

Low rank approximation and decomposition of large matrices using error correcting codes

Shashanka Ubaru, Arya Mazumdar *Member, IEEE*, and Yousef Saad

Abstract—Low rank approximation is an important tool used in many applications of signal processing and machine learning. Recently, randomized sketching algorithms were proposed to effectively construct low rank approximations and obtain approximate singular value decompositions of large matrices. Similar ideas were used to solve least squares regression problems. In this paper, we show how matrices from error correcting codes can be used to find such low rank approximations and matrix decompositions, and extend the framework to linear least squares regression problems.

The benefits of using these code matrices are the following: (i) They are easy to generate and they reduce randomness significantly. (ii) Code matrices with mild properties satisfy the subspace embedding property, and have a better chance of preserving the geometry of an entire subspace of vectors. (iii) For parallel and distributed applications, code matrices have significant advantages over structured random matrices and Gaussian random matrices. (iv) Unlike Fourier or Hadamard transform matrices, which require sampling $O(k \log k)$ columns for a rank- k approximation, the log factor is not necessary for certain types of code matrices. That is, $(1 + \epsilon)$ optimal Frobenius norm error can be achieved for a rank- k approximation with $O(k/\epsilon)$ samples. (v) Fast multiplication is possible with structured code matrices, so fast approximations can be achieved for general dense input matrices. (vi) For least squares regression problem $\min \|Ax - b\|_2$ where $A \in \mathbb{R}^{n \times d}$, the $(1 + \epsilon)$ relative error approximation can be achieved with $O(d/\epsilon)$ samples, with high probability, when certain code matrices are used.

Index Terms—Error correcting codes, low rank approximation, matrix decomposition, randomized sketching algorithms, subspace embedding.

I. INTRODUCTION

MANY scientific computations, signal processing, data analysis and machine learning applications, lead to large dimensional matrices that can be well approximated by a low dimensional (low rank) basis [36], [55], [26], [18]. It is more efficient to solve such computational problems by first transforming these high dimensional matrices into a low dimensional space, while preserving the invariant subspace that captures the essential information of the matrix. Low-rank matrix approximation is an integral component of tools such as principal component analysis (PCA) [30]. It is also an important instrument used in many applications like computer vision (e.g., face recognition) [50], signal processing

(e.g., adaptive beamforming) [42], recommender systems [19], information retrieval and latent semantic indexing [7], [6], web search modeling [31], DNA microarray data [3], [43] and text mining, to name a few examples. Several algorithms have been proposed in the literature for finding low rank approximations of matrices [36], [55], [26], [18], [11], [27]. Recently, research focussed on developing techniques that use randomization for computing low rank approximations and decompositions of such large matrices [26], [47], [33], [53], [41], [12], [35], [52]. It is found that randomness provides an effective way to construct low dimensional bases with high reliability and computational efficiency. Similar ideas based on random sampling have been proposed in the recent literature for solving least squares (ℓ_2) linear regression problems [20], [47], [45], [12], [41], [21], [13].

Randomization techniques for matrix approximations aim to compute a basis that approximately spans the range of an $m \times n$ input matrix A , by sampling¹ the matrix A using random matrices, e.g. i.i.d Gaussian [26]. This task is accomplished by first forming the matrix-matrix product $Y = A\Omega$, where Ω is an $n \times \ell$ random matrix of smaller dimension $\ell \ll \{m, n\}$, and then computing the orthonormal basis of $Y = QR$ that identifies the range of the reduced matrix Y . It can be shown that $A \approx QQ^T A$ with high probability. It has been shown that structured random matrices, like subsampled random Fourier transform (SRFT) and Hadamard transform (SRHT) matrices can also be used in place of fully random matrices [53], [32], [41], [49]. This paper demonstrates how matrices from error correcting codes can be a good choice for computing such low rank approximations.

The input matrices whose low rank approximation is to be computed usually have very large dimensions (e.g., in the order of $10^6 - 10^9$ [26], [54]). In order to form a Gaussian (a fully) random matrix that samples the input matrix, we need to generate a large quantity of random numbers. This could be a serious practical issue (in terms of time complexity and storage). This issue can be addressed by using the structured random matrices, like SRFT and SRHT matrices. An important practical advantage of using these structured random matrices is that their structure allows the computation of matrix-matrix product at a cost of $O(mn \log_2 \ell)$ making the algorithms fast (also known as fast transforms) for general dense input matrices. However, with these matrices mixing of columns might not be as uniform, and there is potential loss in the accuracy.

S. Ubaru and Y. Saad are with the Department of Computer Science and Engineering, University of Minnesota, Twin Cities, MN USA (e-mail: ubaru001@umn.edu, saad@cs.umn.edu). A. Mazumdar is with Department of Electrical and Computer Engineering, University of Minnesota, Twin Cities, MN USA (e-mail: arya@umn.edu). This work was supported by NSF under grant NSF/CCF-1318597 (S. Ubaru and Y. Saad) and NSF/CCF-1318093 (A. Mazumdar). A preliminary version of this work has appeared as conference proceedings in the 32nd International Conference on Machine Learning [51].

¹Sampling is sometimes called ‘sketching’, popular in data streaming model applications [12], [52].

Another drawback with fast transforms is that for parallel and distributed applications, particularly when the input matrix is sparse and/or its columns are distributively stored, it is found that FFT-like algorithms are significantly slower due to communication issues or other machine related issues (machines are optimized for matrix-vector operations) [54]. Similar issues arise when the input matrices are streamed [26], [12], we give more details on these in sec. VI-E and VIII. Also for a rank- k approximation, these matrices require sampling $\ell = O(k \log k)$ columns. Other practical issues arise such as: the Fourier Transform matrices require handling complex numbers and the Hadamard matrices exist only for the sizes which are in powers of 2. All these drawbacks can be overcome if the code matrices presented in this paper are used for sampling the input matrices.

In digital communication, information is encoded (by adding redundancy) to (predominantly binary) vectors or codewords, that are then transmitted over a noisy channel [14]. These codewords are required to be far apart in terms of some distance metric for noise-resilience. Coding schemes usually generate codewords that maintain a fixed minimum Hamming distance between each other, hence they are widespread and act like random vectors. We can define probability measures for matrices formed by stacking up these codewords (see section II-B for details). In this paper, we explore the idea of using subsampled versions of these code matrices as sampling (sketching) matrices in the randomized techniques for matrix approximations.

The idea of using code matrices for such applications is not new in the literature. A class of dual BCH code matrices were used in [2], [32] as Fast Johnson-Lindenstrauss Transform (FJLT) to perform fast dimensionality reduction of vectors. Code matrices have also been used in applications of sparse signal recovery, such as compressed sensing [4] and group testing [22], [38]. For matrix approximations, it is important to show that the sampling matrices used can approximately preserve the geometry of an entire subspace of vectors, i.e., they satisfy the ‘‘subspace embedding’’ property [47], [52]. In section VI-B, we show that the subsampled code matrices with certain mild properties satisfy this subspace embedding property with high probability. Similar to Fourier and Hadamard sampling matrices, fast multiplication is possible with code matrices from certain class of codes due to their structure (see section V for details). Hence, fast approximations can be achieved for general dense input matrices, since the matrix-matrix product $A\Omega$ can be computed in $O(mn \log_2 \ell)$ cost with such code matrices.

In addition, the shortcomings of SRFT/SRHT matrices in parallel and distributed environments, and in data streaming models can be overcome by using code matrices (details in secs. V, VIII). For certain code matrices, the logarithmic factor in the number of samples is not required (see sec. VI-F for an explanation). This is a significant theoretical result that shows that order optimality can be achieved in the number of samples required with partially random matrices. Similar improvements were posed as an open problem in [17] and in [41]. In the context of sparse approximations such improvements appear as main results in many places, see Table 1 of [5].

A preliminary version of part of this paper has appeared in the conference proceedings of the 32nd International Conference on Machine Learning [51]. In this paper, we improve the theoretical results obtained in [51]. In particular, we show that ‘ $(1+\epsilon)$ ’ optimal Frobenius norm error can be achieved for low rank approximation using code matrices with a mild condition. We also improve the spectral norm error bound obtained in [51]. Furthermore, we give several additional details that were omitted in the shorter version [51] such as, details on the computational cost, the choice of the code matrices for different scenarios (type of input matrices, computational environments, etc.) and others. We also extend the theory to show how code matrices can be used to solve linear least squares (ℓ_2) regression problems (see below).

One of the key applications where the randomized approximation (or sketching) algorithms are used is in approximately solving overdetermined least squares regression problem faster [20], [47], [21], [13]. Here, we are given a matrix $A \in \mathbb{R}^{n \times d}$ and a vector $b \in \mathbb{R}^n$, with $n \gg d$. The goal is to solve the least squares regression problem $\min_x \|Ax - b\|_2$ faster, (where $\|\cdot\|_2$ is ℓ_2 norm) and output a vector x' such that, with high probability,

$$\|Ax' - b\|_2 \leq (1 + \epsilon) \|A\hat{x} - b\|_2,$$

where \hat{x} is the ℓ_2 minimizer given by the Moore-Penrose pseudo inverse of A , i.e., $\hat{x} = A^\dagger b$ [24]. For details on the applications where we encounter such extremely overdetermined linear system of equations, we refer to [54]. The idea of randomized approximation [20] is to use a sampling (sketching) matrix to reduce the dimensions of A and b , and then solve the smaller problem to obtain x' . In section VII, we show how the code matrices can be used as the sampling matrix in such least squares regression problems.

The organization of the rest of this paper is as follows: Section II gives the notation and key definitions used, the problem set up and a brief introduction to error correcting coding techniques. Section III discusses the construction of the subsampled code matrices and the intuition behind the construction. The algorithm of the present paper is described in section IV and its computational cost is discussed in section V. Section VI discusses the error analysis for the algorithm, by deriving bounds for the Frobenius norm error, spectral norm error and the singular values obtained from the algorithm. We show that $(1 + \epsilon)$ relative Frobenius norm error approximation can be achieved with code matrices. For this, the code matrices need to satisfy two key properties which are discussed in section VI-B. The bounds for the approximation errors and the singular values obtained are derived in section VI-D. In section VII, we extend the framework to linear least squares (ℓ_2) regression problem and in section VIII, we discuss the choice of error correcting codes for different types of input matrices and computational environments. Section IX illustrates the performance of code matrices via few numerical experiments.

