

Exponential Inapproximability of Selecting a Maximum Volume Sub-matrix

Ali Çivril · Malik Magdon-Ismail

the date of receipt and acceptance should be inserted later

Abstract Given a matrix $A \in \mathbb{R}^{m \times n}$ (n vectors in m dimensions), and a positive integer $k < n$, we consider the problem of selecting k column vectors from A such that the volume of the parallelepiped they define is maximum over all possible choices. We prove that there exists $\delta < 1$ and $c > 0$ such that this problem is not approximable within 2^{-ck} for $k = \delta n$, unless $P = NP$.

Keywords matrices · volume · complexity · inapproximability

1 Introduction

Given a matrix $A \in \mathbb{R}^{m \times n}$, it is of practical importance to obtain the “significant information” contained in A . It becomes especially important to have a compact representation of A when A is large and has low numerical rank, as is typical of modern data. Thus, in a broad sense, we are interested in concise representations of matrices. Besides the tremendous practical impact of linear algebraic algorithms designed to this aim, they also come up in different theoretical forms and paradigms. Specifically, the formalization of “significant information” can be done in several ways and to a great extent, it depends on how a matrix is interpreted.

From a conceptual point of view, rather than interpreting a matrix as a block of numbers, we view it as a set of vectors (specifically, column vectors) which are indivisible entities. Thus, the formalization of “significant information” is essentially related to finding a subset of columns of the matrix which satisfies some certain spectral conditions or orthogonality requirements. From a purely combinatorial perspective, treating vectors as elements of a set, one

Ali Çivril

Melikşah University, Computer Engineering Department, Talas, Kayseri 38280 Turkey

Malik Magdon-Ismail

Rensselaer Polytechnic Institute, Computer Science Department, 110 8th Street, Troy, NY 12180 USA

can also view subset selection in matrices as a generalization of the usual subset selection problem where the elements contain little or no information. To give a specific example, the well known Set Cover problem asks for a smallest cardinality subset of a set system which covers a universal set. Likewise, the problem we are interested in essentially asks for a small number of column vectors to “cover” the whole matrix. In this paper, we state a measure of quality for this problem, namely the volume, and we prove an exponential inapproximability result for the problem of selecting a maximum volume sub-matrix of a matrix.

Several problems in matrix analysis require to construct a more concise version of a matrix generally performed by a re-ordering of the columns [15], such that the new smaller matrix is as good a representative of the original as possible. One of the criteria that defines the quality of a subset of columns of a matrix is how well-conditioned the sub-matrix that they define is. To motivate the discussion, consider the set of three vectors

$$\left\{ e_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, e_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, u = \begin{pmatrix} \sqrt{1-\epsilon^2} \\ \epsilon \end{pmatrix} \right\},$$

which are clearly dependent, and any two of which are a basis. Thus any pair can serve to reconstruct all vectors. Suppose we choose e_1, u as the basis, then $e_2 = (1/\epsilon)u - (\sqrt{1-\epsilon^2}/\epsilon)e_1$, and we have a numerical instability in this representation as $\epsilon \rightarrow 0$. Such problems get more severe as the dimensionality of the space gets large (curse of dimensionality), and it is natural to ask the representatives to be “as far away from each other as possible”. From this simple example, we see that two orthogonal vectors will capture more information about a superset of columns than two that have an acute angle between each other. Hence, in its generality, this vaguely stated problem can be stated as finding a subset of columns with the maximum volume possible or equivalently with the maximum determinant. A similar (but not equivalent) problem is to find a subset with the maximum smallest singular value. Indeed, in one of the early works studying Rank Revealing QR (RRQR) factorizations [21], while discussing different options on how to choose a good sub-matrix, it was noted that it turns out that “the selection of the sub-matrix with the maximum smallest singular value suggested in [14] can be replaced by the selection of a sub-matrix with maximum determinant”, which heuristically proposes to maximize the volume of the sub-matrix instead of using more complicated functions. Several algorithms have been designed following this intuition [3, 5–7, 9, 19, 21, 25]. The optimization problem of finding a maximum volume sub-matrix of a matrix was only recently studied by Çivril and Magdon-Ismail:

Definition 1 [8] *Given a matrix $A \in \mathbb{R}^{m \times n}$ of rank at least k , MAX-VOL is the problem of finding a sub-matrix $C \in \mathbb{R}^{m \times k}$ of A such that the volume of the k dimensional parallelepiped defined by the column vectors in C is maximum over all possible choices.*

Theorem 2 [8] *MAX-VOL is NP-hard. Further, it is NP-hard to approximate to within $2\sqrt{2}/3 + \epsilon$ for arbitrarily small $\epsilon > 0$.*

Since MAX-VOL is NP-hard, it is natural to ask for an algorithm to approximate the maximum volume. The first thing one might try is a simple greedy algorithm for approximating MAX-VOL:

Algorithm 1 Greedy

```

1:  $C \leftarrow \emptyset$ 
2: while  $|C| < k$  do
3:   Select the largest norm vector  $v \in A$ 
4:   Remove the projection of  $v$  from every element of  $A$ 
5:    $C \leftarrow C \cup v$ 
6: end while

```

The analysis of the approximation ratio of this algorithm and a lower bound was also provided in [8]. Specifically, let $Vol(Gr)$ be the volume of the column vectors chosen by Greedy and let $Vol(Opt)$ be the optimum volume. Then, we have

Theorem 3 [8] $Vol(Gr) \geq \frac{1}{k!} \cdot Vol(Opt)$.

Theorem 4 [8] *There exists an instance of MAX-VOL for which $Vol(Gr) \leq \frac{1}{2^{k-1}}(1 - \epsilon) \cdot Vol(Opt)$ for arbitrarily small $\epsilon > 0$. Furthermore, this instance can explicitly be constructed.*

Note that there is a gap between the proven approximation ratio and the lower bound implied by the explicit example. The analysis yielding the ratio $1/k!$ is essentially a product of k different mutually exclusive analyses related to each step of the algorithm. However, it is not clear whether the overall contribution of these different steps to the approximation ratio is actually better than their products. Indeed, the lower bound of $1/2^{k-1}$ pertains to such a peculiar construction that we have conjectured a $1/2^{k-1}$ approximation ratio for the greedy algorithm. Hence, in general, proving an exponential inapproximability for this problem is an important step towards characterizing its approximability properties. It will show that the greedy algorithm is almost the best one can hope for.

This work takes a step towards this goal and proves exponential inapproximability for MAX-VOL via a gap preserving reduction from the well known Label-Cover problem using the Parallel Repetition Theorem [26]. In doing so, we will establish that the greedy algorithm is asymptotically optimal up to a logarithm in the exponent. Specifically, we prove the following theorem:

Theorem 5 *There exists $\delta < 1$ and $c > 0$ such that the problem MAX-VOL is not approximable within 2^{-ck} for $k = \delta n$, unless $P = NP$.*

Our reduction which can be used to prove inapproximability results for other matrix approximation problems with different objective functions, may also be of independent interest.

