

Strong Coresets for k -Median and Subspace Approximation: Goodbye Dimension

Christian Sohler

Department of Computer Science

TU Dortmund

Dortmund, Germany

Email: christian.sohler@tu-dortmund.de

David P. Woodruff

Department of Computer Science

Carnegie Mellon University

Pittsburgh, USA

dwoodruf@cs.cmu.edu

Abstract—We obtain the first strong coresets for the k -median and subspace approximation problems with sum of distances objective function, on n points in d dimensions, with a number of weighted points that is independent of both n and d ; namely, our coresets have size $\text{poly}(k/\epsilon)$. A strong coreset $(1 + \epsilon)$ -approximates the cost function for all possible sets of centers simultaneously. We also give efficient $\text{nnz}(A) + (n + d)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon))$ time algorithms for computing these coresets.

We obtain the result by introducing a new dimensionality reduction technique for coresets that significantly generalizes an earlier result of Feldman, Sohler and Schmidt [13] for squared Euclidean distances to sums of p -th powers of Euclidean distances for constant $p \geq 1$.

Keywords-clustering, dimensionality reduction, subspace approximation

I. INTRODUCTION

Coresets are a technique for data size reduction, which have been developed for a large family of problems in machine learning and statistics. Given a set P of n points p_1, \dots, p_n each in \mathbb{R}^d , loosely speaking a *coreset* is a low-memory data structure D which can be used in place of P to approximate the cost of any query Q on P . For example, in the Frobenius norm subspace approximation problem, one may be interested in computing an approximation to $\sum_{i=1}^n \|p_i - p_i P_V\|_2^2$, where P_V is the orthogonal projection onto a k -dimensional subspace V which corresponds to the query Q . As another important example, in the k -means problem one may be given a query $Q = \{q_1, \dots, q_k\}$ of k points, and one may be interested in computing an approximation to $\sum_{i=1}^n \|p_i - n(p_i, Q)\|_2^2$, where $n(p_i, Q)$ denotes the closest point in Q to point p_i . In these examples, the notion of approximation is a $(1 + \epsilon)$ -relative error approximation, that is, a value $(1 \pm \epsilon) \sum_{i=1}^n \|p_i - p_i P_V\|_2^2$ for the subspace approximation problem, and a value $(1 \pm \epsilon) \sum_{i=1}^n \|p_i - n(p_i, Q)\|_2^2$ for the k -means problem.

Often in these problems one seeks a *strong coreset*, which means that with high probability, the data struc-

ture D should work *simultaneously* for all queries Q . That is, one may use random choices in the construction of D , but after forming D it should be the case that D can be used to provide a $(1 + \epsilon)$ -relative error approximation for every possible query simultaneously. An advantage of such a coreset is that for *any objective function* for which a table of $(1 + \epsilon)$ -approximate values to all possible queries can be used to provide a $(1 + O(\epsilon))$ -approximation to the objective function, one can throw away the original set of points and instead just retain the data structure D . For example, note that the above coreset for subspace approximation contains enough information to approximately solve principal component analysis (PCA), since if one finds the query k -dimensional subspace with minimum approximate value, this provides a k -dimensional subspace providing a $(1 + \epsilon)$ -approximation to the space spanned by the top k principal components. However, the above coreset for subspace approximation can also be used to solve the k -means problem, since the latter can be rewritten as a *constrained* low rank approximation problem [2], [7]. Thus, given that a strong coreset approximately preserves the cost of any query, it can be used in place of the original point set in any application which depends only on the answers to the queries. Note that if the coreset were instead to only approximately preserve the cost of any fixed query with high probability, then it might not be possible to solve the problem using the coreset since one may need to adaptively query the data structure, and outputs to successive queries may no longer be correct since the inputs depend on outputs to previous queries.

Another advantage of a coreset is if it small, then it leads to considerable efficiency gains. For example, in distributed settings, each machine which has a subset of input points can compress its input points to a coreset, and then communicate the coreset to a central coordinator. The central coordinator, who often has

more resources available, can then combine the coresets and use them to optimize the desired function. As communication is a bottleneck, a small coreset gives rise to more efficient protocols. Similarly, when processing a data stream, a common technique is the merge-and-reduce framework, in which one partitions the stream into chunks, computes a coreset on each chunk, and merges the coresets in a binary tree like structure as one processes successive chunks of the data stream. A small coreset thus leads to small space streaming algorithms.

A long line of work has focused on developing strong coresets for both the subspace approximation problem [9], [10], [23], [12], [11], [24], [13] and the k -means problem [1], [18], [15], [16], [17], [4], [19], [11], [14], [13]. Prior to the work of [13], all previous coresets stored a weighted set of points, and the query just consisted of evaluating the same objective function on these weighted points. Moreover all such works required storing a number of points that was at least d , and an important question was to obtain coresets with a number of points *independent* of d . In [13], by taking the top $O(k/\epsilon)$ principal components of the input points, arranged as an $n \times d$ matrix A , it was shown how to obtain the first strong coresets for the subspace approximation problem with a number of points independent of d and n , namely, the authors achieved a coreset size of $O(k/\epsilon)$ points. An important idea to obtain this result was that the cost of projecting the points on the first $O(k/\epsilon)$ principal components is approximately present for every candidate subspace and therefore can be dealt with as an additive constant. The authors also extend this result to the k -means problem by proving that the projection on the first $O(k/\epsilon^2)$ principal components together with an appropriate constant will provide a coreset (of linear size but smaller dimension) for the k -means problem. Combining this with existing constructions they achieved a coreset size of $\text{poly}(k/\epsilon)$ points.

The $O(k/\epsilon^2)$ bound for the k -means problem was improved in [7] by using the fact that the k -means problem can be viewed as a constrained subspace approximation problem. In [7] the authors also find such a coreset in $\text{nnz}(A)$ time, where $\text{nnz}(A)$ is the number of non-zero entries of A .

A major open question was if one could obtain strong coresets *independent of d (and n)* for k -median and the subspace approximation problem with *sum of distances* $\sum_{i=1}^n \|p_i - p_i P_V\|_2$, as opposed to the sum of squares of distances. Unlike the k -means and sum of squares objective for subspace approximation, the k -median and sum of distances measures are much less amenable to algebraic manipulation; indeed there is no singular

value decomposition (SVD) which was the driving force behind previous results. Notably, this version of the subspace problem is NP-hard [6], unlike minimizing the sum of squares.

A. Our Contributions

Our main result is the construction of the first strong coresets independent of the dimension d and number n of input points for the k -median problem, as well as for the subspace approximation problem with sum of distances $\sum_{i=1}^n \|p_i - p_i P_V\|_2$. Our coresets have size $\text{poly}(k/\epsilon)$ for both problems, and consist of a weighted set of points with a small twist. We add a single extra dimension to each of our points! We explain this more below.

Our main new technique is a dimensionality reduction that generalizes a result of [13] for sum of squared distances to p -th powers of distances for any constant $p \geq 1$. We also show how to build a strong coreset for subspace approximation with p -th powers of Euclidean distance cost measure, for constant $p \geq 1$. Finally, we show how to find such coresets in time $\tilde{O}(\text{nnz}(A) + (n+d)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon)))$ for the k -median problem and for the subspace approximation problem with $p \in [1, 2)$, and in $\text{nnz}(A)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon))$ time for the subspace approximation problem with $p > 2$.

1) *Dimensionality Reduction*: We start by outlining our dimensionality reduction technique for the sum of distances objective function. A natural approach to try is to find a low dimensional subspace S of \mathbb{R}^d so that for any rank- k subspace V , $\|A - AP_V\|_{1,2} \approx \|B - BP_V\|_{1,2} + \|A - B\|_{1,2}$, where for a matrix C , $\|C\|_{1,2}$ denotes the sum of Euclidean norms of rows of C , and here B is the projection of the rows of A (our initial points) onto the subspace S . Indeed, this is exactly the approach taken by [13], [7] for the subspace approximation problem with sum of squares of distances, where among other constructions, S can be chosen to be the span of the top $O(k/\epsilon)$ singular vectors of A . It can be shown that the sum of squared distances to any object that is contained in a k -dimensional subspace is roughly the projection cost onto the optimal $O(k/\epsilon)$ -dimensional subspace plus the cost of the projected points. One way to think of this approach is to split the cost into a structured part (the low dimensional point set) and a “pseudorandom” part (the projection cost), where the pseudorandom part essentially acts like a random point set, as its cost will occur for any k -dimensional object, while the structured part is the one that can differ. What significantly helps in the case of squared distances is the Pythagorean theorem, which often allows to easily

express distances as the sum of “independent” distances. For example, if our object is contained in the optimal $O(k/\epsilon)$ -dimensional subspace, then the cost of each point is the squared distance of the projection plus the squared distance from the projected point to the object.