II. PRELIMINARIES

First, we present some of the notation used, some key definitions and give a brief description of error correcting codes used in communication systems and information theory.

A. Notation and Definitions

Throughout the paper, $\|\cdot\|_2$ refers to the ℓ_2 or spectral norm. We use $\|\cdot\|_F$ for the Frobenius norm. The singular value decomposition (SVD) of a matrix A is denoted by $A = U\Sigma V^\top$ and the singular values by $\sigma_j(A)$. We use e_j for the j th standard basis vector. Given a subset T of indices in $\{1, \dots, 2^r\}$ with size n and $r \geq \lceil \log_2 n \rceil$, we define a restriction (sampling) operator $S_T : \mathbb{R}^{2^r} \rightarrow \mathbb{R}^T$ to be the projection of a vectors on to the coordinates of T :

$$S_T \mathbf{x} = (x_j)_{j \in T}.$$

A Rademacher random variable takes values ± 1 with equal probability. We write ε for a Rademacher variable.

Next, we define two key properties which will be used frequently in our theoretical analysis. In order to use the subsampled code matrices as sampling matrices, they need to satisfy these two properties, see section VI. First is the Johnson-Lindenstrauss Transform (JLT) [29] which played a key role in the development of embedding-based randomized sampling. Sarlos [47] developed the important relation between JLT and random matrix sampling (also known as subspace embedding), see details in section VI-B. The JLT property is defined as [47]:

Definition 1 (Johnson-Lindenstrauss Transform): A matrix $\Omega \in \mathbb{R}^{n \times \ell}$ forms a Johnson-Lindenstrauss Transform with parameters ϵ, δ, d or JLT(ϵ, δ, d) for any $0 < \epsilon, \delta < 1$, if for a d -element subspace $V \subset \mathbb{R}^n$, and for all $v \in V$ it holds

$$(1 - \epsilon)\|v\|_2^2 \leq \|\Omega^\top v\|_2^2 \leq (1 + \epsilon)\|v\|_2^2$$

with probability $1 - \delta$.

The other key property which the code matrices need to satisfy is the subspace embedding property defined below.

Definition 2 (Subspace Embedding): A matrix $\Omega \in \mathbb{R}^{n \times \ell}$ is a $(1 \pm \epsilon)$ ℓ_2 -subspace embedding for the row space of an $m \times n$ matrix A , if for an orthonormal basis $V \in \mathbb{R}^{n \times k}$ that spans the row space of A , for all $x \in \mathbb{R}^k$

$$\|\Omega^\top Vx\|_2^2 = (1 \pm \epsilon)\|Vx\|_2^2 = (1 \pm \epsilon)\|x\|_2^2.$$

The above definition is useful when the sampling is achieved column-wise. A similar definition for row-wise sampling holds for an orthonormal matrix $U \in \mathbb{R}^{m \times k}$ which spans the column space of A , see [52]. The above definition simplifies to the following condition:

$$\|V^\top \Omega \Omega^\top V - I\|_2 \leq \epsilon. \quad (1)$$

The matrix Ω with subspace embedding property, satisfying the above condition is said to approximately preserve the geometry of an entire subspace of vectors [49].

In low rank approximation methods, we compute an orthonormal basis that approximately spans the range of an $m \times n$ input matrix A . That is, a matrix Q having orthonormal columns such that $A \approx QQ^\top A$. The basis matrix Q must contain as few columns as possible, but it needs to be an accurate approximation of the input matrix. I.e., we seek a matrix Q with k orthonormal columns such that

$$\|A - QQ^\top A\|_\xi \leq e_k, \quad (2)$$

for a positive error tolerance e_k and $\xi \in \{2, F\}$.

The best rank- k approximation of A with respect to both Frobenius and spectral norm is given by the Eckart-Young theorem [23], and it is $A_k = U_k \Sigma_k V_k^\top$, where U_k and V_k are the k -dominant left and right singular vectors of A , respectively and diagonal Σ_k contains the top k singular values of A . So, the optimal Q in (2) will be U_k for $\xi \in \{2, F\}$, and $e_k = \sigma_{k+1}$ for $\xi = 2$ and $e_k = \sum_{j=k+1}^n \sigma_j$ for $\xi = F$.

B. Error Correcting Codes

In communication systems, data are transmitted from a source (transmitter) to a destination (receiver) through physical channels. These channels are usually noisy, causing errors in the data received. In order to facilitate detection and correction of these errors in the receiver, error correcting codes are used [34]. A block of information (data) symbols are encoded into a binary vector², also called a codeword. Error correcting coding methods check the correctness of the codeword received. The set of codewords corresponding to a set of data vectors (or symbols) that can possibly be transmitted is called the *code*. As per our definition a code \mathcal{C} is a subset of \mathbb{F}_2^ℓ , ℓ being an integer.

A code is said to be linear when adding two codewords of the code coordinate-wise using modulo-2 arithmetic results in a third codeword of the code. Usually a linear code \mathcal{C} is represented by the tuple $[\ell, r]$, where ℓ represents the codeword length and $r = \log_2 |\mathcal{C}|$ is the number of information bits that can be encoded by the code. There are $\ell - r$ redundant bits in the codeword, which are sometimes called parity check bits, generated from messages using an appropriate rule. It is not necessary for a codeword to have the corresponding information bits as r of its coordinates, but the information must be uniquely recoverable from the codeword.

It is perhaps obvious that a linear code \mathcal{C} is a linear subspace of dimension r in the vector space \mathbb{F}_2^ℓ . The basis of \mathcal{C} can be written as the rows of a matrix, which is known as the generator matrix of the code. The size of the generator matrix G is $r \times \ell$, and for any information vector $\mathbf{m} \in \mathbb{F}_2^r$, the corresponding codeword is found by the following linear map:

$$\mathbf{c} = \mathbf{m}G.$$

Note that all the arithmetic operations above are over the binary field \mathbb{F}_2 .

To encode r bits, we must have 2^r unique codewords. Then, we may form a matrix of size $2^r \times \ell$ by stacking up all codewords that are formed by the generator matrix of a given linear coding scheme,

$$\underbrace{C}_{2^r \times \ell} = \underbrace{M}_{2^r \times r} \underbrace{G}_{r \times \ell}. \quad (3)$$

For a given tuple $[\ell, r]$, different error correcting coding schemes have different generator matrices and the resulting codes have different properties. For example, for any two integers t and q , a BCH code [9] has length $\ell = 2^q - 1$ and dimension $r = 2^q - 1 - tq$. Any two codewords in this

²Here, and in the rest of the text, we are considering only binary codes. In practice, codes over other alphabets are also quite common.

BCH code maintain a minimum (Hamming) distance of at least $2t + 1$ between them. The minimum pairwise distance is an important parameter of a code and is called just the distance of the code.

As a linear code \mathcal{C} is a subspace of a vector space, the null space \mathcal{C}^\perp of the code is another well defined subspace. This is called the *dual* of the code. For example, the dual of the $[2^q - 1, 2^q - 1 - tq]$ -BCH code is a code with length $2^q - 1$, dimension tq and minimum distance at least $2^{q-1} - (t-1)2^{q/2}$. The minimum distance of the dual code is called the *dual distance* of the code.

Depending on the coding schemes used, the codeword matrix C will have a variety of favorable properties, e.g., low coherence which is useful in compressed sensing [37], [4]. Since the codewords need to be far apart, they show some properties of random vectors. We can define probability measures for codes generated from a given coding scheme. If $\mathcal{C} \subset \{0, 1\}^\ell$ is an \mathbb{F}_2 -linear code whose dual \mathcal{C}^\perp has a minimum distance above k (dual distance $> k$), then the code matrix is an *orthogonal array* of strength k [16]. This means, in such a code \mathcal{C} , for any k entries of a randomly and uniformly chosen codeword \mathbf{c} say $\mathbf{c}' = \{c_{i_1}, c_{i_2}, \dots, c_{i_k}\}$ and for any k bit binary string α , we have

$$\Pr[\mathbf{c}' = \alpha] = 2^{-k}.$$

This is called the k -wise independence property of codes. We will use this property of codes in our theoretical analysis (see section VI for details).

The codeword matrix C has 2^r codewords each of length ℓ (a $2^r \times \ell$ matrix), i.e., a set of 2^r vectors in $\{0, 1\}^\ell$. Given a codeword $\mathbf{c} \in \mathcal{C}$, let us map it to a vector $\phi \in \mathbb{R}^\ell$ by setting $1 \rightarrow \frac{-1}{\sqrt{2^r}}$ and $0 \rightarrow \frac{1}{\sqrt{2^r}}$. In this way, a binary code \mathcal{C} gives rise to a code matrix $\Phi = (\phi_1^\top, \dots, \phi_{2^r}^\top)^\top$. Such a mapping is called binary phase-shift keying (BPSK) and appeared in the context of sparse recovery (e.g., p. 66 [37]). For codes with dual distance ≥ 3 , this code matrix Φ will have orthonormal columns, see lemma 9. In section VI-B, we will show that these code matrices with certain mild properties satisfy the JLT and the subspace embedding properties and preserve the geometry of vector subspaces with high probability. In the randomized techniques for matrix approximations, we can use a subsampled and scaled version of this matrix Φ to sample a given input matrix and find its active subspace.

III. CONSTRUCTION OF SUBSAMPLED CODE MATRIX

For an input matrix A of size $m \times n$ and a target rank k , we choose $r \geq \lceil \log_2 n \rceil$ as the dimension of the code (length of the message vector) and $\ell > k$ as the length of the code. The value of ℓ will depend on the coding scheme used, particularly on the dual distance of the code (details in section VI-B). We consider an $[\ell, r]$ -linear coding scheme and form the sampling matrix as follows: We draw the sampling test matrix say Ω as

$$\Omega = \sqrt{\frac{2^r}{\ell}} DS\Phi, \quad (4)$$

where

- D is a random $n \times n$ diagonal matrix whose entries are independent random signs, i.e., random variables uniformly distributed on $\{\pm 1\}$.
- S is a uniformly random downsampler, an $n \times 2^r$ matrix whose n rows are randomly selected from a $2^r \times 2^r$ identity matrix.
- Φ is the $2^r \times \ell$ code matrix, generated using an $[\ell, r]$ -linear coding scheme, with BPSK mapping and scaled by $2^{-r/2}$ such that all columns have unit norm.

