

On Selecting a Maximum Volume Sub-matrix of a Matrix and Related Problems

A.Çivril^{*,a}, M.Magdon-Ismail^a

^a*Rensselaer Polytechnic Institute, Computer Science Department, 110 8th Street Troy, NY 12180-3590 USA*

Abstract

Given a matrix $A \in \mathbb{R}^{m \times n}$ (n vectors in m dimensions), we consider the problem of selecting a subset of its columns such that its elements are as linearly independent as possible. This notion turned out to be important in low-rank approximations to matrices and rank revealing QR factorizations which have been investigated in the linear algebra community and can be quantified in a few different ways. In this paper, from a complexity theoretic point of view, we propose four related problems in which we try to find a sub-matrix $C \in \mathbb{R}^{m \times k}$ of a given matrix $A \in \mathbb{R}^{m \times n}$ such that (i) $\sigma_{max}(C)$ (the largest singular value of C) is minimum, (ii) $\sigma_{min}(C)$ (the smallest singular value of C) is maximum, (iii) $\kappa(C) = \sigma_{max}(C)/\sigma_{min}(C)$ (the condition number of C) is minimum, and (iv) the volume of the parallelepiped defined by the column vectors of C is maximum. We establish the NP-hardness of these problems and further show that they do not admit PTAS. We then study a natural greedy heuristic for the maximum volume problem and show that it has approximation ratio $2^{-O(k \log k)}$. Our analysis of the greedy heuristic is tight to within a logarithmic factor in the exponent, which we show by explicitly constructing an instance for which the greedy heuristic is $2^{-\Omega(k)}$ from optimal. When A has unit norm columns, a related problem is to select the maximum number of vectors with a given volume. We show that if the optimal solution selects k columns, then greedy will select $\Omega(k/\log k)$ columns, providing a $\log k$ -approximation.

Key words: subset selection, condition number, maximum volume sub-matrix, complexity, approximation

1. Introduction

To motivate the discussion, consider the set of three vectors

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

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 for the representatives to be “as far away

*Corresponding author

Email addresses: civria@cs.rpi.edu (A.Çivril), magdon@cs.rpi.edu (M.Magdon-Ismail)

from each other as possible”. A natural formalization of this problem is to find the representatives which span the largest volume, since the volume is a quantification of how far the vectors are from each other. Another would be to choose them so as the matrix that they form is well-conditioned, i.e. its condition number is small which intuitively means that the matrix is as close as possible to an orthogonal one and its smallest singular value is large with respect to the largest singular value. Thus, given a set of n vectors in \mathbb{R}^m represented as a matrix $A \in \mathbb{R}^{m \times n}$ and a positive integer k , we discuss four distinct problems in which we ask for a subset $C \in \mathbb{R}^{m \times k}$ satisfying some spectral optimality condition:

1. MinMaxSingularValue: $\sigma_1(C)$ is minimum;
2. MaxMinSingularValue: $\sigma_k(C)$ is maximum;
3. MinSingularSubset: $\kappa(C) = \sigma_1(C)/\sigma_k(C)$ is minimum;
4. MAX-VOL: $Vol(C) = \prod_{i=1}^k \sigma_i(C)$, the volume of the parallelepiped defined by the column vectors of C is maximum.

In all cases, the optimization is over all possible choices of C , and $\sigma_1(C) \geq \sigma_2(C) \geq \dots \geq \sigma_k(C)$ are the singular values of the sub-matrix defined by C . Before presenting the main results of the paper, we will first briefly review how these concepts are related to low-rank approximations to matrices and rank revealing QR factorizations.

1.1. Low-rank approximations to matrices

The notion of volume has already received some interest in the algorithmic aspects of linear algebra. In the past decade, the problem of matrix reconstruction and finding low-rank approximations to matrices using a small sample of columns has received much attention (See for example [1, 2, 3?]). Ideally, one has to choose the columns to be as independent as possible when trying to reconstruct a matrix using a few columns. Along these lines, in [1], the authors introduce ‘volume sampling’ to find low-rank approximation to a matrix where one picks a subset of columns with probability proportional to their volume squared. Improving the existence results in [1], [4] also provides an adaptive randomized algorithm which includes repetitively choosing a small number of columns in a matrix to find a low-rank approximation. The authors show that if one samples columns proportional to the volume squared, then one obtains a provably good matrix reconstruction (randomized). Thus, sampling larger volume columns is good. A natural question is to ask what happens when one uses the columns with largest volume (deterministic). The problem MAX-VOL is the algorithmic problem of obtaining the columns with largest volume and we rely on [4] as the qualitative intuition behind why obtaining the maximum volume sub-matrix should play an important role in matrix reconstruction.

Goreinov and Tyrtshnikov [5] provide a more explicit statement of how volume is related to low-rank approximations in the following theorem:

Theorem 1.1. [5] *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 $\|A_{22} - A_{21}A_{11}^{-1}A_{12}\|_{\infty} \leq \mu(k+1)\sigma_{k+1}(A)$.

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. If $\sigma_{k+1}(A)$ is small for some small k , then this yields a low-rank approximation to A . Thus, finding maximum volume sub-matrices is important for matrix reconstruction. We take a first step in this direction by considering the problem of choosing an $m \times k$ sub-matrix of maximum volume. Relating maximum volume $k \times k$ sub-matrices to maximum volume $m \times k$ matrices or obtaining an analogue of Theorem 1.1 for $m \times k$ sub-matrices is beyond the scope of this paper.

1.2. Rank revealing QR factorizations

QR factorization, which has many practical applications [6], is another approach to finding an orthonormal basis for the space spanned by the columns of a matrix. The task is to express a given matrix $A \in \mathbb{R}^{n \times n}$ as the product of an orthogonal matrix $Q \in \mathbb{R}^{n \times n}$ and an upper-triangular matrix $R \in \mathbb{R}^{n \times n}$. For $1 \leq k \leq n$, the first k columns of Q spans the same space as that spanned by the first k columns of A . A naive approach to finding such a factorization might yield linearly dependent columns in Q , if the structure of A is disregarded. Hence, one might try to consider permuting the columns of A so as to find a QR factorization which reveals ‘‘important’’ information about the matrix. Along these lines, rank revealing QR factorizations were introduced by Chan [7].

Given a matrix $A \in \mathbb{R}^{n \times n}$, consider the QR factorization of the form

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

where $R_{11} \in \mathbb{R}^{k \times k}$ and $\Pi \in \mathbb{R}^{n \times n}$ is a permutation matrix. One can easily see that, by the interlacing property of singular values (see [6]), $\sigma_k(R_{11}) \leq \sigma_k(A)$ and $\sigma_1(R_{22}) \geq \sigma_{k+1}(A)$. If the numerical rank of A is k , i.e. $\sigma_k(A) \gg \sigma_{k+1}(A)$, then one naturally would like to find a permutation Π for which $\sigma_k(R_{11})$ is sufficiently large and $\sigma_1(R_{22})$ is sufficiently small. A QR factorization is said to be a rank revealing QR (RRQR) factorization if $\sigma_k(R_{11}) \geq \sigma_k(A)/p(k, n)$ and $\sigma_1(R_{22}) \leq \sigma_{k+1}(A)p(k, n)$ where $p(k, n)$ is a low degree polynomial in k and n .

