A_k = \arg \min_{C: \text{rank}(C) \leq k} \|A - C\|_{2,F}$. In places where we use $\text{SVD}(A, k)$, we mean rank k SVD of A .

The notation $\pi_B(A)$ denotes the projection of the rows of A on the span of the rows of B . In other words, $\pi_B(A) = AB^\dagger B$ where $(\cdot)^\dagger$ indicates taking the Moore-Penrose pseudoinverse. Alternatively, setting $[U, \Sigma, V] = \text{SVD}(B)$, we have $\pi_B(A) = AVV^T$. We also denote $\pi_B^k(A) = AV_k V_k^T$, the right projection of A on the top k right singular vectors of B .

2. MATRIX SKETCHING PRIOR ART

This section reviews only matrix sketching techniques that run in input sparsity time and whose output sketch is independent of the number of rows in the matrix. We categorize all known results into three main approaches (1) column/row subset selection (2) random projection based techniques and (3) iterative sketching techniques.

Column selection techniques. These techniques, which are also studied under the Column Subset Selection Problem (CSSP) in literature [16, 10, 3, 8, 15, 2], form the sketch B by selecting a subset of ‘‘important’’ columns of the input matrix A . They maintain the sparsity of A and make the sketch B

to be more interpretable. These methods are not typically streaming, nor running in input sparsity time. The only method of this group which achieves both is [11] by Drineas *et al.* that uses reservoir sampling to become streaming. They select $O(k/\varepsilon^2)$ columns proportional to their squared norm and achieve the Frobenius norm error bound $\|A - \pi_{B_k}(A)\|_F^2 \leq \|A - A_k\|_F^2 + \varepsilon \|A\|_F^2$ with time complexity of $O((k^2/\varepsilon^4)(d + k/\varepsilon^2) + \text{nnz}(A))$. In addition, they show that the spectral norm error bound $\|A - \pi_{B_k}(A)\|_2^2 \leq \|A - A_k\|_2^2 + \varepsilon \|A\|_F^2$ holds if one selects $O(1/\varepsilon^2)$ columns. Rudelson *et al.* [34] improved the latter error bound to $\|A - \pi_{B_k}(A)\|_2^2 \leq \|A - A_k\|_2^2 + \varepsilon \|A\|_2^2$ by selecting $O(r/\varepsilon^4 \log(r/\varepsilon^4))$ columns, where $r = \|A\|_F^2/\|A\|_2^2$ is the numeric rank of A . Note that in the result by [11], one would need $O(r^2/\varepsilon^2)$ columns to obtain the same bound.

Another similar line of work is the CUR factorization [4, 10, 12, 14, 27] where methods select c columns and r rows of A to form matrices $C \in \mathbb{R}^{n \times c}$, $R \in \mathbb{R}^{r \times d}$ and $U \in \mathbb{R}^{c \times r}$, and constructs the sketch as $B = CUR$. The only instance of this group that runs in input sparsity time is [4] by Boutsidis and Woodruff, where they select $r = c = O(k/\varepsilon)$ rows and columns of A and construct matrices C, U and R with $\text{rank}(U) = k$ such that with constant probability $\|A - CUR\|_F^2 \leq (1 + \varepsilon)\|A - A_k\|_F^2$. Their algorithm runs in $O(\text{nnz}(A) \log n + (n + d) \text{poly}(\log n, k, 1/\varepsilon))$ time.

Random projection techniques. These techniques [31, 36, 35, 26] operate data-obliviously and maintain a $r \times d$ matrix $B = SA$ using a $r \times n$ random matrix S which has the Johnson-Lindenstrauss Transform (JLT) property [28]. Random projection methods work in the streaming model, are computationally efficient, and sufficiently accurate in practice [7]. The state-of-the-art method of this approach is by Clarkson and Woodruff [6] which was later improved slightly in [30]. It uses a hashing matrix S with only one non-zero entry in each column. Constructing this sketch takes only $O(\text{nnz}(A) + n \cdot \text{poly}(k/\varepsilon) + \text{poly}(dk/\varepsilon))$ time, and guarantees that for any unit vector x that $(1 - \varepsilon)\|Ax\| \leq \|Bx\| \leq (1 + \varepsilon)\|Ax\|$. For these sparsity-efficient sketches using $r = O(d^2/\varepsilon^2)$ also guarantees that $\|A - \pi_B(A)\|_F \leq (1 + \varepsilon)\|A - A_k\|_F$.

Iterative sketching techniques. These operate in streaming model, where random access to the matrix is not available. They maintain the sketch B as a linear combination of rows of A , and update it as new rows are received in the stream. Examples of these methods include different version of iterative SVD [19, 21, 23, 5, 33]. These, however, do not have theoretical guarantees [7]. The FREQUENTDIRECTIONS algorithm [24] is a unique in this group in that it offers strong error guarantees. It is a deterministic sketching technique that processes rows of an $n \times d$ matrix A in a stream and maintains a $\ell \times d$ sketch B (for $\ell < \min(n, d)$) such that the following two error bounds hold for any $0 \leq k < \ell$

$$\|A^T A - B^T B\|_2 \leq \frac{1}{\ell - k} \|A - A_k\|_F^2$$

and

$$\|A - \pi_B(A)\|_F^2 \leq \frac{\ell}{\ell - k} \|A - A_k\|_F^2.$$

Setting $\ell = k + 1/\varepsilon$ and $\ell = k + k/\varepsilon$, respectively, achieves bounds $\|A^T A - B^T B\|_2 \leq \varepsilon \|A - A_k\|_F^2$ and $\|A - \pi_B(A)\|_F^2 \leq (1 + \varepsilon)\|A - A_k\|_F^2$. Although FREQUENTDIRECTIONS does not

run in input sparsity time, we will explain it in detail in the next section, as it is an important building block for the algorithm we introduce.

2.1 Main Results

We present a randomized version of FREQUENTDIRECTIONS, called as SPARSEFREQUENTDIRECTIONS that receives an $n \times d$ sparse matrix A as a stream of its rows. It computes a $\ell \times d$ sketch B in $O(d\ell)$ space.

