

Smaller Coresets for k -Median and k -Means Clustering*

Sariel Har-Peled and Akash Kushal

Department of Computer Science, University of Illinois,
201 N. Goodwin Avenue, Urbana, IL 61801, USA
{sariel,kushal}@uiuc.edu

Abstract. In this paper we show that there exists a (k, ε) -coreset for k -median and k -means clustering of n points in \mathfrak{R}^d , which is of size independent of n . In particular, we construct a (k, ε) -coreset of size $O(k^2/\varepsilon^d)$ for k -median clustering, and of size $O(k^3/\varepsilon^{d+1})$ for k -means clustering.

1. Introduction

Clustering is a widely used technique in Computer Science with applications to unsupervised learning, classification, data mining, and other fields. We study two variants of the clustering problem in the geometric setting. The *geometric k -median clustering* problem is the following: Given a set P of n points in \mathfrak{R}^d , compute a set of k points (i.e., medians) such that the sum of the distances of the points in P to their respective nearest median is minimized. The k -means differs from the above in that instead of the sum of distances, we minimize the sum of squares of distances. Interestingly the 1-mean is the center of mass of the points, while the 1-median problem, also known as the Fermat–Weber problem, has no such closed form. As such the problems have usually been studied separately from each other even in the approximate setting.

An important question underlying approximation algorithms is What portion of the data is necessary to compute (approximately) a certain quantity? The smaller this portion is, the more efficient the resulting algorithm would be. A coreset is a small portion of the data, such that running a clustering algorithm on it, generates a clustering for the whole data, which is approximately optimal. In particular, a small coreset indicates that

* See http://www.uiuc.edu/~sariel/papers/04/small_coreset/ for the most recent version of this paper. Alternative titles for this paper include “Finding Space/Time Stream Coresets” and “Improved Homeland Security Using Grid Computing via Mobile ad-hoc Trustable Coresets.” Work on this paper by the first author was partially supported by NSF CAREER Award CCR-0132901.

a problem is easy to approximate. Furthermore, it implies that one can summarize and sketch the data efficiently. This is useful for database applications, where one can store such sketches efficiently, and perform efficient clustering on a database, or portions of it using the sketches.

In particular, the size of the smallest coreset needed is a fundamental combinatorial property of the clustering problem at hand. Among other things, coresets of size independent of n imply “strong” fixed parameter algorithms [DF]. Those are algorithms with the running time $O(n + \text{poly}(k, \log n, 1/\varepsilon) + \text{func}(k, \varepsilon))$, where poly denotes a polynomial, and $\text{func}(k, \varepsilon)$ denotes a function that depends only on k and ε (and the dimension d).

k-Median Clustering. The k -median problem is nontrivial even in low dimensions and achieving a good approximation proved to be a challenge. Motivated by the work of Arora [A], which proposed a new technique for geometric approximation algorithms, Arora et al. [ARR] presented a $O(n^{O(1/\varepsilon)+1})$ time $(1 + \varepsilon)$ -approximation algorithm for points in the plane. This was significantly improved by Kolliopoulos and Rao [KR] who proposed an algorithm with a running time of $O(\varrho n \log n \log k)$ for the discrete version of the problem, where the medians must belong to the input set and $\varrho = \exp[O((1 + \log 1/\varepsilon)/\varepsilon)^{d-1}]$. The k -median problem has been studied extensively for arbitrary metric spaces and is closely related to the uncapacitated facility location problem. See [CGTS], [GNMO], and [MP] for more information.

The running time of the algorithm of Kolliopoulos and Rao [KR] was further improved to $O(n + \varrho k^{O(1)} \log^{O(1)} n)$ by Har-Peled and Mazumdar [HM] by using coresets. Formally, a weighted subset $S \subseteq \mathbb{R}^d$ is a (k, ε) -coreset for the k -median problem if, for any set C of k centers in the \mathbb{R}^d , the price of clustering P using C , and the price of clustering S using C , is the same up to $1 \pm \varepsilon$ (usually a coreset is a subset of the input set P but in this paper we lift this restriction). Har-Peled and Mazumdar [HM] showed that there exists a coreset of P of size $O(k\varepsilon^{-d} \log n)$, and by computing such a coreset quickly and running the algorithm on this coreset, one gets the aforementioned fast approximation algorithm.

k-Means Clustering. Inaba et al. [IKI] observed that the number of Voronoi partitions of k points in \mathbb{R}^d is n^{kd} and can be done exactly in time $O(n^{kd+1})$. They also proposed approximation algorithms for the 2-means clustering problem with time complexity $O(n^{O(d)})$. de la Vega et al. [dVKKR] proposed a $(1 + \varepsilon)$ -approximation algorithm, for high dimensions, with running time $O(g(k, \varepsilon)dn \log^k n)$, where $g(k, \varepsilon) = \exp[(k^3/\varepsilon^8)(\ln(k/\varepsilon)) \ln k]$ (they refer to it as $\ell_2^2 k$ -median clustering). This was improved to a $O(h(k, \varepsilon)dn)$ time algorithm by Kumar et al. [KSS], where $h(k, \varepsilon) = 2^{(k/\varepsilon)^{O(1)}}$ (as such, this algorithm is only appropriate when the data is high-dimensional). Matoušek [M1] proposed a $(1 + \varepsilon)$ -approximation algorithm for the geometric k -means problem with the running time $O(n\varepsilon^{-2k^2d} \log^k n)$. Again, by constructing coresets of size $O(k\varepsilon^{-d} \log n)$, Har-Peled and Mazumdar [HM] presented an algorithm with the running time $O(n + k^{k+2}\varepsilon^{-(2d+1)k} \log^{k+1} n \log^k(1/\varepsilon))$, which is linear for fixed k and ε . Effros and Schulman [ES] showed that there exists a centroid set of size independent of n . A centroid set is a set that contains at least one k -tuple, which forms (approximately) optimal centers for k -means clustering. While the resulting algorithm is slower than the algorithm of Har-Peled and Mazumdar it does hint to the possibility that a coreset of size independent of n should exist for the k -means problem.

Our Results. In light of the aforementioned results, it is natural to ask what is the smallest coreset one can extract for a set of n points in \mathfrak{R}^d and compute approximate clustering using it? In particular, can one compute a coreset of size independent of n ?

In this paper we answer this question positively, by showing a coreset of size $O(k^2/\varepsilon^d)$ for k -median and $O(k^3/\varepsilon^{d+1})$ for k -means. Interestingly, unlike the previous results, while the intuition for the two cases is similar, the proof and construction are fundamentally different. In particular, the coreset construction for the k -means case is slightly easier than the k -median case.

The previous construction of coresets for clustering relied on first computing a set of k centers which were a constant factor approximation to the optimal clustering. Next, using an exponential grid of $O(\log n)$ levels around each center, and snapping the points to this grid (approximating each point with the closest point on the grid) resulted in the required coreset. The correctness of the above coreset follows since the price of snapping the points to the exponential grid is smaller than $\varepsilon v_{\text{opt}}(P, k)$, where $v_{\text{opt}}(P, k)$ is the price of the optimal k -median clustering of P . In the Appendix we show that any such approach of computing a small set C of points such that snapping the points of P to C is “cheap” (i.e., $\leq \varepsilon v_{\text{opt}}(P, k)$) is doomed, as such a set must have size $\Omega(\log n)$. To overcome this, we need to be considerably more careful in picking C , such that the errors introduced by the snapping *cancel each other out*.

