

Optimal ℓ_1 Column Subset Selection and a Fast PTAS for Low Rank Approximation

Arvind V. Mahankali
 CMU
 amahanka@andrew.cmu.edu

David P. Woodruff
 CMU
 dwoodruf@cs.cmu.edu

Abstract

We study the problem of entrywise ℓ_1 low rank approximation. We give the first polynomial time column subset selection-based ℓ_1 low rank approximation algorithm sampling $\tilde{O}(k)$ columns and achieving an $\tilde{O}(k^{1/2})$ -approximation for any k , improving upon the previous best $O(k)$ -approximation and matching a prior lower bound for column subset selection-based ℓ_1 -low rank approximation which holds for any $\text{poly}(k)$ number of columns. We extend our results to obtain tight upper and lower bounds for column subset selection-based ℓ_p low rank approximation for any $1 < p < 2$, closing a long line of work on this problem.

We next give a $(1 + \varepsilon)$ -approximation algorithm for entrywise ℓ_p low rank approximation of an $n \times d$ matrix, for $1 \leq p < 2$, that is not a column subset selection algorithm. First, we obtain an algorithm which, given a matrix $A \in \mathbb{R}^{n \times d}$, returns a rank- k matrix \hat{A} in $2^{\text{poly}(k/\varepsilon)} + \text{poly}(nd)$ running time that achieves the following guarantee:

$$\|A - \hat{A}\|_p \leq (1 + \varepsilon) \cdot OPT + \frac{\varepsilon}{\text{poly}(k)} \|A\|_p$$

where $OPT = \min_{A_k \text{ rank } k} \|A - A_k\|_p$. Using this algorithm, in the same running time we give an algorithm which obtains error at most $(1 + \varepsilon) \cdot OPT$ and outputs a matrix of rank at most $3k$ — these algorithms significantly improve upon all previous $(1 + \varepsilon)$ - and $O(1)$ -approximation algorithms for the ℓ_p low rank approximation problem, which required at least $n^{\text{poly}(k/\varepsilon)}$ or $n^{\text{poly}(k)}$ running time, and either required strong bit complexity assumptions (our algorithms do not) or had bicriteria rank $3k$. Finally, we show hardness results which nearly match our $2^{\text{poly}(k)} + \text{poly}(nd)$ running time and the above additive error guarantee.

Contents

1	Introduction	1
1.1	Our Results	2
1.1.1	ℓ_p Column Subset Selection for $1 \leq p < 2$	2
1.1.2	General ℓ_p Low Rank Approximation for $1 \leq p < 2$	4
1.1.3	Hardness for ℓ_p Low Rank Approximation with Additive Error - Appendix A	4
1.2	Our Techniques	5
1.2.1	ℓ_1 Column Subset Selection	5
1.2.2	General ℓ_1 Low Rank Approximation	7
1.3	Paper Outline	15
1.3.1	Main Results - Algorithms	15
1.3.2	Appendices - Additional Results	15
2	Optimal ℓ_1 Column Subset Selection via Random Sampling	17
2.1	Preliminaries: Notation and Lewis Weight Sampling	17
2.2	ℓ_p Column Subset Selection Algorithm and Analysis	17
3	$(1 + \varepsilon)$-Approximation in FPT Time with Bicriteria Rank $3k$	22
3.1	Preliminaries: Median-Based Estimator for ℓ_p -norm Dimension Reduction from [2]	22
3.2	$(1 + \varepsilon)$ -Approximation Algorithm and Analysis	25
A	Hardness for ℓ_p Low Rank Approximation with Additive Error, based on [2]	35
A.1	Background: Small Set Expansion Hypothesis	35
A.2	Background: Hardness Proof from [2] — ℓ_p Low Rank Approximation for $p \in (1, 2)$	35
A.3	Deterministic Embedding from ℓ_p^n to $\ell_1^{n^{O(\log n)}}$	36
A.4	Hardness for ℓ_p and $\ell_{1,p}$ Low Rank Approximation with Additive Error, $p \in (1, 2)$	44
A.5	Hardness for Constrained ℓ_1 Low Rank Approximation with Additive Error	49
B	ℓ_p Column Subset Selection Lower Bound	51
C	poly(k)-Approximation Algorithms with Bicriteria Rank Independent of n and d	58
C.1	A poly(k)-Approximation with poly(k) Bicriteria Rank	58
C.2	Reducing the Rank to At Most k	61

1 Introduction

Low rank approximation is one of the most fundamental problems in data science and randomized numerical linear algebra. In this problem, one is given an $n \times d$ matrix A and a rank parameter k , and one would like to approximately decompose A as $U \cdot V$, where $U \in \mathbb{R}^{n \times k}$ and $V \in \mathbb{R}^{k \times d}$ are the low rank factors. This gives a form of compression, since rather than storing the nd parameters needed to represent A , we can store only $(n + d)k$ parameters to store the low rank factors. One can also multiply $U \cdot V$ by a vector in $O((n + d)k)$ time, rather than the $O(nd)$ time needed for multiplication by A . The singular value decomposition (SVD) can be used to find the best low rank approximation of A with respect to the sum of squares of differences, i.e., the Frobenius norm error measure, but this measure is often not robust enough in applications, since the need to fit single large outliers in A is exacerbated by the squared error measure, causing U and V to overfit such outliers and not capture enough of the remaining entries of A .

To overcome this, a large body of work has studied other, more robust error measures, with a notable one being the entrywise ℓ_1 -low rank approximation problem: given $A \in \mathbb{R}^{n \times d}$, a rank parameter k , and an approximation parameter $\alpha \geq 1$, find $U \in \mathbb{R}^{n \times k}$ and $V \in \mathbb{R}^{k \times d}$ so that:

$$\|U \cdot V - A\|_1 \leq \alpha \cdot \min_{U' \in \mathbb{R}^{n \times k}, V' \in \mathbb{R}^{k \times d}} \|U'V' - A\|_1,$$

where for a matrix $B \in \mathbb{R}^{n \times d}$, $\|B\|_1 = \sum_{i=1}^n \sum_{j=1}^d |B_{i,j}|$ is its entrywise 1-norm. Since we take the sum of absolute values of differences of entries of A and corresponding entries of $U \cdot V$, this is often considered more robust than the Frobenius norm error measure, which takes the squared differences. There is a large body of work on applications of ℓ_1 matrix factorization, as well as practical improvements obtained by optimizing the ℓ_1 error measure rather than the more well-understood Frobenius norm error measure, in areas such as computer vision and machine learning. For instance, in [30] and later in [54] it was shown that optimizing the ℓ_1 -based objective above, or a regularized version of it, yields much better performance than an IRLS-based approach or other ℓ_2 -based methods on the structure-from-motion problem. Various other works mentioned below, also motivated by problems in machine learning and signal processing, developed heuristics for ℓ_1 low rank approximation and similar problems.

The ℓ_1 -low rank approximation problem, for $\alpha = 1$, was shown to be NP-hard in [22], and assuming the Exponential Time Hypothesis, it requires $2^{\Omega(1/\epsilon)}$ time for $\alpha = 1 + \epsilon$. While these results rule out extremely accurate solutions to the ℓ_1 -low rank approximation problem in polynomial time, they leave open the possibility of larger approximation factors. Such approximation factors are of considerable interest, and we note that a low rank approximation $U \cdot V$ corresponding to an approximation factor $\alpha \ll \sqrt{nd}$ for the ℓ_1 -low rank approximation problem may be much better in applications than a low rank approximation corresponding to an exact solution to the Frobenius norm error measure, since the error measures are incomparable¹. A number of heuristics were proposed for the ℓ_1 -low rank approximation problem in [29, 30, 31, 33, 54, 9, 7, 8, 39, 36, 37, 35, 42]. The first rigorous approximation factors were proven in [48], where it was shown that $\alpha = \text{poly}(k) \log n$ -approximation is achievable in $\text{poly}(ndk)$ time². Later, in [11] an alternative $\text{poly}(ndk)$ time algorithm with $\text{poly}(k \log(nd))$ -approximation factor was given which held for every entrywise ℓ_p norm, for any constant $p \geq 1$.

In fact, each of the algorithms above is, or can be used to obtain a $\text{poly}(dk)$ time *column subset selection* algorithm for the ℓ_1 -low rank approximation problem with a $\text{poly}(k \log(nd))$ approximation factor. Column subset selection is a special type of low rank approximation in which the left factor U corresponds to a subset of columns of A itself. Column subset selection is a widely studied and extremely important special case of low rank approximation (see, e.g., [6, 17, 5] and the references therein); it has a number of advantages - for example, if the columns of A are sparse, then the columns in the left factor U are also sparse. One might argue that the column subset selection problem, also known as feature selection (since the columns of U can be thought of as the important features), is sometimes more important than the low rank approximation problem itself. Due to these advantages, it has been of interest to determine how well column subset selection algorithms can do compared to the optimal low-rank approximation error, in the Frobenius norm [20, 17, 6], spectral norm [17, 6] and in the ℓ_1 norm [11, 49] and even for more general loss functions [51].

¹We note that the exact solution to the Frobenius norm error measure gives a \sqrt{nd} -approximation to ℓ_1 -low rank approximation by relating the entrywise 2-norm to the entrywise 1-norm.

²Note that the problem is symmetric in n and d , so if $d < n$, we can instead write this as $\text{poly}(k) \log d$.

Often for column subset selection, we allow for a *bicriteria* approximation, namely, for outputting a low rank matrix of rank r with r a bit larger than k . Bicriteria approximations are common in the low rank approximation literature [19, 21, 12], and correspond to the case when k is not known or is not a hard constraint; note that bicriteria approximations still capture the original compression and fast multiplication motivations of low rank approximation discussed above. Moreover, bicriteria approximations are necessary in the context of column subset selection for some norms; indeed, for Frobenius norm it is known [20] that with exactly k columns the best approximation factor possible is $\Theta(k)$, while with $O(k)$ columns an $O(1)$ approximation is possible.

A natural question is what the limits of column subset selection algorithms for ℓ_1 -low rank approximation are. It was shown in [51] that the $O(k \log(nd))$ -approximation can be improved to $O(k \log k)$ with bicriteria rank $r = O(k \log n)$, via a polynomial time algorithm. Moreover, if one is willing to spend $n^{O(k \log k)}$ time, it was shown in [48] how to obtain an approximation factor³ of $O(\sqrt{k \log k})$. From a purely combinatorial perspective, this is close to best possible as [48] shows that there exist matrices for which any subset of at most $\text{poly}(k)$ columns provides at best a $k^{1/2-\gamma}$ -approximation, for an arbitrarily small constant $\gamma > 0$. We note that for every $p \geq 2$, there are tight bounds on the size of the best column subset selection algorithms for entrywise ℓ_p -low rank approximation known⁴, and there are polynomial time algorithms achieving these (see Theorem 4.1 of [14]). For $1 \leq p \leq 2$, however, the best known upper bounds have an approximation factor of roughly $k^{1/p}$, while the known lower bounds are only $k^{1-1/p}$ for $1 < p < 2$ [14]. In addition, the lower bounds of [14] only hold when k columns are selected, and do not rule out smaller approximation factors through slightly larger bicriteria ranks. While these results impose a limit on the approximation error achievable through column subset selection, the current gap of an $O(k \log k)$ approximation factor versus a $k^{1/2-\gamma}$ lower bound for column subset selection algorithms for $p = 1$ has remained elusive:

Question 1: *What is the best approximation factor for column subset selection for ℓ_1 -low rank approximation achievable by a polynomial time algorithm? What about for $1 < p < 2$?*

While the above approximation factors are highly non-trivial, the lower bound of [48] rules out better than a $k^{1/2-\gamma}$ -approximation with column subset selection methods, for arbitrarily small constant $\gamma > 0$. A natural question is if one can improve these approximation factors to $O(1)$ or even $(1 + \varepsilon)$ in polynomial time. This was the main question underlying [2], which, motivated by the fact that k is often small, took a parameterized complexity approach and showed how to obtain a $(1 + \varepsilon)$ -approximation in $n^{\text{poly}(k/\varepsilon)}$ time under the assumption that the entries of the input matrix A are integers in the range $\{-\text{poly}(n), \dots, \text{poly}(n)\}$, or alternatively, outputting a matrix of rank $3k$ instead of k with the same approximation factor guarantee, while removing this assumption on the entries of A . This algorithm has several drawbacks though: (1) the $n^{\Omega(k)}$ running time even for constant ε is prohibitive and makes the algorithm super-polynomial time for any $k = \omega(1)$, i.e., if k is larger than any constant, and (2) the $O(\log n)$ bit complexity assumption on the entries of A may not be realistic; ideally one should allow $\text{poly}(nd)$ bit complexity. This motivates the following central question:

Question 2: *Is it possible to obtain a polynomial time algorithm for ℓ_1 -low rank approximation for $k = \omega(1)$ with a small bicriteria rank r and with a $(1 + \varepsilon)$ approximation factor?*

1.1 Our Results

In this work we obtain results for both column subset selection and general low rank approximation.

1.1.1 ℓ_p Column Subset Selection for $1 \leq p < 2$

We resolve Question 1 above up to small factors by giving a $\text{poly}(ndk)$ time algorithm for finding a subset of $O(k \log k \log d)$ columns providing an $O(\sqrt{k \log^3 k})$ -approximation for ℓ_1 -low rank approximation. This nearly matches the lower bound of [48] which shows that there exist matrices for which the span of *any* subset of $\text{poly}(k)$ columns has approximation factor at least $k^{1/2-\gamma}$ for an arbitrarily small constant $\gamma > 0$. See

³See Theorem C.1, part III, of [48]. Note that part I of that theorem also obtains an $O(\sqrt{k \log k})$ -approximation, but it is not a column subset selection algorithm.

⁴For entrywise ℓ_p -low rank approximation, we seek to find rank- k factors U and V so as to minimize $\sum_{i,j} |(U \cdot V - A)_{i,j}|^p$.

Table 1 for a table containing our results and a comparison to prior work on column subset selection-based algorithms for ℓ_1 -low rank approximation.

Paper	Number of Columns	Approximation Factor	Time Complexity
[48]	$O(k \log k)$	$O(k \log k \log d)$	$\text{poly}(ndk)$
[48]	$O(k \log k)$	$O(\sqrt{k \log k})$	$d^{O(k \log k)}$
[11] and [14]	$O(k \log d)$	$O(k \log d)$	$\text{poly}(ndk)$
[11] and [14]	k	$O(k)$	$d^{O(k)}$
[51]	$O(k \log d)$	$O(k \log k)$	$\text{poly}(ndk)$
This paper	$O(k \log k \log d)$	$O(\sqrt{k}(\log k)^{\frac{3}{2}})$	$\text{poly}(ndk)$

Table 1: Summary of results for ℓ_1 column subset selection. The result in this paper is nearly optimal as it was shown in [48] that there exist matrices for which any subset of columns of size at least $\text{poly}(k)$ spans at best a $k^{1/2-\gamma}$ -approximation, for arbitrarily small constant $\gamma > 0$. We note that [14] obtains the same results as [11] for $p = 1$, but obtains strict improvements for every $p > 1$. Another related work is [49] which obtains a column subset giving a $(1 + \epsilon)$ -approximation, but assumes the entries of the input matrix are i.i.d. from a distribution with a certain moment assumption.

We next extend our result to any $1 < p < 2$, giving a $\text{poly}(ndk)$ time column subset selection-based algorithm for ℓ_p low rank approximation, that chooses $O(k \log k (\log \log k)^2 \log d)$ columns and provides an $O(k^{1/p-1/2} \log^{2/p-1/2} k (\log \log k)^{2/p-1})$ -approximation. We are not aware of any matching lower bounds for column subset selection for $1 < p < 2$ (the tight lower bounds in [14] are specifically for $p > 2$ and the lower bound in [48] is only for $p = 1$), so we also prove the first lower bounds for $1 < p < 2$, showing that there exist matrices for which the span of any column subset of size $k \cdot \text{polylog}(k)$ has approximation factor at least $k^{1/p-1/2-\gamma}$ for an arbitrarily small constant $\gamma > 0$. This shows that our algorithm is nearly optimal for $1 < p < 2$ as well.

Our results complement the results of [14] for the case of $p > 2$; together with our results we obtain optimal bounds on column subset selection, up to small factors, for every ℓ_p -norm, $p \geq 1$, closing a line of work on this problem [48, 51, 49, 11, 14]. We note that all the column subset selection algorithms in this line of work are bicriteria algorithms, as motivated and defined above, meaning that the actual number of columns returned is $k \cdot \text{poly}(\log(nk))$, which is optimal up to polylogarithmic factors.

We use the above algorithm as a subroutine to achieve two additional results, a polynomial time $\tilde{O}(k)$ -approximation algorithm with bi-criteria rank $\tilde{O}(k^2)$ (and no dependence on n and d , even logarithmic, in either the approximation factor or rank), and a $\text{poly}(k)$ -approximation algorithm with rank k (and no dependence on n and d , even logarithmic, in the approximation factor) with running time $\text{poly}(nd)$ — the first of these algorithms is a column subset selection algorithm, while the second is not. These algorithms may be of independent interest — note that all previously known algorithms for getting a $\text{poly}(k)$ -approximation factor independent of n or d require either d^k time or a $\log d$ term in the bi-criteria rank.

We note that there are also works studying ℓ_1 column subset selection, which obtain guarantees in terms of the error from the optimal *column subset* — two such works are [4, 24]. Our column subset selection results are not comparable to these. [4] studies the problem of fitting the columns of a matrix B using a subset of the columns of a (potentially different) matrix A , and considers the case where both A and B have nonnegative entries. The following guarantee is obtained by [4]: if A has a column subset of size k obtaining error at most $\epsilon \|B\|_1$, then the algorithm of [4] (given a $\delta > 0$) finds a subset of columns of A of size $O(k \log(1/\delta)/\delta^2)$ that fits B with error at most $\sqrt{\delta + \epsilon} \|B\|_1$. This error could potentially be larger than that of our algorithm which obtains relative error guarantees, if $\|B\|_1$ is much larger than OPT — in that case, even if there is a column subset obtaining error at most OPT , the algorithm of [4] could obtain error up to $\sqrt{\text{OPT} \cdot \|B\|_1}$, according to this guarantee. Hence, our algorithm could potentially be significantly better in the case where $\|B\|_1 \gg \text{poly}(k)\text{OPT}$. [24] gives a protocol for distributed column subset selection in the ℓ_p -norm, which obtains an $\tilde{O}(\sqrt{k} \log(d))$ -approximation *relative to the error of the best column subset* when $p = 1$. This could be larger than the error from our algorithm in cases where the error due to the best column subset is significantly larger than OPT — note that by the lower bound of [48], the error of the best column subset could be $\Omega(k^{\frac{1}{2}-\alpha} \cdot \text{OPT})$ for an arbitrarily small constant α .

1.1.2 General ℓ_p Low Rank Approximation for $1 \leq p < 2$

We next consider Question 2 and give a new bicriteria algorithm for entrywise ℓ_1 low rank approximation which goes beyond column subset selection. We improve upon the previous best bicriteria algorithms achieving $O(1)$ and $(1 + \varepsilon)$ -approximation, which require at least $n^{\text{poly}(k/\varepsilon)}$ time and are not polynomial time when $k = \omega(1)$. In contrast, our algorithm is polynomial time even for slow-growing functions k of n , such as $k = \Theta(\log^c n)$, for an absolute constant $c > 0$.

As mentioned above, bicriteria algorithms are natural and well-studied in this context. We note though that the main algorithm of [2] outputs a rank- k matrix while ours is bicriteria; however, no $(1 + \varepsilon)$ or $O(1)$ approximations in less than $n^{\text{poly}(k)}$ time were known even for bicriteria algorithms. In particular, even the bicriteria algorithm of [2] requires $n^{\text{poly}(k/\varepsilon)}$ time.

We give a $(1 + \varepsilon)$ -approximation algorithm with $2^{\text{poly}(k/\varepsilon)} + \text{poly}(nd)$ running time and output rank at most $3k$. See Table 2 for a table containing our result and a comparison to prior work for ℓ_1 -low rank approximation, where we list all previous $O(1)$ and $(1 + \varepsilon)$ approximation algorithms. Our algorithm also works for ℓ_p low rank approximation for $1 \leq p < 2$.

An interesting aspect of our result is that it shows entrywise ℓ_p -low rank approximation for $p \in [1, 2)$, with the above approximation factor and bicriteria rank, is *not* $W[1]$ -hard, in the language of parameterized complexity, which a priori may have been the case as all previous algorithms required $n^{\text{poly}(k)}$ time.

Paper	Bicriteria Rank	Approximation Factor	Time Complexity	Notes
[48], Lemma C.10	k	$O(1)$	$n^{\text{poly}(k)}$	BCA
[48], Theorem C.9	$3k$ ⁵	$O(1)$	$n^{\text{poly}(k)}$	
[2]	k	$1 + \varepsilon$	$n^{\text{poly}(k/\varepsilon)}$	BCA
[2]	$3k$	$1 + \varepsilon$	$n^{\text{poly}(k/\varepsilon)}$	
This paper	$3k$	$1 + \varepsilon$	$2^{\text{poly}(k/\varepsilon)} + \text{poly}(nd)$	

Table 2: Summary of our results for general ℓ_1 -Low Rank Approximation. We write BCA in the Notes next to each algorithm if it makes bit-complexity assumptions on the entries of the input matrix. We note that the non-bicriteria algorithm of [2] actually has an $(Mn)^{\text{poly}(k)}$ running time, where M is the maximum value of an entry of the input matrix, and thus is not even polynomial time for constant k if the entries of the input matrix are expressed with more than $\log n$ bits. In contrast, our algorithm is polynomial in the input description length, and thus for example, can handle entries as large as $M = 2^{\text{poly}(n)}$. The prior algorithms from [48] and [2] require time $n^{\text{poly}(k)}$, and thus are not polynomial time for any $k = \omega(1)$. In contrast, our algorithms are all polynomial time even if k is a slow-growing function of n , such as $\Theta(\log^c n)$ for an absolute constant $c > 0$. Another work [34] obtains a $(1 + \varepsilon)$ -approximation algorithm via gradient descent. This algorithm is not polynomial time or fixed-parameter tractable in the worst case, since the number of iterations of gradient descent required is $O(\frac{\sigma_k(\hat{X}_k^*)}{\varepsilon \text{OPT}})$, where \hat{X}_k^* is the optimal rank- k approximation to A in the Frobenius norm, and $\sigma_k(\hat{X}_k^*)$ is its k^{th} singular value. This convergence rate can be arbitrarily large: for instance, when the input matrix $A \in \mathbb{R}^{n \times n}$ is a diagonal matrix with $2^{\text{poly}(n)}$ for its first k entries, and 1 for the remaining $n - k$ entries, it is $\frac{2^{\text{poly}(n)}}{n-k}$.

1.1.3 Hardness for ℓ_p Low Rank Approximation with Additive Error - Appendix A

As an intermediate step for ℓ_p low rank approximation, we obtain an algorithm (Algorithm 2 for $p = 1$ and Algorithm 5 for general p) which achieves the following guarantee: given $A \in \mathbb{R}^{n \times d}$ and $k \in \mathbb{N}$, it obtains a matrix \hat{A} of rank at most k such that

$$\|\hat{A} - A\|_p \leq (1 + \varepsilon) \min_{A_k \text{ rank } k} \|A_k - A\|_p + \frac{\varepsilon}{f} \|A\|_p \quad (1)$$

in $f^{\text{poly}(k/\varepsilon)} + \text{poly}(nd)$ time, where f can be any desired number greater than 1. To our knowledge, there do not exist hardness results for such a guarantee (or for bi-criteria approximations). It is known, due to [2], that

⁵There is a slight typo in [48] in Theorem C.9 and its proof, where the bicriteria rank is said to be $2k$ — the bicriteria rank of that algorithm is actually $3k$, since once the $\text{poly}(n)$ -approximation B is subtracted off from the target matrix A , a good rank- $2k$ approximation M for $B - A$ is needed to recover the original optimum for A , from $B - A$.

achieving an $O(1)$ -approximation for ℓ_p low rank approximation ($p \in (1, 2)$) requires at least $2^{k^{\Omega(1)}}$ running time assuming the Small Set Expansion Hypothesis (SSEH) [43] and Exponential-Time Hypothesis (ETH) [23]. We extend the techniques of [2] to show that even achieving the guarantee in Equation 1 requires $2^{k^{\Omega(1)}}$ time assuming SSEH and ETH, for $p \in (1, 2)$, if $f \geq 2^{\text{poly}(k)}$, even when $\varepsilon = \Theta(1)$. Hence, our algorithm is close to optimal in a sense, since it can achieve that guarantee in $2^{\text{poly}(k)} + \text{poly}(nd)$ time.

We also show that it is optimal in the following related sense: it can achieve a similar guarantee in $2^{\text{poly}(k)} + \text{poly}(nd)$ time for the related problem of constrained ℓ_1 low rank approximation.

Problem 1 (Constrained ℓ_1 Low Rank Approximation). *Given a matrix $A \in \mathbb{R}^{n \times d}$ and a subspace $V \subset \mathbb{R}^n$, find a matrix \hat{A} of rank at most k minimizing $\|\hat{A} - A\|_1$, such that the columns of \hat{A} are in V .*

Our algorithm for $p = 1$ (see Algorithm 2) can be modified very slightly so that it achieves the same guarantee for this problem as well — the modified algorithm can compute, in $2^{\text{poly}(k)} + \text{poly}(nd)$ time, a matrix \hat{A} of rank at most k such that

$$\|\hat{A} - A\|_1 \leq O(1) \min_{A_k} \|A_k - A\|_1 + \frac{1}{2^{\text{poly}(k)}} \|A\|_1$$

and the columns of \hat{A} are contained in V — here, the minimum on the right-hand side is also taken over A_k with rank at most k , such that the columns of A_k are in V . Assuming the SSEH and randomized ETH (used for instance in [16]), we show that achieving this guarantee also requires at least $2^{k^{\Omega(1)}}$ time.

1.2 Our Techniques

We give an overview of our arguments for ℓ_1 -column subset selection, then for general ℓ_1 -low rank approximation. The arguments for ℓ_p -column subset selection and ℓ_p -low rank approximation, for $p \in (1, 2)$, are similar and are included in Section 2 and Section 3 respectively.

1.2.1 ℓ_1 Column Subset Selection

Algorithm and Initial $\tilde{O}(\sqrt{k}) \log(d)$ Approximation Factor. In our algorithm, we uniformly sample a column subset $S^{(0)}$ of size $t = O(k \log k)$ of our input matrix A and argue that we can approximately span a constant fraction of remaining columns of A using $S^{(0)}$, where to approximately span the i -th column means to obtain a column vector v^i for which $\|v^i - A_i\|_1 = O(\sqrt{k \log k})OPT/d$, where OPT is the cost of the optimal rank- k approximation to A . Thus, our total cost to cover a constant fraction of columns will be $O(\sqrt{k \log k})OPT$. We then recurse on the remaining fraction of columns. If there are d_1 columns in the next recursive call, we argue we approximately span a constant fraction of the remaining columns, where now to approximately span the i -th column means to obtain a column vector v^i for which $\|v^i - A_i\|_1 = O(\sqrt{k \log k})OPT/d_1$. Again, the total cost to cover a constant fraction of remaining columns is $O(\sqrt{k \log k})OPT$. We then recurse on the columns still remaining. Since we approximately span a constant fraction of columns in each recursive step, after $O(\log d)$ recursive steps we will have spanned all d columns. The total number of columns we will have chosen is $O(k \log k \log d)$ and the overall approximation factor will be $O(\sqrt{k \log k} \log d)$.

The algorithm described above is simple, and reminiscent of column sampling algorithms [11, 14, 51] in prior work. However, our analysis is completely new and does not involve going through maximum determinant subsets, as in each of these previous column sampling algorithms. These works argued that if one uniformly samples a set S of $2k$ columns of $V^* \in \mathbb{R}^{k \times n}$, where $A_k = U^*V^*$ is the best rank- k approximation to A , and considers a random additional column c , then with probability at least $1/2$, the maximum determinant subset (which is of size k) of columns of $V_{S \cup \{c\}}^* \in \mathbb{R}^{k \times (2k+1)}$ does not contain c , and consequently, by Cramer's rule, c can be expressed as a linear combination of columns in our sample set S with coefficients of absolute value at most 1, and thus by the triangle inequality one pays a cost at most what the subset S pays to approximate the c -th column of A , which since S was chosen uniformly at random, is at most $O(k)OPT/d$ with constant probability.

We do not know how to reduce the approximation factor in the analyses of all of these previous algorithms; intuitively, the difficulty stems from the fact that the maximum determinant subset may not be the best

subset to look at for ℓ_1 ; indeed, it could be that for a random c , one needs coefficients of absolute value 1 to span it using the columns in the set S , and it is unclear how the error propagates other than through the triangle inequality. We thus give the first analysis of the above sampling framework that *does not go through maximum determinant subsets*.

Instead, we argue that if one had V^* , then one could sample columns using its so-called *Lewis weights* [13], creating a sampling and rescaling matrix R with $t/2 = O(k \log k)$ columns, so that the solution $U = \operatorname{argmin}_U \|A_{S \cup c} R - U(V^*)_{S \cup c} R\|_1$ would be an $O(1)$ -approximate rank- k left factor for the submatrix of A indexed by $S \cup c$. By relating ℓ_1 and ℓ_2 -norms of the rows of $A_{S \cup c} R - U(V^*)_{S \cup c} R$ and using the normal equations for least squares regression, we get that $U' = A_{S \cup c} R ((V^*)_{S \cup c} R)^+$ provides an $O(\sqrt{k \log k})$ -approximate rank- k left factor for the submatrix of A indexed by $S \cup c$. The advantage of U' is that it is in the column span of $A_{S \cup c} R$. Note that this subroutine should be reminiscent of the algorithm of [48], which argued one could obtain an $O(\sqrt{k \log k})$ -approximate column subset selection algorithm by enumerating over all subsets of $t/2$ columns of A ; however [48] actually does this and suffers $d^{\Omega(k \log k)}$ time.

Instead, we argue as follows. As noted by [48], since $A_{S \cup c} R$ is an $O(\sqrt{k \log k})$ -approximate left factor for $A_{S \cup c}$ with high probability (over the randomness of R), in particular, there exists a fixed matrix R_0 , with one nonzero entry per column, such that $A_{S \cup c} R_0$ is an $O(\sqrt{k \log k})$ -approximate left factor for $A_{S \cup c}$ (and in particular, $A_{S \cup c} R_0$ is a column subset of $A_{S \cup c}$). Recall that A_c is a uniformly random column of $A_{S \cup c}$, $A_{S \cup c}$ has $t + 1$ columns, and $A_{S \cup c} R_0$ has $t/2$ columns (where $t/2 = O(k \log k)$ is the number of columns required by [13] to perform ℓ_1 Lewis weight sampling). Hence, we can now simply argue as in [11, 51] that with probability at least $\frac{1}{2}$ over S and c , A_c is not in $A_{S \cup c} R_0$.

In summary, we use the facts that (1) column A_c is not chosen by the matrix R_0 with probability at least $\frac{1}{2}$ over S and c , (2) the overall cost of using the columns of $A_{S \cup c} R_0$ to approximately span all the columns of $A_{S \cup c}$ is $O(\sqrt{k \log k}) OPT_{S \cup c}$, where $OPT_{S \cup c}$ is the total cost on $A_{S \cup c}$ under the optimal rank- k approximation, (3) since $S \cup \{c\}$ is itself a uniformly random subset of $[d]$, the value $OPT_{S \cup c}$ is $O(t/d) OPT$ in expectation, and (4) A_c is among $t/2$ random columns of $A_{S \cup c}$ which are not in the chosen subset of columns $A_{S \cup c} R_0$, and thus has at most a $2/t$ fraction of the total cost on $A_{S \cup c}$, in expectation. Combining these statements gives us that with large constant probability, the cost of approximately spanning A_c using our sampled set is $O(\sqrt{k \log k}) OPT/d$, completing the argument.

Remark 1.1. *In fact, by considering the properties of the Lewis weight sampling matrix R , we can even show $A_{S \cup c} R$ is unlikely to contain the column A_c . The subtlety with such an approach is that the distribution of the sampling matrix R may itself depend on A_c . In short, we can argue that since c is a uniformly random column in $S \cup c$, this means that in expectation (and with constant probability), the Lewis weight of the column A_c is small. Precisely, the probability that any individual column of R selects A_c is $\Theta(\frac{1}{r}) = \Theta(\frac{1}{k \log k})$. Thus, with constant probability, A_c is not sampled by R , since R samples $r = O(k \log k)$ columns of A independently with replacement (where the sampling distribution is determined by the Lewis weights of V^*). In summary, with constant probability, it is simultaneously true that (i) $A_{S \cup c} R$ covers $A_{S \cup c}$, meaning it incurs an error of $O(\frac{\sqrt{k \log k}}{d}) OPT$ on A_c in expectation, since A_c is a uniformly random column of $A_{S \cup c}$, and (ii) it does not contain A_c . The argument is the same from here onwards. This is the approach we took in the previous version of this work on arXiv.*

We generalize this approach to obtain optimal column subset selection algorithms for entrywise ℓ_p low rank approximation, for every $1 < p < 2$, replacing the ℓ_1 Lewis weights with the ℓ_p Lewis weights in the above analysis. Note we cannot use earlier sampling distributions, such as ℓ_p leverage scores or total sensitivities, as it is also important in the argument above that one only needs to sample $O(k \log k)$ columns for a rank- k space and these latter sampling distributions would require a larger k^{1+c} samples for a constant $c > 0$ (see, e.g., [53] for a survey); this is important not only for the overall number of sampled columns but also for the approximation factor, since we also relate the p -norm to the 2-norm through this number.

Removing $\log d$ from the Approximation Factor. An improved version of the above argument, inspired by [51], gives an $O(\sqrt{k}(\log k)^{\frac{3}{2}})$ -approximation rather than $O(\sqrt{k \log k} \log d)$ — to achieve this, we note that in each round, we can condition on the event that the columns being chosen are not among the $\frac{1}{t}$ -fraction of columns which have the highest cost, under the optimal ℓ_1 rank- k approximation. This event occurs with constant probability. Moreover, for each of the other columns which are not in this top $\frac{1}{t}$ -fraction of columns

(which are indexed by a subset $F \subset [d]$), we can bound the cost by $O(\sqrt{k \log k}) \frac{OPT_{-F}}{d}$, where OPT_{-F} denotes the cost under the optimal rank- k approximation, excluding the errors from the top $\frac{1}{t}$ -fraction of columns F . Finally, we show that over the course of the $O(\log d)$ recursive rounds, a particular column of A can contribute to OPT_{-F} in at most $O(\log k)$ rounds — that is, it cannot be outside of F for more than $O(\log k)$ rounds without being approximately covered and discarded. A similar technique was used in [51] to obtain an $O(k \log k)$ -approximation factor independent of $\log d$.

Lower Bound for ℓ_p Column Subset Selection, $1 \leq p < 2$. Our nearly matching lower bound for entrywise ℓ_p low rank approximation is a technical generalization of that for ℓ_1 -low rank approximation given in [48] and we defer the details to Appendix B. For $p \in (1, 2)$, we show that a proof strategy similar to that of [48] can be used to show that any column subset selection algorithm that selects at most $O(k(\log k)^{c_1})$ columns (where c_1 can be any constant) achieves no better than an $\Omega(k^{\frac{1}{p}-\frac{1}{2}-\alpha})$ approximation factor in the worst case, where α can be an arbitrary number in $(0, \frac{1}{p} - \frac{1}{2})$.

Additional Column Subset Selection Results (Appendix C): Decreasing the Bicriteria Rank.

Our polynomial time $O(k(\log k)^2)$ -approximation algorithm with $O(k^2(\log k)^2)$ bicriteria rank makes use of the improved approximation factor that is independent of $\log d$, and relies on the following simple observation. Let $U \in \mathbb{R}^{n \times O(k \log k \log d)}$ be the left factor ultimately returned by our main column subset selection algorithm. In each recursive round, if S is the set of $t = O(k \log k)$ columns which are sampled, then for each column A_i which is discarded during that round, A_i can be approximately covered using only the $t = O(k \log k)$ columns belonging to S . The implication of this is that there exists M having rank $O(k \log k \log d)$, which provides an $O(\sqrt{k}(\log k)^{\frac{3}{2}})$ -approximation for A , such that each column of M can be written exactly as a linear combination of $O(k \log k)$ columns of U (more specifically, $O(k \log k)$ columns of U that were obtained in a single round of sampling from A).

Now, as mentioned above, it was shown in [48] that any matrix has a column subset of size $O(k \log k)$ spanning an $O(\sqrt{k \log k})$ -approximation — hence, M has a column subset of size $O(k \log k)$ which spans an $O(\sqrt{k \log k})$ -approximation to M . By the triangle inequality, one can show that since M is an $O(\sqrt{k}(\log k)^{\frac{3}{2}})$ -approximation for A , this $O(k \log k)$ -sized column subset of M spans an $O(k(\log k)^2)$ -approximation for A . To form our left factor, for each of these columns A_i , we could collect the $O(k \log k)$ columns that were sampled in the round when A_i was covered. Since we do not actually know this poly(k)-approximate subset of columns of M , we could naïvely try all of them — however, rather than checking all column subsets of M of size $O(k \log k)$, which takes time $d^{O(k \log k)}$, it suffices to check all $O(k \log k)$ -sized subsets of the $O(\log d)$ rounds of sampling done by our main algorithm, and take the best subset. There are $\binom{O(\log d)}{O(k \log k)} \leq d^{O(1)}$ such subsets, so this algorithm is polynomial time. Finally, to obtain a poly(k)-approximate matrix with rank at most k , we combine this algorithm with an algorithm from [48] (analyzed in Theorem C.19 of that work), which takes a bi-criteria solution of rank r as input and reduces the rank to k , at the cost of an increase in the error by a factor of poly(r). Since the bi-criteria rank of our solution is only poly(k), the approximation error only increases by a factor of poly(k).⁶

1.2.2 General ℓ_1 Low Rank Approximation

We next turn to general low rank approximation, where it is possible to obtain much smaller approximation factors than with column subset selection. A crucial novelty in our algorithm is the use of a randomized rounding technique for solving an integer linear program (ILP) — to our knowledge, such an approach was not previously considered in the context of ℓ_1 or ℓ_p low rank approximation. Randomized rounding of relaxations has been previously used in other subspace optimization problems, such as by [18] in the related problem of subspace approximation (in the $\ell_{p,2}$ norm). However, [18] uses it to select random linear combinations of singular vectors of a matrix obtained by solving a convex relaxation, while we use randomized rounding of an LP to choose columns satisfying multiple linear constraints. One appealing aspect of our algorithm is

⁶In an earlier version of this work, we instead used an algorithm based on Algorithm 4 of [11] to reduce the rank to at most k . This older version of our algorithm had a running time of $2^{O(k \log k)} + \text{poly}(nd)$. Our new algorithm that returns a matrix of rank at most k , shown in Algorithm 8 and using an algorithm of [48], is polynomial time.

that it does not use polynomial system solvers, which are somewhat impractical — these have been used for several other NP-hard matrix factorization problems, such as in [45, 50].

Background: The Algorithm of [2] and its Bicriteria Variant. As a starting point, we recall the bicriteria variant of the main algorithm of [2] (shown in Algorithm 1), which uses a median-based sketch to obtain a $(1 + \varepsilon)$ -approximation for ℓ_1 low rank approximation. This sketch was previously considered in [1], where it was shown that for a k -dimensional subspace V of \mathbb{R}^n , if $S \in \mathbb{R}^{\text{poly}(k/\varepsilon) \times n}$ is a random matrix with i.i.d. standard Cauchy entries, then with constant probability, for all $v \in V$, $\text{med}(Sv)$ is within a $(1 + \varepsilon)$ factor of $\|v\|_1$, where $\text{med}(v)$ is the median of the absolute values of the entries of the vector v . This sketch was then considered in the context of ℓ_1 low rank approximation by [2], where the following “one-sided embedding” property of this sketch is shown: for a given matrix $U \in \mathbb{R}^{n \times k}$, and a fixed matrix $A \in \mathbb{R}^{n \times d}$, if $S \in \mathbb{R}^{\text{poly}(k/\varepsilon) \times n}$, then with constant probability $\text{med}(SUV - SA) \geq (1 - \varepsilon)\|UV - A\|_1$ for all matrices $V \in \mathbb{R}^{k \times d}$, where if M is a matrix with d columns, then $\text{med}(M) := \sum_{i=1}^d \text{med}(M_i)$. These properties make this sketch useful in ℓ_1 low rank approximation, as we now see.