The QR algorithm proposed by Businger and Golub [8], which is essentially the algorithm we will analyze for maximizing the volume, works well in practice. But, as is pointed out by Kahan [9], there are matrices

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

where it fails to satisfy the requirements of an RRQR factorization yielding exponential $p(k, n)$. Much research on finding RRQR factorizations has yielded improved results for $p(k, n)$ [7, 10, 11, 12, 13, 14]. It was noted in [10] that it turns out that “the selection of the sub-matrix with the maximum smallest singular value suggested in [15] can be replaced by the selection of a sub-matrix with maximum determinant”. (Our hardness results for all the problems we consider also make a justification of how similar they are). Along these lines, the effort has been trying to find a sub-matrix with a volume as large as possible. Pan [16] unifies the main approaches by defining the concept of *local maximum volume* and then gives a theorem relating it to $p(k, n)$.

Definition 1.2. [16] *Let $A \in \mathbb{R}^{m \times n}$ and C be a sub-matrix of A formed by any k columns of A . $Vol(C) (\neq 0)$ is said to be local μ -maximum volume in A , if $\mu Vol(C) \geq Vol(C')$ for any C' that is obtained by replacing one column of C by a column of A which is not in C .*

Theorem 1.3. [16] *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 would like to 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 requiring local optimality. Nevertheless, it is clear that an approximation ratio for MAX-VOL translates to a result in the context of RRQR factorizations. Because, if one could obtain a subset which is μ -maximum (as opposed to local μ -maximum) then the same theorem would hold, since the volume of *any* new set of vectors which is a result of exchanging a column between the current set and the rest of the columns is smaller than the largest possible volume. However, the result obtained via the approximation factor we provide for Greedy is already inferior to a mathematically different analysis which proves $p(k, n) = \sqrt{n-k}2^k$ [13] and we do not state it explicitly.

1.3. Our contributions

First, we establish the NP-hardness of the problems we consider. In fact, we prove that no PTAS for them exists by showing that they are inapproximable to within some factor. Specifically, we obtain the following inapproximability results:

1. MinMaxSingularValue is inapproximable to within $2/\sqrt{3} - \epsilon$;
2. MaxMinSingularValue is inapproximable to within $(2/3)^{1/2(k-1)} + \epsilon$;
3. MinSingularSubset is inapproximable to within $(2^{2k-3}/3^{k-2})^{1/2(k-1)} - \epsilon$;

4. MAX-VOL is inapproximable to within $2\sqrt{2}/3 + \epsilon$.

Next, we consider a simple (deterministic) greedy algorithm for the last problem and show that it achieves a $1/k!$ approximation to the optimal volume when selecting k columns. We also construct an explicit example for which the greedy algorithm gives no better than a $1/2^{k-1}$ approximation ratio, thus proving that our analysis of the greedy algorithm is almost tight (to within a logarithmic factor in the exponent). An important property of the approximation ratio for the greedy algorithm is that it is independent of n , and depends only on the number of columns one wishes to select.

We then consider the related problem of choosing the maximum number of vectors with a given volume, in the case when all columns in A have unit norm. If the optimal algorithm loses a constant factor with every additional vector selected (which is a reasonable situation), then the optimal volume will be $2^{-\Omega(k)}$. When the optimal volume for k vectors is $2^{-\Omega(k)}$ as motivated above, we prove that the greedy algorithm chooses $\Omega(k/\log k)$ columns having at least that much volume. Thus, the greedy algorithm is within a $\log k$ -factor of the maximum number of vectors which can be selected given a target volume.

1.4. Preliminaries and Notation

For a matrix $A \in \mathbb{R}^{m \times n}$ where $n \leq m$, $\sigma_i(A)$ is the i^{th} largest singular value of A for $1 \leq i \leq n$. Let $A = \{v_1, v_2, \dots, v_n\}$ be given in column notation. The volume of A , $Vol(A)$ can be recursively defined as follows: if A contains one column, i.e. $A = \{v_1\}$, then $Vol(A) = \|v_1\|$, where $\|\cdot\|$ is the Euclidean norm. If A has more than one column, $Vol(A) = \|v - \pi_{(A-\{v\})}(v)\| \cdot Vol(A - \{v\})$ for any $v \in A$, where $\pi_A(v)$ is the projection of v onto the space spanned by the column vectors of A . 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 [6]). Using this recursive expression, we have

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

where $A_i = [v_1 \dots v_i]$ for $1 \leq i \leq n-1$.

1.5. Organization of the paper

The remainder of the paper is structured as follows: In Section 2, we provide hardness results for the four problems. The approximation ratio of a greedy algorithm for MAX-VOL is analyzed in Section 3 where we also show tightness of the analysis. Finally, some open questions and comments are outlined in Section 4.

2. Hardness of Subset Selection Problems

We define four decision problems:

Problem: Min-MaxSingularValue

Instance: A matrix $A \in \mathbb{R}^{m \times n}$ of rank at least k , and $M \in \mathbb{R}$.

Question: Does there exist a sub-matrix $C \in \mathbb{R}^{m \times k}$ of A such that $\sigma_1(C) \leq M$.

Problem: **Max-MinSingularValue**

Instance: A matrix $A \in \mathbb{R}^{m \times n}$ of rank at least k , and $M \in \mathbb{R}$.

Question: Does there exist a sub-matrix $C \in \mathbb{R}^{m \times k}$ of A such that $\sigma_k(C) \geq M$.

Problem: **Min-SingularSubset**

Instance: A matrix $A \in \mathbb{R}^{m \times n}$ of rank at least k , and $M \in \mathbb{R}$.

Question: Does there exist a sub-matrix $C \in \mathbb{R}^{m \times k}$ of A such that $\sigma_1(C)/\sigma_k(C) \leq M$.

Problem: **MAX-VOL**

Instance: A matrix $A \in \mathbb{R}^{m \times n}$ with normalized columns and of rank at least k , and $M \in [0, 1]$.

Question: Does there exist a sub-matrix $C \in \mathbb{R}^{m \times k}$ of A such that $\text{Vol}(C) \geq M$?

Theorem 2.1. *Max-MinSingularValue, Min-MaxSingularValue, Min-SingularSubset and MAX-VOL are NP-Hard.*

PROOF. We give a reduction from ‘exact cover by 3-sets’, which is known to be NP-complete (See for example [17, 18]). This reduction will provide the NP-hardness result for all the problems.

Problem: **Exact cover by 3-sets (X3C)**

Instance: A set Q and a collection C of 3-element subsets of Q .

Question: Does there exist an exact cover for Q , i.e. a sub-collection $C' \subseteq C$ such that every element in Q appears exactly once in C' ?