It guarantees that with probability at least $1 - \delta$ (for $\delta \in (0, 1)$ being the failure probability), for $\alpha = 6/41$ and any $0 \leq k < \alpha\ell$,

$$\|A^T A - B^T B\|_2 \leq \frac{1}{\alpha\ell - k} \|A - A_k\|_F^2$$

and

$$\|A - \pi_{B_k}(A)\|_F^2 \leq \frac{\ell}{\ell - k/\alpha} \|A - A_k\|_F^2.$$

Note that setting $\ell = \lceil 1/(\varepsilon\alpha) + k/\alpha \rceil$ yields

$$\|A^T A - B^T B\|_2 \leq \varepsilon \|A - A_k\|_F^2$$

and setting $\ell = \lceil k/(\varepsilon\alpha) + k/\alpha \rceil$ yields

$$\|A - \pi_{B_k}(A)\|_F^2 \leq (1 + \varepsilon) \|A - A_k\|_F^2.$$

The expected running time of the algorithm is

$$O(\text{nnz}(A)\ell \log(d) + \text{nnz}(A) \log(n/\delta) + n\ell^2 + n\ell \log(n/\delta))$$

In the likely case where $\text{nnz}(A) = \Omega(n\ell)$ and $n/\delta < d^{O(\ell)}$, the runtime is dominated by $O(\text{nnz}(A)\ell \log(d))$. We also experimentally validate this theory, demonstrating these runtime improvements on sparse data without sacrificing accuracy.

3. PRELIMINARIES

In this section we review some important properties about FREQUENTDIRECTIONS and SIMULTANEOUSITERATION which will be necessary for understanding and proving bounds on SPARSEFREQUENTDIRECTIONS.

3.1 Frequent Directions

The FREQUENTDIRECTIONS algorithm was introduced by Liberty [24] and received an improved analysis by Ghashami et al.[17]. The algorithm operates by collecting several rows of the input matrix and letting the sketch grow. Once the sketch doubles in size, a lossy DENSESHRINK operation reduces its size by a half. This process repeats throughout the stream. The running time of FREQUENTDIRECTIONS and its error analysis are strongly coupled with the properties of the SVD used to perform the DENSESHRINK step.

An analysis of [7] slightly generalized the one in [17]. Let B be the sketch resulting in applying FREQUENTDIRECTIONS with a potentially different shrink operation to A . Then, the FREQUENTDIRECTIONS asymptotic guarantees hold as long as the shrink operation exhibits three properties, for any positive Δ and a constant $\alpha \in (0, 1)$.

1. Property 1:
For any vector $x \in \mathbb{R}^d$, $\|Ax\|^2 - \|Bx\|^2 \geq 0$.
2. Property 2:
For any unit vector $x \in \mathbb{R}^d$, $\|Ax\|^2 - \|Bx\|^2 \leq \Delta$.
3. Property 3: $\|A\|_F^2 - \|B\|_F^2 \geq \alpha\Delta\ell$.

For completeness, the exact guarantee is stated in Lemma 3.1.

LEMMA 3.1 (LEMMA 3.1 IN [7]). *Given an input $n \times d$ matrix A and an integer parameter ℓ , any sketch $\ell \times d$ matrix B which satisfies the three properties above (for some any $\alpha \in (0, 1]$ and $\Delta > 0$), guarantees the following error bounds*

$$0 \leq \|A^T A - B^T B\|_2 \leq \frac{1}{\alpha\ell - k} \|A - A_k\|_F^2,$$

and

$$\|A - \pi_B^k(A)\|_F^2 \leq \frac{\ell}{\ell - k/\alpha} \|A - A_k\|_F^2,$$

where $\pi_B^k(\cdot)$ represents the projection operator onto B_k , the top k singular vectors of B .

Another important property of FREQUENTDIRECTIONS is that its sketches are mergeable [17]. To clarify, consider partitioning a matrix A into t blocks A_1, A_2, \dots, A_t so that $A = [A_1; A_2; \dots; A_t]$. Let $B_i = \text{FD}(A_i, \ell)$ denotes the FD sketch of the matrix block A_i , and $B' = \text{FD}([B_1; B_2; \dots; B_t], \ell)$ denotes the FD sketch of all B_i s combined together. It is shown that B' has at most as much covariance and projection error as $B = \text{FD}(A, \ell)$, i.e. the sketch of the whole matrix A . It follows that this divide-sketch-and-merging can also be applied recursively on each matrix block without increasing the error.

The runtime of FREQUENTDIRECTIONS is determined by the number of shrinking steps. Each of those computes an SVD of B which takes $O(d\ell^2)$ time. Since the SVD is called only every $O(\ell)$ rows this yields a total runtime $O(d\ell^2 \cdot n/\ell) = O(nd\ell)$. This effectively means that on average we are spending $O(d\ell)$ operations per row, even if the row is sparse.

In the present paper, we introduce a new method called as SPARSEFREQUENTDIRECTIONS that uses randomized SVD methods instead of the exact SVD to approximate the singular vectors and values of intermediate matrices B . We show how this new method tolerates the extra approximation error and runs in time proportional to $\text{nnz}(A)$. Moreover, since it received sparse matrix rows, it can observe more the ℓ rows until the size of the sketch doubles. As a remark, Ghashami and Phillips [18] showed that maintaining any rescaled set of ℓ rows of A over a stream is not a feasible approach to obtain sparsity in FREQUENTDIRECTIONS. It was left as an open problem to produce some version of FREQUENTDIRECTIONS that took advantage of the sparsity of A .

3.2 Simultaneous Iteration

Efficiently computing the singular vectors of matrices is one of the most well studied problems in scientific computing. Recent results give very strong approximation guarantees for block power method techniques [32][38][26][20]. Several variants of this algorithm were studied under different names in the literature e.g. Simultaneous Iteration, Subspace Iteration, or Orthogonal Iteration [19]. In this paper, we refer to this group of algorithms collectively as SIMULTANEOUSITERATION. A generic version of SIMULTANEOUSITERATION for rectangular matrices is described in Algorithm 1.

Algorithm 1 SIMULTANEOUSITERATION

Input: $A \in \mathbb{R}^{n \times d}$, rank $k \leq \min(n, d)$, and error $\varepsilon \in (0, 1)$
 $q = \Theta(\log(n/\varepsilon)/\varepsilon)$
 $G \sim \mathcal{N}(0, 1)^{d \times k}$
 $Z = \text{GramSchmidt}(A(A^T A)^q G)$
return Z # $Z \in \mathbb{R}^{n \times k}$

While this algorithm was already analyzed by [19], the proofs of [32, 20, 29, 37] manage to prove stable results that hold for any matrix independent of spectral gap issues. Unfortunately, an in depth discussion of these algorithms and their proof techniques is beyond the scope of this paper.