To this end, we replace each exponential grid around a center point by a set of $O(1/\varepsilon^{d-1})$ lines. We now snap the points to the lines. We end up with $O(k/\varepsilon^{d-1})$ point sets, each one of them is one-dimensional (although the centers are not necessarily on the line). We compute a coreset for each such line separately, and we take the union of those coresets to form the resulting coreset of the whole set.

To figure out how to pick our coreset on each such line, we first solve the toy problem of computing a coreset for a set of points on a line, where the centers are also on the line. This is done by breaking the line into chunks of small error. (This idea is somewhat similar to the analysis of Effros and Schulman [ES], although our analysis is considerably simpler as we apply it only in one dimension. Effros and Schulman, on the other hand, use a rather involved partition scheme to break their input into d -dimensional chunks with low error.) We then extend it to the case where the centers are not necessarily on the line. We do this analysis for the k -median and k -means cases separately, since the analysis is substantially different.

Note that we reduced the question of computing a d -dimensional coreset to a one- and two-dimensional problem (since the Voronoi diagram on a line of k points in \mathfrak{R}^d , can always be simulated by k points in two dimensions). This reduction considerably simplifies our analysis.

The paper is organized as follows. In Section 2 we present the coreset construction for the k -median case. In Section 3 we handle the k -means case. We conclude in Section 4.

2. Coreset for k -Median

2.1. Preliminaries

For a point set X and a point p , both in \mathfrak{R}^d , let $\mathbf{d}(p, X) = \min_{x \in X} \|xp\|$ denote the *distance of p from X* . For a set B of points on a line in \mathfrak{R}^d , let $\mathcal{I}(B)$ denote the smallest closed segment containing all the points of B .

Definition 2.1 (Weighted Set). A *weighted* point set P is a set of points where every point $p \in P$ is assigned a *weight* w_p , which is a real positive number. We denote by $w(P) = \sum_{p \in P} w_p$ the total weight of the set P . We use (p, w_p) to denote a weighted point at p with weight w_p .

Definition 2.2 (k -Median Clustering). For a weighted point set P with points from \mathfrak{N}^d , with an associated weight function $w : P \rightarrow \mathfrak{N}^+$ and any point set C , we define $v_C(P) = \sum_{p \in P} w_p \cdot \mathbf{d}(p, C)$ as the *price* of the k -median clustering provided by C . Further, let $v_{\text{opt}}(P, k) = \min_{C \subseteq \mathfrak{N}^d, |C|=k} v_C(P)$ denote the price of the *optimal k -median* clustering for P . In the following we abuse notation and, for $x \in \mathfrak{N}^d$, we denote $v_{\{x\}}(P)$ by $v_x(P)$.

Definition 2.3 ((k, ε) -Coreset for k -Median). For a weighted point set $P \subseteq \mathfrak{N}^d$, a subset $S \subseteq \mathfrak{N}^d$ is a (k, ε) -*coreset* of P for k -median clustering if, for all sets C of k points in \mathfrak{N}^d , we have $(1 - \varepsilon)v_C(P) \leq v_C(S) \leq (1 + \varepsilon)v_C(P)$. Note that S is not necessarily a subset of P . We abuse terminology and also use the term *coreset* (without mention of k or ε) to denote any set of weighted points being used to approximate some other weighted point set.

Definition 2.4 (Mean/Center of Mass). For a weighted point set P in \mathfrak{N}^d , let $\bar{\mathbf{m}}(P) = \sum_{p \in P} (w_p/w(P))p$ denote the *mean* of P (this is also known as the center of mass of P). We define the *cumulative error* (or just the *error*) for a weighted point set P in \mathfrak{N} as $\mathcal{E}_v(P) = v_{\bar{\mathbf{m}}}(P) = \sum_{p \in P} w_p \|p\bar{\mathbf{m}}\|$, where $\bar{\mathbf{m}} = \bar{\mathbf{m}}(P)$.

2.2. One Dimension

The basic idea for the coreset construction in one dimension (here, both the points and the centers lie in one dimension) is to break the point set into smaller sets, and use the mean point of every subset as the representative for the coreset. We first prove, in Lemma 2.5, that the cumulative error of a point set bounds the error that it might contribute if we use the mean point as the coreset. In Lemma 2.6 we show that cumulative error is a 2-approximation to the optimal 1-median clustering of a point set. Hence, we can use the mean of a point set as its coreset representative. In Lemma 2.7 we extend this observation to several point sets. Then, in Theorem 2.8, we describe the construction and prove that it works.

Lemma 2.5. *Let P be a set of n weighted points on an oriented line ℓ in \mathfrak{N}^d , and let $\bar{\mathbf{m}} = \bar{\mathbf{m}}(P)$. We have:*

- (i) $\sum_L w_p \|\bar{\mathbf{m}}p\| = \sum_R w_p \|\bar{\mathbf{m}}p\|$, where L (resp. R) are the points of P left (resp. right) of $\bar{\mathbf{m}}$ on ℓ .
- (ii) For a point $q \in \ell$ such that $q \notin \mathcal{I}(P)$, we have that $v_q(P) = w(P)\|q\bar{\mathbf{m}}\|$.
- (iii) For any set of points $C \subseteq \mathfrak{N}^d$, we have $|v_C(P) - v_C(S)| \leq \mathcal{E}_v(P)$, where S is a coreset made out of the single point $\bar{\mathbf{m}}$ with weight $w(P)$.

Proof. (i) Rotate space, such that ℓ becomes the x -axis. Then we have $\sum_{p \in P, x_p < x_{\bar{m}}} w_p(x_{\bar{m}} - x_p) = \sum_{p \in P, x_p \geq x_{\bar{m}}} w_p(x_p - x_{\bar{m}})$, since \bar{m} is the mean point of $P \subseteq \ell$, where x_p denotes the x -coordinate of a point $p \in \mathfrak{R}^d$. Now, $\sum_{p \in L} w_p \|\bar{m}p\| = \sum_{p \in L} w_p(x_{\bar{m}} - x_p) = \sum_{p \in R} w_p(x_p - x_{\bar{m}}) = \sum_{p \in R} w_p \|\bar{m}p\|$.

(ii) Assume that $x_q < x_{\bar{m}}$, and then we have

$$\begin{aligned} v_q(P) &= \sum_{p \in P} w_p \|pq\| = \sum_{p \in P, x_p < x_{\bar{m}}} w_p (\|q\bar{m}\| - \|p\bar{m}\|) + \sum_{p \in P, x_p \geq x_{\bar{m}}} w_p (\|q\bar{m}\| + \|p\bar{m}\|) \\ &= w(P) \|q\bar{m}\| + \left(\sum_{p \in P, x_p < \bar{m}} -w_p \|p\bar{m}\| + \sum_{p \in P, x_p \geq \bar{m}} w_p \|p\bar{m}\| \right) = w(P) \|q\bar{m}\| \end{aligned}$$

by the first claim. The case $x_q > x_{\bar{m}}$ follows by symmetry.