Unfortunately, we do not have such a simple formula for exponents other than 2. For the sum of distances one can show that an analogous approach does not work. Consider a set of n points in \mathbb{R}^d where each coordinate of each point is drawn independently from a Gaussian distribution with expectation 0 and variance $1/d$, i.e., the expected squared length of each point is 1. We will assume that d is large, which implies that the squared length is sharply concentrated and the expected length of the vector is close to 1. Assume now that similarly to the case of squared distances we project our input point set to a low dimensional subspace that minimizes the sum of squared projection lengths and we would like to use the projected point set together with the sum of projection lengths as a coresets (where the sum is used as an additive constant in the costs). We will now argue that this cannot work for sufficiently large n and d . Assume that the dimension of the low dimensional subspace is ℓ , a value independent of n and d . In order to understand the properties of the optimal subspace, we first consider an arbitrary fixed subspace of dimension ℓ . If we now consider a random vector $x = (x_1, \dots, x_d)$ where each x_i is chosen from the Gaussian distribution as described above, we notice that since the Gaussian distribution is invariant under rotation, that the expected squared length of the projected point is equal to the expected squared length of the random vector $x' = (x_1, \dots, x_\ell, 0, \dots, 0)$, which is ℓ/d and the expected squared length of the projection is $1 - \ell/d$. For sufficiently large d this length approaches 0 and the expected squared length of the projection approaches 1. Thus, for a fixed subspace, the expected sum of squared distances is n and for $n \rightarrow \infty$ we get that the sum of squared projection lengths is sharply concentrated. Using a union bound over a net of all subspaces we conclude that every ℓ -dimensional subspace will have cost roughly n .

Now consider the cost of an arbitrary point q at distance 1 from the origin. The expected distance of an input point to this point is roughly $\sqrt{2}$ and so the sum of distances will approach $\sqrt{2}n$ as $n \rightarrow \infty$. Now recall that the length of the projection of the input points goes to 0. Thus their distance to q will be roughly 1. Thus, the sum of distances of the projected points is roughly n plus the projection cost, which is roughly n , and thus will give an estimate of $2n$, which is not a

$(1 + \epsilon)$ -approximation. Hence, we cannot simply work with a single additive weight as in the case of squared distances.

Instead, we proceed as follows. We first describe the existence of a coresets and then how to find it efficiently. We start with a k -dimensional subspace S of \mathbb{R}^d for which $\|A(I - P_S)\|_{1,2} = \min_{\text{rank-}k \text{ subspaces } S} \|A(I - P_S)\|_{1,2} = \text{opt}$, where P_S denotes the orthogonal projection onto S . We iteratively augment S by k -dimensional subspaces until the cost no longer drops by $\epsilon^2 \text{opt}$. That is, in the first step, we try to find a $2k$ -dimensional subspace S' containing S for which $\|A(I - P_{S'})\|_{1,2} \leq \text{opt} - \epsilon^2 \text{opt}$. We then replace S with S' . In the second step, we try to find a $3k$ -dimensional subspace S'' containing S for which $\|A(I - P_{S''})\|_{1,2} \leq \text{opt}' - \epsilon^2 \text{opt}'$, where $\|A(I - P_S)\|_{1,2} = \text{opt}' \leq \text{opt} - \epsilon^2 \text{opt}$, etc. This process repeats for at most ϵ^{-2} steps, at which point we have an at most k/ϵ^2 -dimensional subspace S for which for any k -dimensional subspace V , $\|A(I - P_{V \cup S})\|_{1,2} \leq (1 + \epsilon)\|A(I - P_S)\|_{1,2}$. The latter property can be shown to imply that $\|AP_{V \cup S} - AP_S\|_{1,2} \leq \epsilon \cdot \text{opt}$, that is, the projections of the n points onto S are close to the projections of the n points onto $V \cup S$, for any V . See Lemma 6.

Next, since we can “move” each of the rows of $AP_{V \cup S}$ to the corresponding rows in AP_S by paying a total sum of distances cost of $\epsilon \cdot \text{opt}$, it follows by the triangle inequality that for any set C of points that is contained in a k -dimensional subspace V , the sum of distances from the rows of AP_S to their corresponding closest points in C is within $\epsilon \cdot \text{opt}$ of the sum of distances from the rows of $AP_{V \cup S}$ to their corresponding closest points in C .

Now we want to replace our original points (the rows of A) with their projections onto S , namely, replace A with AP_S . Although this step by itself does not reduce the number n of points, each of the n points after projection lives in a much lower k/ϵ^2 dimensional subspace (rather than the initial space which has dimension d), and we will then be able to apply coresets construction techniques which depend on this much smaller dimension. For any set C of points contained in a k -dimensional subspace V , by the Pythagorean theorem we can write the distance of a row p of A to C as $\sqrt{a^2 + b^2}$, where a is the distance of p to $V \cup S$, and b is the distance of the projection of p onto $V \cup S$, to C . We instead try to approximate $\sqrt{a^2 + b^2}$ by $\sqrt{f^2 + g^2}$, where f is the distance of p to S and g is the distance of the projection of p onto S , to C . We observe in Lemma 4 that $|\sqrt{a^2 + b^2} - \sqrt{f^2 + g^2}| \leq |a - f| + |b - g|$, and we know that the average values (over the n points) of

$|a - f|$ and $|b - g|$ are small by Lemma 6 combined with the triangle inequality.

Note that the value f , which is different for each of the n rows of A , *does not depend* on V or C , whereas the value g is exactly the distance of the row of AP_S to C . If we were simply to define $B = AP_S$, then $\|B - B_C\|_{1,2}$, where the i -th row of B_C contains the closest point (of the closure) of C to the corresponding row of B , would fail to capture the distances of the n rows of A to S . Further, unlike for the $\|\cdot\|_{2,2}$ norm, we cannot add a single number $\|A(I - P_S)\|_{2,2}$ to account for this, which is a technique used in [13], [7]; this is precisely the difficulty of the $\|\cdot\|_{1,2}$ norm that we must deal with. Instead, a crucial idea is to *append one additional coordinate* to each row of B , where in the i -th row we append $\|A_{i*}(I - P_S)\|_2$, where P_S is the orthogonal projection onto S . Then, to compute the distance to a k -dimensional subspace V , instead of approximating $\|A - A_C\|_{1,2}$ by $\|B - B_C\|_{1,2}$, where the i -th rows of A_C and B_C contain the closest points (of the closure) of C to the corresponding row of A and the first d coordinates of the corresponding row of B , respectively, we approximate $\|A - AP_V\|_{1,2}$ by $\|B - B_C I^T\|_{1,2}$, where $B_C I^T$ is the matrix which appends an all 0 column to B_C . Thus, the norm of the i -th row of $B - B_C I^T$ is $\sqrt{f^2 + g^2}$, where f is the distance of A_i to S , captured by the $(d + 1)$ -st coordinate of B , and g is the distance of $A_i P_S$ to C . Thus, we have “encoded” the distances of the rows of A to AP_S in the coreset this way. Note that this appended additional coordinate cannot be taken out of each row and combined into a single number, as in [13], [7], because for each row, its square is added to the squared distance of a point to its projection onto S , and then a square root is taken, so it occurs “under the square root” in the distance computations.

Optimizing the Running Time. We next implement the steps above in $\tilde{O}(\text{nnz}(A)) + (n + d)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon))$ time. We first show how to reformulate our algorithm in Section III-C so that it suffices to run any algorithm for finding an ik -dimensional subspace S of \mathbb{R}^d which a $(1 + \epsilon^2/2)$ -approximate ik -dimensional subspace with respect to the $\|\cdot\|_{1,2}$ norm, for a random integer i in $\{1, 2, \dots, 10/\epsilon^2\}$. It is known how to find such a subspace in $\text{nnz}(A) + (n + d)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon))$ time [6]. This is done in our DIMENSIONALITYREDUCTIONII algorithm.