We use the following reduction from X3C to the problems: let $Q = \{q_1, q_2, \dots, q_m\}$ and $C = \{c_1, c_2, \dots, c_n\}$ be given as an instance of X3C. We construct the matrix $A \in \mathbb{R}^{m \times n}$, in which each column $A^{(j)}$ corresponds to the 3-element set c_j . The non-zero entries in $A^{(j)}$ correspond to the elements in c_j . Specifically, set

$$A_{ij} = \begin{cases} 1/\sqrt{3} & \text{if } q_i \in c_j \\ 0 & \text{otherwise} \end{cases}$$

(Note that every $A^{(j)}$ has exactly 3 non-zero entries and has unit norm.) For the reduced instances, we set $k = m/3$ and $M = 1$.

It is clear that the reduction is polynomial time. All that remains is to show that the instance of X3C is true if and only if the corresponding instances of the four decision problems are true.

Suppose the instance of X3C is true. Then, there is a collection $C' = \{c_{i_1}, c_{i_2}, \dots, c_{i_{m/3}}\}$ of cardinality $m/3$, which exactly covers Q . (Note that, m should be a multiple of 3, otherwise no solution exists.) Consider the sub-matrix C of A corresponding to the 3-element sets in C' . Since the cover is exact, $c_{i_j} \cap c_{i_k} = \emptyset \forall j, k \in \{1, \dots, m/3\}$ where $j \neq k$, which means that $A^{(i_j)} \cdot A^{(i_k)} = 0$. Hence, C is orthonormal and all its singular values are 1, which makes the instances of all four problems we consider true.

Conversely, suppose the instance of Min-MaxSingularValue is true, i.e. there exists C such that $\sigma_1(C) \leq 1$. We have $\sigma_1(C) = \|C\|_2 \geq \|C\|_F / \sqrt{k} = 1$, which gives $\sigma_1(C) = 1$. On the other hand, $\sum_{i=1}^k \sigma_i(C)^2 = \|C\|_F^2 = k$. Thus, all the singular values of C are equal to 1, i.e. C is an orthogonal matrix. Now, suppose the instance of Max-MinSingularValue is true, namely there exists C such that $\sigma_k(C) \geq 1$. Then, the volume defined by the vectors in C , $Vol(C) = \prod_{i=1}^k \sigma_i(C) \geq 1$. Since the vectors are all normalized, we also have $Vol(C) \leq 1$, which gives $\prod_{i=1}^k \sigma_i(C) = 1$. Thus, all the singular values of C are equal to 1, which means that C is an orthogonal matrix. If the instance of Min-SingularSubset is true, i.e. there exists C such that $\sigma_1(C)/\sigma_k(C) \leq 1$, we immediately have that C is an orthogonal matrix. Finally, that the instance of MAX-VOL is true means that the columns are pair-wise orthonormal and we have the desired result.

Thus, if any of the reduced instances are true, then there is a C in A whose columns are pairwise orthonormal. We will now show that if such a C exists, then the instance of X3C is true. Let u, v be two columns in C ; we have $u \cdot v = 0$. Since the entries in C are all non-negative, $u_i \cdot v_i = 0 \forall i \in [1, m]$, i.e. u and v correspond to 3-element sets which are disjoint. Hence, the columns in C correspond to a sub-collection C' of 3-element sets, which are pair-wise disjoint. Therefore, every element of Q appears at most once in C' . C' contains m elements corresponding to the m non-zero entries in C . It follows that every element of Q appears exactly once in C' , concluding the proof.

Our reduction in the NP-hardness proofs yields gaps, which also provides hardness of approximation results for the optimization versions of the problems.

Theorem 2.2. *Min-MaxSingularValue(k) is NP-hard to approximate within $2/\sqrt{3} - \epsilon$.*

PROOF. We will provide a lower bound for $\sigma_1(C)$ for the reduced instance of X3C when it is false, which will establish the hardness result. Assume the X3C instance is not true. Then any collection of size $m/3$ has at least two sets which have non-empty intersection. Let s_i and s_j be two sets such that $|s_i \cap s_j| = 1$. And, let v_i, v_j be the corresponding vectors in the Min-MaxSingularValue instance. Then, we have $v_1 \cdot v_2 = 1/3$. Hence, v_1 and v_2 correspond to the following matrix V upto rotation:

$$V = \begin{pmatrix} 1 & \frac{1}{3} \\ 0 & \frac{2\sqrt{2}}{3} \end{pmatrix}$$

Note that, if $|s_i \cap s_j| > 1$, then the largest singular value of the corresponding matrix will be greater than that of V as $v_1 \cdot v_2$ will have a greater value. Also, it is a well known fact that the largest eigenvalue of any symmetric matrix A is greater than that of any principal sub-matrix of A . Thus, if we consider a matrix W

of more than two vectors which also contain v_i and v_j , its largest singular value (which is the square root of the largest eigenvalue of $W^T W$) will be greater than that of V . Hence, in order to find a lower bound for $\sigma_1(C)$, it suffices to analyze V . This amounts to finding the square roots of the eigenvalues of $V^T V$. Hence, we are seeking λ such that

$$\det(V^T V - \lambda I) = \begin{vmatrix} \frac{10}{9} - \lambda & \frac{2\sqrt{2}}{9} \\ \frac{2\sqrt{2}}{9} & \frac{8}{9} - \lambda \end{vmatrix} = 0. \quad (1)$$

$\lambda = 4/3$ and $\lambda = 2/3$ satisfy (1). Hence, $\sigma_1(C) \geq 2/\sqrt{3}$, which concludes the proof.

Theorem 2.3. *MAX-VOL(k) is NP-Hard to approximate within $2\sqrt{2}/3 + \epsilon$.*

PROOF. Assume the X3C instance is not true. Then, we have at least one overlapping element between two sets. Any collection of size $m/3$ will have two sets v_1, v_2 which have non-zero intersection. The corresponding columns in A' have $d(v_1, v_2) = \|v_1 - (v_1 \cdot v_2)v_2\| = \|v_1 - (1/3)v_2\| \leq 2\sqrt{2}/3$, where $d(v_1, v_2)$ is the orthogonal part of v_1 with respect to v_2 . Since $\text{Vol}(A') \leq d(v_1, v_2)$, we have $\text{Vol}(A') \leq 2\sqrt{2}/3$. A polynomial time algorithm with a $2\sqrt{2}/3 + \epsilon$ approximation factor for MAX-VOL would thus decide X3C, which would imply $P = NP$.

Theorem 2.4. *Max-MinSingularValue is NP-hard to approximate within $(2/3)^{1/2(k-1)} + \epsilon$.*

PROOF. If the X3C instance is false, from the proof in Theorem 2.3 we have $\prod_{i=1}^k \sigma_i(C) \leq 2\sqrt{2}/3$. Combining this with $\sigma_1(C) \geq 2/\sqrt{3}$ from the proof in Theorem 2.2, we get $\prod_{i=2}^k \sigma_i(C) \leq \sqrt{6}/3$, which gives $\sigma_k(C) \leq (\sqrt{6}/3)^{1/(k-1)} = (2/3)^{1/2(k-1)}$.

Theorem 2.5. *Min-SingularSubset is NP-hard to approximate within $(2^{2k-3}/3^{k-2})^{1/2(k-1)} - \epsilon$.*

PROOF. Assuming that the X3C instance is false, from the proofs in Theorem 2.2 and Theorem 2.4, we have $\sigma_1(C)/\sigma_k(C) \geq (2^{2k-3}/3^{k-2})^{1/2(k-1)} - \epsilon$.