(iii) We have

$$|v_C(P) - v_C(\mathcal{S})| = \left| \sum_{p \in P} w_p (\mathbf{d}(p, C) - \mathbf{d}(\bar{m}, C)) \right| \leq \sum_{p \in P} w_p \|p\bar{m}\| = \mathcal{E}_v(P),$$

since $\mathbf{d}(q, C) - \|pq\| \leq \mathbf{d}(p, C) \leq \mathbf{d}(q, C) + \|pq\|$, for any $p, q \in \mathfrak{R}^d$. \square

Lemma 2.6. *Let $P \subseteq \mathfrak{R}$ be a set of weighted points. Then $\mathcal{E}_v(P) \leq 2v_{\text{opt}}(P, 1)$.*

Proof. The function $v_x(P) = v_x(L_x) + v_x(R_x)$, where L_x and R_x are all the points of P to the left and right of x , respectively. As x increases, the function $v_x(L_x)$ increases, while $v_x(R_x)$ decreases. Those two functions are equal for $x = \bar{m}(P)$. As such, $v_y(P) = v_y(L_y) + v_y(R_y) \geq v_{\bar{m}}(P)/2 = \mathcal{E}_v(P)/2$, for all $y \in \mathfrak{R}$ and $\bar{m} = \bar{m}(P)$. In particular, this holds for the optimal point y that realizes $v_{\text{opt}}(P, 1)$. \square

Lemma 2.7. *Let P be a set of weighted points in \mathfrak{R} , and let P_1, \dots, P_k be a partition of P into k sets. Then $v_{\text{opt}}(P, 1) \geq (\mathcal{E}_v(P_1) + \mathcal{E}_v(P_2) + \dots + \mathcal{E}_v(P_k)) / 2$, where $\mathcal{E}_v(P_i) = v_{\bar{m}(P_i)}(P_i)$.*

Proof. Let $\xi = \text{Median}(P)$ and

$$\begin{aligned} v_{\text{opt}}(P, 1) &= \sum_{p \in P} w_p \|p\xi\| = \sum_{p \in P_1} w_p \|p\xi\| + \sum_{p \in P_2} w_p \|p\xi\| + \dots + \sum_{p \in P_k} w_p \|p\xi\| \\ &\geq v_{\text{opt}}(P_1, 1) + v_{\text{opt}}(P_2, 1) + \dots + v_{\text{opt}}(P_k, 1) \\ &\geq \frac{1}{2}(\mathcal{E}_v(P_1) + \mathcal{E}_v(P_2) + \dots + \mathcal{E}_v(P_k)), \end{aligned}$$

by Lemma 2.6. \square

Theorem 2.8. *Let P be a weighted point set in \mathfrak{R} , k and $\varepsilon > 0$ parameters. Then there exists a (k, ε) -coreset \mathcal{S} of P of size $O(k/\varepsilon)$.*

Proof. Assume that we have an approximation V , such that $v_{\text{opt}}(P, k) \leq V \leq c \cdot v_{\text{opt}}(P, k)$, where c is a constant (this can be done efficiently in linear time for small k [HM]). We scan the points from left to right and group them into batches with the cumulative error equal to $\varphi = (\varepsilon/10ck)V$. This is done by allowing the first and last points in the batch to be a fraction of a point of P (i.e., a point p of P might appear in two consecutive batches, as two points with total weight w_p). The last batch is of weight $\leq \varphi$. Observe that $\varphi \geq (\varepsilon/10ck)v_{\text{opt}}(P, k)$. Let $\mathcal{B} = \{\mathcal{B}_\infty, \dots, \mathcal{B}_\top\}$ denote the resulting batches.

It is now straightforward to verify that $|\mathcal{B}| = O(k/\varepsilon)$. Indeed, let C_{opt} be the set of k medians realizing $v_{\text{opt}}(P, k)$. Since P is a one-dimensional point set, there are at most $k - 1$ batches that are being served by more than one center in C_{opt} . For any other batch $B \in \mathcal{B}$, B is being served by a single center of C_{opt} . We call this set of batches \mathcal{B}' . Now, $v_{\text{opt}}(P, k) = \sum_i v_{\text{opt}}(P_i, 1)$ where P_i is the subset of P served by center $c_i \in C_{\text{opt}}$. Now, using Lemma 2.7 we have $\sum_{B_i \in \mathcal{B}'} \mathcal{E}_v(B_i)/2 \leq \sum_i v_{\text{opt}}(P_i, 1) = v_{\text{opt}}(P, k)$. Since we have $\mathcal{E}_v(B_i) = \varphi$ except for the last batch, the number of batches in \mathcal{B}' is bounded by $O(1 + v_{\text{opt}}(P, k)/\varphi)$. Hence, $|\mathcal{B}|$ is $O(k + v_{\text{opt}}(P, k)/\varphi) = O(k/\varepsilon)$.

Next, for the coresets construction, we set $\bar{m}(B_i)$ to be the representative point for B_i with weight $w(B_i)$. Let \mathcal{S} be the resulting coreset. We claim that this is a (k, ε) -coreset. Indeed, consider any point set $C = \{x_1, x_2, \dots, x_k\}$. For a point $x_i \in C$, let I_i denote the interval on the real line that it serves. For a batch B , let $\mathcal{I}(B)$ denote the smallest interval containing B . If a batch $B \subseteq I_i$, and $x_i \notin \mathcal{I}(B)$, then, by Lemma 2.5, we have $v_{x_i}(B) = w(B) \cdot \|\bar{m}(B)x_i\|$. Namely, the contribution of the points of B to $v_C(P)$ and $v_C(\mathcal{S})$ are identical.

Thus, the only batches that might contribute to the error are the ones that contain an endpoint of I_1, \dots, I_k (there are at most $k - 1$ such batches), and batches that contain a point of C in their interior (there are at most k such batches). By Lemma 2.5(iii), every such batch B contributes at most $\mathcal{E}_v(B)$ to the overall error. Let B'_1, \dots, B'_{2k-1} be those “problematic” batches. We have that

$$|v_C(P) - v_C(\mathcal{S})| \leq \sum_{i=1}^{2k-1} \mathcal{E}_v(B'_i) \leq (2k - 1) \cdot \varphi \leq \varepsilon v_{\text{opt}}(P, k). \quad \square$$

2.3. Extending to Higher Dimensions

In this section we show how to construct the coreset for a set P of n points in \mathbb{R}^d . We need the following technical lemma.

Lemma 2.9. *Let $\mathbf{c} = (0, \alpha)$ be a point in the plane, let L and R be two weighted sets of points on the positive portion of the x -axis such that all the points of L have a smaller x -axis value than the points of R , and let \mathfrak{l} and \mathfrak{r} be two points on the x -axis such that $v_{\mathfrak{l}}(L) = v_{\mathfrak{r}}(R)$. Furthermore, let $\mathcal{S}_L = \{\mathfrak{l}, w(L)\}$ and $\mathcal{S}_R = \{\mathfrak{r}, w(R)\}$ be the coresets formed by assigning sum of the weights on the points in sets L and R to \mathfrak{l} and \mathfrak{r} , respectively. Also, let $\mathcal{E} = v_{\mathbf{c}}(L) + v_{\mathbf{c}}(R) - v_{\mathbf{c}}(\mathcal{S}_L) - v_{\mathbf{c}}(\mathcal{S}_R)$ be the error caused*




Fig. 1. Case (i) of Lemma 2.9.

by using the coresets \mathcal{S}_L and \mathcal{S}_R instead of the sets L and R , respectively, in relation to the center \mathbf{c} . Then:

- (i) If $x_l \leq x_{l'} \leq x_{r'} \leq x_r$, for all $l' \in L$ and $r' \in R$, then $\mathcal{E} \leq 0$. See Fig. 1.
- (ii) If $x_{l'} \leq x_l \leq x_r \leq x_{r'}$, for all $l' \in L$ and $r' \in R$, then $\mathcal{E} \geq 0$.