A. Intuition

The design of a Subsampled Code Matrix (SCM) is similar to the design of SRFT and SRHT matrices. The intuition for using such a design is well established in [49], [26]. The matrix Φ has entries with magnitude $\pm 2^{-r/2}$ and has orthonormal columns when a coding scheme with dual distance of the code ≥ 3 is used.

The scaling $\sqrt{\frac{2^r}{\ell}}$ is used to make the energy of the sampling matrix equal to unity, i.e., to make the rows of Ω unit vectors. The objective of multiplying by the matrix D is twofold. The first purpose is to flatten out the magnitudes of input vectors, see [49] for the details. For a fixed unit vector \mathbf{x} , the first component of $\mathbf{x}^\top DS\Phi$ is given by $(\mathbf{x}^\top DS\Phi)_1 = \sum_{i=1}^n x_i \varepsilon_i \phi_{j_1}$, where ϕ_{j_1} are components of the first column of the code matrix Φ , with the indices j 's are such that $S_{ij} = 1$ for $i = 1, \dots, n$ and ε_i is the Rademacher variable from D . This sum has zero mean and since entries of Φ have magnitude $2^{-r/2}$, the variance of the sum is 2^{-r} . The Hoeffding inequality [28] shows that

$$\mathbb{P}\{|\mathbf{x}^\top DS\Phi)_1| \geq \tilde{t}\} \leq 2e^{-2^r \tilde{t}^2/2}.$$

That is, the magnitude of the first component of $\mathbf{x}^\top DS\Phi$ is about $2^{-r/2}$. Similarly, the argument holds for the remaining entries. Therefore, it is unlikely that any one of the ℓ components of $\mathbf{x}^\top DS\Phi$ is larger than $\sqrt{2 \log(2\ell)/2^r}$, (with a failure probability of ℓ^{-1}).

The second purpose of multiplying by D is as follows: The code matrix Φ with a dual distance $> k$ forms a deterministic k -wise independent matrix. Multiplying this Φ matrix by D (with independent random signs on the diagonal) results in a k -wise independent random matrix. Note that uniform downsampling of the matrix will not affect this property. Hence, the subsampled code matrix SCM Ω will be a k -wise independent random matrix. This is a key property of SCM Ω that we will use to prove the JLT and the subspace embedding properties for SCM, see section VI-B.

The downsampler S is a formal way of saying, if $n < 2^r$, we choose n out of 2^r possible codewords to form the sampling matrix Ω . Uniform downsampling is used in the theoretical analysis to get an upper bound for the singular values of Ω (see sec. VI-D). In practice, we choose n numbers between 1 to 2^r , use the binary representation of these numbers as the message vectors (form M) and use the generator matrix G of the coding scheme selected to form the sampling matrix Ω , using (3) and BPSK mapping. For dense input matrices, it is advantageous to choose these numbers (message vectors) to be 1 to $2^{\lceil \log_2 n \rceil}$, to exploit the availability of fast multiplication (see details in section V).

IV. ALGORITHM

We use the same prototype algorithm as discussed in [26] for the low rank approximation and decomposition of an input matrix A . The subsampled code matrix (SCM) Ω given in (4), generated from a chosen coding scheme is used as the sampling test matrix. The algorithm is as follows:

Algorithm 1 Prototype Algorithm

Input: An $m \times n$ matrix A , a target rank k .

Output: Rank- k factors U, Σ , and V in an approximate SVD $A \approx U\Sigma V^\top$.

1. Form an $n \times \ell$ subsampled code matrix Ω , as described in Section III and (4), using an $[\ell, r]$ -linear coding scheme, where $\ell > k$ and $r \geq \lceil \log_2 n \rceil$.
 2. Form the $m \times \ell$ sample matrix $Y = A\Omega$.
 3. Form an $m \times \ell$ orthonormal matrix Q such that $Y = QR$.
 4. Form the $\ell \times n$ matrix $B = Q^\top A$.
 5. Compute the SVD of the small matrix $B = \hat{U}\Sigma V^\top$.
 6. Form the matrix $U = Q\hat{U}$.
-

The prototype algorithm requires only two passes over the input matrix (single pass algorithms can also be developed [26, §5.5]), as opposed to $O(k)$ passes required for classical algorithms. This is particularly significant when the input matrix is very large to fit in fast memory (RAM) or when the matrices are streamed [26]. It is known that, the randomized techniques allow us to reorganize the calculations required to exploit the input matrix properties and the modern computer architecture more efficiently. The algorithm is also well suited for implementation in parallel and distributed environments, see [54]. For more details on all the advantages of randomized methods over classical techniques, we refer to [47], [53], [26].

Several algorithms have been developed in the literature which build on the above prototype algorithm. An important requirement (rather a drawback) of the prototype algorithm is that, to obtain a good approximation, the algorithm requires the singular values of the input matrix to decay rapidly [26]. Methods such as randomized power method [26], [44], [25] and randomized block Krylov subspace methods [40] have been proposed to improve the performance (accuracy) of the prototype algorithm, particularly when the singular values of the input matrix decay slowly. In these methods, step 2 in Algorithm 1 is replaced by $Y = (AA^\top)^{q'} A\Omega$, where q' is a small integer, or a block Krylov subspace [24]. However, these algorithms require $2(q' + 1)$ passes over A . Use of structured random matrices like SRFT and SRHT are proposed for a faster computation of the matrix product $Y = A\Omega$ [53], [32], for dense input matrices. The use of sparse random matrices, e.g. CountSketch matrix [13], [52], is proposed to achieve faster computations when the input matrix is sparse.

Algorithm 1 can also be modified to obtain the eigenvalue decompositions of square input matrices [26]. In all the above mentioned modified algorithms, we can use our subsampled code matrix as the random sampling (sketching) matrix. For the analysis in the following sections, we shall consider the prototype algorithm 1.

V. COMPUTATIONAL COST

One of the key advantages of using structured random matrices (SRFT or SRHT) in the randomized sketching algorithms is that, for a general dense matrix, we can compute the matrix-matrix product $Y = A\Omega$ in $O(mn \log_2 \ell)$ time exploiting the structure of Fourier/Hadamard matrices [47], [53], [41], [32], [26]. The idea of fast multiplications was inspired by articles on Fast Johnson-Lindenstrauss Transform (FJLT) [1], [2] where it was shown that matrix-vector products with such structured matrices can be computed in $O(n \log_2 \ell)$ time. Interestingly, Ailon and Liberty [2] give dual BCH code matrices and Hadamard matrices (that are actually a special codes called 1st order Reed-Muller codes) as examples for such structured matrices.

Many, if not most of the structured codes can be decoded using the Fast Fourier Transform (FFT) [8]. The corresponding $2^r \times \ell$ code matrix Φ of such structured codes (after BPSK mapping) will have every column of Φ equal to some column of a $2^r \times 2^r$ Hadamard matrix, see definition 2.2 in [2]. Hence, for a general dense matrix in RAM, the matrix-matrix product $Y = A\Omega$ with these structure code matrices can be computed in $O(mn \log_2 \ell)$ time using the ‘Trimmed Hadamard transform’ technique described in [2], [32]. If $n < 2^r$, we choose the top $2^{\lceil \log_2 n \rceil}$ codewords of Φ as the rows of Ω such that the columns of Ω are some columns of a $2^{\lceil \log_2 n \rceil} \times 2^{\lceil \log_2 n \rceil}$ Hadamard matrix.

Fast multiplications are possible with matrices from another class of codes known as cyclic codes. In cyclic codes, a circular shift of a codeword results in another codeword of that code. So, a $2^r \times \ell$ code matrix Φ generated using an $[\ell, r]$ -cyclic code scheme will consist of $2^r/\ell$ blocks of circulant matrices of size $\ell \times \ell$ (when appropriately rearranged). It is known that the matrix-vector products with circulant matrices can be computed in $O(\ell \log_2 \ell)$ operations via FFT [24]. So, for a general dense input matrix, the matrix-matrix product $Y = A\Omega$ with such cyclic code matrices can be computed in $O(mn \log_2 \ell)$ time.

The remaining steps (steps 3 – 6) of the algorithm can be computed in $O((m + n)k^2)$ time using the row extraction method described in [26]. Therefore, for a general dense input matrix in RAM, the total computational cost of Algorithm 1 using SCM is $O(mn \log_2 \ell + (m + n)k^2)$ for structured and cyclic codes.

For sparse input matrices or when the columns of A are distributively stored, we can choose codewords at random from a desired code (as described earlier) making Ω unstructured and $Y = A\Omega$ a dense transform, similar to a random sampling matrix. The computational cost of the algorithm for such cases is $O(nz(A)\ell + (m + n)k^2)$, where $nz(A)$ is the number of nonzero entries in the input matrix A . We will see that, for code matrices with certain properties, $\ell = O(k/\epsilon)$ which will be advantageous in these cases (compared to SRFT/SRHT which require $\ell = O(k \log k/\epsilon)$). Additional details of the choice of the code matrix for different types of input matrices and computational environments are given in section VIII.

VI. ANALYSIS

This section discusses the performance (error) analysis of the subsampled code matrices (SCM) as sampling matrices in Algorithm 1. We will prove that an approximation error of $(1 + \epsilon)$ times the best rank- k approximation (Frobenius norm error) possible for a given matrix A can be achieved with code matrices. That is,

$$\|A - \hat{A}_k\|_F \leq (1 + \epsilon)\|A - A_k\|_F,$$

where \hat{A}_k is the rank- k approximation obtained from Algorithm 1 and A_k is the best rank- k approximation as defined in section II. In order to prove this, we show that SCM satisfies the Johnson Lindenstrauss Transforms (JLT) and the subspace embedding properties via the k -wise independence property of the codes (the relation between these two properties and $(1+\epsilon)$ approximation is given in sec. VI-D, Lemma 7). We also derive the bounds for the spectral norm error and the singular values obtained, based on the deterministic error bounds in the literature for the algorithm for a given sampling matrix Ω .