3. The Greedy Approximation Algorithm for MAX-VOL

Having shown that the decision problem MAX-VOL is NP-hard, it has two natural interpretations as an optimization problem for a given matrix A :

1. *MAX-VOL(k)*: Given k , find a subset of size k with maximum volume.
2. *MaxSubset(V)*: Given V and that A has unit norm vectors, find the largest subset $C \subseteq A$ with volume at least V .

The natural question is whether there exists a simple heuristic with some approximation guarantee. One obvious strategy is the following greedy algorithm which was also proposed in [8] to construct QR factorizations of matrices:

We would like to note that one can obtain a result related to the approximation ratio of Greedy which is implicit in [13] via the following theorem:

Theorem 3.1. [13] For a matrix $A \in \mathbb{R}^{n \times n}$ and an integer $k(1 \leq k < n)$, let the first k columns of $A\Pi$ be the columns chosen by Greedy where $\Pi \in \mathbb{R}^{n \times n}$ is a permutation matrix and

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

Then, $\sigma_i(R_{11}) \geq \sigma_i(A)/(\sqrt{n-i}2^i)$ for $1 \leq i \leq k$.

Based on this theorem, one can easily derive the following result.

Theorem 3.2. Greedy has approximation ratio $O(2^{-k(k-1)/2}n^{-k/2})$.

PROOF. Let C be the first k columns of $A\Pi$, i.e. the columns chosen by Greedy. Since, Q is orthogonal, we have

$$\begin{aligned} Vol(C) = Vol(R_{11}) &= \prod_{i=1}^k \sigma_i(R_{11}) \geq \prod_{i=1}^k \sigma_i(A)/(\sqrt{n-i}2^i) \\ &\geq 2^{-k(k-1)/2} \left(\prod_{i=1}^k \sigma_i(A)/n^{1/2} \right) \\ &\geq 2^{-k(k-1)/2} n^{-k/2} \left(\prod_{i=1}^k \sigma_i(A) \right) \\ &\geq 2^{-k(k-1)/2} n^{-k/2} \cdot Vol_{max} \end{aligned}$$

where Vol_{max} is the maximum possible volume a subset can attain.

This analysis is loose as the volume $\prod_{i=1}^k \sigma_i(A)$ may not be attainable using k columns of A . One major problem with this bound is that it has exponential dependence on n . Our (almost) tight analysis will provide an improvement on the approximation ratio of the theorem above in two ways: first, we will remove the dependence on n , and second we get better than quadratic dependence on k in the exponent. The outline of the remainder of this section is as follows: In Section 3.1, we analyze the performance ratio of Greedy. Section 3.2 presents an explicit example for which Greedy is bad. We analyze Greedy for MaxSubset(V) in Section 3.3 where we require the columns of the matrix be unit norm, in which case the volume is monotonically non-increasing or non-decreasing in the number of vectors chosen by any algorithm.

3.1. Approximation Ratio of Greedy

We consider Greedy after k steps. First, we assume that the dimension of the space spanned by the column vectors in A is at least k , since otherwise there is nothing to prove. Let $span(S)$ denote the space spanned by the vectors in the set S and let $\pi_S(v)$ be the projection of v onto $span(S)$. In this section, let $d(v, S) = \|v - \pi_S(v)\|$ be the norm of the part of v orthogonal to $span(S)$. Let $V_k = \{v_1, \dots, v_k\}$ be the set of vectors in order that have been chosen by the greedy algorithm at the end of the k^{th} step. Let

$W_k = \{w_1, \dots, w_k\}$ be a set of k vectors of maximum volume. Our main result in this subsection is the following theorem:

Theorem 3.3. $Vol(V_k) \geq 1/k! \cdot Vol(W_k)$.

We prove the theorem through a sequence of lemmas. The basic idea is to show that at the j^{th} step, Greedy loses a factor of at most j to the optimal. Theorem 3.3 then follows by an elementary induction. First, define $\alpha_i = \pi_{(V_{k-1})}(w_i)$ for $i = 1, \dots, k$. α_i is the projection of w_i onto $span(V_{k-1})$ where $V_{k-1} = \{v_1, \dots, v_{k-1}\}$. Let $\beta_i = w_i - \pi_{(V_{k-1})}(w_i)$. Hence, we have

$$w_i = \alpha_i + \beta_i \quad \text{for } i = 1, \dots, k. \quad (2)$$

Note that the dimension of $span(V_{k-1})$ is $k - 1$, which means that the α_i 's are linearly dependent. We will need some stronger properties of the α_i 's.

Definition 3.4. A set of m vectors is said to be in general position, if they are linearly dependent and any $m - 1$ element subset of them are linearly independent.

It's immediate from Definition 3.4 that

Remark 3.5. Let $U = \{\gamma_1, \dots, \gamma_m\}$ be a set of m vectors in general position. Then, γ_i can be written as a linear combination of the other vectors in U , i.e.

$$\gamma_i = \sum_{l \neq i} \lambda_l^i \gamma_l \quad (3)$$

for $i = 1, \dots, m$. λ_l^i 's are the coefficients of γ_l in the expansion of γ_i .

Lemma 3.6. Let $U = \{\gamma_1, \dots, \gamma_m\}$ be a set of m vectors in general position. Then, there exists a γ_i such that $|\lambda_j^i| \leq 1$ for all $j \neq i$.

PROOF. Assume, without loss of generality that $A = \{\gamma_2, \gamma_3, \dots, \gamma_m\}$ has the greatest volume among all possible $m - 1$ element subsets of U . We claim that γ_1 has the desired property. Consider the set $B_j = \{\gamma_1, \dots, \gamma_{j-1}, \gamma_{j+1}, \dots, \gamma_m\}$ for $2 \leq j \leq m$. Let $C_j = A - \{\gamma_j\} = B_j - \{\gamma_1\}$. Then, since A has the greatest volume, $Vol(A) = Vol(C_j) \cdot d(\gamma_j, C_j) \geq Vol(B_j) = Vol(C_j) \cdot d(\gamma_1, C_j)$. Hence, we have $d(\gamma_j, C_j) \geq d(\gamma_1, C_j)$. Then, using (3), we can write

$$\gamma_1 = \lambda_j^1 \gamma_j + \sum_{l \neq j, l \neq 1} \lambda_l^1 \gamma_l \quad (4)$$

Denoting $\delta_j = \pi_{C_j}(\gamma_j)$ and $\theta_j = \gamma_j - \delta_j$, (4) becomes

$$\gamma_1 = \left(\lambda_j^1 \delta_j + \sum_{l \neq j, l \neq 1} \lambda_l^1 \gamma_l \right) + \lambda_j^1 \theta_j$$

where the term in parentheses is in $\text{span}(C_j)$. Hence, the part of γ_1 which is not in $\text{span}(C_j)$, $\theta_1 = \gamma_1 - \pi_{C_j}(\gamma_1) = \lambda_j^1 \theta_j$ and so $\|\theta_1\| = |\lambda_j^1| \|\theta_j\|$. Note that $\|\theta_1\| = d(\gamma_1, C_j)$ and $\|\theta_j\| = d(\gamma_j, C_j)$, so $d(\gamma_1, C_j) = |\lambda_j^1| d(\gamma_j, C_j)$. Since $d(\gamma_1, C_j) \leq d(\gamma_j, C_j)$, we have $|\lambda_j^1| \leq 1$.