Proof. (i) For two points p, q on the x -axis, such that $x_p \leq x_q$, let $e(p, q) = \|q\mathbf{c}\| - \|p\mathbf{c}\|$. In particular, for any four points a, b, c, d on the x -axis, such that $x_a \leq x_b \leq x_c \leq x_d$, we have $e(a, b)/\|ab\| \leq e(c, d)/\|cd\|$. This follows since the function $f(x) = \|\mathbf{c} - (x, 0)\|$ is a convex function with a positive second derivative, as can be easily verified. In particular, for any $a \leq b$ we have $f'(a) \leq e(a, b)/\|ab\| \leq f'(b)$. Thus, for a point z on the real line between R and L , we have

$$\begin{aligned}
 \mathcal{E} &= v_{\mathbf{c}}(L) + v_{\mathbf{c}}(R) - v_{\mathbf{c}}(\mathcal{S}_L) - v_{\mathbf{c}}(\mathcal{S}_R) \\
 &= \sum_{p \in L} w_p (\|\mathbf{c}p\| - \|\mathbf{c}l\|) + \sum_{p \in R} w_p (\|\mathbf{c}p\| - \|\mathbf{c}r\|) \\
 &= \sum_{p \in L} w_p e(l, p) - \sum_{p \in R} w_p e(p, r) = \sum_{p \in L} w_p \|p\| \frac{e(l, p)}{\|p\|} - \sum_{p \in R} w_p \|r\| \frac{e(p, r)}{\|r\|} \\
 &\leq \sum_{p \in L} w_p \|p\| f'(z) - \sum_{p \in R} w_p \|r\| f'(z) = f'(z)(v_l(L) - v_r(R)) = 0,
 \end{aligned}$$

since $e(l, p)/\|p\| \leq f'(z) \leq e(p, r)/\|r\|$ for any $p \in L$ and $q \in R$.

The second claim follows by similar argumentation. \square

2.3.1. *Construction.* The following technical lemma states the existence of an ε -net for the sphere. See Lemma 13.1.1 of [M2] for details.

Lemma 2.10. *There exists a set of points Q on a sphere of unit radius in d -dimensions (S^{d-1}) centered at the origin with the following properties:*

- (i) Q has $O(\varepsilon^{-(d-1)})$ points, and
- (ii) $\forall p$ that lie on the unit radius ball, $\exists q \in Q$ such that $\|pq\| \leq \varepsilon$. Furthermore, Q can be computed in $O(\varepsilon^{-(d-1)})$ time.

We compute a set $\mathbf{C} = \{\mathbf{c}_1, \dots, \mathbf{c}_k\}$ of centers which is a c -approximation to $v_{\text{opt}}(P, k)$; namely, $v_{\mathbf{C}}(P) \leq c \cdot v_{\text{opt}}(P, k)$, where c is a constant (as was done in [HM]). Now we divide the set of points P into k sets based on which point in \mathbf{C} is nearest to them. This gives us a partition of P into k subsets P_1, P_2, \dots, P_k where P_i is closest to $\mathbf{c}_i \in \mathbf{C}$. Around each of the points $\mathbf{c}_i \in \mathbf{C}$ we place a fan \mathcal{L}_i of lines passing through it. This is




Fig. 2. Computing the coresets.

done by taking a unit sphere centered at c_i , and placing an $\varepsilon/(3c)$ -net N_{c_i} on this sphere, using Lemma 2.10. For every $p \in N_{c_i}$, we generate the line spanning the segment $c_i p$.

For each point of $p \in P_i$, let $l_i(p)$ be its closest line in \mathcal{L}_i , and let p' be the projection of p into $l_i(p)$. Let P' be the set of these snapped (projected) points. Also, let P_ℓ be the set of points projected onto the line ℓ . Next, we compute a coresets \mathcal{S}_ℓ , for each line $\ell \in \mathcal{L}_\infty \cup \dots \cup \mathcal{L}_\parallel$, using the one-dimensional method. Namely, we scan every line ℓ , and break the point set, P_ℓ , along it into batches, such that for each batch B (except the last one), we have $\mathcal{E}_v(B) = (\varepsilon A_\ell)/(20ck)$, where A_ℓ is a c -approximation to $v_{\text{opt}}(P_\ell, k)$ (again, allowing a boundary point to appear in two batches with a fractional weight). See Fig. 2. (In the following, for the sake of simplicity of exposition, we ignore the fact that a batch contains weighted points. This is a minor technicality, and it can be easily handled.)

Hence, we get $O(k/\varepsilon)$ points selected in the coresets on each of the lines through c_i , and hence $O(k/\varepsilon^d)$ coresets points for each P_i . Thus, the total number of points in the coresets \mathcal{S} is $O(k^2/\varepsilon^d)$.

2.3.2. Correctness

Observation 2.11. Let p be a point of P , and let c_i be its nearest point in C , let p' be the corresponding point in P' . We have $\|pp'\| \leq \|pc_i\|\varepsilon/(3c)$.

Lemma 2.12. Let P be a set of points on a line ℓ , and let \mathcal{S}_ℓ be the coresets constructed for it. Also, let C be a set of k points in \mathbb{R}^d . Then $|v_C(P) - v_C(\mathcal{S}_\ell)| \leq (\varepsilon/3)v_{\text{opt}}(P, k)$.

Proof. The proof follows the one-dimensional case (i.e., Theorem 2.8), although the analysis is somewhat more involved. We rotate and translate all the points so that the line ℓ coincides with the x -axis. Let $C = \{c_1, \dots, c_k\}$, and let c'_1, \dots, c'_k denote the projection of c_1, \dots, c_k into ℓ , respectively. Next, we partition the line into k intervals

$\mathcal{I}_1, \dots, \mathcal{I}_k$, such that \mathcal{I}_i is the portion of ℓ closer to \mathbf{c}_i than any other point of C (note that the points of C are not necessarily on ℓ). Then every point in the coreset \mathcal{S}_ℓ corresponds to a subset (i.e., batch) of P . By construction, all the batches have the same cumulative error (except the last batch, which might have a smaller cumulative error). In particular, for any batch B we have that the cumulative error $\mathcal{E}_v(B) \leq (\varepsilon/20k)v_{\text{opt}}(P, k)$.

Let \widehat{B} be the union of the set of all batches which are served by more than one center of C and the set of all batches B such that the interval $\mathcal{I}(B)$ contains the projection of a center point of C to ℓ . We also add the last batch on ℓ to \widehat{B} . Clearly, $|\widehat{B}| \leq 2k$. Let $U = \bigcup_{B \in \widehat{B}} B$ be the points of P in \widehat{B} , and let $\mathcal{S}_U \subseteq \mathcal{S}_\ell$ be the corresponding coreset. It follows that the total error contributed by the points of U is

$$E^* = |v_C(U) - v_C(\mathcal{S}_U)| \leq \sum_{B \in \widehat{B}} \mathcal{E}_v(B) \leq 2k \frac{\varepsilon}{20k} v_{\text{opt}}(P, k) \leq \frac{\varepsilon}{10} v_{\text{opt}}(P, k),$$

by Lemma 2.5(iii).