Algorithm 1 $(1 + \varepsilon)$ -approximation algorithm from [2] that gives bi-criteria rank $3k$. This is adapted from Algorithm 1 of [2] and Theorems 10 and 23 of [2].

Require: $A \in \mathbb{R}^{n \times d}$, $k \in \mathbb{N}$, $\varepsilon > 0$ with $n \geq d$

Ensure: $\hat{A} \in \mathbb{R}^{n \times d}$

procedure PREVIOUSONEPLUSEPSAPPROXIMATION(A, k, ε)

$B \leftarrow$ The rank- k SVD of A

$C \leftarrow A - B$

$S \leftarrow$ A $\text{poly}(k/\varepsilon) \times n$ matrix of i.i.d. standard Cauchy random variables

$U, V \leftarrow$ The $n \times 2k$ and $2k \times d$ zero matrices

$\tilde{S} \leftarrow S$ with each entry rounded to the nearest integer multiple of $\frac{\varepsilon}{\text{poly}(n)}$

 Guess all possible values of $\tilde{S}U^*$ with each entry rounded to the nearest integer multiple of $\frac{\varepsilon^3}{\text{poly}(n)}\|C\|_1$.

for each guessed value M of $\tilde{S}U^*$ **do**

$V_{\text{guess}} \leftarrow \text{argmin}_{V'} \text{med}(MV' - \tilde{S}C)$ such that $\|V'\|_\infty \leq \text{poly}(k)$

$U_{\text{guess}} \leftarrow \text{argmin}_{U'} \|U'V_{\text{guess}} - C\|_1$

if $\|U_{\text{guess}}V_{\text{guess}} - C\|_1 \leq \|UV - C\|_1$ **then**

$U \leftarrow U_{\text{guess}}, V \leftarrow V_{\text{guess}}$

end if

end for

return $B + UV$

end procedure

Algorithm 1 is the bicriteria variant of the main algorithm of [2] (we cannot directly modify the main algorithm since it requires bit complexity assumptions). The main algorithm is given in Algorithm 1 of [2], and Theorem 10 of [2] gives the analysis of that algorithm, while Theorem 23 of [2] describes how the algorithm should be modified to remove the need for bit complexity assumptions at the cost of a bicriteria rank of $3k$. Hence, our summary of the analysis largely follows Theorem 10 of [2], with minor modifications as given by Theorem 23 of [2].

Briefly, its analysis proceeds as follows. Define $U^* \in \mathbb{R}^{n \times 2k}$ and $V^* \in \mathbb{R}^{2k \times d}$ so that U^*V^* is the optimal rank- $2k$ approximation for C . First, V^* is assumed without loss of generality to be a $\text{poly}(k)$ -well-conditioned basis, meaning for all row vectors $x \in \mathbb{R}^{2k}$, $\frac{1}{\text{poly}(k)}\|x\|_1 \leq \|x^T V^*\|_1 \leq \text{poly}(k)\|x\|_1$. Then, U^* is assumed to have all of its entries rounded to the nearest integer multiple of $\frac{\varepsilon\|C\|_1}{\text{poly}(n)}$. This is not an issue, because if \tilde{U} is the rounded version of U^* , then

$$\|(\tilde{U} - U^*)V^*\|_1 \leq \text{poly}(k) \cdot \|\tilde{U} - U^*\|_1 \leq \text{poly}(k) \cdot O(nk) \cdot \frac{\varepsilon\|C\|_1}{n \cdot \text{poly}(n)} = \frac{\varepsilon}{\text{poly}(n)}\|C\|_1 \leq \varepsilon \cdot OPT$$

where the second inequality is because \tilde{U}, U^* have $2nk$ entries, and the last inequality is because the rank- k SVD of A gives an $O(n)$ -approximation, assuming $n \geq d$.

In addition, S is also assumed to be discretized, and the discretized version is written as \tilde{S} . This is done as follows. First, because V^* is a $\text{poly}(k)$ -well-conditioned basis, each of its entries is at most $\text{poly}(k)$ (note that in Theorem 10 of [2], it is mentioned that we can assume each entry of V^* is at most $\text{poly}(ndk/\varepsilon)$, but this can be decreased further to $\text{poly}(k)$). Hence, we can restrict ourselves to right factors V' for which each of its entries is at most $\text{poly}(k)$. Now, if $V \in \mathbb{R}^{2k \times d}$ has each entry at most $\text{poly}(k)$, then

$$\|U^*V - C\|_1 \leq \|U^*V\|_1 + \|C\|_1 = \sum_{i=1}^d \|U^*V_i\|_1 + \|C\|_1 \leq \text{poly}(k)d\|U^*\|_1 + \|C\|_1 = \text{poly}(k) \cdot d \cdot \|C\|_1$$

where the first inequality is by the triangle inequality, the second inequality is because $\|V_i\|_\infty \leq \text{poly}(k)$, and the last equality is because $\|U^*\|_1 \leq \text{poly}(k)\|U^*V^*\|_1 = \text{poly}(k)\|C\|_1$ since V^* is a well-conditioned basis.

Now, the reason for rounding the entries of S to the nearest integer multiple of $\frac{\varepsilon}{\text{poly}(n)}$ (and why this does not significantly increase the error) is that, for any V with $\|V\|_\infty \leq \text{poly}(k)$,

$$\begin{aligned} |\text{med}(\tilde{S}U^*V - \tilde{S}C) - \text{med}(SU^*V - SC)| &\leq \sum_{i=1}^d \|(\tilde{S}U^*V_i - \tilde{S}C_i) - (SU^*V_i - SC_i)\|_\infty \\ &= \sum_{i=1}^d \|(\tilde{S} - S)(U^*V_i - C_i)\|_\infty \\ &\leq \sum_{i=1}^d \|\tilde{S} - S\|_\infty \|U^*V_i - C_i\|_1 \\ &= \|S - \tilde{S}\|_\infty \|U^*V - C\|_1 \end{aligned} \tag{2}$$

where the first inequality is due to the fact that, if $v_1, v_2 \in \mathbb{R}^n$, then $|\text{med}(v_1 + v_2) - \text{med}(v_1)| \leq \|v_2\|_\infty$, and the second is because, for a matrix B and a vector v , $\|Bv\|_\infty \leq \|B\|_\infty \|v\|_1$.

Since the entries of S are rounded to the nearest multiple of $\frac{\varepsilon}{\text{poly}(n)}$, $\|S - \tilde{S}\|_\infty \leq \frac{\varepsilon}{\text{poly}(n)}$ and $\|S - \tilde{S}\|_\infty \|U^*V - C\|_1$ is at most $\frac{\varepsilon}{\text{poly}(n)}\|C\|_1 = \varepsilon \cdot \text{OPT}$. Therefore, not much additional error is incurred when minimizing $\text{med}(\tilde{S}U^*V - \tilde{S}C)$, subject to the constraint that $\|V\|_\infty \leq \text{poly}(k)$, as opposed to minimizing $\text{med}(SU^*V - SC)$.

In summary, the algorithm works by guessing all possible values of $\tilde{S}U^*$. By well-known properties of Cauchy matrices, the entries of S are bounded above by $\text{poly}(n)$, and those of U^* can also be bounded above by $\text{poly}(n)\|C\|_1$, meaning there are $\text{poly}(n/\varepsilon)$ choices for each entry of $\tilde{S}U^*$, and $\tilde{S}U^*$ has $\text{poly}(k/\varepsilon)$ entries, meaning the overall running time is $n^{\text{poly}(k/\varepsilon)}$.

Our Approach: Achieving FPT Time by Reducing the Number of Guesses Per Entry. Our approach, like those of [2] and Theorem C.9 of [48], follows the general strategy of first taking a good initialization B , subtracting it from A , and finding a good rank $2k$ approximation for the residual $C := B - A$. Now, the running time of Algorithm 1 is dominated by the time it takes to guess $\tilde{S}U^*$, and there are $\text{poly}(n/\varepsilon)$ guesses per entry of $\tilde{S}U^*$. We now describe how with our approach, we reduce the number of possibilities per entry to $\text{poly}(k/\varepsilon)$, while still obtaining a $(1 + \varepsilon)$ -approximation.

Perhaps the most obvious optimization to make is to use a better initialization — rather than letting B be the rank- k SVD of A , we could run a $\text{poly}(k)$ -approximation algorithm on A to obtain B , such as our Algorithm 8 which gives a $\text{poly}(k)$ -approximation with rank at most k in $\text{poly}(nd)$ time. A $\text{poly}(k) \log(d)$ -approximation algorithm, such as that of [48], would also suffice.

Using an initialization algorithm with a better approximation factor reduces the number of guesses per entry of $\tilde{S}U^*$, but the number of guesses remains $\text{poly}(n/\varepsilon)$, rather than $\text{poly}(k/\varepsilon)$, mainly for the following reasons:

- When U^* is being discretized (recall that this is not explicitly done in the algorithm, but the analysis assumes U^* is discretized in order to have a finite number of entries to guess) the entries still need to be rounded to the nearest integer multiple of $\frac{\varepsilon\|C\|_1}{n \cdot \text{poly}(k)}$. This is because U^* has n rows, and therefore,

if \tilde{U} denotes the rounded version of U^* , then the additional error from using \tilde{U} instead of U^* can still only be upper bounded by

$$\|(\tilde{U} - U^*)V^*\|_1 \leq \text{poly}(k)\|\tilde{U} - U^*\|_1 \leq \text{poly}(k) \cdot O(nk) \cdot \|\tilde{U} - U^*\|_\infty$$

in the worst case — meaning a rounding granularity of at least $\frac{1}{n}$ is needed.

- When S is being discretized to obtain \tilde{S} , then as mentioned above in Equation 2, the additive error is at most $\|S - \tilde{S}\|_\infty \|U^*V - C\|_1$, where V is the right factor that Algorithm 1 obtains. Recall that the upper bound for $\|U^*V - C\|_1$ is $\text{poly}(k) \cdot d \cdot \|C\|_1$ for *all* V with no entry larger than $\text{poly}(k)$ (in our exposition of the algorithm of [2], we showed that we just need to consider V with no entry larger than $\text{poly}(k)$, while the original algorithm in [2] in fact allowed V to have entries at most $\text{poly}(n/\varepsilon)$ — this small change can only improve the upper bound on the additive error). Hence, when rounding S , a granularity of at least $\frac{1}{d}$ seems to be needed.

It is not clear how to circumvent these issues if we round S and U^* separately. Instead, to avoid rounding with granularities of $\frac{\varepsilon}{n}$ or $\frac{\varepsilon}{d}$, we round SU^* itself in our analysis. Specifically, we show the following. If M is SU^* , but with each entry rounded to the nearest power of $1 + \frac{1}{\text{poly}(k/\varepsilon)}$, or set to 0 if it is below $\text{poly}(\varepsilon/k) \cdot \text{OPT}$, then we obtain a small additive error by solving the following problem instead:

$$\min_{V'} \text{med}(MV' - SC) \text{ subject to } \|V'\|_1 \leq \text{poly}(k)$$

and then again finding a good left factor U' for V' through linear programming. The number of choices for each entry of M is then $\text{poly}(k/\varepsilon)$, and since M is a $\text{poly}(k/\varepsilon) \times k$ matrix (in fact, a $k \cdot \text{poly}(1/\varepsilon) \times k$ matrix) the number of guesses for M is $2^{O(k^2 \cdot \text{poly}(1/\varepsilon) \cdot \text{polylog}(k/\varepsilon))}$. Note that the constraint on V' is now different — instead of having the constraint that $\|V'\|_\infty \leq \frac{\text{poly}(n)}{\varepsilon}$, as in the main algorithm of [2], or $\|V'\|_\infty \leq \text{poly}(k)$, as in our presentation of that paper's algorithm, we instead enforce a constraint on the ℓ_1 -norm of V' . This has the benefit that the additive error obtained by minimizing $\text{med}(MV' - SC)$ instead of $\text{med}(SU^*V' - SC)$ is small — this is necessary because we are now rounding SU^* to a coarser granularity, $\frac{1}{\text{poly}(k/\varepsilon)}$ instead of $\frac{\varepsilon^2}{\text{poly}(n)}$. However, enforcing the constraint that $\|V'\|_1 \leq \text{poly}(k)$ is nontrivial, as we now see.

Remark 1.2. *We round each entry of SU^* to the nearest power of $1 + \frac{1}{\text{poly}(k/\varepsilon)}$ — a similar alternative approach that seems to work is rounding the entries to the nearest multiple of $\frac{1}{\text{poly}(k/\varepsilon)} \cdot \|C\|_1$, which is $\frac{1}{\text{poly}(k/\varepsilon)} \cdot \text{OPT}$ if C is the residual from a $\text{poly}(k)$ -approximation algorithm.*

Ensuring that the Candidate Right Factor Has Norm At Most $\text{poly}(k)$ Through Randomized Rounding of an ILP. How do we enforce the ℓ_1 -norm constraint on V' ? First, let us discuss how the original median-based problem is solved in [2]. For each column index $i \in [d]$, Algorithm 1 finds $V_i \in \mathbb{R}^{2k}$ such that V_i minimizes $\text{med}(\tilde{S}U^*V_i - \tilde{S}C_i)$ subject to the constraint that $\|V_i\|_\infty$ is small. Note that there are $r!$ orderings of the coordinates of $\tilde{S}U^*V_i - \tilde{S}C_i$, where r is the number of rows in S , meaning that all of those orderings can be tried — for a fixed ordering of the coordinates, the ordering turns into a linear constraint, and the ℓ_∞ norm constraint can also be written as a linear constraint, meaning this can be solved with linear programming, and the overall running time is $r! \text{poly}(nd) = 2^{O(r \log r)} \text{poly}(nd)$, and this fits within the $n^{\text{poly}(k/\varepsilon)}$ running time of Algorithm 1.

Enforcing the constraint that $\|V'\|_1 \leq \text{poly}(k)$ is more subtle. If we solve a similar problem on each $i \in [d]$ — for instance, minimizing $\text{med}(MV'_i - SC_i)$ such that $\|V'_i\|_1 \leq \text{poly}(k)$ — then the overall norm of $\|V'\|_1$ could still depend on d in the worst case, if each minimizer V'_i has norm roughly equal to $\text{poly}(k)$. It is also not easy to directly include this constraint inside a median-based optimization problem that includes information from all the columns. For instance, one naïve way of minimizing $\text{med}(MV' - SC)$ such that $\|V'\|_1 \leq \text{poly}(k)$ is to do the following: simultaneously try all possible orderings, for each $i \in [d]$, of the coordinates of $MV'_i - SC_i$. The number of such orderings is $(r!)^d$, and this does not lead to an FPT running time.

Instead, we still solve separate median-based optimization problems for each $i \in [d]$, and combine the results for different i using a relaxation of a suitable ILP. Instead of finding V_i minimizing $\text{med}(MV_i - SC_i)$,

we instead seek to minimize the ℓ_1 norm of V_i . At the same time, we would like the cost $\text{med}(MV_i - SC_i)$ to be small enough. Hence, for each column index $i \in [d]$, we find a column $V_{i,c}$ minimizing $\|V_{i,c}\|_1$, subject to the constraint that $\text{med}(MV_{i,c} - SC_i) \leq c$ for a well-chosen c . For any c , the running time of this step is $r! \text{poly}(nd)$ (by trying all orderings of the coordinates of $MV_{i,c} - SC_i$ and including the orderings as linear constraints in the LP), and this fits in our desired FPT running time.

Here, c should be chosen so that it is not much higher than the cost $\text{med}(MV_i^* - SC_i)$ of V_i^* — precisely, it should be within a $(1 + O(\varepsilon))$ factor of $\text{med}(MV_i^* - SC_i)$. Although we do not know $\text{med}(MV_i^* - SC_i)$, we can guess all powers of $(1 + \varepsilon)$ less than $O(1)\|C\|_1$ and greater than $O(\frac{\varepsilon^2}{\text{poly}(k)d})\|C\|_1$ in the place of c . (The lower bound is chosen so that, even if the cost on some of the columns is $O(\frac{\varepsilon^2}{\text{poly}(k)d})\|C\|_1$, the overall additive error of these columns is at most $O(\frac{\varepsilon^2}{\text{poly}(k)})\|C\|_1$ which is acceptable.) The number of such cost bounds c is thus polynomial in d, k and $1/\varepsilon$.

For each $i \in [d]$, we now have several minimizers $V_{i,c}$ for each possible cost bound c . Now, the question is, which cost bound c should we pick for each $i \in [d]$? We can decide this through the following 0–1 integer linear program. For each $i \in [d]$ and each possible cost bound c , we create a variable $x_{i,c}$ which can be 0 or 1 (1 representing the minimizer $V_{i,c}$ being chosen as the i^{th} column of V' , and 0 representing $V_{i,c}$ not being chosen). It is then natural to add the constraint that $\sum_c x_{i,c} = 1$ for each $i \in [d]$, since only one $V_{i,c}$ can be chosen as the i^{th} column of V' .

In addition, we wish to have $\|V'\|_1 \leq \text{poly}(k)$ and $\text{med}(MV' - SC) \leq (1 + O(\varepsilon))OPT_{C,2k} + O(\varepsilon^2/f)\|C\|_1$ (where $OPT_{C,2k}$ is the optimal rank- $2k$ approximation error for C). Note that these can be made to hold if V' is taken to be V^* , since for each $i \in [d]$, there is at least one cost bound c for which V_i^* is feasible. These can be represented as constraints that are linear in the $x_{i,c}$, since

$$\|V'\|_1 = \sum_{i=1}^d \sum_c x_{i,c} \|V_{i,c}\|_1$$

and

$$\text{med}(MV' - SC) = \sum_{i=1}^d \sum_c x_{i,c} \text{med}(MV_{i,c} - SC_i)$$

Now, solving this ILP will again take at least $2^{\Omega(kd)}$ time — instead, we can relax the 0–1 constraint on the $x_{i,c}$, so that we now have the constraints $x_{i,c} \in [0, 1]$ for all i, c . Since, for each $i \in [d]$, we also have the constraint $\sum_c x_{i,c} = 1$, this means that for each column V'_i of V' , the $x_{i,c}$ give us a probability distribution on the cost bound c to be chosen for V'_i . By the constraints of the new LP, if we sample for each $i \in [d]$ a single $V_{i,c}$ to be the i^{th} column of V' , according to the distribution given by the $x_{i,c}$ (i.e., for each $i \in [d]$, $V_{i,c}$ is chosen with probability $x_{i,c}$) then the expectation of $\|V'\|_1$ is $\text{poly}(k)$, while the expectation of $\text{med}(MV' - SC)$ is at most $(1 + O(\varepsilon))OPT + \frac{\varepsilon}{\text{poly}(k)}\|C\|_1$.

This gives the desired result, but a few subtleties arise when sampling according to the $x_{i,c}$ and analyzing this using Markov's inequality. To obtain a $(1 + O(\varepsilon))$ -approximation, we need $\text{med}(MV' - SC)$ to be at most $(1 + O(\varepsilon))$ times its expectation. By Markov's inequality, $\text{med}(MV' - SC)$ is at most $(1 + 2\varepsilon)$ -times its expectation with probability at least ε (meaning this fails to occur with probability at most $1 - \varepsilon$). To apply a union bound to control $\|V'\|_1$ as well, we note that $\|V'\|_1$ is at most $\frac{2}{\varepsilon}$ times its expectation with failure probability at most $\frac{\varepsilon}{2}$ — having $\|V'\|_1 \leq \frac{\text{poly}(k)}{\varepsilon}$ instead of $\|V'\|_1 \leq \text{poly}(k)$ is enough for our purposes, since we round SU^* with a granularity of $\frac{1}{\text{poly}(k/\varepsilon)}$. By a union bound, $\|V'\|_1$ and $\text{med}(MV' - SC)$ are both small enough with probability $\frac{\varepsilon}{2}$ (i.e., failure probability $1 - \frac{\varepsilon}{2}$), and we can simply sample V' a total of $O(1/\varepsilon)$ times, choosing the best solution found, to reduce this failure probability to a small constant independent of ε .

Finally, since V' has norm at most $\frac{\text{poly}(k)}{\varepsilon}$, the difference between $\text{med}(MV' - SC)$ and $\text{med}(SU^*V' - SC)$ is at most $\frac{1}{\text{poly}(k/\varepsilon)}\|C\|_1$, meaning $\text{med}(SU^*V' - SC)$ is also small enough, and so is $\|U^*V' - C\|_1$. At this point, we can find an appropriate left factor U' for V' through linear programming.

As a summary of this discussion, we show our algorithm in the ℓ_1 -case in Algorithms 2 and 3. Algorithm 3 simply shows the process of obtaining an initial crude approximation B and subtracting it from A to obtain

C , and Algorithm 2 shows how we obtain a matrix UV such that

$$\|UV - C\|_1 \leq (1 + O(\varepsilon)) \min_{C^* \text{ rank } 2k} \|C^* - C\|_1 + O\left(\frac{\varepsilon}{\text{poly}(k)}\right) \|C\|_1$$

meaning that $UV + B$ is a $(1 + \varepsilon)$ -approximation to the optimal rank- k approximation error for A .

Remark 1.3. *Note that we need to know OPT in order to enforce the linear constraint that $\text{med}(MV' - SC)$ is at most $(1 + O(\varepsilon))OPT + \frac{\varepsilon}{\text{poly}(k)}\|C\|_1$. Observe that it suffices to have an estimate \widehat{OPT} of OPT that is accurate within a $(1 + \varepsilon)$ -factor. We can obtain such an \widehat{OPT} as follows — if E is the error achieved by the rank- k SVD of C , then E is within a \sqrt{nd} factor of OPT , meaning it suffices to guess all powers of $(1 + \varepsilon)$ that are between OPT and $\frac{1}{\sqrt{nd}}OPT$, and one of these will give a $(1 + O(\varepsilon))$ -approximate factorization with additive $\frac{\varepsilon}{\text{poly}(k)}\|C\|_1$ error.*

Remark 1.4. *For each column index i and cost bound c , we minimize the norm of $V_{i,c}$ such that $\text{med}(MV_{i,c} - SC_i) \leq c$. The argument would also proceed similarly if we minimized $\text{med}(MV_{i,c} - SC_i)$ while having a constraint on the norm of $V_{i,c}$. In particular, we can try all powers of $(1 + \varepsilon)$ between $\frac{\text{poly}(k)}{n}$ and $\text{poly}(k)$, and the linear program will still be feasible because $\|V^*\|_1 \leq \text{poly}(k)$.*

Hardness for Additive Error - Appendix A. Our techniques for our hardness results are based on the proof by [2] that, assuming the Small-Set Expansion Hypothesis and the Exponential Time Hypothesis, finding a constant-factor approximation for the optimal rank- k approximation error takes at least 2^{k^c} time for some constant $c > 0$. To obtain our first result, that computing a matrix \widehat{A} with rank at most k such that

$$\|\widehat{A} - A\|_p \leq O(1) \min_{A_k \text{ rank } k} \|A_k - A\|_p + \frac{1}{2^{\text{poly}(k)}} \|A\|_p$$

requires 2^{k^c} time for $p \in (1, 2)$ with the same hardness assumptions, we show that the reduction of [2] from the Small Set Expansion problem to ℓ_p Low Rank Approximation can be performed in such a way that each entry of the input matrix A ultimately has $\text{poly}(k)$ bits in both its numerator and denominator. If this holds, then we can assume without loss of generality that the entries of A are in fact integers with at most $\text{poly}(k)$ bits, meaning $\|A\|_p \leq 2^{\text{poly}(k)} \min_{A_k \text{ rank } k} \|A_k - A\|_p$, and the above guarantee is in fact equivalent to obtaining an $O(1)$ -approximation.

The hardness results of [2], and our above hardness result, apply to ℓ_p low rank approximation for $p \in (1, 2)$, but not to ℓ_1 low rank approximation. Intuitively, this is because the reduction of [2] ultimately shows that finding the best rank- $(d - 1)$ approximation to an $n \times d$ matrix is NP-hard, assuming the Small Set Expansion Hypothesis. However, this is not true in the $p = 1$ case: finding the best rank- $(d - 1)$ subspace can actually be done in polynomial time when $p = 1$ [7, 47]. Instead, we show that a similar hardness result holds for the somewhat more general constrained ℓ_1 low rank approximation problem. We do this by combining the techniques of [2] with the following theorem:

Theorem 1.5 (Embedding ℓ_p^n into $\ell_1^{n^{O(\log n)}}$ Deterministically). *Let $n \in \mathbb{N}$, and $p \in (1, 2)$. Then, there exists a matrix $R \in \mathbb{R}^{k \times n}$, for some $k = n^{O(\log n)}$, such that for all $x \in \mathbb{R}^n$,*

$$\Omega(1)\|x\|_p \leq \|Rx\|_1 \leq O(1)\|x\|_p$$

R can be computed deterministically in $n^{O(\log n)}$ time.

Below, we describe how we construct this embedding, and compare it with existing results related to embeddings of ℓ_p^n into ℓ_1^m , for different values of m . Using this embedding, we first adapt the techniques of [2] to show a similar hardness result for $\ell_{1,p}$ norm low-rank approximation:

Theorem 1.6 (Hardness for Sum of Column ℓ_p Norms). *Let $A \in \mathbb{R}^{n \times d}$, $k \in \mathbb{N}$ and $p \in (1, 2)$. Then, assuming the Small-Set Expansion Hypothesis and the Exponential-Time Hypothesis, at least $2^{k^{\Omega(1)}}$ time is required to find a matrix $\widehat{A} \in \mathbb{R}^{n \times d}$ of rank at most k such that*

$$\|\widehat{A} - A\|_{1,p} \leq O(1) \min_{A_k \text{ rank } k} \|A_k - A\|_{1,p} + \frac{1}{2^{\text{poly}(k)}} \|A\|_{1,p}$$

Algorithm 2 Obtaining a matrix \widehat{A} such that $\|\widehat{A} - A\|_1 \leq (1 + \varepsilon)OPT_{A,k} + \frac{\varepsilon}{f}\|A\|_1$, where $OPT_{A,k} = \min_{A_k \text{ rank } k} \|A - A_k\|_1$. We first guess a sketched left factor SU , then find an appropriate right factor V with norm at most $\text{poly}(k)$. The argument ε is assumed to be at most c for some absolute constant c .

Require: $A \in \mathbb{R}^{n \times d}$, $k \in \mathbb{N}$, $\varepsilon \in (0, c)$, $f > 1$, $\widehat{OPT} \geq 0$

Ensure: $U \in \mathbb{R}^{n \times k}$, $V \in \mathbb{R}^{k \times d}$

procedure GUESSINGADDITIVEEPSAPPROXIMATION($A, k, \varepsilon, f, \widehat{OPT}$)

If A has rank k , return A .

$r \leftarrow O(\max(k/\varepsilon^6 \log(k/\varepsilon), 1/\varepsilon^9))$

$q \leftarrow \text{poly}(k)$

$S \leftarrow$ An $r \times n$ matrix of i.i.d. standard Cauchy random variables

$\mathcal{I} \leftarrow \{0\} \cup \left\{ \sigma \cdot \left(1 + \frac{1}{f \text{poly}(k/\varepsilon)}\right)^t \mid t \in \mathbb{Z}, \frac{1}{f \text{poly}(k/\varepsilon)}\|A\|_1 \leq \left(1 + \frac{1}{f \text{poly}(k/\varepsilon)}\right)^t \leq \text{poly}(k/\varepsilon)\|A\|_1, \sigma = \pm 1 \right\}$

$\mathcal{C} \leftarrow \left\{ M \in \mathbb{R}^{r \times k} \mid M_{i,j} \in \mathcal{I} \forall i \in [r], j \in [k] \right\}$ — This is the set of (sketched) left factors we will guess.

// Guess possible rounded (sketched) left factors and find a good right factor V , with $\|V\|_1 \leq \text{poly}(k)$.

$U_{best} \leftarrow 0, V_{best} \leftarrow 0$

for $M \in \mathcal{C}$ **do**

$\text{COSTBOUNDS} \leftarrow \left\{ \frac{\varepsilon^2 \|A\|_1}{fd} \leq c \leq O(\|A\|_1) \text{ and } c \text{ is an integer power of } (1 + \varepsilon) \right\}$

for $i \in [d]$, $c \in \text{COSTBOUNDS}$ **do**

$V_{i,c} \leftarrow \text{argmin}_{V_i} \|V_i\|_1$ subject to the constraint that $\text{med}(MV_i - SA_i) \leq c$

$C_{i,c} \leftarrow \text{med}(MV_{i,c} - SA_i)$

end for

// Create LP to find a good distribution over $c \in \text{COSTBOUNDS}$ for each $i \in [d]$.

$\text{VARIABLES} \leftarrow \{x_{i,c} \mid i \in [d], c \in \text{COSTBOUNDS}\}$

$\text{CONSTRAINTS} \leftarrow \left\{ 0 \leq x_{i,c} \mid i \in [d], c \in \text{COSTBOUNDS} \text{ and } \sum_{c \in \text{COSTBOUNDS}} x_{i,c} = 1 \mid i \in [d] \right\}$

$\text{CONSTRAINTS} \leftarrow \text{CONSTRAINTS} \cup \left\{ \sum_{i \in [d], c \in \text{COSTBOUNDS}} x_{i,c} \|V_{i,c}\|_1 \leq kq = \text{poly}(k) \right\}$

$\Delta \leftarrow (1 + O(\varepsilon))\widehat{OPT} + O\left(\frac{\varepsilon^2}{f}\right)\|A\|_1$

$\text{CONSTRAINTS} \leftarrow \text{CONSTRAINTS} \cup \left\{ \sum_{i \in [d], c \in \text{COSTBOUNDS}} x_{i,c} C_{i,c} \leq \Delta \right\}$

$x_{i,c} \leftarrow$ Solution to the LP given by VARIABLES and CONSTRAINTS , for all $i \in [d]$, $c \in \text{COSTBOUNDS}$

If LP is infeasible, then **continue** to next $M \in \mathcal{C}$.

// For each column, sample an appropriate cost bound. Do this $O(1/\varepsilon)$ times, then V' meets both // the cost and norm constraints with constant probability.

for $t = 1 \rightarrow 10/\varepsilon$ **do**

$c_i \leftarrow$ An element $c \in \text{COSTBOUNDS}$ sampled according to the distribution on COSTBOUNDS given by $\{x_{i,c} \mid c \in \text{COSTBOUNDS}\}$

$V'_i \leftarrow V_{i,c_i}$ for all $i \in [d]$

Break if $\|V'\|_1 \leq \frac{2kq}{\varepsilon}$ and $\text{med}(MV' - SA) \leq (1 + 2\varepsilon)\Delta$

end for

$U' \leftarrow \text{argmin}_U \|UV' - A\|_1$

If $\|U'V' - A\|_1 \leq \|U_{best}V_{best} - A\|_1$ then $U_{best} \leftarrow U'$ and $V_{best} \leftarrow V'$.

end for

return U', V'

end procedure

Algorithm 3 Obtaining a $(1 + \varepsilon)$ -approximation with bicriteria rank at most $3k$. First apply POLYKERRORNOTBI CRITERIA APPROXIMATION from Algorithm 8 to A to obtain a $\text{poly}(k)$ -approximation B — then, apply Algorithm 2 to the residual to obtain an approximation UV with additive error $\varepsilon/\text{poly}(k)\|A - B\|_1$. Finally, $B + UV$ gives a $(1 + \varepsilon)$ -approximation with rank $3k$.

Require: $A \in \mathbb{R}^{n \times d}$, $k \in \mathbb{N}$, $\varepsilon \in (0, c)$

Ensure: $\hat{A} \in \mathbb{R}^{n \times d}$ having rank $3k$

```

procedure ROUNDINGGUESSINGEPSAPPROXIMATION( $A, k, \varepsilon$ )
   $W, Z \leftarrow$  POLYKERRORNOTBI CRITERIA APPROXIMATION( $A, k$ )
   $B \leftarrow WZ$ 
   $C \leftarrow A - B$ 
   $f \leftarrow \text{poly}(k)$ , the approximation factor of Algorithm 8

  // Guess all  $O((\log nd)/\varepsilon)$  possible values for  $\widehat{OPT}$  and try them for Algorithm 2
  // as described in Remark 1.3.
   $C_{SVD,2k} \leftarrow$  The optimal rank- $2k$  approximation for  $C$  under the  $\ell_2$  norm.
   $SVDERROR \leftarrow \|C - C_{SVD,2k}\|_1$ 
   $\hat{A} \leftarrow 0 \in \mathbb{R}^{n \times d}$ 
  for  $t = 0 \rightarrow O(\frac{\log nd}{\varepsilon})$  do
     $\widehat{OPT} \leftarrow SVDERROR/(1 + \varepsilon)^t$ 
     $U, V \leftarrow$  GUESSINGADDITIVEEPSAPPROXIMATION( $C, 2k, \varepsilon, f, \widehat{OPT}$ )
    If  $\|(B + UV) - A\|_1 \leq \|\hat{A} - A\|_1$  then  $\hat{A} \leftarrow B + UV$ 
  end for
return  $\hat{A}$ 
end procedure

```

where for a matrix $M \in \mathbb{R}^{n \times d}$, its $\ell_{1,p}$ norm is $\|M\|_{1,p} = \sum_{j=1}^d \|M_j\|_p$.

In fact, finding the best rank- k approximation to A in the $\ell_{1,p}$ norm requires $2^{k^{\Omega(1)}}$ time even in the special case when $A \in \mathbb{R}^{(k+1) \times k^{O(\log k)}}$ — these are the dimensions of the hard instances for the above theorem. Thus, to reduce from $\ell_{1,p}$ low rank approximation (on this family of hard instances A) to constrained ℓ_1 low rank approximation, we can simply multiply on the left by an embedding matrix $R \in \mathbb{R}^{k^{O(\log k)} \times (k+1)}$. In other words, to find \hat{A} such that

$$\|\hat{A} - A\|_{1,p} \leq O(1) \min_{A_k \text{ rank } k} \|A_k - A\|_{1,p} + \frac{1}{2^{\text{poly}(k)}} \|A\|_{1,p}$$

it suffices to find $\hat{A} \in \mathbb{R}^{(k+1) \times k^{O(\log k)}}$ with rank at most k such that

$$\|R\hat{A} - RA\|_1 \leq O(1) \min_{A_k \text{ rank } k} \|RA_k - RA\|_1 + \frac{1}{2^{\text{poly}(k)}} \|RA\|_1$$

This is exactly constrained ℓ_1 low rank approximation: the input matrix is RA , and the goal is to find a matrix M of rank at most k , such that its columns are in the column span of R (i.e. can be written as $R\hat{A}$ for \hat{A} with rank at most k). The overall running time of the reduction is $k^{O(\log k)} = 2^{O(\log^2 k)} \ll 2^{k^{\Omega(1)}}$, meaning the same running time lower bound applies for constrained ℓ_1 low rank approximation:

Theorem 1.7 (Hardness for Constrained ℓ_1 Low Rank Approximation). *Let $A \in \mathbb{R}^{m \times d}$, $R \in \mathbb{R}^{m \times n}$ and $k \in \mathbb{N}$. Assuming the Small-Set Expansion Hypothesis and Exponential-Time Hypothesis, at least $2^{k^{\Omega(1)}}$ time is required to find a matrix $\hat{A} \in \mathbb{R}^{m \times d}$ with rank at most k , such that the columns of \hat{A} are in the column span of R and, with constant probability,*

$$\|\hat{A} - A\|_1 \leq O(1) \min_{A_k} \|A_k - A\|_1 + \frac{1}{2^{\text{poly}(k)}} \|A\|_1$$

The minimum on the right-hand side is taken over all matrices A_k with rank at most k , whose columns are contained in the span of R .

Note that Algorithm 2 can be modified to work for constrained ℓ_1 low rank approximation, with the same running time guarantee. The only modification to the algorithm itself would be that, once the right factor V' is obtained, then the left factor U is set to be $R \cdot (\operatorname{argmin}_{U_0} \|RU_0V' - A\|_1)$ instead of $\operatorname{argmin}_U \|UV' - A\|_1$, where R is the subspace to which the low-rank solution is constrained — note that this new minimizer can also be found using linear programming. The only change that will be made to the analysis in the introduction and in Theorem 3.15 is that U^* will be the optimal rank- k left factor whose columns are contained in the span of R , rather than the optimal (un-constrained) rank- k left factor.

Remark 1.8. *This reduction corrects an error in an earlier version of this work. Previously we claimed that ℓ_p low rank approximation could be directly reduced to constrained ℓ_1 low rank approximation through the use of an embedding matrix. However, this line of reasoning leads to a reduction from ℓ_p low rank approximation to constrained $\ell_{p,1}$ low rank approximation, and it is unclear how to obtain a reduction from constrained $\ell_{p,1}$ low rank approximation to constrained ℓ_1 low rank approximation. In this version, we fix this issue by adding the intermediate step of reducing to $\ell_{1,p}$ low rank approximation.*

Deterministic Embedding of ℓ_p into ℓ_1 We deterministically construct the embedding matrix mentioned in Theorem 1.5, using an observation from [38]:

Fact 1 (Observation 2.7.2 of [38]). *Let R_1, R_2, \dots, R_n be real random variables on a probability space that has k elements $\omega_1, \omega_2, \dots, \omega_k$, and let $A \in \mathbb{R}^{k \times n}$ such that $A_{i,j} = \Pr[\omega_i]R_j(\omega_i)$. For $x \in \mathbb{R}^n$ let us set $X := \sum_{j=1}^n R_j x_j$. Then, $E[|X|] = \|Ax\|_1$.*

By this observation, if $E[|X|] = \Theta(1)\|x\|_p$, then the corresponding matrix A gives an embedding of the ℓ_p norm into the ℓ_1 norm. For this, we can select the random variables R_1, R_2, \dots, R_n in the observation above to be standard p -stable random variables, for $p \in (1, 2)$ — to show that $E[|X|] = \Theta(1)\|x\|_p$, we can make use of results on p -stable random variables from [48]. For the sample space to have a finite number of elements k , we can round the R_i to the nearest multiple of $\operatorname{poly}(n)$ and truncate them so that they are not more than $\operatorname{poly}(n)$.

However, if we let the R_i be fully independent, then the sample space may have a very large size k (i.e. it may be exponential in n). Instead, we let the R_i be $O(\log n)$ -wise independent — by a result of [32], the R_i can be constructed using a sample space of size $k = n^{O(\log n)}$, and we can use lemmas from [27] and [26] on p -stable random variables with limited independence to show that $E[|X|] = \Theta(1)\|x\|_p$ even when the R_i are $O(\log n)$ -wise independent.

Remark 1.9. *In an earlier version of this work, we instead used the randomized embedding of [25]. To obtain a running time lower bound when this randomized embedding is used, we assumed a randomized version of ETH that was previously used in [16], for instance. In the current version, we use the deterministic embedding mentioned in Theorem 1.5 in its place. The randomized embedding had the advantage that it provided an embedding of ℓ_p^n into $\ell_1^{O(n)}$, rather than ℓ_p^n into $\ell_1^{n^{O(\log n)}}$. One issue with our previous proof is that we assumed that the entries of the embedding of [25] can be computed in $\operatorname{poly}(n)$ time — while we believe this is true, it is not immediately clear how to prove this. Another work [28] deterministically constructs $(1 - \varepsilon)n$ -dimensional subspaces of \mathbb{R}^n in which the ℓ_1 norm and the ℓ_p norm are the same up to normalization and constant factors. However, [28] does not provide a matrix R which can be used to embed $\mathbb{R}^{(1-\varepsilon)n}$ with the ℓ_p norm into \mathbb{R}^n with the ℓ_1 norm with low distortion, and hence is not applicable for our reduction.*

1.3 Paper Outline

1.3.1 Main Results - Algorithms

In Section 2, we describe our $\tilde{O}(k^{1/p-1/2})$ -approximate algorithm for ℓ_p column subset selection for $p \in [1, 2)$, and its analysis. In Section 3, we analyze our $(1 + \varepsilon)$ -approximation algorithm for ℓ_p low rank approximation, for $p \in [1, 2)$, which returns a matrix of rank at most $3k$ in $2^{\operatorname{poly}(k/\varepsilon)} + \operatorname{poly}(nd)$ time.

1.3.2 Appendices - Additional Results

In Appendix A, we show how to extend the hardness results of [2] to show that even obtaining an $O(1)$ -approximation for ℓ_p low rank approximation with $\frac{1}{2^{\operatorname{poly}(k)}}\|A\|_p$ additive error is hard, when $p \in (1, 2)$.