Lemma 3.7. *If $\|\alpha_i\| > 0$ for $i = 1, \dots, k$ and $k \geq 2$, then there exists a set of m vectors $U = \{\alpha_{i_1}, \dots, \alpha_{i_m}\} \subseteq \{\alpha_1, \dots, \alpha_k\}$ with $m \geq 2$ that are in general position.*

PROOF. Note that the cardinality of a set U with the desired properties should be at least 2, since otherwise there is nothing to prove. We argue by induction on k . For the base case $k = 2$, we have two vectors α_1 and α_2 spanning a 1-dimensional space and clearly any one of them is linearly independent since neither is 0. Assume that, as the induction hypothesis, any set of $k \geq 2$ non-zero vectors $\{\alpha_1, \dots, \alpha_k\}$ spanning at most a $k - 1$ dimensional space has a non-trivial subset in general position. Consider a $k + 1$ element set $A = \{\alpha_1, \dots, \alpha_{k+1}\}$ with $\dim(\text{span}(A)) \leq k$. If the vectors in A are not in general position, then there is a k element subset A' of A which is linearly dependent. Hence, $\dim(\text{span}(A')) \leq k - 1$ for which, by the induction hypothesis, we know that there exists a non-trivial subset in general position.

The existence of a subset in general position guaranteed by Lemma 3.7 will be needed when we apply the next lemma.

Lemma 3.8. *Assume $\|\alpha_i\| > 0$ for $i = 1, \dots, k$. Then, there exists an α_{i_j} such that $d(\alpha_{i_j}, W'_{k-1}) \leq (m - 1) \cdot d(v_k, V_{k-1})$, where $W'_{k-1} = W_k - \{w_{i_j}\}$.*

PROOF. Let $U = \{\alpha_{i_1}, \dots, \alpha_{i_m}\} \subseteq \{\alpha_1, \dots, \alpha_k\}$ be in general position where $m \geq 2$ (the existence of U is given by Lemma 3.7). Assume α_{i_1} has the property given by Lemma 3.6. Let $U' = \{w_{i_2}, \dots, w_{i_m}\}$. We claim that α_{i_1} has the desired property. First, note that $d(\alpha_{i_1}, W'_{k-1}) \leq d(\alpha_{i_1}, U')$, since $\text{span}(U')$ is a subspace of $\text{span}(W'_{k-1})$. We seek a bound on $d(\alpha_{i_1}, W'_{k-1})$. Using (3) and (2), we have

$$\alpha_{i_1} = \sum_{l \neq 1} \lambda_{i_l}^1 \alpha_{i_l} = \sum_{l \neq 1} \lambda_{i_l}^1 (w_{i_l} - \beta_{i_l}).$$

where α_{i_l} 's are the vectors in U and β_{i_l} 's are their orthogonal parts. Rearranging,

$$\sum_{l \neq 1} \lambda_{i_l}^1 \beta_{i_l} = \left(\sum_{l \neq 1} \lambda_{i_l}^1 w_{i_l} \right) - \alpha_{i_1}.$$

Note that the right hand side is an expression for the difference between a vector in $\text{span}(U')$ and α_{i_1} . Hence,

$$\begin{aligned}
d(\alpha_{i_1}, W'_{k-1}) &\leq d(\alpha_{i_1}, U') = \min_{v \in \text{span}(U')} \|v - \alpha_{i_1}\| \\
&\leq \left\| \sum_{l \neq 1} \lambda_{i_l}^1 w_{i_l} - \alpha_{i_1} \right\| \\
&= \left\| \sum_{l \neq 1} \lambda_{i_l}^1 \beta_{i_l} \right\| \\
&\leq \sum_{l \neq 1} \lambda_{i_l}^1 \|\beta_{i_l}\| \\
&\leq (m-1) \cdot \max_{1 \leq l \leq m} \|\beta_{i_l}\| \\
&\leq (m-1) \cdot d(v_k, V_{k-1}).
\end{aligned}$$

where the last two inequalities follow from Lemma 2.1 and the greedy property of the algorithm, respectively.

Before stating the final lemma, which gives the approximation factor of Greedy at each round, we need the following simple observation.

Lemma 3.9. *Let u be a vector, V and W be subspaces and $\alpha = \pi_V(u)$. Then $d(u, W) \leq d(u, V) + d(\alpha, W)$.*

PROOF. Let $\gamma = \pi_W(\alpha)$. By triangle inequality for vector addition, we have

$$\|u - \gamma\| \leq \|u - \alpha\| + \|\alpha - \gamma\| = d(u, V) + d(\alpha, W). \text{ The result follows since } d(u, W) \leq \|u - \gamma\|.$$

Lemma 3.10. *At the k^{th} step of Greedy, there exists a w_i such that $d(w_i, W'_{k-1}) \leq k \cdot d(v_k, V_{k-1})$ where $W'_{k-1} = W_k - \{w_i\}$.*

PROOF. For $k = 1$, there's nothing to prove. For $k \geq 2$, there are two cases.

1. One of the w_i 's is orthogonal to V_{k-1} ($\|\alpha_i\| = 0$). In this case, by the greedy property, $d(v_k, V_{k-1}) \geq \|w_i\| \geq d(w_i, W'_{k-1})$, which gives the result.
2. For all w_i , $\|\alpha_i\| > 0$, i.e., all w_i have non-zero projection on V_{k-1} . Assuming that $\alpha_1 = \pi_{V_{k-1}}(w_1)$ has the desired property proved in Lemma 3.8, we have for the corresponding w_1

$$\begin{aligned}
d(w_1, W'_{k-1}) &\leq d(w_1, V_{k-1}) + d(\alpha_1, W'_{k-1}) \\
&\leq \|\beta_1\| + d(\alpha_1, W'_{k-1}) \\
&\leq \|\beta_1\| + (m-1) \cdot d(v_k, V_{k-1}) \\
&\leq m \cdot d(v_k, V_{k-1}).
\end{aligned}$$

The first inequality is due to Lemma 3.9. The last inequality follows from the greedy property of the algorithm, i.e. the fact that $d(v_k, V_{k-1}) \geq \|\beta_1\|$. The lemma follows since $m \leq k$.