1.1 Preliminaries and Notation

We introduce some preliminary notation and definitions. Let a matrix A be given in column notation as: $A = \{v_1, v_2, \dots, v_n\}$. The volume of A , $Vol(A)$ can be recursively defined as follows: if A contains one column, i.e. $A = \{v_1\}$, then $Vol(A) = \|v\|_2$, where $\|\cdot\|_2$ is the Euclidean norm. If A has more than one column, $Vol(A) = \|v - \pi_{(A-\{v\})}(v)\|_2 \cdot Vol(A - \{v\})$ for any $v \in A$, where $\pi_X(v)$ is the projection of v onto the space spanned by the column vectors of X . It is well known that $\pi_{(A-\{v\})}(v) = A_v A_v^+ v$, where A_v is the matrix whose columns are the vectors in $A - \{v\}$, and A_v^+ is the pseudo-inverse of A_v (see for example [15]). Using this recursive expression, we have

$$Vol(S) = Vol(A) = \|v_1\|_2 \cdot \prod_{i=1}^{n-1} \|v_{i+1} - A_i A_i^+ v_{i+1}\|_2$$

where $A_i = \{v_1 \dots v_i\}$ for $\leq i \leq n - 1$.

We observe a simple fact about the “distance” of a vector to a subspace in the following lemma, which will be useful in the final proof. Given two sets of vectors P and $Q = \{q_1, \dots, q_m\}$, let $d(q, P) = \|q - \pi_P(q)\|_2$ denote the distance of $q \in Q$ to the space spanned by the vectors in P .

Lemma 6 (Union Lemma) $Vol(P \cup Q) \leq Vol(P) \cdot \prod_{i=1}^n d(q_i, P)$.

Proof We argue by induction on m . For $m = 1$, Q has one element and the statement trivially holds. Assume that it is true for $n = k$ where $Q = \{q_1, \dots, q_k\}$. Then, for any q_{k+1}

$$\begin{aligned} Vol(P \cup Q \cup \{q_{k+1}\}) &= Vol(P \cup Q) \cdot d(q_{k+1}, P \cup Q) \\ &\leq_{(a)} Vol(P \cup Q) \cdot d(q_{k+1}, P) \\ &\leq_{(b)} Vol(P) \cdot \prod_{i=1}^k d(q_i, P) \cdot d(q_{k+1}, P) \\ &= Vol(P) \cdot \prod_{i=1}^{k+1} d(q_i, P). \end{aligned}$$

(a) follows because $d(q, A \cup B) \leq d(q, A)$ for any A, B and (b) follows by the induction hypothesis.

1.2 Related Work

The concept of volume has been closely related to matrix approximation and mainly studied from a linear algebraic perspective. There are a few results revealing the relationship between the volume of a subset of columns of a matrix and its approximation. In [11], the authors introduced *volume sampling* to find low-rank approximation to a matrix where one picks a subset of columns

with probability proportional to the volume of the simplex they define. In volume sampling, one picks a subset of columns S of size k with probability

$$P_S = \frac{\text{Vol}(S)^2}{\sum_{T:|T|=k} \text{Vol}(T)^2},$$

where the summation in the denominator is over all subsets of size k . This sampling provides an almost tight low-rank approximation of a matrix in Frobenius norm. Improving this existence result, Deshpande and Vempala [12] provided an adaptive randomized algorithm for the low-rank approximation problem, which includes a sub-procedure that repetitively chooses a small number of columns by approximating volume sampling. This algorithm is essentially a greedy algorithm and can be regarded as a randomized version of the greedy algorithm we have analyzed for MAX-VOL [8]. They show that, if \tilde{P}_S is the probability that this algorithm chooses a subset of columns S of size k , then

$$\tilde{P}_S \leq k! \cdot P_S. \quad (1)$$

Thus, not only is sampling larger volume columns good, but approximately sampling columns with large volume can prove useful for matrix approximation. A natural question is to ask what happens when one finds a set of columns with the largest volume (deterministic), which is our problem MAX-VOL. Note that, the last expression (1) is reminiscent of the approximation ratio we have proved for MAX-VOL in [8], but its analysis relies on a linear algebraic identity whereas the result in [8] is derived via combinatorial means. MAX-VOL and volume sampling seem to be related, but they have different characteristics. MAX-VOL is proven to be intractable by using complexity theoretic tools, whereas according to a recent result by Deshpande and Rademacher [10], volume sampling can be exactly implemented in polynomial time. This work together with [10] reveals the fact that, although one can exactly sample the columns of a matrix with probability proportional to their volumes, identifying a subset with the maximum volume is hard.

Goreinov and Tyrtshnikov [16] provided explicit statements of how MAX-VOL, in particular, is related to low-rank approximations in the following theorem:

Theorem 7 [16] *Suppose that A is an $m \times n$ block matrix of the form*

$$A = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}$$

where A_{11} is nonsingular, $k \times k$, whose volume is at least μ^{-1} times the maximum volume among all $k \times k$ sub-matrices. Then ¹

$$\|A_{22} - A_{21}A_{11}^{-1}A_{12}\|_{\infty} \leq (k+1)\mu\sigma_{k+1}(A).$$

¹ $\|B\|_{\infty}$ denotes the maximum modulus of the entries of a matrix B .

This theorem implies that if one has a good approximation to the maximum volume $k \times k$ sub-matrix, then the rows and columns corresponding to this sub-matrix can be used to obtain a good approximation to the entire matrix in the ∞ -norm. If $\sigma_{k+1}(A)$ is small for some small k , then this yields a low-rank approximation to A . [17] also proves a similar result to Theorem 7.

Pan [24] unifies the main approaches developed for finding RRQR factorizations by defining the concept of *local maximum volume* and then gives a theorem relating it to the quality of approximation.

Definition 8 [24] *Let $A \in \mathbb{R}^{m \times n}$ and C be a sub-matrix of A formed by any k columns of A . $\text{Vol}(C) (\neq 0)$ is said to be local μ -maximum volume in A , if $\mu \text{Vol}(C) \geq \text{Vol}(C')$ for any C' that is obtained by replacing one column of C by a column of A which is not in C .*

Theorem 9 [24] *For a matrix $A \in \mathbb{R}^{n \times n}$, an integer k ($1 \leq k < n$) and $\mu \geq 1$, let $\Pi \in \mathbb{R}^{n \times n}$ be a permutation matrix such that the first k columns of $A\Pi$ is a local μ -maximum in A . Then, for the QR factorization*

$$A\Pi = Q \begin{pmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{pmatrix},$$

we have

$$\sigma_{\min}(R_{11}) \geq (1/\sqrt{k(n-k)\mu^2 + 1})\sigma_k(A)$$

and

$$\sigma_1(R_{22}) \leq \sqrt{k(n-k)\mu^2 + 1}\sigma_{k+1}(A).$$

We note that, MAX-VOL asks for a stronger property of the set of vectors to be chosen, i.e. it asks for a “good” set of vectors in a global sense rather than only requiring local optimality. Obviously, a solution to MAX-VOL provides a set of vectors with local maximum volume.