Let us fix a center $\mathbf{c} \in C$, and let \mathcal{I} be its Voronoi cell on ℓ . Next, consider the set R (resp. L) of the batches to the right (resp. left) of \mathbf{c}' that lie in its interval \mathcal{I} . Let B^1, B^2, \dots, B^t denote the batches of R sorted from left to right. Furthermore, let L^i and R^i be the set of points of B^i to the left and right of the mean $\bar{m}(B^i)$, respectively, for $i = 1, \dots, t$. Finally, let \mathcal{S}_l^i and \mathcal{S}_r^i denote the coresets formed by placing a point at $\bar{m}(B^i)$ with weight $w(L^i)$ and $w(R^i)$, respectively. Let $\mathcal{S}^i = \mathcal{S}_l^i \cup \mathcal{S}_r^i$ be the one point coreset placed at $\bar{m}(B^i)$ with weight $w(L^i) + w(R^i)$. Let \mathcal{S}_R denote the resulting coreset for all the batches in R . Let P_R denote the points in R . By Lemma 2.9(ii) we have that

$$v_C(B^i) - v_C(\mathcal{S}^i) = v_C(L^i) - v_C(\mathcal{S}_l^i) + v_C(R^i) - v_C(\mathcal{S}_r^i) \geq 0.$$

Thus, the error contributed by the coreset of R is

$$E = v_C(P_R) - v_C(\mathcal{S}_R) = \sum_{i=1}^t (v_C(B^i) - v_C(\mathcal{S}^i)) \geq 0.$$

On the other hand, by Lemma 2.9(i) we have

$$E' = \sum_{i=1}^{t-1} (v_C(R^i) - v_C(\mathcal{S}_r^i) + v_C(L^{i+1}) - v_C(\mathcal{S}_l^{i+1})) \leq 0.$$

Thus,

$$\begin{aligned} 0 \leq E &= E' + v_C(L^1) - v_C(\mathcal{S}_l^1) + v_C(R^t) - v_C(\mathcal{S}_r^t) \leq \mathcal{E}_v(B^1) + \mathcal{E}_v(B^t) \\ &= 2(\varepsilon/20k)v_{\text{opt}}(P, k). \end{aligned}$$

Namely, the total error induced by using the coreset for batches in R is bounded by $2(\varepsilon/20k)v_{\text{opt}}(P, k)$. By symmetry, the same holds for the batches in L . Thus, the total error induced by such batches is at most $k \cdot 2 \cdot 2(\varepsilon/20k)v_{\text{opt}}(P, k) \leq (\varepsilon/5)v_{\text{opt}}(P, k)$. Thus, we have

$$|v_C(P) - v_C(\mathcal{S}_\ell)| \leq (\varepsilon/5)v_{\text{opt}}(P, k) + (\varepsilon/10)v_{\text{opt}}(P, k) \leq (\varepsilon/3)v_{\text{opt}}(P, k)$$

as desired. \square

Theorem 2.13. *Let P be a point set of n points in \mathfrak{R}^d , and let \mathcal{S} be the coresset constructed for it in Section 2.3.1. Then \mathcal{S} is a weighted set of size $O(k^2/\varepsilon^d)$ and it is a (k, ε) -coreset of P for k -median clustering.*

Proof. By Lemma 2.12 we know that the error between the distance of any set C of size k and the snapped points P' on the fans can be well approximated using the coresset. Furthermore, the error introduced by the snapping is bounded by

$$E = \sum_{p \in P} \|pp'\| \leq \sum_{i=1}^k \left(\sum_{p \in P_i} \left(\frac{\varepsilon}{3c} \right) \|px_i\| \right) \leq (\varepsilon/g)3v_{\text{opt}}(P, k),$$

by Observation 2.11.

So, $|\nu_C(P) - \nu_C(P')| \leq \sum_{p \in P} \|pp'\| \leq (\varepsilon/3)v_{\text{opt}}(P, k)$. Thus, for any set C of k points, we have

$$\begin{aligned} |\nu_C(P) - \nu_C(\mathcal{S})| &\leq |\nu_C(P) - \nu_C(P')| + |\nu_C(P') - \nu_C(\mathcal{S})| \\ &\leq (\varepsilon/3)v_{\text{opt}}(P, k) + (\varepsilon/3)v_{\text{opt}}(P', k) \\ &\leq (\varepsilon/3)v_{\text{opt}}(P, k) + (\varepsilon/3)v_{\text{opt}}(P, k)(1 + \varepsilon/3) \leq \varepsilon v_{\text{opt}}(P, k), \end{aligned}$$

by Lemma 2.12. □

3. Coreset for k -Means

3.1. Preliminaries

Definition 3.1 (k -Means Clustering). Let $\mu_C(P) = \sum_{p \in P} w_p \cdot (\mathbf{d}(p, C))^2$ denote the price of the k -means clustering of P as provided by the set of centers C . Let $\mu_{\text{opt}}(P, k) = \min_{C \subseteq \mathfrak{R}^d, |C|=k} \mu_C(P)$ denote the price of the *optimal k -means clustering* of P . Again, for $x \in \mathfrak{R}^d$, we use $\mu_x(P)$ to denote the quantity $\mu_{\{x\}}(P)$.

Definition 3.2 ((k, ε) -Coreset for k -Means). A weighted set \mathcal{S} is a (k, ε) -coreset of P for k -means clustering if, for any set C of k points in \mathfrak{R}^d , we have $(1 - \varepsilon)\mu_C(P) \leq \mu_C(\mathcal{S}) \leq (1 + \varepsilon)\mu_C(P)$.

Definition 3.3. For a point set P , the *error* of P is $\widehat{\mathcal{E}}(P) = \sum_{p \in P} \|p\bar{\mathbf{m}}\|^2$, where $\bar{\mathbf{m}} = \bar{\mathbf{m}}(P)$.

3.2. The One-Dimensional Case

3.2.1. Construction. Let P be a given set of n points on the real line. The procedure is similar to the k -median case except for the fact that just picking the mean point of each batch as its representative does not suffice and we will need two appropriately placed representative points for each batch. We consider the points from left to right and group

them into batches, such that a batch B has $\widehat{\mathcal{E}}(B) \leq \xi$, and for two consecutive batches B and B' we have $\widehat{\mathcal{E}}(B \cup B') \geq \xi$, where $\xi \leq (\varepsilon^2/100k^2)\mu_{\text{opt}}(P, k)$. Arguing in a similar fashion to the k -median case, it follows that the number of batches is $O(k^2/\varepsilon^2)$. Let $\mathcal{B}(P)$ denote the resulting set of batches.

Lemma 3.4. *Let B be a set of points on a line. There exist two weighted points (q_1, w_1) and (q_2, w_2) both lying completely within $\mathcal{I}(B)$, such that*

- $w_1 + w_2 = |B|$,
- $(q_1 w_1 + q_2 w_2)/(w_1 + w_2) = \bar{m}$, where $\bar{m} = \bar{m}(P)$,
- and $w_1 \|q_1 \bar{m}\|^2 + w_2 \|q_2 \bar{m}\|^2 = \sum_{p \in B} \|p \bar{m}\|^2$.

Let $\mathcal{T}(B) = \{(q_1, w_1), (q_2, w_2)\}$ denote this coreset.

Proof. We construct these weighted points through a sequence of steps. Let the leftmost point in B be p_l and the rightmost point be p_r .

- For every point $p \in B$ to the right of \bar{m} we add a point at the rightmost extreme of B with weight $\|p \bar{m}\|/\|p_r \bar{m}\|$. Clearly, $(\|p \bar{m}\|/\|p_r \bar{m}\|)\|p_r \bar{m}\|^2 \geq \|p \bar{m}\|^2$. Similarly for every point $p \in B$ to the left of \bar{m} we add a point at the leftmost extreme of B with weight $\|p \bar{m}\|/\|p_l \bar{m}\|$. This results in two weighted points p_l and p_r . Furthermore, we have $\bar{m}(\{p_l, p_r\}) = \bar{m}$, $\widehat{\mathcal{E}}(\{p_l, p_r\}) \geq \widehat{\mathcal{E}}(B)$, and $w_{p_l} + w_{p_r} \leq |B|$.
- Now we scale up the weights so that $w_{p_l} + w_{p_r} = |B|$. Note that this does not change the mean, and only increases $\widehat{\mathcal{E}}(\{p_l, p_r\})$.
- Finally, consider the scaled set

$$C(t) = \{(p_l \cdot t + (1-t)\bar{m}, w_{p_l}), (p_r \cdot t + (1-t)\bar{m}, w_{p_r})\}.$$

Clearly, $C(t)$ for $t \in [0, 1]$ has $\bar{m}(C(t)) = \bar{m}$. Furthermore, $C(1)$ is just the current two weighed points, and $C(0)$ is just one point at \bar{m} . Thus, pick $t^* \in [0, 1]$, such that $\widehat{\mathcal{E}}(C(t^*)) = \widehat{\mathcal{E}}(B)$. This is possible, since $\widehat{\mathcal{E}}(C(1)) \geq \widehat{\mathcal{E}}(B)$.