The last lemma immediately leads to the result of Theorem 3.3, with a simple inductive argument as follows:

PROOF. The base case is easily established since $\text{Vol}(V_1) = \text{Vol}(W_1)$. Assume that $\text{Vol}(V_{k-1}) \geq 1/(k-1)! \cdot \text{Vol}(W_{k-1})$ for some $k > 2$. By Lemma 3.10, we have a w_i such that $d(w_i, W'_{k-1}) \leq k \cdot d(v_k, V_{k-1})$ where $W'_{k-1} = W_k - \{w_i\}$. It follows that

$$\begin{aligned} \text{Vol}(V_k) &= d(v_k, V_{k-1}) \cdot \text{Vol}(V_{k-1}) \\ &\geq \frac{d(w_i, W'_{k-1})}{k} \cdot \frac{\text{Vol}(W_{k-1})}{(k-1)!} \\ &\geq \frac{d(w_i, W'_{k-1})}{k!} \cdot \text{Vol}(W'_{k-1}) \\ &= \frac{\text{Vol}(W_k)}{k!}. \end{aligned}$$

3.2. Lower Bound for Greedy

We give a lower bound of $1/2^{k-1}$ for the approximation factor of Greedy by explicitly constructing a bad example. We will inductively construct a set of unit vectors satisfying this lower bound. It will be the case that the space spanned by the vectors in the optimal solution is the same as the space spanned by the vectors chosen by Greedy. An interesting property of our construction is that both the optimal volume and the volume of the vectors chosen by Greedy approach 0 in the limit of a parameter δ , whereas their ratio approaches to $1/2^{k-1}$.

We will first consider the base case $k = 2$: let the matrix $A = [v_1 w_1 w_2]$ where $\dim(A) = 2$ and $d(v_1, w_1) = d(v_1, w_2) = \delta$ for some $1 > \delta > 0$ such that θ , the angle between w_1 and w_2 is twice the angle between v_1 and w_1 , i.e. v_1 is ‘between’ w_1 and w_2 . If the greedy algorithm first chooses v_1 , then $\lim_{\delta \rightarrow 0} \text{Vol}(V_2)/\text{Vol}(W_2) = 1/2 \cos \theta/2 = 1/2$. Hence, for $k = 2$, there’s a set of vectors for which $\text{Vol}(W_2) = (2 - \epsilon) \cdot \text{Vol}(V_2)$ for arbitrarily small $\epsilon > 0$.

For arbitrarily small $\epsilon > 0$, assume that there is an optimal set of k vectors $W_k = \{w_1, \dots, w_k\}$ such that $\text{Vol}(W_k) = (1 - \epsilon)2^{k-1} \cdot \text{Vol}(V_k)$ where $V_k = \{v_1, \dots, v_k\}$ is the set of k vectors chosen by Greedy. The vectors in W_k and V_k span a subspace of dimension k , and assume $w_i \in \mathbb{R}^d$ where $d > k$. Let $d(v_2, V_1) = \epsilon_1 = \delta$ for some $1 > \delta > 0$, and $d(v_{i+1}, V_i) = \epsilon_i = \delta \epsilon_{i-1}$ for $i = 2, \dots, k-1$. Thus, $\text{Vol}(V_k) = \delta^{k(k-1)/2}$ and $\text{Vol}(W_k) = (1 - \epsilon)2^{k-1} \delta^{k(k-1)/2}$. Assume further that for all w_i in W_k , $d(w_i, V_j) \leq \epsilon_j$ for $j = 1, \dots, k-2$ and $d(w_i, V_{k-1}) = \epsilon_{k-1}$ so that there exists an execution of Greedy where no $\{w_1, \dots, w_k\}$ is chosen.

We will now construct a new set of vectors $W_{k+1} = W'_k \cup \{w_{k+1}\} = \{w'_1, \dots, w'_k, w_{k+1}\}$ which will be the optimal solution. Let $w'_i = \pi_{V_j}(w_i)$, and let $e_i^j = \pi_{V_j}(w_i) - \pi_{V_{j-1}}(w_i)$ for $j = 2, \dots, k$ and $e_i^1 = w_i^1$. Namely, e_i^j is the component of w_i which is in V_j , but perpendicular to V_{j-1} and e_i^1 is the component of w_i which is in the span of v_1 . (Note that $\|e_i^k\| = \epsilon_{k-1}$.) Let u be a unit vector perpendicular to $\text{span}(W_k)$. For each w_i we define a new vector $w'_i = (\sum_{j=1}^{k-1} e_i^j) + \sqrt{1 - \delta^2} e_i^k + \delta \epsilon_{k-1} u$. Intuitively, we are defining a set of new

vectors which are first rotated towards V_{k-1} and then towards u such that they are $\delta\epsilon_{k-1}$ away from V_k . Introduce another vector $w_{k+1} = \sqrt{1-\delta^2}v_1 - \delta\epsilon_{k-1}u$. Intuitively, this new vector is v_1 rotated towards the negative direction of u . Note that, in this setting $\epsilon_k = \delta\epsilon_{k-1}$. We finally choose $v_{k+1} = w_{k+1}$.

Lemma 3.11. *For any $w \in W_{k+1}$, $d(w, V_j) \leq \epsilon_j$ for $j = 1, \dots, k-1$ and $d(w, V_k) = \epsilon_k$.*

PROOF. For $w = w_{k+1}$, $d(w_{k+1}, V_j) = \epsilon_k \leq \epsilon_j$ for $j = 1, \dots, k$. Let $w = w'_i$ for some $1 \leq i \leq k$. Then, for any $1 \leq j \leq k-1$, we have $d(w'_i, V_j)^2 = \sum_{l=j+1}^{k-1} \|e_i^l\|^2 + (1-\delta^2)\|e_i^k\|^2 + \delta^2\|e_i^k\|^2 = \sum_{l=j+1}^k \|e_i^l\|^2 = d(w_i, V_j)^2 \leq \epsilon_j^2$ by the induction hypothesis.

Lemma 3.11 ensures that $\{v_1, \dots, v_{k+1}\}$ is a valid output of Greedy. What remains is to show that for any $\epsilon > 0$, we can choose δ sufficiently small so that $\text{Vol}(W_{k+1}) \geq (1-\epsilon)2^k \cdot \text{Vol}(W_k)$. In order to show this, we will need the following lemmas.

Lemma 3.12. $\lim_{\delta \rightarrow 0} \text{Vol}(W_{k+1}) = 2\epsilon_k \cdot \text{Vol}(W_k)$.

PROOF. With a little abuse of notation, let W_{k+1} denote the matrix of coordinates for the vectors in the set W_{k+1} .

$$W_{k+1} = \begin{pmatrix} w_{1,1} & w_{1,2} & \cdots & w_{1,k} & \sqrt{1-\delta^{2k}} \\ w_{2,1} & w_{2,2} & \cdots & w_{2,k} & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \sqrt{1-\delta^2}w_{k,1} & \sqrt{1-\delta^2}w_{k,2} & \cdots & \sqrt{1-\delta^2}w_{k,k} & 0 \\ \delta^k & \delta^k & \cdots & \delta^k & -\delta^k \end{pmatrix}$$

where $w_{i,j}$ is the i^{th} coordinate of w_j , which is in W_k . (Note that this is exactly how U is constructed in the inductive step). Expanding on the right-most column of the matrix, we have

$$\text{Vol}(W_{k+1}) = |\det(W_{k+1})| = |\sqrt{1-\delta^{2k}} \cdot \det(A) + (-1)^{k+1}\delta^k \cdot \det(B)| \quad (5)$$

where A and B are the corresponding minors of the coefficients, i.e. the left-most lower and upper $k \times k$ sub-matrices of W_{k+1} , respectively. Clearly, we have $\det(B) = \sqrt{1-\delta^2} \cdot \det(W_k)$ where W_k is the matrix of coordinates for the vectors in the set W_k . Let C be the matrix obtained by replacing each $w_{1,i}$ by 1 in W_k . Then, using row interchange operations on A , we can move the last row of A to the top. This gives a sign change of $(-1)^{k-1}$. Then, factoring out $\sqrt{1-\delta^2}$ and δ^k in the first and last rows respectively, we have $\det(A) = (-1)^{k-1}\delta^k\sqrt{1-\delta^2} \cdot \det(C)$. Hence, (5) becomes

$$|\det(W_{k+1})| = (\delta^k\sqrt{1-\delta^2})|\sqrt{1-\delta^{2k}} \cdot \det(C) + \det(W_k)| \quad (6)$$

We will need the following lemma to compare $\det(W_k)$ and $\det(C)$.

