

k -Center Clustering with Outliers in the Sliding-Window Model

Mark de Berg ✉

Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands

Morteza Monemizadeh ✉

Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands

Yu Zhong ✉

Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands

Abstract

The k -center problem for a point set P asks for a collection of k congruent balls (that is, balls of equal radius) that together cover all the points in P and whose radius is minimized. The k -center problem with outliers is defined similarly, except that z of the points in P do need not to be covered, for a given parameter z . We study the k -center problem with outliers in data streams in the sliding-window model. In this model we are given a possibly infinite stream $P = \langle p_1, p_2, p_3, \dots \rangle$ of points and a time window of length W , and we want to maintain a small sketch of the set $P(t)$ of points currently in the window such that using the sketch we can approximately solve the problem on $P(t)$.

We present the first algorithm for the k -center problem with outliers in the sliding-window model. The algorithm works for the case where the points come from a space of bounded doubling dimension and it maintains a set $S(t)$ such that an optimal solution on $S(t)$ gives a $(1 + \varepsilon)$ -approximate solution on $P(t)$. The algorithm uses $O((kz/\varepsilon^d) \log \sigma)$ storage, where d is the doubling dimension of the underlying space and σ is the spread of the points in the stream. Algorithms providing a $(1 + \varepsilon)$ -approximation were not even known in the setting without outliers or in the insertion-only setting with outliers. We also present a lower bound showing that any algorithm that provides a $(1 + \varepsilon)$ -approximation must use $\Omega((kz/\varepsilon) \log \sigma)$ storage.

2012 ACM Subject Classification Theory of computation \rightarrow Design and analysis of algorithms

Keywords and phrases Streaming algorithms, k -center problem, sliding window, bounded doubling dimension

Digital Object Identifier 10.4230/LIPIcs.ESA.2021.13

Funding MdB is supported by the Dutch Research Council (NWO) through Gravitation-grant NETWORKS-024.002.003.

1 Introduction

Clustering is one of the most important tools to analyze large data sets. A well-known class of clustering algorithms is formed by centroid-based algorithms, which include k -means clustering, k -median clustering and k -center clustering. The latter type of clustering is the topic of our paper. In the k -center problem one is given a set P of points from a metric space and a parameter k , and the goal is to find k congruent balls (that is, balls of equal radius) that together cover the points from P and whose radius is minimized. Note that the special case $k = 1$ corresponds to the minimum-enclosing ball problem. Data sets in practice often contain outliers, leading to the *k -center problem with outliers*. Here we are given, besides P and k , a parameter z that indicates the allowed number of outliers. Thus the radius of the balls in an optimal solution is given by

$$\text{OPT}_{k,z}(P) := \text{the smallest radius } \rho \text{ such that we can cover all points from } P, \\ \text{except for at most } z \text{ outliers, by } k \text{ balls of radius } \rho.$$


© Mark de Berg, Morteza Monemizadeh, and Yu Zhong;
licensed under Creative Commons License CC-BY 4.0

29th Annual European Symposium on Algorithms (ESA 2021).

Editors: Petra Mutzel, Rasmus Pagh, and Grzegorz Herman; Article No. 13; pp. 13:1–13:13

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In this paper we study the k -center problem with outliers in data streams, where the input is a possibly infinite stream $P = \langle p_1, p_2, \dots \rangle$ of points. The goal is to maintain a solution to the k -center problem as the points arrive over time, without any knowledge of future arrivals and using limited (sub-linear) storage. Since we cannot store all the points in the stream, we cannot expect to maintain an optimal solution. Hence, the two main quality criteria of a streaming algorithm are its approximation ratio and the amount of storage it uses. We will study this problem in the *sliding-window model*. In this model we are given a window length W and we are, at any time t , only interested in the points that arrived in the time window $(t - W, t]$. Working in the sliding-window model is often significantly more difficult than working in the standard (insertion-only) streaming model.

Previous work

Charikar et al. [6] were the first to study the metric k -center problem in data streams. They developed an algorithm in the insertion-only model that computes an 8-approximation for the k -center problem using $\Theta(k)$ space. Later McCutchen and Khuller [13] improved the approximation ratio to $2 + \varepsilon$ at the cost of increasing the storage to $O((k/\varepsilon) \log(1/\varepsilon))$. McCutchen and Khuller also studied the k -center problem with $z \geq 1$ outliers, for which they gave a $(4 + \varepsilon)$ -approximation algorithm that requires $O(kz/\varepsilon)$ space.

The above results are for general metric spaces. In spaces of *bounded doubling dimension*¹ better bounds are possible. Indeed, Ceccarello et al. [3] gave a $(3 + \varepsilon)$ -approximation algorithm for the k -center problem with z outliers, thus improving the approximation ratio $(4 + \varepsilon)$ for general metrics. Their algorithm requires $O((k + z)(1/\varepsilon)^d)$ storage, where d is the doubling dimension of the underlying space (which is assumed to be a fixed constant).

The algorithms mentioned so far are deterministic. Charikar et al. [8] and Ding et al. [10] studied sampling-based streaming algorithms for the Euclidean k -center problem with outliers, showing that if one allows slightly more than z outliers then randomization can reduce the storage requirements. Our focus, however, is on deterministic algorithms.

For the k -center problem in the sliding-window model, the only result we are aware of is due to Cohen-Addad et al. [9]. They deal with the k -center problem in general metric spaces, but without outliers, and they propose a $(6 + \varepsilon)$ -approximation algorithm using $O((k/\varepsilon) \log \sigma)$ storage, and a $(4 + \varepsilon)$ -approximation for the special case $k = 2$. Here σ denotes (an upper bound on) the spread of the points in the stream; in other words, $\sigma := \Delta_{\max}/\Delta_{\min}$, where Δ_{\max} is an upper bound on the maximum distance and Δ_{\min} is a lower bound on the minimum distance between any two distinct points. It is assumed that Δ_{\max} and Δ_{\min} are known to the algorithm. They also prove that any algorithm for the 2-center problem with outliers in general metric spaces that achieves an approximation ratio of less than 4 requires $\Omega(W^{1/3})$ space, where W is the size² of the window. Table 1 gives an overview of the known results on the k -center problem in the insertion-only and the sliding-window model.