Independently of our work, there are some results in computational geometry which are related to the ability to construct large simplices embedded in V-polytopes. Essentially, the problem we consider is a more general version of finding a large simplex in a V-polytope, where the vertices of the polytope are the column vectors. The results in this area are similar to ours in spirit but using different techniques [18, 22, 23]. The most relevant work to ours is that of Koutis [22], which shows exponential inapproximability for finding a large simplex in a V-polytope. Koutis provides a reduction from set packing using an inapproximability result of [20], whereas our reduction is directly from the Label Cover problem.

2 The Label-Cover Problem

Our reduction will be from the Label Cover problem. Label Cover combinatorially captures the expressive power of a 2-prover 1-round proof system for

the problem Max-3SAT(5). Specifically, there exists a reduction from Max-3SAT(5) to Label Cover, so that using the well known parallel repetition technique for the specified proof system yields a new k -fold Label Cover instance. For simplicity, we prefer to state our reduction from Label Cover and for the sake of completeness, we provide a canonical reduction from Max-3SAT(5) to Label Cover.

Max-3SAT(5) is defined as follows: Given a set of $5n/3$ variables and n clauses in conjunctive normal form where each clause contains three distinct variables and each variable appears in exactly five clauses, find an assignment of variables such that it maximizes the fraction of satisfied clauses. The following result is well known [1, 2]:

Theorem 10 *There is a constant $\epsilon > 0$, such that it is NP-hard to distinguish between the instances of Max-3SAT(5) having optimal value 1 and optimal value at most $(1 - \epsilon)$.*

Although this result was proved for general 3CNF formulas, without the requirement that each variable appears exactly 5 times, there is a standard reduction from Max-3SAT to Max-3SAT(5) [13], which only results in a difference in the constant ϵ .

A Label Cover instance L is defined as follows:

$$L = (G(V, W, E), (\Sigma_V, \Sigma_W), \Pi)$$

where

- $G(V, W, E)$ is a regular bipartite graph with vertex sets V and W , and the edge set E .
- Σ_V and Σ_W are the label sets associated with V and W , respectively.
- Π is the collection of constraints on the edge set, where the constraint on an edge e is defined as a function $\Pi_e : \Sigma_V \rightarrow \Sigma_W$.

A labeling is an assignment to the vertices of the graph, $\sigma : \{V \rightarrow \Sigma_V\} \cup \{W \rightarrow \Sigma_W\}$. It is said to satisfy an edge $e = (v, w)$ if $\Pi_e(\sigma(v)) = \sigma(w)$. The Label Cover problem asks for an assignment σ such that the fraction of the satisfied edges is maximum.

A standard reduction from Max-3SAT(5) to Label Cover reveals that

Theorem 11 *There is a constant $\epsilon' > 0$, such that it is NP-hard to distinguish between the instances of Label Cover having optimal value 1 and optimal value at most $(1 - \epsilon')$.*

In order to amplify the gap, one can define a new Label Cover instance for which the vertex set is essentially a set Cartesian product of the original one. This instance, as follows, captures a standard 2-prover 1-round protocol with parallel repetition applied ℓ times. We first note that for a given set $S = \{s_1, \dots, s_n\}$, S^ℓ consists of all ℓ -tuples of the form $(s_{i_1}, \dots, s_{i_\ell})$ where $s_{i_j} \in S$ and i_j runs over $\{1, \dots, n\}$ for $\ell \geq j \geq 1$. Given the original Label

Cover instance $L = (G(V, W, E), (\Sigma_V, \Sigma_W), \Pi)$ reduced from Max-3SAT(5), let

$$L^\ell = (G^\ell(V^\ell, W^\ell, E^\ell), \Sigma_V^\ell, \Sigma_W^\ell, \Pi^\ell),$$

where V^ℓ , W^ℓ , Σ_V^ℓ and Σ_W^ℓ are the ℓ times Cartesian products of the sets V , W , Σ_V and Σ_W , respectively as defined above. Let

- E^ℓ consist of all edges of the form $e = (v, w)$ where $v = (v_{i_1}, \dots, v_{i_\ell})$ and $w = (w_{i_1}, \dots, w_{i_\ell})$ satisfying $(v_{i_j}, w_{i_j}) \in E$ and for all $\ell \geq j \geq 1$.
- Π^ℓ be the collection of constraints on the edge set E^ℓ . The constraint on an edge $e = (v, w)$ where $v = (v_{i_1}, \dots, v_{i_\ell})$ and $w = (w_{i_1}, \dots, w_{i_\ell})$ is a function $\Pi_e^\ell : \Sigma_V^\ell \rightarrow \Sigma_W^\ell$ which is essentially an ℓ -tuple constraint $(\Pi_{e_1}^\ell, \dots, \Pi_{e_\ell}^\ell)$, where $\Pi_{e_j}^\ell = \Pi_{(v_{i_j}, w_{i_j})}$ for $\ell \geq j \geq 1$.

A labeling σ of the vertices V^ℓ and W^ℓ is said to satisfy an edge $e = (v, w)$ where $v = (v_{i_1}, \dots, v_{i_\ell})$ and $w = (w_{i_1}, \dots, w_{i_\ell})$, if $\Pi_e^\ell(\sigma(v)) = \sigma(w)$. Note that this requirement is equal to $\Pi_{(v_{i_j}, w_{i_j})}(\sigma(v_{i_j})) = \sigma(w_{i_j})$ for all $\ell \geq j \geq 1$. It is easy to see that, in this new Label Cover instance, $|V| = (5n/3)^\ell$, $|W| = n^\ell$, $|E| = (5n)^\ell$, $|\Sigma_V^\ell| = 7^\ell$ and $|\Sigma_W^\ell| = 2^\ell$; the degrees of the vertices in V and W is 3^ℓ and 5^ℓ , respectively. The following theorem is a well known result by Raz [26]:

Theorem 12 *There is an absolute constant $\alpha > 0$, such that it is NP-hard to distinguish between the case that $\text{OPT}(L^\ell) = 1$ and $\text{OPT}(L^\ell) \leq 2^{-\alpha\ell}$.*

3 Exponential Inapproximability of MAX-VOL

3.1 The Basic Gadget

At the heart of our analysis is a set of vectors with a special property. We will use a set of vectors (composed of binary entries for simplicity of construction) such that any two of them have large dot-product. We will also require that the dot product of a vector and the binary complement of any other vector is large. More specifically, we need these dot products be proportional to the Euclidean norms squared of the vectors.