Lemma 3.13. $\lim_{\delta \rightarrow 0} \det(C)/\det(W_k) = 1$.

PROOF. For $i > 1$, the elements of the i^{th} rows of both W_k and C has δ^{i-1} as a common coefficient by construction. Factoring out these common coefficients, we have $\det(W_k) = \delta^{k(k-1)/2} \cdot \det(U)$ and $\det(C) = \delta^{k(k-1)/2} \cdot \det(U')$ where U and U' are matrices with non-zero determinants as δ approaches 0. Furthermore, $\lim_{\delta \rightarrow 0} \det(U) = \det(U')$ as the elements in the first row of U approaches 1. The result then follows.

Using Lemma 3.13 and (6), we have

$$\lim_{\delta \rightarrow 0} \text{Vol}(W_{k+1}) = \lim_{\delta \rightarrow 0} |\det(W_{k+1})| = 2\delta^k |\det(W_k)| = 2\epsilon_k \cdot \text{Vol}(W_k)$$

Theorem 3.14. $\text{Vol}(W_{k+1}) \geq (1 - \epsilon)2^k \cdot \text{Vol}(V_{k+1})$ for arbitrarily small $\epsilon > 0$.

PROOF. Given any $\epsilon' > 0$ we can choose δ small enough so that $\text{Vol}(W_{k+1}) \geq 2\epsilon_k(1 - \epsilon') \cdot \text{Vol}(W_k)$, which is always possible by Lemma 3.12. Given any ϵ'' , we can apply induction hypothesis to obtain V_k and W_k such that $\text{Vol}(W_k) \geq (1 - \epsilon'')2^{k-1} \cdot \text{Vol}(V_k)$. Thus,

$$\begin{aligned} \text{Vol}(W_{k+1}) &\geq 2\epsilon_k(1 - \epsilon') \cdot \text{Vol}(W_k) \\ &\geq 2\epsilon_k(1 - \epsilon')(1 - \epsilon'')2^{k-1} \cdot \text{Vol}(V_k) \\ &= (1 - \epsilon')(1 - \epsilon'')2^k \cdot \text{Vol}(V_{k+1}), \end{aligned}$$

where we have used $\text{Vol}(V_{k+1}) = \epsilon_k \cdot \text{Vol}(V_k)$. Choosing ϵ' and ϵ'' small enough such that $(1 - \epsilon')(1 - \epsilon'') > 1 - \epsilon$ gives the result.

3.3. Maximizing the Number of Unit norm Vectors Attaining A Given Volume

In this section, we give a result on approximating the maximum number of unit norm vectors which can be chosen to have at least a certain volume. This result is essentially a consequence of the previous approximation result. We assume that all the vectors in A have unit norm, hence the volume is non-increasing in the number of vectors chosen by Greedy. Let OPT_k denote the optimal volume for k vectors. Note that $OPT_k \geq OPT_{k+1}$ and the number of vectors m , chosen by Greedy attaining volume at least OPT_k is not greater than k . Our main result states that, if the optimal volume of k vectors is $2^{-\Omega(k)}$, then Greedy chooses $\Omega(k/\log k)$ vectors having at least that volume. Thus, Greedy gives a $\log k$ -approximation to the optimal number of vectors. We prove the result through a sequence of lemmas. The following lemma is an immediate consequence of applying Greedy on W_k .

Lemma 3.15. *Let $W_k = \{w_1, \dots, w_k\}$ be a set of k vectors of optimal volume OPT_k . Then there exists a permutation π of the vectors in W_k such that $d_{\pi(k)} \leq d_{\pi(k-1)} \leq \dots \leq d_{\pi(2)}$ where $d_{\pi_i} = d(w_{\pi_i}, \{w_{\pi_1}, \dots, w_{\pi_{i-1}}\})$ for $k \geq i \geq 2$.*

We use this existence result to prove the following lemma.

Lemma 3.16. $OPT_m \geq (OPT_k)^{(m-1)/(k-1)}$ where $m \leq k$.

PROOF. Let $W_k = \{w_1, \dots, w_k\}$ be a set of vectors of optimal volume OPT_k . By Lemma 3.15, we know that there exists an ordering of vectors in W_k such that $d_{\pi(k)} \leq d_{\pi(k-1)} \leq \dots \leq d_{\pi(2)}$ where $d_{\pi_i} = d(w_{\pi_i}, \{w_{\pi_1}, \dots, w_{\pi_{i-1}}\})$ for $k \geq i \geq 2$. Let $W_m' = \{w_{\pi(1)}, \dots, w_{\pi(m)}\}$. Then, we have $OPT_m \geq Vol(W_m') = \prod_{i=2}^m d_{\pi_i} \geq (\prod_{i=2}^k d_{\pi_i})^{(m-1)/(k-1)} = (OPT_k)^{(m-1)/(k-1)}$.

Lemma 3.17. Suppose $OPT_k \leq 2^{(k-1)m \log m / (m-k)}$. Then, the greedy algorithm chooses at least m vectors whose volume is at least OPT_k .

PROOF. We are seeking a condition for OPT_k which will provide a lower bound for m such that $OPT_m/m! \geq OPT_k$. If this holds, then $Vol(Greedy_m) \geq Opt_m/m! \geq OPT_k$ and so Greedy can choose at least m vectors which have volume at least OPT_k . It suffices to find such an m satisfying $(OPT_k)^{(m-1)/(k-1)}/m! \geq OPT_k$ by Lemma 3.16. This amounts to $1/m! \geq (OPT_k)^{1-(m-1)/(k-1)}$. Since $1/m! \geq 1/m^m$ for $m \geq 1$, we require $1/m^m \geq (OPT_k)^{1-(m-1)/(k-1)}$. Taking logarithms of both sides and rearranging, we have $-(k-1)m \log m / (k-m) \geq \log OPT_k$. Taking exponents of both sides yields $2^{(k-1)m \log m / (m-k)} \geq OPT_k$.