While our main interest is in the k -center problem for $k > 1$, our results also apply when $k = 1$. Hence, we also briefly discuss previous results for the 1-center problem.

¹ The *doubling dimension* of a space X is the smallest number d such that any ball B in the space can be covered by 2^d balls of radius $\text{radius}(B)/2$.

² Here the window size W is defined in terms of the number of points in the window, that is, the window consists of the W most recent points. We define the window in a slightly more general manner, by defining W to be the length (that is, duration) of the window. Note that if we assume that the i -th point arrives at time $t = i$, then the two models are the same.

■ **Table 1** Results for the k -center problem with and without outliers in the insertion-only and the sliding-window model. Bounds on the storage are asymptotic. In the papers where the metric space has bounded doubling dimension or is Euclidean, the dimension d is considered a constant.

model	metric space	approx.	storage	outliers	ref.
insertion-only	general	8	k	no	[6]
	general	$2 + \varepsilon$	$(k/\varepsilon) \log(1/\varepsilon)$	no	[13]
	general	$4 + \varepsilon$	kz/ε	yes	[13]
	bounded doubling	$3 + \varepsilon$	$(k+z)/\varepsilon^d$	yes	[3]
sliding window	general	$6 + \varepsilon$	$(k/\varepsilon) \log \sigma$	no	[9]
	bounded doubling	$1 + \varepsilon$	$(kz/\varepsilon^d) \log \sigma$	yes	here

For the 1-center problem in d -dimensional Euclidean space, streaming algorithms that maintain an ε -kernel give a $(1 + \varepsilon)$ -approximation. An example is the algorithm of Zarrabi-Zadeh [14] which maintains an ε -kernel of size $O(1/\varepsilon^{(d-1)/2} \log(1/\varepsilon))$. Moreover, using only $O(d)$ storage one can obtain a 1.22-approximation for the 1-center problem without outliers [2, 4]. For the 1-center problem with outliers, one can obtain a $(1 + \varepsilon)$ -approximation algorithm that uses $z/\varepsilon^{O(d)}$ storage by the technique of Agarwal et al. [1]. Zarrabi-Zadeh and Mukhopadhyay [15] studied the 1-center problem with z outliers in high-dimensional Euclidean spaces, where d is not considered constant, giving a 1.73-approximation algorithm that requires $O(d^3 z)$ storage. Recently, Hatami and Zarrabi-Zadeh [12] extended this result to 2-center problem with z outliers, obtaining a $(1.8 + \varepsilon)$ -approximation using $O(d^3 z^2 + dz^4/\varepsilon)$ storage. None of the 1-center algorithms discussed above works in the sliding-window model.

A problem that is closely related to the 1-center problem is the *diameter problem*, where the goal is to maintain an approximation of the diameter of the points in the stream. This problem has been studied in the sliding-window model by Feigenbaum, Kannan, Zhang [11] and later by Chan and Sadjadj [5], who gave a $(1 + \varepsilon)$ -approximation for the diameter problem (without outliers) in the sliding window model, using $O((1/\varepsilon)^{(d+1)/2} \log(\sigma/\varepsilon))$ storage.

Our contribution

We present the first algorithm for the k -center problem with z outliers in the sliding-window model. It works in spaces of bounded doubling dimension and yields a $(1 + \varepsilon)$ -approximation. So far a $(1 + \varepsilon)$ -approximation was not even known for the k -center problem without outliers in the insertion-only model. Our algorithm uses $O((kz/\varepsilon^d) \log \sigma)$ storage,³ where d is the doubling dimension and σ is the spread, as defined above. Thus for the 1-center problem we obtain a solution that uses $O((z/\varepsilon^d) \log \sigma)$ storage. This solution also works for the diameter problem. Note that also for the 1-center problem with outliers (and the diameter problem with outliers) an algorithm for the sliding-window model was not yet known. A useful property of the sketch⁴ maintained by our algorithm for the k -center problem, is that it can also be used for the k' -center problem for any $k' < k$, as well as for the diameter problem. As in the previous papers on the k -center problem (or the diameter problem) in the sliding-window model [5, 9, 11], we assume that Δ_{\max} and Δ_{\min} (upper and lower bounds on the maximum and minimum distance between any two points in the stream) are known to the algorithm.

³ To correctly state the bound for the case $z = 0$ as well, we should actually write $O((k(z+1)/\varepsilon^d) \log \sigma)$. To keep the bound concise, we will just write $O((kz/\varepsilon^d) \log \sigma)$, however.

⁴ We use the word *sketch* even though we do not study how to compose the sketches for two separate streams into a sketch for the concatenation of the stream, as the term *sketch* seems more appropriate than *data structure*, for example.

As mentioned above, our algorithm provides a $(1 + \varepsilon)$ -approximation. More precisely, it maintains a set $S(t) \subset P(t)$ such that computing an optimal solution for $S(t)$ and then suitably expanding the balls in the solution, gives a $(1 + \varepsilon)$ -approximate solution for $P(t)$. Thus, to obtain a $(1 + \varepsilon)$ -approximate solution for $P(t)$ one needs to compute an optimal (or $(1 + \varepsilon')$ -approximate, for a suitable ε') solution for $S(t)$. Depending on the underlying space this can be slow, and it may be preferable to compute, say, a 2-approximation for $S(t)$, which would then give a $(2 + \varepsilon)$ -approximation for $P(t)$. The various options in Euclidean spaces and general spaces of bounded doubling dimension are discussed in Section 2.3.