Given a vector $v = (v_1 \dots v_m)$ where $v_i \in \{0, 1\}$ for $m \geq i \geq 1$, we denote the binary complement of v by $\bar{v} = (\bar{v}_1 \dots \bar{v}_m)$ where $\bar{v}_i = 1$ if $v_i = 0$, and $\bar{v}_i = 0$ otherwise. We begin with the following lemma:

Lemma 13 *For $m \geq 2$, there exists a set of vectors $B = \{b_1, \dots, b_{2m-1}\}$ of dimension 2^m with binary entries such that the following three conditions hold:*

1. $\|b_i\|_2 = 2^{(m-1)/2}$ for $2m-1 \geq i \geq 1$
2. $b_i \cdot \bar{b}_j = 2^{m-2}$ for $2m-1 \geq i > j \geq 1$.
3. $b_i \cdot b_j = 2^{m-2}$ for $2m-1 \geq i > j \geq 1$.

Proof Consider the Hadamard matrix H of dimension $2^m \times 2^m$ with entries -1 and 1 , constructed recursively by Sylvester's method. Let B be the $(2^m - 1) \times 2^m$ matrix consisting of the rows of H for which we replace -1 's with 0 's, excluding the all 1 's row. We claim that the rows of B satisfy the requirements. Indeed, by the properties of Hadamard matrices, each row of B has exactly 2^{m-1} 1 's which satisfies the first requirement. Note also that, for $m \geq 2$, two distinct rows of H (excluding the all 1 's vector) have exactly 2^{m-2} element-wise dot-products of the following four types: $1 \cdot 1$, $1 \cdot (-1)$, $(-1) \cdot 1$, $(-1) \cdot (-1)$. Considering the construction of B , we have that the dot-product of any two of its rows is 2^{m-2} since all the products in H involving -1 vanishes for B . Similarly the dot-product of a row with the binary complement of another row is 2^{m-2} by symmetry. Thus, the second and the third requirement also hold.

3.2 The Reduction

Lemma 13 guarantees the existence of a set of binary vectors $B = \{b_1, \dots, b_{2^\ell}\}$ of dimension $2^{\ell+1}$ such that the following three conditions hold:

1. $\|b_i\|_2 = 2^{\ell/2}$ for $2^\ell \geq i \geq 1$
2. $b_i \cdot \bar{b}_j = 2^{\ell-1}$ for $2^\ell \geq i > j \geq 1$.
3. $b_i \cdot b_j = 2^{\ell-1}$ for $2^\ell \geq i > j \geq 1$.

B can be constructed in time $O(2^{2\ell})$. In our reduction, ℓ will be a constant (to be exactly determined later) inversely proportional to α which is the constant in Raz' Theorem. Hence, one can construct B in constant time. For the sake of simplicity of our argument, we normalize the vectors in B , which then clearly satisfies

1. $\|b_i\|_2 = 1$ for $2^\ell \geq i \geq 1$
2. $b_i \cdot \bar{b}_j = 1/2$ for $2^\ell \geq i > j \geq 1$.
3. $b_i \cdot b_j = 1/2$ for $2^\ell \geq i > j \geq 1$.

Given a Max-3SAT(5) instance and the reduction described in the previous section, we will define a column vector for each vertex-label pair in L^ℓ , making $(35n/3)^\ell + (2n)^\ell$ vectors in total. (Note that $|V^\ell| = (5n/3)^\ell$, $|W^\ell| = n^\ell$, $\Sigma_V^\ell = \{1, \dots, 7^\ell\}$ and $\Sigma_W^\ell = \{1, \dots, 2^\ell\}$). Each vector will be composed of $|E^\ell| = (5n)^\ell$ "blocks" which are either vectors from the set B or the zero vector according to the adjacency information. More specifically, let $A_{v,i}$ be the vector for the vertex label pair $v \in V^\ell$ and $i \in \Sigma_V^\ell$. Similarly let $A_{w,j}$ be the vector for the pair $w \in W$ and $j \in \Sigma_W^\ell$. Both of these vectors are $(5n)^\ell 2^{\ell+1}$ dimensional. The block of $A_{v,i}$ corresponding to an edge $e \in E^\ell$ is denoted by $A_{v,i}(e)$. The block of $A_{w,j}$ corresponding to an edge $e \in E^\ell$ is denoted by $A_{w,j}(e)$. We define

$$A_{v,i}(e) = \begin{cases} \frac{\overline{b_{\Pi_e^\ell(i)}}}{3^{\ell/2}} & \text{if } e \text{ is incident to } v \\ \vec{0} & \text{if } e \text{ is not incident to } v. \end{cases}$$

$$A_{w,j}(e) = \begin{cases} \frac{b_j}{5^{\ell/2}} & \text{if } e \text{ is incident to } w \\ \vec{0} & \text{if } e \text{ is not incident to } w \end{cases}$$

In order to show how our reduction works, we present a part of a simple bipartite graph in Figure 1 with all the edges drawn between two pairs of nodes, and the corresponding (row) vectors computed by the reduction in Figure 2. Note that $A_{v,i}$ has exactly 3^ℓ non-zero blocks, and $A_{w,j}$ has 5^ℓ non-zero blocks. Hence, according to the definition above, their Euclidean norm is 1. The column vector set for the MAX-VOL instance is defined as

$$A \in \mathbb{R}^{M \times N} = \{A_{v,i} | v \in V^\ell, i \in \Sigma_V^\ell\} \cup \{A_{w,j} | w \in W^\ell, j \in \Sigma_W^\ell\}.$$

Note that $M = (5n)^\ell 2^{\ell+1}$ and $N = (35n/3)^\ell + (2n)^\ell$, both having polynomial size in n for constant ℓ . From an intuitive point of view, we define mutually orthogonal subspaces for each edge, and then we “spread” the Euclidean norm of each vector to the subspaces corresponding to the edges incident to the vertex corresponding to the vector. A crucial observation for this construction is that, vectors A_{v_1, i_1} and A_{v_2, i_2} are orthogonal to each other for all $v_1, v_2 \in V^\ell$, and $i_1, i_2 \in \Sigma_V^\ell$, since there are no edges between the vertices in V^ℓ . The same result holds for the vertices in W^ℓ . From now on, this fact will be used frequently without explicit reference. We set the number of column vectors k to be chosen in the MAX-VOL instance to $|V^\ell| + |W^\ell| = (5n/3)^\ell + n^\ell$. Note that k is a constant fraction of N , the total number of columns, i.e. there exists a constant $\delta < 1$ such that $k = \delta N$.