After finding such a subspace S , which is an ik -dimensional subspace of \mathbb{R}^d , we need the distance of each row of A to S . To do so, we set up n regression problems, the i -th being: $\min_x \|A_i - xV^T\|_2$, where $P_S = VV^T$ and V^T is an orthonormal basis for

S , which can be computed in $d \cdot \text{poly}(k/\epsilon)$ time. Using input sparsity time algorithms for regression [5], [21], [22], if we choose a CountSketch matrix S with $\text{poly}(k/\epsilon)$ rows, then with probability at least $9/10$, $\|A_i S - xV^T S\|_2 = (1 \pm \epsilon)\|A_i - xV^T\|_2$ for all x . We can compute AS in $\text{nnz}(A)$ time and $V^T S$ in $d \cdot \text{poly}(k/\epsilon)$ time, at which point we can solve the n regression problems in $n \cdot \text{poly}(k/\epsilon)$ time, so $\text{nnz}(A) + (n + d)\text{poly}(k/\epsilon)$ total time. We repeat the entire procedure $O(\log n)$ times, and take the median of our $O(\log n)$ estimates, for each $i \in \{1, 2, \dots, n\}$, giving us $O(\text{nnz}(A) \log n + (n + d)\text{poly}(k/\epsilon) \log n) = \tilde{O}(\text{nnz}(A) + (n + d)\text{poly}(k/\epsilon))$ total time to obtain $(1 + \epsilon)$ -approximations of the distances for each row of A to S . Here, for a function f , $\tilde{O}(f) = f \cdot \text{poly}(\log f)$. We note that such approximations, when used as the last coordinates of the rows of matrix B , only change $\|B - BIP_I^T\|_{1,2}$ by a $(1 + \epsilon)$ -factor, where P is any rank- k orthogonal projection matrix. See Lemma 11.

We also need to project each of the rows of A onto S , which would take more than $\text{nnz}(A)$ time; indeed, just writing down such projections could take $\Omega(nd)$ time. Fortunately, for our coreset constructions we describe next, we never need to explicitly perform this projection. We first show how to avoid this for the subspace approximation problem.

2) *Coreset Construction for Subspace Approximation:* We first explain the construction for $p = 1$, then how to optimize the running time. At this point we have a rank- $\text{poly}(k/\epsilon)$ matrix B (or more precisely a rank- $(ik + 1)$ matrix B for some $i = O(1/\epsilon^2)$) for which for all k -dimensional orthogonal projection matrices P , $\| \|A - AP\|_{1,2} - \|B - BIP_I^T\|_{1,2} \| \leq \epsilon \|A - AP\|_{1,2}$. Unfortunately the dimension of B is still $n \times d$. Viewing the rows as points in \mathbb{R}^d , we would like to reduce the number n of points. We first observe that since ℓ_2 embeds linearly into ℓ_1 with distortion at most $(1 + \epsilon)$ via multiplication by a Gaussian matrix G (this is a special case Dvoretzky’s theorem), we have $\| \|BG - BIP_I^T G\|_{1,1} = (1 \pm \epsilon) \|B - BIP_I^T\|_{1,2}$ for all k -dimensional orthogonal projection matrices P . Here $\|\cdot\|_{1,1}$ denotes the entrywise 1-norm of a matrix. The intuition here is that G maps \mathbb{R}^d to $\mathbb{R}^{d'}$ for a $d' > d$ for which the image of \mathbb{R}^d in $\mathbb{R}^{d'}$ consists of only flat vectors, so the ℓ_1 -norm of every vector coincides with its ℓ_2 -norm, up to $(1 \pm \epsilon)$ -factor, after scaling by the square root of the dimension. Here, $d' = O(d(\log(1/\epsilon))/\epsilon^2)$, but it does not matter for our purposes since we will never actually instantiate G .

Next, we appeal to the Lewis weight sampling result of [8], which says that one can find a sam-

pling and rescaling matrix T (a matrix which just samples rows and rescales them by positive weights) with $O(r \log(r)/\epsilon^2)$ rows for which for any rank- r space C , $\|TCx\|_1 = (1 \pm \epsilon)\|Cx\|_1$ simultaneously for all x . Noting that for every P , each column of $BG - BIPIT^T G$ is in the column span of B , which is a $\text{poly}(k/\epsilon)$ -dimensional subspace, we have $\|TBG - TBIPIT^T G\|_{1,1} = (1 \pm \epsilon)\|BG - BIPIT^T G\|_{1,1}$ for all rank- k orthogonal projection matrices P , where T has $\text{poly}(k/\epsilon)$ rows. Finally, noting that the rows of $TBG - TBIPIT^T G$ are still in the row span of G , we can apply Dvoretzky's theorem one more time to conclude that $\|TBG - TBIPIT^T G\|_{1,1} = (1 \pm \epsilon)\|TB - TBIPIT^T\|_{1,2}$ for all rank- k orthogonal projection matrices P . Stringing the inequalities together, we obtain $\|TB - TBIPIT^T\|_{1,2} = (1 \pm \Theta(\epsilon))\|B - BIPIT^T\|_{1,2}$ for all rank- k orthogonal projection matrices. Note that we never need to multiply by G . Rather G is a tool in the analysis which shows the sampling procedure of [8] works for sums of Euclidean norms.

Consequently, our strong coresets consists of the rows of TB , so $\text{poly}(k/\epsilon)$ points in \mathbb{R}^d , i.e., the rows of T . These are the analogue of the k/ϵ right singular vectors of [13], [7] used to obtain a strong coresets for the $\|\cdot\|_{2,2}$ error measure. We discuss a similar argument for p -th powers below.

Optimizing the Running Time. We now obtain $\tilde{O}(\text{nnz}(A) + (n+d)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon)))$ running time. In this running time we can find the $\text{poly}(k/\epsilon)$ -dimensional subspace S for which $B = [AP_S, v]$, where P_S is the orthogonal projection onto the $\text{poly}(k/\epsilon)$ -dimensional subspace found by our DIMENSIONALITYREDUCTIONII algorithm, and $v_i = (1 \pm \epsilon)\|A_{i*} - A'_{i*}\|_2$, where $A' = AP_S$ for $i = 1, 2, \dots, n$.

As above, let $V \in \mathbb{R}^{d \times \text{poly}(k/\epsilon)}$ have columns which form an orthonormal basis for the column span of S . To find the sampling and rescaling matrix T , the procedure in Theorem 1.1 of [8] takes time equal to that of $O(\log \log n)$ invocations of constant factor ℓ_2 -leverage score approximations of matrices of the form WAV , where W is a non-negative diagonal matrix. We use the input sparsity time approximate leverage score samplers of [5], [21], [22], which compute SWA for a CountSketch matrix S with $\text{poly}(k/\epsilon)$ rows. This procedure computes SWA in $O(\text{nnz}(A))$ time, then computes $SWAV$ in $d \cdot \text{poly}(k/\epsilon)$ time, then a QR factorization in $\text{poly}(k/\epsilon)$ time, then $(WAV)(R^{-1}G)$ for a Gaussian matrix G with $O(\log n)$ columns. The row norms of $WAV(R^{-1}G)$ can be computed in $\text{nnz}(A) \log n + d \cdot \text{poly}(k/\epsilon) \log n$ time using that G has only $O(\log n)$ columns. Since the procedure reduces

to $O(\log \log n)$ invocations of this, in total this gives $\tilde{O}(\text{nnz}(A) + (n+d)\text{poly}(k/\epsilon))$ time to find the matrix T . Finally, T selects $\text{poly}(k/\epsilon)$ rows of A , and for each we compute its projection onto S , taking $d \cdot \text{poly}(k/\epsilon)$ time in total. We also output the corresponding entry of v . We thus obtain our coresets TB in $\tilde{O}(\text{nnz}(A) + (n+d)\text{poly}(k/\epsilon))$ total time.

Our coresets for subspace approximation with sum of p -th powers error measure follows via similar techniques. The running time is slightly worse for $p > 2$ due to the fact that we can only implement our DIMENSIONALITYREDUCTIONII algorithm in $\tilde{O}(\text{nnz}(A) + (n+d)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon)))$ time if $p \in [1, 2)$. For $p > 2$ we use a slower algorithm running in $O(\text{nnz}(A)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon)))$ due to [10] (they state their algorithm as $O(nd\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon)))$ but if A is sparse, the $nd \cdot \text{poly}(k/\epsilon)$ can be replaced with an $\text{nnz}(A) \cdot \text{poly}(k/\epsilon)$ given that their algorithm just requires computing projections).