In order to interpret this result, we will need to restrict OPT_k . Otherwise, for example if $OPT_k = 1$, the greedy algorithm may never get more than 1 vector to guarantee a volume of at least OPT_k since it might be possible to misguess the first vector. In essence, the number of vectors chosen by the algorithm depends on OPT_k . First, we discuss what is a reasonable condition on OPT_k . Consider n vectors in m dimensions which defines a point in $\mathbb{R}^{m \times n}$. The set of points in which any two vectors are orthogonal has measure 0. Thus, define $2^{-\alpha} = \max_{ij} d(v_i, v_j)$. Then, it is reasonable to assume that $\alpha > 0$, in which case $OPT_k \leq 2^{-\alpha k} = 2^{-\Omega(k)}$. Hence, we provide the following theorem which follows from the last lemma under the reasonable assumption that the optimal volume decreases by at least a constant factor with the addition of one more vector.

Theorem 3.18. If $OPT_k \leq 2^{-\Omega(k)}$, then the greedy algorithm chooses $\Omega(k/\log k)$ vectors having volume at least OPT_k .

PROOF. For some α , $OPT_k \leq 2^{-\alpha k}$. Thus, we solve for m such that $2^{-\alpha k} \leq 2^{(k-1)m \log m / (m-k)}$. Suitable rearrangements yield

$$m \leq \frac{\alpha k(k-m)}{(k-1) \log m} \leq \frac{2\alpha k}{\log m}$$

.

For m , the largest integer such that $m \leq 2\alpha k / \log m$, we have

$$m \approx \frac{2\alpha k}{\log(2\alpha k / \log m)} = \frac{2\alpha k}{\log(2\alpha k) - \log \log m} = \Omega\left(\frac{k}{\log k}\right).$$

In reality, for a random selection of n vectors in m dimensions, α will depend on n and so the result is not as strong as it appears.

4. Discussion

Our analysis of the approximation ratio relies on finding the approximation factor at each round of Greedy. Indeed, we have found examples for which the volume of the vectors chosen by Greedy falls behind the optimal volume by as large a factor as $1/k$, making Lemma 3.10 tight. But it might be possible to improve the analysis by correlating the ‘gains’ of the algorithm between different steps. Hence, one of the immediate questions is that whether one can close the gap between the approximation ratio and the lower bound for Greedy. We conjecture that the approximation ratio is $1/2^{k-1}$.

We list other open problems as follows:

- Do there exist efficient non-greedy algorithms with better guarantees for MAX-VOL?
- There is a huge gap between the approximation ratio of the algorithm we have analyzed and the inapproximability result. Can this gap be closed on the inapproximability side by using more advanced techniques?
- Volume seems to play an important role in constructing a low-rank approximation to a matrix. Can the result of [5] be extended to yield a direct relationship between low-rank approximations and large volume $m \times k$ sub-matrices of a matrix? Or, can we establish a result stating that there must exist a large volume $k \times k$ sub-matrix of a large volume $m \times k$ sub-matrix such that one can find an approximation to the maximum volume $k \times k$ sub-matrix by running the same algorithm again on the $m \times k$ sub-matrix? Solutions proposed for low-rank approximations thus far consider only randomized algorithms. Can this work be extended to find a deterministic algorithm for matrix reconstruction? Establishing the relationship between the maximum $k \times k$ sub-matrix and some $k \times k$ sub-matrix of a maximum volume $m \times k$ sub-matrix would give a deterministic algorithm for matrix reconstruction with provable guarantees.

We would like to note that the approximation ratio of Greedy algorithm is considerably small because of the ‘multiplicative’ nature of the problem. Another important problem which resembles MAX-VOL in terms of behavior (but not necessarily in nature) is the Shortest Vector Problem (SVP), which is not known to have a polynomial factor approximation algorithm. Indeed, the most common algorithm which works well in practice has a $2^{O(n)}$ approximation ratio [19] and non-trivial hardness results for this problem are difficult to find.

Acknowledgments: We would like to thank Christos Boutsidis for useful discussions and pointing out Theorem 3.1, and the anonymous referees for their helpful comments.

References

- [1] A. Deshpande, L. Rademacher, S. Vempala, G. Wang, Matrix approximation and projective clustering via volume sampling, in: Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, 2006, pp. 1117–1126.
- [2] P. Drineas, R. Kannan, M. W. Mahoney, Fast monte carlo algorithms for matrices II: Computing a low-rank approximation to a matrix, *SIAM Journal on Computing* 36 (1) (2006) 158–183.
- [3] P. Drineas, R. Kannan, M. W. Mahoney, Fast Monte Carlo algorithms for matrices III: Computing a compressed approximate matrix decomposition, *SIAM Journal on Computing* 36 (1) (2006) 184–206.
- [4] A. Deshpande, L. Rademacher, S. Vempala, G. Wang, Matrix approximation and projective clustering via volume sampling, in: Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, ACM Press, 2006, pp. 1117–1126.
- [5] S. A. Goreinov, E. E. Tyrtshnikov, The maximal-volume concept in approximation by low-rank matrices, in: Contemporary Mathematics, Vol. 280, AMS, 2001, pp. 47–51.
- [6] G. H. Golub, C. V. Loan, *Matrix Computations*, Johns Hopkins U. Press, 1996.
- [7] T. F. Chan, Rank revealing QR factorizations, *Linear Algebra and Its Applications* (88/89) (1987) 67–82.
- [8] P. A. Businger, G. H. Golub, Linear least squares solutions by Householder transformations, *Numerische Mathematik* (7) (1965) 269–276.
- [9] W. Kahan, Numerical linear algebra, *Canadian Mathematical Bulletin* 9 (1966) 757–801.
- [10] Y. P. Hong, C. T. Pan, Rank-revealing QR factorizations and the singular value decomposition, *Mathematics of Computation* 58 (1992) 213–232.
- [11] I. C. F. Ipsen, C. T. Kelley, S. R. Pope, Rank-deficient nonlinear least squares problems and subset selection, *SIAM Journal on Numerical Analysis* 49 (2011) 1244–1266.
- [12] C. T. Pan, P. T. P. Tang, Bounds on singular values revealed by QR factorizations, *BIT Numerical Mathematics* 39 (1999) 740–756.
- [13] M. Gu, S. C. Eisenstat, Efficient algorithms for computing a strong rank-revealing QR factorization, *SIAM Journal on Scientific Computing* 17 (4) (1996) 848–869. doi:<http://dx.doi.org/10.1137/0917055>.
- [14] F. R. de Hoog, R. M. M. Mattheijb, Subset selection for matrices, *Linear Algebra and its Applications* (422) (2007) 349–359.

- [15] G. H. Golub, V. Klema, G. W. Stewart, Rank degeneracy and least squares problems, Tech. rep., Dept. of Computer Science, Univ. of Maryland (1976).
- [16] C. T. Pan, On the existence and computation of rank-revealing LU factorizations, *Linear Algebra and its Applications* 316 (1-3) (2000) 199–222.
- [17] M. R. Garey, D. S. Johnson, *Computers and Intractability*, W. H. Freeman, 1979.
- [18] R. M. Karp, Reducibility among combinatorial problems, in: R. E. Miller, J. W. Thatcher (Eds.), *Complexity of Computer Computations*, Plenum Press, 1972, pp. 85–103.
- [19] A. K. Lenstra, H. W. Lenstra, L. Lovasz, Factoring polynomials with rational coefficients, *Mathematische Annalen* (261) (1982) 515–534.