Our second contribution is a lower bound for the k -center problem with outliers in the sliding-window model. Our lower bound shows that any algorithm that provides a $(1 + \varepsilon)$ -approximation must use $\Omega((kz/\varepsilon) \log \sigma)$ storage. This matches our upper bound up to the dependency on ε and d . The lower-bound construction uses points in \mathbb{R}^1 , so it shows that our algorithm is optimal in this case. The lower bound model is very general. It allows the algorithm to store points, or weighted points, or balls, or whatever it wants so that it can approximate an optimal solution for $P(t)$ at any time t . The main condition is that each object comes with an expiration time which corresponds to the expiration time of some point in the stream, and that the algorithm can only change its state when an object expires or a new object arrives; see Section 3 for details.

2 The algorithm

Let $P := \langle p_1, p_2, \dots \rangle$ be a possibly infinite stream of points from a metric space X of doubling dimension d and spread σ , where d is considered to be a fixed constant. We denote the arrival time of a point p_i by $t_{\text{arr}}(p_i)$. We say that p_i *expires* at time $t_{\text{exp}}(p_i) := t_{\text{arr}}(p_i) + W$, where W is the given length of the time window. To simplify the exposition, we assume that all arrival times and departure times (that is, times at which a point expires) are distinct. For a time t we define $P(t)$ to be the set⁵ of points currently in the window. In other words, $P(t) := \{p_i : t_{\text{arr}}(p_i) \leq t < t_{\text{exp}}(p_i)\}$. For a point $q \in X$ and a parameter $r \geq 0$, we use $\text{ball}(q, r)$ to denote the ball with center q and radius r .

In the following we show how to maintain a sketch $\Gamma(t)$ of $P(t)$ for the k -center problem with outliers. The idea behind our algorithm is as follows. Consider an optimal solution for $P(t)$ consisting of k balls B_1, \dots, B_k of radius ρ_{opt} . Suppose we cover the points in each ball B_i by a number of smaller balls of radius $\varepsilon \rho_{\text{opt}}$; we will call these *mini-balls*. Since the underlying space has doubling dimension d , we can do this using $O(1/\varepsilon^d)$ mini-balls for each ball B_i . To obtain a $(1 + \varepsilon)$ -approximation, we do not need to know all the points: it suffices to keep $z + 1$ points in each mini-ball, since then we cannot designate all points in the mini-ball as outliers. There are several challenges to overcome to make this idea work. For instance, we do not know ρ_{opt} . Hence, we will develop an algorithm for the decision version of the problem, which we will then apply “in parallel” to different possible values for the optimal radius. But when ρ is much smaller than ρ_{opt} we have another challenge, namely that there can be many outliers (which may need to be stored, because at some later moment in time they may become inliers). Finally, since points arrive and expire, the balls in an optimal solution move over time, which makes it hard to maintain the mini-balls correctly. Below we show how to overcome these challenges.

⁵ We allow the same point from X to occur multiple times in the stream, so $P(t)$ is actually a multi-set. Whenever we refer to “sets” in the remainder of the paper we mean “multi-sets”. The distance Δ_{min} is defined with respect to distinct points, however.

2.1 A sketch for the decision problem

Recall that for a set Q of points, $\text{OPT}_{k,z}(Q)$ denotes the smallest radius ρ such that we can cover all points from Q , except for at most z outliers, by k balls of radius ρ , and that Δ_{\min} and Δ_{\max} are bounds on the minimum and maximum distance between any two points in the stream. Let ρ be a parameter with $\Delta_{\min} \leq \rho \leq \Delta_{\max}$. We will design a sketch $\Gamma(t)$ and a corresponding decision algorithm TRYTOCOVER with the following properties:

- The sketch $\Gamma(t)$ uses $O(kz/\varepsilon^d)$ storage.
- TRYTOCOVER either reports a collection \mathcal{C}^* of k balls of radius $\text{OPT}_{k,z}(P(t)) + 2\varepsilon\rho$ that together cover all points in $P(t)$ except for at most z outliers, or it reports NO. In latter case we have $\text{OPT}_{k,z}(P(t)) > 2\rho$.

Our sketch $\Gamma(t)$ is a tuple $[\tau(t), \mathcal{B}(t), \mathcal{R}(t), P_{\text{out}}(t)]$, where τ is a *timestamp*, $\mathcal{B}(t)$ is a collection of mini-balls, $\mathcal{R}(t)$ is a collection of so-called *representative sets*, and $P_{\text{out}}(t)$ is a set of outliers. We will maintain the following invariants.

(Inv-1) For any time t' with $t \leq t' < \tau(t)$ we are guaranteed that $\text{OPT}_{k,z}(P(t')) > 2\rho$. The idea is that when we find $k + z + 1$ points with pairwise distances greater 2ρ , then we can set τ to the expiration time of the oldest of these points, and we can delete this point (as well as any older points) from the sketch.

(Inv-2) Each mini-ball $B \in \mathcal{B}(t)$ has radius $\varepsilon\rho$ and the center of B , denoted by $\text{center}(B)$, is a point from the stream. Note that the center does not need to be a point from $P(t)$. The mini-balls in $\mathcal{B}(t)$ are *well spread* in the sense that no mini-ball contains the center of any other mini-ball. In other words, we have $\text{dist}(\text{center}(B), \text{center}(B')) > \varepsilon\rho$ for any two mini-balls $B, B' \in \mathcal{B}(t)$.

(Inv-3) For each mini-ball $B \in \mathcal{B}(t)$ the set $\mathcal{R}(t)$ contains a *representative set* $R(B) \subseteq B \cap P(t)$, and these are the only sets in $\mathcal{R}(t)$. The representative sets $R(B)$ are pairwise disjoint, and each set $R(B)$ contains at most $z + 1$ points. When $|R(B)| = z + 1$, we say that the mini-ball B is *full*.