3) Coresets Construction for k -Median: To obtain a coresets for k -median, we first apply our dimensionality reduction to get a matrix B such that for every set of k -centers C we have $|\|A - A_C\|_{1,2} - \|B - B_C I^T\|_{1,2}| \leq \epsilon \cdot \|A - A_C\|_{1,2}$, where A_C and B_C denote the matrices that contain in the i -th row the closest center of C to the i -th row of A and B , respectively. We note that B can be viewed as a point set in $O(k/\epsilon^2)$ dimensions. We can then use an arbitrary coresets construction for this low dimensional point set where we append k arbitrary dimensions to the space. Thus, the effect of the construction will be to replace the d in a coresets construction by $O(k/\epsilon^2)$. We claim that a coresets for this enlarged space is also a coresets for the d -dimensional space. The reason is that any set of k -centers in the d -dimensional space is either in the span of B (in which case the coresets guarantee holds) or there is an orthogonal transformation that does not change B and maps the remaining centers to the k added dimensions. This implies that the coresets property holds for the full space. Thus, the cost of the coresets approximates the cost of B upto a factor of $1 \pm \epsilon$. Combining this with the error bound of $|\|A - A_C\|_{1,2} - \|B - B_C\|_{1,2}| \leq \epsilon \cdot \|A - A_C\|_{1,2}$ gives that the resulting set will be a $1 + O(\epsilon)$ coresets and the result follows by rescaling ϵ by a constant. Notice that the guarantee the coresets provides is slightly stronger than what we need as our centers will always have the last (special) coordinate equal to 0.

Plugging in the k -median coresets of [11] or [3], which are both of size $O(\frac{dk \log k}{\epsilon^2})$ (the first one has negative weights, which may be undesirable in some situations), we obtain a coresets of size $O(\frac{k^2 \log k}{\epsilon^4})$.

In order to get a running time of $O(\text{nnz}(A) + (n + d)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon)))$ we approximate the matrix B of projections with a factored low rank matrix of approximate projections, see Lemma 13.

4) *Outline:* In Section II, we give preliminaries. In Section III, we provide our main dimensionality reduction technique. In Section IV, we obtain our coresets for subspace approximation. Finally, in Section V, we obtain our coreset for k -median. Due to space constraints, we defer a number of proofs to the full version of the paper.

II. PRELIMINARIES

We use $A \in \mathbb{R}^{n \times d}$ to denote a point set of n points in d dimensions (the rows of A). A_{i*} denotes the i -th row of A and A_{*j} denotes the j -th column. For a matrix $A \in \mathbb{R}^{n \times d}$ we use $\|A\|_{p,2} = (\sum_{i=1}^n \|A_{i*}\|_2^p)^{1/p}$. In particular, we have $\|A\|_{2,2}^2 = \|A\|_F^2$, where $\|A\|_F$ denotes the Frobenius norm of A . For a subspace S we use $\text{cost}_p(A, S)$ to denote the sum of p -th powers of the l_2 -distances from the rows of A to S . For a non-empty set of points $C \subseteq \mathbb{R}^d$ we define $\text{dist}(p, C) = \inf_{q \in C} \|p - q\|_2$. The closure of a set $C \subseteq \mathbb{R}^d$ is $\{x \in \mathbb{R}^d : B(x, r) \cap C \neq \emptyset \forall r > 0\}$, where $B(x, r)$ is a ball with center x and radius r . We start with a few claims that will be useful to deal with norms and powers of norms. These are elementary properties about numbers and the proofs can be found in the full version of the paper.

Claim 1: Let $a, b, c \geq 0$ such that $a^2 = b^2 - c^2$. For $p \geq 2$ we have $a^p \leq b^p - c^p$.

Claim 2: Let $a, b, c \geq 0$ such that $a^2 = b^2 - c^2$, $a^p \geq \epsilon b^p$ and $b^p \geq c^p$. Let $1 \leq p \leq 2$ and $1 \geq \epsilon > 0$. Then $a^p \leq 10 \cdot \epsilon^{\frac{p-2}{p}} \cdot (b^p - c^p)$.

Claim 3: Let $a, b, x \geq 0$. Let $1 \geq p > 0$. Then $|(a+x)^p - (b+x)^p| \leq |a^p - b^p|$.

Lemma 4: Let $a, b, f, g \geq 0$. Then we have $|\sqrt{a^2 + b^2} - \sqrt{f^2 + g^2}| \leq |a - f| + |b - g|$.

Claim 5: Let $a, b \geq 0$ and $1 \geq \epsilon > 0$ and $p \geq 1$. Then $(a+b)^p \leq (1+\epsilon)a^p + (1+\frac{2p}{\epsilon})^p b^p$.

III. DIMENSIONALITY REDUCTION

Our first result is a dimensionality reduction lemma for clustering problems where the cluster centers are contained in a low-dimensional subspace such as, for example, k -median clustering.

Algorithm 1 Dimensionality Reduction Algorithm

- 1: **procedure** DIMREDUCTION(A, n, d, k, ϵ, p)
 - 2: Compute a $(1 + \epsilon)$ -approximation S to the k -subspace problem with cost function sum of p -th powers of l_2 -distances
 - 3: Let opt denote the cost of an optimal solution to the above problem
 - 4: Let $k^* = k$
 - 5: **while** there exists a subspace $S' \supseteq S$ of rank $k^* + k$ such that $\text{cost}_p(A, S') \leq \text{cost}_p(A, S) - \epsilon^{\max\{\frac{2}{p}, 1\}} \text{opt}/80$ **do**
 - 6: $k^* = k^* + k$
 - 7: $S = S'$
 - 8: **end while**
 - 9: Let A' be the projection of A on S
 - 10: Let $B \in \mathbb{R}^{n \times (d+1)}$ be a matrix whose entry at position $1 \leq i, j \leq d$ equals the entry of A' and whose entries in the last column are $\|A_{i*} - A'_{i*}\|_2$
 - 11: **return** B
 - 12: **end procedure**
-

In the next lemma we show for the output space S of dimension ℓ of the above algorithm and any subspace S^* of dimension $\ell + k$ that contains S that the corresponding projections of the rows of A onto S and S^* have small distance on average.

Lemma 6: Let $1 \geq \epsilon > 0$ and $p \geq 1$. Let A be the input matrix of algorithm DIMREDUCTION. Let opt be the cost of an optimal solution to the linear k -subspace problem with respect to the sum of p -th powers of l_2 -distances. Let $S \subseteq \mathbb{R}^d$ be the subspace in the last iteration of the while loop and let ℓ be its dimension. Let $S^* \subseteq \mathbb{R}^d$ be an arbitrary subspace of dimension $k + \ell$ that contains S . Let P and P^* be orthogonal projection matrices onto S and S^* . Then we have

$$\|AP - AP^*\|_{p,2}^p \leq \epsilon \cdot \text{opt}.$$

Proof: We know from the algorithm that $\|A - AP\|_{p,2}^p - \|A - AP^*\|_{p,2}^p \leq \epsilon^{\max\{\frac{2}{p}, 1\}} \text{opt}/80$. Furthermore, we have $\|A - AP\|_{p,2}^p \leq (1 + \epsilon) \cdot \text{opt}$ by the way S is computed. We first consider the case when $p = 2$. Since $A_{i*} - A_{i*}P^*$ is orthogonal to $A_{i*}P^* - A_{i*}P$ we know that in this case

$$\|A_{i*}P - A_{i*}P^*\|_2^2 = \|A_{i*} - A_{i*}P\|_2^2 - \|A_{i*} - A_{i*}P^*\|_2^2$$

Applying the above equality row wise we obtain

$$\begin{aligned} \|AP - AP^*\|_{2,2}^2 &= \sum_{i=1}^n \|A_{i*}P - A_{i*}P^*\|_2^2 \\ &= \sum_{i=1}^n (\|A_{i*} - A_{i*}P\|_2^2 - \|A_{i*} - A_{i*}P^*\|_2^2) \\ &\leq \epsilon \cdot \text{opt}. \end{aligned}$$

Next we consider $p > 2$. We define $a = \|A_{i*}P - A_{i*}P^*\|_2$, $b = \|A_{i*} - A_{i*}P\|_2$ and $c = \|A_{i*} - A_{i*}P^*\|_2$. We observe that $a^2 = b^2 - c^2$ and so by Claim 1 we obtain that

$$\|A_{i*}P - A_{i*}P^*\|_2^p \leq \|A_{i*} - A_{i*}P\|_2^p - \|A_{i*} - A_{i*}P^*\|_2^p.$$

Again we can apply the inequality row-wise and obtain

$$\begin{aligned} \|AP - AP^*\|_{p,2}^p &= \sum_{i=1}^n \|A_{i*}P - A_{i*}P^*\|_2^p \\ &\leq \sum_{i=1}^n (\|A_{i*} - A_{i*}P\|_2^p - \|A_{i*} - A_{i*}P^*\|_2^p) \\ &\leq \epsilon \cdot \text{opt}. \end{aligned}$$