Using this, we show that obtaining an $O(1)$ -approximation for constrained ℓ_1 low rank approximation with $\frac{1}{2^{\text{poly}(k)}}\|A\|_1$ additive error is hard. Next, in Appendix B we show how the lower bound for ℓ_1 column subset selection due to [48] can be extended to obtain a lower bound for ℓ_p column subset selection, for $p \in (1, 2)$. Finally, in Appendix C, we show how to obtain a $\text{poly}(k)$ -approximation algorithm with running time $\text{poly}(nd)$, based on Algorithm 4 and techniques from [11, 48].

2 Optimal ℓ_1 Column Subset Selection via Random Sampling

2.1 Preliminaries: Notation and Lewis Weight Sampling

Suppose $A \in \mathbb{R}^{n \times d}$. We use the following notation for the rows, columns, and submatrices of A . For $i \in [n]$, we let A^i be the i^{th} row of A , and for $j \in [d]$, we let A_j be the j^{th} column of A . Moreover, for $S \subset [d]$, we let A_S be the submatrix of A , such that its columns are those of A whose indices are in S , and for $R \subset [n]$, we let A^R be the submatrix of A , such that its rows are those of A whose indices are in R .

We now recall basic facts about row sampling with Lewis weights. We use Lewis weight sampling as a black box, for more details we refer the reader to [13, 48].

Lemma 2.1 (Sampling and Rescaling Matrix Based on Lewis weights - Adapted from Theorem 7.1 of [13]). *Let $A \in \mathbb{R}^{n \times d}$, $1 \leq p < 2$, and let $r = O(d \log d)$ if $p = 1$ and $r = O(d \log d (\log \log d)^2)$ otherwise. There exists a distribution (p_1, p_2, \dots, p_n) on the rows of A such that if S is a matrix with r rows, each row chosen independently as the i^{th} standard basis vector times $\frac{1}{(rp_i)^{\frac{1}{p}}}$ with probability p_i , then with probability 0.9,*

$$\Omega(1)\|Ax\|_p \leq \|SAx\|_p \leq O(1)\|Ax\|_p$$

for all $x \in \mathbb{R}^d$. The distribution (p_1, p_2, \dots, p_n) can be computed in $\text{nnz}(A) + \text{poly}(d)$ time.

Lemma 2.2 ($O(1)$ Dilation and Contraction for Lewis Weights - Lemma D.11 ($p = 1$) and E.11 ($p \in (1, 2)$) of [48]). *Let $M \in \mathbb{R}^{n \times d}$ and $U \in \mathbb{R}^{n \times t}$. Let $r = O(t \log t)$ if $p = 1$ and $O(t \log t (\log \log t)^2)$ if $p \in (1, 2)$, and suppose $S \in \mathbb{R}^{r \times n}$ is a sampling and rescaling matrix generated according to the ℓ_p Lewis weights of U . Then, with probability 0.9, $\|SM\|_p^p \leq O(1)\|M\|_p^p$, and with probability 0.9, for all $x \in \mathbb{R}^t$, $\|SUX\|_p^p \geq \Omega(1)\|UX\|_p^p$.*

Lemma 2.3 ($O(1)$ Contraction on Affine Subspace for Lewis Weights - From Lemmas D.11 and D.7 of [48] for $p = 1$, and Lemmas E.11 and E.7 of [48] for $p \in (1, 2)$). *Let $A \in \mathbb{R}^{n \times d}$, $U \in \mathbb{R}^{n \times k}$, and let $V^* = \text{argmin}_{V \in \mathbb{R}^{k \times d}} \|UV - A\|_p$. Let $r = O(k \log k)$ if $p = 1$ and $r = O(k \log k (\log \log k)^2)$ if $p \in (1, 2)$. Let $S \in \mathbb{R}^{r \times n}$ be a sampling and rescaling matrix generated according to the ℓ_p Lewis weights of U . Then, with probability 0.9, simultaneously for all $V \in \mathbb{R}^{k \times d}$,*

$$\|SUV - SA\|_p^p \geq \Omega(1)\|UV - A\|_p^p - O(1)\|UV^* - A\|_p^p$$

2.2 ℓ_p Column Subset Selection Algorithm and Analysis

In this section, we analyze Algorithm 4. Note that the algorithm itself is nearly the same as Algorithm 1 of [51]⁷, with the difference being that $2r = O(k \log k)$ columns are sampled per iteration, instead of $2k$ — however, its analysis is significantly different, and does not use maximum-determinant column subsets. Instead, we rely on the following key existence result, shown in Theorem C.1 of [48]:

Theorem 2.4. (*Existence of a Good Column Subset*) *Let $A \in \mathbb{R}^{n \times d}$, $p \in [1, 2)$, $k \in \mathbb{N}$, and $r = O(k \log k)$ if $p = 1$ and $r = O(k \log k (\log \log k)^2)$ if $1 < p < 2$. Then, there exist matrices $U \in \mathbb{R}^{n \times r}$ and $V \in \mathbb{R}^{r \times d}$, such that the columns of U are columns of A , and*

$$\|UV - A\|_p \leq O\left(r^{\frac{1}{p} - \frac{1}{2}}\right) \min_{A_k \text{ rank } k} \|A - A_k\|_p$$

Moreover, this occurs with probability $\frac{9}{10}$ if the columns of U are sampled from the columns of A according to the Lewis weights of V^* , where U^*V^* is an optimal rank- k approximation to A .

Proof. This result was essentially shown in Theorem C.1 of [48] for the case $p = 1$ — we give a brief sketch here. The key point is that $U = AR$, where R is given by (column) Lewis weights of V^* (here we first generate a sampling matrix based on the Lewis weights of $(V^*)^T$, then transpose the sampling matrix).

Indeed, if R is chosen according to the column Lewis weights of V^* , then letting U^i be the minimizer of $\|U^iV^*R - A^iR\|_2$ instead of $\|U^iV^*R - A^iR\|_1$ for all $i \in [n]$ gives an $O(r^{\frac{1}{2}})$ -approximation. This minimizer

⁷Our pseudocode in Algorithm 4 is based on Algorithm 2 of <https://arxiv.org/pdf/1811.01442v1.pdf>, which is version 1 of [51] on arXiv.

U^{ii} is given by $U^{ii} = A^i R (V^* R)^+$, where $(V^* R)^+$ is the pseudo-inverse of $V^* R$. Hence we can choose $U' = AR(V^* R)^+$, and hence there exists an $O(r^{\frac{1}{2}})$ -approximation for A with left factor AR .

For $1 < p < 2$, the proof is nearly identical [48]. Instead of Lemma D.11 of [48] we can use part (III) of Lemma E.11 of the same work to extend the result to $p \neq 1$. Similarly, instead of Lemma D.8, we can use Lemma E.8 of [48]. Finally, we apply Lemma B.10 of [48] to convert between the ℓ_p -norm and the ℓ_2 -norm, rather than between the ℓ_1 -norm and the ℓ_2 -norm as was done in [48] — now we obtain a distortion of $r^{\frac{1}{p} - \frac{1}{2}}$. \square

Algorithm 4 Randomly sample columns of A repeatedly, to obtain $O(k \cdot \text{polylog}(k) \log(d))$ columns of A spanning a good approximation. This is a variant of Algorithm 1 in [51], with the difference being that we sample $k \cdot \text{polylog}(k)$ columns in each round instead of $2k$ columns. Here, MULTIPLE REGRESSION-SOLVER(n, d, m, U, B) (where $U \in \mathbb{R}^{n \times d}$, $B \in \mathbb{R}^{n \times m}$) is a subroutine which computes $\min_x \|Ux - B\|_p^p$ for each $j \in [m]$. The call BOTTOMK(SORT($cost$), $\Omega(m)$) serves to find the $\Omega(m)$ column indices in $[m]$ having the smallest regression cost.

Require: $A \in \mathbb{R}^{n \times d}$, $k \in \mathbb{N}$, $p \in [1, 2)$

Ensure: $S \subset [d]$, $|S| = O(k \log k \log d)$ if $p = 1$ and $O(k \log k (\log \log k)^2 \log d)$ otherwise

procedure RANDOMCOLUMNSUBSETSELECTION(A, k, p)

 samples $\leftarrow O(\log d)$

$T_0 \leftarrow [d]$

$r \leftarrow O(k \log k)$ if $p = 1$ and $O(k \log k (\log \log k)^2)$ otherwise

for $i = 1 \rightarrow \text{samples}$ **do**

$m \leftarrow |T_{i-1}|$

for $j = 1 \rightarrow O(\log d)$ **do**

 Sample $S^{(j)}$ from $\binom{T_{i-1}}{2r}$ uniformly at random

$m \leftarrow |T_{i-1} \setminus S^{(j)}|$

$\{\text{cost}_t\}_{t \in T_{i-1} \setminus S_j} \leftarrow \text{MULTIPLE REGRESSION SOLVER}(n, 2r, m, A_{S^{(j)}}, A_{T_{i-1} \setminus S^{(j)}})$

$R^{(j)} \leftarrow \text{BOTTOMK}(\text{SORT}(\text{cost}), \Omega(m))$

$c_j \leftarrow \sum_{t \in R^{(j)}} \text{cost}_t$

end for

$j^* \leftarrow \min_{j \in [O(\log d)]} c_j$

$S_i \leftarrow S^{(j^*)} \cup R^{(j^*)}$

$T_i \leftarrow T_{i-1} \setminus S_i$

end for

$S \leftarrow \cup_i S_i$

return $S, S_1, S_2, \dots, S_{O(\log d)}$

end procedure

Now, we analyze Algorithm 4:

Theorem 2.5 (Column Sampling Approximation Factor). *Let $A \in \mathbb{R}^{n \times d}$, $p \in [1, 2)$, $k \in \mathbb{N}$, and $r = O(k \log k)$ if $p = 1$ and $r = O(k \log k (\log \log k)^2)$ otherwise. Let $S \subset [d]$ be the set of columns of A that is returned by Algorithm 4. Let $U = A_S \in \mathbb{R}^{n \times O(r \log d)}$, and let $V = \text{argmin}_{V' \in \mathbb{R}^{O(r \log d) \times d}} \|UV' - A\|_p$, i.e., V is the optimal right factor for A_S . Then, with probability $1 - o(1)$,*

$$\|UV - A\|_p \leq O(r^{\frac{1}{p} - \frac{1}{2}} (\log d)^{\frac{1}{p}}) \min_{A_k \text{ rank } k} \|A - A_k\|_p$$

Proof. We first show the following claim (whose role is analogous to Lemma 6 of [11]), which is that after sampling $2r$ columns of A , at least a constant fraction of the remaining columns are covered, up to our desired approximation factor, with constant probability. The main theorem is then a consequence of this claim, together with a Markov bound (to show that a constant fraction of the columns of A are covered with constant probability) and a union bound over the samples in each of the iterations (note that we perform $O(\log d)$ repetitions in each iteration, and use the result of the best repetition, to boost the probability of a constant fraction of columns being covered).

Claim 2.6. Let $B \in \mathbb{R}^{n \times 2r}$ be a submatrix of A , whose columns are a uniformly random subset of those of A , of size $2r$. Furthermore, let A_i be an additional, uniformly random, column of A not among those of B . Then, with constant probability (where the probability is taken over B and A_i),

$$\min_{x \in \mathbb{R}^{2r}} \|Bx - A_i\|_p^p \leq O(r^{1-\frac{p}{2}}) \frac{OPT^p}{d}$$

where $OPT = \min_{A_k \text{ rank-}k} \|A - A_k\|_p$.

Proof. The proof is somewhat along the lines of Lemma 6 of [11], but the analysis is based on the existence result in Theorem 2.4 rather than the existence of an $O(k)$ -approximate column subset, as in [11]. Let $B' \in \mathbb{R}^{n \times (2r+1)}$ be B with A_i adjoined. Then, the columns of B' also form a uniformly random subset of the columns of A . Define $\Delta \in \mathbb{R}^{n \times d}$ so that if U^*V^* is the optimal rank- k approximation to A , then $\Delta = A - U^*V^*$. For a subset $S \subset [d]$, let Δ_S denote the submatrix of Δ containing those columns whose indices are in S . Finally, let $T \subset [d]$ such that $B' = A_T$, and let B'_k be the best rank- k approximation to B' .

Then, by the definition of B'_k , $\|B' - B'_k\|_p^p \leq \|\Delta_T\|_p^p$. Hence, taking expectations gives

$$E \left[\|B' - B'_k\|_p^p \right] \leq E \left[\|\Delta_T\|_p^p \right] = \frac{|T|}{d} OPT^p = O\left(\frac{r}{d} OPT^p\right) \quad (3)$$

where the last equality is because T is a uniformly random subset of $[d]$ of size $2r+1$. Hence, by Markov's inequality, with probability $19/20$, $\|B' - B'_k\|_p^p \leq O\left(\frac{r}{d} OPT^p\right)$ — denote this event by \mathcal{E}_1 .

Now we use Theorem 2.4, by which there exists a subset S of column indices of B' of size r such that, if we let B'_S denote the corresponding submatrix of B' , and $V = \operatorname{argmin}_{V' \in \mathbb{R}^{r \times |T|}} \|B'_S V' - B'\|_p$, then

$$\|B'_S V - B'\|_p \leq O(r^{\frac{1}{p}-\frac{1}{2}}) \|B'_k - B'\|_p$$

i.e. $\|B'_S V - B'\|_p^p \leq O(r^{1-\frac{p}{2}}) \|B'_k - B'\|_p^p$. If \mathcal{E}_1 holds, then this implies that

$$\|B'_S V - B'\|_p^p \leq O(r^{1-\frac{p}{2}}) \frac{r \cdot OPT^p}{d}$$

Finally, let $i \in T$ be a uniformly random column index — for instance, the uniformly random column of B' that is not in B . Then,

$$E \left[\|B'_S V_i - B'_i\|_p^p \mid \mathcal{E}_1 \right] \leq \frac{1}{2r+1} O(r^{1-\frac{p}{2}}) \frac{r \cdot OPT^p}{d} = O(r^{1-\frac{p}{2}}) \frac{OPT^p}{d} \quad (4)$$

and hence, assuming \mathcal{E}_1 , with probability at least $19/20$, $\|B'_S V_i - B'_i\|_p^p \leq O(r^{1-\frac{p}{2}}) OPT^p/d$. Let \mathcal{E}_2 denote the event that this holds — then, the probability of $\mathcal{E}_1 \cap \mathcal{E}_2$ is at least $18/20$.

Finally, let \mathcal{E}_3 be the event that A_i is not in the $O(r^{\frac{1}{p}-\frac{1}{2}})$ -approximate column subset S . Note that S has size r , while B' has $2r+1$ columns. Moreover, S is determined entirely by B' — therefore, since A_i is also a uniformly random column of B' , with probability at least $\frac{1}{2}$, A_i is not in S , meaning that it is covered well by B (since S is contained entirely in B).

In summary, if \mathcal{E}_1 , \mathcal{E}_2 and \mathcal{E}_3 all hold, and x is obtained by performing ℓ_p regression, using B to fit A_i , then

$$\|Bx - A_i\|_p^p \leq \|B'_S V_i - B'_i\|_p^p \leq O(r^{1-\frac{p}{2}}) \frac{OPT^p}{d}$$

Moreover, the failure probability of $\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3$ is at most $\frac{1}{2} + \frac{2}{10} = \frac{7}{10}$. This proves the lemma. \square

We can combine the above claim with a Markov bound, as follows. Let $B \in \mathbb{R}^{n \times 2r}$ be a submatrix of A , such that the column indices of B form a uniformly random subset of T_{i-1} (here, we are using the notation of Algorithm 4 — T_i is the set of indices of columns of A which have not been discarded after i iterations). For $i \in [m]$, let Z_i be equal to 1 if A_i is approximately covered by B , i.e.,

$$\min_{x \in \mathbb{R}^{2r}} \|Bx - A_i\|_p^p \leq O\left(r^{1-\frac{p}{2}} \frac{OPT^p}{d}\right)$$

and 0 otherwise. Then, $E_{B,i}[Z_i] = c$, where $0 < c < 1$, and the expectation is taken over uniformly random B and i . Hence, if we let $Z := \sum_i Z_i$ be the number of approximately covered columns, then $E[Z] = cm$, and $E[m-Z] = (1-c)m$. Therefore, by Markov's inequality, with constant probability $F > 0$, $(m-Z) \leq (1-\frac{c}{2})m$, meaning $Z \geq \frac{c}{2}m$. In other words, with probability F , there exists $R \subset T_{i-1}$ with $|R| \geq \frac{c}{2}|T_{i-1}|$ such that for each $j \in R$,

$$\min_{V_j \in \mathbb{R}^k} \|BV_j - A_j\|_p^p \leq O\left(r^{1-\frac{p}{2}} \frac{OPT^p}{|T_{i-1}|}\right)$$

meaning

$$\min_{V \in \mathbb{R}^{k \times |R|}} \|BV - A_R\|_p^p \leq O\left(r^{1-\frac{p}{2}} \frac{OPT^p \cdot |T_{i-1}|}{|T_{i-1}|}\right) = O\left(r^{1-\frac{p}{2}} OPT^p\right)$$

Denote this event by \mathcal{E} , meaning \mathcal{E} occurs with probability F . Then, as done in Algorithm 4, it is sufficient to sample the submatrix B in $O(\log d)$ independent iterations, and take the sampled submatrix which minimizes the sum of the lowest $\frac{c}{2}|T_{i-1}|$ residuals. By choosing B in this way, we ensure that \mathcal{E} has failure probability at most $(1-F)^{O(\log d)} = \frac{1}{d^{O(1)}}$. Since we perform $O(\log d)$ iterations, we only have to perform a union bound over $O(\log d)$ such events, meaning the overall failure probability of Algorithm 4 is $O(\log d) \cdot \frac{1}{d^{O(1)}} = o(1)$. \square

Remark 2.7. Note that our application of linearity of expectation in the above proof is valid. For a fixed $i \in [m]$, it may not seem that we can take the expectation $E_{B,i}[Z_i]$ over a uniformly random column index i , since i is determined by Z_i . However, we could for instance shuffle the indices $i \in [m]$, and sum the Z_i in the shuffled order — then, each index $i \in [m]$ is a uniformly random column index, and we can use linearity of expectation.

We can also remove the $O((\log d)^{\frac{1}{p}})$ term in the approximation factor from Theorem 2.5 with the following refined analysis, reminiscent of one performed in [51] — we examine the number of times a column of A can remain “uncovered” before it is removed:

Theorem 2.8 (Column Sampling - Better Approximation Factor). *Let $A \in \mathbb{R}^{n \times d}$, $p \in [1, 2)$, $k \in \mathbb{N}$ and $r = O(k \log k)$ if $p = 1$ and $r = O(k \log k (\log \log k)^2)$ otherwise. Let $S \subset [d]$ be the set of columns of A that is returned by Algorithm 4. Let $U = A_S \in \mathbb{R}^{n \times O(r \log d)}$ and let $V = \operatorname{argmin}_{V' \in \mathbb{R}^{O(r \log d) \times d}} \|UV' - A\|_p$, i.e., V is the optimal right factor for A_S . Then, with probability $1 - o(1)$,*

$$\|UV - A\|_p \leq O\left(r^{\frac{1}{p}-\frac{1}{2}} (\log k)^{\frac{1}{p}}\right) \min_{A_k \text{ rank } k} \|A - A_k\|_p$$

Proof. Define Δ as in the proof of Theorem 2.5. Consider the i^{th} iteration of Algorithm 4, and recall that T_{i-1} is the set of remaining column indices. Let $m = |T_{i-1}|$. Finally, let $T_{i-1, \text{big}} \subset T_{i-1}$ consist of the $\frac{m}{r}$ indices of T_{i-1} with greatest cost, rounding down if $\frac{m}{r}$ is not an integer (i.e. $j \in T_{i-1, \text{big}}$ if $\|\Delta_j\|_p$ is among the top $\frac{m}{r}$ column norms of Δ).

Note that with constant probability, $S^{(j)}$ will be disjoint from $T_{i-1, \text{big}}$ (provided r is multiplied by a sufficiently large constant and $m \geq \Omega(r)$), since the probability of not selecting an element of $T_{i-1, \text{big}}$ in a uniformly random subset of size $2r$ is at least $(1 - O(\frac{1}{r}))^{2r} \geq \Omega(1)$ for r sufficiently large. Since we take $O(\log d)$ samples per iteration and choose the best one, this occurs on all iterations with probability $1 - o(1)$.

Condition on this event (which we can call \mathcal{E}_1) occurring — then, for uniformly random $j \in T_{i-1} \setminus T_{i-1, \text{big}}$, with constant probability, if $B = A_{S^{(j)}}$,

$$\min_{x \in \mathbb{R}^{2r}} \|Bx - A_j\|_p^p \leq O\left(r^{1-\frac{p}{2}} \frac{OPT_{T_{i-1} \setminus T_{i-1, \text{big}}}^p}{m}\right)$$

This is by Claim 2.6, since $S^{(j)}$ is a uniformly random subset of $T_{i-1} \setminus T_{i-1, \text{big}}$. Therefore, conditioning on \mathcal{E}_1 , with constant probability, the smallest $\Omega(m)$ columns have a cost of $O\left(r^{1-\frac{p}{2}} OPT_{T_{i-1} \setminus T_{i-1, \text{big}}}^p\right)$. Finally, \mathcal{E}_1 occurs in the i^{th} iteration with probability at least $1 - \frac{1}{\text{poly}(d)}$ since we repeat the sampling process $O(\log d)$ times and take the sample which gives the smallest cost on the lowest $\Omega(m)$ columns — hence, with probability $1 - \frac{1}{\text{poly}(d)}$, the smallest $\Omega(m)$ columns have a cost of $O\left(r^{1-\frac{p}{2}} OPT_{T_{i-1} \setminus T_{i-1, \text{big}}}^p\right)$, and this occurs on every iteration i with probability $1 - o(1)$ by a union bound (since there are $O(\log d)$ iterations).

It remains to bound $\sum_i OPT_{T_{i-1} \setminus T_{i-1, big}}^p$ where i ranges across all of the iterations — we show that

$$\sum_i OPT_{T_{i-1} \setminus T_{i-1, big}}^p = O((\log k)OPT^p)$$

First assume that on each iteration i , the column indices j which are “covered” (and hence discarded) have the smallest $\|\Delta_j\|_p$. If we remove this assumption, then this can only decrease $\sum_i OPT_{T_{i-1} \setminus T_{i-1, big}}^p$. To see why, suppose i is the index of the latest iteration on which this does not occur, meaning there exist column indices j_1 and j_2 such that in iteration i , column j_1 is not removed and column j_2 is removed, and $\|\Delta_{j_1}\|_p < \|\Delta_{j_2}\|_p$. Let (A_1, A_2, \dots, A_t) be the $\|\Delta_j\|_p$ s of the remaining columns if column j_2 is removed, and (B_1, B_2, \dots, B_t) be the $\|\Delta_j\|_p$ s of the remaining columns if column j_1 is removed instead. Then, for all $j \in [t]$, $A_j \leq B_j$, and hence, if on all subsequent iterations, the columns are removed in order of their $\|\Delta_j\|_p$ s, then this will only increase $OPT_{T_{i-1} \setminus T_{i-1, big}}$ for all subsequent i .

We can therefore argue that on the last iteration (and by recursing, on all iterations) where the columns are not removed in order of their $\|\Delta_j\|_p$ s, removing them in order of their $\|\Delta_j\|_p$ s instead can only increase $\sum_i OPT_{T_{i-1} \setminus T_{i-1, big}}^p$. Hence, we can assume without loss of generality that in iteration i , the indices j in T_{i-1} with the $\Omega(m)$ lowest values of $\|\Delta_j\|_p$ will be removed. If this holds, then for $j \in T_{i-1} \setminus T_{i-1, big}$, after $O(\log r)$ iterations, i.e. for iteration i' where $i' \geq i + O(\log r)$, j will not be in $T_{i'-1}$, since $T_{i'-1}$ will be equal to $T_{i-1, big}$ (since a constant fraction of the columns in T_{i-1} are removed in each iteration). \square

3 $(1 + \varepsilon)$ -Approximation in FPT Time with Bicriteria Rank $3k$

In this section, we give an algorithm for ℓ_p -low rank approximation, for $p \in [1, 2)$, which runs in $2^{\text{poly}(k/\varepsilon)} + \text{poly}(nd)$ time and outputs a matrix of rank $3k$. Our algorithm discussed in Subsubsection 1.2.2 is a special case of this algorithm, in the case $p = 1$. The analysis here is similar to the analysis given in the introduction, but uses sketching matrices whose entries are p -stable random variables, rather than Cauchy random variables.

3.1 Preliminaries: Median-Based Estimator for ℓ_p -norm Dimension Reduction from [2]

We first recall some concepts from [2] related to sketches based on medians and dense p -stable random matrices (Section 2 of [2] for the $p = 1$ case and Subsection 3.2 of [2] for the $p \in (1, 2)$ case). In the following, we let $B \in \mathbb{R}^{n \times d}$.

Definition 3.1 (p -Stable Random Variables - As Defined in Section 3.2 of [2]). *Suppose Z, Z_1, Z_2, \dots, Z_n are i.i.d. random variables, and $p \in [1, 2]$. We say Z and Z_i are p -stable if, for any $x \in \mathbb{R}^n$, $\|x\|_p Z$ and $\sum_{i=1}^n x_i Z_i$ have the same distribution.*

Note that p -stable random variables only exist for $p \in (0, 2]$ — we consider $p \in [1, 2)$. 1-stable random variables are also called Cauchy random variables, as in our description of our algorithm in the introduction. For $p \in [1, 2)$, we use med_p to denote the median of a half- p -stable random variable — that is, if Z is a p -stable random variable, then $|Z|$ is a half- p -stable random variable. Note that the median of a half-Cauchy random variable is just 1, but for $p \in (1, 2)$, there is no simple closed form for med_p . med_p can be computed up to a $(1 \pm \varepsilon)$ -factor, as described in [27] — this is enough for our purposes. This definition is relevant for the median-based sketch of [2], as we see below.

Definition 3.2 (Medians and Quantiles of Vectors (Definition 4 in [2])). *For a vector $v \in \mathbb{R}^n$, we let $\text{med}(v)$ be the median of $|v_i|$ for $i \in [n]$. In addition, for $\alpha \in [0, 1]$, we let $q_\alpha(v)$ denote the minimum value greater than $\lceil \alpha n \rceil$ of the values $|v_1|, |v_2|, \dots, |v_n|$.*

The following lemmas from [2] allow us to obtain very accurate estimates of the ℓ_p -norms of matrices after first multiplying by a dense p -stable matrix to reduce the dimension.

Lemma 3.3 (p -stable Matrix + Median Preserves Norms (from [2])). *Let S be an $m \times n$ matrix with i.i.d. standard p -stable entries and let M be an $n \times d$ matrix. For $\varepsilon > 0$, with probability $1 - \frac{1}{\Omega(1)}$,*

$$(1 - \varepsilon)\|M\|_p \leq \left(\sum_i \text{med}(SM_i)^p \right)^{\frac{1}{p}} / \text{med}_p \leq (1 + \varepsilon)\|M\|_p$$

as long as $m = \text{poly}(1/\varepsilon)$.

Proof. This is Lemma 6 from [2] in the case $p = 1$, and Lemma 12 of [2] in the case $p \in (1, 2)$. □

Remark 3.4. *From inspecting the proof of the above lemma in the $p = 1$ case, it can be applied as long as $m \geq \frac{1}{\varepsilon^3}$. This is because $\Pr[\text{med}(SM_i) = (1 \pm \varepsilon)\|M_i\|_1] \leq e^{-\Theta(\varepsilon^2)m}$ (by Lemma 4 of [2]) and as long as $m \geq \frac{1}{\varepsilon^3}$, this probability is at most $e^{-\frac{1}{\varepsilon}} \leq \varepsilon$, which is sufficient for this lemma.*

Lemma 3.5 (p -stable Matrix + Top Quantile Does not Cause Dilation (from [2])). *When S is an $m \times n$ matrix with i.i.d. standard p -stable entries, $m = \text{poly}(1/\varepsilon)$, and M is an $n \times d$ matrix, then with probability $1 - \frac{1}{\Omega(1)}$,*

$$\left(\sum_i q_{1-\frac{\varepsilon}{2}}(SM_i)^p \right)^{\frac{1}{p}} / \text{med}_p \leq O\left(\frac{1}{\varepsilon}\right)\|M\|_p$$

Proof. This is Lemma 7 from [2] in the $p = 1$ case, and Lemma 13 of [2] in the $p \in (1, 2)$ case. □

Remark 3.6. From inspecting the proof of the above lemma in the $p = 1$ case, we see that it can be applied as long as $m \geq \frac{1}{\varepsilon^2}$. This is because $\Pr[q_{1-\varepsilon/2}(SM_i) \geq \frac{\varepsilon}{2}\|M_i\|_1] \leq e^{-\Theta(\varepsilon)c \cdot m} \leq e^{-O(1/\varepsilon)} \leq \varepsilon$ (by Lemma 4 of [2]) as long as $m \geq \frac{1}{\varepsilon^2}$.

Remark 3.7. In addition, the failure probabilities in the above two lemmas can be as small as desired, since they are obtained from Markov bounds.

Lemma 3.8 (Quasi-Subspace Embedding with Median Estimator (from [2])). *Let X be a k -dimensional subspace of \mathbb{R}^n and $\varepsilon, \delta > 0$. Let S be an $m \times n$ matrix whose entries are i.i.d. standard p -stable random variables, where $m = O(1/\varepsilon^2 \cdot k \log(k/\varepsilon\delta))$. Then, with probability at least $1 - \Theta(\delta)$, for all $x \in X$,*

$$(1 - \Theta(\varepsilon))\|x\|_p \leq q_{\frac{1}{2}-\varepsilon}(Sx/\text{med}_p) \leq q_{\frac{1}{2}+\varepsilon}(Sx/\text{med}_p) \leq (1 + O(\varepsilon))\|x\|_p$$

Proof. This is Lemma 5 of [2] in the $p = 1$ case, and Lemma 11 of [2] in the $p \in (1, 2)$ case. \square

Remark 3.9. Note that we can select any number of rows that is larger than some $O(\frac{1}{\varepsilon^2}k \log(\frac{k}{\varepsilon\delta}))$, by inspecting the proof of Lemma 5 of [2] — any number of rows larger than the specified $O(\frac{1}{\varepsilon^2}k \log(\frac{k}{\varepsilon\delta}))$ allows the net argument in that proof to work.

The following lemma shown in [2] allows this sketch to serve as a “quasi-affine embedding”:

Lemma 3.10. *Let $U \in \mathbb{R}^{n \times k}$ and $A \in \mathbb{R}^{n \times d}$. Let V^* be chosen to minimize $\|UV^* - A\|_p$. Suppose S is an $m \times n$ random matrix such that:*

1. $q_{\frac{1}{2}-\varepsilon}(SUX/\text{med}_p) \geq (1 - \Theta(\varepsilon))\|Ux\|_p$ for all $x \in \mathbb{R}^k$
2. For each $i \in [d]$, with probability at least $1 - \varepsilon^3$, $\text{med}(S[U, A_i]x/\text{med}_p) \geq (1 - \varepsilon^3)\|[U, A_i]x\|_p$ for all $x \in \mathbb{R}^{k+1}$
3. $(\sum_i \text{med}(SUV_i^* - SA_i)^p)^{\frac{1}{p}}/\text{med}_p \leq (1 + \varepsilon^3)\|UV^* - A\|_p$
4. $(\sum_i q_{1-\varepsilon/2}(SUV_i^* - SA_i)^p)^{\frac{1}{p}}/\text{med}_p \leq O(\frac{1}{\varepsilon})\|UV^* - A\|_p$

If statements 1, 3 and 4 each hold with probability $1 - \frac{1}{\Omega(1)}$, then with probability $1 - \frac{1}{\Omega(1)}$, for all $V \in \mathbb{R}^{k \times d}$,

$$\left(\sum_i \text{med}(SUV_i - SA_i)^p\right)^{\frac{1}{p}}/\text{med}_p \geq (1 - O(\varepsilon))\|UV - A\|_p$$

Proof. This is Theorem 11 of [2] in the $p = 1$ case and Theorem 12 of [2] in the $p \in (1, 2)$ case. Note that in the statements of these Theorems in [2], it is not explicitly stated that statements 1, 3 and 4 only need to hold with constant probability. However, this is true because, by inspecting the proof of Theorem 11 of [2], we see that statement 2 is only used to perform a Markov bound, which needs to hold with constant probability — once that Markov bound is obtained, a union bound can be performed over statements 1, 2, 3 and 4. \square

Specializing this lemma to p -stable matrices gives:

Lemma 3.11 (Lower Bound for One-Sided Embedding with Median Estimator (from [2])). *Let $U \in \mathbb{R}^{n \times k}$ and $A \in \mathbb{R}^{n \times d}$. If S is an $m \times n$ random matrix with i.i.d. standard p -stable entries, where $m = O(\max(k/\varepsilon^6 \log(k/\varepsilon), 1/\varepsilon^9))$, then with probability $1 - \frac{1}{\Omega(1)}$,*

$$\left(\sum_i \text{med}(SUV_i - SA_i)^p\right)^{\frac{1}{p}}/\text{med}_p \geq (1 - O(\varepsilon))\|UV - A\|_p$$

for all $V \in \mathbb{R}^{k \times d}$.

Proof. This is a corollary of all the above lemmas. For statement 1 of the previous lemma to hold, it is sufficient for S to have at least $O(k/\varepsilon^2 \log(k/\varepsilon))$ rows by Lemma 3.8. For statement 2 of the previous lemma to hold, it is sufficient for S to have at least $O(k/\varepsilon^6 \log(k/\varepsilon))$ rows, again by Lemma 3.8. For statement 3 of the previous lemma to hold, it is sufficient for S to have $O(1/\varepsilon^9)$ rows, by Lemma 3.3. Finally, for statement 4 of the previous lemma to hold, it is sufficient for S to have $O(1/\varepsilon^2)$ rows, by Lemma 3.5. \square

We recall another useful lemma on p -stable matrices, which bounds the ℓ_p -norm of SM if S is a p -stable matrix and M is a fixed matrix:

Lemma 3.12 (Distortion in ℓ_p -norm with p -stable Matrices - Lemma E.11 of [48]). *Let $M \in \mathbb{R}^{n \times d}$, and let $S \in \mathbb{R}^{r \times n}$ be a random matrix whose entries are i.i.d. standard p -stable random variables. Then, with probability $1 - \frac{1}{\Omega(1)}$,*

$$\|SM\|_p^p \leq O(r \log d) \|M\|_p^p$$

Note that in the original statement of this lemma in [48], the p -stable matrices are rescaled by $\Theta(1/r^{1/p})$ — since our sketch uses p -stable matrices that are not rescaled, we include this factor in the distortion.

In the course of our analysis, it will also be useful to note that med_p is bounded away from 0 — that is, there exists a constant $K > 0$ such that $\text{med}_p \geq K$. To our knowledge, this fact was not explicitly shown elsewhere, and we prove it below.

Lemma 3.13 (med_p is $\Omega(1)$). *There exists an absolute constant $K > 0$ such that $\text{med}_p \geq K$ for all $p \in [1, 2]$.*

Proof. First, we recall the following formula from [41] for the c.d.f. of a standard p -stable random variable for $p \in (1, 2]$. For $x > 0$, if X is a standard p -stable random variable, then

$$\Pr[X > x] = 1 - \frac{1}{\pi} \int_0^{\frac{\pi}{2}} e^{-x \frac{p}{p-1} \cdot V(\theta;p)} d\theta$$

where

$$V(\theta;p) = \left(\frac{\cos \theta}{\sin p\theta} \right)^{\frac{p}{p-1}} \cdot \frac{\cos(p-1)\theta}{\cos \theta}$$

(This is a corollary of Theorem 1 of [41].) Hence, because X is symmetric, $x > 0$ is less than med_p if and only if

$$\frac{3}{4} > \Pr[X > x] = 1 - \frac{1}{\pi} \int_0^{\frac{\pi}{2}} e^{-x \frac{p}{p-1} V(\theta;p)} d\theta$$

or equivalently

$$I := \int_0^{\frac{\pi}{2}} e^{-x \frac{p}{p-1} V(\theta;p)} d\theta > \frac{\pi}{4}$$

(recall that med_p is the median of $|X|$ rather than X). Now, our goal is to bound the integral I on the left-hand side from below. As a first step, we show the following claim.

Claim 3.14. *Let $c > 0$ be a sufficiently small absolute constant (to be chosen outside of this claim). There exists a constant $D > 0$ (which may depend on c) such that for all $p \in [1, 2]$ and $\theta \in [\frac{\pi}{8}, \frac{\pi}{2} - c]$,*

$$\frac{\cos(\theta)}{\sin(p\theta)} \cdot \left(\frac{\cos((p-1)\theta)}{\cos(\theta)} \right)^{\frac{p-1}{p}} \leq D$$

Proof. Observe that for $p \in [1, 2]$ and $\theta \in [\frac{\pi}{8}, \frac{\pi}{2} - c]$, $\frac{\pi}{8} \leq p\theta \leq \pi - 2c$, meaning $\sin(p\theta) > 0$ and is in fact bounded away from 0 because $[1, 2] \times [\frac{\pi}{8}, \frac{\pi}{2} - c]$ is compact. Let $C_0 > 0$ be such that $\sin(p\theta) \geq C_0$ for all $(p, \theta) \in [1, 2] \times [\frac{\pi}{8}, \frac{\pi}{2} - c]$. Then,

$$\frac{\cos(\theta)}{\sin(p\theta)} \cdot \left(\frac{\cos((p-1)\theta)}{\cos(\theta)} \right)^{\frac{p-1}{p}} \leq \cos(\theta)^{\frac{1}{p}} \cdot (\cos((p-1)\theta))^{\frac{p-1}{p}} \cdot \frac{1}{C_0} \leq \frac{1}{C_0} \quad (5)$$

where the second inequality holds because $\cos(\theta)$ and $\cos((p-1)\theta)$ are at most 1. (Note that $\cos(\theta)^{\frac{1}{p}}$ and $\cos((p-1)\theta)^{\frac{p-1}{p}}$ are well-defined because $\cos(\theta)$ and $\cos((p-1)\theta)$ are nonnegative for this choice of θ .) \square

Hence, for $p \in [1, 2]$ and $\theta \in [\frac{\pi}{8}, \frac{\pi}{2} - c]$,

$$V(\theta;p) \leq \left(\frac{\cos \theta}{\sin p\theta} \right)^{\frac{p}{p-1}} \cdot \frac{\cos(p-1)\theta}{\cos \theta} = \left(\frac{\cos(\theta)}{\sin(p\theta)} \cdot \left(\frac{\cos((p-1)\theta)}{\cos(\theta)} \right)^{\frac{p-1}{p}} \right)^{\frac{p}{p-1}} \leq D^{\frac{p}{p-1}}$$

Hence, we can bound I from below:

$$\begin{aligned} \int_0^{\frac{\pi}{2}} e^{-x \frac{p}{p-1} V(\theta;p)} d\theta &\geq \int_{\frac{\pi}{8}}^{\frac{\pi}{2}-c} e^{-x \frac{p}{p-1} D \frac{p}{p-1}} d\theta \\ &= \left(\frac{3\pi}{4} - c \right) e^{-(xD) \frac{p}{p-1}} \end{aligned} \quad (6)$$

We can choose $c = \frac{\pi}{8}$, meaning $I > \frac{\pi}{4}$ as long as $\frac{5\pi}{8} \cdot e^{-(xD) \frac{p}{p-1}} > \frac{\pi}{4}$ or $e^{-(xD) \frac{p}{p-1}} > \frac{2}{5}$. This holds as long as $x < \frac{1}{D} (-\log(2/5)) \frac{p-1}{p}$. Letting $C = -\ln(2/5)$, we observe that $0 < C < 1$, meaning $\sqrt{C} \leq C^{\frac{p-1}{p}} \leq 1$.

In summary, $x \leq \text{med}_p$ as long as $x < \frac{\sqrt{C}}{D}$, and hence we can take $\frac{\sqrt{C}}{2D}$ to be the desired K . \square

3.2 $(1 + \varepsilon)$ -Approximation Algorithm and Analysis

We now present and analyze our $(1 + \varepsilon)$ -approximation algorithm with bicriteria rank $3k$ for ℓ_p -low rank approximation, where $p \in [1, 2)$.