For the proof of correctness of SPARSEFREQUENTDIRECTIONS, the main lemma proven by [29] suffices. SIMULTANEOUSITERATION (Algorithm 1) guarantees the three following error bounds with high probability:

1. Frobenius norm error bound:

$$\|A - ZZ^T A\|_F \leq (1 + \varepsilon) \|A - A_k\|_F$$

2. Spectral norm error bound:

$$\|A - ZZ^T A\|_2 \leq (1 + \varepsilon) \|A - A_k\|_2$$

3. Per vector error bound:

$$|u_i^T A A^T u_i - z_i^T A A^T z_i| \leq \varepsilon \sigma_{k+1}^2$$

for all i . Here u_i denotes the i th left singular vector of A , and σ_{k+1} is the $(k + 1)$ th singular value of A , and z_i is the i th column of the matrix Z returned by SIMULTANEOUSITERATION.

In addition, for a constant ε , SIMULTANEOUSITERATION runs in $\tilde{O}(\text{nnz}(A))$ time.

In this paper, we show that SPARSEFREQUENTDIRECTIONS can replace the computation of an exact SVD by using the results of [29] with ε being a constant. This alteration does give up the optimal asymptotic accuracy (matching that of FREQUENTDIRECTIONS).

4. SPARSE FREQUENT DIRECTIONS

The SPARSEFREQUENTDIRECTIONS (SFD) algorithm is described in Algorithm 2, and is an extension of FREQUENTDIRECTIONS to sparse matrices. It receives the rows of an input matrix A in a streaming fashion and maintains a sketch B of ℓ rows. Initially B is empty. On receiving rows of A , SFD stores non-zeros in a buffer matrix A' . The buffer is deemed full when it contains ℓd non-zeros or d rows. SFD then calls BOOSTEDSPARSESHRINK to produce its sketch matrix B' of size $\ell \times d$. Then, it updates its ongoing sketch B of the entire stream by merging it with the (dense) sketch B' using DENSESHRINK.

Algorithm 2 SPARSEFREQUENTDIRECTIONS

Input: $A \in \mathbb{R}^{n \times d}$, an integer $\ell \leq d$, failure probability δ
 $B = 0^{\ell \times d}$, $A' = 0^{0 \times d}$
for $a \in A$ **do**
 $A' = [A'; a]$
 if $\text{nnz}(A') \geq \ell d$ **or** $\text{rows}(A') = d$ **then**
 $B' = \text{BOOSTEDSPARSESHRINK}(A', \ell, \delta)$
 $B = \text{DENSESHRINK}([B; B'], \ell)$
 $A' = 0^{0 \times d}$
return B

BOOSTEDSPARSESHRINK amplifies the success probability of another algorithm SPARSESHRINK in Algorithm 3. SPARSESHRINK runs SIMULTANEOUSITERATION instead of a full SVD to take advantage of the sparsity of its input A' . However, as we will discuss, by itself SPARSESHRINK has too high of a probability of failure. Thus we use BOOSTEDSPARSESHRINK which keeps running SPARSESHRINK and probabilistically verifying the correctness of its result using VERIFYSPECTRAL, until it decides that the result is correct with high enough probability. Each of DENSESHRINK, SPARSESHRINK, and BOOSTEDSPARSESHRINK produce sketch matrices of size $\ell \times d$.

Algorithm 3 SPARSESHRINK

Input: $A' \in \mathbb{R}^{m \times d}$, an integer $\ell \leq m$
 $Z = \text{SIMULTANEOUSITERATION}(A', \ell, 1/4)$
 $P = Z^T A'$, $[H, \Lambda, V] = \text{SVD}(P, \ell)$
 $\tilde{\Lambda} = \sqrt{\Lambda^2 - \lambda_\ell^2 I_\ell}$
 $B' = \tilde{\Lambda} V^T$
return B'

Algorithm 4 BOOSTEDSPARSESHRINK

Input: $A' \in \mathbb{R}^{m \times d}$, integer $\ell \leq m$, failure probability δ
while True **do**
 $B' = \text{SPARSESHRINK}(A', \ell)$
 $\Delta = (\|A'\|_F^2 - \|B'\|_F^2) / \alpha \ell$ for $\alpha = 6/41$
 if VERIFYSPECTRAL($(A'^T A' - B'^T B') / (\Delta/2)$, δ) **then**
 return B'

Algorithm 5 DENSESHRINK

Input: $A \in \mathbb{R}^{m \times d}$, an integer $\ell \leq m$
 $[H, \Lambda, V] = \text{SVD}(A, \ell)$
 $\tilde{\Lambda} = \sqrt{\Lambda^2 - \lambda_\ell^2 I_\ell}$
 $B = \tilde{\Lambda} V^T$
Return B

Our main result is stated in the next theorem. It follows from combining the proofs contained in the subsections below.

THEOREM 4.1 (MAIN RESULT). *Given a sparse matrix $A \in \mathbb{R}^{n \times d}$ and an integer $\ell \leq d$, SPARSEFREQUENTDIRECTIONS computes a small sketch $B \in \mathbb{R}^{\ell \times d}$ such that with probability at least $1 - \delta$ for $\alpha = 6/41$ and any $0 \leq k < \alpha \ell$,*

$$\|A^T A - B^T B\|_2 \leq \frac{1}{\alpha \ell - k} \|A - A_k\|_F^2$$

and

$$\|A - \pi_{B_k}(A)\|_F^2 \leq \frac{\ell}{\ell - k/\alpha} \|A - A_k\|_F^2.$$

The total memory footprint of the algorithm is $O(d\ell)$ and its expected running time is

$$O(\text{nnz}(A)\ell \log(d) + \text{nnz}(A) \log(n/\delta) + n\ell^2 + n\ell \log(n/\delta)).$$

4.1 Success Probability

SPARSESHRINK, described in Algorithm 3, calls SIMULTANEOUSITERATION to approximate the top rank ℓ subspace

of A' . As `SIMULTANEOUSITERATION` is randomized, it fails to converge to a good subspace when the initial choice of the random matrix G does not sufficiently align with the top ℓ singular vectors of A' (see Algorithm 1). This occurs with probability at most $\rho_\ell = O(1/\sqrt{\ell})$. In Section 4.3.1, we prove that with probability of at least $1 - \rho_\ell$ that `SPARSESHRINK` satisfies the three properties required for Lemma 3.1 using $\alpha = 6/41$ and $\Delta = 41/8 s_\ell^2$, but replacing Property 2 with a stronger version

- Property 2 (strengthened):

$$\|A'^T A' - B'^T B'\|_2 \leq (\Delta/2) = 41/16 s_\ell^2$$

where s_ℓ denotes the ℓ th singular value of A' .