Now we consider the final case of $1 \leq p < 2$. Here we will make a case distinction. The first case is that $\|A_{i*}P^* - A_{i*}P\|_2^p \leq \frac{\epsilon}{4} \|A_{i*} - A_{i*}P\|_2^p$. Let J be the set of indices for which this inequality is satisfied. It follows by summing up over all rows in J that

$$\begin{aligned} \sum_{i \in J} \|A_{i*}P^* - A_{i*}P\|_2^p &\leq \frac{\epsilon}{4} \|A - AP\|_{p,2}^p \\ &\leq \frac{\epsilon}{4} (1 + \epsilon) \text{opt} \leq \frac{\epsilon}{2} \text{opt}. \end{aligned}$$

For the remaining case we will use Claim 2 with $a = \|A_{i*}P - A_{i*}P^*\|_2$, $b = \|A_{i*} - A_{i*}P\|_2$ and $c = \|A_{i*} - A_{i*}P^*\|_2$. We observe that $a^2 = b^2 - c^2$ and that $a^p \geq \frac{\epsilon}{4} b^p$ since we are in the second case. Furthermore, $b^p \geq c^p$ by the choices of P and P^* . Therefore, Claim 2 implies

$$\|A_{i*}P - A_{i*}P^*\|_2^p \leq 10(\epsilon/4)^{\frac{p-2}{p}} (\|A_{i*} - A_{i*}P\|_2^p - \|A_{i*} - A_{i*}P^*\|_2^p).$$

Applying the above inequality row wise we obtain

$$\sum_{i \notin J} \|A_{i*}P^* - A_{i*}P\|_2^p \leq \frac{\epsilon}{2} \text{opt}.$$

Summing up the two cases yields the lemma. \blacksquare

Remark 7: We observe that in the proof we only used two properties of S . The first one is that $\|A - AP\|_{p,2}^p \leq (1 + \epsilon) \cdot \text{opt}$ and the second one is that $\|A - AP\|_{p,2}^p - \|A - AP^*\|_{p,2}^p \leq \epsilon^{\max\{\frac{2}{p}, 1\}} \text{opt}/80$. Thus, any subspace that satisfies these two properties will also satisfy the

above lemma. We will use this later on when we discuss optimizing the running time of our algorithm.

A. Dimensionality reduction for sums of Euclidean distances

We first consider the case of minimizing sum of distances. This case is technically less tedious and illustrates the underlying ideas.

Theorem 8: Let $1 \geq \epsilon > 0$. Let $A \subseteq \mathbb{R}^{n \times d}$ be a matrix. Let $B \in \mathbb{R}^{n \times (d+1)}$ be the rank $O(k/\epsilon^2)$ matrix output by algorithm DIMREDUCTION with input parameters $A, n, d, k, p = 1$ and $\epsilon/2$. Let $C \subseteq \mathbb{R}^d$ be an arbitrary non-empty set that is contained in a k -dimensional subspace. Let A' and B' be the matrices whose rows contain the closest points in the (closure of) C w.r.t. the rows of A, BI , respectively. Then we have

$$\|A - A'\|_{1,2} - \|B - B'I^T\|_{1,2} \leq \epsilon \|A - A'\|_{1,2},$$

where $I \in \mathbb{R}^{(d+1) \times d}$ has diagonal entries 1 and all other entries are 0.

Proof: Let S be the subspace as in Lemma 6 and let S^* be the span of S and C (if S^* has less than $\ell + k$ dimensions, we can add arbitrary dimensions). Let P, P^* be the corresponding orthogonal projection matrices. We know from Lemma 6 that

$$\|AP - AP^*\|_{1,2} \leq \frac{\epsilon}{2} \cdot \text{opt}$$

where opt is the cost of an optimal solution to the k -subspace problem with sum of distances. By orthogonality, we can write

$$\|A_{i*} - A'_{i*}\|_2 = \sqrt{\|A_{i*} - A_{i*}P^*\|_2^2 + \|A_{i*}P^* - A'_{i*}\|_2^2}.$$

Furthermore, we have

$$\|B_{i*} - B'_{i*}I^T\|_2 = \sqrt{\|A_{i*}P - B'_{i*}\|_2^2 + \|A_{i*} - A_{i*}P\|_2^2}.$$

Using Lemma 4 with $a = \|A_{i*} - A_{i*}P^*\|_2$, $b = \|A_{i*}P^* - A'_{i*}\|_2$, $e = \|A_{i*} - A_{i*}P\|_2$ and $f = \|A_{i*}P - B'_{i*}\|_2$ we obtain that

$$\begin{aligned} &\|A_{i*} - A'_{i*}\|_2 - \|B_{i*} - B'_{i*}I^T\|_2 \\ &\leq \left| \|A_{i*} - A_{i*}P^*\|_2 - \|A_{i*} - A_{i*}P\|_2 \right| \\ &\quad + \left| \|A_{i*}P^* - A'_{i*}\|_2 - \|A_{i*}P - B'_{i*}\|_2 \right|. \end{aligned}$$

Note that $\|A_{i*}P^* - A'_{i*}\|_2 \leq \|A_{i*}P^* - B'_{i*}\|_2$ and $\|A_{i*}P - B'_{i*}\|_2 \leq \|A_{i*}P - A'_{i*}\|_2$; indeed this follows since A'_{i*} and B'_{i*} are the closest points in the closure of C to $A_{i*}P^*$ and $A_{i*}P$, respectively. So we can upper bound $\left| \|A_{i*}P^* - A'_{i*}\|_2 - \|A_{i*}P - B'_{i*}\|_2 \right|$ by the maximum of $\left| \|A_{i*}P^* - B'_{i*}\|_2 - \|A_{i*}P - B'_{i*}\|_2 \right|$

and $|\|A_{i^*}P^* - A'_{i^*}\|_2 - \|A_{i^*}P - A'_{i^*}\|_2|$. Either way, using the triangle inequality we obtain

$$\| \|A_{i^*} - A'_{i^*}\|_2 - \|B_{i^*} - B'_{i^*}I^T\|_2 \| \leq 2 \cdot \|A_{i^*}P - A_{i^*}P^*\|_2.$$

Summing up over all rows and using $\|AP - AP^*\|_{1,2} \leq \frac{\epsilon}{2} \cdot \text{opt}$ together with the fact that the sum of distances to C is at least opt we obtain the result. \blacksquare

B. Dimensionality reduction for powers of Euclidean distances

In order to obtain a dimensionality reduction for powers of Euclidean distances we follow the same approach as before. The main challenge is that some calculations become more difficult as the triangle inequality is replaced by a relaxed triangle inequality. The proof can be found in the full version of the paper.

Theorem 9: Let $p \geq 1$ be a constant. Let $A \subseteq \mathbb{R}^{n \times d}$ be a matrix and $1 \geq \epsilon \geq 0$. Let $B \in \mathbb{R}^{n \times (d+1)}$ be the rank $O(k/\epsilon^{O(p)})$ matrix output by algorithm DIMENSIONALITYREDUCTION with input parameters $A, n, d, k, p \geq 1$ and $\frac{\epsilon^{p+3}}{3(84p)^{2p}}$. Let $C \subseteq \mathbb{R}^d$ be an arbitrary non-empty set that is contained in a k -dimensional subspace. Let A' and B' be the matrices whose rows contain the closest points in the (closure of) C w.r.t. the rows of A, BI , respectively. Then we have

$$\| \|A - A'\|_{p,2}^p - \|B - B'I^T\|_{p,2}^p \| \leq \epsilon \|A - A'\|_{p,2}^p,$$

where $I \in \mathbb{R}^{(d+1) \times d}$ has diagonal entries 1 and all other entries are 0.

C. Optimizing the Running Time

We first give an alternative algorithm to our DIMREDUCTION algorithm. This algorithm can be implemented using a black box call to an algorithm for finding low dimensional subspaces approximately minimizing the $\|\cdot\|_{p,2}^p$ norm.

Algorithm 2 Dimensionality Reduction Algorithm II

- 1: **procedure** DIMREDUCTIONII(A, n, d, k, ϵ, p)
- 2: $\tau = \Theta(\epsilon^{\max(2/p, 1)})$.
- 3: Choose a random $i^* \in \{1, 2, \dots, 10/\tau\}$
- 4: Let S be an i^*k -dimensional subspace E with

$$\begin{aligned} & \|A(I - P_E)\|_{p,2}^p \\ & \leq (1 + \Theta(\epsilon^{\max(2/p, 1)})) \min_{\text{rank-}i^*k E'} \|A(I - P_{E'})\|_{p,2}^p. \end{aligned}$$

Such a space S can be found by Theorem 1 of [6] with the k there equal to our i^*k , and the ϵ there can be set to our τ , if $p \in [1, 2)$. The success probability is at least $9/10$. For $p > 2$, one can use the algorithm in [10].