A. Setup

Let A be an $m \times n$ input matrix with SVD given by $A = U\Sigma V^\top$, and partition its SVD as follows

$$A = U \begin{bmatrix} \Sigma_1 & & \\ & \Sigma_2 & \\ & & \end{bmatrix} \begin{bmatrix} V_1^\top \\ V_2^\top \\ \end{bmatrix} \begin{matrix} k & n-k \\ & n-k \end{matrix}. \quad (5)$$

Let Ω be the $n \times \ell$ test (sampling) matrix, where ℓ is the number of samples. Consider the matrices

$$\Omega_1 = V_1^\top \Omega \quad \text{and} \quad \Omega_2 = V_2^\top \Omega. \quad (6)$$

The objective of any low rank approximation algorithm is to approximate the subspace that spans the top k left singular vectors of A . Hence, for a given sampling matrix Ω , the key challenge is to show that Ω_1 is full rank. That is, we need to show that for any orthonormal matrix V of dimension k , with high probability $V^\top \Omega$ is well conditioned [26]. This is true if the test matrix Ω satisfies the subspace embedding property, and it is said to preserve the geometry of an entire subspace of vectors V .

B. Subsampled Code Matrices, JLT and Subspace Embedding

Recall from section III the construction of the ‘tall and thin’ $n \times \ell$ subsampled error correcting code matrix Ω . The critical requirement to prove the $(1 + \epsilon)$ optimal error bound is to show that these matrices satisfy the two key properties: JLT and subspace embedding. The subspace embedding property will also imply that Ω_1 will be full rank, which will enable us use the deterministic bounds developed in the literature to derive the bounds for the spectral norm error and the singular values obtained (see sec. VI-C).

1) *Johnson-Lindenstrauss Transform*: We saw the definition of JLT in sec. II, which says that a matrix Ω that satisfies $\text{JLT}(\epsilon, \delta, d)$ preserves the norm for any vector v in a d -element subspace $V \subset \mathbb{R}^n$. We will use two key results developed in the literature to show that code matrices with certain mild properties satisfy the JLT property.

The first result is by Ailon and Liberty [2], where they show a matrix Ω which is 4-wise independent will satisfy the JLT property, see Lemma 5.1 in [2]. Interestingly, they give the 2 error correcting dual BCH codes as examples for such 4-wise independent matrices and also demonstrate how fast multiplications can be achieved with these code matrices. However, a minor drawback with using 4-wise independent matrices is that the maximum entries of A need to be restricted.

The second (stronger) result is by Clarkson and Woodruff [12] (see Theorem 2.2), where they show if Ω is a $4\lceil \log(\sqrt{(2)/\delta}) \rceil$ -wise independent matrix, then Ω will satisfy the JLT property. Recall that the SCM matrix Ω defined in eq. (4) will be a random k -wise independent matrix if the dual distance of the code is $> k$. Thus, any error correcting code matrix with a dual distance > 4 (more than 2 error correcting ability) will satisfy the JLT property.

One of the important results related to JLT that is of interest for our theoretical analysis is the matrix multiplication property. This is defined in the following lemma, which is Theorem 2.8 in [52]. We can see similar results in Lemma 6 in [47] and Theorem 2.2 in [12].

Lemma 1: For $\epsilon, \delta \in (0, 1/2)$, let Ω be a random matrix (or from a distribution \mathcal{D}) with n rows that satisfies (ϵ, δ, d) -JLT property. Then for A, B matrices with n rows,

$$\Pr [\|A^\top B - A^\top \Omega \Omega^\top B\|_F \leq 3\epsilon \|A\|_F \|A\|_F] \geq 1 - \delta. \quad (7)$$

We will see in sec. VI-D that the above lemma is one of the two main ingredients required to prove $(1 + \epsilon)$ optimal error bound. The other ingredient is the subspace embedding property.

2) *Subspace Embedding*: One of the primary results developed in the randomized sampling algorithms literature was establishing the relation between the Johnson-Lindenstrauss Transform (JLT) and subspace embedding. The following lemma which is corollary 11 in [47] gives this important relation.

Lemma 2: Let $0 < \epsilon, \delta < 1$ and f be some function. If $\Omega \in \mathbb{R}^{n \times \ell}$ satisfies a $\text{JLT}(\epsilon, \delta, d)$ with $\ell = O(k \log(k/\epsilon)/\epsilon^2 \cdot f(\delta))$, then for any orthonormal matrix $V \in \mathbb{R}^{n \times k}$, $n \geq k$ we have

$$\Pr (\|V^\top \Omega \Omega^\top V - I\|_2 \leq \epsilon) \geq 1 - \delta.$$

The above lemma shows that, any sampling matrix Ω satisfying JLT and having length $\ell = O(k \log(k/\epsilon)/\epsilon^2)$ satisfies the subspace embedding property. Thus, any SCM Ω with a dual distance > 4 will also satisfy the subspace embedding property (since they satisfy JLT as we saw in the previous section). The subspace embedding property implies that the singular values of $V^\top \Omega$ are bounded, i.e., $V^\top \Omega$ is well conditioned with high probability. This result is critical since it shows that the SCM matrices can preserve the geometry of the top k -singular vectors of the input matrix A .

Observe that with the above analysis, we will require $\ell = O(k \log(k/\epsilon))$ number of samples for the subspace embedding property to be satisfied, which is similar to a subsampled Fourier or Hadamard matrix. Next, we show that for the subspace embedding property to be satisfied, we will require only $O(k/\epsilon)$ number of samples for certain types of code matrices.

We know that the code matrices display some of the properties of random matrices, particularly when the distance of the code is high. Indeed a code with dual distance above k supports k -wise independent probability measure and SCM Ω will be a random matrix with k -wise independent rows. This property of SCM helps us use the following lemma given in [12, Lemma 3.4] which states,

Lemma 3: Given an integer k and $\epsilon, \delta > 0$. If $\Omega \in \mathbb{R}^{n \times \ell}$ is $\rho(k + \log(1/\delta))$ -wise independent matrix with an absolute constant $\rho > 1$, then for any orthonormal matrix $V \in \mathbb{R}^{n \times k}$ and $\ell = O(k \log(1/\delta)/\epsilon)$, with probability at least $1 - \delta$ we have

$$\|V^\top \Omega \Omega^\top V - I\|_2 \leq \epsilon.$$

Thus, a sampling SCM matrix Ω which is $\lceil k + \log(1/\delta) \rceil$ -wise independent satisfies the subspace embedding property with the number of samples (length) $\ell = O(k/\epsilon)$. Hence, an SCM Ω with dual distance $> \lceil k + \log(1/\delta) \rceil$ will preserve the geometry of V with $\ell = O(k/\epsilon)$.

In summary, any SCM with dual distance > 4 satisfies the JLT property, and will satisfy the subspace embedding property if $\ell = O(k \log(k/\epsilon))$. If the dual distance is $> k$, then the SCM can preserve the geometry of V with $\ell = O(k/\epsilon)$.

C. Deterministic Error bounds

In order to derive the bounds for the spectral norm error and the singular values obtained, we will use the deterministic error bounds for Algorithm 1 developed in the literature [26], [25]. Algorithm 1 constructs an orthonormal basis Q for the range of Y , and the goal is to quantify how well this basis captures the action of the input matrix A . Let $QQ^\top = P_Y$, where P_Y is the unique orthogonal projector with $\text{range}(P_Y) = \text{range}(Y)$. If Y is full rank, we can express the projector as $P_Y = Y(Y^\top Y)^{-1}Y^\top$. We seek to find an upper bound for the approximation error given by, for $\xi \in \{2, F\}$

$$\|A - QQ^\top A\|_\xi = \|(I - P_Y)A\|_\xi.$$

The deterministic upper bound for the approximation error of Algorithm 1 is given in [26]. We restate theorem 9.1 in [26] below:

Theorem 4 (Deterministic error bound): Let A be $m \times n$ matrix with singular value decomposition given by $A = U\Sigma V^\top$, and fixed $k \geq 0$. Choose a test matrix Ω and construct the sample matrix $Y = A\Omega$. Partition Σ as in (5), and define Ω_1 and Ω_2 via (6). Assuming that Ω_1 is full row rank, the approximation error satisfies for $\xi \in \{2, F\}$

$$\|(I - P_Y)A\|_\xi^2 \leq \|\Sigma_2\|_\xi^2 + \|\Sigma_2 \Omega_2 \Omega_1^\dagger\|_\xi^2. \quad (8)$$

An elaborate proof for the above theorem can be found in [26]. Using the submultiplicative property of the spectral and

Frobenius norms, and the Eckart-Young theorem mentioned earlier, equation (8) can be simplified to

$$\|A - QQ^\top A\|_\xi \leq \|A - A_k\|_\xi \sqrt{1 + \|\Omega_2\|_2^2 \|\Omega_1^\dagger\|_2^2}. \quad (9)$$

Recently, Ming Gu [25] developed deterministic lower bounds for the singular values obtained from randomized algorithms, particularly for the power method [26]. Given below is the modified version of Theorem 4.3 in [25] for Algorithm 1.

Theorem 5 (Deterministic singular value bounds): Let $A = U\Sigma V^\top$ be the SVD of A , for a fixed k , and let $V^\top \Omega$ be partitioned as in (6). Assuming that Ω_1 is full row rank, then Algorithm 1 must satisfy for $j = 1, \dots, k$:

$$\sigma_j \geq \sigma_j(\hat{A}_k) \geq \frac{\sigma_j}{\sqrt{1 + \|\Omega_2\|_2^2 \|\Omega_1^\dagger\|_2^2 \left(\frac{\sigma_{k+1}}{\sigma_j}\right)^2}} \quad (10)$$

where σ_j are the j th singular value of A and \hat{A}_k is the rank- k approximation obtained by our algorithm.

The proof for the above theorem can be seen in [25]. In both the above theorems, the key assumption is that Ω_1 is full row rank. This is indeed true if the sampling matrix Ω satisfies the subspace embedding property.