Clearly, $C(t^*)$ is the required coreset. \square

Let $\mathcal{S}(P) = \bigcup_{B \in \mathcal{B}(P)} \mathcal{T}(B)$ be the constructed coreset for P .

3.2.2. *Correctness.* The following claim is well known (Lemma 2.1 in [KMN⁺]).

Lemma 3.5. *Let B be a set of points in \mathfrak{R}^d , then for any $q \in \mathfrak{R}^d$ we have $\mu_q(B) = |B| \|q \bar{m}\|^2 + \widehat{\mathcal{E}}(B)$.*

Lemma 3.6. *Let B be a set of points in \mathfrak{R}^d lying on a common line, let $\mathcal{T} = \mathcal{T}(B)$ be the coreset constructed for it by the algorithm of Lemma 3.4, and let q be any point in \mathfrak{R}^d . Then $\mu_q(B) = \mu_q(\mathcal{T})$.*

Proof. We have $\mu_q(B) = |B| \|q \bar{m}\|^2 + \widehat{\mathcal{E}}(B)$, and $\mu_q(\mathcal{T}) = w(\mathcal{T}) \|q \bar{m}(\mathcal{T})\|^2 + \widehat{\mathcal{E}}(\mathcal{T}) = |B| \|q \bar{m}\|^2 + \widehat{\mathcal{E}}(B)$, by Lemma 3.4. Thus, $\mu_q(B) = \mu_q(\mathcal{T})$. \square

Theorem 3.7. *Let P be a set of n points in \mathfrak{R}^d such that the points of P all lie on a line ℓ , and let \mathcal{S} be the coresset constructed for it in Section 3.2.1. Then \mathcal{S} is a $(k, \varepsilon/3)$ -coreset for k -means clustering of P , for any set of k centers in \mathfrak{R}^d .*

Proof. The proof is similar to the k -median case. We first rotate space such that ℓ is on the x -axis. Let $\mathbf{C} = \{c_1, \dots, c_k\}$ be a set of k centers, $\mu_{\mathbf{C}} = \mu_{\mathbf{C}}(P)$ and $\mu'_{\mathbf{C}} = \mu_{\mathbf{C}}(\mathcal{S})$. Let $\mathcal{I}_1, \dots, \mathcal{I}_k$ be a partition of the line into intervals such that \mathcal{I}_i is the loci of points closest to c_i out of all the centers in \mathbf{C} , for $i = 1, \dots, k$. Note that the batches of $\mathcal{B}(P)$, and their corresponding coresset points, that lie completely within \mathcal{I}_i , do not contribute to the overall error $|\mu_{\mathbf{C}} - \mu'_{\mathbf{C}}|$ by Lemma 3.6.

Thus, the only problematic batches are the ones that contain an endpoint of $\mathcal{I}_1, \dots, \mathcal{I}_k$. There are at most $k - 1$ such batches. Let B be one such batch. Assume that the interval $\mathcal{I}(B)$ intersects $\mathcal{I}_1, \dots, \mathcal{I}_t$, and let $V_i = \mathcal{I}_i \cap B$, for $i = 1, \dots, t$. Let $\bar{\mathbf{m}} = \bar{\mathbf{m}}(B)$ and let $\mathcal{S}_B = \mathcal{T}(B)$. We partition \mathcal{S}_B into portions corresponding to the sets V_1, \dots, V_t . Formally, \mathcal{S}_i is a set of the two points of \mathcal{S}_B , re-weighted such that $w(\mathcal{S}_i) = |V_i|$, for $i = 1, \dots, t$. We have, by Lemma 3.5, that

$$\begin{aligned} \mu_{\mathbf{C}}(\mathcal{S}_B) &= \sum_i \mu_{\mathbf{C}}(\mathcal{S}_i) \leq \sum_i \mu_{c_i}(\mathcal{S}_i) = \sum_i (\widehat{\mathcal{E}}(\mathcal{S}_i) + |V_i| \|c_i \bar{\mathbf{m}}\|^2) \\ &= \sum_i \left(\widehat{\mathcal{E}}(\mathcal{S}_i) + \sum_{p \in V_i} \|c_i \bar{\mathbf{m}}\|^2 \right) = \widehat{\mathcal{E}}(\mathcal{S}_B) + \sum_i \sum_{p \in V_i} \|c_i \bar{\mathbf{m}}\|^2. \end{aligned}$$

Since $\|c_i p\| - \|p \bar{\mathbf{m}}\| \leq \|c_i \bar{\mathbf{m}}\|$, we have

$$\begin{aligned} \sum_i \sum_{p \in V_i} \|c_i \bar{\mathbf{m}}\|^2 &\geq \sum_i \sum_{p \in V_i} (\|c_i p\| - \|p \bar{\mathbf{m}}\|)^2 \\ &\geq \sum_i \sum_{p \in V_i} \|c_i p\|^2 - 2 \sum_i \sum_{p \in V_i} \|c_i p\| \cdot \|p \bar{\mathbf{m}}\| + \sum_i \sum_{p \in V_i} \|p \bar{\mathbf{m}}\|^2 \\ &\geq \mu_{\mathbf{C}}(B) - 2 \sum_i \sum_{p \in V_i} \|c_i p\| \cdot \|p \bar{\mathbf{m}}\| + \widehat{\mathcal{E}}(B). \end{aligned}$$

We also have $\|c_i p\| + \|p \bar{\mathbf{m}}\| \geq \|c_i \bar{\mathbf{m}}\|$ and so

$$\begin{aligned} \sum_i \sum_{p \in V_i} \|c_i \bar{\mathbf{m}}\|^2 &\leq \sum_i \sum_{p \in V_i} (\|c_i p\| + \|p \bar{\mathbf{m}}\|)^2 \\ &\leq \mu_{\mathbf{C}}(B) + 2 \sum_i \sum_{p \in V_i} \|c_i p\| \cdot \|p \bar{\mathbf{m}}\| + \widehat{\mathcal{E}}(B). \end{aligned}$$

We conclude that

$$\left| \sum_i \sum_{p \in V_i} \|c_i \bar{\mathbf{m}}\|^2 - \mu_{\mathbf{C}}(B) \right| \leq 2 \sum_i \sum_{p \in V_i} \|c_i p\| \cdot \|p \bar{\mathbf{m}}\| + \widehat{\mathcal{E}}(B).$$