- 5: For $i = 1, \dots, n$, output a $(1 \pm \epsilon)$ -approximation to $\|A_i(I - P_S)\|_2$. These n values can be found by solving n regression problems each with probability $1 - 1/n^2$, using the regression algorithm of [5]. See Lemma 11 below.
 - 6: **end procedure**
-

Lemma 10: With probability at least $4/5$, DIMREDUCTIONII finds a $\text{poly}(k/\epsilon)$ -dimensional subspace S for which for all k -dimensional spaces W , $\|A(I - P_S)\|_{p,2}^p - \|A(I - P_{S \cup W})\|_{p,2}^p \leq \Theta(\epsilon^{\max(2/p, 1)}) \text{opt}$. Further, for $p \in [1, 2)$, finding such an S can be done in $O(\text{nnz}(A) + (n + d)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon)))$ time, and for $p > 2$, can be found in $O(\text{nnz}(A)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon)))$ time.

Proof: We condition on the event that the algorithm of [6] or of [10] for computing a $(1 + \tau)$ -approximation S succeeds, which holds with probability at least $9/10$. The algorithm requires time $O(\text{nnz}(A) + (n + d)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon)))$ for $p \in [1, 2)$ and $O(\text{nnz}(A)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon)))$ for $p > 2$. For each $j \in \{1, 2, \dots, 10/\tau + 1\}$, let V^j be the optimal jk -dimensional subspace, and consider a telescoping sum:

$$\begin{aligned} & \text{opt} - \|A(I - P_{V^{10/\tau+1}})\|_{p,2}^p \\ & = \|A(I - P_{V^1})\|_{p,2}^p - \|A(I - P_{V^{10/\tau+1}})\|_{p,2}^p \\ & = \sum_{i=1}^{10/\tau} (\|A(I - P_{V^i})\|_{p,2}^p - \|A(I - P_{V^{i+1}})\|_{p,2}^p) \\ & \geq 0. \end{aligned}$$

There are $10/\tau$ summands in the telescoping sum, and they sum up to at most opt , so a $9/10$ -fraction of them must be at most $\tau \cdot \text{opt}$. Let i^* be the index sampled by the algorithm. Then with probability at least $9/10$, we have $\|A(I - P_{V^{i^*}})\|_{p,2}^p - \|A(I - P_{V^{i^*+1}})\|_{p,2}^p \leq \tau \cdot \text{opt}$, and let us condition on this event.

Now, $\tilde{V}^{i^*} \cup W$ is an $(i^* + 1)k$ -dimensional subspace,

and so we have $\|A(I - P_{\tilde{V}^{i^*} \cup W})\|_{p,2}^p \geq \|A(I - P_{V^{i^*+1}})\|_{p,2}^p$. Also, by the guarantee of \tilde{V}^{i^*} , we have $\|A(I - P_{\tilde{V}^{i^*}})\|_{p,2}^p \leq (1 + \tau)\|A(I - P_{V^{i^*}})\|_{p,2}^p \leq (1 + \tau)(\|A(I - P_{V^{i^*}})\|_{p,2}^p + \tau \cdot \text{opt})$, where opt is the cost of the best k -dimensional subspace. Consequently, for any k -dimensional subspace W ,

$$\begin{aligned} & \|A(I - P_{\tilde{V}^{i^*}})\|_{p,2}^p - \|A(I - P_{\tilde{V}^{i^*} \cup W})\|_{p,2}^p \\ & \leq (1 + \tau)(\|A(I - P_{V^{i^*}})\|_{p,2}^p + \tau \cdot \text{opt}) \\ & \quad - \|A(I - P_{V^{i^*+1}})\|_{p,2}^p \\ & \leq O(\tau)\text{opt}, \end{aligned}$$

where we used that $\|A(I - P_{V^{i^*}})\|_{1,2} - \|A(I - P_{V^{i^*+1}})\|_{1,2} \leq \tau \text{opt}$ for our choice of i^* , and also that $\|A(I - P_{V^{i^*}})\|_{p,2}^p \leq \text{opt}$.

Note that the overall success probability is at least $1 - 1/10 - 1/10 = 4/5$, and the claimed running time follows from [6] and [10]. \blacksquare

We defer the proof of the next three lemmas to the full version of the paper.

Lemma 11: With probability at least $1 - 1/n$, DIMENSIONALITYREDUCTIONII outputs a $(1 \pm \epsilon)$ -approximation to $\|A_i(I - P_S)\|_2$ simultaneously for every $i \in [n]$. Further, this can be done in $O(\text{nnz}(A) \log n + (n + d)\text{poly}(k/\epsilon) \log n)$ time.

We also show the following lemma stating that the approximations returned by Lemma 11 suffice.

Lemma 12: Let $B \in \mathbb{R}^{n \times (d+1)}$ be a matrix for which for any rank k orthogonal projection matrix $P \in \mathbb{R}^{d \times d}$ we have

$$|\|A - AP\|_{p,2}^p - \|B - BIPIT\|_{p,2}^p| \leq \epsilon \|A - AP\|_{p,2}^p.$$

Suppose we replace the last column v of B with a vector v' for which $v_i = (1 \pm \epsilon)v'_i$ for all $i \in [n]$. Then the above inequality continues to hold with ϵ replaced with $O(\epsilon)$.

We cannot afford to compute the projection of A onto S , as this could take longer than $O(\text{nnz}(A))$ time. Fortunately, we show in Section IV that we do not need to compute this. For the k -median problem we need the following additional lemma to approximate matrix B .

Lemma 13: Let S be the subspace guaranteed by Lemma 10. Given S we can compute in time $O(\text{nnz}(A) \log n + (n + d)\text{poly}(k/\epsilon))$ a matrix \tilde{B} of rank $O(k/\epsilon^2)$ such that with probability at least $9/10$ we have for every set C contained in a k -dimensional subspace

$$|\|B - B'I\|_{p,2}^p - \|\tilde{B} - \tilde{B}'I\|_{p,2}^p| \leq \epsilon \|A - A'\|_{p,2}^p.$$

Here B' and \tilde{B}' are matrices that contain in the i -th row in the first d coordinates the point from (the closure

of) C that is closest to the i -th row of BI and $\tilde{B}I$ respectively, and have the $d + 1$ -coordinate 0.

IV. CORESETS FOR SUBSPACE APPROXIMATION

Plugging in the guarantee of Lemma 10 into Remark 7 shows how to obtain an $n \times (d + 1)$ matrix B for which

$$|\|A - AP\|_{p,2}^p - \|B - BIPIT\|_{p,2}^p| \leq \epsilon \|A - AP\|_{p,2}^p,$$

for all rank- k orthogonal projection matrices P . Lemma 10 shows how to efficiently find the S for which $B = AP_S$, and Lemma 11 shows how to efficiently find an approximate vector v of $(d + 1)$ -st coordinates of B . Further, after scaling ϵ by a constant factor, Lemma 12 shows that the above inequality continues to hold with the approximate vector v of $(d + 1)$ -st coordinates of B furnished by Lemma 11.

The main issue is that we cannot afford to compute the projection of A onto S to form the first d columns of B . Although the matrix B is $n \times (d + 1)$, it has rank $O(k/\epsilon^2)$. Thinking of its rows as n points in \mathbb{R}^{d+1} , we would like to find a weighted subset TB of these n points so that $\|TB - TBIPIT\|_{1,2} = (1 \pm \epsilon)\|B - BIPIT\|_{1,2}$ for all rank- k orthogonal projection matrices P . Here, T is called a sampling and rescaling matrix, and our goal will be to find such a T for which each row of T contains a single non-zero non-negative entry, corresponding to the sampled row of B , rescaled by a non-negative value.

Recall that B has a particular form, namely, $B = [AP_S, v]$, where P_S is the orthogonal projection onto the $\text{poly}(k/\epsilon)$ -dimensional subspace found by our DIMENSIONALITYREDUCTION or DIMENSIONALITYREDUCTIONII algorithm, and $v_i = (1 \pm \epsilon)\|A_{i*} - A'_{i*}\|_2$, where $A' = AP_S$ for $i = 1, 2, \dots, n$. Since S is $\text{poly}(k/\epsilon)$ -dimensional, we can write $P_S = UU^T$, where U is a $d \times \text{poly}(k/\epsilon)$ matrix with orthonormal columns.