D. Error Bounds

The following theorem gives the approximation error bounds when the subsampled code matrix (SCM) is used as the sampling matrix Ω in Algorithm 1. The upper and lower bounds for the singular values obtained by the algorithm are also given.

Theorem 6 (Error bounds for code matrix): Let A be $m \times n$ matrix with singular values $\sigma_1 \geq \sigma_2 \geq \sigma_3 \geq \dots$. Generate a subsampled code matrix Ω from a desired coding scheme as in (4) with $r \geq \lceil \log_2(n) \rceil$ as the dimension of the code. For any code matrix Ω with **dual distance** > 4 and **length** $\ell = O(k \log(k/\epsilon)/\epsilon^2 \cdot f(\delta))$ the following three bounds hold with probability at least $1 - \delta$:

- 1) The Frobenius norm error satisfies,

$$\|A - \hat{A}_k\|_F \leq \|A - A_k\|_F (1 + \epsilon). \quad (11)$$

- 2) The spectral norm error satisfies,

$$\|A - \hat{A}_k\|_2 \leq \|A - A_k\|_2 \sqrt{1 + \frac{3n}{\ell}}. \quad (12)$$

- 3) The singular values obtained satisfy:

$$\sigma_j \geq \sigma_j(\hat{A}_k) \geq \frac{\sigma_j}{\sqrt{1 + \left(\frac{3n}{\ell}\right) \left(\frac{\sigma_{k+1}}{\sigma_j}\right)^2}}. \quad (13)$$

If the code matrix Ω has **dual distance** $\geq \lceil k + \log(1/\delta) \rceil$, then the above three bounds hold for **length** $\ell = O(k \log(1/\delta)/\epsilon)$.

Proof - Frobenius norm Error: As we have been alluding to in the previous sections, the $(1 + \epsilon)$ optimal Frobenius norm error given in eq. (11) is related to the JLT and the subspace embedding properties. The following lemma gives this relation which is Lemma 4.2 in Woodruff's monograph [52].

Lemma 7: Let Ω satisfy the subspace embedding property for any fixed k -dimensional subspace M with probability

9/10, so that $\|\Omega^\top y\|_2^2 = (1 \pm 1/3)\|y\|_2^2$ for all $y \in M$. Further, suppose Ω satisfies the $(\sqrt{\epsilon/k}, 9/10, k)$ -JLT property such that the conclusion in Lemma 1 holds, i.e., for any matrices A, B each with n rows,

$$\Pr \left[\|A^\top B - A^\top \Omega \Omega^\top B\|_F \leq 3\sqrt{\epsilon/k} \|A\|_F \|B\|_F \right] \geq 9/10.$$

Then the column space of $A\Omega$ contains a $(1 + \epsilon)$ rank- k approximation to A .

From the analysis in section VI-B (in particular from Lemma 1 and 2), we know that both the conditions in the above lemma are true for SCM with dual distance > 4 and length $\ell = O(k \log(k/\epsilon)/\epsilon^2 \cdot f(\delta))$, when appropriate ϵ and δ are chosen. Since $\hat{A}_k = QQ^\top A$, where Q is the orthonormal matrix spanning the column space of $A\Omega$, we obtain the Frobenius error bound in eq. (11) from the above lemma.

Clarkson and Woodruff [12] gave the Frobenius norm error bound for low rank approximation using k -wise independent sampling matrices. The error bound in (11) for SCM with dual distance $> k$ is straight from the following lemma which is a modification of Theorem 4.2 in [12].

Lemma 8: If $\Omega \in \mathbb{R}^{n \times \ell}$ is a $\rho(k + \log(1/\delta))$ -wise independent sampling matrix, then for $\ell = O(k \log(1/\delta)/\epsilon)$, with probability at least $1 - \delta$, we have

$$\|A - \hat{A}_k\|_F \leq \|A - A_k\|_F (1 + \epsilon). \quad (14)$$

Proof of this lemma is clear from the proof of Theorem 4.2 in [12]. ■

Proof - Spectral norm Error: The proof of the approximate error bounds given in (12) follows from the deterministic bounds given in sec. VI-C. We start from equation (9) in Theorem 4, the terms that depend on the choice of the test matrix Ω are $\|\Omega_2\|_2^2$ and $\|\Omega_1^\dagger\|_2^2$.

We know that the SCM Ω satisfies the subspace embedding property for the respective dual distances and lengths mentioned in the Theorem 6. This also ensures that the spectral norm of Ω_1^\dagger is under control. We have the condition $\|V_k^\top \Omega \Omega^\top V_k - I\|_2 \leq \epsilon_0$, implying

$$\sqrt{1 - \epsilon_0} \leq \sigma_k(V_k^\top \Omega) \leq \sigma_1(V_k^\top \Omega) \leq \sqrt{1 + \epsilon_0}.$$

Then from Lemma 3.6 in [53], we have

$$\|\Omega_1^\dagger\|_2^2 = \frac{1}{\sigma_k^2(\Omega_1)} \leq \frac{1}{(1 - \epsilon_0)}.$$

In Lemma 7, we chose $\epsilon_0 = 1/3$ to prove the $(1 + \epsilon)$ approximation. So, we have

$$\|\Omega_1^\dagger\|_2^2 \leq 3/2.$$

Next, we bound the spectral norm of Ω_2 as follows $\|\Omega_2\|_2^2 = \|V_2^\top \Omega\|_2^2 \leq \|V_2\|_2 \|\Omega\|_2^2 = \|\Omega\|_2^2 = \sigma_1^2(\Omega)$, since V_2 is an orthonormal matrix. So, we need an upper bound on the top singular value of SCM Ω , which we derive from the following two lemmas. The first lemma shows that if a code has dual distance ≥ 3 , the resulting code matrix Φ has orthonormal columns.

Lemma 9 (Code matrix with orthonormal columns): A code matrix Φ generated by a coding scheme which results in codes that have dual distance ≥ 3 , has orthonormal columns.

Proof: If a code has dual distance 3, then the corresponding code matrix (stacked up codewords as rows) is an orthogonal array of strength 2 [16]. This means all the tuples of bits, i.e., $\{0, 0\}, \{0, 1\}, \{1, 0\}, \{1, 1\}$, appear with equal frequencies in any two columns of the codeword matrix C . As a result, the Hamming distance between any two columns of C is exactly 2^{r-1} (half the length of the column). This means after the BPSK mapping, the inner product between any two codewords will be zero. It is easy to see that the columns are unit norm as well. ■

If there is no downsampling in Ω , then the singular values of Ω will simply be $\sqrt{n/\ell}$, due to the scaling in (4) of the orthonormal matrix and since $r = \log_2 n$. If we downsample the rows of Φ to form Ω , then the above fact helps us use Lemma 3.4 from [49] which shows that randomly sampling the rows of a matrix with orthonormal columns results in a well-conditioned matrix, and gives bounds for the singular values. The following lemma is a modification of Lemma 3.4 in [49].

Lemma 10 (Row sampling): Let Φ be a $2^r \times \ell$ code matrix with orthonormal columns and let

$$M = 2^r \cdot \max_{j=1, \dots, 2^r} \|e_j^\top \Phi\|_2^2.$$

For a positive parameter α , select the sample size

$$n \geq \alpha M \log(\ell).$$

Draw a random subset T from $\{1, \dots, 2^r\}$ by sampling n coordinates without replacement. Then

$$\sqrt{\frac{(1 - \nu)n}{2^r}} \leq \sigma_\ell(S_T \Phi) \text{ and } \sigma_1(S_T \Phi) \leq \sqrt{\frac{(1 + \eta)n}{2^r}} \quad (15)$$

with failure probability at most

$$\ell \cdot \left[\frac{e^{-\nu}}{(1 - \nu)^{(1 - \nu)}} \right]^{\alpha \log(\ell)} + \ell \cdot \left[\frac{e^\eta}{(1 + \eta)^{(1 + \eta)}} \right]^{\alpha \log(\ell)},$$

where $\nu \in [0, 1)$ and $\eta > 0$.

The bounds on the singular values of the above lemma are proved in [49] using the matrix Chernoff bounds.

Since n is fixed and $M = \ell$ for code matrices (all the entries of the matrix are $\pm 2^{-r/2}$), we get the condition $n \geq \alpha \ell \log(\ell)$. So, α is less than the ratio $n/\ell \log(\ell)$ and this ratio is typically more than 10 in the low rank approximation applications. For $\alpha = 10$, we choose $\nu = 0.6$ and $\eta = 1$, then the failure probability is at most $2\ell^{-1}$. Since we use the scaling $\sqrt{\frac{2^r}{\ell}}$, the bounds on the singular values of the subsampled code matrix Ω will be

$$\sqrt{\frac{2n}{5\ell}} \leq \sigma_\ell(\Omega) \text{ and } \sigma_1(\Omega) \leq \sqrt{\frac{2n}{\ell}}. \quad (16)$$

Thus, we obtain $\|\Omega_2\|_2^2 \|\Omega_1^\dagger\|_2^2 = 3n/\ell$. We substitute this value in (9) to get the spectral norm error bounds in (12). ■

Similarly, we obtain the bounds on the singular values given in (13) by substituting the above value of $\|\Omega_2\|_2^2 \|\Omega_1^\dagger\|_2^2$ in (10) of Theorem 5.

We observe that the upper bounds for the spectral norm error obtained in (12) for the SCM is similar to the bounds obtained for Gaussian random matrices and structured random

matrices like SRFT/SRHT given in the review article by Halko et.al [26]. For the structured random matrices, $(1 + \epsilon)$ optimal Frobenius norm error has been derived in [41] and [10]. We have a similar $(1 + \epsilon)$ optimal Frobenius norm error obtained for subsampled code matrices with dual distance > 4 in (11). Importantly, we show that this optimal error bound can be achieved with number of samples $\ell = O(k/\epsilon)$ as opposed to $O(k \log k/\epsilon)$ required for structured random matrices when the dual distance of the code is $> k$. Details on how to generate such code matrices with dual distance $> k$ and length $\ell = O(k/\epsilon)$ is given in section VIII.