(Inv-4) Define $P_{\text{in}}(t) := \bigcup_{B \in \mathcal{B}(t)} R(B)$ and let $S(t) := P_{\text{in}}(t) \cup P_{\text{out}}(t)$ be the collection of all points in our sketch. For any point $p_i \in P(t) \setminus S(t)$ – these are exactly the points that have not yet expired but that have been discarded by the algorithm – we have (i) $t_{\text{exp}}(p_i) \leq \tau(t)$, and/or (ii) $p_i \in B$ for a mini-ball $B \in \mathcal{B}(t)$ that is full and such that all points in $R(B)$ arrived after p_i .

(Inv-5) $|\mathcal{B}(t)| = O(k/\varepsilon^d)$ and $|P_{\text{out}}(t)| \leq z$.

At time $t = 0$, before the arrival of the first point, we have $\tau(t) = 0$ and $\mathcal{B}(t) = \mathcal{R}(t) = P_{\text{out}}(t) = \emptyset$. Since $P(0) = \emptyset$ this trivially satisfies the invariants. Before we prove that we can maintain the sketch upon the arrival and departure of points, we present our decision algorithm TRYTOCOVER and prove its correctness.

■ **Algorithm 1** TRYTOCOVER($\Gamma(t)$).

-
- 1: $S(t) \leftarrow \bigcup_{B \in \mathcal{B}(t)} R(B)$
 - 2: Compute $\text{OPT}_{k,z}(S(t))$ and the corresponding collection $\mathcal{C} := \{C_1, \dots, C_k\}$ of balls.
 - 3: **if** $t < \tau(t)$ or $\text{OPT}_{k,z}(S(t)) > 2\rho$ **then**
 - 4: Report NO
 - 5: **else**
 - 6: Increase the radius of each ball $C_i \in \mathcal{C}$ by $2\varepsilon\rho$.
 - 7: Report the collection $\mathcal{C}^* := \{C_1^*, \dots, C_k^*\}$ of expanded balls.
-

In line 2 we compute an optimal solution on the point set $S(t)$. How this can be done depends on the underlying metric space and will be discussed in Section 2.3. The following lemma establishes the correctness of the algorithm.

► **Lemma 1.** *Algorithm TRYTOCOVER either reports a collection \mathcal{C}^* of k balls of radius $\text{OPT}_{k,z}(P(t)) + 2\varepsilon\rho$ that together cover all points in $P(t)$ except for at most z outliers, or it reports NO. In latter case we have $\text{OPT}_{k,z}(P(t)) > 2\rho$.*

Proof. First suppose the algorithm reports NO. If this happens because $t < \tau(t)$ then $\text{OPT}_{k,z}(P(t)) > 2\rho$ by (Inv-1). Otherwise, this happens because $\text{OPT}_{k,z}(S(t)) > 2\rho$. But then $\text{OPT}_{k,z}(P(t)) > 2\rho$, because (Inv-3) implies that $S(t) \subseteq P(t)$.

Now suppose the algorithm reports a collection $\mathcal{C}^* := \{C_1^*, \dots, C_k^*\}$ of balls. Let \mathcal{C} be the corresponding set of balls before they were expanded. Since \mathcal{C} is an optimal solution for $S(t)$, the balls C_i have radius $\text{OPT}_{k,z}(S(t)) \leq \text{OPT}_{k,z}(P(t))$ and together they cover all points in $S(t)$ except for at most z outliers. Now consider a point $p_i \in P(t) \setminus S(t)$. To finish the proof, we must show that p_i is covered by one of the balls in \mathcal{C}^* . To this end, first observe that $t_{\text{exp}}(p_i) > t$ because $p_i \in P(t)$. Since TRYTOCOVER did not report NO, this implies that $t_{\text{exp}}(p_i) > \tau(t)$. Hence, we can conclude from (Inv-4) that $p_i \in B$ for a mini-ball $B \in \mathcal{B}(t)$ that is full. Thus $R(B)$ contains $z + 1$ points, and since we allow only z outliers this implies that at least one point from $R(B)$ is covered by a ball $C_i \in \mathcal{C}$. Because $\text{diam}(B) = 2\varepsilon\rho$, this implies that p_i must be covered by C_i^* , thus finishing the proof. ◀

Next we show how to update the sketch $\Gamma(t)$.

Handling departures

When a point p_j in one of our representative sets $R(B)$ expires, we simply delete it from $R(B)$, and if $R(B)$ then becomes empty we remove $R(B)$ from $\mathcal{R}(t)$ and B from $\mathcal{B}(t)$. Similarly, if p_j was a point in $P_{\text{out}}(t)$ we remove p_j from $P_{\text{out}}(t)$.

It is trivial to verify that (Inv-1)–(Inv-3) and (Inv-5) still hold for the updated sketch. To see that (Inv-4) holds as well, consider a point $p_i \in P(t) \setminus S(t)$. The only reason for (Inv-4) to be violated, would be when $p_i \in B$ for a mini-ball B that was full before the deletion of p_j but is no longer full after the deletion. However, (Inv-4) states that all points in $R(B)$ arrived after p_i . Since p_i did not yet expire, this means that the point p_j that currently expires cannot be a point from $R(B)$.

Handling arrivals

The arrival of a point p_j at time $t := t_{\text{arr}}(p_j)$ is handled by Algorithm 2. We denote the sketch just before the arrival by $\Gamma(t^-)$, and the updated sketch by $\Gamma(t^+)$.

Algorithm 2 HANDLEARRIVAL($\Gamma(t^-), p_j$).