Theorem 3.15 (Correctness and Running Time of Algorithm 5). *Let $A \in \mathbb{R}^{n \times d}$, $k \in \mathbb{N}$, $\varepsilon \in (0, C)$ (where C is a sufficiently small absolute constant), $f > 1$ and $p \in [1, 2)$. Furthermore, if $\text{OPT} = \min_{A_k} \text{rank}_k \|A_k - A\|_p$, then suppose $(1 - O(\varepsilon))\text{OPT} \leq \widehat{\text{OPT}} \leq (1 + O(\varepsilon))\text{OPT}$. Finally, let $U \in \mathbb{R}^{n \times k}$, $V \in \mathbb{R}^{k \times d}$ be the output of GUESSINGEPSAPPROXIMATION($A, k, \varepsilon, f, \widehat{\text{OPT}}, p$) (shown in Algorithm 5). Then,*

$$\|UV - A\|_p \leq (1 + O(\varepsilon))\text{OPT} + O\left(\frac{\varepsilon}{f}\right) \|A\|_p$$

The running time of Algorithm 5 is at most $f^{O(rk)} + 2^{O(rk \log(k/\varepsilon))} + \text{poly}(fnd/\varepsilon)$, where r , the number of rows in the p -stable sketching matrix, is $O(\max(k^2/\varepsilon^6 \log(k/\varepsilon), 1/\varepsilon^9))$.

Remark 3.16. *We can obtain $\widehat{\text{OPT}}$ efficiently as follows. If $A_{\text{SVD},k}$ is the rank- k SVD of A (which can be computed in polynomial time) then $\text{OPT} \leq \|A_{\text{SVD},k} - A\|_p \leq (nd)^{1/p-1/2} \text{OPT}$, meaning we can guess all integer powers of $1 + \varepsilon$ between $\|A_{\text{SVD},k} - A\|_p$ and $\frac{1}{(nd)^{1/p-1/2}} \|A_{\text{SVD},k} - A\|_p$ — the number of guesses is $O(\frac{\log(nd)}{\varepsilon})$. We can input all of those guesses to Algorithm 5, and one of them will produce the right answer. This is done when applying Algorithm 5 within Algorithm 6.*

Remark 3.17. *Note that f represents the approximation factor of the initialization algorithm used to obtain A . In our case f will equal $\text{poly}(k)$, but we will analyze this algorithm for a general f .*

Proof. Let $U^* \in \mathbb{R}^{n \times k}$, $V^* \in \mathbb{R}^{k \times d}$ such that $\|U^*V^* - A\|_p = \text{OPT}$. Without loss of generality, assume V^* is a q -well-conditioned basis (where $q = \text{poly}(k)$), meaning that for all $x \in \mathbb{R}^k$,

$$\frac{\|x\|_p}{q} \leq \|x^T V^*\|_p \leq q \|x\|_p$$

(Note that well-conditioned bases exist for all p , for instance see Lemma 10 of [11].) In particular, this implies that $\|V^*\|_p^p \leq kq^p$ by letting x be each of the standard basis vectors. Now, let $S \in \mathbb{R}^{r \times n}$ be a random matrix where each entry is an i.i.d. p -stable random variable, where $r = O(\max(k/\varepsilon^6 \log(k/\varepsilon), 1/\varepsilon^9))$ as in Algorithm 5.

We first analyze the effect of rounding SU^* , multiplicatively, so that the absolute value of each entry is rounded to the nearest power of $1 + \frac{1}{f \text{poly}(k/\varepsilon)}$, or set to 0 if it is too small (here, $\text{poly}(k/\varepsilon)$ in the denominator is $(k/\varepsilon)^c$ for a sufficiently large constant c). Note that $\|U^*V^*\|_p \leq O(1)\|A\|_p$ by the triangle inequality. Hence, because V^* is a well-conditioned basis, $O(1)\|A\|_p \geq \frac{1}{q}\|U^*\|_p$, and $\|U^*\|_p \leq O(q)\|A\|_p$. Therefore, $\|SU^*\|_\infty \leq \|SU^*\|_p \leq \text{poly}(k/\varepsilon)\|U^*\|_p \leq \text{poly}(k/\varepsilon)\|A\|_p$, where the second inequality is due to Lemma 3.12, since U^* has k columns and S has $\text{poly}(k/\varepsilon)$ rows.

Now, let M_1 be SU^* , but with the absolute value of each entry rounded to the nearest power of $1 + \frac{1}{f \text{poly}(k/\varepsilon)}$. In addition, let M_2 be M_1 , but with each entry having absolute value less than $\frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p$ being replaced by 0. Then,

$$\|M_2 - M_1\|_p^p \leq \left(\frac{1}{f \text{poly}(k/\varepsilon)} \right)^p \|A\|_p^p \cdot O\left(\frac{k^2}{\varepsilon^9} \log(k/\varepsilon) \right) \leq \left(\frac{1}{f \text{poly}(k/\varepsilon)} \right)^p \|A\|_p^p$$

Algorithm 5 Guessing a sketched left factor SU , and finding an appropriate right factor V with norm at most $\text{poly}(k)$, to obtain a $(1 + \varepsilon)$ -approximation with additive $\varepsilon/(f\text{poly}(k))\|A\|_p^p$ error.

Require: $A \in \mathbb{R}^{n \times d}$, $k \in \mathbb{N}$, $\varepsilon \in (0, c)$, $f > 1$, $\widehat{OPT} \geq 0$, $p \in [1, 2)$

Ensure: $U \in \mathbb{R}^{n \times k}$, $V \in \mathbb{R}^{k \times d}$

procedure GUESSINGADDITIVEEPSAPPROXIMATION($A, k, \varepsilon, f, \widehat{OPT}, p$)

If A has rank k , return A .

$r \leftarrow O(\max(k/\varepsilon^6 \log(k/\varepsilon), 1/\varepsilon^9))$

$q \leftarrow \text{poly}(k)$

$S \leftarrow$ An $r \times n$ matrix of i.i.d. standard p -stable random variables

$\mathcal{I} \leftarrow \{0\} \cup \left\{ \sigma \cdot \left(1 + \frac{1}{f\text{poly}(k/\varepsilon)}\right)^t \mid t \in \mathbb{Z}, \frac{1}{f\text{poly}(k/\varepsilon)}\|A\|_p \leq \left(1 + \frac{1}{f\text{poly}(k/\varepsilon)}\right)^t \leq \text{poly}(k/\varepsilon)\|A\|_p, \sigma = \pm 1 \right\}$

$\mathcal{C} \leftarrow \left\{ M \in \mathbb{R}^{r \times k} \mid M_{i,j} \in \mathcal{I} \ \forall i \in [r], j \in [k] \right\}$ — This is the set of (sketched) left factors we will guess.

// Guess possible rounded (sketched) left factors and find a good right factor V , with $\|V\|_p \leq \text{poly}(k)$.

$U_{best} \leftarrow 0, V_{best} \leftarrow 0$

for $M \in \mathcal{C}$ **do**

$\text{COSTBOUNDS} \leftarrow \left\{ \frac{\varepsilon^2}{f}\|A\|_p/d^{\frac{1}{p}} \leq c \leq O(\|A\|_p) \text{ and } c \text{ is an integer power of } (1 + \varepsilon) \right\}$

for $i \in [d]$, $c \in \text{COSTBOUNDS}$ **do**

$V_{i,c} \leftarrow \text{argmin}_{V_i} \|V_i\|_p$ subject to the constraint that $\text{med}(MV_i - SA_i)/\text{med}_p \leq c$

$C_{i,c} \leftarrow \text{med}(MV_{i,c} - SA_i)/\text{med}_p$

end for

// Create LP to find a good distribution over $c \in \text{COSTBOUNDS}$ for each $i \in [d]$.

$\text{VARIABLES} \leftarrow \{x_{i,c} \mid i \in [d], c \in \text{COSTBOUNDS}\}$

$\text{CONSTRAINTS} \leftarrow \left\{ 0 \leq x_{i,c} \ \forall i \in [d], c \in \text{COSTBOUNDS} \text{ and } \sum_{c \in \text{COSTBOUNDS}} x_{i,c} = 1 \ \forall i \in [d] \right\}$

$\text{CONSTRAINTS} \leftarrow \text{CONSTRAINTS} \cup \left\{ \sum_{i \in [d], c \in \text{COSTBOUNDS}} x_{i,c} \|V_{i,c}\|_p^p \leq kq^p \right\}$

$\Delta \leftarrow (1 + O(\varepsilon))^p \left(\widehat{OPT} + \frac{1}{f\text{poly}(k/\varepsilon)}\|A\|_p \right)^p + \left(\frac{\varepsilon^2}{f} \right)^p \|A\|_p^p$

$\text{CONSTRAINTS} \leftarrow \text{CONSTRAINTS} \cup \left\{ \sum_{i \in [d], c \in \text{COSTBOUNDS}} x_{i,c} C_{i,c}^p \leq \Delta \right\}$

$x_{i,c} \leftarrow$ Solution to the LP given by VARIABLES and CONSTRAINTS , for all $i \in [d]$, $c \in \text{COSTBOUNDS}$

If LP is infeasible, then **continue** to next $M \in \mathcal{C}$.

// For each column, sample an appropriate cost bound. Do this $O(1/\varepsilon)$ times, then V' meets both // the cost and norm constraints with constant probability.

for $t = 1 \rightarrow 10/\varepsilon$ **do**

$c_i \leftarrow$ An element $c \in \text{COSTBOUNDS}$ selected according to the distribution on COSTBOUNDS given by $\{x_{i,c} \mid c \in \text{COSTBOUNDS}\}$ (note that for $i \in [d]$, the $x_{i,c}$ are nonnegative and sum to 1)

$V'_i \leftarrow V_{i,c_i}$ for all $i \in [d]$

Break if $\|V'\|_p^p \leq \frac{2kq^p}{\varepsilon}$ and $\sum_{i=1}^d \text{med}(MV'_i - SA_i)^p / \text{med}_p^p \leq (1 + 2\varepsilon)\Delta$

end for

$U' \leftarrow \text{argmin}_U \|UV' - A\|_p$

If $\|U'V' - A\|_p \leq \|U_{best}V_{best} - A\|_p$ then $U_{best} \leftarrow U'$ and $V_{best} \leftarrow V'$.

end for

return U', V'

end procedure

Algorithm 6 First apply POLYKERRORNOTBI CRITERIA APPROXIMATION from Algorithm 8 to A to obtain a $\text{poly}(k)$ -approximation B — then, apply Algorithm 5 to the residual to obtain an approximation UV with additive error $1/\text{poly}(k/\varepsilon)\|A - B\|_p \leq \varepsilon \cdot \text{OPT}$. Finally, $B + UV$ gives a $(1 + \varepsilon)$ -approximation with rank $3k$.

Require: $A \in \mathbb{R}^{n \times d}$, $k \in \mathbb{N}$, $\varepsilon \in (0, c)$, $p \in [1, 2]$

Ensure: $\hat{A} \in \mathbb{R}^{n \times d}$ having rank $3k$

procedure ROUNDINGGUESSINGEPSAPPROXIMATION(A, k, ε, p)
 $W, Z \leftarrow \text{POLYKERRORNOTBI CRITERIA APPROXIMATION}(A, k, p)$
 $B \leftarrow WZ$
 $C \leftarrow A - B$
 $f \leftarrow \text{poly}(k)$, the approximation factor of Algorithm 8

// Guess all $O((\log nd)/\varepsilon)$ possible values for $\widehat{\text{OPT}}$ and try them for Algorithm 5
// as described in Remark 3.16.
 $C_{\text{SVD}, 2k} \leftarrow$ The optimal rank- $2k$ approximation for C under the ℓ_2 norm.
 $\text{SVDERROR} \leftarrow \|C - C_{\text{SVD}, 2k}\|_p$
 $\hat{A} \leftarrow 0 \in \mathbb{R}^{n \times d}$
for $t = 0 \rightarrow O(\frac{\log nd}{\varepsilon})$ **do**
 $\widehat{\text{OPT}} \leftarrow \text{SVDERROR}/(1 + \varepsilon)^t$
 $U, V \leftarrow \text{GUESSINGADDITIVEEPSAPPROXIMATION}(C, 2k, \varepsilon, f, \widehat{\text{OPT}}, p)$
If $\|(B + UV) - A\|_p \leq \|\hat{A} - A\|_p$ then $\hat{A} \leftarrow B + UV$
end for
return \hat{A}
end procedure

where the first inequality holds because M_1 and M_2 have $rk = O(\max(k^2/\varepsilon^6 \log(k/\varepsilon), 1/\varepsilon^9))$ entries. Moreover,

$$\|SU^* - M_1\|_p^p \leq \left(\frac{1}{f \text{poly}(k/\varepsilon)}\right)^p \|SU^*\|_p^p \leq \left(\frac{1}{f \text{poly}(k/\varepsilon)}\right)^p \cdot \text{poly}(k/\varepsilon) \|A\|_p^p \leq \left(\frac{1}{f \text{poly}(k/\varepsilon)}\right)^p \|A\|_p^p$$

where the first inequality is because for any $a \in \mathbb{R}$, and $t \in (0, 1)$, if \hat{a} is a with its absolute value rounded to the nearest power of $(1 + t)$, then $|a - \hat{a}|^p \leq t^p |a|^p$. The second inequality is simply because $\|SU^*\|_p^p \leq \text{poly}(k/\varepsilon) \|A\|_p^p$ as mentioned above. Hence,

$$\begin{aligned} \|SU^* - M_2\|_p &\leq \|SU^* - M_1\|_p + \|M_1 - M_2\|_p \\ &\leq \left(\left(\frac{1}{f \text{poly}(k/\varepsilon)}\right)^p \|A\|_p^p\right)^{\frac{1}{p}} + \left(\left(\frac{1}{f \text{poly}(k/\varepsilon)}\right)^p \|A\|_p^p\right)^{\frac{1}{p}} \\ &\leq \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p + \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p \\ &\leq \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p \end{aligned} \tag{7}$$

where the first inequality is due to the triangle inequality.

Now, observe that Algorithm 5 will guess M_2 at some point — that is, since $M_2 \in \mathcal{C}$, M will be equal to M_2 at some point. Suppose $M = M_2$. Let us condition on the following events involving S . Let \mathcal{E}_1 denote the event that for all $V \in \mathbb{R}^{k \times d}$,

$$\left(\sum_i \text{med}(SU^*V_i - SA_i)^p\right)^{1/p} / \text{med}_p \geq (1 - O(\varepsilon)) \|U^*V - A\|_p$$

\mathcal{E}_2 the event that

$$\left(\sum_i \text{med}(SU^*V_i^* - SA_i)^p\right)^{1/p} / \text{med}_p \leq (1 + \varepsilon) \|U^*V^* - A\|_p = (1 + \varepsilon) \text{OPT}$$

and \mathcal{E}_3 the event that

$$\|SU^*\|_p \leq \text{poly}(k/\varepsilon)\|U^*\|_p$$

By Lemma 3.11, 3.3 and 3.12 respectively, these each occur with probability $1 - O(1)$ (where the constant probability can be made as small as desired by increasing r by a constant factor), and by a union bound they occur simultaneously with probability $1 - O(1)$. First, we use these events to examine the effect of rounding SU^* to M , when the right factor is V^* , and more generally, when the ℓ_p norm of the right factor V is at most $\text{poly}(k/\varepsilon)$.

Claim 3.18. *Suppose \mathcal{E}_1 , \mathcal{E}_2 and \mathcal{E}_3 hold, and suppose $V \in \mathbb{R}^{k \times d}$ such that $\|V\|_p^p \leq \text{poly}(k/\varepsilon)$. Then,*

$$\left| \left(\sum_{i=1}^d \text{med}(MV_i - SA_i)^p \right)^{\frac{1}{p}} / \text{med}_p - \left(\sum_{i=1}^d \text{med}(SU^*V_i - SA_i)^p \right)^{\frac{1}{p}} / \text{med}_p \right| \leq \frac{1}{f \text{poly}(k/\varepsilon) \text{med}_p} \|A\|_p$$

Proof. Let $x \in \mathbb{R}^d$ be the vector whose i^{th} coordinate is $\text{med}(MV_i - SA_i)$, and let $y \in \mathbb{R}^d$ be the vector whose i^{th} coordinate is $\text{med}(SU^*V_i - SA_i)$. Then, observe that

$$\left| \left(\sum_{i=1}^d \text{med}(MV_i - SA_i)^p \right)^{\frac{1}{p}} - \left(\sum_{i=1}^d \text{med}(SU^*V_i - SA_i)^p \right)^{\frac{1}{p}} \right| = \|x\|_p - \|y\|_p$$

Now, by the triangle inequality, this is at most $\|x - y\|_p$, which we can bound from above:

$$\begin{aligned} \|x - y\|_p^p &= \sum_{i=1}^d |\text{med}(MV_i - SA_i) - \text{med}(SU^*V_i - SA_i)|^p \\ &\leq \sum_{i=1}^d \|(M - SU^*)V_i\|_\infty^p \\ &\leq \sum_{i=1}^d \|(M - SU^*)V_i\|_p^p \\ &= \|(M - SU^*)V\|_p^p \end{aligned} \tag{8}$$

Here, the first inequality is because $|\text{med}(v_1 + v_2) - \text{med}(v_1)| \leq \|v_2\|_\infty$ for any two vectors $v_1, v_2 \in \mathbb{R}^n$, and the second is because $\|v\|_\infty \leq \|v\|_p$ for any vector v .

Hence,

$$\|x - y\|_p \leq \|(M - SU^*)V\|_p \leq \text{poly}(k/\varepsilon)\|M - SU^*\|_p \leq \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p$$

Here, the second inequality holds because, even though V is not necessarily a well-conditioned basis, for any vector $x \in \mathbb{R}^k$,

$$\|x^T V\|_p = \left\| \sum_{i=1}^k x_i V^i \right\|_p \leq \sum_{i=1}^k |x_i| \|V^i\|_p \leq k \|x\|_p \|V\|_p = \text{poly}(k/\varepsilon) \|x\|_p$$

where the first inequality is due to the triangle inequality, and the last equality is because $\|V\|_p \leq \text{poly}(k/\varepsilon)$. Hence, for any matrix $D \in \mathbb{R}^{n \times k}$,

$$\|DV\|_p^p = \sum_{i=1}^n \|D^i V\|_p^p \leq \sum_{i=1}^n \text{poly}(k/\varepsilon) \|D^i\|_p^p = \text{poly}(k/\varepsilon) \|D\|_p^p$$

and taking p^{th} roots gives $\|DV\|_p \leq \text{poly}(k/\varepsilon) \|D\|_p$. \square

Claim 3.19. *Suppose \mathcal{E}_1 , \mathcal{E}_2 and \mathcal{E}_3 occur. Then, the linear program constructed in Algorithm 5 with variables in the set VARIABLES, and constraints in the set CONSTRAINTS, is feasible when $M = M_2$ is guessed.*

Proof. Consider a particular $i \in [d]$. Note that

$$\begin{aligned}
\frac{\text{med}(MV_i^* - SA_i)}{\text{med}_p} &\leq \frac{\text{med}(SU^*V_i^* - SA_i)}{\text{med}_p} + \frac{\|(M - SU^*)V_i^*\|_\infty}{\text{med}_p} \\
&\leq \frac{\text{med}(SU^*V_i^* - SA_i)}{\text{med}_p} + \frac{\|(M - SU^*)V_i^*\|_p}{\text{med}_p} \\
&\leq \frac{\text{med}(SU^*V_i^* - SA_i)}{\text{med}_p} + \text{poly}(k) \frac{\|(M - SU^*)\|_p}{\text{med}_p} \\
&\leq \frac{\text{med}(SU^*V_i^* - SA_i)}{\text{med}_p} + \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p \\
&\leq \frac{\left(\sum_{j \in [d]} \text{med}(SU^*V_j^* - SA_j)^p\right)^{\frac{1}{p}}}{\text{med}_p} + \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p \\
&\leq (1 + \varepsilon)OPT + \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p \\
&\leq O(1) \|A\|_p
\end{aligned} \tag{9}$$

The first inequality is true because for $v_1, v_2 \in \mathbb{R}^n$, $|\text{med}(v_1 + v_2) - \text{med}(v_1)| \leq \|v_2\|_\infty$, meaning $|\text{med}(MV_i^* - SA_i) - \text{med}(SU^*V_i^* - SA_i)| \leq \|(M - SU^*)V_i^*\|_\infty$. Here, the second inequality is by $\|x\|_p \geq \|x\|_\infty$, and the third is because V^* is a $\text{poly}(k)$ well-conditioned basis. The fourth inequality is because med_p is nonnegative and bounded away from 0 by Lemma 3.13, and because $\|M - SU^*\|_p \leq \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p$ by Equation 7. Finally, the sixth inequality is because \mathcal{E}_2 holds and the seventh is because $OPT \leq \|A\|_p$.

As a summary, we have shown that for each $i \in [d]$, $\text{med}(MV_i^* - SA_i)/\text{med}_p \leq O(1) \|A\|_p$. Hence, there are two cases — either there exists $c_0 \in \text{COSTBOUNDS}$ such that $c_0 \geq \text{med}(MV_i^* - SA_i)/\text{med}_p \geq \frac{c_0}{1+\varepsilon}$, or $\text{med}(MV_i^* - SA_i)/\text{med}_p \leq (\frac{\varepsilon^2}{f}) \|A\|_p/d^{\frac{1}{p}}$.

- In the first case, we can assign $x_{i,c_0} = 1$ and $x_{i,c'} = 0$ for $c' \neq c_0$. Define V_{i,c_0} as in Algorithm 5, meaning it is equal to $\text{argmin}_{V_i} \|V_i\|_p$ subject to the constraint that $\text{med}(MV_i - SA_i)/\text{med}_p \leq c_0$. Then, $\|V_{i,c_0}\|_p \leq \|V_i\|_p$ and moreover, since $\text{med}(MV_i^* - SA_i)/\text{med}_p \geq \frac{c_0}{1+\varepsilon}$, $\text{med}(MV_{i,c_0} - SA_i)/\text{med}_p \leq c_0 \leq (1 + O(\varepsilon)) \text{med}(MV_i^* - SA_i)/\text{med}_p$.
- In the second case, we can assign $x_{i,c_0} = 1$ and $x_{i,c'} = 0$ for $c' \neq c_0$, where c_0 is now $(\frac{\varepsilon^2}{f}) \|A\|_p/d^{\frac{1}{p}}$. Again, V_{i,c_0} is equal to $\text{argmin}_{V_i} \|V_i\|_p$ subject to the constraint that $\text{med}(MV_i - SA_i)/\text{med}_p \leq c_0$. Hence, since $\text{med}(MV_i^* - SA_i)/\text{med}_p \leq c_0$, $\|V_{i,c_0}\|_p \leq \|V_i^*\|_p$. Moreover, $\text{med}(MV_{i,c_0} - SA_i)/\text{med}_p \leq (\frac{\varepsilon^2}{f}) \|A\|_p/d^{\frac{1}{p}}$, which is sufficient for our purposes, as we see now.

We now conclude that this assignment to the $x_{i,c}$ satisfies all the constraints of the linear program. Clearly the constraints $x_{i,c} \geq 0$ (for all $i \in [d]$, $c \in \text{COSTBOUNDS}$) and $\sum_{c \in \text{COSTBOUNDS}} x_{i,c} = 1$ (for all $i \in [d]$) are satisfied (since for each $i \in [d]$, exactly one $x_{i,c}$ is set to 1 and the rest are 0). For each $i \in [d]$, let c_i be the unique element of COSTBOUNDS such that $x_{i,c_i} = 1$ according to our assignment. Then, for all $i \in [d]$, $\|V_{i,c_i}\|_p^p \leq \|V_i^*\|_p^p$ as mentioned above (where we wrote $\|V_{i,c_0}\|_p^p \leq \|V_i^*\|_p^p$), meaning

$$\sum_{i=1}^d \sum_{c \in \text{COSTBOUNDS}} x_{i,c} \|V_{i,c}\|_p^p = \sum_{i=1}^d \|V_{i,c_i}\|_p^p \leq \sum_{i=1}^d \|V_i^*\|_p^p = \|V^*\|_p^p \leq kq^p = \text{poly}(k)$$

where the last inequality is because each row of V^* has norm at most q . In addition,

$$\begin{aligned}
\sum_{i=1}^d \sum_{c \in \text{COSTBOUNDS}} x_{i,c} C_{i,c}^p &= \sum_{i=1}^d \text{med}(MV_{i,c_i} - SA_i)^p / \text{med}_p^p \\
&\leq \sum_{i=1}^d \left((1 + O(\varepsilon))^p \text{med}(MV_i^* - SA_i)^p / \text{med}_p^p + (\varepsilon^2/f)^p \|A\|_p^p / d \right) \\
&= (1 + O(\varepsilon))^p \frac{\sum_{i=1}^d \text{med}(MV_i^* - SA_i)^p}{\text{med}_p^p} + \left(\frac{\varepsilon^2}{f}\right)^p \|A\|_p^p
\end{aligned} \tag{10}$$

By Claim 3.18, this is at most

$$(1 + O(\varepsilon))^p \left(\left(\sum_{i=1}^d \text{med}(SU^*V_i^* - SA_i)^p \right)^{\frac{1}{p}} / \text{med}_p + \frac{1}{f \text{poly}(k/\varepsilon) \text{med}_p} \|A\|_p \right)^p + \left(\frac{\varepsilon^2}{f} \right)^p \|A\|_p^p$$

and this is in turn at most

$$(1 + O(\varepsilon))^p \left(\widehat{OPT} + \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p \right)^p + \left(\frac{\varepsilon^2}{f} \right)^p \|A\|_p^p$$

Here, we used the fact that \mathcal{E}_2 occurs, meaning $\left(\sum_{i=1}^d \text{med}(SU^*V_i^* - SA_i)^p \right)^{\frac{1}{p}} / \text{med}_p \leq (1 + \varepsilon)OPT$, $OPT \leq \frac{1}{1-O(\varepsilon)} \widehat{OPT} \leq (1 + O(\varepsilon)) \widehat{OPT}$, and moreover, $\text{med}_p = \Omega(1)$ (regardless of p), meaning $\frac{1}{\text{med}_p}$ is bounded above.

Hence, we can conclude that

$$\sum_{i=1}^d \sum_{c \in \text{CostBounds}} x_{i,c} C_{i,c}^p \leq (1 + O(\varepsilon))^p \left(\widehat{OPT} + \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p \right)^p + \left(\frac{\varepsilon^2}{f} \right)^p \|A\|_p^p$$

and we have shown that this assignment to the $x_{i,c}$ satisfies all the constraints of the LP. \square

Claim 3.20. *Suppose \mathcal{E}_1 , \mathcal{E}_2 and \mathcal{E}_3 occur. Then, when $M = M_2$ is guessed, with probability $1 - O(1)$, Algorithm 5 finds V' such that $\|V'\|_p^p \leq \frac{2kq^p}{\varepsilon}$ and $\sum_{i=1}^d \text{med}(MV'_i - SA_i)^p / \text{med}_p^p \leq (1 + 2\varepsilon)\Delta$, where Δ is defined in Algorithm 5.*

Proof. The proof is by Markov's inequality. First, by the previous claim, if \mathcal{E}_1 , \mathcal{E}_2 and \mathcal{E}_3 hold, then solving the LP within Algorithm 5 gives $x_{i,c}$ for $i \in [d]$ and $c \in \text{COSTBOUNDS}$ such that

$$\sum_{i \in [d], c \in \text{COSTBOUNDS}} x_{i,c} \|V_{i,c}\|_p^p \leq kq^p \quad \text{and} \quad \sum_{i \in [d], c \in \text{COSTBOUNDS}} x_{i,c} C_{i,c}^p \leq \Delta$$

Now, for each column $i \in [d]$, suppose we sample a single $c_i \in \text{COSTBOUNDS}$ according to the distribution on COSTBOUNDS given by the $x_{i,c}$, and we let $V'_i = V_{i,c_i}$. Denote this distribution by π_i . Then, conditioning on a *fixed* value of the p -stable matrix S such that \mathcal{E}_1 , \mathcal{E}_2 and \mathcal{E}_3 hold,

$$E_{c_i \sim \pi_i, \forall i \in [d]} \left[\|V'\|_p^p \mid S \right] = \sum_{i=1}^d \sum_{c \in \text{COSTBOUNDS}} x_{i,c} \|V_{i,c}\|_p^p \leq kq^p$$

The first equality is by the linearity of expectation. Note that prior to sampling c_i for each $i \in [d]$, the only source of randomness is S , and the appropriate rounded matrix M , as well as the minimizers $V_{i,c}$ and their costs $C_{i,c}$ are all determined by S . The second inequality holds if S is such that \mathcal{E}_1 , \mathcal{E}_2 and \mathcal{E}_3 are satisfied. By the same argument

$$E_{c_i \sim \pi_i, \forall i \in [d]} \left[\sum_{i=1}^d \text{med}(MV'_i - SA_i)^p / \text{med}_p^p \mid S \right] = \sum_{i \in [d], c \in \text{COSTBOUNDS}} x_{i,c} C_{i,c}^p \leq \Delta$$

Hence, letting $\mathcal{E} = \mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3$,

$$E_{c_i \sim \pi_i, \forall i \in [d]} \left[\|V'\|_p^p \mid \mathcal{E} \right] = \frac{\int_{\mathcal{E}} E_{c_i \sim \pi_i, \forall i \in [d]} \left[\|V'\|_p^p \mid S \right] p(S) dS}{P[\mathcal{E}]} \leq \frac{\int_{\mathcal{E}} kq^p \cdot p(S) dS}{P[\mathcal{E}]} = kq^p$$

Here, the first equality is by the identity $E[X] = E[E[X \mid Y]]$, where $X = \|V'\|_p^p$ and $Y = S$, the expectation is taken using the probability measure obtained by conditioning on \mathcal{E} , and $p(S)$ denotes the p.d.f. of the p -stable matrix S . The second is because the integral is over S for which \mathcal{E}_1 , \mathcal{E}_2 and \mathcal{E}_3 hold. Note that we

take the expectation conditioned on a fixed value of the continuous random variable S . To make this fully rigorous, we can assume the entries of S are rounded to the nearest integer multiples of an arbitrarily small $\delta > 0$. This would then allow us to treat the integral as a discrete sum, and it is easy to see that rounding S in this way does not affect the analysis — see the remark below on how we can justify the above steps, such as taking the expectation conditioned on a fixed value of S , without discretizing S . Similarly,

$$\begin{aligned}
& E_{c_i \sim \pi_i, \forall i \in [d]} \left[\sum_{i=1}^d \text{med}(MV'_i - SA_i)^p / \text{med}_p^p \mid \mathcal{E} \right] \\
&= \frac{\int_{\mathcal{E}} E_{c_i \sim \pi_i, \forall i \in [d]} \left[\sum_{i=1}^d \text{med}(MV'_i - SA_i)^p / \text{med}_p^p \mid S \right] p(S) dS}{P[\mathcal{E}]} \\
&\leq \frac{\int_{\mathcal{E}} \Delta \cdot p(S) dS}{P[\mathcal{E}]} \\
&= \Delta
\end{aligned} \tag{11}$$

Now, by Markov's inequality, if we sample c_i according to π_i for $i \in [d]$, then

$$P \left[\|V'\|_p^p \geq \frac{2kq^p}{\varepsilon} \mid \mathcal{E} \right] \leq \frac{\varepsilon}{2}$$

and

$$P \left[\sum_{i=1}^d \text{med}(MV'_i - SA_i)^p / \text{med}_p^p \geq (1 + 2\varepsilon)\Delta \mid \mathcal{E} \right] \leq \frac{1}{1 + 2\varepsilon} \leq 1 - \varepsilon$$

where the second inequality holds because for $\varepsilon \in (0, \frac{1}{2})$, $(1 + 2\varepsilon)(1 - \varepsilon) = 1 + \varepsilon - 2\varepsilon^2 = 1 + \varepsilon(1 - 2\varepsilon) \geq 1$, meaning $1 - \varepsilon \geq \frac{1}{1 + 2\varepsilon}$. Hence, by a union bound,

$$P \left[\sum_{i=1}^d \text{med}(MV'_i - SA_i)^p / \text{med}_p^p \geq (1 + 2\varepsilon)\Delta \text{ or } \|V'\|_p^p \geq \frac{2kq^p}{\varepsilon} \mid \mathcal{E} \right] \leq 1 - \frac{\varepsilon}{2}$$

Finally, if we repeatedly sample c_i for $i \in [d]$ over the course of $\frac{10}{\varepsilon}$ trials, then the probability that none of the V' satisfy the desired properties is at most

$$\left(1 - \frac{\varepsilon}{2}\right)^{10/\varepsilon} = \left(1 - \frac{\varepsilon}{2}\right)^{(2/\varepsilon) \cdot 5} \leq \frac{1}{e^5} \leq \frac{1}{100}$$

This completes the proof of the claim — if we condition on \mathcal{E}_1 , \mathcal{E}_2 and \mathcal{E}_3 , then with probability $\frac{99}{100}$, we obtain V' such that $\|V'\|_p^p \leq \frac{2kq^p}{\varepsilon}$ and

$$\sum_{i=1}^d \text{med}(MV'_i - SA_i)^p / \text{med}_p^p \leq (1 + 2\varepsilon)\Delta$$

□

Remark 3.21. In the proof of the above lemma, we take the expectation of $\|V'\|_p^p$ and $\sum_{i=1}^d \text{med}(MV'_i - SA_i)^p / \text{med}_p^p$ conditioned on a fixed value of the p -stable matrix S , and then use this to compute $E[\|V'\|_p^p \mid \mathcal{E}]$ and $E[\sum_{i=1}^d \text{med}(MV'_i - SA_i)^p / \text{med}_p^p \mid \mathcal{E}]$, using the identity $E[X] = E[E[X \mid Y]]$. Note that the inner expectation here is conditioned on a fixed value of the continuous random variable S . We can make this fully rigorous by discretizing S , i.e. rounding the entries of S to the nearest integer multiple of $\delta > 0$, where δ is arbitrarily small, as mentioned above (in fact, it suffices to round the entries of S to a multiple of $\frac{\varepsilon}{\text{poly}(nd)}$). Alternatively, we can formalize this using disintegrations, by applying Theorem 3 of [10], with the following parameters:

- The sample space Ω is $\mathbb{R}^{\text{poly}(k/\varepsilon) \times n} \times \mathcal{D}$, where $\mathbb{R}^{\text{poly}(k/\varepsilon) \times n}$ is the set of all possible S and \mathcal{D} is the collection of all possible choices of c_i , for $i \in [d]$.

- λ , in the statement of Theorem 3 of [10], would be the product measure of the Lebesgue measure on $\mathbb{R}^{\text{poly}(k/\varepsilon) \times n}$, with the counting measure on \mathcal{D} .
- ρ would be the probability distribution on Ω which is absolutely continuous with respect to λ , such that its density with respect to λ , at a point (S, c) in the sample space, is $p(S)q_{c|S}(c | S)$, where $p(S)$ is the p.d.f. of a p -stable matrix S , and $q_{c|S}(c | S)$ is the distribution on the choices c_i , for $i \in [d]$, that is given by Algorithm 5. (Note that by Fubini's Theorem, $\rho(\Omega) = \int_{\Omega} p(S)q_{c|S}(c | S)d\lambda = 1$, so ρ is a probability distribution.)
- μ is the probability distribution on $\mathbb{R}^{\text{poly}(k/\varepsilon) \times n}$, such that $S \sim \mu$ has i.i.d. standard p -stable entries.

Note that λ is σ -finite since the Lebesgue measure is σ -finite, and ρ and μ are finite measures since they are probability distributions. Thus, the relevant measures are all σ -finite, and ρ is absolutely continuous w.r.t. λ , meaning the hypotheses of Theorem 3 of [10] are satisfied. With these parameters, from part (v) of Theorem 3 in [10], we find that for $S \in \mathbb{R}^{\text{poly}(k/\varepsilon) \times n}$, the collection of probability measures on \mathcal{D} given by

$$\rho_S(A) = \sum_{c \in A} q_{c|S}(c | S)$$

for $A \subset \mathcal{D}$ is a valid disintegration of ρ with respect to μ . Thus, using this disintegration, we can rigorously define $E[\|V'\|_p^p | S]$ and $E[\sum_{i=1}^d \text{med}(MV'_i - SA_i) / \text{med}_p^p | S]$ for a fixed value of the p -stable matrix S , and use these to compute $E[\|V'\|_p^p | \mathcal{E}]$ and $E[\sum_{i=1}^d \text{med}(MV'_i - SA_i) / \text{med}_p^p | \mathcal{E}]$, using part (iii) of Definition 1 of [10] (which is equivalent to the identity $E[X] = E[E[X|Y]]$).

Now, suppose \mathcal{E}_1 , \mathcal{E}_2 and \mathcal{E}_3 hold, and that we have obtained V' with

$$\|V'\|_p^p \leq \frac{O(kq^p)}{\varepsilon} \text{ and } \sum_{i=1}^d \text{med}(MV'_i - SA_i)^p / \text{med}_p^p \leq (1 + 2\varepsilon)\Delta$$

By the previous claim, this occurs with constant probability. By Claim 3.18, since $\|V'\|_p \leq \text{poly}(k/\varepsilon)$, we know that

$$\begin{aligned} \left(\sum_{i=1}^d \text{med}(MV'_i - SA_i)^p \right)^{\frac{1}{p}} / \text{med}_p &\geq \left(\sum_{i=1}^d \text{med}(SU^*V'_i - SA_i)^p \right)^{\frac{1}{p}} / \text{med}_p - \frac{1}{f \text{poly}(k/\varepsilon) \text{med}_p} \|A\|_p \\ &\geq (1 - O(\varepsilon)) \|U^*V' - A\|_p - \frac{1}{f \text{poly}(k/\varepsilon) \text{med}_p} \|A\|_p \\ &\geq (1 - O(\varepsilon)) \|U'V' - A\|_p - \frac{1}{f \text{poly}(k/\varepsilon) \text{med}_p} \|A\|_p \end{aligned} \quad (12)$$

Here, the first inequality is by Claim 3.18, the second is because we are conditioning on \mathcal{E}_1 (meaning the median-based sketch and S provide a one-sided embedding for U^* and A), and the third inequality is because $U' = \text{argmin}_U \|UV' - A\|_p$. Therefore,

$$\begin{aligned} \|U'V' - A\|_p &\leq (1 + O(\varepsilon)) \left(\sum_{i=1}^d \text{med}(MV'_i - SA_i)^p \right)^{\frac{1}{p}} / \text{med}_p + \frac{1}{f \text{poly}(k/\varepsilon) \text{med}_p} \|A\|_p \\ &\leq (1 + O(\varepsilon))(1 + 2\varepsilon)^{1/p} \Delta^{1/p} + \frac{1}{f \text{poly}(k/\varepsilon) \text{med}_p} \|A\|_p \\ &\leq (1 + O(\varepsilon)) \Delta^{1/p} + \frac{1}{f \text{poly}(k/\varepsilon) \text{med}_p} \|A\|_p \\ &\leq (1 + O(\varepsilon)) \Delta^{1/p} + \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p \end{aligned} \quad (13)$$

where the last inequality is because $\text{med}_p = \Omega(1)$ as p ranges through $[1, 2]$. Finally, we can bound $\Delta^{1/p}$ from above by observing that the function $f : x \rightarrow |x|^{1/p}$ is subadditive:

$$\begin{aligned}
\Delta^{1/p} &= \left((1 + O(\varepsilon))^p \left(\widehat{OPT} + \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p \right)^p + \left(\frac{\varepsilon^2}{f} \right)^p \|A\|_p^p \right)^{1/p} \\
&\leq \left((1 + O(\varepsilon))^p \left(\widehat{OPT} + \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p \right)^p \right)^{1/p} + \left(\left(\frac{\varepsilon^2}{f} \right)^p \|A\|_p^p \right)^{1/p} \\
&= (1 + O(\varepsilon)) \left(\widehat{OPT} + \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p \right) + \frac{\varepsilon^2}{f} \|A\|_p \\
&\leq (1 + O(\varepsilon)) \left(OPT + \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p \right) + \frac{\varepsilon^2}{f} \|A\|_p
\end{aligned} \tag{14}$$

where the first inequality holds because $(|x| + |y|)^{1/p} \leq |x|^{1/p} + |y|^{1/p}$, and the last is because $\widehat{OPT} \leq (1 + O(\varepsilon))OPT$. In summary,

$$\begin{aligned}
\|U'V' - A\|_p &\leq (1 + O(\varepsilon))\Delta^{1/p} + \frac{1}{f \text{poly}(k/\varepsilon)} \|A\|_p \\
&= (1 + O(\varepsilon))OPT + O\left(\frac{\varepsilon^2}{f}\right) \|A\|_p
\end{aligned} \tag{15}$$

Finally, we analyze the running time of Algorithm 5.