E. Differences in the construction

An important difference between the construction of subsampled code matrices SCM given in (4) and the construction of SRHT or SRFT given in [26], [49] is in the way these matrices are subsampled. In the case of SRHT, a Hadamard matrix of size $n \times n$ is applied to input matrix A and ℓ out of n columns are sampled at random (n must be a power of 2). When the input matrix is distributively stored, this procedure introduces communication issues. The subsampling will require additional communication since each of the nodes must sample the same columns (also recall that we need to sample $O(k \log k)$ columns), making the fast transform slow. Similar issues arise when the input matrices are streamed.

In contrast, in the case of SCM, a $2^r \times \ell$ code matrix generated from an $[\ell, r]$ -linear coding scheme is considered, and n out of 2^r codewords are chosen (if $r > \log_2 n$). When the input matrix is distributively stored, different rows/columns of the matrix can be sampled by different codewords locally and hence communicating only the sampled rows/columns. Similarly, when the input matrix is streamed, at a given time instant, the newly arrived rows of the matrix can simply be sampled by new codewords of the code, requiring minimal storage and communication. For details on space required and communication complexity for sketching streaming matrices using random sign matrices, see [12]. The subsampling will not affect the k -wise independent property of the code matrix or the distinctness of rows when uniformly subsampled. This need not be true in the case of SRHT. The importance of the distinctness of rows is discussed next.

F. Logarithmic factor

A crucial advantage of the code matrices is that they have very low *coherence*. Coherence is defined as the maximum inner product between any two rows. This is in particular true when the minimum distance of the code is close to half the length. If the minimum distance of the code is d then the code matrix generated from an $[\ell, r]$ -code has coherence equal to $\frac{\ell - 2d}{2^r}$. For example, if we consider dual BCH code (see sec. II-B) the coherence is $\frac{2(t-1)\sqrt{\ell+1}-1}{2^r}$. Low coherence ensures near orthogonality of rows. This is a desirable property in many applications such as compressed sensing and sparse recovery.

For a rank- k approximation using subsampled Fourier or Hadamard matrices, we need to sample $O(k \log k)$ columns. This logarithmic factor emerges as a necessary condition in the

theoretical proof (given in [49]) that shows that these matrices approximately preserve the geometry of an entire subspace of input vectors (satisfy the subspace embedding property). The log factor is also necessary to handle the worst case input matrices. The discussions in sec. 11 of [26] and sec. 3.3 of [49] give more details. In the case of certain subsampled code matrices, the log factor is not necessary to tackle these worst case input matrices. To see why this is true, let us consider the worst case example for orthonormal matrix V described in Remark 11.2 of [26].

An infinite family of worst case examples of the matrix V is as follows. For a fixed integer k , let $n = k^2$. Form an $n \times k$ orthonormal matrix V by regular decimation of the $n \times n$ identity matrix. That is, V is a matrix whose j th row has a unit entry in column $(j-1)/k$ when $j \equiv 1 \pmod{k}$ and is zero otherwise. This type of matrix is troublesome when DFT or Hadamard matrices are used for sampling.

Suppose that we apply $\Omega = DFR^\top$ to the matrix V^\top , where D is same as in (4), F is an $n \times n$ DFT or Hadamard matrix and R is $\ell \times n$ matrix that samples ℓ coordinates from n uniformly at random. We obtain a matrix $X = V^\top \Omega = WR^\top$, which consists of ℓ random columns sampled from $W = V^\top DF$. Up to scaling and modulation of columns, W consists of k copies of a $k \times k$ DFT or Hadamard matrix concatenated horizontally. To ensure that X is well conditioned, we need $\sigma_k(X) > 0$. That is, we must pick at least one copy of each of the k distinct columns of W . This is the coupon collector's problem [39] in disguise and to obtain a complete set of k columns with non-negligible probability, we must draw at least $k \log(k)$ columns.

In the case of code matrices, we apply a subsampled code matrix $\Omega = DS\Phi$ to the matrix V^\top . We obtain $X = V^\top \Omega = V^\top DS\Phi$, which consists of k randomly selected rows of the code matrix Φ . That is, X consists of k distinct codewords of length ℓ . The code matrix has low coherence and all rows are distinct. If we use a code matrix with dual distance $> k$, then X contains k rows which are k -wise independent (near orthonormal) and $\sigma_k(X) > 0$; as a result the geometry of V is preserved and the log factor is not necessary. Thus, for the worst case scenarios we have an $O(\log k)$ factor improvement over other structured matrices. More importantly, this shows that the order optimal can be achieved with the immediate lower bound of $O(k)$ in the number of samples required for the sampling matrices constructed from deterministic matrices.

VII. LEAST SQUARES REGRESSION PROBLEM

In this section, we extend the framework to solve the least squares (ℓ_2) regression problem. As discussed in the introduction, the idea of randomized approximations is to reduce the dimensions of $A \in \mathbb{R}^{n \times d}$ and $b \in \mathbb{R}^n$ with $n \gg d$, by pre-multiplying them by a sampling matrix $\Omega \in \mathbb{R}^{n \times \ell}$, and then to solve the smaller problem quickly,

$$\min_x \|\Omega^\top Ax - \Omega^\top b\|_2. \quad (17)$$

Let the optimal solution be $x' = (\Omega^\top A)^\dagger \Omega^\top b$. Here we analyze the performance of SCM as the sampling matrix Ω . We require the sampling matrix Ω to satisfy the JLT and the

subspace embedding properties, which are indeed satisfied by any SCM with dual distance > 4 . Hence, we can use the results developed by Sarlos [47], and Clarkson and Woodruff [12] for our analysis.

We know that, any code matrix with dual distance > 4 satisfies the JLT property from our analysis in section VI-B. Sarlos in [47] derived the relation between the JLT matrices and the sampling matrices in the ℓ_2 regression problem (17). The following theorem is a modification of theorem 12 in [47].

Theorem 11: Suppose $A \in \mathbb{R}^{n \times d}, b \in \mathbb{R}^n$. Let $\mathcal{Z} = \min_x \|Ax - b\|_2 = \|A\hat{x} - b\|_2$, where $\hat{x} = A^\dagger b$ is the minimizer. Let $0 < \epsilon, \delta < 1$ and $\Omega \in \mathbb{R}^{n \times \ell}$ be a random matrix satisfying JLT and $\tilde{\mathcal{Z}} = \min_x \|\Omega^\top (Ax - b)\|_2 = \|\Omega^\top (Ax' - b)\|_2$, where $x' = (\Omega^\top A)^\dagger \Omega^\top b$. Then, with probability at least $1 - \delta$, we have

- If $\ell = O(\log(1/\delta)/\epsilon^2)$,

$$\tilde{\mathcal{Z}} \leq (1 + \epsilon)\mathcal{Z}. \quad (18)$$

- If $\ell = O(d \log d \cdot \log(1/\delta)/\epsilon)$,

$$\|Ax' - b\|_2 \leq (1 + \epsilon)\mathcal{Z}. \quad (19)$$

- If $\ell = O(d \log d \cdot \log(1/\delta)/\epsilon^2)$,

$$\|\hat{x} - x'\|_2 \leq \frac{\epsilon}{\sigma_{\min}(A)} \mathcal{Z}. \quad (20)$$

The proof for this theorem can be seen in [47].

If $\sqrt{\|b\|_2^2 - \mathcal{Z}^2} \geq \gamma \|b\|_2$ for some $0 < \gamma \leq 1$, then we can replace the last equation (20) by

$$\|\hat{x} - x'\|_2 \leq \epsilon \left(\kappa(A) \sqrt{\gamma^{-2} - 1} \right) \|\hat{x}\|_2, \quad (21)$$

where $\kappa(A)$ is the 2 norm condition number of A . (This equation is given to be consistent with the results given in the related literature [20], [21], [54].) Thus, any code matrix with dual distance > 4 can be used as the sampling matrix in the least squares regression problem. Again, the performance of such code matrices is very similar to that of structured random matrices (SRHT) given in [21], [10]. Fast multiplication can be used to sample dense input matrix A .

Similar to the earlier analysis, we can expect improved performance when SCM with dual distance $> k$ are used (k -wise independence property of codes). For this, we use the bounds derived by Clarkson and Woodruff [12] for random sign matrices. The following theorem which is a modification of Theorem 3.1 in [12] gives the upper bound for the regression problem in such cases.

Theorem 12: Given $\epsilon, \delta > 0$, suppose $A \in \mathbb{R}^{n \times d}, b \in \mathbb{R}^n$ and A has rank at most k . If Ω is a $\rho(k + \log(1/\delta))$ -wise independent with an absolute constant $\rho > 1$, and x' and \hat{x} are solutions as defined before, then for $\ell = O(k \log(1/\delta)/\epsilon)$, with probability at least $1 - \delta$, we have

$$\|Ax' - b\|_2 \leq (1 + \epsilon) \|A\hat{x} - b\|_2.$$

This theorem shows that, if the code matrix is $\lceil k + \log(1/\delta) \rceil$ -wise independent (i.e., dual distance $> \lceil k + \log(1/\delta) \rceil$), we can get ϵ -approximate solution for the regression problem with $\ell = O(k \log(1/\delta)/\epsilon)$ samples. Thus, for the regression

problem too, we have the $\log k$ factor gain in the number of samples over other structured random matrices (SRHT) given in [21], [10].

Typically, A is full rank. So, we will need a code matrix with dual distance $> d$ and length of the code $\ell = O(d \log(1/\delta)/\epsilon)$. Article [54] discusses the applications where such over-determined system of equations are encountered. In typical applications, n will be in the range of $10^6 - 10^9$ and d in the range of $10^2 - 10^3$ (details in [54]). In the next section, we discuss how to generate such code matrices with dual distance $> d$ and minimum length ℓ , and discuss the choice of the error correcting codes for different type of input matrices and computational environments.

VIII. CHOICE OF ERROR CORRECTING CODES

A. Codes with dual-distance at least $k + 1$

The requirement of k -wise independence of codewords translates to the dual distance of the code being greater than k . Since a smaller code (less number of codewords, i.e., smaller r) leads to less randomness in sampling, we would like to use the smallest code with dual distance greater than k .