However, for the proof of `SPARSEFREQUENTDIRECTIONS` we require that *all* `SPARSESHRINK` runs to be successful. The failure probability of `SPARSESHRINK`, which is upper bounded by $O(1/\sqrt{\ell})$, is high enough that a simple union bound would not give a meaningful bound on the failure probability of `SPARSEFREQUENTDIRECTIONS`. We therefore reduce the failure probability of each `BOOSTEDSPARSESHRINK`, by wrapping each call of `SPARSESHRINK` in the verifier `VERIFYSPECTRAL`. If `VERIFYSPECTRAL` does not verify the correctness, then it reruns `SPARSESHRINK` and tries again until it can verify it. But to perform this verification efficiently, we need to loosen the definition of correctness. In particular, we say `SPARSESHRINK` is successful if the sketch B' computed from its output satisfies $\|A'^T A' - B'^T B'\|_2 \leq \Delta$ (the original Property 2 specification in Section 3.1), where $\Delta = (\|A'\|_F^2 - \|B'\|_F^2)/\alpha\ell$. Combining the two inequalities through Δ , a successful run implies that $\|A'^T A' - B'^T B'\|_2 \leq (\|A'\|_F^2 - \|B'\|_F^2)/\alpha\ell$. `VERIFYSPECTRAL` verifies the success of the algorithm by approximating the spectral norm of $(A'^T A' - B'^T B')/(\Delta/2)$; it does so by running the power method for $c \cdot \log(d/\delta_i)$ steps for some constant c .

Algorithm 6 `VERIFYSPECTRAL`

Initialization persistent $i = 0$ (i retains its state between invocations of this method)

Input: Matrix $C \in \mathbb{R}^{d \times d}$, failure probability δ

$i = i + 1$ and $\delta_i = \delta/2i^2$

Pick x uniformly at random from the unit sphere in \mathbb{R}^d .

if $\|C^{c \cdot \log(d/\delta_i)} x\| \leq 1$ **return** TRUE

else return FALSE

LEMMA 4.1. *The `VERIFYSPECTRAL` algorithm returns TRUE if $\|C\|_2 \leq 1$. If $\|C\|_2 \geq 2$ it returns FALSE with probability at least $1 - \delta_i$.*

PROOF. If $\|C\| \leq 1$ than $\|C^{c \cdot \log(d/\delta_i)} x\| \leq \|C\|^{c \cdot \log(d/\delta_i)} \|x\| \leq 1$. If $\|C\| \geq 2$, consider execution i of the method. Let v_1 denote the top singular vector of C . Then $\|C^{c \cdot \log(d/\delta_i)} x\| \geq |\langle v_1, x \rangle| 2^{c \cdot \log(d/\delta_i)} \geq 1$, for some constant c as long as $|\langle v_1, x \rangle| = \Omega(\text{poly}(\delta_i/d))$. Let $\Phi(t')$ denote the density function of the random variable $t' = \langle v_1, x \rangle$. Then $\Pr[|\langle v_1, x \rangle| \leq t] = \int_{-t}^t \Phi(t') dt' \leq 2t\Phi(0) = O(t\sqrt{d})$. Setting the failure probability to be at most δ_i , we conclude that $|\langle v_1, x \rangle| = \Omega(\delta_i/\sqrt{d})$ with probability at least $1 - \delta_i$. \square

Therefore, `VERIFYSPECTRAL` fails with probability at most δ_i during execution i . If any of `VERIFYSPECTRAL` runs fail, `BOOSTEDSPARSESHRINK` and hence `SPARSEFREQUENTDIRECTIONS` potentially fail. Taking the union bound over

all invocations of `VERIFYSPECTRAL` we obtain that `SPARSEFREQUENTDIRECTIONS` fails with probability at most $\sum_{i=1}^{\infty} \delta/2i^2 \leq \delta$, hence it succeeds with probability at least $1 - \delta$.

4.2 Space Usage and Runtime Analysis

Throughout this manuscript we assume the constant-word-size model. Integers and floating point numbers are represented by a constant number of bits. Random access into memory is assumed to require $O(1)$ time. In this model, multiplying a sparse matrix A' by a dense vector requires $O(\text{nnz}(A'))$ operations and storing A' requires $O(\text{nnz}(A'))$ bits of memory.

FACT 4.1. *The total memory footprint of `SPARSEFREQUENTDIRECTIONS` is $O(d\ell)$.*

PROOF. It is easy to verify that, except for the buffer matrix A' , the algorithm only manipulates $\ell \times d$ matrices; in particular, observe that the $(\text{rows}(A') = d)$ condition in `SPARSEFREQUENTDIRECTIONS` ensures that $m = d$ in `SPARSESHRINK`, and in `DENSESHRINK` also $m = 2\ell$. Each of these $\ell \times d$ matrices clearly require at most $O(d\ell)$ bits of memory. The buffer matrix A' contains at most $O(d\ell)$ non-zeros and therefore does not increase the space complexity of the algorithm. \square

We turn to bounding the expected runtime of `SPARSEFREQUENTDIRECTIONS` which is dominated by the cumulative running times of `DENSESHRINK` and `BOOSTEDSPARSESHRINK`. Denote by T the number of times they are executed. It is easy to verify $T \leq \text{nnz}(A)/d\ell + n/d$. Since `DENSESHRINK` runs in $O(d\ell^2)$ time deterministically, the total time spent by `DENSESHRINK` through T iterations is $O(Td\ell^2) = O(\text{nnz}(A)\ell + n\ell^2)$.

The running time of `BOOSTEDSPARSESHRINK` is dominated by those of `SPARSESHRINK` and `VERIFYSPECTRAL`, and its expected number of iterations. Note that, in expectation, they are each executed on any buffer matrix A'_i a small constant number of times because `VERIFYSPECTRAL` succeeds with probability (much) greater than $1/2$. For asymptotic analysis it is identical to assuming they are each executed once.

Note that the running time of `SPARSESHRINK` on A'_i is $O(\text{nnz}(A'_i)\ell \log(d))$. Since $\sum_i \text{nnz}(A'_i) = \text{nnz}(A)$ we obtain a total running time of $O(\text{nnz}(A)\ell \log(d))$. The i th execution of `VERIFYSPECTRAL` requires $O(d\ell \log(d/\delta_i))$ operations. This, because it multiplies $A'^T A' - B'^T B'$ by a single vector $O(\log(d/\delta_i))$ times, and both $\text{nnz}(A') \leq O(d\ell)$ and $\text{nnz}(B') \leq d\ell$. In expectation `VERIFYSPECTRAL` is executed $O(T)$ times, therefore total running time of it is

$$\begin{aligned} O(d\ell \sum_{i=1}^{O(T)} \log(d/\delta_i)) &= O(d\ell \sum_{i=1}^{O(T)} \log(di^2/\delta)) \\ &= O(Td\ell \log(Td/\delta)) = O((\text{nnz} + n\ell) \log(n/\delta)). \end{aligned}$$

Combining the above contributions to the total running time of the algorithm we obtain Fact 4.2.