Before showing how to find a sampling and rescaling matrix T , we first need the following theorem due to Dvoretzky.

Fact 14: (Special case of Dvoretzky's Theorem, see Variations and Extensions on p.30 of [20]) Let $t \geq Cd \log(1/\epsilon)/\epsilon^2$ for a sufficiently large constant $C > 0$, and suppose G is a $d \times t$ matrix of i.i.d. $N(0, 1/t)$ random variables, where $N(0, 1/t^2)$ denotes a normal random variable with mean 0 and variance $1/t$. Then with probability at least $99/100$, simultaneously for all $x \in \mathbb{R}^d$, $\|x\|_2 = (1 \pm \epsilon)\|xG\|_1$. In particular, there exists such a matrix G for which this property holds for all $x \in \mathbb{R}^d$. For $p > 1$, there is a $d \times t'$ matrix G of i.i.d. normal random variables, suitably

scaled, with $t' = (d/\epsilon)^{O(p)}$ for which for all $x \in \mathbb{R}^d$, $\|x\|_2 = (1 \pm \epsilon)\|xG\|_p$.

Lemma 15: (Sampling Lemma) Given S , in $n \cdot \text{poly}(k(\log n)/\epsilon)$ time it is possible to find a sampling and rescaling matrix T with $O(\text{rank}(S) \log(\text{rank}(S)/\epsilon^2))$ rows for which for all rank- k orthogonal projection matrices P ,

$$\|TB - TBIPIT^T\|_{1,2} = (1 \pm \epsilon)\|B - BIPIT^T\|_{1,2}.$$

Instantiating S with the output of our DIMENSIONALITYREDUCTION algorithm, T would have $O(k \log(k/\epsilon)/\epsilon^4)$ rows. Instantiating S with the output of our DIMENSIONALITYREDUCTIONII algorithm, T would have $\text{poly}(k/\epsilon)$ rows.

For constant $p > 1$, it is possible to find a sampling and rescaling matrix T with $\text{poly}(\text{rank}(S)/\epsilon)$ rows for which for all rank- k orthogonal projection matrices P ,

$$\|TB - TBIPIT^T\|_{p,2}^p = (1 \pm \epsilon)\|B - BIPIT^T\|_{p,2}^p.$$

Instantiating S with the output of our DIMENSIONALITYREDUCTION or DIMENSIONALITYREDUCTIONII algorithms, T would have $\text{poly}(k/\epsilon)$ rows.

Proof: Let t be as in Fact 14, and fix the $d \times t$ matrix G of that fact. Applying the guarantee of Fact 14 to each row of $B - BIPIT^T$,

$$\|B - BIPIT^T\|_{p,2}^p = (1 \pm \Theta(\epsilon))\|BG - BIPIT^TG\|_{p,p}^p,$$

where for a matrix C , $\|C\|_{p,p}^p$ denotes the sum of p -th powers of absolute values of its entries.

We next apply Theorem 1 of [8], which shows how, given a matrix C with f columns, to find a sampling and rescaling matrix T with $O(f\epsilon^{-2} \log f)$ rows for $p = 1$, and $(f\epsilon^{-1})^{O(p)}$ rows for $p > 1$, for which with high probability, simultaneously for all $x \in \mathbb{R}^t$, $\|TCx\|_p = (1 \pm \epsilon)\|Cx\|_p$. Further, the time to find T is $O(\log \log n)$ calls to computing 2-approximate statistical leverage scores of matrices of the form WC , where W is a non-negative diagonal matrix. Using the algorithm of Theorem 29 of [5], T can be computed in $\tilde{O}(\text{nnz}(C))$ time.

Instantiating the matrix C of the previous paragraph with the $n \times O(k/\epsilon^2)$ matrix AU , where recall $P_S = UUT^T$ and U has k/ϵ^2 columns if S is the output of our DIMENSIONALITYREDUCTION algorithm, it follows that $\|TAUx\|_1 = (1 \pm \epsilon)\|AUx\|_1$ for all $x \in \mathbb{R}^{O(k/\epsilon^2)}$, and T has $O(k \log(k/\epsilon)/\epsilon^4)$ rows. If S is the output of our DIMENSIONALITYREDUCTIONII algorithm or $p > 1$, then T has $\text{poly}(k/\epsilon)$ rows. The overall time is $n \cdot \text{poly}(k(\log n)/\epsilon)$.

Consequently,

$$\|TBG - TBIPIT^TG\|_{p,p}^p = (1 \pm \epsilon)\|BG - BIPIT^TG\|_{p,p}^p,$$

since each column of BG is in the column span of AU . Again applying the guarantee of Fact 14 to each row of $TBG - TBIPIT^TG$, we have

$$\|TBG - TBIPIT^TG\|_{p,p}^p = (1 \pm \epsilon)\|TB - TBIPIT^T\|_{p,2}^p.$$

Combining the above inequalities, we have

$$\|B - BIPIT^T\|_{p,2}^p = (1 \pm O(\epsilon))\|TB - TBIPIT^T\|_{p,2}^p,$$

and the guarantee of the lemma follows by rescaling ϵ by a constant factor. \blacksquare

Theorem 16: (Strong Coreset for Subspace Approximation) For $p = 1$, it is possible to find a matrix $TB \in \mathbb{R}^{O(k(\log k)/\epsilon^4) \times d+1}$ for which for all rank- k orthogonal projection matrices P ,

$$\|A - AP\|_{p,2} - \|TB - TBIPIT^T\|_{p,2} \leq \epsilon\|A - AP\|_{p,2}.$$

Further, in $\tilde{O}(\text{nnz}(A) + (n + d)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon)))$ time, it is possible to find a matrix $TB \in \mathbb{R}^{\text{poly}(k/\epsilon) \times d+1}$ satisfying the above condition for $p \in [1, 2)$ for all rank- k orthogonal projection matrices P .

Finally, in $\text{nnz}(A)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon))$ time, it is possible to find a matrix $TB \in \mathbb{R}^{\text{poly}(k/\epsilon) \times d+1}$ satisfying the above condition for $p > 2$ for all rank- k orthogonal projection matrices P .

Proof: We start by proving the structural part of the theorem, and then address the running time.

Let B be the output of DIMENSIONALITYREDUCTION, which has the property mentioned at the beginning of this section. As described above, we can assume DIMENSIONALITYREDUCTION produces $B = [AP_S, v]$, where P_S and v are described above. Further, we can assume AP_S is given in factored form $(AU)U^T$ for $P_S = UU^T$.

By Lemma 15, we can find a sampling and rescaling matrix T for which for all rank- k orthogonal projection matrices P , $\|TB - TBIPIT^T\|_{p,2}^p = (1 \pm \epsilon)\|B - BIPIT^T\|_{p,2}^p$, so

$$\begin{aligned} & \left| \|A - AP\|_{p,2}^p - \|TB - TBIPIT^T\|_{p,2}^p \right| \\ &= \left| \|A - AP\|_{p,2}^p - \|B - BIPIT^T\|_{p,2}^p \right| \pm \epsilon \|B - BIPIT^T\|_{p,2}^p \\ &\leq \epsilon \|A - AP\|_{p,2}^p + \epsilon \|B - BIPIT^T\|_{p,2}^p \\ &\leq \epsilon \|A - AP\|_{p,2}^p + \epsilon (\|A - AP\|_{p,2}^p + \epsilon \|A - AP\|_{p,2}^p) \\ &\leq (2\epsilon + \epsilon^2) \|A - AP\|_{p,2}^p, \end{aligned}$$

and rescaling ϵ by a constant factor gives the desired guarantee. We note that By Lemma 15, T will have $O(k(\log k)/\epsilon^4)$ rows for $p = 1$ if we run DIMENSIONALITYREDUCTION.