3.3 Analysis

We start with the completeness of the reduction:




Fig. 1 A part of a simple bipartite graph representing a Label-Cover instance

	e_1	e_2	e_3	\dots	
$A_{v_1,1}$	$\overline{a_{e_1(1)}}$	$\overline{a_{e_2(1)}}$	$\vec{0}$	$\vec{0}$	
$A_{v_1,2}$	$\overline{a_{e_1(2)}}$	$\overline{a_{e_2(2)}}$	$\vec{0}$	$\vec{0}$	
\vdots					
$A_{v_2,1}$	$\vec{0}$	$\vec{0}$	$\overline{a_{e_3(1)}}$	$\vec{0}$	$\overline{a_e(i)} = \frac{\overline{b_{\Pi_e^\ell(i)}}}{3^{\ell/2}}$
\vdots					$a(j) = \frac{b_j}{5^{\ell/2}}$
$A_{w_1,1}$	$a(1)$	$\vec{0}$	$\vec{0}$	$\vec{0}$	
\vdots					
$A_{w_2,1}$	$\vec{0}$	$a(1)$	$a(1)$	$\vec{0}$	

Fig. 2 The resulting (row) vectors in MAX-VOL instance computed from the graph in Figure 1 by our reduction

Theorem 14 *If the Label Cover instance L^ℓ has a labeling that satisfies all the edges, then in the MAX-VOL instance, there exist k column vectors with volume 1.*

Proof We show that there are at least k orthogonal vectors. For an edge $e = (v, w)$, let $i \in \Sigma_V^\ell$ and $j \in \Sigma_W^\ell$ be the labeling of v and w assigned by the optimal labeling which satisfies all the edges. Then, in the MAX-VOL instance the dot product of the vectors $A_{v,i}$ and $A_{w,j}$ is

$$A_{v,i} \cdot A_{w,j} = \sum_{e \in E^\ell} A_{v,i}(e) \cdot A_{w,j}(e) = \overline{b_{\Pi_e^\ell(i)}} \cdot b_j = \overline{b_j} \cdot b_j = 0. \quad (2)$$

This is due to the fact that the labeling satisfies e , i.e. $b_{\Pi_e^\ell(i)} = b_j$. Since all the edges are satisfied, and there exists a vector from each vertex corresponding to the optimal labeling satisfying the equation (2), we have $|V^\ell| + |W^\ell|$ orthogonal vectors, i.e. we have k orthogonal vectors.

Before proving the soundness of the reduction, which will prove hardness of approximation, we first give the intuition for the argument. According to our construction of the MAX-VOL instance, there is a set of vectors corresponding to each node in V^ℓ and W^ℓ . The set of vectors defined for a specific node has

high pair-wise dot products whereas a vector from a node $v_1 \in V^\ell$ and another from $v_2 \in V^\ell$ are orthogonal to each other. The same goes for the vectors defined for W^ℓ . Hence, if vectors are chosen from the same set corresponding to a single node, the total volume will decrease exponentially with respect to the number of such vectors. Let us call these vectors *duplicates* in V^ℓ and W^ℓ . The more intricate part of the analysis is due to the dot products *between* the vectors defined for V^ℓ and W^ℓ , which is enforced to be non-zero by the unsatisfied edges in the Label-Cover instance. We will show that, in case the Label-Cover instance has few satisfied edges, any k vectors chosen in the MAX-VOL instance should satisfy the following: either the number of duplicates in V^ℓ and W^ℓ is large enough so that the total volume is small, or the dot products between V^ℓ and W^ℓ leads to a small volume.

Theorem 15 *There exist absolute constants α and c such that, if the Label Cover instance L^ℓ does not have any labeling that satisfies more than $2^{-\alpha^\ell}$ of the edges, then the volume of any k vectors in the MAX-VOL instance is at most 2^{-ck} .*

Proof Let $V^\ell = \{v_1, \dots, v_{(5n/3)^\ell}\}$ and $W^\ell = \{w_1, \dots, w_{n^\ell}\}$. Let A_v be the vectors corresponding to the vertex $v \in V^\ell$: $A_v = \{A_{v,i} | i \in \Sigma_V^\ell\}$. Similarly, let $A_w = \{A_{w,j} | j \in \Sigma_W^\ell\}$ for $w \in W^\ell$. Let A_{V^ℓ} be the set of all vectors corresponding to the nodes in V^ℓ , and A_{W^ℓ} be the set of all vectors corresponding to the nodes in W^ℓ , i.e.

$$A_{V^\ell} = \bigcup_{i=1}^{(5n/3)^\ell} A_{v_i}, \quad A_{W^\ell} = \bigcup_{i=1}^{n^\ell} A_{w_i}.$$

For a set of vectors C of size k , let $C_u = C \cap A_u$ for all $u \in \{V^\ell \cup W^\ell\}$, $C_{V^\ell} = C \cap A_{V^\ell}$ and $C_{W^\ell} = C \cap A_{W^\ell}$. Let $V^\ell(C)$ and $W^\ell(C)$ be the set of vectors for which C “selects” at least one vector from V^ℓ and W^ℓ , respectively.

$$V^\ell(C) = \{v \in V^\ell | C_v \neq \emptyset\}, \quad W^\ell(C) = \{w \in W^\ell | C(A_w) \neq \emptyset\}.$$

For ease of notation, we let $k_{V_C} = |C_{V^\ell}|$, $k_{W_C} = |C_{W^\ell}|$, $d_{V_C} = k_{V_C} - |V^\ell(C)|$, $d_{W_C} = k_{W_C} - |W^\ell(C)|$. Note that k_{V_C} and k_{W_C} denote how many vectors are chosen by C from V^ℓ and W^ℓ , respectively. Whereas d_{V_C} and d_{W_C} are the total number of duplicates in C_{V^ℓ} and C_{W^ℓ} , respectively. The following lemma relates the number of duplicates on one side with its volume.

Lemma 16 $Vol(C_{V^\ell}) \leq (\sqrt{3}/2)^{d_{V_C}}$ and $Vol(C_{W^\ell}) \leq (\sqrt{3}/2)^{d_{W_C}}$.

Proof Let P be the set of $|V^\ell(C)|$ elements which contains exactly one vector of the form $A_{v,i}$ for each $v \in V^\ell(C)$. In words, we consider the vectors of C corresponding to the nodes in the Label-Cover instance minus all the duplicates. For the duplicate vector $A_{v,j}$, we have $A_{v,i} \cdot A_{v,j} = 1/2$. Hence, $d(A_{v,j}, P) \leq d(A_{v,j}, A_{v,i}) = \sqrt{3}/2$. By the definition of d_{V_C} and by the Union Lemma, we get $Vol(C_{V^\ell}) \leq (\sqrt{3}/2)^{d_{V_C}}$. The argument for $Vol(C_{W^\ell})$ is similar.