This gives us

$$\begin{aligned}
|\mu_{\mathcal{C}}(\mathcal{S}_B) - \mu_{\mathcal{C}}(B)| &= \left| \widehat{\mathcal{E}}(\mathcal{S}_B) + \sum_i \sum_{p \in V_i} \|\mathbf{c}_i \bar{\mathbf{m}}\|^2 - \mu_{\mathcal{C}}(B) \right| \\
&\leq 2 \sum_i \sum_{p \in V_i} \|\mathbf{c}_i p\| \cdot \|p \bar{\mathbf{m}}\| + 2\widehat{\mathcal{E}}(B) \\
&\leq 2\widehat{\mathcal{E}}(B) + 2 \sqrt{\sum_i \sum_{p \in V_i} \|\mathbf{c}_i p\|^2} \sqrt{\sum_i \sum_{p \in V_i} \|p \bar{\mathbf{m}}\|^2} \\
&\leq 2\widehat{\mathcal{E}}(B) + 2\sqrt{\mu_{\mathcal{C}}(B)} \sqrt{\widehat{\mathcal{E}}(B)},
\end{aligned}$$

by the Cauchy–Swartz inequality. By construction $\widehat{\mathcal{E}}(B) \leq (\varepsilon^2/100k^2)\mu_{\text{opt}}(P, k)$. Thus,

$$\begin{aligned}
|\mu_{\mathcal{C}}(\mathcal{S}_B) - \mu_{\mathcal{C}}(B)| &\leq 2 \frac{\varepsilon^2}{100k^2} \mu_{\text{opt}}(P, k) + 2 \frac{\varepsilon}{10k} \sqrt{\mu_{\mathcal{C}}(B) \mu_{\text{opt}}(P, k)} \\
&\leq 2 \frac{\varepsilon^2}{100k^2} \mu_{\text{opt}}(P, k) + 2 \frac{\varepsilon}{10k} \cdot \frac{\mu_{\mathcal{C}}(B) + \mu_{\text{opt}}(P, k)}{2} \\
&\leq \frac{\varepsilon}{5k} \mu_{\text{opt}}(P, k) + \frac{\varepsilon}{10k} \mu_{\mathcal{C}}(B).
\end{aligned}$$

Since there are $k - 1$ border batches, we conclude that

$$|\mu_{\mathcal{C}}(\mathcal{S}) - \mu_{\mathcal{C}}(P)| \leq \frac{\varepsilon}{5} \mu_{\text{opt}}(P, k) + \frac{\varepsilon}{10} \mu_{\mathcal{C}}(P) \leq \frac{\varepsilon}{3} \mu_{\mathcal{C}}(P),$$

as required. \square

3.3. Extending to Higher Dimensions

Again we use a similar approach to the one we used for the k -median case. We calculate an approximation $\mu_{\text{opt}}(P, k) \leq A \leq c \cdot \mu_{\text{opt}}(P, k)$, where $c > 1$ is a constant. Then we partition the point set P into sets P_1, P_2, \dots, P_k with P_i consisting of points in the area of control of $\mathbf{c}_i \in A$. Then we draw $O(1/\varepsilon^{d-1})$ lines through each of the centers of A as before and snap the points of P_i onto the closest line around \mathbf{c}_i . We compute a coreset for every line using the algorithm of Section 3.2.1. This gives us $O(k^2/\varepsilon^2)$ points selected for the coreset on every line, thus the total size of the resulting coreset \mathcal{S} is $O(k^3/\varepsilon^{d+1})$.

This set is indeed a coreset. The proof is an easy extension of the one-dimensional case. The snapping into the lines introduces a multiplicative error smaller than $\varepsilon/3$. The coreset construction introduces an error of similar magnitude, by Theorem 3.7. Since this is a straightforward extension of our previous discussion, we omit any further details.

Theorem 3.8. *Given a set P of n points in \mathfrak{R}^d , one can compute a (k, ε) -coreset for P for k -means clustering of size $O(k^3/\varepsilon^{d+1})$.*

3.3.1. Centroid Set. Given a set P of n points in \mathfrak{R}^d , a set $\mathcal{D} \subseteq \mathfrak{R}^d$ is a (k, ε) -approximate centroid set for P , if there exists a subset $C \subseteq \mathcal{D}$ of size k , such that $\mu_C(P) \leq (1 + \varepsilon)\mu_{\text{opt}}(P, k)$.

Matoušek showed that there exists an ε -approximate centroid set of size $O(n\varepsilon^{-d} \log(1/\varepsilon))$ [M1]. Interestingly enough, his construction is weight insensitive. In particular, using a $(k, \varepsilon/2)$ -coreset \mathcal{S} in his construction results in an ε -approximate centroid set of size $O(|\mathcal{S}| \varepsilon^{-d} \log(1/\varepsilon))$.

Theorem 3.9. *Given a set P of n points in \mathfrak{R}^d , one can compute a (k, ε) -centroid set for P for k -means clustering of size $O((k^3/\varepsilon^{2d+1}) \log(1/\varepsilon))$.*

Theorem 3.9 slightly improves (as far as the dependency of k is concerned) over the result of Effros and Schulman [ES] that showed that there exists a centroid set of size $O(\varepsilon^{-d-1}(k^4 + k^2\varepsilon^{-2}))$. We conjecture that the dependency on ε in the bound on the coreset size in Theorem 3.9 can be further improved by constructing the centroid set for each line separately. Since this would lead to only minor improvements over the result of Effros and Schulman [ES], we do not investigate this direction any further.

3.3.2. Running Time. The running time of the resulting $(1 + \varepsilon)$ -approximate k -means clustering algorithm is $O(n + \text{poly}_d(k, \log n, 1/\varepsilon) + (k^3/\varepsilon^{d+1})((k^3/\varepsilon^{2d+1}) \log(1/\varepsilon))^{k+1})$, where $\text{poly}_d(\dots)$ is a polynomial with degree that depends on d . This follows easily by performing our coreset construction on the coreset generated by the construction of Har-Peled and Mazumdar [HM], and we omit the easy but tedious details. This compares quite favorably with the Effros and Schulman [ES] algorithm, that has the running time

$$O\left(k^2 n \log \log n + \frac{k^8}{\varepsilon^{3(d+1)}} n + \frac{1}{\varepsilon^{d+1}} \left(\frac{k^4}{\varepsilon^{d+1}}\right)^{k+2}\right),$$

which has worse dependency on n and k . Unfortunately, it does not improve over the algorithm of Har-Peled and Mazumdar [HM], which has the running time $O(n + k^{k+2} \varepsilon^{-(2d+1)k} \log^{k+1} n \log^k(1/\varepsilon))$, which already has linear running time for small values of k and better dependency on k .

4. Conclusions

In this paper we showed the existence of small coresets for the k -means and k -median clustering in \mathfrak{R}^d , with size independent of n . We believe that this result is quite surprising.

Our techniques do not yield any significant improvement in performance over the approximation algorithms of Har-Peled and Mazumdar [HM]. As mentioned in the Introduction, the results in this paper imply algorithms with the running time $O(n + \text{poly}(k, \log n, 1/\varepsilon) + \text{func}(k, \varepsilon))$, where poly denotes a polynomial, and $\text{func}(k, \varepsilon)$ denotes a function that depends only on k and ε (and the dimension d). This, however, improves upon the results of Har-Peled and Mazumdar [HM] only for a very narrow interval of values of k in the k -means case.

One improvement implied by our work is that one can stream points for $(1 + \varepsilon)$ -approximate k -median clustering in \mathfrak{R}^d using $O((k^2/\varepsilon^d) \log^{d+1} n)$ space, using the standard techniques which are also used by Har-Peled and Mazumdar [HM] (this improves by one log factor over the best bound implied by their techniques). The new result for k -means coreset does not imply any improvement for the k -means case.

At this point, there are numerous problems for further research. In particular:

1. Can the running time of approximate k -means clustering be improved to be similar to the k -median bounds? Can one do FPTAS for k -median and k -means (in both k and $1/\varepsilon$)? Currently, we can only compute the (k, ε) -coreset in fully polynomial time, but cannot extract the approximation itself from it.
2. Does a coreset exist for the problems of k -median and k -means clustering with only polynomial dependency on the dimension and no dependency on n ? There are some partial relevant results [BHI].
3. Can one improve the bounds on the size of the coresets for k -median and k -means clustering?