For an efficient algorithm, we instead run DIMENSIONALITYREDUCTIONII in $\tilde{O}(\text{nnz}(A) + (n +$

d) $\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon))$ time for $p \in [1, 2)$, and in $\text{nnz}(A)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon))$ time for $p > 2$, to obtain S , which is $\text{poly}(k/\epsilon)$ -dimensional, and in $n \cdot \text{poly}(k/\epsilon)$ time we can compute U , where $P_S = UU^T$. The correctness is given by Lemma 10. We also compute the $(d + 1)$ -st column of B in $\tilde{O}(\text{nnz}(A) + (n + d)\text{poly}(k/\epsilon))$ time. By Lemma 15, in $n \cdot \text{poly}(k(\log n)/\epsilon)$ time we can find a sampling and rescaling matrix T . Finally, T selects $\text{poly}(k/\epsilon)$ rows of A , and for each we compute its projection onto S , taking $d \cdot \text{poly}(k/\epsilon)$ time in total. We also output the corresponding entry of v . We thus obtain our coresets TB in the stated running times. ■

V. CORESETS FOR k -MEDIAN

We can combine our dimensionality reduction with coresets computations. We first compute the matrix B using our dimensionality reduction. Recall that this matrix has rank $O(k/\epsilon^2)$. We can therefore compute an orthogonal basis for the span of B and add k arbitrary dimensions. Then we apply an arbitrary coreset construction in the resulting space. We will only be interested in the space spanned by the first d dimensions of B (recall that the last dimension is only needed to “adjust” the distances). Since the Euclidean distance does not change under orthogonal transformations we can rotate any set of centers to our subspace without changing the distances. Therefore, the coreset will be a coreset for the whole input space.

Theorem 17: Let $1 \geq \epsilon > 0$. Given a matrix $A \in \mathbb{R}^{n \times d}$. We can compute in time $\tilde{O}(\text{nnz}(A) + (n + d)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon)))$ a matrix $B \in \mathbb{R}^{s \times (d+1)}$, $s = O(\frac{k^2 \log k}{\epsilon^4})$, and non-negative weights w_1, \dots, w_s such that with probability at least $3/5$, simultaneously for every set C of k centers we have

$$\| \|A - A^C\|_{1,2} - \sum_{i=1}^s w_i \|S_{i*} - S_{i*}^C\|_2 \| \leq \epsilon \|A - A^C\|_{1,2}$$

Here A^C contains for each row of A the nearest center of C , and S^C contains for each row of S the nearest center of C with respect to ST^T .

Proof: We use Lemma 10 (plugging it into Remark 7) to compute with probability at least $9/10$ in time $O(\text{nnz}(A) + (n + d)\text{poly}(k/\epsilon) + \exp(\text{poly}(k/\epsilon)))$ the subspace S of rank $O(k/\epsilon^2)$ using parameter $\epsilon/10$. We then apply Lemma 13 to obtain with probability at least $4/5$ matrix \tilde{B} of rank $O(k/\epsilon^2)$ from matrix B defined by S in time $O(\text{nnz}(A) \log n + (n + d)\text{poly}(k/\epsilon))$. Next we apply the coreset construction of [11] or [3] to obtain in time $\tilde{O}(n\text{poly}(k \log(1/\delta)/\epsilon))$ (recall that the dimension of our space is $\text{poly}(k/\epsilon)$) a coreset S^* of

size $O(k^2 \log k/\epsilon^4)$, where we set the error probability $\delta = 1/10$. Finally we compute for each coreset point the corresponding coordinates in the original space in $O(d \cdot \text{poly}(k/\epsilon))$ time to obtain the matrix S .

We get the following guarantees. Theorem 8 implies

$$\| \| \|A - A'\|_{1,2} - \|B - B'I\|_{1,2} \| \leq \epsilon \|A - A'\|_{1,2}.$$

By Lemma 13 we know that

$$\| \|B - B'I\|_{1,2} - \| \tilde{B} - (\tilde{B}I)^C \|_{1,2} \| \leq \epsilon \|A - A'\|_{1,2}$$

For the coreset S we know that

$$\| \| \tilde{B} - (\tilde{B}I)^C \|_{1,2} - \sum_{i=1}^{|S|} w_i \|S_{i*} - S_{i*}^C\|_2 \| \leq \epsilon \| \tilde{B} - (\tilde{B}I)^C \|_{1,2}.$$

Combining these statements gives

$$\begin{aligned} & \| \| \|A - A'\|_{1,2} - \sum_{i=1}^{|S|} w_i \|S_{i*} - S_{i*}^C\|_2 \| \\ & \leq 2\epsilon \|A - A'\|_{1,2} + \epsilon \| \tilde{B} - (\tilde{B}I)^C \|_{1,2} \\ & \leq 2\epsilon \|A - A'\|_{1,2} + \epsilon((1 + 2\epsilon)\epsilon \|A - A'\|_{1,2}) \\ & \leq 5\epsilon \|A - A'\|_{1,2} \end{aligned}$$

Rescaling ϵ gives the result. ■

ACKNOWLEDGMENT

Christian Sohler acknowledges the support of the German Science Foundation (DFG) Collaborative Research Center SFB 876 “Providing Information by Resource-Constrained Analysis”, project A2. David Woodruff acknowledges support in part by an Office of Naval Research (ONR) grant N00014-18-1-2562.

REFERENCES

- [1] Mihai Badoiu, Sariel Har-Peled, and Piotr Indyk. Approximate clustering via core-sets. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 250–257, 2002.
- [2] Christos Boutsidis, Anastasios Zouzias, Michael W. Mahoney, and Petros Drineas. Randomized dimensionality reduction for k-means clustering. *IEEE Trans. Information Theory*, 61(2):1045–1062, 2015.
- [3] Vladimir Braverman, Dan Feldman, and Harry Lang. New frameworks for offline and streaming coresets constructions. *CoRR*, abs/1612.00889, 2016.
- [4] Ke Chen. On coresets for k-median and k-means clustering in metric and euclidean spaces and their applications. *SIAM J. Comput.*, 39(3):923–947, 2009.

- [5] Kenneth L. Clarkson and David P. Woodruff. Low rank approximation and regression in input sparsity time. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 81–90, 2013.
- [6] Kenneth L. Clarkson and David P. Woodruff. Input sparsity and hardness for robust subspace approximation. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 310–329, 2015.
- [7] Michael B. Cohen, Sam Elder, Cameron Musco, Christopher Musco, and Madalina Persu. Dimensionality reduction for k-means clustering and low rank approximation. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 163–172, 2015.
- [8] Michael B. Cohen and Richard Peng. L_p row sampling by lewis weights. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 183–192, 2015.
- [9] Amit Deshpande, Luis Rademacher, Santosh Vempala, and Grant Wang. Matrix approximation and projective clustering via volume sampling. *Theory of Computing*, 2(12):225–247, 2006.
- [10] 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.
- [11] Dan Feldman and Michael Langberg. A unified framework for approximating and clustering data. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 569–578, 2011.
- [12] Dan Feldman, Morteza Monemizadeh, Christian Sohler, and David P. Woodruff. Coresets and sketches for high dimensional subspace approximation problems. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 630–649, 2010.
- [13] Dan Feldman, Melanie Schmidt, and Christian Sohler. Turning big data into tiny data: Constant-size coresets for k -means, PCA and projective clustering. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 1434–1453, 2013.
- [14] Dan Feldman and Leonard J. Schulman. Data reduction for weighted and outlier-resistant clustering. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 1343–1354, 2012.
- [15] Gideon Frawling and Christian Sohler. Coresets in dynamic geometric data streams. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 209–217, 2005.
- [16] Gideon Frawling and Christian Sohler. A fast k -means implementation using coresets. In *Proceedings of the 22nd ACM Symposium on Computational Geometry, Sedona, Arizona, USA, June 5-7, 2006*, pages 135–143, 2006.
- [17] Sarel Har-Peled and Akash Kushal. Smaller coresets for k -median and k -means clustering. *Discrete & Computational Geometry*, 37(1):3–19, 2007.
- [18] Sarel Har-Peled and Soham Mazumdar. On coresets for k -means and k -median clustering. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*, pages 291–300, 2004.
- [19] Michael Langberg and Leonard J. Schulman. Universal epsilon-approximators for integrals. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 598–607, 2010.
- [20] Jiri Matoušek. Lecture notes on metric embeddings. Technical report, 2013.
- [21] Xiangrui Meng and Michael W. Mahoney. Low-distortion subspace embeddings in input-sparsity time and applications to robust linear regression. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 91–100, 2013.
- [22] Jelani Nelson and Huy L. Nguyen. OSNAP: faster numerical linear algebra algorithms via sparser subspace embeddings. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 117–126, 2013.
- [23] Nariankadu D. Shyamalkumar and Kasturi R. Varadarajan. Efficient subspace approximation algorithms. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, New Orleans, Louisiana, USA, January 7-9, 2007*, pages 532–540, 2007.
- [24] Kasturi R. Varadarajan and Xin Xiao. On the sensitivity of shape fitting problems. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2012, December 15-17, 2012, Hyderabad, India*, pages 486–497, 2012.