Let the constant $c = 1/(3 \cdot 5^{\ell+1})$. Recall that, our reduction will require ℓ to be inversely proportional to α in Raz' Theorem. Hence, although having an exponential dependence on α , c is a constant. We will show that Theorem 15 holds for this value of c ; we will prove that $\text{Vol}(C) \leq 2^{-ck}$ for any set C of k vectors. To this aim, we argue by contradiction. The next lemma roughly states that if the volume of C is large enough, then its vectors are almost equally distributed among the nodes of the Label-Cover instance. This condition will in turn imply a small volume completing our argument.

Claim 17 *If $\text{Vol}(C) \geq 2^{-ck}$ for $c = 1/(3 \cdot 5^{\ell+1})$, then*

$$(1 - \epsilon_1)(5n/3)^\ell < k_{V_C} < (1 + \epsilon_1)(5n/3)^\ell, \quad (3)$$

$$(1 - \epsilon_2)n^\ell < k_{W_C} < (1 + \epsilon_2)n^\ell, \quad (4)$$

where $\epsilon_1 = \frac{1}{3^{\ell+1}}((3/5)^\ell + (3/5)^{2\ell})$ and $\epsilon_2 = \frac{1}{3^{\ell+1}}((3/5)^\ell + 1)$.

Proof First, we note that $\text{Vol}(C) \leq \text{Vol}(C_{V^\ell})$ since all the vectors in the MAX-VOL instance have unit norm. Similarly, $\text{Vol}(C) \leq \text{Vol}(C_{W^\ell})$. Thus, by the premise of the claim, we have $\text{Vol}(C_{V^\ell}) \geq 2^{-ck}$ and $\text{Vol}(C_{W^\ell}) \geq 2^{-ck}$. By Lemma 16, we get

$$(\sqrt{3}/2)^{d_{V_C}} = 2^{d_{V_C}(-1+\log 3/2)} \geq \text{Vol}(C_{V^\ell}) \geq 2^{-ck}$$

which implies $d_{V_C} \leq ck/(1 - \log 3/2) < 5ck$ since $\log 3 < 1.6$. The analysis for d_{W_C} along exactly the same lines also yields $d_{W_C} < 5ck$. Noting the expressions for c and k , and following the definitions, we obtain

$$\begin{aligned} k_{V_C} &= |V^\ell(C)| + d_{V_C} < |V^\ell| + 5ck \\ &= (5n/3)^\ell + \frac{1}{3 \cdot 5^\ell}((5n/3)^\ell + n^\ell) \\ &= (1 + \epsilon_1)(5n/3)^\ell. \end{aligned}$$

Similarly,

$$\begin{aligned} k_{W_C} &= |W^\ell(C)| + d_{W_C} < |W^\ell| + 5ck \\ &= n^\ell + \frac{1}{3 \cdot 5^\ell}((5n/3)^\ell + n^\ell) \\ &= (1 + \epsilon_2)n^\ell \end{aligned}$$

which proves the right hand sides of (3) and (4). Noting that $k_{V_C} + k_{W_C} = k = (5n/3)^\ell + n^\ell$, we get

$$\begin{aligned}
k_{V_C} &= k - k_{W_C} > (5n/3)^\ell + n^\ell - (1 + \epsilon_2)n^\ell \\
&= (5n/3)^\ell - \frac{1}{3^{\ell+1}}((3n/5)^\ell + n^\ell) \\
&= (1 - \epsilon_1)(5n/3)^\ell
\end{aligned}$$

and

$$\begin{aligned}
k_{W_C} &= k - k_{V_C} > (5n/3)^\ell + n^\ell - (1 + \epsilon_1)(5n/3)^\ell \\
&= n^\ell - \frac{1}{3^{\ell+1}}((3n/5)^\ell + n^\ell) \\
&= (1 - \epsilon_2)n^\ell
\end{aligned}$$

which proves the left hand sides.

Claim 17 ensures that if the volume of a set of k vectors exceeds 2^{-ck} , then some certain concentration result should hold, namely Equation (3) and Equation (4). We will now show that, these equations imply $\text{Vol}(C) < 2^{-ck}$ which is our contradiction.

Without loss of generality, let $V^\ell(C) = \{v_1, \dots, v_q\}$, $W^\ell(C) = \{w_1, \dots, w_p\}$. Note that these sets contain the nodes of the Label-Cover instance from which C “selects” at least one vector. Let $Q = \{A_{v_1, i_1}, \dots, A_{v_q, i_q}\}$ where $A_{v_s, i_s} \in C_{v_s}$ for $s = 1, \dots, q$. Let $P = \{A_{w_1, j_1}, \dots, A_{w_p, j_p}\}$ where $A_{v_s, i_s} \in C_{v_s}$ for $s = 1, \dots, p$. By definition,

$$q = k_{V_C} - d_{V_C} > (1 - 2\epsilon_1)(5n/3)^\ell, \quad p = k_{W_C} - d_{W_C} > (1 - 2\epsilon_2)n^\ell.$$

In words, the set of nodes from which C selects at least one vector essentially covers V^ℓ and W^ℓ . These vectors are all orthogonal. From this point of view, $V^\ell(C)$ and $W^\ell(C)$ play an important role in our argument. Since C “covers” V^ℓ and W^ℓ and since the Label-Cover instance has many unsatisfied edges, it means that the dot products of many vectors in C_{V^ℓ} with many vectors in C_{W^ℓ} will be large. This will lead to small volume. Hence, we are essentially interested in the number of unsatisfied edges between $V^\ell(C)$ and $W^\ell(C)$. Since there are at most $2^{-\alpha\ell}$ satisfied edges in the Label-Cover instance, and there are exactly 3^ℓ edges incident to a node in V^ℓ , the number of unsatisfied edges incident to $V^\ell(C)$ is greater than $(1 - 2\epsilon_1 - 2^{-\alpha\ell})(5n)^\ell$. Similarly, the number of unsatisfied edges incident to $W^\ell(C)$ is greater than $(1 - 2\epsilon_2 - 2^{-\alpha\ell})(5n)^\ell$. Thus, the number of unsatisfied edges whose end points are in $V^\ell(C)$ and $W^\ell(C)$, is greater than $(1 - 2\epsilon_1 - 2\epsilon_2 - 2^{-\alpha\ell+1})(5n)^\ell$.