FACT 4.2. *Algorithm `SPARSEFREQUENTDIRECTIONS` runs in expected time of*

$$O(\text{nnz}(A)\ell \log(d) + \text{nnz}(A) \log(n/\delta) + n\ell^2 + n\ell \log(n/\delta)).$$

4.3 Error Analysis

We turn to proving the error bounds of Theorem 4.1. Our proof is divided into three parts. We first show that SPARSESHRINK obtains the three properties needed for Lemma 3.1 with probability at least $1 - \rho_\ell$, and with the constraint on Property 2 strengthened by a factor $1/2$. Then we show how loosening Property 2 back to its original bound enables BOOSTEDSPARSESHRINK to succeed with probability $1 - \delta_i$ for some $\delta_i \ll \rho_\ell$. Finally we show that due to the mergeability of FREQUENTDIRECTIONS [25], discussed in Section 3.1, the SPARSEFREQUENTDIRECTIONS algorithm obtains the same error guarantees as BOOSTEDSPARSESHRINK with probability $1 - \delta$ for a small δ of our choice.

In what follows, we mainly consider a single run of SPARSESHRINK or BOOSTEDSPARSESHRINK and let s_ℓ and u_ℓ denote the ℓ th singular value and ℓ th left singular vector of A' , respectively.

4.3.1 Error Analysis: SPARSESHRINK

Here we show that with probability at least $1 - \rho_\ell$ that B' computed from SPARSESHRINK(A', ℓ) satisfies the three properties discussed in Section 3.1 required for Lemma 3.1.

- Property 1: For any unit vector $x \in \mathbb{R}^d$, $\|A'x\|^2 - \|B'x\|^2 \geq 0$,
- Property 2 (strengthened): For any unit vector $x \in \mathbb{R}^d$, $\|A'x\|^2 - \|B'x\|^2 \leq \Delta/2 = (41/16)s_\ell^2$,
- Property 3: $\|A'\|_F^2 - \|B'\|_F^2 \geq \ell\alpha\Delta = \ell(3/4)s_\ell^2$.

LEMMA 4.2. *Property 1 holds deterministically for SPARSESHRINK: $\|A'x\|^2 - \|B'x\|^2 \geq 0$ for all x .*

PROOF. Let $P = Z^T A'$ be as defined in SPARSESHRINK. Consider an arbitrary unit vector $x \in \mathbb{R}^d$, and let $y = A'x$.

$$\begin{aligned} \|A'x\|^2 - \|Px\|^2 &= \|A'x\|^2 - \|Z^T A'x\|^2 = \|y\|^2 - \|Z^T y\|^2 \\ &= \|(I - ZZ^T)y\|^2 \geq 0 \end{aligned}$$

and

$$\|Px\|^2 - \|B'x\|^2 = \lambda_\ell^2 \sum_{i=1}^{\ell} \langle x, v_i \rangle^2 \geq 0,$$

therefore $\|A'x\|^2 - \|B'x\|^2 = (\|A'x\|^2 - \|Px\|^2) + (\|Px\|^2 - \|B'x\|^2) \geq 0$. \square

LEMMA 4.3. *With probability at least $1 - \rho_\ell$, Property 2 holds for SPARSESHRINK: for any unit vector $x \in \mathbb{R}^d$, $\|A'x\|^2 - \|B'x\|^2 \leq 41/16 s_\ell^2$.*

PROOF. Consider an arbitrary unit vector $x \in \mathbb{R}^d$, and note that

$$\|A'x\|^2 - \|B'x\|^2 = (\|A'x\|^2 - \|Px\|^2) + (\|Px\|^2 - \|B'x\|^2).$$

We bound each term individually. The first term is bounded as

$$\|A'x\|^2 - \|Px\|^2 = x^T (A'^T A' - P^T P)x \quad (1)$$

$$\leq \|A'^T A' - P^T P\|_2 \quad (2)$$

$$= \|A'^T A' - A'^T Z Z^T A'\|_2 \quad (3)$$

$$= \|A'^T (I - Z Z^T) A'\|_2 \quad (4)$$

$$= \|A'^T (I - Z Z^T)^T (I - Z Z^T) A'\|_2 \quad (5)$$

$$= \|(I - Z Z^T) A'\|_2^2 \quad (6)$$

$$\leq 25/16 s_{\ell+1}^2 \leq 25/16 s_\ell^2. \quad (7)$$

Where transition 5 is true because $(I - Z Z^T)$ is a projection. Transition 7 also holds by the spectral norm error bound of [29] for $\varepsilon = 1/4$. To bound the second term, note that $\|Px\| = \|Z^T A'x\| = \|\Lambda V^T x\|$, since $[H, \Lambda, V] = \text{SVD}(P, \ell)$ as defined in SPARSESHRINK.

$$\begin{aligned} \|Px\|^2 - \|B'x\|^2 &= \sum_{i=1}^{\ell} \lambda_i^2 \langle x, v_i \rangle^2 - \sum_{i=1}^{\ell} \tilde{\lambda}_i^2 \langle x, v_i \rangle^2 = \\ &= \sum_{i=1}^{\ell} (\lambda_i^2 - \tilde{\lambda}_i^2) \langle x, v_i \rangle^2 = \sum_{i=1}^{\ell} \lambda_\ell^2 \langle x, v_i \rangle^2 \leq \lambda_\ell^2 \leq s_\ell^2, \end{aligned}$$

where last inequality follows by the Courant-Fischer min-max principle, i.e. as λ_ℓ is the ℓ th singular value of the projection of A' onto Z , then $\lambda_\ell \leq s_\ell$. Summing the two terms yields $\|A'x\|^2 - \|B'x\|^2 \leq 41/16 s_\ell^2$. \square

The original bound $\|A'^T A' - B'^T B'\|_2 \leq \Delta = 41/8 s_\ell^2$ discussed in Section 4.1 is also immediately satisfied.