1. If p_j lies in an existing mini-ball $B \in \mathcal{B}(t^-)$ then add p_j to $R(B)$. If p_j did not lie in an existing mini-ball then add p_j to $P_{\text{out}}(t^-)$.
 2. Let $Q := P_{\text{out}}(t^-) \cup \bigcup_{B \in \mathcal{B}(t^-)} R(B)$ be the set of all points in the current sketch, including the just inserted point. Sort the points in Q by decreasing arrival time, and let $q_1, q_2 \dots$ denote the sorted sequence of points – note that q_1 is the newly inserted point p_j . Let $Q_i := \{q_1, \dots, q_i\}$ denote the set containing the first i points from Q . Find the largest index i^* such that the following holds:
 - The 3-approximation algorithm of Charikar et al. [7] for the k -center clustering algorithm with z outliers, when run on Q_{i^*} , reports a collection \mathcal{B}^* of k balls of radius at most 6ρ . (In fact, any other constant-factor approximation algorithm can be used if we adjust the value 6ρ appropriately. For concreteness, we use the one by Charikar et al.)
 3. If $Q_{i^*} \neq Q$ (and so the point q_{i^*+1} exists) then set $\tau(t^+) := \max(\tau(t^-, t_{\text{exp}}(q_{i^*+1}))$. Remove the points in $Q \setminus Q_{i^*}$ from the representative sets they are in; note that these points all have expiration time at least $\tau(t^+)$.
 4. Let $\mathcal{B}^* := \{B_1, \dots, B_k\}$ be the balls in the solution reported for Q_{i^*} . Note that all points from Q_{i^*} , except for at most z outliers, lie in a ball from \mathcal{B}^* . The new set of mini-balls and outliers is now computed as follows. Increase the radius of the balls in \mathcal{B}^* by $\varepsilon\rho$. Denote the set of expanded balls by $\overline{\mathcal{B}}^*$
 - a. Add each mini-ball $B \in \mathcal{B}(t^-)$ whose center lies inside a ball from $\overline{\mathcal{B}}^*$ to the set $\mathcal{B}(t^+)$ and add their representative sets to $\mathcal{R}(t^+)$.
 - b. Let $Z \subset Q_{i^*}$ be the set of points that are not in a representative set added to $\mathcal{R}(t^+)$ in Step 4a. Let $Z_{\text{in}} \subseteq Z$ be the points that lie inside a ball from $\overline{\mathcal{B}}^*$. Go over the points in Z_{in} one by one, in arbitrary order, and handle each point q as follows: If there is a mini-ball in the current set $\mathcal{B}(t^+)$ that contains q then add q to $R(B)$ for an arbitrary such mini-ball B ; otherwise create a mini-ball B with center q and radius $\varepsilon\rho$, set $R(B) := \{q\}$ and add B to $\mathcal{B}(t^+)$.
 - c. For all mini-balls $B \in \mathcal{B}(t^+)$ such that $|R(B)| > z+1$, remove the oldest $|R(B)| - (z+1)$ points from $R(B)$.
 - d. Set $P_{\text{out}}(t^+) := Z \setminus Z_{\text{in}}$. Note that these are exactly the points from Q_{i^*} that have not been added to a representative set
-

► **Lemma 2.** *After the arrival of point p_j has been handled, the invariants (Inv-1)–(Inv-5) are restored.*

Proof. (Inv-1) If $\tau(t^+) = \tau(t^-)$ then obviously (Inv-1) still holds, so assume that τ is updated by the algorithm in Step 3. Thus the algorithm of Charikar et al. [7] returned a solution of radius more than 6ρ on Q_{i^*+1} . Since the algorithm is a 3-approximation, this means $\text{OPT}_{k,z}(Q_{i^*+1}) > 2\rho$. Because we sorted the points on expiration time, we must have $\text{OPT}_{k,z}(P(t')) > 2\rho$ for all t' with $t \leq t' < t_{\text{exp}}(q_{i^*+1})$. Hence, (Inv-1) still holds.

(Inv-2) By construction, the mini-balls in $\mathcal{B}(t^+)$ have radius $\varepsilon\rho$ and are centered at a point from the stream. The mini-balls from $\mathcal{B}(t^-)$ are well spread by induction, so we only need to prove that any mini-balls created in Step 4b are well spread. This is true because we only create a new mini-ball when its center does not lie inside an existing mini-ball.

(Inv-3) The representative sets $R(B)$ are disjoint by construction, and $|R(B)| \leq z+1$ is guaranteed by Step 4c. Hence (Inv-3) holds.

- (Inv-4) Take a point $p_i \in P(t^+) \setminus S(t^+)$ and assume $t_{\text{exp}}(p_i) > \tau(t^+)$. We have two cases.
- The first case is that p_i was discarded due to the arrival of p_j . Then it was either removed in Step 3, which implies $t_{\text{exp}}(p_i) \geq \tau(t^+)$; or it was removed in Step 4c, which implies there is a full mini-ball $B \in \mathcal{B}(t^+)$ such that all points in $R(B)$ are newer than p_i . Hence, (Inv-4) holds for p_i .
 - The second case is that p_i was already discarded earlier. We can assume that (Inv-4) holds before the arrival of p_j . If $t_{\text{exp}}(p_i) \leq \tau(t^-)$ then we also have $t_{\text{exp}}(p_i) \leq \tau(t^+)$ and we are done. So assume that $p_i \in B$ for a mini-ball $B \in \mathcal{B}(t^-)$ that was full and such that all points in $R(B)$ arrived after p_i . If $R(B)$ contained a point $p_{i'} \in Q \setminus Q_{i^*}$ then $t_{\text{exp}}(p_i) < t_{\text{exp}}(p_{i'}) \leq \tau(t^+)$, and so (Inv-4) holds. If $R(B)$ only contained points from Q_{i^*} then we have two sub-cases.
 The first sub-case is that B is still present in $\mathcal{B}(t^+)$. Then B must still be full, because $R(B)$ only contained points from Q_{i^*} . Hence, (Inv-4) holds.
 The second sub-case is that B is not present in $\mathcal{B}(t^+)$. Thus its center must lie outside all balls in $\overline{\mathcal{B}}^*$. But then the points in $R(B)$, which are all in Q_{i^*} , are outside any ball in \mathcal{B}^* . But this is impossible, since $|R(B)| = z + 1$ and \mathcal{B}^* has only z outliers.
- (Inv-5) All mini-balls in $\mathcal{B}(t^+)$ have their center inside one of the k balls in $\overline{\mathcal{B}}^*$, which have radius at most 6ρ . Moreover, the mini-balls in $\mathcal{B}(t^+)$ have radius $\varepsilon\rho$ and they are well spread. Since the underlying space has doubling dimension d , this implies that $|\mathcal{B}(t^+)| = O(k/\varepsilon^d)$.
 The outlier set $P_{\text{out}}(t^-)$ consists of the points from Q_{i^*} that lie outside the balls in $\overline{\mathcal{B}}^*$. Hence, these points also lie outside the balls in \mathcal{B}^* . Since \mathcal{B}^* is a valid solution to the k -center problem with z outliers, there are at most z such points. Thus $|P_{\text{out}}(t^+)| \leq z$. ◀