We now give an upper bound for the distance of the vectors in Q to P , namely $\|A_{v_s, i_s} - \pi_P(A_{v_s, i_s})\|_2$ for each $A_{v_s, i_s} \in Q$. To this end, we define the set $N(A_{v_s, i_s}) = \{C_w | e = (A_{v_s, i_s}, w) \text{ is unsatisfied}\}$. Note that the vectors in different sets are mutually orthogonal, and by the reduction we have

$$A_{v_s, i_s} \cdot A_{w, j} = \sum_{e \in E^\ell} A_{v_s, i_s}(e) \cdot A_{w, j}(e) = \overline{b_{\Pi_e^\ell(i)}} \cdot b_j = \frac{1}{2 \cdot 3^{\ell/2} \cdot 5^{\ell/2}}$$

for $A_{w, j} \in N(A_{v_s, i_s})$ since $e = (A_{v_s, i_s}, A_{w, j})$ is unsatisfied. Thus, by the Pythagoras Theorem, we obtain

$$d(A_{v_s, i_s}, P) = \|A_{v_s, i_s} - \pi_P(A_{v_s, i_s})\|_2 < \left(1 - \frac{|N(A_{v_s, i_s})|}{4 \cdot 3^\ell \cdot 5^\ell}\right)^{\frac{1}{2}}.$$

Using the Union Lemma, we get

$$\begin{aligned} \text{Vol}(P \cup Q) &\leq \text{Vol}(P) \cdot \prod_{s=1}^q d(A_{v_s, i_s}, P) \\ &< \text{Vol}(P) \cdot \prod_{s=1}^q \left(1 - \frac{|N(A_{v_s, i_s})|}{4 \cdot 3^\ell \cdot 5^\ell}\right)^{\frac{1}{2}}. \end{aligned}$$

The product in the last expression is maximized when all the factors are equal to each other. We also previously showed that $\sum_{s=1}^q |N(A_{v_s, i_s})| > (1 - 2\epsilon_1 - 2\epsilon_2 - 2^{-\alpha\ell+1})(5n)^\ell$ and that q , the number of distinct nodes hit in V^ℓ satisfies, $q > (1 - 2\epsilon_1)(5n/3)^\ell$. Hence, we obtain

$$\begin{aligned} \text{Vol}(P \cup Q) &< \text{Vol}(P) \cdot \prod_{s=1}^q \left(1 - \frac{\sum_{s=1}^q |N(A_{v_s, i_s})|}{q \cdot 4 \cdot 3^\ell \cdot 5^\ell}\right)^{\frac{1}{2}} \\ &< \text{Vol}(P) \cdot \prod_{s=1}^q \left(1 - \frac{(1 - 2\epsilon_1 - 2\epsilon_2 - 2^{-\alpha\ell+1})(5n)^\ell}{(5n/3)^\ell \cdot 4 \cdot 3^\ell \cdot 5^\ell}\right)^{\frac{1}{2}} \\ &= \text{Vol}(P) \cdot \left(1 - \frac{(1 - 2\epsilon_1 - 2\epsilon_2 - 2^{-\alpha\ell+1})}{4 \cdot 5^\ell}\right)^{\frac{q}{2}} \\ &< \text{Vol}(P) \cdot \left(1 - \frac{(1 - 2\epsilon_1 - 2\epsilon_2 - 2^{-\alpha\ell+1})}{4 \cdot 5^\ell}\right)^{\frac{(1-2\epsilon_1)(5n/3)^\ell}{2}}. \end{aligned}$$

To simplify, let $t = \frac{4 \cdot 5^\ell}{(1-2\epsilon_1-2\epsilon_2-2^{-\alpha\ell+1})}$. For $\ell \geq 1$, we have

$$\epsilon_1 = \frac{1}{3^{\ell+1}} \left((3/5)^\ell + (3/5)^{2\ell} \right) \leq \frac{1}{3^2} \left((3/5) + (3/5)^2 \right) < 3/20.$$

Noting that $\log e \geq 10/7$, we obtain $\log e \cdot (1 - 2\epsilon_1) \geq 10/7 \cdot 7/10 = 1$. Then, we get

$$\begin{aligned}
\text{Vol}(P \cup Q) &< \text{Vol}(P) \cdot \left(1 - \frac{1}{t}\right)^{t \cdot \frac{(1-2\epsilon_1)(5n/3)^\ell}{2t}} \\
&\leq e^{-\frac{(1-2\epsilon_1)(5n/3)^\ell}{2t}} \\
&= 2^{-\log e \cdot \frac{(1-2\epsilon_1)(5n/3)^\ell}{2t}} \\
&\leq 2^{-\frac{(5n/3)^\ell}{2t}},
\end{aligned}$$

where e is the base of the natural logarithm. In the second inequality, we have used the fact that $\text{Vol}(P) \leq 1$ and $(1 - \frac{1}{t})^t \leq e^{-1}$ for $t > 1$.

We will now provide an upper bound for t to further simplify the last expression. To this aim, let ℓ' be the smallest integer such that $2^{-\alpha\ell'+1} \leq 11/27$. Taking logarithms and rearranging, it is easy to see that $\ell' = \left\lceil \frac{\log(\frac{54}{11})}{\alpha} \right\rceil$. Note also that for $\ell \geq 2$, we have $2\epsilon_1 < 1/27$ and $2\epsilon_2 < 3/27$. Then, for $\ell = \ell'$, we get

$$t = \frac{4 \cdot 5^\ell}{(1 - 2\epsilon_1 - 2\epsilon_2 - 2^{-\alpha\ell+1})} < \frac{4 \cdot 5^\ell}{(1 - \frac{1}{27} - \frac{3}{27} - \frac{11}{27})} = \frac{4 \cdot 5^\ell}{(4/9)} = 9 \cdot 5^\ell.$$

Since $k = (5n/3)^\ell + n^\ell$, we also have

$$n^\ell = k/(1 + (5/3)^\ell) > k/(5/3)^{\ell+1},$$

which yields

$$\text{Vol}(P \cup Q) < 2^{-\frac{(5n/3)^\ell}{9 \cdot 5^\ell}} = 2^{-\frac{n^\ell}{3^{\ell+2}}} < 2^{-\frac{k}{3 \cdot 5^{\ell+1}}} = 2^{-ck},$$

which is our contradiction. Thus, the volume of a set of k vectors in a negative instance of MAX-VOL cannot exceed 2^{-ck} for $c = \frac{1}{3 \cdot 5^{\ell+1}}$.

We have shown that

- if the optimal value of the Label Cover instance is 1, then the optimal value of the MAX-VOL instance is 1.
- if the optimal value of the ℓ -fold Label Cover instance is less than $2^{-\alpha\ell}$, then the optimal value of the MAX-VOL instance is less than 2^{-ck} .