LEMMA 4.4. *With probability at least $1 - \rho_\ell$, Property 3 holds for SPARSESHRINK: $\|A'\|_F^2 - \|B'\|_F^2 \geq \ell(3/4)s_\ell^2$.*

PROOF.

$$\|A'\|_F^2 - \|P\|_F^2 = \|A'\|_F^2 - \|Z^T A'\|_F^2 = \|A' - Z Z^T A'\|_F^2 \geq 0$$

In addition,

$$\|P\|_F^2 - \|B'\|_F^2 = \ell \lambda_\ell^2 \geq \ell(3/4)s_\ell^2.$$

The last inequality holds by the per vector error bound of [29] for $i = \ell$ and $\varepsilon = 1/4$, i.e. $|u_\ell^T A' A'^T u_\ell - z_\ell^T A' A'^T z_\ell| = |s_\ell^2 - \lambda_\ell^2| \leq 1/4 s_{\ell+1}^2 \leq 1/4 s_\ell^2$, which means $\lambda_\ell^2 \geq 3/4 s_\ell^2$. Therefore

$$\|A'\|_F^2 - \|B'\|_F^2 = (\|A'\|_F^2 - \|P\|_F^2) + (\|P\|_F^2 - \|B'\|_F^2) \geq \ell(3/4)s_\ell^2.$$

\square

4.3.2 Error Analysis: BOOSTEDSPARSESHRINK and SPARSEFREQUENTDIRECTIONS

We now consider the BOOSTEDSPARSESHRINK algorithm, and the looser version of Property 2 (the original version) as

- Property 2: For any unit vector $x \in \mathbb{R}^d$, $\|A'x\|^2 - \|B'x\|^2 \leq \Delta = (41/8)s_\ell^2$.

By invoking VERIFYSPECTRAL($(A'^T A' - B'^T B')/(\Delta/2), \delta$), then VERIFYSPECTRAL always returns TRUE if $\|A'^T A' - B'^T B'\|_2 \leq \Delta/2$ (as is true of the input with probability at least $1 - \rho_\ell$ by Lemma 4.3), and VERIFYSPECTRAL catches a failure event where $\|A'^T A' - B'^T B'\|_2 \geq \Delta$ with probability at least $1 - \delta_i$ by Lemma 4.1. As discussed in Section 4.1 all invocations of VERIFYSPECTRAL succeed with probability at most $1 - \delta$, hence all runs of BOOSTEDSPARSESHRINK succeed and satisfy Property 2 (as well as Properties 1 and 3) with $\alpha = 6/41$ and $\Delta = 41/8 s_\ell^2$, and with probability at least $1 - \delta$. Finally, we can invoke the mergeability property of FREQUENTDIRECTIONS [25] and Lemma 3.1 to obtain the error bounds in our main result, Theorem 4.1.

5. EXPERIMENTS

In this section we empirically validate that SPARSEFREQUENTDIRECTIONS matches (and often improves upon) the accuracy of FREQUENTDIRECTIONS, while running significantly faster on sparse real and synthetic datasets.

We do not implement SPARSEFREQUENTDIRECTIONS exactly as described above. Instead we directly call SPARSESHRINK in Algorithm 2 in place of BOOSTEDSPARSESHRINK. The randomized error analysis of SIMULTANEOUSITERATION indicates that we may occasionally miss a subspace within a call of SIMULTANEOUSITERATION and hence SPARSESHRINK; but in practice this is not a catastrophic event, and as we will observe, does not prevent SPARSEFREQUENTDIRECTIONS from obtaining small empirical error.

The empirical comparison of FREQUENTDIRECTIONS to other matrix sketching techniques is now well-trodden [17, 7]. FREQUENTDIRECTIONS (and, as we observe, by association SPARSEFREQUENTDIRECTIONS) has much smaller error than other sketching techniques which operate in a stream. However, FREQUENTDIRECTIONS is somewhat slower by a factor of the sketch size ℓ up to some leading coefficients. We do not repeat these comparison experiments here.

Setup. We ran all the algorithms under a common implementation framework to test their relative performance as accurately as possible. We ran the experiments on an Intel(R) Core(TM) 2.60 GHz CPU with 64GB of RAM running Ubuntu 14.04.3. All algorithms were coded in C, and compiled using gcc 4.8.4. All linear algebra operation on dense matrices (such as SVD) invoked those implemented in LAPACK.

Datasets. We compare the performance of the two algorithms on both synthetic and real datasets. Each dataset is an $n \times d$ matrix A containing n datapoints in d dimensions.

The real dataset is part of the 20 Newsgroups dataset [22], that is a collection of approximately 20,000 documents, partitioned across 20 different newsgroups. However we use the ‘by date’ version of the data, where features (columns) are tokens and rows correspond to documents. This data matrix is a zero-one matrix with 11,314 rows and 117,759 columns. In our experiment, we use the transpose of the data and picked the first $d = 3000$ columns, hence the subset matrix has $n = 117,759$ rows and $d = 3000$ columns; roughly 0.15% of the subset matrix is non-zeros.

The synthetic data generates n rows i.i.d. Each row receives exactly $z \ll d$ non-zeros (with default $z = 100$ and $d = 1000$), with the remaining entries as 0. The non-zeros are chosen as either 1 or -1 at random. Each non-zero location is chosen without duplicates among the columns. The first $1.5z$ columns (e.g., 150), the ‘head’, have a higher probability of receiving a non-zero than the last $d - 1.5z$ columns, the ‘tail’. The process to place a non-zero first chooses the head with probability 0.9 or the tail with probability 0.1. For whichever set of columns it chooses (head or tail), it places the non-zero uniformly at random among those columns.

Measurements. Each algorithm outputs a sketch matrix B of ℓ rows. For each of our experiments, we measure the efficiency of algorithms against one parameter and keep others fixed at a default value. Table 1 lists all parameters along with their default value and the range they vary in for synthetic dataset. We measure the accuracy of the algorithms with respect to:

- Projection Error:

$$\text{proj-err} = \|A - \pi_{B_k}(A)\|_F^2 / \|A - A_k\|_F^2,$$

- Covariance Error:
 $\text{cov-err} = \|A^T A - B^T B\|_2 / \|A\|_F^2,$
- Runtime in seconds.

In all experiments, we have set $k = 10$. Note that proj-err is always larger than 1, and for FREQUENTDIRECTIONS and SPARSEFREQUENTDIRECTIONS the cov-err is always smaller than $1/(\frac{6}{41}\ell - k)$ due to our error guarantees.

5.1 Observations

By considering Table 2 on synthetic data and Figure 1 on the real data, we can vary and learn many aspects of the runtime and accuracy of SPARSEFREQUENTDIRECTIONS and FREQUENTDIRECTIONS.