One of the choices of the code can be the family of dual BCH codes. As mentioned earlier, this family has length ℓ , dimension $t \log(\ell + 1)$ and dual distance at least $2t + 1$. Hence, to guarantee dual distance at least k , the size of the code must be $2^{\frac{k \log(\ell+1)}{2}} = (\ell + 1)^{k/2}$. We can choose n vectors of length $\frac{k \log(\ell+1)}{2}$ and form the codewords by simply multiplying these with the generator matrix (over \mathbb{F}_2) to form the subsampled code matrix. Therefore, forming these code matrices will be much faster than generating $n \times \ell$ i.i.d Gaussian random matrices or random sign matrices which have k -wise independent rows.

In general, from the Gilbert-Varshamov bound of coding theory [34], it is known that linear codes of size $\sim \sum_{i=0}^k \binom{\ell}{i}$ exist that have length ℓ and dual distance greater than k . The construction of these code families are still randomized. However, when $k = O(\ell)$, or the dual distance is linearly growing with the code length, the above construction of dual BCH code does not hold in general. Infinite families of codes that have distance proportional to the length are called *asymptotically good codes*. The Gilbert-Varshamov bound implies that asymptotically good linear codes of size $\sim 2^{\ell h(\frac{k}{\ell})}$ exist³, that have length ℓ and dual distance greater than k .

B. Choice of the code matrices

Depending on the types of input matrices and the computational environments, we can choose different types of code matrices that best suit the applications. If the input matrix is a general dense matrix which can be stored in the fast memory (RAM), we can choose any structured code matrix with dual distance > 4 , $r = \lceil \log_2 n \rceil$ (or choose message vectors to be 1 to $2^{\lceil \log_2 n \rceil}$) and $\ell = O(k \log k)$ (eg., dual BCH codes), so that the fast multiplication technique can be exploited (the log factor will not be an issue). This will be similar to using any other structured random matrices like

³ $h(x) \equiv -x \log_2 x - (1-x) \log_2 (1-x)$ is the binary entropy function

SRFT or SRHT. In fact, Hadamard matrices are also a class of linear codes, with variants known as Hadamard codes, Simplex codes or 1st-order Reed-Muller codes. The dual distance of Hadamard code is 3. However, with code matrices (say dual BCH codes), subsampling of columns is not required, thus reducing randomness and cost.

If the input matrix is sparse and/or is distributively stored, and for parallel implementation, we can choose a code matrix with dual distance $> k$ and generate them as mentioned earlier and as in section III. These code matrices are not structured and we can treat them as dense transforms (any random matrices), a method to sample such distributively stored matrices was described in sec. VI-E. For SRFT/SRHT sampling matrices, we need to communicate $O(k \log k)$ columns, but for code matrices with dual distance $> k$, the log factor is not necessary. This will help us overcome the issues with SRFT/SRHT for sparse input matrices and in parallel and distributed applications. These code matrices are easy to generate (than i.i.d Gaussian random matrices), the log factor in the number of samples is not necessary, and thus, using code matrices in these applications will reduce randomness and cost significantly.

A strategy to sample streaming data was also described in sec. VI-E that requires minimal storage and communication. For details on the cost, space required and communication complexity for sketching streaming matrices using random sign matrices, we refer [12] (observe that the SCM matrices are equivalent to random sign matrices without the scaling $1/\sqrt{\ell}$). If the log factor is not an issue (for smaller k), then we can choose any code matrix with dual distance > 4 and $r = \lceil \log_2 n \rceil$, and form $Y = A\Omega$ as a dense transform. These code matrices are almost deterministic and unlike SRFT/SRHT, subsampling of columns is not required.

In practice, code matrices generated by any linear coding scheme can be used in place of Gaussian random matrices. As there are many available classes of algebraic and combinatorial codes, we have a large pool of candidate matrices. In this paper we chose dual BCH codes for our numerical experiments as they particularly have low coherence, and turn out to perform quite well in practice.

IX. NUMERICAL EXPERIMENTS

The following experiments will illustrate the performance of subsampled code matrices as sampling matrices in algorithm 1. We compare the performance of dual BCH code matrices against the performance of random Gaussian matrices and subsampled Fourier transform (SRFT) matrices for different input matrices from various applications.

Our first experiment is with a 4770×4770 matrix named Kohonen from the Pajek network (a directed graph's matrix representation), available from the UFL Sparse Matrix Collection [15]. Such graph Laplacian matrices are commonly encountered in machine learning and image processing applications. The performance of the dual BCH code matrix, Gaussian matrix, subsampled Fourier transform (SRFT) and Hadamard (SRHT) matrices are compared as sampling matrices Ω in algorithm 1. For SRHT, we have to subsample the

TABLE I
COMPARISON OF ERRORS

| MATRIX | ℓ | DUAL BCH | GAUSSIAN | SRFT |
|------------|--------|----------|----------|---------|
| LPICERIA3D | 31 | 21.8779 | 23.7234 | 23.3688 |
| S80PI | 63 | 3.8148 | 3.8492 | 3.7975 |
| DELAUNAY | 63 | 6.3864 | 6.3988 | 6.3829 |
| LPICERIA3D | 63 | 15.4865 | 18.3882 | 16.3619 |
| DETER3 | 127 | 9.2602 | 9.2658 | 9.2984 |
| EPA | 255 | 5.5518 | 5.5872 | 5.4096 |
| KOHONEN | 511 | 4.2977 | 4.2934 | 4.2610 |

rows as well (similar to code matrices) since the input size is not a power of 2. All experiments were implemented in matlab v8.1.

Figure 1(A) gives the actual error $e_\ell = \|A - Q^{(\ell)}(Q^{(\ell)})^\top A\|$ for each ℓ number of samples when a subsampled dual BCH code matrix, a Gaussian matrix, SRFT and SRHT matrices are used as sampling matrices in algorithm 1, respectively. The best rank- ℓ approximation error $\sigma_{\ell+1}$ is also given. Figure 1(B) plots the singular values obtained from algorithm 1, for $\ell = 255$ and different sampling matrices Ω used. The top 255 exact singular values of the matrix (available in the UFL database) are also plotted. We observe that, in practice, the performance of all four sampling matrices are similar.

Table I compares the errors e_ℓ for ℓ number of samples, obtained for a variety of input matrices from different applications when subsampled dual BCH code, Gaussian and SRFT matrices were used. All matrices were obtained from the UFL database [15]. Matrices lpi_ceria3d (4400×3576) and deter3 (21777×7647) are from linear programming problems. S80PI_n1 (4028×4028) is from an eigenvalue/model reduction problem. Delaunay (4096×4096), EPA (4772×4772) and Kohonen are graph Laplacian matrices. We observe that, for small ℓ , in the first five examples the error performance of code matrices is slightly better than that of Gaussian matrices. For higher ℓ , the error remains similar to the error for Gaussian matrices. Therefore, in practice, we can use code matrices in place of fully random (Gaussian) matrices or structured random matrices due to the advantages of code matrices over the other sampling matrices, as discussed in the previous sections. Next, we illustrate the performance of algorithm 1 with different sampling matrices in a practical application.

Eigenfaces: Eigenfaces is a popular method for face recognition that is based on Principal Component Analysis (PCA) [50], [48]. In this experiment (chosen as a verifiable comparison with results in [25]), we demonstrate the performance of randomized algorithm with different sampling matrices on face recognition. The face dataset is obtained from the AT&T Labs Cambridge database of faces [46]. There are ten different images of each of 40 distinct subjects. The size of each image is 92×112 pixels, with 256 gray levels per pixel. 200 of these faces, 5 from each individual are used as training images and the remaining 200 as test images to classify.

In the first step, we compute the principal components (dimensionality reduction) of mean shifted training image dataset using algorithm 1, with different sampling matrix Ω and different $p = \ell - k$ values (oversampling used). Next, we project the mean-shifted images into the singular vector

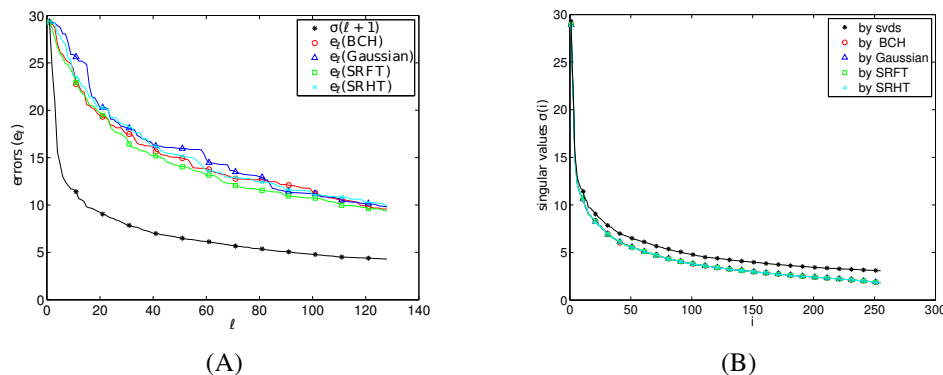


Fig. 1. (A) The theoretical minimum $\sigma_{\ell+1}$ and approximate error as a function of the number of random samples ℓ using dual BCH code, Gaussian, SRFT and SRHT matrices as sampling matrix in Algorithm 1 for input matrix Kohonen. (B) Estimates for top 255 singular values computed by Algorithm 1 using dual BCH code, Gaussian, SRFT and SRHT matrices and the exact singular values by svds function.

TABLE II
COMPARISON OF THE NUMBER OF INCORRECT MATCHES

| RANK k | DUAL BCH | | GAUSSIAN | | SRFT | | T-SVD |
|-------------|----------|------|----------|------|------|------|-------|
| | p | 20 | p | 20 | p | 20 | |
| 10 | 18 | 13 | 19 | 15 | 21 | 18 | 26 |
| 20 | 14 | 11 | 14 | 12 | 16 | 12 | 13 |
| 30 | 10 | 08 | 13 | 08 | 12 | 09 | 10 |
| 40 | 09 | 08 | 08 | 07 | 08 | 10 | 06 |

space using the singular vectors obtained from the first step. The projections are called feature vectors and are used to train the classifier. To classify a new face, we mean-shift the image and project it onto the singular vector space obtained in the first step, obtaining a new feature vector. The new feature vector is classified using a classifier which is trained on the feature vectors from the training images. We used the in-built MATLAB function `classify` for feature training and classification. We compare the performance of the dual BCH code matrix, Gaussian matrix and SRFT matrix against exact truncated SVD (T-SVD). The results are summarized in Table II. For $p = 10$ dual BCH code matrices give results that are similar to those of truncated SVD, and for rank $k < 40$, $p = 20$ our results are superior.