Claim 3.22. *The running time of Algorithm 5 is at most $f^{O(rk)} + 2^{O(rk \log(k/\varepsilon))} + \text{poly}(fnd/\varepsilon)$, where r , the number of rows in the p -stable matrix S , is at most $O(\max(k/\varepsilon^6 \log(k/\varepsilon), 1/\varepsilon^9))$*

Proof. First, let us find the number of matrices in $|\mathcal{C}|$. Note that each $M \in \mathcal{C}$ has rk entries (recall that r is the number of rows of S). For each entry, there are $|\mathcal{Z}|$ choices. The magnitude of the largest possible guess is $\text{poly}(k/\varepsilon)\|A\|_p$, while that of the smallest possible guess is $\frac{1}{f \text{poly}(k/\varepsilon)}\|A\|_p$. Therefore, the number of guesses for each entry is $\frac{O(\log(fk/\varepsilon))}{\log(1 + \frac{1}{f \text{poly}(k/\varepsilon)})}$. Since for $x < 1$, $\log(1 + x) \geq \frac{x}{2}$, the number of guesses per each entry of M is in fact $O(f \text{poly}(k/\varepsilon) \log(fk/\varepsilon))$. In summary,

$$|\mathcal{C}| = (f \text{poly}(k/\varepsilon) \log(fk/\varepsilon))^{rk} = f^{O(rk)} 2^{O(rk \log(k/\varepsilon))}$$

Let us now calculate the running time needed for each guess. The size of COSTBOUNDS is

$$\log\left(\frac{fd}{\varepsilon}\right) \cdot \frac{1}{\log(1 + \varepsilon)} \leq O\left(\frac{1}{\varepsilon} \log\left(\frac{fd}{\varepsilon}\right)\right)$$

For each $i \in [d]$ and $c \in \text{COSTBOUNDS}$, we solve the problem of minimizing $\|V_i\|_p$ subject to the constraint that $\text{med}(MV_i - SA_i)/\text{med}_p \leq c$. Computing SA_i takes $O(nr)$ time. If $V_{i,c,best}$ is the solution to this problem, then there are at most $r!$ possible orderings for the coordinates of $MV_{i,c,best} - SA_i$. Therefore, we can make this optimization problem into a convex program by trying all orderings of the coordinates of $MV_i - SA_i$ (meaning that the orderings are added in the form of $\text{poly}(k/\varepsilon)$ additional constraints). Once we do this, $\text{med}(MV_i - SA_i)/\text{med}_p \leq c$ becomes a linear constraint. Hence, we solve $r!$ distinct convex programs. Since M is an $r \times k$ matrix, and SA_i is an r -dimensional vector, finding $V_{i,c,best}$ takes at most $r! \cdot \text{poly}(r)$ running time. In summary, the total time taken to find $V_{i,c}$ and $C_{i,c}$ for all $i \in [d]$ and $c \in \text{COSTBOUNDS}$ is

$$\begin{aligned}
d \cdot |\text{COSTBOUNDS}| \cdot r! \cdot \text{poly}(r) + O(nr) &\leq O\left(\frac{d}{\varepsilon} \log\left(\frac{fd}{\varepsilon}\right)\right) \cdot r^{O(r)} + O(nr) \\
&\leq \text{poly}(fnd/\varepsilon) 2^{O(r \log r)}
\end{aligned} \tag{16}$$

Next comes the running time for the convex program used to find the $x_{i,c}$. As calculated above, the number of variables is $d \cdot |\text{COSTBOUNDS}| = O\left(\frac{d}{\varepsilon} \log\left(\frac{fd}{\varepsilon}\right)\right)$ and the number of constraints is $d|\text{COSTBOUNDS}| + d + 2 = O\left(\frac{d}{\varepsilon} \log\left(\frac{fd}{\varepsilon}\right)\right)$. Therefore, the convex program can be solved in $\text{poly}(fd/\varepsilon)$ time.

Finally, we calculate the running time of the stage where we sample c_i for every $i \in [d]$ in order to find a right factor V' . The time needed to sample c_i is $d \cdot |\text{COSTBOUNDS}| = O\left(\frac{d}{\varepsilon} \log\left(\frac{fd}{\varepsilon}\right)\right)$, and for each sample of the c_i , computing the norm of the new V' takes $O(nk)$ time, and computing $\sum_{i \in [d]} \text{med}(MV'_i - SA_i)^p / \text{med}_p^p$ takes $n \cdot \text{poly}(k)$ time. The number of samples is $O(1/\varepsilon)$, meaning the total running time of this stage is $\text{poly}(fnd/\varepsilon)$. After V' has been found, the running time needed to find U' by solving $\min_U \|UV' - A\|_p$ using convex programming is $\text{poly}(nd)$. In summary, the total running time of Algorithm 5 is

$$f^{O(rk)} 2^{O(rk \log(k/\varepsilon))} \left(\text{poly}(fnd/\varepsilon) 2^{O(r \log r)} + \text{poly}(fnd/\varepsilon) \right)$$

and this is at most $f^{O(rk)} 2^{O(rk \log(k/\varepsilon))} \text{poly}(fnd/\varepsilon)$. The runtime in the theorem statement is due to the inequality $abc \leq \frac{a^3+b^3+c^3}{3}$, for $a, b, c \geq 0$. \square

\square

We use this to obtain a $(1 + \varepsilon)$ -approximation algorithm with bicriteria rank $3k$.

Theorem 3.23 (Correctness and Running Time of Algorithm 6). *Let $A \in \mathbb{R}^{n \times d}$, $k \in \mathbb{N}$, and $\varepsilon \in (0, c)$ where c is a sufficiently small absolute constant. Then, Algorithm 6, with these inputs, returns $\widehat{A} \in \mathbb{R}^{n \times d}$ such that \widehat{A} has rank $3k$ and*

$$\|\widehat{A} - A\|_p \leq (1 + O(\varepsilon)) \min_{A_k \text{ rank } k} \|A - A_k\|_p$$

The running time of Algorithm 6 is $2^{O(rk \log(k/\varepsilon))} + \text{poly}(nd/\varepsilon)$, where r is the same as in the statement of Theorem 3.15.

Proof. Let A_k be the optimal rank- k approximation for A , and suppose

$$(1 - O(\varepsilon)) \text{OPT}_{C,2k} \leq \widehat{\text{OPT}} \leq (1 + O(\varepsilon)) \text{OPT}_{C,2k}$$

where $\text{OPT}_{C,2k}$ is the error from the optimal rank- $2k$ approximation for C (as the algorithm will indeed guess such an $\widehat{\text{OPT}}$ at some point). If this is the case, then when we call

$$\text{GUESSINGADDITIVEEPSAPPROXIMATION}(C, 2k, \varepsilon, f, \widehat{\text{OPT}}, p)$$

we obtain U, V such that

$$\begin{aligned} \|UV - C\|_p &\leq (1 + O(\varepsilon)) \min_{C_{2k} \text{ rank } 2k} \|C_{2k} - C\|_p + O\left(\frac{\varepsilon}{f}\right) \|C\|_p \\ &\leq (1 + O(\varepsilon)) \|(A_k - B) - (A - B)\|_p + O\left(\frac{\varepsilon}{\text{poly}(k)}\right) \cdot \text{poly}(k) \|A - A_k\|_p \\ &\leq (1 + O(\varepsilon)) \|A_k - A\|_p + O(\varepsilon) \|A - A_k\|_p \\ &\leq (1 + O(\varepsilon)) \|A_k - A\|_p \end{aligned} \tag{17}$$

where the second inequality is because $A_k - B$ has rank at most $2k$, and because B is the result of a $\text{poly}(k)$ -approximation for A . Since $\|UV - C\|_p = \|UV - (A - B)\|_p = \|(UV + B) - A\|_p$, this completes the proof of correctness for Algorithm 6. Note that $UV + B$ has rank at most $3k$ because UV has rank at most $2k$ and B has rank at most k .

Now, we analyze the running time of Algorithm 6. First, the running time of Algorithm 8 (used to find B) is $\text{poly}(nd)$. The running time of `GUESSINGADDITIVEEPSAPPROXIMATION` is at most $f^{O(rk)} + 2^{O(rk \log(k/\varepsilon))} + \text{poly}(fnd/\varepsilon)$, and since $f = \text{poly}(k)$, this is

$$2^{O(rk \log k)} + 2^{O(rk \log(k/\varepsilon))} + \text{poly}(nd/\varepsilon) = \boxed{2^{O(rk \log(k/\varepsilon))} + \text{poly}(nd/\varepsilon)}$$

The number of times Algorithm 6 is called is $O((\log nd)/\varepsilon)$, meaning the overall running time of Algorithm 6 is also the above. All other steps in Algorithm 6 can be done in polynomial time. \square

A Hardness for ℓ_p Low Rank Approximation with Additive Error, based on [2]

A.1 Background: Small Set Expansion Hypothesis

The hardness proof in [2] proceeds by a reduction from the Small Set Expansion problem — our presentation of this problem follows that of [2].

Problem 2 (Small Set Expansion Problem - As Presented in [2]). *Let $G = (V, E)$ be a regular graph. For any subset $S \subset V$, the measure of S is defined to be $\mu(S) := |S|/|V|$. The distribution $G(S)$ over the vertices of G is generated as follows: first a uniformly random vertex $x \in S$ is selected, then a uniformly random neighbor y of x is selected (as a sample of $G(S)$). For $S \subset V$, the expansion of S is defined to be $\Phi_G(S) := \Pr_{y \sim G(S)}[y \notin S]$. Finally, for $\delta \in (0, 1)$, $\Phi_G(\delta) := \min_{S \subset V | \mu(S) \leq \delta} \Phi_G(S)$.*

The Small Set Expansion Problem is as follows: given a graph $G = (V, E)$ and $\varepsilon, \delta \in (0, 1)$, the goal is to decide whether $\Phi_G(\delta) \leq \varepsilon$ or $\Phi_G(\delta) \geq 1 - \varepsilon$.

In other words, $\Phi_G(S)$ is the proportion of neighbors that S has, which do not belong to S — this can be considered as the “expansion” of S — and $\Phi_G(\delta)$ is the smallest expansion among all subsets of V which have at most $\delta|V|$ vertices. The Small Set Expansion Hypothesis is as follows:

Conjecture 1 (Small Set Expansion Hypothesis — Conjecture 1.3 of [43]). *For any fixed $\varepsilon \in (0, 1)$, there exists $\delta \in (0, 1)$ such that it is NP-hard to decide whether $\Phi_G(\delta) \leq \varepsilon$ or $\Phi_G(\delta) \geq 1 - \varepsilon$.*

A.2 Background: Hardness Proof from [2] — ℓ_p Low Rank Approximation for $p \in (1, 2)$

Now, we summarize the reduction due to [2] from the Small Set Expansion Problem to ℓ_p low rank approximation for $p \in (1, 2)$. The reduction in Section 5 of [2] shows that given $k \in \mathbb{N}$ and $p \in (1, 2)$, and given a matrix $A \in \mathbb{R}^{n \times d}$, it is NP-hard to find an $O(1)$ -approximation for the error from the best rank- k approximation for A in the ℓ_p -norm. Ultimately, [2] reduces from Problem 2 to finding the best rank- $(k - 1)$ approximation for a $k \times k$ matrix. The first step is to reduce from Problem 2 to computing the $2 \rightarrow p^*$ norm, where p^* is the Hölder conjugate of p :

Theorem A.1 (Theorem 2.4 of [3] for $q \in 2\mathbb{Z} \setminus \{2\}$, and Theorem 21 of [2] for general $q \in (2, \infty)$). *Let G be a regular graph, $\lambda \in (0, 1)$ and $q \in (2, \infty)$. Let M be the normalized adjacency matrix of G , and let $V_{\geq \lambda}(G)$ be the subspace spanned by the eigenvectors of M with eigenvalue at least λ . Finally, let $P_{\geq \lambda}(G)$ be the orthogonal projection matrix onto $V_{\geq \lambda}(G)$. Then,*

- For all $\delta > 0$, $\varepsilon > 0$, $\|P_{\geq \lambda}(G)\|_{2 \rightarrow q} \leq \varepsilon/\delta^{(q-2)/2q}$ implies that $\Phi_G(\delta) \geq 1 - \lambda - \varepsilon^2$.
- There is a constant $a = a(q)$ such that for all $\delta > 0$, $\Phi_G(\delta) > 1 - a\lambda^{2q}$ implies $\|P_{\geq \lambda}(G)\|_{2 \rightarrow q} \leq 2/\sqrt{\delta}$.

This implies that the $2 \rightarrow q$ norm is hard to approximate — for completeness, we include the proof of this from [2]. In the next subsection, we show that this proof can be modified to obtain hardness for a multiplicative $O(1)$ -approximation for ℓ_p low rank approximation with additive $\frac{1}{2^{\text{poly}(k)}}\|A\|_p$ error.

Theorem A.2 (Theorem 7 of [2]). *Assuming the Small Set Expansion Hypothesis, for any $q \in (2, \infty)$, and $r > 1$, it is NP-hard to approximate the $\|\cdot\|_{2 \rightarrow q}$ norm within a factor r .*

Proof. (From [2], Page 46). Using [44], the Small Set Expansion Hypothesis implies that for any sufficiently small numbers $0 < \delta \leq \delta'$, there is no polynomial time algorithm that can distinguish between the following cases for a given graph G :

- Yes case: $\Phi_G(\delta) < 0.1$
- No case: $\Phi_G(\delta') > 1 - 2^{-a' \log(1/\delta')}$

In particular, for all $\eta > 0$, if we let $\delta' = \delta^{(q-2)/8q}$ and make δ small enough, then in the No case $\Phi_G(\delta^{(q-2)/8q}) > 1 - \eta$. (Since $q > 2$, $\delta' \rightarrow 0$ as $\delta \rightarrow 0$.)

Using Theorem A.1, in the Yes case we know $\|P_{\geq 1/2}(G)\|_{2 \rightarrow q} \geq 1/(10\delta^{(q-2)/2q})$, while in the No case, if we choose δ sufficiently small so that η is smaller than $a(1/2)^{2q}$, then we know that $\|P_{\geq 1/2}(G)\|_{2 \rightarrow q} \leq 2/\sqrt{\delta'} = 2/\delta^{(q-2)/4q}$. The gap between the Yes case and the No case is at least $\delta^{-(q-2)/4q}/20$, which goes to ∞ as δ decreases. \square

This proof shows that computing $\|P_{\geq 1/2}(G)\|_{2 \rightarrow p^*}$ within a constant factor is sufficient to decide the Small Set Expansion Problem. Now (if G is assumed to be a graph with k vertices) the following steps are used to reduce the problem of computing the $2 \rightarrow p^*$ norm of the $k \times k$ matrix $P_{\geq 1/2}(G)$ to the problem of rank- $(k-1)$ ℓ_p low rank approximation. For convenience, let $P := P_{\geq 1/2}(G)$.

First, computing the $2 \rightarrow p^*$ norm of P is equivalent to computing the $2 \rightarrow p^*$ norm of some invertible matrix P_1 which can be constructed in poly(k) time:

Lemma A.3 (Claim 14 of [2]). *Let A be a nonzero $n \times d$ matrix. For any $p, q \in (1, \infty)$ and any $\varepsilon > 0$, there is an invertible and polynomial-time computable $\max(n, d) \times \max(n, d)$ matrix B such that $(1 - \varepsilon)\|A\|_{p \rightarrow q} \leq \|B\|_{p \rightarrow q} \leq (1 + \varepsilon)\|A\|_{p \rightarrow q}$.*

The proof of the above lemma proceeds as follows: first, A is made into a square matrix by adding rows/columns of zeros — then, $\frac{\varepsilon M}{\|I\|_{p \rightarrow q}}$ is added to every diagonal entry to make it invertible, where M is the absolute value of the largest entry of A . Hence, in our case, we will add $\frac{\varepsilon M}{\|I\|_{2 \rightarrow p^*}}$ to every diagonal entry of P to obtain P_1 . The next step, after obtaining the invertible matrix P_1 , is to compute a matrix P_2 such that finding $\|P_2\|_{p \rightarrow p^*}$ allows us to find $\|P_1\|_{2 \rightarrow p^*}$:

Lemma A.4 (Claim 13 of [2]). *For any $A \in \mathbb{R}^{n \times d}$ and $p \in (1, \infty)$, if p^* is its Hölder conjugate, then $\|AA^T\|_{p \rightarrow p^*} = \|A\|_{2 \rightarrow p^*}^2$.*

Hence, by this claim, the appropriate P_2 will simply be $P_1 P_1^T$. Finally, it is useful to reduce this problem to computing $\min_{p^* \rightarrow p}(P_3)$ of some well-chosen matrix P_3 , since computing $\min_{p^* \rightarrow p}$ was shown by [2] to be equivalent to rank- $(k-1)$ ℓ_p -low rank approximation:

Lemma A.5 (Fact 4 of [2]). *For $p, q \in (1, \infty)$, if A is an invertible matrix, then $\min_{p \rightarrow q}(A^{-1}) = (\|A\|_{q \rightarrow p})^{-1}$.*

Lemma A.6 (Lemma 1/Lemma 27 of [2] - Equivalence of $\min_{p^* \rightarrow p}$ and Rank- $(k-1)$ ℓ_p Low Rank Approximation). *Let $p \in (1, \infty)$ and let p^* be the Hölder conjugate of p . Let $A \in \mathbb{R}^{n \times d}$ with $n \geq d$ and $k = d - 1$. Then,*

$$\min_{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}} \|UV - A\|_p = \min_{x \in \mathbb{R}^d, \|x\|_{p^*} = 1} \|Ax\|_p$$

Hence, P_3 is in fact P_2^{-1} , and the $2 \rightarrow p^*$ norm of $P_{\geq 1/2}(G)$ can be approximated up to a constant factor if $\min_{p^* \rightarrow p}(P_3) = \min_M \text{rank } k-1 \|M - P_3\|_p$ can be approximated up to a constant factor. In [2], it is also mentioned that assuming the Exponential-Time Hypothesis [23] together with SSEH implies that the running time required to obtain an $O(1)$ -approximation is $2^{k^{\Omega(1)}}$ — this is also true for our hardness result, as we mention in the proof of Theorem A.24.

A.3 Deterministic Embedding from ℓ_p^n to $\ell_1^{n^{O(\log n)}}$

We can construct a matrix that can be used to embed ℓ_p^n into ℓ_1^r , where $r = n^{O(\log n)}$ — the matrix can be constructed deterministically, in $n^{O(\log n)}$ time. This embedding is then used to reduce ℓ_p low rank approximation to $\ell_{1,p}$ low rank approximation, and then to reduce $\ell_{1,p}$ low rank approximation to constrained ℓ_1 low rank approximation. Throughout this section, we will use the following notation: for two functions $f, g > 0$, $f = O_p(g)$ if $f \leq C_p g$ for some constant C_p depending on p . We also define Ω_p and Θ_p similarly. In addition, using the notation of [27], we say that $x \approx_\varepsilon y$ for two $x, y \in \mathbb{R}$ if $|x - y| \leq \varepsilon$.

Theorem A.7 (Deterministic Embedding of ℓ_p^n into ℓ_1^r). *Let $n \in \mathbb{N}$, and $p \in (1, 2)$. Then, there exists a matrix $R \in \mathbb{R}^{r \times n}$, where $r = n^{O(\log n)}$, such that for all $x \in \mathbb{R}^n$,*

$$\Omega_p(1)\|x\|_p \leq \|Rx\|_1 \leq O_p(1)\|x\|_p$$

R can be constructed deterministically in $n^{O(\log n)}$ time.

Our proof of Theorem A.7 will be based on the following observation from [38]:

Fact 2 (Observation 2.7.2 of [38]). *Let R_1, R_2, \dots, R_n be real random variables on a probability space that has k elements $\omega_1, \omega_2, \dots, \omega_k$, and let $A \in \mathbb{R}^{k \times n}$ such that $A_{i,j} = \Pr[\omega_i]R_j(\omega_i)$. For $x \in \mathbb{R}^n$ let us set $X := \sum_{j=1}^n R_j x_j$. Then, $E[|X|] = \|Ax\|_1$.*

Our embedding matrix R will be constructed in the same way A is constructed in the above observation. The random variables R_i will be standard p -stable random variables (discretized appropriately) and will be selected so that they are $O(\log n)$ -wise independent, rather than fully independent, so that the sample space only has size $n^{O(\log n)}$. We use lemmas from [27] which are useful when dealing with identically distributed random variables that have limited independence. The first of these lemmas gives smooth approximations to indicator functions of an interval:

Lemma A.8 (Lemma 2.5 of [27]). *There exist constants $c', \varepsilon_0 > 0$ such that for all $c > 0$ and $0 < \varepsilon < \varepsilon_0$, and for all $[a, b] \subset \mathbb{R}$, there exists a function $J_{[a,b]}^c : \mathbb{R} \rightarrow \mathbb{R}$ satisfying:*

- $\|(J_{[a,b]}^c)^{(\ell)}\|_\infty = O(c^\ell)$ for all $\ell \geq 0$.
- For all x such that $a, b \notin [x - \varepsilon, x + \varepsilon]$, and as long as $c > c'\varepsilon^{-1} \log^3(1/\varepsilon)$, $|J_{[a,b]}^c(x) - I_{[a,b]}(x)| < \varepsilon$.

Here, in the first item above, for a function f , $f^{(\ell)}$ denotes the ℓ^{th} derivative of f .

The following lemma from [27] allows us to show that, if X is a linear combination of n fully independent p -stable random variables, Y is a linear combination of n p -stable random variables which are only $O(1)$ -wise independent, and f is a smooth function, then $E[f(X)]$ and $E[f(Y)]$ are not too far apart. When we apply this lemma, f will be $J_{[a,b]}^c$ for some c and ε .

Lemma A.9 (Lemma 2.2 of [27]). *There exists an $\varepsilon_0 > 0$ such that the following holds. Let n be a positive integer and $0 < \varepsilon < \varepsilon_0$, $0 < p < 2$ be given. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ satisfy $\|f^{(\ell)}\|_\infty = O(\alpha^\ell)$ for all $\ell \geq 0$, for some α satisfying $\alpha^p \geq \log(1/\varepsilon)$. Let $k = \alpha^p$. Let $a \in \mathbb{R}^n$ satisfy $\|a\|_p = O(1)$. Let X_i be a $3Ck$ independent family of p -stable random variables for C a suitably large even constant. Let Y_i be a fully independent family of p -stable random variables. Let $X = \sum_i a_i X_i$ and $Y = \sum_i a_i Y_i$. Then, $E[f(X)] = E[f(Y)] + O(\varepsilon)$.*

Remark A.10. *Note that this lemma only requires that the X_i are $3Ck$ -wise independent for some constant C and some $k = O(\log 1/\varepsilon)$ — since we will anyways choose $\varepsilon = \Theta(1)$, this does not affect the amount of independence that we need.*

It is also useful to note that p -stable random variables have finite expectation for $p \in (1, 2)$:

Lemma A.11. *Let $p \in (1, 2)$ and suppose X is a standard p -stable random variable. Then, $E[|X|] = O_p(1)$.*

Proof. Note that the p.d.f. of X , which we denote $p_X(x)$, is $\Theta(1/x^{p+1})$ for large x :

Theorem A.12 (Theorem 1.12 of [40]). *Let X be a p -stable random variable with p.d.f. $p_X(x)$. Then,*

$$\lim_{x \rightarrow \infty} \frac{p_X(x)}{(pC_p)/x^{p+1}} = 1$$

for some constant C_p depending on p .

Remark A.13. *Theorem 1.12 in [40] is more general and also applies when X is not symmetric.*

Hence, $p_X(x) \leq O_p(\frac{1}{x^{p+1}})$ (for some constant K and all x with $|x| \geq K$), meaning

$$\begin{aligned} E[|X|] &\leq \int_{-K}^K |x| p_X(x) dx + 2 \int_K^\infty x p_X(x) dx \\ &\leq 2K + 2 \int_K^\infty x \cdot O\left(\frac{1}{x^{p+1}}\right) dx \\ &\leq 2K + 2 \int_K^\infty O\left(\frac{1}{x^p}\right) dx \\ &\leq O_p(1) \end{aligned} \tag{18}$$

where the last inequality holds because the integral $\int_K^\infty \frac{1}{x^p} dx$ converges, since $p > 1$. \square

Also, we note that linear combinations of $O(1)$ -wise independent p -stable random variables are anti-concentrated:

Lemma A.14 (Anti-Concentration of $O(1)$ -wise Independent p -stable Random Variables (Due to [27])). *Let R_1, R_2, \dots, R_n be an $O(\log(1/\varepsilon))$ -wise independent family of p -stable random variables, and $x \in \mathbb{R}^n$ with $\|x\|_p = O(1)$. Then, for any $t \in \mathbb{R}$,*

$$\Pr\left[\left|\left(\sum_{i=1}^n R_i x_i\right) - t\right| \leq \varepsilon\right] \leq O(\varepsilon)$$

In particular, if X is a standard p -stable random variable, then $\Pr[X \in [t - \varepsilon, t + \varepsilon]] \leq O(\varepsilon)$ for $\varepsilon \in (0, \varepsilon_0)$ and any $t \in \mathbb{R}$.

Proof. Let $R = \sum_{i=1}^n R_i x_i$. The main tool in showing this lemma is the following result from [27]:

Lemma A.15 ((Shown in Section A.4 of [27])). *For some constant ε_0 and any $t \in \mathbb{R}$ and $\varepsilon \in (0, \varepsilon_0)$, there exists a nonnegative function $f_{t,\varepsilon} : \mathbb{R} \rightarrow \mathbb{R}$ such that:*

- For some $\alpha = O(1/\varepsilon)$, $\|f_{t,\varepsilon}^{(\ell)}\|_\infty = O(\alpha^\ell)$ for all $\ell \geq 0$.
- If X is a standard p -stable random variable, then $E[f_{t,\varepsilon}(X)] = O(\varepsilon)$.
- $f_{t,\varepsilon}(t + \varepsilon) = \Omega(1)$
- As $|x - t|$ increases, $f_{t,\varepsilon}(x)$ decreases.

The rest of the proof follows that of a similar claim within Theorem 2.1 of [27]. Let S_1, S_2, \dots, S_n be a fully independent family of p -stable random variables. Note that $S = \sum_{i=1}^n S_i x_i$ is a standard p -stable random variable, and by Lemma A.15, $E[f_{t,\varepsilon}(S)] = O(\varepsilon)$. Now, by Lemma A.9, since the R_i are $O(\log(1/\varepsilon))$ -wise independent,

$$E[f_{t,\varepsilon}(R)] \leq E[f_{t,\varepsilon}(S)] + O(\varepsilon) \leq O(\varepsilon)$$

where $E[f_{t,\varepsilon}(S)] \leq O(\varepsilon)$ by Lemma A.15. Finally, for all $r \in [t - \varepsilon, t + \varepsilon]$, $f_{t,\varepsilon}(r) \geq f_{t,\varepsilon}(t + \varepsilon)$, since $f_{t,\varepsilon}(x)$ is decreasing in $|x - t|$. Hence,

$$E[f_{t,\varepsilon}(R)] \geq \int_{t-\varepsilon}^{t+\varepsilon} f_{t,\varepsilon}(r) p_R(r) dr \geq f_{t,\varepsilon}(t + \varepsilon) \Pr[R \in [t - \varepsilon, t + \varepsilon]]$$

where p_R is the p.d.f. of R . Since $f_{t,\varepsilon}(t + \varepsilon) = \Omega(1)$, combining this with the previous chain of inequalities gives

$$O(\varepsilon) \geq \Omega(1) \Pr[R \in [t - \varepsilon, t + \varepsilon]]$$

and this gives the desired result. \square

We will study $E[|\langle Z, x \rangle|]$ for a fixed $x \in \mathbb{R}^n$, where Z is a vector whose entries are $O(\log n)$ -wise independent p -stable random variables, which are truncated at $\text{poly}(n)$, i.e. Z_i follows a p -stable distribution conditioned on $|Z_i| \leq \text{poly}(n)$. To analyze the tails of $|\langle Z, x \rangle|$, we use the following lemma from [26]:

Lemma A.16 (Lemma 23 of [26]). *Let $p \in (0, 2]$. Suppose $x \in \mathbb{R}^n$, $\|x\|_p = 1$, $0 < \varepsilon < 1$ is given, and R_1, \dots, R_n are k -wise independent p -stable random variables for $k \geq 2$. Let Q be a standard p -stable random variable. Then, for all $t \geq 0$, $R = \sum_{i=1}^n R_i x_i$ satisfies*

$$\left| \Pr[|Q| \geq t] - \Pr[|R| \geq t] \right| = O\left(\frac{k^{-1/p}}{1 + t^{p+1}} + \frac{k^{-2/p}}{1 + t^2} + 2^{-\Omega(k)}\right)$$

Now we show that $E[|\langle Z, x \rangle|] = \Theta(1)\|x\|_p$:

Lemma A.17 (Expected Value of $|\langle Z, x \rangle|$). *Let $n \in \mathbb{N}$ and $x \in \mathbb{R}^n$. For $i \in [n]$, let Z_i be a standard p -stable random variable truncated at $D = \text{poly}(n)$, and suppose the Z_i are $O(\log n)$ -wise independent. Let $Z \in \mathbb{R}^n$ so that the i^{th} coordinate of Z is Z_i . Then, $E[|\langle Z, x \rangle|] = \Theta_p(1)\|x\|_p$.*

Proof. Let $x \in \mathbb{R}^n$ with $\|x\|_p = 1$, and define Z as in the statement of this lemma. Here, the amount of independence of the coordinates of Z is actually $O(\max(\log n, \log(1/\varepsilon)))$, where ε is a constant that will be chosen to be sufficiently small later. We will prove the lemma in two steps:

- With constant probability, $|\langle Z, x \rangle| \geq \Omega(1)\text{med}_p$ which is $\Omega(1)$ by Lemma 3.13. Therefore, $E[|\langle Z, x \rangle|] \geq \Pr[|\langle Z, x \rangle| \geq \Omega(1)] \cdot \Omega(1) = \Omega(1)$.
- $E[|\langle Z, x \rangle|] \leq O_p(1)$.

Then, we can extend to general x by multiplying both sides by $\|x\|_p$.

The proof of the first step is similar to that of Theorem 2.1 of [27], and Lemma I.3 of [48]. First assume for simplicity that the coordinates of Z are standard p -stable random variables that are not truncated at $D = \text{poly}(n)$ — for $D = n^c$ and c sufficiently large, the probability of $\|Z\|_\infty \leq D$ is $1 - o(1)$, so this assumption does not affect the probability calculations in the proof of the first step. Let $b > 0$ be a small constant, chosen so that if X is a p -stable random variable, then with probability 0.999, $|X| \geq b$. In addition, define $Y \in \mathbb{R}^n$ so that the coordinates of Y are i.i.d. p -stable random variables. Note that $\langle Y, x \rangle$ is equal to $\|x\|_p X = X$, where X is a p -stable random variable. Hence, if $1_{[-b, b]}$ is the indicator function of $[-b, b]$, then

$$E[1_{[-b, b]}(\langle Y, x \rangle)] = \Pr[|\langle Y, x \rangle| \leq b] \leq 0.001$$

For convenience, let $f_0 = 1_{[-b, b]}$, and define $f = J_{[-b, b]}^c$, as in Lemma A.8, for some sufficiently small $\varepsilon > 0$ and sufficiently large c . Now, we wish to show that $E[1_{[-b, b]}(\langle Z, x \rangle)] = \Pr[|\langle Z, x \rangle| \leq b]$ is small. Following the proofs of Theorem 2.1 of [27] and Lemma I.3 of [48], we will proceed by showing that

$$E[f_0(\langle Y, x \rangle)] \approx_{O(\varepsilon)} E[f(\langle Y, x \rangle)] \approx_{O(\varepsilon)} E[f(\langle Z, x \rangle)] \approx_{O(\varepsilon)} E[f_0(\langle Z, x \rangle)]$$

First, we show that $|E[f_0(\langle Y, x \rangle)] - E[f(\langle Y, x \rangle)]| \leq O_p(\varepsilon)$ as follows. If $\langle Y, x \rangle$ is not in the intervals $[-b - \varepsilon, -b + \varepsilon]$ or $[b - \varepsilon, b + \varepsilon]$, then $|f_0(\langle Y, x \rangle) - f(\langle Y, x \rangle)| \leq \varepsilon$ by the construction in Lemma A.8. On the other hand, since $\|x\|_p = 1$, $\langle Y, x \rangle$ is equal to X , where X is a standard p -stable random variable. By Lemma A.14, X is anti-concentrated, meaning

$$\Pr[X \in [-b - \varepsilon, -b + \varepsilon]] = \Pr[X \in [b - \varepsilon, b + \varepsilon]] = O_p(\varepsilon)$$

Moreover, $|f_0| \leq 1$ and $|f| \leq O(1)$ by the construction in Lemma A.8. Hence, if we let \mathcal{E}_1 and \mathcal{E}_2 be the events that $\langle Y, x \rangle$ is in $[-b - \varepsilon, -b + \varepsilon]$ and $[b - \varepsilon, b + \varepsilon]$ respectively, then

$$\begin{aligned} \left| E[f_0(\langle Y, x \rangle)] - E[f(\langle Y, x \rangle)] \right| &\leq E\left[|f_0(\langle Y, x \rangle) - f(\langle Y, x \rangle)| \right] \\ &\leq E\left[|f_0(\langle Y, x \rangle) - f(\langle Y, x \rangle)| \mid \neg\mathcal{E}_1 \cap \neg\mathcal{E}_2 \right] \\ &\quad + E\left[|f_0(\langle Y, x \rangle) - f(\langle Y, x \rangle)| \mid \mathcal{E}_1 \cup \mathcal{E}_2 \right] \\ &\leq O_p(\varepsilon) + \Pr[\mathcal{E}_1 \cup \mathcal{E}_2] \cdot O(1) \\ &\leq O_p(\varepsilon) \end{aligned} \tag{19}$$

where the first inequality is by Jensen's inequality. Next, observe that $|E[f(\langle Y, x \rangle)] - E[f(\langle Z, x \rangle)]| \leq O(\varepsilon)$ due to Lemma A.9.

Finally, we show that $|E[f(\langle Z, x \rangle)] - E[f_0(\langle Z, x \rangle)]| \leq O(\varepsilon)$. By Lemma A.14, $\langle Z, x \rangle$ is anti-concentrated, and in particular, $\Pr[\langle Z, x \rangle \in [-b - \varepsilon, -b + \varepsilon]] \leq O(\varepsilon)$ and $\Pr[\langle Z, x \rangle \in [b - \varepsilon, b + \varepsilon]] \leq O(\varepsilon)$. Moreover, outside of the intervals $[-b - \varepsilon, -b + \varepsilon]$ and $[b - \varepsilon, b + \varepsilon]$, $|f - f_0| \leq O(\varepsilon)$. Hence, by the same reasoning used to show that $|E[f(\langle Y, x \rangle)] - E[f_0(\langle Y, x \rangle)]| \leq O(\varepsilon)$, we know that $|E[f(\langle Z, x \rangle)] - E[f_0(\langle Z, x \rangle)]| \leq O(\varepsilon)$. In summary, we have shown that

$$E[f_0(\langle Y, x \rangle)] \approx_{O(\varepsilon)} E[f(\langle Y, x \rangle)] \approx_{O(\varepsilon)} E[f(\langle Z, x \rangle)] \approx_{O(\varepsilon)} E[f_0(\langle Z, x \rangle)]$$

meaning

$$E[1_{[-b,b]}(\langle Z, x \rangle)] \leq E[1_{[-b,b]}(\langle Y, x \rangle)] + O(\varepsilon) \leq 0.001 + O(\varepsilon)$$

Letting ε be a sufficiently small constant, we find that $\Pr[|\langle Z, x \rangle| \geq b] \geq \Omega(1)$, meaning $E[|\langle Z, x \rangle|] \geq \Omega_p(1)$.

Now, we show that $E[|\langle Z, x \rangle|] \leq O_p(1)$. We now drop the assumption that the Z_i are not truncated — in other words, let the Z_i follow a standard p -stable distribution, conditioned on $|Z_i| \leq D = \text{poly}(n)$. Also, define $Z' \in \mathbb{R}^n$, so that the coordinates of Z' are standard (un-truncated) p -stable random variables, and the coordinates of Z' are $O(\log n)$ -wise independent. Finally, let X be a standard p -stable random variable.

First, we write the expectation of $|\langle Z, x \rangle|$ in terms of the tails of $|\langle Z', x \rangle|$ — then, we can relate the tails of $|\langle Z', x \rangle|$ to those of $|X|$ using Lemma A.16. Note that

$$\begin{aligned} E[|\langle Z, x \rangle|] &= E[|\langle Z', x \rangle| \mid \|Z'\|_\infty \leq D] \\ &= \int_0^\infty \Pr[|\langle Z', x \rangle| \geq t \mid \|Z'\|_\infty \leq D] dt \\ &= \int_0^{Dn} \Pr[|\langle Z', x \rangle| \geq t \mid \|Z'\|_\infty \leq D] dt \\ &= \int_0^{Dn} \frac{\Pr[|\langle Z', x \rangle| \geq t \text{ and } \|Z'\|_\infty \leq D]}{\Pr[\|Z'\|_\infty \leq D]} dt \\ &\leq O(1) \int_0^{Dn} \Pr[|\langle Z', x \rangle| \geq t] dt \end{aligned} \tag{20}$$

Here, the third equality is because, if $\|Z'\|_\infty \leq D$, then $\|Z'\|_2 \leq D\sqrt{n}$, and by the Cauchy-Schwarz inequality, $|\langle Z', x \rangle| \leq \|Z'\|_2 \|x\|_2 \leq D\sqrt{n}\sqrt{n} = Dn$. The first inequality holds because with probability $1 - o(1)$, $\|Z'\|_\infty \leq D$ as long as $D = n^c$ for sufficiently large c .