Runtime. Consider the last row of Table 2, the ‘Run Time’ row, and the last column of Figure 1. SPARSEFREQUENTDIRECTIONS is clearly faster than FREQUENTDIRECTIONS for all datasets, except when the synthetic data becomes dense in the last column of the ‘Run Time’ row, where $d = 1000$ and $\text{nnz per row} = 500$ in the right-most data point. For the default values the improvement is between about a factor of $1.5x$ and $2x$, but when the matrix is very sparse the improvement is $10x$ or more. Very sparse synthetic examples are seen in the left data points of the last column, and in the right data points of the second column, of the ‘Run Time’ row.

In particular, these two plots (the second and fourth columns of the ‘Run Time’ row) really demonstrate the dependence of SPARSEFREQUENTDIRECTIONS on $\text{nnz}(A)$ and of FREQUENTDIRECTIONS on $n \cdot d$. In the last column, we fix the matrix size n and d , but increase the number of non-zeros $\text{nnz}(A)$; the runtime of FREQUENTDIRECTIONS is basically constant, while for SPARSEFREQUENTDIRECTIONS it grows linearly. In the second column, we fix n and $\text{nnz}(A)$, but increase the number of columns d ; the runtime of FREQUENTDIRECTIONS grows linearly while the runtime for SPARSEFREQUENTDIRECTIONS is basically constant.

These algorithms are designed for datasets with extremely large values of n ; yet we only run on datasets with n up to 60,000 in Table 2, and 117,759 in Figure 1. However, both FREQUENTDIRECTIONS and SPARSEFREQUENTDIRECTIONS have runtime that grows linearly with respect to the number of rows (assuming the sparsity is at an expected fixed rate per row for SPARSEFREQUENTDIRECTIONS). This can also be seen empirically in the first column of the ‘Run Time’ row where, after a small start-up cost, both FREQUENTDIRECTIONS and SPARSEFREQUENTDIRECTIONS grow linearly as a function of the number of data points n . Hence, it is valid to directly extrapolate these results for datasets of increased n .

Accuracy. We will next discuss the accuracy, as measured in Projection Error in the top row of Table 2 and left plot of Figure 1, and in Covariance Error in the middle row of Table 2 and middle plot of Figure 1.

We observe that both FREQUENTDIRECTIONS and SPARSEFREQUENTDIRECTIONS obtain very small error (much smaller than upper bounded by the theory), as has been observed elsewhere [17, 7]. Moreover, the error for SPARSEFREQUENTDIRECTIONS always nearly matches, or improves over FRE-

Table 1: Parameter values

	default	range
datapoints (n)	10000	$[10^4 - 6 \times 10^4]$
dimension (d)	1000	$[10^3 - 6 \times 10^3]$
sketch size (ℓ)	50	$[5 - 100]$
nnz per row	100	$[5 - 500]$

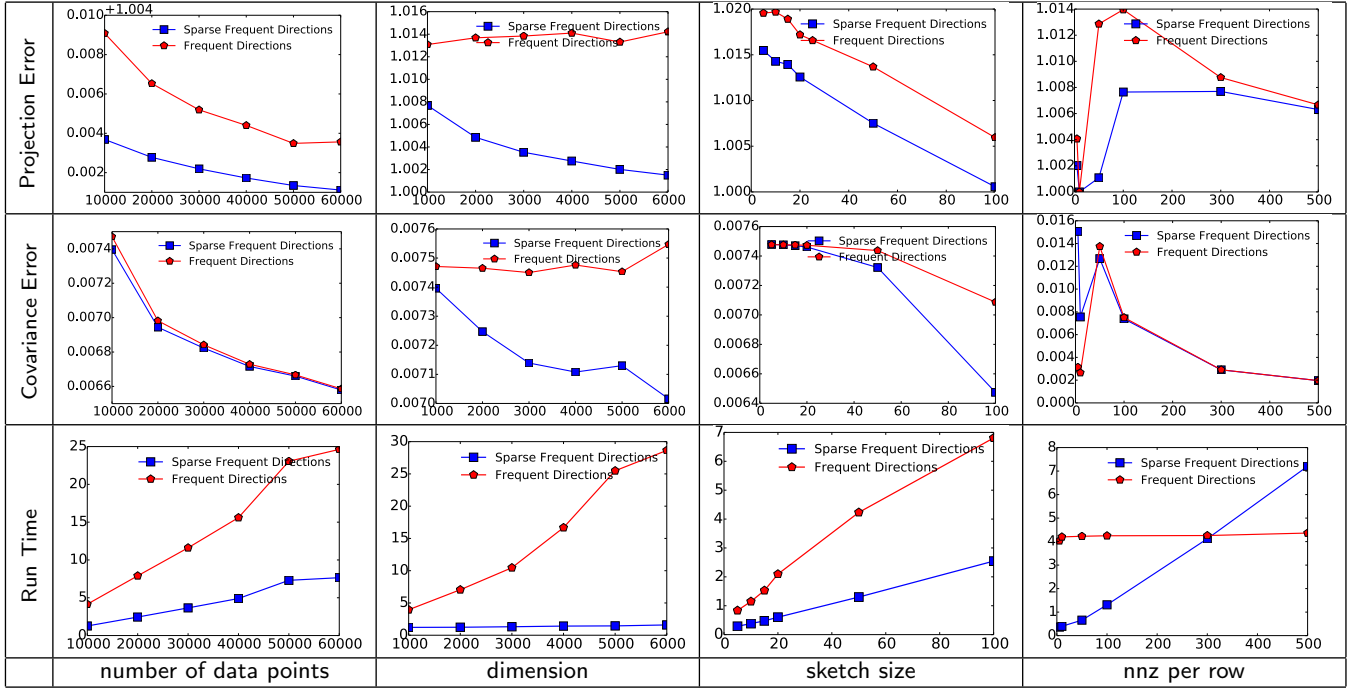


Table 2: Comparing performance of FrequentDirections and SparseFrequentDirections on synthetic data. Each column reports the measurement against one parameter; ordered from left to right it is number of datapoints (n), dimension (d), sketch size (ℓ), and number of non-zeros (nnz) per row. Table 1 lists default value of all parameters.




Figure 1: Comparing performance of FrequentDirections and SparseFrequentDirections on 20 Newsgroups dataset. We plot Projection Error, Covariance Error, and Run Time as a function of sketch size (ℓ).

QUENTDIRECTIONS. We can likely attribute this improvement to being able to process more rows in each batch, and hence needing to perform the shrinking operation fewer overall times. The one small exception to SPARSEFREQUENTDIRECTIONS having less Covariance Error than FREQUENTDIRECTIONS is for extreme sparse datasets in the leftmost data points of Table 2, last column – we attribute this to some peculiar orthogonality of columns with near equal norms due to extreme sparsity.