2.2 A sketch for the optimization problem

Above we presented a sketch for a decision version of the problem, for given parameters ρ and ε . The sketch uses $O(kz/\varepsilon^d)$ storage. We also gave an algorithm `TRYTOCOVER` that either reports a collection \mathcal{C}^* of k balls of radius $\text{OPT}_{k,z}(P(t)) + 2\varepsilon\rho$ covering all points in $P(t)$ except at most z outliers, or that reports `NO`. In latter case we know that $\text{OPT}_{k,z}(P(t)) > 2\rho$. To make the parameter ρ and ε explicit, we will from now on denote the sketch by $\Gamma_{\rho,\varepsilon}$.

In the optimization version of the problem we wish to find k congruent balls of minimum radius that together cover all points in $P(t)$ except for at most z outliers. To obtain a $(1 + \varepsilon)$ -approximation for the optimization problem, for a given $\varepsilon > 0$, we maintain a sketch $\Gamma_{\rho_i,\varepsilon/2}$ for every $0 \leq i \leq \lfloor \log \sigma \rfloor$, where $\rho_i := 2^i \cdot \Delta_{\min}$. We can then obtain a $(1 + \varepsilon)$ -approximate solution with the algorithm shown in Algorithm 3.

■ **Algorithm 3** `FINDAPPROXIMATECENTERS`($\Gamma(t)$).

```

1:  $i \leftarrow 0$ 
2: repeat
3:    $answer \leftarrow \text{TRYTOCOVER}(\Gamma_{\rho_i,\varepsilon/2}); i \leftarrow i + 1$ 
4: until  $answer \neq \text{NO}$ 
5: if the radii of the balls in  $answer$  is less than  $\Delta_{\min}$  then
6:   Reduce the radii of the balls to zero
7: Report  $answer$ 
    
```

We obtain the following theorem, whose proof can be found in the full version of the paper.

► **Theorem 3.** *The sketch above uses $O((kz/\varepsilon^d)\log\sigma)$ storage and that allows us to report at any time t a collection \mathcal{C}^* of balls of radius at most $(1+\varepsilon)\cdot\text{OPT}_{k,z}(P(t))$ covering all points from $P(t)$, except at most z outliers.*

Our sketch for the k -center problem can also be used for the k' -center problem for $k' < k$. Moreover, a sketch for the k -center problem for $k \geq 1$ can be used for the diameter problem. Recall that the diameter problem with outliers for the set $P(t)$ asks for the value

$$\text{diam}_z(P(t)) := \min\{\text{diam}(P(t) \setminus Q) : |Q| = z\},$$

that is, $\text{diam}_z(P(t))$ is the smallest diameter one can obtain by deleting z outliers from $P(t)$. We say that an algorithm reports a $(1-\varepsilon)$ -approximation to $\text{diam}_z(P(t))$ if it reports a value D with $(1-\varepsilon)\cdot\text{diam}_z(P(t)) \leq D \leq \text{diam}_z(P(t))$. The proof of the following theorem is in the full version of the paper.

► **Theorem 4.** *The sketch for the k -center problem with outliers as presented above can also be used to provide a $(1+\varepsilon)$ -approximation for the k' -center problem with outliers, for any $1 \leq k' \leq k$. Moreover, it can be used to provide a $(1-2\varepsilon)$ -approximation for the diameter problem with outliers.*

Theorem 4 implies that there exists a sketch for the diameter problem with outliers in the sliding-window model that gives a $(1+\varepsilon)$ -approximation to $\text{diam}_z(P(t))$ using $O((z/\varepsilon^d)\log\sigma)$ storage, namely the sketch for the 1-center problem.

2.3 Time complexity

Above we focused on the storage used by our sketch, and on the approximation ratio it can potentially provide. We now discuss the time complexity.

The algorithm that handles departures trivially runs in $O(1)$ time. The most time-consuming step in the algorithm that handles arrivals is Step 2, where we run the algorithm of Charikar et al. [7] on a set of $O(kz/\varepsilon^d)$ points. This takes $O((kz/\varepsilon^d)^3)$ time. Note that, both for departures and arrivals, we need to execute the respective algorithms for each of the $\log\sigma$ sketches $\Gamma_{\rho_i, \varepsilon/2}$ we maintain.

The time needed for FINDAPPROXIMATECENTERS is more interesting, since it calls TRYTOCOVER, which is supposed to compute an optimal solution for the k -center problem with outliers on a given set $S(t)$ of $O(kz/\varepsilon)$ points. It is routine to check that if TRYTOCOVER would use a c -approximate solution on $S(t)$, instead of an optimal solution, our final result would be a $c(1+\varepsilon)$ -approximation. In d -dimensional Euclidean space we can solve the k -center problem with outliers in time polynomial in $n := |S(t)|$ (for constant k and d), as follows: first generate all $O(n^{d+1})$ potential centers – there are $O(n^{d+1})$ potential centers because the smallest enclosing ball of a given point set is defined by at most $d+1$ points – then generate all possible collections of k such centers, and then for each of the $O(n^{(d+1)k})$ such collections find the minimum radius ρ such that we can cover all except for z points. A similar approach is possible for other metrics in \mathbb{R}^d that are sufficiently well-behaved (such as ℓ_p metrics, for instance).