Hence, our task is now to show that the integral $\int_0^{Dn} \Pr[|\langle Z', x \rangle| \geq t] dt$ is at most $O_p(1)$. Since the coordinates of Z' are $O(\log n)$ -wise independent p -stable random variables (which are in particular not truncated), Lemma A.16 is now applicable, and we can use it to relate the integral to the expectation of $|X|$. First, note that

$$0 \leq \int_0^{Dn} \Pr[|X| \geq t] dt \leq \int_0^\infty \Pr[|X| \geq t] dt = E[|X|] \leq O_p(1)$$

and therefore, it suffices to show that

$$\left| \int_0^{Dn} \Pr[|\langle Z', x \rangle| \geq t] dt - \int_0^{Dn} \Pr[|X| \geq t] dt \right| \leq O_p(1)$$

as follows:

$$\begin{aligned} &\left| \int_0^{Dn} \Pr[|\langle Z', x \rangle| \geq t] dt - \int_0^{Dn} \Pr[|X| \geq t] dt \right| \\ &\leq \int_0^{Dn} \left| \Pr[|\langle Z', x \rangle| \geq t] - \Pr[|X| \geq t] \right| dt \\ &\leq \int_0^{Dn} O\left(\frac{(\log n)^{-1/p}}{1+t^{p+1}} + \frac{(\log n)^{-2/p}}{1+t^2} + 2^{-\Omega(\log n)}\right) dt \\ &\leq o(1) \int_0^{Dn} \frac{1}{1+t^{p+1}} dt + o(1) \int_0^{Dn} \frac{1}{1+t^2} dt + \int_0^{Dn} 2^{-\Omega(\log n)} dt \\ &\leq o(1) \int_0^\infty \frac{1}{1+t^{p+1}} dt + o(1) \int_0^\infty \frac{1}{1+t^2} dt + \int_0^{Dn} 2^{-\Omega(\log n)} dt \\ &\leq o(1) + \int_0^{Dn} 2^{-\Omega(\log n)} dt \\ &\leq o(1) \end{aligned} \tag{21}$$

Here, the second inequality is by Lemma A.16, and the last inequality holds if the coordinates of Z and Z' are $(C \log n)$ -wise independent for some sufficiently large C , since $Dn = \text{poly}(n)$. \square

In order to have finitely many values for the Z_i (and hence a finite sample space) we can round the Z_i to the nearest integer multiple of $\frac{1}{\text{poly}(n)}$:

Lemma A.18 (Expected Value of $|\langle Z, x \rangle|$ After Rounding Z). *Let $n \in \mathbb{N}$ and $x \in \mathbb{R}^n$. For $i \in [n]$, let Z_i be a standard p -stable random variable truncated at $D = \text{poly}(n)$, and rounded to the nearest multiple of $\frac{1}{F}$, where $F = \text{poly}(n)$. Suppose the Z_i are $O(\log n)$ -wise independent. Let $Z \in \mathbb{R}^n$ so that the i^{th} coordinate of Z is Z_i . Then, $E[|\langle Z, x \rangle|] = \Theta_p(1)\|x\|_p$.*

Proof. Let $Y \in \mathbb{R}^n$ so that the coordinates of Y are $O(\log n)$ -wise independent p -stable random variables truncated at $D = \text{poly}(n)$, and suppose $Z \in \mathbb{R}^n$ so that the coordinates of Z are those of Y , but rounded to the nearest multiple of $\frac{1}{F}$ where $F = \text{poly}(n)$. Finally, fix $x \in \mathbb{R}^n$ with $\|x\|_p = 1$. Then, by Lemma A.17, $E[|\langle Y, x \rangle|] = \Theta_p(1)$. Therefore,

$$\begin{aligned} \left| E[|\langle Z, x \rangle|] - E[|\langle Y, x \rangle|] \right| &\leq E\left[\left| |\langle Z, x \rangle| - |\langle Y, x \rangle| \right| \right] \\ &\leq E\left[|\langle Z - Y, x \rangle| \right] \\ &\leq \frac{n}{F} \end{aligned} \tag{22}$$

and to have $E[|\langle Z, x \rangle|] = \Theta_p(1)$, it suffices to choose $F = Cn$ for a sufficiently large absolute constant C . \square

Hence, we can assume that each of the Z_i has at most $\text{poly}(n)$ values. Now, since we only need the Z_i to be $O(\log n)$ -wise independent, a sample space of size $n^{O(\log n)}$ suffices to construct the embedding matrix in Fact 2, and this sample space can be efficiently constructed by the following theorem of [32]:

Theorem A.19 (Theorem 3.2 of [32]). *Let $X = (X_1, X_2, \dots, X_n)$ be a random vector, such that the X_i are identically distributed and take values in $[r]$. In addition, suppose the X_i satisfy all independence constraints belonging to a set \mathcal{C} , i.e., \mathcal{C} consists of subsets of $[n]$ such that, for all $S \in \mathcal{C}$,*

$$\Pr[X_i = x_i \forall i \in S] = \prod_{i \in S} \Pr[X_i = x_i]$$

Then, it is possible to construct a joint distribution over the X_i with sample space Ω , such that $|\Omega| \leq |\mathcal{C}|$, in running time $O(rn|\mathcal{C}|^{2.62})$.

In particular, if we wish for the X_i to be k -wise independent (corresponding to the case when $\mathcal{C} = \binom{[n]}{k}$) then $|\Omega| \leq \binom{n}{k}$ and the running time is $O(rn^{O(k)})$.

In summary, combining Fact 2, Lemma A.18 and Lemma A.19, we obtain the following:

Theorem A.20 (Deterministic Embedding of ℓ_p^n into $\ell_1^{n^{O(\log n)}}$). *Let $n \in \mathbb{N}$, and $p \in (1, 2)$. Then, there exists a matrix $R \in \mathbb{R}^{r \times n}$, where $r = n^{O(\log n)}$, such that for all $x \in \mathbb{R}^n$,*

$$\Omega_p(1)\|x\|_p \leq \|Rx\|_1 \leq O_p(1)\|x\|_p$$

R can be constructed deterministically in $n^{O(\log n)}$ time.

Remark A.21. *The running time depends on p as well, and goes to infinity as $p \rightarrow 1$. In particular, the number of terms in the Taylor series of e^x needed below depends on p . This does not affect our reductions later on, since we can treat p as a constant when reducing $\ell_{1,p}$ low rank approximation to constrained ℓ_1 low rank approximation for a fixed p (and similarly when reducing ℓ_p low rank approximation to $\ell_{1,p}$ low rank approximation).*

Proof. We apply Theorem A.19 with the X_i being $O(\log n)$ -wise independent. Here, X has the same distribution as the Z specified in the statement of Lemma A.18. Each of the coordinates Z_i of Z takes on values between $-D$ and D , where $D = \text{poly}(n)$, in increments of $\frac{1}{F}$ where $F = \text{poly}(n)$. To run the procedure described in Theorem A.19, we must approximately compute the distribution of Z_i , meaning we must find $\Pr[t \leq Z_i \leq t + F]$, for all integer multiples t of F that are between $-D$ and D .

Claim A.22. *Each of the probabilities $\Pr[t \leq Z_i \leq t + F]$ can be computed with error at most $\frac{1}{\text{poly}(n)}$ in $\text{poly}(n)$ time.*

Proof. Computing these probabilities reduces to the problem of computing the c.d.f. of a standard p -stable random variable at t , where the numerator and denominator of t are at most $\text{poly}(n)$. Let X be a standard p -stable random variable. Then, as noted in the proof of Lemma 3.13, it is known due to [41] that for $x > 0$

$$\Pr[X > x] = 1 - \frac{1}{\pi} \int_0^{\frac{\pi}{2}} e^{-x^{\frac{p}{p-1}} \cdot V(\theta;p)} d\theta$$

where

$$V(\theta;p) = \left(\frac{\cos \theta}{\sin p\theta} \right)^{\frac{p}{p-1}} \cdot \frac{\cos(p-1)\theta}{\cos \theta}$$

Hence,

$$\Pr[X \leq x] = \frac{1}{\pi} \int_0^{\frac{\pi}{2}} e^{-x^{\frac{p}{p-1}} \cdot V(\theta;p)} d\theta$$

Hence, we must compute the integral $\int_0^{\frac{\pi}{2}} e^{-x^{\frac{p}{p-1}} \cdot V(\theta;p)} d\theta$ within an accuracy of $O(\frac{1}{K})$, where $K = \text{poly}(n)$. Observe that for $\theta \in [0, \frac{\pi}{2}]$, $\cos \theta \geq 0$, and for $p \in [1, 2]$, $0 \leq p\theta \leq \pi$, meaning $\sin p\theta \geq 0$, and $0 \leq (p-1)\theta \leq \frac{\pi}{2}$, meaning $\cos(p-1)\theta \geq 0$. Therefore, $e^{-x^{\frac{p}{p-1}} V(\theta;p)} \leq 1$, meaning it suffices to evaluate

$$I := \int_{\frac{1}{K}}^{\frac{\pi}{2} - \frac{1}{K}} e^{-x^{\frac{p}{p-1}} \cdot V(\theta;p)} d\theta$$

which is at most $\frac{2}{K}$ away from the original integral.

As a first step, we show how to compute the exponent, $-x^{\frac{p}{p-1}} V(\theta;p)$, within an error of $\frac{1}{2\text{poly}(n)}$ in $\text{poly}(n)$ time. First, note that $|x^{\frac{p}{p-1}}| \leq \text{poly}(n)$, so we must simply compute $V(\theta;p)$ within an error of $\frac{1}{2\text{poly}(n)}$. However, observe that $\cos \theta$, $\sin p\theta$, $\cos((p-1)\theta)$, and $\cos \theta$ are bounded away from 0 on $[\frac{1}{K}, \frac{\pi}{2} - \frac{1}{K}]$, since their only singularities are at 0 and $\frac{\pi}{2}$. Furthermore, we can show that $\cos \theta, \sin p\theta \geq \Omega(\frac{1}{K})$ on this interval, by noting that $\cos \theta \geq \frac{1}{2}(\frac{\pi}{2} - x)$, $\sin p\theta \geq \min(\frac{x}{2}, \frac{\pi/2 - x}{2})$, and $\cos(p-1)\theta \geq \cos \theta$.

This means that all of the factors in $V(\theta;p)$ are bounded above by 1 and below by $\frac{1}{\text{poly}(n)}$. Hence, it suffices to compute each of those factors up to an additive $\frac{1}{2\text{poly}(n)}$ error, in order to compute them up to a multiplicative $(1 \pm \frac{1}{2\text{poly}(n)})$ error, since they are all at least $\frac{1}{\text{poly}(n)}$. This suffices to compute $V(\theta;p)$ up to a multiplicative $(1 \pm \frac{1}{2\text{poly}(n)})$ error, and hence up to an additive $\frac{1}{2\text{poly}(n)}$ error (since it is at most $\text{poly}(n)$).

Now, let us compute each of the factors in $V(\theta;p)$ by approximating them with their Taylor series. First, let P_T be the degree- T Taylor polynomial for $f(\theta) = \sin p\theta$. Then, by Taylor's theorem, for all $x \in [0, C]$ (where $C = O(1)$)

$$|P_T(x) - \sin(px)| \leq \max_{x \in [0, C]} |f^{(T)}(x)| \cdot \left| \frac{x^{T+1}}{(T+1)!} \right| \leq \frac{p^T x^T}{(T+1)!}$$

and this expression is at most $\frac{1}{2^{O(n \log n)}}$ for $T = O(n)$. We can analyze the other factors of $V(\theta;p)$ similarly.

Hence, for any $\theta \in [\frac{1}{K}, \frac{\pi}{2} - \frac{1}{K}]$, if we let $E := x^{\frac{p}{p-1}} V(\theta;p)$, then we can compute E up to an additive $\frac{1}{2\text{poly}(n)}$ error. Since the function $f(x) = e^{-x}$ is Lipschitz continuous on $(-\infty, 0]$ with Lipschitz constant 1, $|e^{-E} - e^{-E'}| \leq \frac{1}{2\text{poly}(n)}$, where E' is our estimate of E . In addition, we can compute $e^{-E'}$ up to additive $\frac{1}{2\text{poly}(n)}$ error, by using the first $O(n)$ terms of the Taylor series of e^{-x} , through similar reasoning as above: again let P_T be the degree T Taylor polynomial of e^x , centered at 0. We know that $0 \leq x^{\frac{p}{p-1}} V(\theta;p) \leq \text{poly}(n)$. Now, for $x \in [-\text{poly}(K), 0]$,

$$|e^x - P_T(x)| \leq \left(\max_{y \in [-\text{poly}(n), 0]} |e^y| \right) \left| \frac{x^{T+1}}{(T+1)!} \right| \leq \left| \frac{x^{T+1}}{(T+1)!} \right|$$

Since $x \leq \text{poly}(n)$, the right-hand side is at most $\frac{1}{2\text{poly}(n)}$ if $T = n^c$ for a sufficiently large $c > 0$.

In summary, we can evaluate the integrand $e^{-x^{\frac{p}{p-1}}V(\theta;p)}$ at any $\theta \in [\frac{1}{K}, \frac{\pi}{2} - \frac{1}{K}]$, up to an additive $\frac{1}{2\text{poly}(n)}$ error, in $\text{poly}(n)$ time. The final step is to show that we can compute $\int_{\frac{1}{K}}^{\frac{\pi}{2}-\frac{1}{K}} e^{-x^{\frac{p}{p-1}}V(\theta;p)} d\theta$ by evaluating the integrand at $\text{poly}(n)$ many points. We can do this with the trapezoidal rule, by bounding the second derivative of the integrand. Let

$$f(\theta) = e^{-x^{\frac{p}{p-1}}V(\theta;p)}$$

Then,

$$f'(\theta) = e^{-x^{\frac{p}{p-1}}V(\theta;p)} \cdot (-x^{\frac{p}{p-1}} \cdot V'(\theta;p))$$

and

$$f''(\theta) = e^{-x^{\frac{p}{p-1}}V(\theta;p)} \cdot (-x^{\frac{p}{p-1}} \cdot V'(\theta;p))^2 + e^{-x^{\frac{p}{p-1}}V(\theta;p)} \cdot (-x^{\frac{p}{p-1}} \cdot V''(\theta;p))$$

Observe that $e^{-x^{\frac{p}{p-1}}V(\theta;p)} \leq 1$, and $x^{\frac{p}{p-1}} \leq \text{poly}(n)$. Finally, note that the numerators of $V'(\theta;p)$ and $V''(\theta;p)$ will be at most 1, and the denominators will only involve $\sin p\theta$ and $\cos \theta$, which are at least $\frac{1}{\text{poly}(n)}$ — therefore, $V'(\theta;p)$ and $V''(\theta;p)$ are at most $\text{poly}(n)$. Therefore, the second derivative of the integrand is at most $\text{poly}(n)$, and it suffices to evaluate the integrand at $\text{poly}(n)$ many points to get an estimate of $\int_{\frac{1}{K}}^{\frac{\pi}{2}-\frac{1}{K}} e^{-x^{\frac{p}{p-1}}V(\theta;p)} d\theta$ with $\frac{1}{\text{poly}(n)}$ additive error. \square

Observe that if we compute the distribution of the Z_i approximately rather than exactly, then this may affect $E[\langle Z, x \rangle]$. We now show that this is not an issue:

Claim A.23. *Suppose Z_1, Z_2, \dots, Z_n are as specified in the statement of Lemma A.18, and W_1, W_2, \dots, W_n are $O(\log n)$ -wise independent random variables, such that for all $z \in \mathbb{R}$, $|\Pr[Z_i = z] - \Pr[W_i = z]| \leq \frac{1}{\text{poly}(n)}$, and W_i takes on the same values as Z_i . Then, $E[\langle W, x \rangle] = \Theta_p(1)$.*

Proof. Let Z be a real discrete random variable whose distribution is as specified in the statement of Lemma A.18, and W be a random variable so that for all $z \in \mathbb{R}$, $|\Pr[Z = z] - \Pr[W = z]| \leq \frac{1}{F}$, where $F = n^c$ for a sufficiently large constant c . Suppose Z takes on the values z_1, z_2, \dots, z_r (and W takes on these values). For convenience, define $p_{Z,i} = \Pr[Z = z_i]$ and $p_{W,i} = \Pr[W = z_i]$, and $s_{Z,i} = \sum_{j=1}^i p_{Z,i}$ and $s_{W,i} = \sum_{j=1}^i p_{W,i}$.

Let P_U be the uniform distribution on $[0, 1]$. Define the random variables Z_U and W_U as follows. Suppose U is drawn from P_U . If U is in the interval $[s_{Z,i-1}, s_{Z,i}]$ (where $s_{Z,0} = 0$), then define $Z_U = z_i$. Similarly, if U is in the interval $[s_{W,i-1}, s_{W,i}]$, then define $W_U = z_i$. Note that Z_U and Z have the same distributions, and W_U and W have the same distributions. Finally, we say U is *bad* if $U \in [s_{Z,i-1}, s_{Z,i}]$ and $U \in [s_{W,j-1}, s_{W,j}]$ for $i \neq j$. Note that $\Pr[U \text{ is bad}] \leq \frac{\text{poly}(n)}{F}$, since for each $i \in [n]$,

$$|s_{Z,i} - s_{W,i}| \leq \sum_{j=1}^i |p_{Z,i} - p_{W,i}| \leq \frac{n}{F}$$

meaning $\Pr[U \in [s_{Z,i-1}, s_{Z,i}] \cap [s_{W,j-1}, s_{W,j}]]$ is $O(\frac{n}{F})$ if $i \neq j$, and we can perform a union bound over all such pairs i, j to find that $\Pr[U \text{ is bad}] \leq O(\frac{n^3}{F})$.

Now, let U_1, U_2, \dots, U_n be $O(\log n)$ -wise independent random variables drawn from P_U , and define $Z_{U_1}, Z_{U_2}, \dots, Z_{U_n}$ and $W_{U_1}, W_{U_2}, \dots, W_{U_n}$ as described above. Let $\mathbf{Z} = (Z_{U_1}, Z_{U_2}, \dots, Z_{U_n})$ and $\mathbf{W} = (W_{U_1}, W_{U_2}, \dots, W_{U_n})$. Finally, let $x \in \mathbb{R}^n$ so that $\|x\|_p = 1$. By Lemma A.18, we know that $E[\langle \mathbf{Z}, x \rangle] = \Theta_p(1)$, and we wish to show that $E[\langle \mathbf{W}, x \rangle] = \Theta_p(1)$. Let \mathcal{E} be the event that U_i is bad for some i — observe that if none of the U_i were bad, then $\mathbf{Z} = \mathbf{W}$. However,

$$\Pr[\mathcal{E}] \leq \sum_{i=1}^n \Pr[U_i \text{ is bad}] \leq O\left(\frac{n^4}{F}\right) \leq \frac{1}{n^{10}}$$

if the constant c is chosen large enough. Hence,

$$\Theta_p(1) = E[\langle \mathbf{Z}, x \rangle] = E[\langle \mathbf{Z}, x \rangle \mid \mathcal{E}] \Pr[\mathcal{E}] + E[\langle \mathbf{Z}, x \rangle \mid \neg \mathcal{E}] \Pr[\neg \mathcal{E}]$$

and

$$E[|\langle \mathbf{W}, x \rangle|] = E[|\langle \mathbf{W}, x \rangle| \mid \mathcal{E}] \Pr[\mathcal{E}] + E[|\langle \mathbf{W}, x \rangle| \mid \neg \mathcal{E}] \Pr[\neg \mathcal{E}]$$

Note that with probability 1, $\|\mathbf{Z}\|_2, \|\mathbf{W}\|_2 \leq \text{poly}(n)$, meaning $E[|\langle \mathbf{W}, x \rangle| \mid \mathcal{E}] \Pr[\mathcal{E}] \leq \frac{1}{\text{poly}(n)}$ and $E[|\langle \mathbf{Z}, x \rangle| \mid \mathcal{E}] \Pr[\mathcal{E}] \leq \frac{1}{\text{poly}(n)}$ if F is chosen large enough. In addition, if \mathcal{E} does not hold, then $\mathbf{W} = \mathbf{Z}$. Hence,

$$\left| E[|\langle \mathbf{W}, x \rangle|] - E[|\langle \mathbf{Z}, x \rangle|] \right| \leq \frac{1}{\text{poly}(n)}$$

and $E[|\langle \mathbf{W}, x \rangle|] = \Theta_p(1)$. This proves the claim. \square

If we let $Z \in \mathbb{R}^n$ have coordinates given by the statement of Lemma A.18, then in $n^{O(\log n)}$ time, we can construct a sample space Ω for the Z_i , of size at most $n^{O(\log n)}$, such that the $O(\log n)$ -wise independence of the Z_i holds. We can then construct the embedding matrix R as described in Fact 2, using the Z_i in place of the random variables R_i in the statement of Fact 2, and letting R be the resulting matrix A given in Fact 2. \square

A.4 Hardness for ℓ_p and $\ell_{1,p}$ Low Rank Approximation with Additive Error, $p \in (1, 2)$

The goal of this section is to use the techniques of [2] to show the following stronger hardness result.

Theorem A.24 (Hardness for ℓ_p and $\ell_{1,p}$ Low Rank Approximation with Additive Error). *Suppose the SSEH and ETH hold. Then, for $p \in (1, 2)$, at least $2^{k^{\Omega(1)}}$ time is required to achieve the following guarantees:*

- Given a matrix $A \in \mathbb{R}^{n \times d}$, $k \in \mathbb{N}$, find a matrix $\hat{A} \in \mathbb{R}^{n \times d}$ of rank at most k such that

$$\|\hat{A} - A\|_p \leq O(1) \min_{A_k \text{ rank } k} \|A - A_k\|_p + \frac{1}{2^{\text{poly}(k)}} \|A\|_p$$

- Given a matrix $A \in \mathbb{R}^{n \times d}$, $k \in \mathbb{N}$, find a matrix $\hat{A} \in \mathbb{R}^{n \times d}$ of rank at most k such that

$$\|\hat{A} - A\|_{1,p} \leq O(1) \min_{A_k \text{ rank } k} \|A - A_k\|_{1,p} + \frac{1}{2^{\text{poly}(k)}} \|A\|_{1,p}$$

where for a matrix $M \in \mathbb{R}^{n \times d}$, $\|M\|_{1,p} = \sum_{j=1}^d \|M_j\|_p$.

Using the hardness of the second guarantee, we will show hardness for constrained ℓ_1 low rank approximation as well. It is worth noting that Algorithm 5 can in fact be used to achieve the first guarantee in $2^{\text{poly}(k)} + \text{poly}(nd)$ time, by letting $f = 2^{\text{poly}(k)}$ (instead of letting $f = \text{poly}(k)$ as is done when Algorithm 5 is called by Algorithm 6).

First we show the following lemma, similar to Lemma A.6, which will be used to show hardness for $\ell_{1,p}$ -norm low rank approximation:

Lemma A.25 (Equivalence of $\min_{p^* \rightarrow 1}$ and Rank— $(k-1)$ $\ell_{1,p}$ Low Rank Approximation). *Let $p \in (1, \infty)$ and let p^* be the Hölder conjugate of p . Let $A \in \mathbb{R}^{n \times d}$ with $n \geq d$ and $k = d - 1$. Then,*

$$\min_{U \in \mathbb{R}^{d \times k}, V \in \mathbb{R}^{k \times n}} \|UV - A^T\|_{1,p} = \min_{x \in \mathbb{R}^d, \|x\|_{p^*} = 1} \|Ax\|_1 =: \min_{p^* \rightarrow 1}(A)$$

Proof. The proof is the same as that of Lemma 27 of [2], but we make small modifications now that the objective is the $\ell_{1,p}$ -norm. First, assume the rank of A is d , since both sides are 0 otherwise. In addition, for a matrix $M \in \mathbb{R}^{n \times d}$, define

$$\|M\|_{p,1} = \sum_{i=1}^n \|M^i\|_p$$

In other words, $\|M\|_{p,1} = \|M^T\|_{1,p}$. Now, let $x \in \mathbb{R}^d$ such that $\|x\|_{p^*} = 1$, and take $V \in \mathbb{R}^{k \times d}$ so that the rows of V span a $(k-1)$ -dimensional subspace of \mathbb{R}^d orthogonal to x . In addition, define $U = \text{argmin}_{U \in \mathbb{R}^{n \times k}} \|UV - A\|_{p,1}$, meaning that each row U^i of U is the minimizer $\text{argmin}_{U^i} \|U^i V - A^i\|_p$. Then, we can write

$$\|U^i V - A^i\|_p = \min_{\langle y, x \rangle = 0} \|y - A^i\|_p = \min_{\langle z, x \rangle = -\langle A^i, x \rangle} \|z\|_p$$

Claim A.26. $\min_{\langle z, x \rangle = -\langle A^i, x \rangle} \|z\|_p = |\langle A^i, x \rangle|$

Proof. First, by Hölder's inequality, for all $z \in \mathbb{R}^d$,

$$|\langle x, z \rangle| \leq \|x\|_{p^*} \|z\|_p = \|z\|_p$$

since $\|x\|_p^* = 1$. Hence, if $\langle z, x \rangle = -\langle A^i, x \rangle$, then

$$\|z\|_p \geq |\langle x, z \rangle| = |\langle A^i, x \rangle|$$

On the other hand, if we choose z so that the j^{th} coordinate of z is $z_j = -\langle x, A^i \rangle (\text{sgn}(x_j) |x_j|^{\frac{p^*}{p}})$ then

$$\langle z, x \rangle = -\langle x, A^i \rangle \sum_j (\text{sgn}(x_j) |x_j|^{\frac{p^*}{p}}) \cdot x_j = -\langle x, A^i \rangle \sum_j |x_j|^{\frac{p^*+p}{p}} = -\langle x, A^i \rangle \sum_j |x_j|^{p^*} = -\langle x, A^i \rangle$$

where the third equality is because $\frac{1}{p^*} + \frac{1}{p} = 1$, meaning $p^* + p = pp^*$, and the last equality is because $\|x\|_{p^*} = 1$. In addition,

$$\|z\|_p^p = |\langle x, A^i \rangle|^p \sum_j |x_j|^{p^*} = |\langle x, A^i \rangle|^p$$

meaning $\|z\|_p = |\langle x, A^i \rangle|$. This proves the claim. \square

By this claim, $\|U^i V - A^i\|_p = |\langle A^i, x \rangle|$, meaning

$$\|UV - A\|_{p,1} = \sum_{i=1}^n \|U^i V - A^i\|_p = \sum_{i=1}^n |\langle A^i, x \rangle| = \|Ax\|_1$$

Hence, for every $x \in \mathbb{R}^d$ with $\|x\|_{p^*} = 1$, if the rows of $V \in \mathbb{R}^{k \times d}$ span the subspace orthogonal to x , then

$$\|Ax\|_1 = \|UV - A\|_{p,1} \geq \min_{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}} \|UV - A\|_{p,1}$$

Similarly, if $V \in \mathbb{R}^{k \times d}$, we can assume without loss of generality that V has rank exactly k , and choose x perpendicular to the rows of V such that $\|x\|_{p^*} = 1$. Then, as we have shown,

$$\|UV - A\|_{p,1} = \|Ax\|_1 \geq \min_{p^* \rightarrow 1} (A)$$

meaning

$$\min_{U \in \mathbb{R}^{n \times k}, V \in \mathbb{R}^{k \times d}} \|UV - A\|_{p,1} \geq \min_{p^* \rightarrow 1} (A)$$

Finally, to complete the proof, note that $\|UV - A\|_{p,1}$ is equal to $\|(UV - A)^T\|_{1,p} = \|V^T U^T - A^T\|_{1,p}$. \square

Now we show the main result of this subsection.

Theorem A.27 (Hardness for ℓ_p and $\ell_{1,p}$ Low Rank Approximation with Additive Error). *Suppose the SSEH and ETH hold. Then, for $p \in (1, 2)$, at least $2^{k^{\Omega(1)}}$ time is required to achieve the following guarantees:*

- Given a matrix $A \in \mathbb{R}^{n \times d}$, $k \in \mathbb{N}$, find a matrix $\hat{A} \in \mathbb{R}^{n \times d}$ of rank at most k such that

$$\|\hat{A} - A\|_p \leq O(1) \min_{A_k \text{ rank } k} \|A - A_k\|_p + \frac{1}{2^{\text{poly}(k)}} \|A\|_p$$

- Given a matrix $A \in \mathbb{R}^{n \times d}$, $k \in \mathbb{N}$, find a matrix $\hat{A} \in \mathbb{R}^{n \times d}$ of rank at most k such that

$$\|\hat{A} - A\|_{1,p} \leq O(1) \min_{A_k \text{ rank } k} \|A - A_k\|_{1,p} + \frac{1}{2^{\text{poly}(k)}} \|A\|_{1,p}$$

Proof. The main idea for showing hardness of the first guarantee is that while performing the reduction due to [2], we can ensure that the matrix A has entries with at most $\text{poly}(k)$ bits. As a result of this, we show that the first guarantee is in fact equivalent to achieving an $O(1)$ -approximation. To show hardness for the second guarantee, we use this observation, and also combine the reduction of [2] with Theorem A.7 above. We use Theorem A.7 to reduce the computation of $\min_{p^* \rightarrow p}(A)$ to that of $\min_{p^* \rightarrow 1}(RA)$, and this is equivalent to finding the error from the best rank— $(k-1)$ approximation to $A^T R^T$ in the $\ell_{1,p}$ -norm, by Lemma A.25. We now begin the proof.

As in the reduction of [2], let G be a regular graph that has k vertices, so that $P := P_{\geq 1/2}(G)$ is a $k \times k$ matrix. Note that we can write $P = UU^T$, where $U \in \mathbb{R}^{k \times t}$ (for some $t \in \mathbb{N}$) has orthonormal columns. The entries of U have absolute value at most 1, meaning the entries of P have absolute value at most k by the triangle inequality. In addition, the entries of P can be rounded to the nearest integer multiple of $\frac{1}{D}$, where $D = \text{poly}(k)$ is a sufficiently large power of 2, without significantly changing $\|P\|_{2 \rightarrow p^*}$. We can see this as follows. Let \hat{P} be P with its entries rounded to the nearest integer multiple of $\frac{1}{D}$. Then,

$$\begin{aligned} \left| \|P\|_{2 \rightarrow p^*} - \|\hat{P}\|_{2 \rightarrow p^*} \right| &\leq \|P - \hat{P}\|_{2 \rightarrow p^*} \\ &\leq \|P - \hat{P}\|_{2 \rightarrow 2} \\ &\leq \|P - \hat{P}\|_F \\ &\leq \frac{k}{D} \end{aligned} \tag{23}$$

Here, the first inequality is by the triangle inequality. The second is because $\|x\|_{p^*} \leq \|x\|_2$ for any $x \in \mathbb{R}^n$, and the third is because the spectral norm is at most the Frobenius norm. Finally, the last inequality is because $\|P - \hat{P}\|_\infty \leq \frac{1}{D}$, and P and \hat{P} are $k \times k$ matrices.

Now, we must simply choose a large enough D so that we can distinguish between the Yes and No cases of the Small Set Expansion Problem using a constant-factor approximation to $\|\hat{P}\|_{2 \rightarrow p^*}$, where the Yes and No cases are as specified in the proof of Theorem A.2. As mentioned in the proof of Theorem A.2, in the Yes case,

$$\|P\|_{2 \rightarrow p^*} \geq 1/(10\delta^{(p^*-2)/2p^*}) =: C_1$$

while in the No case

$$\|P\|_{2 \rightarrow p^*} \leq 2/\delta^{(p^*-2)/4p^*} =: C_2$$

(where $\delta \in (0, 1)$ can be any sufficiently small number). Hence, it suffices to choose $D \geq 50k\delta^{(p^*-2)/4p^*}$. To see why, note that in the No case,

$$\|\hat{P}\|_{2 \rightarrow p^*} \leq \|P\|_{2 \rightarrow p^*} + \frac{k}{D} \leq \frac{2}{\delta^{(p^*-2)/4p^*}} + \frac{1}{50\delta^{(p^*-2)/4p^*}} \leq \frac{3}{\delta^{(p^*-2)/4p^*}}$$

while in the Yes case,

$$\|\hat{P}\|_{2 \rightarrow p^*} \geq \|P\|_{2 \rightarrow p^*} - \frac{k}{D} \geq \frac{1}{10\delta^{(p^*-2)/2p^*}} - \frac{1}{50\delta^{(p^*-2)/4p^*}} \geq \frac{1}{20\delta^{(p^*-2)/2p^*}}$$

where the last inequality holds because $\frac{1}{\delta} > 1$ and $(p^*-2)/2p^* > (p^*-2)/4p^*$, meaning $\frac{1}{20\delta^{(p^*-2)/2p^*}} \geq \frac{1}{50\delta^{(p^*-2)/4p^*}}$. Note that the gap between the Yes and No case is still

$$\frac{1/(20\delta^{(p^*-2)/2p^*})}{3/\delta^{(p^*-2)/4p^*}} = \frac{1}{60\delta^{(p^*-2)/4p^*}}$$

and this can be made arbitrarily large as $\delta \rightarrow 0$, meaning computing $\|\hat{P}\|_{2 \rightarrow p^*}$ within an arbitrary constant factor is NP-hard. We now reduce the problem of computing $\|\hat{P}\|_{2 \rightarrow p^*}$ to that of computing the optimal rank— $(k-1)$ approximation of a well-chosen matrix where each entry has at most $\text{poly}(k)$ bits.

First, observe that each entry of \hat{P} has at most $O(\log k)$ bits in its numerator and denominator, since each entry is at most k in absolute value, and each entry has denominator D , which has $O(\log k)$ bits. Now, we show that each of the steps in the reduction of [2] preserves the bit complexity of \hat{P} , meaning that none

of the P_i mentioned in the previous subsection will have more than $\text{poly}(k)$ bits. First, we can construct P_1 so that $(1 - \varepsilon)\|\widehat{P}\|_{2 \rightarrow p^*} \leq \|P_1\|_{2 \rightarrow p^*} \leq (1 + \varepsilon)\|\widehat{P}\|_{2 \rightarrow p^*}$ by Lemma A.3 — recall that this can be done by adding $\frac{\varepsilon M}{\|I\|_{2 \rightarrow p^*}}$ to each diagonal entry of \widehat{P} , where I is the identity matrix. Since only a constant-factor approximation to $\|\widehat{P}\|_{2 \rightarrow p^*}$ is needed to decide the Small Set Expansion problem, we can let $\varepsilon = \Theta(1)$. In addition, recall that M has at most $O(\log k)$ bits in both its numerator and denominator (where M is the largest entry of \widehat{P}). Finally, $\|I\|_{2 \rightarrow p^*} = 1$ since $p^* > 2$, meaning $\|x\|_p^* \leq \|x\|_2$ for all x . Hence, all entries of P_1 have at most $O(\log k)$ bits in the numerator and denominator.

Furthermore, recall that all entries of \widehat{P} are integer multiples of $\frac{1}{D}$, meaning all entries of P_1 are integer multiples of $\frac{1}{K}$ for some $K = \text{poly}(k)$. Hence, we can compute $P_2 = P_1 P_1^T$, and all of the entries of P_2 are at most $\text{poly}(k)$ and are integer multiples of $\frac{1}{K^2}$, meaning all entries of P_2 have at most $O(\log k)$ bits in their numerators and denominators. Note that by Lemma A.4, $\|P_2\|_{p \rightarrow p^*} = \|P_1\|_{2 \rightarrow p^*}^2$, meaning in order to get a constant-factor approximation to $\|P_1\|_{2 \rightarrow p^*}$ it suffices to get a constant-factor approximation to $\|P_2\|_{p \rightarrow p^*}$.

Finally, we compute $P_3 = P_2^{-1}$ as described in the previous subsection. Each entry of P_3 is a cofactor of P_2 , divided by the determinant of P_2 . Each cofactor is the sum of at most $k!$ terms, each of which is a product of $(k - 1)$ entries of P_2 . Similarly, the determinant of P_2 is the sum of $k!$ products of k entries of P_2 . Since each entry of P_2 is an integer multiple of $\frac{1}{D^2}$, each of these products has a denominator of D^{2k} , which is at most $2^{\text{poly}(k)}$ (meaning it has $\text{poly}(k)$ bits). The numerator of each of these products has absolute value at most $k! \cdot \text{poly}(k)$, meaning it has at most $\text{poly}(k)$ bits. In summary, for each entry of P_3 , its numerator and denominator have at most $\text{poly}(k)$ bits, and our argument so far establishes the following claim:

Claim A.28. *There exists a matrix $P_3 \in \mathbb{R}^{k \times k}$, which can be deterministically computed in $\text{poly}(k)$ time, such that its entries have at most $\text{poly}(k)$ bits in their numerators and denominators, and computing a constant-factor approximation for $\|\widehat{P}\|_{2 \rightarrow p^*}$ reduces to computing a constant-factor approximation for $\min_{p^* \rightarrow p}(P_3)$.*

We now show the following claim, which we use for showing hardness both for ℓ_p low rank approximation and $\ell_{1,p}$ low rank approximation:

Claim A.29. $\min_{p^* \rightarrow p}(P_3) \geq \frac{1}{\text{poly}(k)}$

Proof. Since each entry of P has absolute value at most k , $\|P\|_{2 \rightarrow p^*} \leq \text{poly}(k)$, meaning $\|\widehat{P}\|_{2 \rightarrow p^*} \leq \text{poly}(k)$. Next, note that by Lemma A.3,

$$\|P_1\|_{2 \rightarrow p^*} \leq O(1)\|\widehat{P}\|_{2 \rightarrow p^*} \leq \text{poly}(k)$$

Moreover, by Lemma A.4,

$$\|P_2\|_{p \rightarrow p^*} = \|P_1\|_{2 \rightarrow p^*}^2 \leq \text{poly}(k)$$

and finally, by Lemma A.5,

$$\min_{p^* \rightarrow p}(P_3) = \min_{p^* \rightarrow p}(P_2^{-1}) = (\|P_2\|_{p \rightarrow p^*})^{-1} \geq \frac{1}{\text{poly}(k)}$$

□

Now we show hardness for ℓ_p low rank approximation. Let $A = P_3$ — by Lemma A.6, $\min_{p^* \rightarrow p}(A) = \min_{M_{k-1} \text{ rank } (k-1)} \|M_{k-1} - A\|_p$. To show the desired result, it suffices to show the following claim:

Claim A.30. $\frac{1}{2^{\text{poly}(k)}}\|A\|_p \leq \text{OPT}_p := \min_{M_{k-1} \text{ rank } (k-1)} \|M_{k-1} - A\|_p$

Proof. Note that by the previous claim, $\text{OPT}_p = \min_{p^* \rightarrow p}(A) \geq \frac{1}{\text{poly}(k)}$. On the other hand, $\|A\|_p \leq 2^{\text{poly}(k)}$, since the entries of A have at most $\text{poly}(k)$ bits, meaning they are at most $2^{\text{poly}(k)}$. Therefore, $\|A\|_p \leq 2^{\text{poly}(k)}$, and this proves the claim. Alternatively, it is possible to prove this claim using an argument similar to Claim 1 on page 15 of [2]. □

Hence, if we find \widehat{A} of rank at most $k-1$ such that

$$\|\widehat{A} - A\|_p \leq O(1) \min_{M_{k-1} \text{ rank } (k-1)} \|M_{k-1} - A\|_p + \frac{1}{2^{\text{poly}(k)}} \|A\|_p$$

then the claim above in fact shows that $\|\widehat{A} - A\|_p \leq O(1) \min_{M_{k-1} \text{ rank } (k-1)} \|M_{k-1} - A\|_p$ — this is enough to obtain a constant-factor approximation to $\|\widehat{P}\|_{2 \rightarrow q}$, and decide the Small Set Expansion instance. We now show a similar hardness result for low rank approximation in the $\ell_{1,p}$ norm by using Theorem A.7.

Claim A.31. *There exists a matrix $B \in \mathbb{R}^{k \times k^{O(\log k)}}$ which can be computed in $k^{O(\log k)}$ time such that the following properties hold:*

- *The entries of B have at most $\text{poly}(k)$ bits in their numerators and denominators.*
- *Computing a constant-factor approximation for $\|\widehat{P}\|_{2 \rightarrow p^*}$ reduces to computing a constant-factor approximation for $\text{OPT}_{1,p} := \min_{M_{k-1} \text{ rank } (k-1)} \|M_{k-1} - B\|_{1,p}$, and*
- *$\text{OPT}_{1,p} \geq \frac{1}{\text{poly}(k)}$*

Proof. Recall that to compute a constant-factor approximation for $\|\widehat{P}\|_{2 \rightarrow p^*}$, it suffices to find a constant-factor approximation for

$$\min_{p^* \rightarrow p} (P_3) = \min_{\|x\|_{p^*}=1} \|P_3 x\|_p$$

By Theorem A.7, there exists a matrix $R \in \mathbb{R}^{k^{O(\log k)} \times k}$ such that for all $y \in \mathbb{R}^k$,

$$\Omega(1)\|y\|_p \leq \|Ry\|_1 \leq O(1)\|y\|_p$$

and R can be constructed in $k^{O(\log k)}$ time so that the above holds with constant probability. Hence,

$$\min_{p^* \rightarrow 1} (RP_3) = \min_{\|x\|_{p^*}=1} \|RP_3 x\|_1$$

is within a constant factor of $\min_{p^* \rightarrow p} (P_3)$, and it suffices to compute a constant factor approximation to $\min_{p^* \rightarrow 1} (RP_3)$, which by Lemma A.25 is equal to $\min_{M_{k-1} \text{ rank } (k-1)} \|M_{k-1} - P_3^T R^T\|_{1,p}$.

Finally, to complete the proof of the claim, we define B to be $P_3^T R^T$, but with each of its entries rounded to the nearest integer multiple of $\frac{1}{2^{\text{poly}(k)}}$. Let

$$\text{OPT}_{1,p}^0 = \min_{M_{k-1} \text{ rank } (k-1)} \|M_{k-1} - P_3^T R^T\|_{1,p}$$

with M^0 being the corresponding minimizer, and

$$\text{OPT}_{1,p} = \min_{M_{k-1} \text{ rank } (k-1)} \|M_{k-1} - B\|_{1,p}$$

with M being the corresponding minimizer. Then, note that

$$\|M - B\|_{1,p} \leq \|M^0 - B\|_{1,p} \leq \|M^0 - P_3^T R^T\|_{1,p} + \frac{1}{2^{\text{poly}(k)}}$$

where the last inequality is because

$$\|B - P_3^T R^T\|_{1,p} = \sum_{j=1}^{k^{O(\log k)}} \|B_j - (P_3^T R^T)_j\|_p \leq k^{O(\log k)} \cdot \frac{1}{2^{\text{poly}(k)}} \leq \frac{1}{2^{\text{poly}(k)}}$$

Similarly,

$$\|M^0 - P_3^T R^T\|_{1,p} \leq \|M - P_3^T R^T\|_{1,p} \leq \|M - B\|_{1,p} + \frac{1}{2^{\text{poly}(k)}}$$

meaning that $\left| \|M - B\|_{1,p} - \|M^0 - P_3^T R^T\|_{1,p} \right| \leq \frac{1}{2^{\text{poly}(k)}}$. Since $\|M^0 - P_3^T R^T\|_{1,p} \geq \frac{1}{\text{poly}(k)}$ by Claim A.29, this means that $\|M - B\|_{1,p}$ is within a constant factor of $\|M^0 - P_3^T R^T\|_{1,p}$.