6. REFERENCES

- [1] Nick Asendorf, Madison McGaffin, Matt Prelee, and Ben Schwartz. Algorithms for completing a user ratings matrix.
- [2] Christos Boutsidis, Petros Drineas, and Malik Magdon-Ismail. Near optimal column-based matrix reconstruction. In *Foundations of Computer Science, 2011 IEEE 52nd Annual Symposium on*, pages 305–314. IEEE, 2011.

- [3] Christos Boutsidis, Michael W Mahoney, and Petros Drineas. An improved approximation algorithm for the column subset selection problem. In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2009.
- [4] Christos Boutsidis and David P Woodruff. Optimal cur matrix decompositions. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 353–362. ACM, 2014.
- [5] Matthew Brand. Incremental singular value decomposition of uncertain data with missing values. In *Computer Vision—ECCV 2002*, pages 707–720. Springer, 2002.
- [6] Kenneth L. Clarkson and David P. Woodruff. Low rank approximation and regression in input sparsity time. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, 2013.
- [7] Amey Desai, Mina Ghashami, and Jeff M Phillips. Improved practical matrix sketching with guarantees. *arXiv preprint arXiv:1501.06561*, 2015.
- [8] Amit Deshpande and Santosh Vempala. Adaptive sampling and fast low-rank matrix approximation. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 2006.
- [9] Inderjit S Dhillon and Dharmendra S Modha. Concept decompositions for large sparse text data using clustering. *Machine learning*, 42(1-2):143–175, 2001.
- [10] Petros Drineas and Ravi Kannan. Pass efficient algorithms for approximating large matrices. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2003.
- [11] Petros Drineas, Ravi Kannan, and Michael W. Mahoney. Fast monte carlo algorithms for matrices ii: Computing a low-rank approximation to a matrix. *SIAM Journal on Computing*, 36(1):158–183, 2006.
- [12] Petros Drineas, Ravi Kannan, and Michael W Mahoney. Fast monte carlo algorithms for matrices iii: Computing a compressed approximate matrix decomposition. *SIAM Journal on Computing*, 36(1):184–206, 2006.
- [13] Petros Drineas, Iordanis Kerenidis, and Prabhakar Raghavan. Competitive recommendation systems. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 82–90. ACM, 2002.
- [14] Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Relative-error cur matrix decompositions. *SIAM Journal on Matrix Analysis and Applications*, 30(2):844–881, 2008.
- [15] Petros Drineas, Michael W Mahoney, S Muthukrishnan, and Tamás Sarlós. Faster least squares approximation. *Numerische Mathematik*, 117:219–249, 2011.
- [16] Alan Frieze, Ravi Kannan, and Santosh Vempala. Fast monte-carlo algorithms for finding low-rank approximations. *Journal of the ACM (JACM)*, 51(6):1025–1041, 2004.
- [17] Mina Ghashami, Edo Liberty, Jeff M Phillips, and David P Woodruff. Frequent directions: Simple and deterministic matrix sketching. *arXiv preprint arXiv:1501.01711*, 2015.
- [18] Mina Ghashami and Jeff M Phillips. Relative errors for deterministic low-rank matrix approximations. In *Proceedings of 25th ACM-SIAM Symposium on Discrete Algorithms*, 2014.
- [19] Gene H Golub and Charles F Van Loan. *Matrix computations*, volume 3. JHUP, 2012.
- [20] Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- [21] Peter Hall, David Marshall, and Ralph Martin. Incremental eigenanalysis for classification. In *Proceedings of the British Machine Vision Conference*, 1998.
- [22] Ken Lang. Newsweeder: Learning to filter netnews. In *Proceedings of the Twelfth International Conference on Machine Learning*, pages 331–339, 1995.
- [23] A Levey and Michael Lindenbaum. Sequential karhunen-loeve basis extraction and its application to images. *Image Processing, IEEE Transactions on*, 9(8):1371–1374, 2000.
- [24] Edo Liberty. Simple and deterministic matrix sketching. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2013.
- [25] Edo Liberty. Simple and deterministic matrix sketching. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2013.
- [26] Edo Liberty, Franco Woolfe, Per-Gunnar Martinsson, Vladimir Rokhlin, and Mark Tygert. Randomized algorithms for the low-rank approximation of matrices. *Proceedings of the National Academy of Sciences*, 104(51):20167–20172, 2007.
- [27] Michael W Mahoney and Petros Drineas. Cur matrix decompositions for improved data analysis. *Proceedings of the National Academy of Sciences*, 106(3):697–702, 2009.
- [28] Jiří Matoušek. On variants of the johnson–lindenstrauss lemma. *Random Structures & Algorithms*, 33(2):142–156, 2008.
- [29] Cameron Musco and Christopher Musco. Stronger approximate singular value decomposition via the block lanczos and power methods. *arXiv preprint arXiv:1504.05477*, 2015.
- [30] Jelani Nelson and Huy L. Nguyen. OSNAP: Faster numerical linear algebra algorithms via sparser subspace embeddings. In *Proceedings of 54th IEEE Symposium on Foundations of Computer Science*, 2013.
- [31] Christos H. Papadimitriou, Hisao Tamaki, Prabhakar Raghavan, and Santosh Vempala. Latent semantic indexing: A probabilistic analysis. In *Proceedings of the 17th ACM Symposium on Principles of Database Systems*, 1998.
- [32] Vladimir Rokhlin, Arthur Szlam, and Mark Tygert. A randomized algorithm for principal component analysis. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1100–1124, 2009.
- [33] David A Ross, Jongwoo Lim, Rwei-Sung Lin, and Ming-Hsuan Yang. Incremental learning for robust

- visual tracking. *International Journal of Computer Vision*, 77(1-3):125–141, 2008.
- [34] Mark Rudelson and Roman Vershynin. Sampling from large matrices: An approach through geometric functional analysis. *Journal of the ACM*, 54(4):21, 2007.
- [35] Tamas Sarlos. Improved approximation algorithms for large matrices via random projections. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, 2006.
- [36] Santosh S Vempala. *The random projection method*, volume 65. American Mathematical Soc., 2004.
- [37] Rafi Witten and Emmanuel Candès. Randomized algorithms for low-rank matrix factorizations: sharp performance bounds. *Algorithmica*, 72(1):264–281, 2013.
- [38] Franco Woolfe, Edo Liberty, Vladimir Rokhlin, and Mark Tygert. A fast randomized algorithm for the approximation of matrices. *Applied and Computational Harmonic Analysis*, 25(3):335–366, 2008.