For arbitrary metric spaces the situation is more tricky. The standard assumption, which we also make, is that we can compute the distance between any two given points from the metric space in $O(1)$ time. Still, computing an optimal solution on the set $S(t)$ can be quite slow or perhaps even infeasible. For example, when an optimal center is a point from the metric space that does not occur in the stream, then the algorithm may be unable to retrieve this center. So unless we are in Euclidean space, or some other “well-defined” space, it

13:10 k -Center Clustering with Outliers in the Sliding-Window Model

seems most natural to define the underlying space as consisting of exactly the points in the stream. Note that this can still be somewhat problematic, as it may require the algorithm to use points that have already expired as centers for the current set $S(t)$. These issues seem unavoidable for any algorithm that maintains a sketch in the sliding-window model. Also note that, using just the points in the current set $P(t)$ as centers, one can always obtain a 2-approximation to the optimal clustering. Invariant (Inv-4) thus implies that by using the set $S(t)$ as potential centers we can obtain a $(2 + \varepsilon)$ approximation. Also note that the optimal k -center clustering with z outliers on $S(t)$, under the restriction that the centers come from $S(t)$, can be trivially computed in $O(|S(t)|^{k+2})$ time.

The following theorem summarizes the performance of our sketch.

► **Theorem 5.** *Consider the k -center problem with z outliers in the sliding-window model, for streams of points from a space with doubling dimension d . Let Δ_{\min} and Δ_{\max} be given (lower and upper) bounds on the minimum and maximum distance between any two points in the stream, and let $\sigma := \Delta_{\max}/\Delta_{\min}$ be (an upper bound on) the spread of the stream. Suppose we have a c -approximation algorithm for the static version of the problem that runs in $T_{k,z}(n)$ time on a set of n points. Then for any $0 < \varepsilon < 1$ we can maintain a sketch with the following properties:*

- The sketch uses $O((kz/\varepsilon^d) \log \sigma)$ storage.
- Departures and arrivals can be handled in $O(\log \sigma)$ and $O((kz/\varepsilon^d)^3 \log \sigma)$ time, respectively.
- At any time t , the algorithm can report a valid solution for $P(t)$ of cost at most $(1 + \varepsilon)c \cdot \text{OPT}_{k,z}(P(t))$. The time needed to compute the solution is $T_{k,z}(n)$, where $n := O(kz/\varepsilon^d)$.

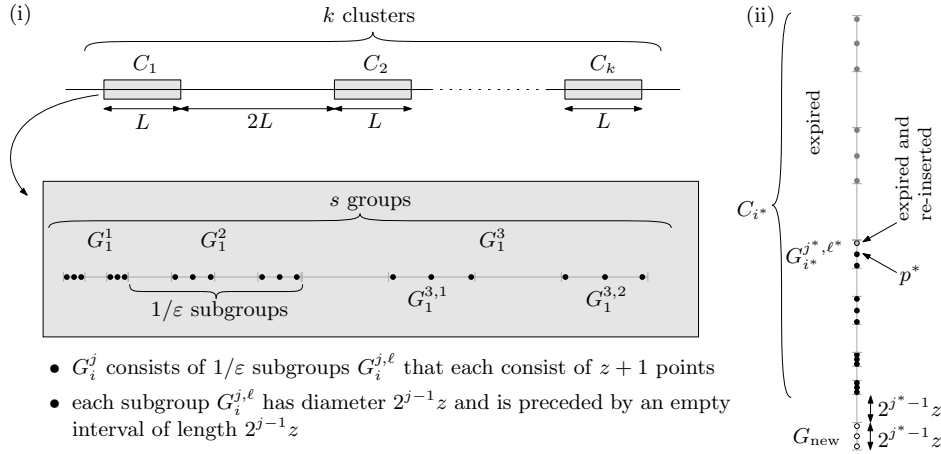
3 A lower bound

Above we presented a sketch of size $O((kz/\varepsilon^d) \log \sigma)$ that provides a $(1 + \varepsilon)$ -approximation for the k -center problem with z outliers in the sliding-window model. In this section we show that this is tight, up to the dependency on d . In particular, we prove that any sketch giving a $(1 + \varepsilon)$ -approximation for the problem in \mathbb{R}^1 must use $\Omega((kz/\varepsilon) \log \sigma)$ storage. In the full paper we show that to obtain any constant approximation ratio, a sketch must use $\Omega(kz)$ storage.

The lower-bound model

Our lower-bound is extremely simple and general. We allow the algorithm to store points, or weighted points, or balls, or whatever it wants so that it can approximate an optimal solution for $P(t)$ at any time t . We only make the following assumptions. Let $S(t)$ be the collection of objects being stored at time t .

- Each object in $S(t)$ is accompanied by an expiration time, which is equal to the expiration time of some point $p_i \in P(t)$.
- Let $p_i \in P(t)$. If no object in $S(t)$ uses $t_{\text{exp}}(p_i)$ as its expiration time, then no object in $S(t')$ with $t' > t$ can use $t_{\text{exp}}(p_i)$ as its expiration time.
- The solution reported by the algorithm is uniquely determined by $S(t)$, and the algorithm only modifies $S(t)$ when a new point arrives or when an object in $S(t)$ expires.
- The algorithm is deterministic and oblivious of future arrivals. In other words, the set $S(t)$ is uniquely determined by the sequence of arrivals up to time t , and the solution reported for $P(t)$ is uniquely determined by $S(t)$.