In summary, to obtain a constant-factor approximation to $\|\widehat{P}\|_{2 \rightarrow p^*}$ with constant probability, it suffices to obtain a constant-factor approximation to $\min_{M_{k-1} \text{ rank } (k-1)} \|M_{k-1} - B\|_{1,p}$. Moreover, the entries of B have at most $\text{poly}(k)$ bits in their numerators and denominators. We can see this as follows. First, recall that P_3 has entries with at most $\text{poly}(k)$ bits in their numerators and denominators, meaning its entries are at most $2^{\text{poly}(k)}$. In addition, the entries of R are at most $O(1)$, since for any standard basis vector e_i , $\|Re_i\|_1 = \Theta(1)$. Therefore, the entries of $P_3^T R^T$ are at most $2^{\text{poly}(k)}$. Finally, after we round the entries of $P_3^T R^T$ to the nearest integer multiple of $\frac{1}{2^{\text{poly}(k)}}$ to obtain B , the entries of B are at most $2^{\text{poly}(k)}$ and are multiples of $\frac{1}{2^{\text{poly}(k)}}$, meaning they have at most $\text{poly}(k)$ bits. Finally, $\text{OPT}_{1,p} = \Theta(1) \cdot \text{OPT}_{1,p}^0 \geq \frac{1}{\text{poly}(k)}$. \square

Since the entries of B have at most $\text{poly}(k)$ bits in their numerators and denominators, $\|B\|_{1,p} \leq 2^{\text{poly}(k)}$. In addition, $\text{OPT}_{1,p} \geq \frac{1}{\text{poly}(k)}$, meaning $\|B\|_{1,p} \leq 2^{\text{poly}(k)} \text{OPT}_{1,p}$, and if we find \widehat{B} of rank at most $k-1$ such that

$$\|\widehat{B} - B\|_{1,p} \leq O(1) \min_{M_{k-1} \text{ rank } (k-1)} \|M_{k-1} - B\|_{1,p} + \frac{1}{2^{\text{poly}(k)}} \|B\|_{1,p}$$

then $\|\widehat{B} - B\|_{1,p}$ is a constant factor approximation to $\text{OPT}_{1,p}$.

In summary, we have shown that achieving either of the guarantees in Theorem A.24 is sufficient to decide Problem 2. Now, suppose we assume the Exponential-Time Hypothesis (ETH) in addition to the Small Set Expansion Hypothesis — ETH is the assumption that any algorithm for solving 3-SAT takes at least $2^{\Omega(n)}$ time, where n is the number of variables in the 3-SAT instance. Since SSEH is the assumption that the Small Set Expansion Problem is NP-hard, we can assume there is a reduction from 3-SAT to Small Set Expansion that takes an instance with n variables to a graph with at most $m = n^{\frac{1}{c}}$ vertices for some $c > 0$ (since the reduction takes polynomial time, it creates a graph with size at most $\text{poly}(n)$). Hence, assuming SSEH and ETH, there is no algorithm which can decide any Small Set Expansion instance in $2^{o(m^c)}$ time, since this would imply that any 3-SAT instance could be decided in $2^{o(n)}$ time. Therefore, an algorithm achieving either of the guarantees in Theorem A.24 requires 2^{k^c} time. \square

A.5 Hardness for Constrained ℓ_1 Low Rank Approximation with Additive Error

We recall the constrained ℓ_1 low rank approximation problem:

Problem 3 (Constrained ℓ_1 Low Rank Approximation). *Given a matrix $A \in \mathbb{R}^{n \times d}$ and a subspace $V \subset \mathbb{R}^n$, find a matrix \widehat{A} of rank at most k minimizing $\|\widehat{A} - A\|_1$, such that the columns of \widehat{A} are in V .*

Theorem A.32 (Hardness for Constrained ℓ_1 Low Rank Approximation with Additive Error). *Suppose the SSEH and ETH hold. Then, for $k \in \mathbb{N}$, at least $2^{k^{\Omega(1)}}$ time is required to achieve the following guarantee with constant probability: given a matrix $A \in \mathbb{R}^{n \times d}$, $Y \in \mathbb{R}^{n \times t}$ for some $t \leq n$, and $k \in \mathbb{N}$, find a matrix \widehat{A} of rank at most k such that*

$$\|\widehat{A} - A\|_1 \leq O(1) \min_{A_k} \|A_k - A\|_1 + \frac{1}{2^{\text{poly}(k)}} \|A\|_1$$

and the columns of \widehat{A} are contained in the column span of Y — here, the minimum on the right-hand side is taken over A_k having rank k , such that the columns of A_k are contained in the column span of Y .

Proof. We reduce $\ell_{1,p}$ low rank approximation (more specifically, the problem of obtaining the guarantee in Theorem A.24) to the above guarantee for constrained ℓ_1 low rank approximation. For convenience, let $p = \frac{3}{2}$. Recall that by Claim A.31 and the discussion following its proof, given any instance of Problem 2 on a graph of size k , a matrix $B \in \mathbb{R}^{k \times k^{O(\log k)}}$ can be computed in $k^{O(\log k)}$ time such that, in order to decide the instance of Problem 2, it suffices to compute a matrix $\widehat{B} \in \mathbb{R}^{k \times k^{O(\log k)}}$ such that

$$\|\widehat{B} - B\|_{1,p} \leq O(1) \min_{M_{k-1} \text{ rank } (k-1)} \|M_{k-1} - B\|_{1,p} + \frac{1}{2^{\text{poly}(k)}} \|B\|_{1,p} \quad (24)$$

Now, by Theorem A.7, there exists a matrix $R \in \mathbb{R}^{k^{O(\log k)} \times k}$, which can be computed in $k^{O(\log k)}$ time, such that for all $x \in \mathbb{R}^k$,

$$\Omega(1)\|x\|_p \leq \|Rx\|_1 \leq O(1)\|x\|_p$$

Now, suppose we could find $M \in \mathbb{R}^{k \times k^{O(\log k)}}$, with rank at most $k - 1$, such that

$$\|RM - RB\|_1 \leq O(1) \min_{M_{k-1}} \min_{\text{rank}(k-1)} \|RM_{k-1} - RB\|_1 + \frac{1}{2^{\text{poly}(k)}} \|RB\|_1 \quad (25)$$

Note that this is a solution to the constrained ℓ_1 low rank approximation problem, that achieves the guarantee described in the statement of this theorem (replacing A with RB , the subspace $Y \in \mathbb{R}^{n \times t}$ with R , and the target rank k with $k - 1$). However, for any matrix $M \in \mathbb{R}^{k \times k^{O(\log k)}}$,

$$\|RM\|_1 = \sum_{i=1}^d \|RM_i\|_1 = \sum_{i=1}^d \Theta(1)\|M_i\|_p = \Theta(1)\|M\|_{1,p}$$

where the second equality is because for any $x \in \mathbb{R}^n$, $\|Rx\|_1$ is within a constant factor of $\|x\|_p$. Hence, Equation 25 implies that

$$\|M - B\|_{1,p} \leq O(1) \min_{M_{k-1}} \min_{\text{rank}(k-1)} \|M_{k-1} - B\|_{1,p} + \frac{1}{2^{\text{poly}(k)}} \|B\|_{1,p}$$

i.e. M satisfies the guarantee of Equation 24. Therefore, achieving the guarantee in Equation 25 is enough to decide the original instance of Problem 2, and therefore requires $2^{k^{\Omega(1)}}$ time to achieve. \square

B ℓ_p Column Subset Selection Lower Bound

We use a construction similar to [48], to show that column subset selection algorithms cannot give an approximation factor better than $O(k^{\frac{1}{p}-\frac{1}{2}})$ for ℓ_p -low rank approximation, for $1 < p < 2$, if $k \cdot \text{polylog}(k)$ columns are chosen. The hard distribution for ℓ_1 column subset selection from [48], which we use here for ℓ_p column subset selection, is as follows: $A \in \mathbb{R}^{(k+n) \times n}$ is a random matrix where each of the first k rows has i.i.d. $N(0, 1)$ entries, and the remaining $n \times n$ submatrix is the identity matrix. We show that if $n = k \cdot \text{poly}(\log k)$, then any subset of $r \leq k \cdot \text{poly}(\log k)$ columns cannot give an approximation factor better than $O(k^{\frac{1}{p}-\frac{1}{2}})$, for sufficiently large k , unless $n - r = o(n)$. The proofs in this section follow those of the analogous lemmas in Section G.3 of [48] with slight modifications — for each of our lemmas, we note the corresponding lemma in [48]. First, we state some definitions.

Definition B.1. Suppose $\beta, \gamma \in \mathbb{R}$ and $p \in (1, 2)$. Then, we define

$$Y_{\beta, \gamma, p} = \left\{ y \in \mathbb{R}^n \mid \|y\|_p \leq O(k^\gamma), |y_i| \leq \frac{1}{k^\beta} \leq 1 \forall i \in [n] \right\}$$

Definition B.2. (Similar to Definition G.19 of [48]) Let $V \in \mathbb{R}^{n \times r}$ have orthonormal columns, and $A \in \mathbb{R}^{k \times r}$ be a random matrix, for which each entry is i.i.d. $N(0, 1)$. Then, we define the event $\widehat{\mathcal{E}}(A, V, \beta, \gamma, p)$ as follows: $\forall y \in Y_{\beta, \gamma, p}$, $AV^T y$ has at most $O(\frac{k}{\log k})$ coordinates with absolute value at least $\Omega(\frac{1}{\log k})$, and $\|A\|_2 = O(\sqrt{r})$.

Lemma B.3 (Gaussian Matrices Have Small Operator Norm - Due to [46], as stated in [48]). Let $A \in \mathbb{R}^{r \times k}$ be a random matrix for which each entry is i.i.d. $N(0, 1)$. With probability at least $1 - e^{-\Theta(r)}$, the maximum singular value of $\|A\|_2$ is at most $O(\sqrt{r})$.

Lemma B.4 (Flat Vectors Stay Flat - Similar to Lemma G.20 of [48]). Suppose $k \geq 1$, $c_2 \geq c_1 \geq 1$, and $k \leq r = O(k(\log k)^{c_1})$, $r \leq n = O(k(\log k)^{c_2})$. Let $V \in \mathbb{R}^{n \times r}$ have orthonormal columns, and let $A \in \mathbb{R}^{k \times r}$ be a random matrix with each entry being i.i.d. $N(0, 1)$. Finally, suppose $\beta, \gamma \in \mathbb{R}$ such that $0 < \gamma \leq p\gamma < \beta(2 - p) \leq \beta$. Then,

$$\Pr\left[\widehat{\mathcal{E}}(A, V, \beta, \gamma, p)\right] \geq 1 - 2^{-\Theta(k)}$$

Proof. The proof is the same as that of Lemma G.20 in [48] — it uses a net argument and a union bound. Rather than constructing a net for all of $Y_{\beta, \gamma, p}$, the coordinates of a point $y \in Y_{\beta, \gamma, p}$ are first divided between points y^j with disjoint supports, such that y^j has coordinates between $\frac{1}{2^{j+1}}$ and $\frac{1}{2^j}$. Then, for each j , an ε -net \mathcal{N}_j (for a suitable ε) is constructed for all points of the same form as y^j , with coordinates between $\frac{1}{2^{j+1}}$ and $\frac{1}{2^j}$. It is shown that with high probability, all points in \mathcal{N}_j have at most $O(k/\log^2 k)$ coordinates which are at least $\Omega(\frac{1}{\log^2 k})$ — if this occurs, then this implies that $\widehat{\mathcal{E}}(A, V, \beta, \gamma, p)$ occurs. Moreover, it is only necessary to consider $O(\log k)$ distinct values of j — for sufficiently large values of $j \geq j^*$, $AV^T y^j$ makes a very small contribution to the coordinates of $AV^T y$, and in fact, $\sum_{j \geq j^*} AV^T y^j$ makes a very small contribution to the coordinates of $AV^T y$ (the proof of this last statement is where Lemma B.3 is used).

We now begin the proof. Let $Y_{\beta, \gamma, p}$ be as in Definition B.1, and take $y \in Y_{\beta, \gamma, p}$. As mentioned before, write $y = \sum_{j=j_0}^{\infty} y^j$, for $j \geq j_0$, y^j has coordinates in the interval $[\frac{1}{2^{j+1}}, \frac{1}{2^j}]$. Note that the y^j have disjoint supports, and we can take j_0 so that $\frac{1}{2^{j_0}} = \frac{1}{k^\beta}$ by the definition of $Y_{\beta, \gamma}$. Let s_j be the support size of y^j . Then,

$$\left(\frac{s_j}{2^{p(j+1)}}\right)^{\frac{1}{p}} \leq \|y^j\|_p \leq \|y\|_p \leq O(k^\gamma)$$

meaning

$$s_j \leq O(2^{p(j+1)} k^{p\gamma})$$

In addition, it will later be useful to bound from above $\|y^j\|_2$:

$$\|y^j\|_2^2 \leq \frac{1}{2^{2j}} \cdot s_j \leq O(k^{p\gamma} \cdot 2^{p(j+1)-2j})$$

We use this to construct an ε -net \mathcal{N}_j for points of the same form as y^j , and bound its size:

$$\mathcal{N}_j := \{x \in \mathbb{R}^n \mid \exists x' \in \mathbb{Z}^n, x = \varepsilon x', \|x\|_p \leq O(k^\gamma), \forall i \in [n], \text{ either } \frac{1}{2^{j+1}} \leq |x_i| < \frac{1}{2^j} \text{ or } x_i = 0\}$$

In other words, it is the integer grid, but contracted by a factor of ε , and with coordinates and ℓ_p -norm in the same range as x . This is an ε -net in the ℓ_∞ norm, meaning it is an $\varepsilon\sqrt{n}$ -net in the ℓ_2 norm. We will choose $\varepsilon = O(\frac{1}{nrk^3}) = O(\frac{1}{k^{c_1+c_2+3}})$, meaning \mathcal{N}_j is an $O(\frac{1}{k^{c_1+c_2/2+3}})$ -net for the ℓ_2 norm. Now, we bound the size of \mathcal{N}_j — for each coordinate x_i , since it is between $\frac{1}{2^{j+1}}$ and $\frac{1}{2^j}$ and the net has a granularity of ε , the number of choices for x_i is $1 + \frac{1}{2^j\varepsilon} \leq 1 + \frac{2}{\varepsilon}$ since $\frac{1}{2^j} \leq \frac{1}{k^\beta} \leq 2$. The number of coordinates of y^j is at most $n = O(k(\log k)^{c_2})$, meaning

$$\begin{aligned} |\mathcal{N}_j| &\leq \left(1 + \frac{2}{\varepsilon}\right)^{O(k(\log k)^{c_2})} \\ &\leq 2^{O(k(\log k)^{c_2} \log \frac{1}{\varepsilon})} \\ &\leq 2^{O(k(\log k)^{c_2+1})} \end{aligned} \tag{26}$$

In preparation for the union bound over points in \mathcal{N}_j , we consider the event $\mathcal{E}(y^j)$ that for a single y^j , $AV^T y^j$ has at least $O(\frac{k}{\log^2 k})$ coordinates which are at least $O(\frac{1}{\log^2 k})$ in absolute value. Notice that a single coordinate of $AV^T y^j$ is $N(0, \|V^T y^j\|_2^2)$ since the rows of A are i.i.d. $N(0, 1)$. Hence, the probability q that a particular coordinate of $AV^T y^j$ is greater than $O(\frac{1}{\log^2 k})$ is, by properties of the Gaussian distribution, at most

$$\exp\left(-\frac{1}{\|V^T y^j\|_2^2} \cdot \frac{1}{\log^4 k}\right)$$

since the probability that a Gaussian random variable $N(0, \sigma^2)$ has absolute value greater than t is at most $e^{O(-\frac{t^2}{\sigma^2})}$. Hence, letting $i_0 = O(\frac{k}{\log^2 k})$, the probability that $AV^T y^j$ has at least $O(\frac{k}{\log^2 k})$ coordinates greater than $O(\frac{1}{\log^2 k})$ in absolute value is at most

$$\begin{aligned} \sum_{i=i_0}^k q^i (1-q)^i \binom{k}{i} &\leq k 2^k q^i \\ &\leq k 2^k \exp\left(-\frac{i}{\|V^T y^j\|_2^2} \cdot \frac{1}{\log^4 k}\right) \\ &\leq k 2^k \exp\left(-\Theta\left(\frac{k}{\|V^T y^j\|_2^2 \log^6 k}\right)\right) \\ &\leq k 2^k \exp\left(-\Theta\left(\frac{k^{1-p\gamma}}{2^{p-(2-p)j} \cdot \log^6 k}\right)\right) \\ &\leq k 2^k \exp\left(-\Theta\left(\frac{k^{1+\beta(2-p)-p\gamma}}{\log^6 k}\right)\right) \exp\left(-\Theta\left(\frac{k^{1-p\gamma}}{2^{p-(2-p)j} \cdot \log^6 k}\right)\right) \\ &\leq \exp\left(-\Theta\left(\frac{k^{1-p\gamma}}{2^{p-(2-p)j} \cdot \log^6 k}\right)\right) \end{aligned} \tag{27}$$

The first inequality is because $(1-q)^i \leq 1$, $\binom{k}{i} \leq 2^k$, and there are at most k summands. The third inequality is because $i \geq i_0 = \Theta(\frac{k}{\log^2 k})$. The fourth inequality is because $\|V^T y^j\|_2^2 \leq \|V^T\|_2^2 \|y^j\|_2^2 \leq O(k^{p\gamma} \cdot 2^{p(j+1)-2j})$ because V has orthonormal columns and because of the upper bound on $\|y^j\|_2$ mentioned above. The fifth inequality is because $2^{(2-p)j} \geq 2^{(2-p)j_0} = \frac{1}{k^{\beta(2-p)}}$. Finally, the sixth inequality is because $\beta(2-p) > p\gamma$, meaning $k^{\beta(2-p)-p\gamma}/(\log^6 k) = \omega(1)$, and $k 2^k$ is absorbed in the remaining exponential — this is where we use the hypothesis that $p\gamma < (2-p)\beta$.

Now we show that we can ignore the effect of y^j for $j \geq \Theta(\log k)$ — we then perform a union bound over the remaining $O(\log k)$ nets \mathcal{N}_j to show that $\mathcal{E}(y^j)$ holds for all y^j in those \mathcal{N}_j . In particular, as in [48], let $j_1 = \lceil 100(c_1 + c_2 + 1) \log k \rceil$. Then, for $j \geq j_1$, $\|\sum_{j=j_1}^\infty y^j\|_2 \leq \Theta(2^{-j_1} \sqrt{n}) \leq \frac{1}{k^{100(1+c_1)}}$ since all coordinates

of $\sum_{j=j_1}^{\infty} y^j$ are at most $\frac{1}{2^j}$ and the y^j have disjoint supports. Hence,

$$\begin{aligned}
\|AV^T y^j\|_{\infty} &\leq \|AV^T y^j\|_2 \\
&\leq \|A\|_2 \|V^T\|_2 \|y^j\|_2 \\
&\leq O(\sqrt{r}) \cdot \frac{1}{k^{100(1+c_1)}} \\
&\leq O\left(\frac{\sqrt{k}(\log k)^{c_1/2}}{k^{100(1+c_1)}}\right) \\
&\leq \frac{1}{k^{100}}
\end{aligned} \tag{28}$$

where the third inequality is because A has top singular value at most $O(\sqrt{r})$ and V has orthonormal columns. Hence, if $y^{---j_1} := \sum_{j=j_0}^{j_1} y^j$, then y has at least $O(\frac{k}{\log k})$ coordinates with absolute value at least $O(\frac{1}{\log k})$, if and only if this is the case for y^{---j_1} .

Finally, by a union bound over $O(\log k)$ nets $\mathcal{N}_{j_0}, \dots, \mathcal{N}_{j_1}$, we can show that for all j between j_0 and j_1 , with high probability, $AV^T y^j$ has at most $O(k/\log^2 k)$ coordinates with absolute value greater than or equal to $O(1/\log^2 k)$ — therefore, outside of a set of $O(\log k) \cdot O(k/\log^2 k) = O(k/\log k)$ coordinates, each coordinate of $AV^T y$ is at most $O(\log k) \cdot O(1/\log^2 k) = O(1/\log k)$.

First, the probability that there exists $y^j \in \mathcal{N}_j$ for some j between j_0 and j_1 , such that $\mathcal{E}(\widehat{y}^j)$ occurs, is

$$\begin{aligned}
P\left[\exists y^j \in \bigcup_{j=j_0}^{j_1} \mathcal{N}_j, \mathcal{E}(y^j) \text{ happens}\right] &\leq \sum_{j=j_0}^{j_1} |\mathcal{N}_j| \cdot \exp\left(-\Theta\left(\frac{k^{1-p\gamma}}{2^{p-(2-p)j} \cdot \log^6 k}\right)\right) \\
&\leq \sum_{j=j_0}^{j_1} 2^{O(k(\log k)^{c_2+1})} \cdot \exp\left(-\Theta\left(\frac{k^{1-p\gamma} \cdot 2^{(2-p)j}}{\log^6 k}\right)\right) \\
&\leq O(\log k) \cdot 2^{O(k(\log k)^{c_2+1})} \cdot \exp\left(-\Theta\left(\frac{k^{1-p\gamma} \cdot k^{\beta(2-p)}}{\log^6 k}\right)\right) \\
&\leq O(\log k) \cdot 2^{-\Theta(k)} \\
&\leq 2^{-\Theta(k)}
\end{aligned} \tag{29}$$

The first inequality above is by a union bound and our upper bound on the probability of $\mathcal{E}(\widehat{y}^j)$ for an individual \widehat{y}^j in \mathcal{N}_j . The third inequality is because $2^j \geq k^\beta$. The fourth inequality is because $\beta(2-p) - p\gamma > 0$, meaning that $k^{\text{poly}(\log k)} = o(k^{1-p\gamma+\beta(2-p)}/\log^6 k)$.

Finally, given a vector y^j with each entry having absolute value in $[\frac{1}{2^{j+1}}, \frac{1}{2^j})$ such that $\|y^j\|_p \leq k^\gamma$ and each entry is at most $\frac{1}{k^\beta}$, there exists $\widehat{y}^j \in \mathcal{N}_j$ such that $\|y^j - \widehat{y}^j\|_2 \leq O(\frac{1}{k^{c_1+c_2/2+3}})$. Hence, assuming that $\mathcal{E}(\widehat{y}^j)$ does not occur for any \widehat{y}^j in \mathcal{N}_j for j between j_0 and j_1 ,

$$\begin{aligned}
\|AV^T y^j - AV^T \widehat{y}^j\|_{\infty} &\leq \|AV^T y^j - AV^T \widehat{y}^j\|_2 \\
&\leq O(\sqrt{r}) \|y^j - \widehat{y}^j\|_2 \\
&\leq O\left(\frac{\sqrt{r}}{k^{c_1+c_2/2+3}}\right) \\
&\leq O\left(\frac{1}{k^3}\right)
\end{aligned} \tag{30}$$

since A has operator norm $O(\sqrt{r})$ and V^T has operator norm ≤ 1 . Hence, with probability $1 - 2^{-\Theta(k)}$, for all $y \in Y_{\beta, \gamma, p}$, and all j between j_0 and j_1 , $AV^T y^j$ has at most $O(\frac{k}{\log^2 k})$ coordinates with absolute value more than $O(\frac{1}{\log^2 k})$.

In summary, as discussed above, this implies that with probability $1 - 2^{-\Theta(k)}$, for all $y \in Y_{\beta, \gamma, p}$, $AV^T y$ has at most $O(k/\log k)$ coordinates with absolute value more than $O(1/\log k)$. This completes the proof. \square

Lemma B.5 (Paying with Either Regression Cost or Norm - Similar to Lemma G.22 of [48]). *For any $t, k \geq 1$, and any constants $c_2 \geq c_1 \geq 1$, let $k \leq r = O(k(\log k)^{c_1})$, and $r \leq n = O(k(\log k)^{c_2})$. Let $V \in \mathbb{R}^{n \times r}$ be a matrix with orthonormal columns. For an arbitrary constant $\alpha \in (0, 0.5)$, if $A \in \mathbb{R}^{k \times r}$ is such that $\widehat{\mathcal{E}}(A, V, \frac{1+\alpha}{2}, \frac{1}{p} - \frac{1}{2} - \alpha, p)$ holds, then with probability $1 - 2^{-\Theta(tk)}$, there are at least $\lceil \frac{t}{10} \rceil$ such $j \in [t]$ that $\forall x \in \mathbb{R}^r$, either $\|Ax - v_j\|_p \geq \Omega(k^{\frac{1}{p} - \frac{1}{2} - \alpha})$ or $\|Vx\|_p \geq \Omega(k^{\frac{1}{p} - \frac{1}{2} - \alpha})$.*

Remark B.6. *The meaning of this lemma is that when a subset A_S of columns of the hard instance $A \in \mathbb{R}^{(k+n) \times n}$ is chosen, then there exists a large enough subset of the remaining columns such that for each of those columns v_j , the regression coefficient vector y_j used to fit those v_j (using A_S as the left factor) either leads to a large regression error on the top k rows, or has a large norm, in which case it leads to a large regression cost on the bottom n rows, which are simply the identity matrix.*

Proof. The proof is similar to that of the analogous Lemma G.22 of [48]. Here we use a simplified version of that argument, that was given to us by Peilin Zhong.

We show the desired statement using a net argument. For convenience, let $\gamma = \frac{1}{p} - \frac{1}{2} - \alpha$, and let $\beta = \frac{1+\alpha}{2}$. Then, note that $p\gamma = 1 - \frac{p}{2} - p\alpha$, which is less than $\beta(2-p) = 1 - \frac{p}{2} + \alpha \cdot \frac{2-p}{2}$. Hence, by Lemma B.4, $\widehat{\mathcal{E}}(A, V, \frac{1+\alpha}{2}, \frac{1}{p} - \frac{1}{2} - \alpha, p)$ holds with probability $1 - 2^{-\Theta(k)}$. For a particular $x \in \mathbb{R}^r$, if we let $y := Vx$, meaning $x = V^T y$, then the statement becomes equivalent to showing that there are at least $\lceil \frac{t}{10} \rceil$ indices $j \in [t]$ such that $\forall y \in \mathbb{R}^n$, either $\|y\|_p \geq \Omega(k^\gamma)$, or $\|AV^T y - v_j\|_p \geq \Omega(k^\gamma)$. Throughout this proof, we consider a *fixed* matrix $A \in \mathbb{R}^{k \times r}$ such that $\widehat{\mathcal{E}}(A, V, \beta, \gamma, p)$ holds — all randomness will be over the coordinates of the v_j .

Let $B_{\gamma, p}$ denote the ℓ_p -ball of radius $O(k^\gamma)$ with center 0. Then, we must show that for all $y \in B_{\gamma, p}$, $\|AV^T y - v\|_p \geq \Omega(k^\gamma)$ with high probability. For each such y , we can write it as $y_0 + y_1$, where all the coordinates of y_0 are either 0 or greater than $\frac{1}{k^\beta}$ in absolute value, all the coordinates of y_1 are at most $\frac{1}{k^\beta}$, and the supports of y_0 and y_1 are disjoint. Recall that this means $y_1 \in Y_{\beta, \gamma, p}$ as defined above. Since $\widehat{\mathcal{E}}(A, V, \beta, \gamma, p)$ holds, for each $y \in B_{\gamma, p}$, its corresponding y_1 is such that $AV^T y_0$ has at most $O(k/\log k)$ coordinates that are at least $O(1/\log k)$ in absolute value.

We first build a net for all possible y_0 . For convenience, let $T_{\beta, \gamma, p}$ be the set of all possible y_0 , i.e. $T_{\beta, \gamma, p} := \{x \in \mathbb{R}^n \mid \forall i \in [n], \text{ either } |x_i| \geq \frac{1}{k^\beta} \text{ or } x_i = 0\}$. Then, for any $\varepsilon > 0$, the following is an ε -net in the ℓ_∞ norm for $T_{\beta, \gamma, p}$ (and hence an $\varepsilon\sqrt{n}$ -net in the ℓ_2 norm):

$$\mathcal{N} := \{x \in \mathbb{R}^n \mid \exists x' \in \mathbb{Z}^n, x = \varepsilon x', \|x\|_p \leq O(k^\gamma), \forall i \in [n], \text{ either } |x_i| \geq \frac{1}{k^\beta} \text{ or } x_i = 0\}$$

Let us calculate the size of $|\mathcal{N}|$. Note that if s is the support size of y_0 , then $\frac{1}{k^{\beta p}} s \leq \|y_0\|_p^p \leq \|y\|_p^p \leq k^{\gamma p}$ meaning $s \leq k^{\beta p + \gamma p}$. Moreover, for each coordinate of $x \in \mathcal{N}$, the number of choices is at most $O(\frac{k^\gamma}{\varepsilon})$ (since each coordinate is at most k^γ and has a granularity of ε). Hence, if we let $\varepsilon = O(1/rnk^3)$ as in [48], then

$$|\mathcal{N}| \leq \binom{n}{s} \cdot (k^\gamma/\varepsilon)^{O(s)} = 2^{O(s) \log k} = 2^{k^{\beta p + \gamma p} \log k}$$

since $n = O(k(\log k)^{c_2})$.

We will perform a union bound over all $y_0 \in \mathcal{N}$, to show that the failure event $\mathcal{E}_1(y_0)$ does not occur, where $\mathcal{E}_1(y_0)$ is the event that for some $y_1 \in Y_{\beta, \gamma, p}$ with disjoint support from y_0 , if $y := y_0 + y_1$, then $\|AV^T y - v\|_p \leq O(k^\gamma)$. First, let $\mathcal{E}_2(y_0)$ be the event that $AV^T y_0 - v$ has at most $O(k/\log k)$ coordinates that are at least $O(1/\log k)$ in absolute value.

Claim B.7 (Similar to Claim G.23 of [48]). *Assume $\widehat{\mathcal{E}}(A, V, \beta, \gamma, p)$ holds. Then, for all $y_0 \in T_{\beta, \gamma, p}$, $\mathcal{E}_1(y_0)$ implies $\mathcal{E}_2(y_0)$.*

Proof. The proof is similar to that of Claim G.23 of [48]. Suppose $\mathcal{E}_1(y_0)$ occurs, meaning there is $y_1 \in Y_{\beta, \gamma, p}$ such that if $y := y_0 + y_1$, then $\|AV^T y - v\|_p \leq O(k^\gamma)$. Let s be the number of coordinates of $AV^T y - v$ that are greater than $\frac{1}{\log k}$ in absolute value — then, $\frac{s}{\log^p k} \leq \|AV^T y - v\|_p^p \leq k^{p\gamma}$, meaning $s \leq k^{p\gamma} \log^p k$, and this is $o(k/\log k)$, since $p\gamma < p \cdot (\frac{1}{p} - \frac{1}{2}) < 1$.

In addition, observe that

$$AV^T y_0 - v = (AV^T y - v) - AV^T y_1$$

since $y = y_0 + y_1$. Since $\widehat{\mathcal{E}}(A, V, \beta, \gamma, p)$ occurs, $AV^T y_1$ has at most $O(k/\log k)$ coordinates that are greater than $\frac{1}{\log k}$ in absolute value, by the previous lemma. Hence, since $AV^T y - v$ has at most $o(k/\log k)$ coordinates that are greater than $\Omega(1/\log k)$ in absolute value, $AV^T y_0 - v$ also has at most $O(k/\log k)$ coordinates that are greater than $\Omega(1/\log k)$ in absolute value, meaning $\mathcal{E}_2(y_0)$ holds. \square

Now, we show that the probability of $\mathcal{E}_2(y_0)$ is small, where the randomness is over the entries of a vector $v \in \mathbb{R}^k$ with i.i.d. $N(0, 1)$ entries.

Claim B.8. *Let $z \in \mathbb{R}^k$, and let v be a k -dimensional vector with i.i.d. $N(0, 1)$ entries. Then, with probability $1 - 2^{-\Theta(k)}$, there exist at most $O(k/\log k)$ coordinates $i \in [k]$ such that $|z_i - v_i| \leq O(1/\log k)$.*

Proof. Let $i \in [k]$. Then,

$$\begin{aligned} \Pr[|z_i - v_i| \leq O(1/\log k)] &= \int_{v_i - O(1/\log k)}^{v_i + O(1/\log k)} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \\ &\leq O(1/\log k) \end{aligned} \quad (31)$$

Hence, if Z_i is 1 if $|z_i - v_i| \leq O(1/\log k)$ and 0 otherwise, and $Z = \sum Z_i$ is the number of coordinates on which the difference between z and v is at most $O(1/\log k)$, then $E[Z] \leq O(k/\log k)$. Since the Z_i are independent, by a Chernoff bound, with probability $1 - 2^{-\Theta(k)}$, $Z \leq O(k/\log k)$. This proves the claim. \square

By applying the claim above with $z = AV^T y_0$ for any particular $y_0 \in \mathbb{R}^n$, the probability that $AV^T y_0 - v$ has at least $\frac{k}{10}$ coordinates on which the difference is at most $O(1/\log k)$ (meaning $AV^T y_0 - v$ has less than $\frac{9k}{10}$ coordinates on which the difference is greater than $O(1/\log k)$) is at most $2^{-\Theta(k)}$ — since this implies $\mathcal{E}_2(y_0)$, the probability of $\mathcal{E}_2(y_0)$ is also at most $2^{-\Theta(k)}$. Now, we union bound over the net \mathcal{N} to show that this occurs for *all* $y_0 \in \mathcal{N}$ with high probability:

$$\begin{aligned} \Pr[\exists y_0 \in \mathcal{N}, \mathcal{E}_2(y_0)] &\leq |\mathcal{N}| \cdot 2^{-\Theta(k)} \\ &\leq 2^{k(\beta+\gamma)p \log k} 2^{-\Theta(k)} \\ &\leq 2^{-\Theta(k)} \end{aligned} \quad (32)$$

Here we use the hypothesis that $\beta + \gamma < \frac{1}{p}$, meaning $k^{(\beta+\gamma)p} = o(k)$.

The above argument shows that for all $y_0 \in \mathcal{N}$, with probability at least $1 - 2^{-\Theta(k)}$, the event \mathcal{F} holds that there does not exist $y_1 \in Y_{\beta, \gamma, p}$ such that $\|AV^T y - v\|_p \leq O(k^\gamma)$ where $y = y_0 + y_1$. Assume that \mathcal{F} holds. Now, suppose $y_0 \in T_{\beta, \gamma, p}$, and there exists $y_1 \in Y_{\beta, \gamma, p}$ such that, if $y = y_0 + y_1$, then $\|AV^T y - v\|_p \leq O(k^\gamma)$. Let $\widehat{y}_0 \in \mathcal{N}$ such that $\|\widehat{y}_0 - y_0\|_2 \leq \varepsilon\sqrt{n} = O(1/r\sqrt{n}k^3)$. Then, letting $\widehat{y} := \widehat{y}_0 + y_1$, we obtain

$$\|AV^T \widehat{y} - v\|_p \leq \|AV^T y - v\|_p + \|AV^T (\widehat{y} - y)\|_p \leq O(k^\gamma) + k^{\frac{1}{p} - \frac{1}{2}} \|AV^T (\widehat{y}_0 - y_0)\|_2$$

Here, the first inequality is by the triangle inequality, and the second is because $\widehat{y}_0 \in \mathcal{N}$, meaning $\|AV^T \widehat{y}_0 - v\|_p \leq O(k^\gamma)$. Finally, because of our choice of \widehat{y}_0 as the net vector closest to y_0 , and because $k^{\frac{1}{p} - \frac{1}{2}} \leq \sqrt{k}$, we obtain $\|AV^T \widehat{y} - v\|_p \leq O(k^\gamma) + \sqrt{k}\sqrt{r} \cdot O\left(\frac{1}{r\sqrt{n}k^3}\right) = O(k^\gamma)$ since $\|A\|_2 \leq O(\sqrt{r})$, and $\|V\|_2 \leq 1$.

Therefore, if for all $y_0 \in \mathcal{N}$, there does not exist a corresponding y_1 for which $y := y_0 + y_1$ satisfies $\|AV^T y - v\|_p$, then the same holds for all $y_0 \in T_{\beta, \gamma, p}$. This means that, with probability $1 - 2^{-\Theta(k)}$ over the entries of v , for all $y_0 \in T_{\beta, \gamma, p}$, there does not exist $y_1 \in Y_{\beta, \gamma, p}$ with disjoint support from y_0 such that $\|AV^T y - v\|_p \leq O(k^\gamma)$, meaning for all $y \in B_{\gamma, p}$, $\|AV^T y - v\|_p \geq \Omega(k^\gamma)$. We have shown that for any $j \in [t]$, the event $\mathcal{E}_3(v_j)$ holds with probability $1 - 2^{-\Theta(k)}$, where $\mathcal{E}_3(v_j)$ is the event that with probability $1 - 2^{-\Theta(k)}$, for all $y \in \mathbb{R}^n$, either $\|y\|_p \geq \Omega(k^\gamma)$ or $\|AV^T y - v_j\|_p \geq \Omega(k^\gamma)$.

Note that the v_j are independent for different j . Hence, we can show that $\mathcal{E}_3(v_j)$ holds simultaneously for $\Omega(t)$ indices $j \in [t]$ with the desired probability, as follows. Let $Z_j = 1$ if $\mathcal{E}_3(v_j)$ does not hold. Then, the probability that at least $\frac{9t}{10}$ of the Z_j are equal to 1 is at most

$$\sum_{d=\frac{9t}{10}}^t \binom{t}{d} (2^{-\Theta(k)})^d \leq O(t) \cdot 2^t \cdot 2^{-\Theta(tk)} = 2^{-\Theta(tk)}$$

meaning that with probability $1 - 2^{-\Theta(tk)}$, $\mathcal{E}_3(v_j)$ holds for at least $\frac{t}{10}$ indices $j \in [t]$. \square

Now we prove the main result of this appendix:

Theorem B.9 (Lower Bound for ℓ_p Low Rank Approximation through Column Subset Selection - Similar to Theorem G.28 of [48]). *For a sufficiently large $k \in \mathbb{N}$, and any constant $c \geq 1$, let $n = O(k(\log k)^c)$ and let $M \in \mathbb{R}^{(k+n) \times n}$ be a random matrix such that the top $k \times n$ submatrix has i.i.d. $N(0, 1)$ entries, and the bottom $n \times n$ submatrix is the identity matrix. Then, with probability $1 - 2^{-\Theta(k)}$, for any subset $S \subset [n]$ with $|S| \leq \frac{n}{2} =: r$,*

$$\min_{X \in \mathbb{R}^{r \times n}} \|M_S X - M\|_p \geq \Omega(k^{\frac{1}{p} - \frac{1}{2} - \alpha}) \min_{M_k \text{ rank } k} \|M_k - M\|_p$$

where $\alpha \in (0, \frac{1}{p} - \frac{1}{2})$ can be arbitrary.

Proof. Throughout the proof, fix $\alpha \in (0, \frac{1}{p} - \frac{1}{2})$, $\gamma = \frac{1}{p} - \frac{1}{2} - \alpha$, and $\beta = \frac{1+\alpha}{2}$. In addition, let $A \in \mathbb{R}^{k \times n}$ be the submatrix of M consisting of its top k rows. The proof of this theorem follows that of Theorem G.28 of [48]: we apply Lemma B.5 with A_S in the place of A , to show that each subset S of size at most r incurs a large cost, then perform a union bound over all such S . For convenience, we consider the p^{th} power of the ℓ_p norm throughout the proof.