Appendix. Lower Bound on Coreset If Bounding Snapping Error

Let P be a set of n points in \mathfrak{R}^d . The previous construction of coresets for k -means and k -median clustering by Har-Peled and Mazumdar [HM], worked by finding a set \mathcal{S} such that $v_{\mathcal{S}}(P) \leq \varepsilon v_{\text{opt}}(P, k)$. This property by itself is sufficient to guarantee that \mathcal{S} is a (k, ε) -coreset for P . Surprisingly, the following theorem shows that, in the worst case, any set with this property must be large (i.e., size dependent on n).

Theorem A.1. *For any $\varepsilon < \frac{1}{32}$, there exists a set P of n points in \mathfrak{R} such that, for any set \mathcal{S} , if $v_{\mathcal{S}}(P) \leq \varepsilon v_{\text{opt}}(P, 1)$, then $|\mathcal{S}|$ is $\Omega(\varepsilon^{-1} \log n)$.*

Proof. Consider the $n = 2^t$ points in P placed on the real line in the following way. There are $n/2^i$ points placed uniformly in the intervals $\mathcal{I}_i = (-2^{i+1}, -2^i) \cup (2^i, 2^{i+1})$, for $i = 1, 2, \dots, t = \lg n$. Also, place a single point on the origin forming \mathcal{I}_0 . Now, let \mathcal{S} be any weighted coreset for the points of P . Let s_i denote the number of points in $\mathcal{S} \cap \mathcal{I}_i$. Each point in \mathcal{I}_i can be served by (snapped to) either one of these s_i points, the outermost points of \mathcal{I}_{i-1} or the innermost points of \mathcal{I}_{i+1} . Hence, there are at most $s_i + 4$ representatives that can serve the points of \mathcal{I}_i .

Now consider the contribution of the points of \mathcal{I}_i to $v_{\mathcal{S}}(P)$. We derive a lower bound on this contribution by considering the two intervals of \mathcal{I}_i to be a single interval of length 2^{i+1} , which are served using $s_i + 4$ centers. Clearly, the minimum is achieved when those centers serve intervals of equal length, which is $L_i = 2^{i+1}/(s_i + 4)$. The distance of points of P on this interval to the points of \mathcal{S} is going to be on average about $L_i/4$. Thus, the contribution of the points of P of \mathcal{I}_i to $v_{\mathcal{S}}(P)$ is $(n/2^i)(L_i/4) \geq n/(2(s_i + 4))$. Note that the origin is a median in this case and

$$v_{\text{opt}}(P, 1) \leq \sum_{i=1}^t \frac{n}{2^i} \cdot 2^{i+1} \leq 2nt.$$

Hence,

$$\frac{1}{2} \sum_{i=1}^t \frac{n}{s_i + 4} \leq v_{\mathcal{S}}(P) \leq \varepsilon v_{\text{opt}}(P, 1) \leq \varepsilon 2nt.$$

This gives us that $\sum_{i=1}^t (1/(s_i + 4)) \leq 4\varepsilon \lg n$.

We observe that $|\mathcal{S}| \geq \alpha = \sum_{i=1}^t s_i$. The quantity α is minimized, subject to the above constraint, when all the s_i 's are equal. Thus, the above minimum is achieved when $s_1 = \dots = s_t = \beta$, and $\lg n/(\beta + 4) = 4\varepsilon \lg n$. Thus, $\beta = 1/(4\varepsilon) - 4$, implying that

$$|\mathcal{S}| \geq \sum_{i=1}^{\lg n} s_i \geq \beta \lg n = \Omega\left(\frac{\lg n}{\varepsilon}\right). \quad \square$$

This testifies that our more involved analysis (i.e., Theorem 2.13) to get a better coresets of size independent of n is indeed necessary. In particular, our improved coresets construction works since it guarantees that the errors introduced by snapping the points to the coresets cancel themselves out when considering any set of k medians.

References

- [A] S. Arora. Polynomial time approximation schemes for euclidean tsp and other geometric problems. *J. Assoc. Comput. Mach.*, 45(5):753–782, Sept. 1998.
- [ARR] S. Arora, P. Raghavan, and S. Rao. Approximation schemes for Euclidean k -median and related problems. In *Proc. 30th Annu. ACM Sympos. Theory Comput.*, pages 106–113, 1998.
- [BHI] M. Bădoiu, S. Har-Peled, and P. Indyk. Approximate clustering via coresets. In *Proc. 34th Annu. ACM Sympos. Theory Comput.*, pages 250–257, 2002.
- [CGTS] M. Charikar, S. Guha, E. Tardos, and D. B. Shmoys. A constant-factor approximation algorithm for the k -median problem. In *Proc. 31st Annu. ACM Sympos. Theory Comput.*, pages 1–10, 1999.
- [DF] R. G. Downey and M. R. Fellows. Fixed-parameter tractability and completeness, I: Basic results. *SIAM J. Comput.*, 24(4):873–921, 1995.
- [dlVKKR] W. F. de la Vega, M. Karpinski, C. Kenyon, and Y. Rabani. Approximation schemes for clustering problems. In *Proc. 35th Annu. ACM Sympos. Theory Comput.*, pages 50–58, 2003.
- [ES] M. Effros and L. J. Schulman. Deterministic clustering with data nets. Technical Report TR04-085, Elec. Colloq. Comp. Complexity, 2003. <http://www.eccc.uni-trier.de/eccc-reports/2004/TR04-085/>.
- [GNMO] S. Guha, R. Motwani, N. Mishra, and L. O’Callaghan. Clustering data streams. In *Proc. 41st Annu. IEEE Sympos. Found. Comput. Sci.*, pages 359–366, 2000.
- [HM] S. Har-Peled and S. Mazumdar. Coresets for k -means and k -median clustering and their applications. In *Proc. 36th Annu. ACM Sympos. Theory Comput.*, pages 291–300, 2004.
- [IKI] M. Inaba, N. Katoh, and H. Imai. Applications of weighted voronoi diagrams and randomization to variance-based k -clustering. In *Proc. 10th Annu. ACM Sympos. Comput. Geom.*, pages 332–339, 1994.
- [KMN⁺] T. Kanungo, D. M. Mount, N. S. Netanyahu, C. D. Piatko, R. Silverman, and A. Y. Wu. A local search approximation algorithm for k -means clustering. In *Proc. 18th Annu. ACM Sympos. Comput. Geom.*, pages 10–18, 2002.
- [KR] S. G. Kolliopoulos and S. Rao. A nearly linear-time approximation scheme for the euclidean k -median problem. In *Proc. 7th Annu. European Sympos. Algorithms*, pages 378–389, 1999.

- [KSS] A. Kumar, Y. Sabharwal, and S. Sen. Linear time algorithms for clustering problems in any dimension. In *Proc. 32nd Internat. Colloq. Automata Lang. Prog.*, 2005.
- [M1] J. Matoušek. On approximate geometric k -clustering. *Discrete Comput. Geom.*, 24:61–84, 2000.
- [M2] J. Matoušek. *Lectures on Discrete Geometry*. Springer, New York, 2002.
- [MP] R. Mettu and C. G. Plaxton. Optimal time bounds for approximate clustering. In *Proc. 18th Conf. Uncertainty Artificial Intell.*, pages 344–351, 2002.

Received May 14, 2006. Online publication December 5, 2006.