X. CONCLUSION

This paper advocated the use of matrices generated by error correcting codes as an alternative to random Gaussian or subsampled Fourier/Hadamard matrices for computing low rank matrix approximations. Among the attractive properties of the proposed approach are the numerous choices of parameters available, ease of generation, reduced randomness and cost, and the near-orthogonality of rows. We showed that any code matrix with dual distance > 4 preserves the geometry of an entire subspace of vectors. Indeed if the dual distance of the code matrix is $> k$, then the length of the code (sampling complexity) required is in $O(k)$, thus leading to an order optimal in the worst-case guaranteed sampling complexity, an improvement by a factor of $O(\log k)$ over other known deterministic matrices. We saw that fast multiplication is possible with structured code matrices, resulting in fast approximations for general dense input matrices. The

implementation issues of FFT-like structured random matrices in the parallel and distributed environments can be overcome by using code matrices as sampling matrices.

It is known that Gaussian matrices perform much better in practice compared to their theoretical analysis [26]. Our code matrices (a) are almost deterministic, and (b) have ± 1 entries. Still, they perform equally well (as illustrated by experiments) compared to random real Gaussian matrices and complex Fourier matrices. Because of the availability of different families of classical codes in the rich literature of coding theory, many possible choices of code matrices are at hand. One of the contributions of this paper is to open up these options for use as structured sampling operators in low-rank approximations and least squares regression problem.

Interesting future works include improving the approximation error bounds obtained in this paper, and extending the framework of code matrices to other similar applications. The connections between code matrices and JLT and random sign matrices might lead to improved analysis in other applications of codes such as sparse recovery [4].

REFERENCES

- [1] N. Ailon and B. Chazelle. Approximate nearest neighbors and the fast johnson-lindenstrauss transform. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 557–563. ACM, 2006.
- [2] N. Ailon and E. Liberty. Fast dimension reduction using rademacher series on dual bch codes. *Discrete & Computational Geometry*, 42(4):615–630, 2009.
- [3] O. Alter, P. O. Brown, and D. Botstein. Singular value decomposition for genome-wide expression data processing and modeling. *Proceedings of the National Academy of Sciences*, 97(18):10101–10106, 2000.
- [4] A. Barg, A. Mazumdar, and R. Wang. Restricted isometry property of random subdictionaries. *IEEE Transactions on Information Theory*, 61(8), 2015.
- [5] R. Berinde, P. Indyk, and M. Ruzic. Practical near-optimal sparse recovery in the l_1 norm. In *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*, pages 198–205. IEEE, 2008.
- [6] M. Berry, Z. Drmac, and E. Jessup. Matrices, vector spaces, and information retrieval. *SIAM Review*, 41(2):335–362, 1999.
- [7] M. W. Berry, S. T. Dumais, and G. W. O’Brien. Using linear algebra for intelligent information retrieval. *SIAM review*, 37(4):573–595, 1995.
- [8] R. E. Blahut. Transform techniques for error control codes. *IBM Journal of Research and development*, 23(3):299–315, 1979.
- [9] R. C. Bose and D. K. Ray-Chaudhuri. On a class of error correcting binary group codes. *Information and control*, 3(1):68–79, 1960.
- [10] C. Boutsidis and A. Gittens. Improved matrix algorithms via the subsampled randomized hadamard transform. *SIAM Journal on Matrix Analysis and Applications*, 34(3):1301–1340, 2013.

- [11] C. Boutsidis, M. W. Mahoney, and P. Drineas. An improved approximation algorithm for the column subset selection problem. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 968–977. Society for Industrial and Applied Mathematics, 2009.
- [12] K. L. Clarkson and D. P. Woodruff. Numerical linear algebra in the streaming model. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 205–214. ACM, 2009.
- [13] K. L. Clarkson and D. P. Woodruff. Low rank approximation and regression in input sparsity time. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 81–90. ACM, 2013.
- [14] T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- [15] T. A. Davis and Y. Hu. The University of Florida sparse matrix collection. *ACM Transactions on Mathematical Software (TOMS)*, 38(1):1, 2011.
- [16] P. Delsarte and V. I. Levenshtein. Association schemes and coding theory. *Information Theory, IEEE Transactions on*, 44(6):2477–2504, 1998.
- [17] A. Deshpande and S. Vempala. Adaptive sampling and fast low-rank matrix approximation. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 292–303. Springer, 2006.
- [18] P. Drineas, R. Kannan, and M. 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.
- [19] P. Drineas, I. Kerenidis, and P. Raghavan. Competitive recommendation systems. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 82–90. ACM, 2002.
- [20] P. Drineas, M. W. Mahoney, and S. Muthukrishnan. Sampling algorithms for l_2 regression and applications. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1127–1136. Society for Industrial and Applied Mathematics, 2006.
- [21] P. Drineas, M. W. Mahoney, S. Muthukrishnan, and T. Sarlós. Faster least squares approximation. *Numerische Mathematik*, 117(2):219–249, 2011.
- [22] D.-Z. Du and F. K. Hwang. Combinatorial group testing and its applications (applied mathematics). 2000.
- [23] C. Eckart and G. Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.
- [24] G. H. Golub and C. F. V. Loan. *Matrix Computations, 4th edition*. Johns Hopkins University Press, Baltimore, MD, 4th edition, 2013.
- [25] M. Gu. Subspace iteration randomization and singular value problems. *SIAM Journal on Scientific Computing*, 37(3):A1139–A1173, 2015.
- [26] N. Halko, P. Martinsson, and J. Tropp. Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions. *SIAM Review*, 53(2):217–288, 2011.
- [27] S. Har-Peled. Low rank matrix approximation in linear time. *arXiv preprint arXiv:1410.8802*, 2014.
- [28] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
- [29] W. B. Johnson and J. Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.
- [30] I. Jolliffe. *Principal component analysis*. Wiley Online Library.
- [31] J. M. Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)*, 46(5):604–632, 1999.
- [32] E. Liberty. *Accelerated Dense Random Projections*. PhD thesis, Yale University, 2009.
- [33] E. Liberty, F. Woolfe, P.-G. Martinsson, V. Rokhlin, and M. Tygert. Randomized algorithms for the low-rank approximation of matrices. *Proceedings of the National Academy of Sciences*, 104(51):20167–20172, 2007.
- [34] F. J. MacWilliams and N. J. A. Sloane. *The theory of error-correcting codes*, volume 16. Elsevier, 1977.
- [35] M. W. Mahoney. Randomized algorithms for matrices and data. *Foundations and Trends® in Machine Learning*, 3(2):123–224, 2011.
- [36] I. Markovsky. *Low Rank Approximation: Algorithms, Implementation, Applications*. Springer, 2012.
- [37] A. Mazumdar. *Combinatorial methods in coding theory*. PhD thesis, University of Maryland, 2011.
- [38] A. Mazumdar. On almost disjoint matrices for group testing. In *Algorithms and Computation*, pages 649–658. Springer, 2012.
- [39] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge International Series on Parallel Computation. Cambridge University Press, 1995.
- [40] C. Musco and C. Musco. Stronger approximate singular value decomposition via the block lanczos and power methods. *arXiv preprint arXiv:1504.05477*, 2015.
- [41] N. H. Nguyen, T. T. Do, and T. D. Tran. A fast and efficient algorithm for low-rank approximation of a matrix. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 215–224. ACM, 2009.
- [42] P. Parker, P. J. Wolfe, and V. Tarokh. A signal processing application of randomized low-rank approximations. In *Statistical Signal Processing, 2005 IEEE/SP 13th Workshop on*, pages 345–350. IEEE, 2005.
- [43] S. Raychaudhuri, J. M. Stuart, and R. B. Altman. Principal components analysis to summarize microarray experiments: application to sporulation time series. In *Pacific Symposium on Biocomputing. Pacific Symposium on Biocomputing*, page 455. NIH Public Access, 2000.
- [44] 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.
- [45] V. Rokhlin and M. Tygert. A fast randomized algorithm for overdetermined linear least-squares regression. *Proceedings of the National Academy of Sciences*, 105(36):13212–13217, 2008.
- [46] F. S. Samaria and A. C. Harter. Parameterisation of a stochastic model for human face identification. In *Applications of Computer Vision, 1994., Proceedings of the Second IEEE Workshop on*, pages 138–142. IEEE, 1994.
- [47] T. Sarlos. Improved approximation algorithms for large matrices via random projections. In *Foundations of Computer Science, 2006. FOCS’06. 47th Annual IEEE Symposium on*, pages 143–152. IEEE, 2006.
- [48] L. Sirovich and M. Meytlis. Symmetry, probability, and recognition in face space. *Proceedings of the National Academy of Sciences*, 106(17):6895–6899, 2009.
- [49] J. A. Tropp. Improved analysis of the subsampled randomized hadamard transform. *Advances in Adaptive Data Analysis*, 3(01n02):115–126, 2011.
- [50] M. Turk and A. Pentland. Eigenfaces for recognition. *Journal of cognitive neuroscience*, 3(1):71–86, 1991.
- [51] S. Ubaru, A. Mazumdar, and Y. Saad. Low rank approximation using error correcting coding matrices. In *Proceedings of The 32nd International Conference on Machine Learning*, pages 702–710, 2015.
- [52] D. P. Woodruff. Sketching as a tool for numerical linear algebra. *Theoretical Computer Science*, 10(1-2):1–157, 2014.
- [53] F. Woolfe, E. Liberty, V. Rokhlin, and M. Tygert. A fast randomized algorithm for the approximation of matrices. *Applied and Computational Harmonic Analysis*, 25(3):335–366, 2008.
- [54] J. Yang, X. Meng, and M. W. Mahoney. Implementing randomized matrix algorithms in parallel and distributed environments. *arXiv preprint arXiv:1502.03032*, 2015.
- [55] J. Ye. Generalized Low Rank Approximations of Matrices. *Machine Learning*, 61(1-3):167–191, 2005.