First, notice that $\min_{M_k \text{ rank } k} \|M_k - M\|_p^p \leq n$, since we could take M_k to be the $(k+n) \times n$ matrix whose first k rows are the same as those of M , and whose last n rows are 0. Hence, it suffices to show that for any subset $S \subset [n]$ of size at most r , $\min_{X \in \mathbb{R}^{r \times n}} \|M_S X - M\|_p^p \geq k^{p\gamma} n$. First, fix a subset $S \subset [n]$ with $|S| \leq r$. Observe that the cost of fitting a single column M_l of M using M_S is

$$\text{cost}(S, l) := \min_{x_l \in \mathbb{R}^r} \left(\|A_S x_l - A_l\|_p^p + \|x_l\|_p^p - 1 \right)$$

(where the notation $\text{cost}(S, l)$ is from the proof of Theorem G.28 in [48]). This is because the first k entries of M_l are given by A_l , and one of the last n entries of M_l is 1 and the others are 0. We can apply Lemma B.5 now, as follows. Suppose $\widehat{\mathcal{E}}(A_S, I_r, \beta, \gamma, p)$ occurs. Then, since the columns of A are independent, and A_l is independent from A_S for $l \notin S$, by Lemma B.5, with probability $1 - 2^{-\Theta(nk)}$, for at least $\Omega(n)$ of the indices $l \in [n] \setminus S$, either $\|A_S x_l - A_l\|_p^p \geq \Omega(k^{\gamma p})$ or $\|x_l\|_p^p \geq \Omega(k^{\gamma p})$ (since while applying that lemma, we can take $n = r$ and the matrix $V \in \mathbb{R}^{n \times r}$ to be I_r , which has orthonormal columns — note that we are taking $t = n$ in that lemma). Hence,

$$\min_{X \in \mathbb{R}^{r \times n}} \|M_S X - M\|_p^p = \sum_{l \in [n] \setminus S} \text{cost}(S, l) \geq \Omega(n \cdot k^{\gamma p}) \geq \Omega(k^{\gamma p} \cdot \text{OPT})$$

Now, instead of conditioning on $\widehat{\mathcal{E}}(A_S, I_r, \beta, \gamma, p)$ simultaneously for all S , we can simply condition on $\widehat{\mathcal{E}}(A, I_n, \beta, \gamma, p)$:

Claim B.10 (Similar to Claim G.29 of [48]). $\widehat{\mathcal{E}}(A, I_n, \beta, \gamma, p)$ implies $\widehat{\mathcal{E}}(A_S, I_r, \beta, \gamma, p)$.

Proof. In the proof of this claim, we use $Y_{\beta, \gamma, p, t}$ to denote the instance of $Y_{\beta, \gamma, p}$ in \mathbb{R}^t . Suppose $\widehat{\mathcal{E}}(A, I_n, \beta, \gamma, p)$ occurs. Then, for all $y \in Y_{\beta, \gamma, p, n} \subset \mathbb{R}^n$, Ay has at most $O(k/\log k)$ coordinates with absolute value at least $\Omega(1/\log k)$. In particular, for any subset $S \subset [n]$, let $y_S \in \mathbb{R}^n$ be a point whose support is contained in S , and let \widetilde{y}_S be y_S , with all coordinates not indexed by S removed. Observe that $A_S I_r \widetilde{y}_S = A I_n y_S$, and since $\widehat{\mathcal{E}}(A, I_n, \beta, \gamma, p)$ holds, $A I_n y_S$ has at most $O(k/\log k)$ coordinates with absolute value at least $\Omega(1/\log k)$. This implies that for any $y \in Y_{\beta, \gamma, p, r}$, $A_S I_r y$ has at most $O(k/\log k)$ coordinates with absolute value at least $\Omega(1/\log k)$. Moreover, since A_S is a subset of columns of A , the operator norm of A_S is at most $\|A\|_2 \leq \sqrt{n} = O(\sqrt{r})$, since $r = \frac{n}{2}$. \square

Hence, we can perform a union bound over all subsets $S \subset [n]$ with $|S| \leq r$, while conditioning on $\widehat{\mathcal{E}}(A, I_n, \beta, \gamma, p)$. For convenience, let $\widetilde{\mathcal{E}} = \widehat{\mathcal{E}}(A, I_n, \beta, \gamma, p)$. Then,

$$\begin{aligned}
& \Pr \left[\exists S \subset [n], |S| \leq r \text{ s.t. } \min_{X \in \mathbb{R}^{r \times n}} \|M_S X - M\|_p^p < O(k^{\gamma p} \cdot OPT) \right] \\
&= \Pr \left[\exists S \subset [n], |S| \leq r \text{ s.t. } \min_{X \in \mathbb{R}^{r \times n}} \|M_S X - M\|_p^p < O(k^{\gamma p} \cdot OPT) \mid \widetilde{\mathcal{E}} \right] \Pr[\widetilde{\mathcal{E}}] \\
&+ \Pr \left[\exists S \subset [n], |S| \leq r \text{ s.t. } \min_{X \in \mathbb{R}^{r \times n}} \|M_S X - M\|_p^p < O(k^{\gamma p} \cdot OPT) \mid \neg \widetilde{\mathcal{E}} \right] \Pr[\neg \widetilde{\mathcal{E}}] \\
&\leq \Pr \left[\exists S \subset [n], |S| \leq r \text{ s.t. } \min_{X \in \mathbb{R}^{r \times n}} \|M_S X - M\|_p^p < O(k^{\gamma p} \cdot OPT) \mid \widetilde{\mathcal{E}} \right] + \Pr[\neg \widetilde{\mathcal{E}}] \\
&\leq \sum_{S \subset [n], |S| \leq r} \Pr \left[\min_{X \in \mathbb{R}^{r \times n}} \|M_S X - M\|_p^p < O(k^{\gamma p} \cdot OPT) \mid \widetilde{\mathcal{E}} \right] + 2^{-\Theta(k)} \\
&\leq \sum_{S \subset [n], |S| \leq r} 2^{-\Theta(nk)} + 2^{-\Theta(k)} \\
&= \binom{n+1}{r+1} 2^{-\Theta(nk)} + 2^{-\Theta(k)} \\
&\leq 2^{O(r \log n)} 2^{-\Theta(nk)} + 2^{-\Theta(k)} \\
&= 2^{O(n \log k)} 2^{-\Theta(nk)} + 2^{-\Theta(k)} \\
&= 2^{-\Theta(k)}
\end{aligned} \tag{33}$$

Here, the first inequality is because probabilities are at most 1. The second is because $\widetilde{\mathcal{E}}$ occurs with probability at least $1 - 2^{-\Theta(k)}$, and by a union bound over $S \subset [n]$ with $|S| \leq r$. The third is because of our observation above, that $\widetilde{\mathcal{E}}$ implies $\widehat{\mathcal{E}}(A_S, I_r, \beta, \gamma, p)$, and if A_S is such that $\widehat{\mathcal{E}}(A_S, I_r, \beta, \gamma, p)$ holds, then the probability that $\min_{X \in \mathbb{R}^{r \times n}} \|M_S X - M\|_p^p < O(k^{\gamma p} \cdot OPT)$ is at most $2^{-\Theta(nk)}$. After that, the second equality is because the number of subsets of $[n]$ of size at most r is $\binom{n+1}{r+1}$.

In summary, the probability that there is a subset achieving error less than $O(k^{\gamma p} \cdot OPT)$ is at most $2^{-\Theta(k)}$. \square

C poly(k)-Approximation Algorithms with Bicriteria Rank Independent of n and d

C.1 A poly(k)-Approximation with poly(k) Bicriteria Rank

We can remove the $O(\log d)$ -factor in the bicriteria rank of Algorithm 4, at the cost of an increase in the approximation factor to $O(k^{\frac{2}{p}-1}\text{poly}(\log k))$. This works as follows: note that Algorithm 4 selects columns in $O(\log d)$ blocks, each having size r (which is $O(k \log k)$ if $p = 1$ and $O(k \log k \log \log k)$ if $p \in (1, 2)$) — out of these $O(\log d)$ blocks, we show that there exists a subset of blocks of size r (hence giving a column subset of size r^2) which spans an $O(k^{\frac{2}{p}-1}\text{poly}(\log k))$ -approximation to A . An advantage of this algorithm is that both the rank *and* the approximation factor are polynomial in k , and do not depend on n and d .

Algorithm 7 Obtaining a poly(k)-approximation of rank $O(r^2)$, where $r = O(k \log k)$ if $p = 1$ and $O(k \log k \log \log k)$ if $p \in (1, 2)$. This algorithm simply finds the left factor, and the right factor can be obtained using linear programming.

Require: $A \in \mathbb{R}^{n \times d}$, $k \in \mathbb{N}$, $p \in [1, 2)$

Ensure: $U \in \mathbb{R}^{n \times O(r^2)}$

procedure POLYKERRORANDRANKAPPROXIMATION(A, k, p)

$S_1, S_2, \dots, S_b \leftarrow \text{RANDOMCOLUMNSSUBSETSELECTION}(A, k, p)$ (Note that the output S is being discarded, since all we need are the blocks. Also note that the number of blocks b is $O(\log d)$.)

$U \leftarrow$ The $n \times O(r^2)$ zero matrix

MINERROR $\leftarrow 0$

if $b < \frac{C}{\varepsilon} r$ **then**

// Here C is a sufficiently large absolute constant. In this case RANDOMCOLUMNSSUBSETSELECTION

// already gives $O(k^2 \log^2 k / \varepsilon)$ columns.

$S \leftarrow \cup_{i=1}^b S_i$

$U \leftarrow A_S$

else

for $I \subset [b], |I| = r$ **do**

$T \leftarrow \cup_{i \in I} S_i$

$U_{temp} \leftarrow A_T$

$V_{temp} \leftarrow \text{argmin}_{V \in \mathbb{R}^{r^2 \times d}} \|U_{temp} V - A\|_p$

if $\|U_{temp} V_{temp} - A\|_p \leq \text{MINERROR}$ **then**

MINERROR $\leftarrow \|U_{temp} V_{temp} - A\|_p$

$U \leftarrow U_{temp}$

end if

end for

end if

end procedure

Theorem C.1 (Column Subset Selection - Removing Dependence on $\log d$). *Let $A \in \mathbb{R}^{n \times d}$, $p \in [1, 2)$, and $k \in \mathbb{N}$. In addition, let $r = O(k \log k)$ if $p = 1$ and $r = O(k \log k \log \log k)$ otherwise. Then, with constant probability, Algorithm 7 returns $U \in \mathbb{R}^{n \times O(r^2/\varepsilon)}$ and $V \in \mathbb{R}^{O(r^2/\varepsilon) \times d}$ such that*

$$\|UV - A\|_p \leq O(r^{\frac{2}{p}-1} (\log k)^{\frac{1}{p}}) \min_{A_k \text{ rank } k} \|A_k - A\|_p$$

with running time $\text{nnz}(A) + d^{1+\varepsilon} \text{poly}(k \log(d))$, aside from the time required for Algorithm 4.

This running time can be achieved if the multiple-response ℓ_p -regression is performed using a single sampling/rescaling matrix based on the ℓ_p -Lewis weights of A_T , where T is the subset of columns returned by Algorithm 4.

Proof. Let $T \subset [d]$ be the subset of columns returned by Algorithm 4. We will show that with constant probability, T contains a subset of size r^2 which provides a left factor for A giving a poly(k)-approximation.

Note that A_T is a left factor which already gives an $O(r^{\frac{1}{p}-\frac{1}{2}}(\log k)^{\frac{1}{p}})$ -approximation for A . By the proof of Theorem 2.8, to obtain this approximation factor, it suffices to fit each column A_j of A using at most r columns of A_T — i.e. there is $B_j \in \mathbb{R}^{|T|}$ having at most r nonzero entries, such that we can use $A_T B_j$ as our approximation to A_j . This is because in the proof of Theorem 2.8, to bound the cost on each discarded column A_j , we only consider the ℓ_p -regression error obtained by fitting A_j using the block of r columns that were sampled on the round where A_j is discarded.

Let $B \in \mathbb{R}^{|T| \times d}$ such that the j^{th} column of B is B_j . Then, as a restatement of Theorem 2.8, $\|A_T B - A\|_p \leq O(r^{\frac{1}{p}-\frac{1}{2}}(\log k)^{\frac{1}{p}}) \min_{A_k \text{ rank } k} \|A - A_k\|_p$.

Let $M = A_T B$ — we will find a $\text{poly}(k)$ -approximation for M and show that it is also a $\text{poly}(k)$ -approximation for A by the triangle inequality. By Theorem 2.4, there exists a subset of columns of M , of size r , giving rise to a left factor $U_0 \in \mathbb{R}^{n \times r}$ such that

$$\min_{V \in \mathbb{R}^{r \times d}} \|U_0 V - M\|_p \leq O(r^{\frac{1}{p}-\frac{1}{2}}) \min_{M_k \text{ rank } k} \|M_k - M\|_p$$

Let V_0 be the optimal right factor for U_0 . Then, letting M_k and A_k be the optimal rank- k approximations in the ℓ_p -norm for M and A respectively,

$$\begin{aligned} \|U_0 V_0 - A\|_p &\leq \|U_0 V_0 - M\|_p + \|M - A\|_p \\ &\leq O(r^{\frac{1}{p}-\frac{1}{2}}) \|M_k - M\|_p + \|M - A\|_p \\ &\leq O(r^{\frac{1}{p}-\frac{1}{2}}) \|A_k - M\|_p + \|M - A\|_p \\ &\leq O(r^{\frac{1}{p}-\frac{1}{2}}) (\|A_k - A\|_p + \|A - M\|_p) + \|M - A\|_p \\ &\leq O(r^{\frac{1}{p}-\frac{1}{2}}) \|A - M\|_p + O(r^{\frac{1}{p}-\frac{1}{2}}) \|A - A_k\|_p \\ &\leq O(r^{\frac{1}{p}-\frac{1}{2}}) \cdot O(r^{\frac{1}{p}-\frac{1}{2}}(\log k)^{\frac{1}{p}}) \|A - A_k\|_p \\ &= O(r^{\frac{2}{p}-1}(\log k)^{\frac{1}{p}}) \|A - A_k\|_p \end{aligned} \tag{34}$$

meaning U_0 is a left factor giving an $O(r^{\frac{2}{p}-1}(\log k)^{\frac{1}{p}})$ -approximation for A .

Finally, for each column v of U_0 , v is a column of M , meaning it can be written as a linear combination of at most r columns of A_T — these are the r columns in the block which was sampled in the round when v was approximately covered and discarded. Hence, A_T has a subset of columns of size r^2 , which is the union of r blocks out of the $O(\log d)$ blocks which were sampled over the course of Algorithm 4, such that this subset spans an $O(r^{\frac{2}{p}-1}(\log k)^{\frac{1}{p}})$ approximation for A . Since Algorithm 7 checks all subsets of this size, it gives an $O(r^{\frac{2}{p}-1}(\log k)^{\frac{1}{p}})$ -approximation for A .

Let us now analyze the runtime of this algorithm, aside from the time required to perform RANDOMCOLUMNSUBSETSELECTION. If the number of blocks $b < \frac{C}{\varepsilon} r$ (C is a sufficiently large absolute constant mentioned in Algorithm 7 and will be chosen appropriately later), then RANDOMCOLUMNSUBSETSELECTION already returns a subset of $O(r^2/\varepsilon)$ columns.

Otherwise, $r \leq \frac{\varepsilon b}{C}$, and Algorithm 7 checks all subsets of $[b]$ of size r . In this case, the number of subsets checked is

$$\binom{b}{r} \leq \binom{b}{\frac{\varepsilon b}{C}} \leq 2^{H_2(\frac{\varepsilon}{C})b} = d^{O(1) \cdot H_2(\frac{\varepsilon}{C})} \leq d^\varepsilon$$

where the first inequality holds as long as $r \leq \frac{\varepsilon b}{C} \leq \frac{b}{2}$, and in the second inequality H_2 is the binary entropy function. The first equality is because $b = O(\log d)$, meaning $2^b = d^{O(1)}$. Finally, the fourth inequality is because we can choose C to be sufficiently large.

Hence, the number of subsets checked is at most d^ε if C is chosen to be sufficiently large. For each subset, we perform multiple-response ℓ_p -regression. We can speed up this ℓ_p -regression as follows. Let $S \in \mathbb{R}^{\widehat{O}(k \log d) \times n}$ be a sampling and rescaling matrix generated according to the Lewis weights of A_T (recall that T is the subset of columns returned by Algorithm 4). Recall that by Lemma 2.3, with probability $1 - O(1)$, for all $V \in \mathbb{R}^{|T| \times d}$,

$$\|S A_T V - S A\|_p^p \geq \Omega(1) \|A_T V - A\|_p^p - O(1) \|A_T V^* - A\|_p^p \tag{35}$$

where $V^* = \operatorname{argmin}_{V'} \|A_T V' - A\|_p$.

Now, out of all the subsets tried by Algorithm 7, let U_{temp}^* be the optimal subset of blocks, and let U_{temp} be an arbitrary subset of blocks that is tried. We make the following definitions: $V_{temp}^* = \operatorname{argmin}_V \|U_{temp}^* V - A\|_p$, $V_{temp,S}^* = \operatorname{argmin}_V \|SU_{temp}^* V - SA\|_p$, and $V_{temp,S} = \operatorname{argmin}_V \|SU_{temp} V - SA\|_p$. Then, since U_{temp} and U_{temp}^* are both given by subsets of columns of A , we can apply Equation 35. Applying it to U_{temp} gives us

$$\|SU_{temp} V_{temp,S} - SA\|_p^p \geq \Omega(1) \|U_{temp} V_{temp,S} - A\|_p^p - O(1) \|A_T V^* - A\|_p^p \quad (36)$$

Now, let U_{temp}^{**} be the subset of blocks tried by Algorithm 7 which minimizes the sketched error, i.e. it is the subset U_{temp} minimizing $\|SU_{temp} V_{temp,S} - SA\|_p$. Define $V_{temp,S}^{**} = \operatorname{argmin}_V \|SU_{temp}^{**} V - SA\|_p$. Then,

$$\begin{aligned} \|SU_{temp}^* V_{temp}^* - SA\|_p^p &\geq \|SU_{temp}^* V_{temp,S}^* - SA\|_p^p \\ &\geq \|SU_{temp}^{**} V_{temp,S}^{**} - SA\|_p^p \\ &\geq \Omega(1) \|U_{temp}^{**} V_{temp,S}^{**} - A\|_p^p - O(1) \|A_T V^* - A\|_p^p \end{aligned} \quad (37)$$

where the first inequality is because $V_{temp,S}^*$ is the minimizer of the sketched error for the left factor U_{temp}^* , the second is because U_{temp}^{**} is the subset of blocks minimizing the sketched error, and the third is by Equation 36.

Finally, by Lemma 2.2, with probability $1 - O(1)$,

$$\|SU_{temp}^* V_{temp}^* - SA\|_p^p \leq O(1) \|U_{temp}^* V_{temp}^* - A\|_p^p$$

with probability $1 - O(1)$.

Putting this together, if U_{temp}^{**} is the subset of blocks minimizing the sketched ℓ_p regression error, and $V_{temp,S}^{**}$ is its corresponding right factor (that minimizes the sketched ℓ_p regression error), then with probability $1 - O(1)$,

$$\begin{aligned} \|U_{temp}^{**} V_{temp,S}^{**} - A\|_p^p &\leq O(1) \|U_{temp}^* V_{temp}^* - A\|_p^p + O(1) \|A_T V^* - A\|_p^p \\ &\leq \left(O(r^{\frac{2}{p}-1} (\log k)^{\frac{1}{p}}) \right)^p \min_{A_k \text{ rank } k} \|A - A_k\|_p^p \end{aligned} \quad (38)$$

where the second inequality is because U_{temp}^* is an $O(r^{\frac{2}{p}-1} (\log k)^{\frac{1}{p}})$ -approximate column subset as shown above, while A_T provides an $O(r^{\frac{1}{p}-\frac{1}{2}} (\log k)^{\frac{1}{p}})$ -approximation by our analysis of Algorithm 4.

Hence, by taking p^{th} powers, we find that performing multiple-response ℓ_p regression for each of the subsets of blocks, while reusing a single sampling and rescaling matrix S generated using the ℓ_p Lewis weights of A_T , only worsens the approximation guarantees by an $O(1)$ factor.

Now we finish analyzing the running time when Algorithm 7 is implemented using ℓ_p Lewis weights. Computing the ℓ_p Lewis weights of A_T , and generating the sampling matrix S , takes at most $\operatorname{nnz}(A) + \operatorname{poly}(k)$ time, by Lemma 2.1, and performing the multiplication SA takes time at most $\operatorname{nnz}(A)$. Now, for each of the d^ε subsets we try, we perform multiple-response ℓ_p regression on the $\tilde{O}(k \log d) \times d$ matrix SA , fitting the columns of SA using a $\tilde{O}(k \log d) \times r^2$ matrix SU_{temp} . Each of the d ℓ_p -regression steps takes time $\operatorname{poly}(k \log d)$. Therefore, each multiple-response ℓ_p regression step takes $d \operatorname{poly}(k \log d)$ time, and overall, trying all subsets takes $d^{1+\varepsilon} \operatorname{poly}(k \log d)$ time. This completes the proof. \square

Remark C.2. Note that in [11], it was shown that every $A \in \mathbb{R}^{n \times d}$ has a subset of k columns spanning an $O(k)$ -approximation to the optimum. Using this in our analysis, instead of Theorem 2.4 allows us to get an $O(r^{\frac{1}{p}-\frac{1}{2}} \cdot k \cdot (\log k)^{\frac{1}{p}})$ -approximation instead (in particular, for $p = 1$, this is an $\tilde{O}(k^{\frac{3}{2}})$ -approximation) while giving us a rank $O(\operatorname{rk})$ -solution, in particular removing one $\log(k)$ factor from the rank in the $p = 1$ case.

Remark C.3. In the analysis of our algorithm with rank at most k , we will take ε from Theorem C.1 to just be a sufficiently small constant to optimize the running time, but to also keep the rank at most $O(k^2 \log^2 k)$.

Remark C.4. Note that this kind of analysis is not applicable to any bi-criteria column subset selection algorithm, and it uses special properties of Algorithm 4 — given an arbitrary bi-criteria column subset

selection algorithm, we cannot use this analysis to reduce the rank. The key property is that for each column that is discarded by Algorithm 4, it can be fit using at most r columns belonging to the left factor that Algorithm 4 returns. Note that Algorithm 1 of [51] also has this property, and hence it can also be used to obtain a $\text{poly}(k)$ -approximation algorithm with $\text{poly}(k)$ bicriteria rank, albeit with a somewhat larger approximation factor.

C.2 Reducing the Rank to At Most k

We can combine our algorithm from the previous subsection with Theorem C.19 of [48]:

Theorem C.5 (Theorem C.19 of [48] — Reducing the Rank). *Given a rank r matrix $B = U_B V_B \in \mathbb{R}^{n \times d}$, where $U_B \in \mathbb{R}^{n \times r}$ and $V_B \in \mathbb{R}^{r \times d}$, $p \in [1, 2)$, and $k \in [r]$ there exists an algorithm which takes $(n+d) \cdot \text{poly}(k)$ time to output two matrices $U \in \mathbb{R}^{n \times k}$ and $V \in \mathbb{R}^{k \times d}$ such that*

$$\|UV - B\|_p \leq \text{poly}(r) \min_{B_k \text{ rank } k} \|B_k - B\|_p$$

Proof. In the case $p = 1$, this is Theorem C.19 of [48]. Theorem C.19 of [48] also generalizes to $p \in (1, 2)$. This is because all of the lemmas used in the proof of this theorem can be extended to the ℓ_p norm case. Claim C.18 of [48], which is that multiplying any vector in a low-dimensional subspace by a dense Cauchy matrix leads to at most a constant-factor decrease in its ℓ_1 -norm, can be extended to the ℓ_p -norm by considering matrices with p -stable entries instead. Similarly, Lemma D.23 of [48] (multiplying a fixed matrix by a dense Cauchy matrix does not increase its ℓ_1 -norm significantly) can be generalized to the ℓ_p -norm case by using matrices with p -stable entries. Claims B.10 and B.11 of [48] already apply to the ℓ_p -norm case, and Lemmas D.8 and D.10 of [48] are generalized to the ℓ_p -norm case in Lemma E.8 and E.10 respectively. \square

Denote the algorithm described in Theorem C.19 of [48] by $\text{REMOVEBICRITERIARANK}(U_B, V_B, k, p)$, where $B = U_B V_B$ is a matrix of rank at most r and k is the desired rank. Then, we obtain a $\text{poly}(k)$ -approximation algorithm, shown in Algorithm 8:

Algorithm 8 Reducing the bicriteria rank of Algorithm 7.

Require: $A \in \mathbb{R}^{n \times d}$, $k \in \mathbb{N}$, $p \in [1, 2)$

Ensure: $W \in \mathbb{R}^{n \times k}$, $Z \in \mathbb{R}^{k \times d}$

procedure POLYKERRORNOTBICRITERIAAPPROXIMATION(A, k, p)

$U \leftarrow \text{POLYKERRORANDRANKAPPROXIMATION}(A, k, p)$

$V \leftarrow \text{argmin}_{V'} \|UV' - A\|_p$

$W, Z \leftarrow \text{REMOVEBICRITERIARANK}(U, V, k, p)$

end procedure

Theorem C.6 (Analysis of Algorithm 8). *Let $A \in \mathbb{R}^{n \times d}$ and $k \in \mathbb{N}$. Then, with constant probability, Algorithm 8 returns $W \in \mathbb{R}^{n \times k}$ and $Z \in \mathbb{R}^{k \times d}$ such that*

$$\|WZ - A\|_p \leq \text{poly}(k) \min_{A_k \text{ rank } k} \|A_k - A\|_p$$

The running time of Algorithm 8 is $\text{nnz}(A) + (d^{1.1} + n)\text{poly}(k \log d)$, aside from the running time of Algorithm 4.

Proof. For a matrix $M \in \mathbb{R}^{n \times d}$ and $k \in \mathbb{N}$, let $\text{OPT}_{M,k} := \min_{M_k \text{ rank } k} \|M_k - M\|_p$. Now, by the guarantees on Algorithm 7, $\|UV - A\|_p \leq \text{poly}(k)\text{OPT}_{A,k}$. Let $r = O(k^2 \log^2 k)$ be the bicriteria rank of Algorithm 7. Then, by Theorem C.5, $\|WZ - UV\|_p \leq \text{poly}(r)\text{OPT}_{UV,k} = \text{poly}(k)\text{OPT}_{UV,k}$. Therefore,

$$\|WZ - A\|_p \leq \|WZ - UV\|_p + \|UV - A\|_p \leq \text{poly}(k)\text{OPT}_{UV,k} + \text{poly}(k)\text{OPT}_{A,k}$$

and it suffices to show that $\text{OPT}_{UV,k} \leq \text{poly}(k)\text{OPT}_{A,k}$. However, if $A_k = \text{argmin}_{A' \text{ rank } k} \|A' - A\|_p$, then

$$\text{OPT}_{UV,k} \leq \|A_k - UV\|_p \leq \|A_k - A\|_p + \|UV - A\|_p \leq \text{OPT}_{A,k} + \text{poly}(k)\text{OPT}_{A,k} \leq \text{poly}(k)\text{OPT}_{A,k}$$

Here, the first inequality is because A_k has rank at most k . The second inequality is by the triangle inequality, and the third is because of the approximation guarantees on Algorithm 7. In summary, $OPT_{UV,k} \leq \text{poly}(k)OPT_{A,k}$, meaning $\|WZ - A\|_p \leq \text{poly}(k)OPT_{A,k}$.

Now, let us analyze the running time of Algorithm 8. Taking $\varepsilon = 0.1$ in Algorithm 7, we find that the running time needed to obtain U is $\text{nnz}(A) + d^{1.1}\text{poly}(k \log d)$, aside from the time taken to run Algorithm 4. We can also find V in $\text{nnz}(A) + d \cdot \text{poly}(k)$ time. To see this, note that we can multiply U and A by a sampling matrix S corresponding to the Lewis weights of U — by Lemma 2.1, S can be generated in $\text{nnz}(U) + \text{poly}(k)$ time, and computing SU and SA takes at most $\text{nnz}(A)$ time. By Lemma 2.3, performing ℓ_p regression using SU and SA instead of U and A leads to an $O(1)$ -approximate right factor V . Each ℓ_p regression takes $\text{poly}(k)$ time, meaning finding V takes $d \cdot \text{poly}(k)$ time. Finally, by Theorem C.19 of [48], finding W and Z takes $(n+d) \cdot \text{poly}(k)$ time. The overall running time is therefore $\text{nnz}(A) + (d^{1.1} + n)\text{poly}(k)$, aside from the time taken to run Algorithm 4. \square

Remark C.7. Note that the bicriteria rank r has to be $\text{poly}(k)$, rather than $\text{poly}(k, \log d)$, before applying the algorithm described in Theorem C.19 of [48], since the approximation factor is $\text{poly}(r)$.

Remark C.8. It is also possible to combine our Algorithm 7 with Algorithm 4 of [11] to reduce the bicriteria rank to at most k , and obtain a $\text{poly}(k)$ -approximate matrix, although this requires $2^{O(k \log k)} + \text{poly}(nd)$ running time. Algorithm 4 of [11] as originally presented instead obtains a $\text{poly}(k, \log d)$ -approximate matrix, since the initialization for Algorithm 4 in [11] has bicriteria rank $O(k \log d)$, and the approximation factor of Algorithm 4 of [11] is proportional to the bicriteria rank of the initialization (due to the use of well-conditioned bases [15]). This is further discussed in Appendix E of [14]. We could remove the $\log d$, in the approximation factor of Algorithm 4 of [11], by using our Algorithm 7 as the initialization instead.

Acknowledgements

We are grateful to Peilin Zhong for suggesting the proof of Lemma B.5, which is part of the ℓ_p column subset selection lower bound. A. Mahankali was supported by the SURF award from CMU's Undergraduate Research Office. D. Woodruff was supported by the Office of Naval Research (ONR) grant N00014-18-1-256.

References

- [1] Arturs Backurs, Piotr Indyk, Eric Price, Ilya P. Razenshteyn, and David P. Woodruff. Nearly-optimal bounds for sparse recovery in generic norms, with applications to k -median sketching. *CoRR*, abs/1504.01076, 2015.
- [2] Frank Ban, Vijay Bhattiprolu, Karl Bringmann, Pavel Kolev, Euiwoong Lee, and David P. Woodruff. A ptas for ℓ_p -low rank approximation. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '19, page 747–766, USA, 2019. Society for Industrial and Applied Mathematics.
- [3] Boaz Barak, Fernando G. S. L. Brandão, Aram Wettroth Harrow, Jonathan A. Kelner, David Steurer, and Yuan Zhou. Hypercontractivity, sum-of-squares proofs, and their applications. In Howard J. Karloff and Toniann Pitassi, editors, *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 307–326. ACM, 2012.
- [4] Aditya Bhaskara and Silvio Lattanzi. Non-negative sparse regression and column subset selection with L1 error. In Anna R. Karlin, editor, *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, volume 94 of *LIPIcs*, pages 7:1–7:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- [5] Christos Boutsidis, Petros Drineas, and Malik Magdon-Ismail. Near optimal column-based matrix reconstruction. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 305–314, 2011.
- [6] Christos Boutsidis, Michael W Mahoney, and Petros Drineas. An improved approximation algorithm for the column subset selection problem. In *Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 968–977. Society for Industrial and Applied Mathematics, 2009.
- [7] J. Paul Brooks and José H. Dulá. The ℓ_1 -norm best-fit hyperplane problem. *Appl. Math. Lett.*, 26(1):51–55, 2013.
- [8] J. Paul Brooks, José H. Dulá, and Edward L Boone. A pure ℓ_1 -norm principal component analysis. *Computational statistics & data analysis*, 61:83–98, 2013.
- [9] J. Paul Brooks and Sapan Jot. Pcall: An implementation in r of three methods for ℓ_1 -norm principal component analysis. *Optimization Online preprint*, 2012.
- [10] J. T. Chang and D. Pollard. Conditioning as disintegration. *Statistica Neerlandica*, 51(3):287–317, 1997.
- [11] Flavio Chierichetti, Sreenivas Gollapudi, Ravi Kumar, Silvio Lattanzi, Rina Panigrahy, and David P. Woodruff. Algorithms for ℓ_p low-rank approximation. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 806–814, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- [12] Kenneth L Clarkson and David P Woodruff. Input sparsity and hardness for robust subspace approximation. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 310–329. IEEE, <https://arxiv.org/pdf/1510.06073>, 2015.
- [13] Michael B. Cohen and Richard Peng. ℓ_p row sampling by lewis weights. *CoRR*, abs/1412.0588, 2014.

- [14] Chen Dan, Hong Wang, Hongyang Zhang, Yuchen Zhou, and Pradeep Ravikumar. Optimal analysis of subset-selection based ℓ_p low-rank approximation. In Wallach et al. [52], pages 2537–2548.
- [15] Anirban Dasgupta, Petros Drineas, Boulos Harb, Ravi Kumar, and Michael W. Mahoney. Sampling algorithms and coresets for ℓ_p regression. *SIAM J. Comput.*, 38(5):2060–2078, 2009.
- [16] Holger Dell, Thore Husfeldt, Dániel Marx, Nina Taslaman, and Martin Wahlén. Exponential time complexity of the permanent and the tutte polynomial. *ACM Trans. Algorithms*, 10(4), August 2014.
- [17] Amit Deshpande and Luis Rademacher. Efficient volume sampling for row/column subset selection. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 329–338. IEEE, 2010.
- [18] Amit Deshpande, Madhur Tulsiani, and Nisheeth K. Vishnoi. Algorithms and hardness for subspace approximation. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, page 482–496, USA, 2011. Society for Industrial and Applied Mathematics.
- [19] Amit Deshpande and Kasturi R. Varadarajan. Sampling-based dimension reduction for subspace approximation. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007*, pages 641–650, 2007.
- [20] Amit Deshpande and Santosh Vempala. Adaptive sampling and fast low-rank matrix approximation. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 292–303. Springer, 2006.
- [21] Dan Feldman, Amos Fiat, Micha Sharir, and Danny Segev. Bi-criteria linear-time approximations for generalized k-mean/median/center. In *Proceedings of the 23rd ACM Symposium on Computational Geometry, Gyeongju, South Korea, June 6-8, 2007*, pages 19–26, 2007.
- [22] Nicolas Gillis and Stephen A Vavasis. On the complexity of robust pca and ℓ_1 -norm low-rank matrix approximation. *arXiv preprint arXiv:1509.09236*, 2015.
- [23] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *Journal of Computer and System Sciences*, 62(2):367 – 375, 2001.
- [24] Shuli Jiang, Dennis Li, Irene Li, Arvind Mahankali, and David P. Woodruff. An efficient protocol for distributed column subset selection in the entrywise ℓ_p -norm. Unpublished Manuscript, 2020.
- [25] William B. Johnson and Gideon Schechtman. Embedding l_p^m into l_1^n . *Acta Math.*, 149:71–85, 1982.
- [26] Daniel M. Kane, Jelani Nelson, Ely Porat, and David P. Woodruff. Fast moment estimation in data streams in optimal space. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, page 745–754, New York, NY, USA, 2011. Association for Computing Machinery.
- [27] Daniel M. Kane, Jelani Nelson, and David P. Woodruff. On the exact space complexity of sketching and streaming small norms. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '10, page 1161–1178, USA, 2010. Society for Industrial and Applied Mathematics.
- [28] Zohar Shay Karnin. Deterministic construction of a high dimensional l_p section in l_1^n for any $p < 2$. In Lance Fortnow and Salil P. Vadhan, editors, *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 645–654. ACM, 2011.
- [29] Qifa Ke and Takeo Kanade. Robust subspace computation using ℓ_1 norm. *Technical Report CMU-CS-03-172, Carnegie Mellon University, Pittsburgh, PA.*, 2003.
- [30] Qifa Ke and Takeo Kanade. Robust ℓ_1 norm factorization in the presence of outliers and missing data by alternative convex programming. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 739–746. IEEE, 2005.

- [31] Eunwoo Kim, Minsik Lee, Chong-Ho Choi, Nojun Kwak, and Songhwai Oh. Efficient-norm-based low-rank matrix approximations for large-scale problems using alternating rectified gradient method. *IEEE transactions on neural networks and learning systems*, 26(2):237–251, 2015.
- [32] Daphne Koller and Nimrod Megiddo. Constructing small sample spaces satisfying given constraints. In S. Rao Kosaraju, David S. Johnson, and Alok Aggarwal, editors, *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA*, pages 268–277. ACM, 1993.
- [33] Nojun Kwak. Principal component analysis based on ℓ_1 -norm maximization. *IEEE transactions on pattern analysis and machine intelligence*, 30(9):1672–1680, 2008.
- [34] Anastasios Kyrillidis. Simple and practical algorithms for ℓ_p -norm low-rank approximation. In Amir Globerson and Ricardo Silva, editors, *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI 2018, Monterey, California, USA, August 6-10, 2018*, pages 414–424. AUAI Press, 2018.
- [35] P. P. Markopoulos, S. Kundu, S. Chamadia, and D. A. Pados. Efficient ℓ_1 -Norm Principal-Component Analysis via Bit Flipping. *ArXiv e-prints*, 2016.
- [36] Panos P. Markopoulos, George N. Karystinos, and Dimitrios A. Pados. Some options for ℓ_1 -subspace signal processing. In *ISWCS 2013, The Tenth International Symposium on Wireless Communication Systems, Ilmenau, TU Ilmenau, Germany, August 27-30, 2013*, pages 1–5, 2013.
- [37] Panos P. Markopoulos, George N. Karystinos, and Dimitrios A. Pados. Optimal algorithms for ℓ_1 -subspace signal processing. *IEEE Trans. Signal Processing*, 62(19):5046–5058, 2014.
- [38] Jiri Matousek. Lecture notes on metric embeddings, January 2013. Available online at <https://kam.mff.cuni.cz/matousek/ba-a4.pdf>.
- [39] Deyu Meng, Zongben Xu, Lei Zhang, and Ji Zhao. A cyclic weighted median method for ℓ_1 low-rank matrix factorization with missing entries. In *AAAI*, volume 4, page 6, 2013.
- [40] J. P. Nolan. *Stable Distributions - Models for Heavy Tailed Data*. Birkhauser, Boston, 2018. In progress, Chapter 1 online at <http://fs2.american.edu/jpnolan/www/stable/stable.html>.
- [41] John P. Nolan. Numerical calculation of stable densities and distribution functions. *Communications in Statistics. Stochastic Models*, 13(4):759–774, 1997.
- [42] Young Woong Park and Diego Klabjan. Iteratively reweighted least squares algorithms for ℓ_1 -norm principal component analysis. *arXiv preprint arXiv:1609.02997*, 2016.
- [43] Prasad Raghavendra and David Steurer. Graph expansion and the unique games conjecture. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing, STOC '10*, page 755–764, New York, NY, USA, 2010. Association for Computing Machinery.
- [44] Prasad Raghavendra, David Steurer, and Prasad Tetali. Approximations for the isoperimetric and spectral profile of graphs and related parameters. In Leonard J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 631–640. ACM, 2010.
- [45] Ilya Razenshteyn, Zhao Song, and David P. Woodruff. Weighted low rank approximations with provable guarantees. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing, STOC '16*, page 250–263, New York, NY, USA, 2016. Association for Computing Machinery.
- [46] Mark Rudelson and Roman Vershynin. Non-asymptotic theory of random matrices: extreme singular values, 2010.

- [47] Christian Sohler and David P. Woodruff. Subspace embeddings for the l_1 -norm with applications. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, page 755–764, New York, NY, USA, 2011. Association for Computing Machinery.
- [48] Zhao Song, David P. Woodruff, and Peilin Zhong. Low rank approximation with entrywise ℓ_1 -norm error. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 688–701. ACM, 2017.
- [49] Zhao Song, David P. Woodruff, and Peilin Zhong. Average case column subset selection for entrywise ℓ_1 -norm loss. In Wallach et al. [52], pages 10111–10121.
- [50] Zhao Song, David P. Woodruff, and Peilin Zhong. Relative error tensor low rank approximation. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '19, page 2772–2789, USA, 2019. Society for Industrial and Applied Mathematics.
- [51] Zhao Song, David P. Woodruff, and Peilin Zhong. Towards a zero-one law for column subset selection. In Wallach et al. [52], pages 6120–6131.
- [52] Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett, editors. *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, 2019.
- [53] David P. Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends in Theoretical Computer Science*, 10(1-2):1–157, 2014.
- [54] Yinqiang Zheng, Guangcan Liu, Shigeki Sugimoto, Shuicheng Yan, and Masatoshi Okutomi. Practical low-rank matrix approximation under robust ℓ_1 -norm. In *2012 IEEE Conference on Computer Vision and Pattern Recognition, Providence, RI, USA, June 16-21, 2012*, pages 1410–1417, 2012.