By the combination of Theorem 10 and Theorem 12, we know that there exists a gap producing reduction from SAT to ℓ -fold Label Cover with parameters 1 and $2^{-\alpha\ell}$. This means that there is a polynomial time reduction from SAT to MAX-VOL such that, given a formula ϕ

- if ϕ is satisfiable, then $\text{OPT}(\text{MAX-VOL}) = 1$.
- if ϕ is not satisfiable, then $\text{OPT}(\text{MAX-VOL}) < 2^{-ck}$.

Thus, unless $P = NP$, MAX-VOL is inapproximable within 2^{-ck} for some constant $c > 0$.

4 Discussion

Our reduction heavily relies on the Raz’ Parallel Repetition Theorem [26]. Indeed, it doesn’t seem possible to get an exponential inapproximability result without parallel repetition. But, since the degrees of the vertices in the Label-Cover instance exponentially increases with respect to the number of repetitions, our constant c depends on the constant α in Raz’ result. It might be possible to improve this constant by making use of more sophisticated parallel repetition theorems, but we did not proceed so far. Indeed, the exact analysis is irrelevant as the constant will be too small in all cases. Overall, the strength of our result is directly related to the underlying theorems for the inapproximability of Label-Cover.

Another way of getting a stronger hardness result is to find a more sophisticated reduction. In our MAX-VOL instance, the subspaces “reserved” for each edge in the Label-Cover instance are orthogonal to each other. This dramatically simplifies the analysis, yielding perfect completeness, i.e. volume 1 in MAX-VOL. It might be possible to construct a MAX-VOL instance for which these subspaces have some pair-wise angle, so that we sacrifice the perfect completeness, but at the same time get a much smaller soundness. This would improve the inapproximability result.

The obvious open problem is whether the inapproximability can be strengthened to 2^{-k+1} . Recall that this is the lower bound for the greedy algorithm for MAX-VOL. Considering the multiplicative nature of the problem yielding a very small approximation ratio for the obvious greedy algorithm, a significant improvement of the upper bound would be expected to provide asymptotically better approximations in the exponent. This suggests that the inherent hardness of MAX-VOL might be very close to the performance of the greedy algorithm. However, with the techniques we have used, it is not possible to break the dependence of c on the constant in the parallel repetition theorems.

We would finally like to point out that the reduction and the analysis provided in this paper might be a good starting point for studying hardness of other matrix approximation problems in general (e.g. [4, 11]) for which no technique related to the PCP theorem have been used. Such an extension to other matrix approximation problems is not trivial. Indeed, computing the volume is already difficult although purely geometric intuition is used. Relating this to other linear algebraic functions (e.g. singular values) which continuously depend on the entries will put even more strain on the analysis.

Acknowledgments: We would like to thank the anonymous referees for their helpful comments and Ioannis Koutis who, in the final stages of this paper, pointed out to us the relevant lines of research in V-polytope theory [22].

References

1. S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.
2. S. Arora and S. Safra. Probabilistic checking of proofs: a new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998.
3. C. Boutsidis, P. Drineas, and M. Magdon-Ismail. Near-optimal column based matrix reconstruction. In *FOCS '11: Proceedings of the 52nd IEEE Symposium on Foundations of Computer Science*. to appear.
4. C. Boutsidis, M. W. Mahoney, and P. Drineas. An improved approximation algorithm for the column subset selection problem. In *SODA '09: Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 968–977, 2009.
5. T. F. Chan. Rank revealing QR factorizations. *Linear Algebra and Its Applications*, (88/89):67–82, 1987.
6. T. F. Chan and P. Hansen. Low-rank revealing QR factorizations. *Numerical Linear Algebra with Applications*, 1:33–44, 1994.
7. S. Chandrasekaran and I. C. F. Ipsen. On rank-revealing factorizations. *SIAM Journal on Matrix Analysis and Applications*, 15:592–622, 1994.
8. A. Çivril and M. Magdon-Ismail. On selecting a maximum volume sub-matrix of a matrix and related problems. *Theoretical Computer Science*, 410(47-49):4801–4811, 2009.
9. F. R. de Hoog and R. M. M. Mattheijb. Subset selection for matrices. *Linear Algebra and its Applications*, (422):349–359, 2007.
10. A. Deshpande and L. Rademacher. Efficient volume sampling for row/column subset selection. In *FOCS '10: Proceedings of 51st Annual Symposium on Foundations of Computer Science*, pages 329–338, 2010.
11. A. Deshpande, L. Rademacher, S. Vempala, and G. Wang. Matrix approximation and projective clustering via volume sampling. *Theory of Computing*, 2(1):225–247, 2006.
12. A. Deshpande and S. Vempala. Adaptive sampling and fast low-rank matrix approximation. In *RANDOM'06: 10th International Workshop on Randomization and Computation*, pages 292–303, 2006.
13. U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.
14. G. H. Golub, V. Klema, and G. W. Stewart. Rank degeneracy and least squares problems. Technical report, Dept. of Computer Science, Univ. of Maryland, 1976.
15. G. H. Golub and C. V. Loan. *Matrix Computations*. Johns Hopkins U. Press, 1996.
16. S. A. Goreinov and E. E. Tyrtshnikov. The maximal-volume concept in approximation by low-rank matrices. In *Contemporary Mathematics*, volume 280, pages 47–51. AMS, 2001.
17. S. A. Goreinov, N. L. Zamarashkin, and E. E. Tyrtshnikov. Pseudo-skeleton approximations by matrices of maximal volume. *Matematicheskie Zametki*, 62:619–623, 1997.
18. P. Gritzmann, V. Klee, and D. G. Larman. Largest j -simplices n -polytopes. *Discrete and Computational Geometry*, 13:477–515, 1995.
19. M. Gu and S. C. Eisenstat. Efficient algorithms for computing a strong rank-revealing QR factorization. *SIAM Journal on Scientific Computing*, 17(4):848–869, 1996.
20. E. Hazan, S. Safra, and O. Schwartz. On the complexity of approximating k -set packing. *Computational Complexity*, 15:20–39, 2006.
21. Y. P. Hong and C. T. Pan. Rank-revealing QR factorizations and the singular value decomposition. *Mathematics of Computation*, 58:213–232, 1992.
22. I. Koutis. Parameterized complexity and improved inapproximability for computing the largest j -simplex in a V -polytope. *Information Processing Letters*, 100:8–13, 2006.
23. A. Packer. Polynomial-time approximation of largest simplices in V -polytopes. *Discrete Applied Mathematics*, 134:213–237, 2004.
24. C. T. Pan. On the existence and computation of rank-revealing LU factorizations. *Linear Algebra and its Applications*, 316(1-3):199–222, 2000.
25. C. T. Pan and P. T. P. Tang. Bounds on singular values revealed by QR factorizations. *BIT Numerical Mathematics*, 39:740–756, 1999.
26. R. Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, 1998.