■ **Figure 1** The lower-bound construction for a $(1 + \varepsilon)$ -approximate sketch. In the example, $s = 3$, $k = 3$, $\varepsilon = 1/2$ and $z = 2$. (i) The initial configuration. (ii) Rotated view of the cluster containing p^* .

The storage used by the algorithm is defined as the number of objects in $S(t)$. The algorithm can decide which objects to keep in $S(t)$ anyway it wants; it may even keep an unbounded amount of extra information in order to make its decisions. The algorithm can also derive a solution for $P(t)$ in any way it wants, as long as the solution is valid and uniquely determined by $S(t)$. Clearly, the sketch from the previous section adheres to the model.

A lower bound for a $(1 + \varepsilon)$ -approximate sketch

Let ALG be a $(1 + \varepsilon')$ -approximation algorithm, for some $0 < \varepsilon' < 1$, where we assume for simplicity that $1/\varepsilon'$ is an integer. The lower-bound instance, which consists of points in \mathbb{R}^1 , is as follows.

Let $\varepsilon := \varepsilon'/8$ and consider the configuration shown in Figure 1(i). The configuration consists of k clusters C_1, \dots, C_k , placed next to each other from left to right. Each cluster C_i consists of s groups, G_i^1, \dots, G_i^s , where $s = \Theta(\log \sigma)$. Each group G_i^j consists of $1/\varepsilon$ subgroups, $G_i^{j,1}, \dots, G_i^{j,1/\varepsilon}$. A subgroup $G_i^{j,\ell}$ consists of $z + 1$ points at distance 2^{j-1} apart; the diameter of a subgroup – that is, the distance between its leftmost and rightmost point – is thus $2^{j-1}z$. Subgroup $G_i^{j,\ell}$ is preceded by an empty interval of length $2^{j-1}z$. Hence, the total diameter of the group G_i , including the empty interval preceding its leftmost subgroup, is $2^{i-1}z/\varepsilon$. (The exception is G_1 , for which we do not count the empty interval preceding its leftmost subgroup.) This brings the total diameter of a cluster to

$$L := \sum_{j=1}^s 2^j z/\varepsilon - z/\varepsilon = (2^{s+1} - 3) \cdot z/\varepsilon < 2^{s+1} z/\varepsilon.$$

In between every two consecutive clusters there is an empty interval of length $2L$.

The arrival order of the points in the configuration is as follows. The subgroups within a cluster arrive from right to left, in a round-robin fashion over the clusters: in the first round the subgroups $G_k^{s, \frac{1}{\varepsilon}}, \dots, G_1^{s, \frac{1}{\varepsilon}}$ arrive, in the second round the subgroups $G_k^{s, \frac{1}{\varepsilon}-1}, \dots, G_1^{s, \frac{1}{\varepsilon}-1}$ arrive, and so on. More formally, $G_i^{j,\ell}$ arrives before $G_{i'}^{j',\ell'}$ iff: $j > j'$ or $(j = j'$ and $\ell > \ell')$ or $(j = j'$ and $\ell = \ell'$ and $i > i')$. This finishes the description of the initial part of the instance.

13:12 k -Center Clustering with Outliers in the Sliding-Window Model

Let t be the time at which the last point in the configuration arrives, let $P(t)$ be the set of all points in the clusters C_1, \dots, C_k , and let $S(t)$ be the set of objects stored by ALG. We will argue that $|S(t)| = \Omega(kzs/\varepsilon)$, otherwise an adversary can continue the instance such that at some time t' in the future the algorithm does not give a $(1 + \varepsilon)$ -approximation.

Consider the points in the subgroups $G_i^{j,\ell}$ with $j > 1$, except for the very first point that arrived. Let T_{exp} be the set of expiration times of these points. Note that $|T_{\text{exp}}| = (k(z+1)(s-1)/\varepsilon) - 1$. Suppose at least one of these expiration times, say the expiration time of point p^* , is not used by any object in $S(t)$. Let $t_{p^*}^-$ and $t_{p^*}^+$ be the times immediately before and after $t_{\text{exp}}(p^*)$, respectively. The conditions of the model imply that the algorithm reports the same solution at times $t_{p^*}^-$ and $t_{p^*}^+$. The adversary will now continue the scenario in such that one of these answers is not sufficiently accurate, as described next.

Let $G_{i^*}^{j^*,\ell^*}$ be the subgroup that p^* belongs to. First, the adversary waits until all points that arrived before p^* have expired. The situation is then as follows: clusters C_i with $i < i^*$ consist of the subgroups up to $G_{j^*,\ell^*}^{(i)}$, while clusters C_i with $i > i^*$ are missing the last of these subgroups. Cluster C_{i^*} has the points up to point p^* (which expires next) in $G_{i^*}^{j^*,\ell^*}$. Before p^* expires, the adversary now adds a new subgroup G_{new} of $z+1$ points to C_{i^*} . The diameter of G_{new} is $z^{j^*-1}z$ and its distance to C_{i^*} is $z^{j^*-1}z$ as well; see Figure 1(ii). (The construction has been rotated by 90 degrees in the figure, to make it fit.) In addition, the adversary adds the points from $G_{j^*,\ell^*}^{(i)}$ that have already expired back.

We now analyze $\text{OPT}_{k,z}(P(t_{p^*}^-))$. Since the distance between any two of the k clusters is larger than the diameter of the clusters, and any cluster still contains at least $z+1$ points – the latter follows because $j^* > 1$ – an optimal solution will use a separate ball for each cluster. The largest cluster is C_{i^*} , because G_{new} was added to it. Since each subgroup in C_{i^*} has $z+1$ points and we can designate only z points as outliers, it is optimal to designate the z rightmost points (topmost in Figure 1(ii)) from $G_{i^*}^{j^*,\ell^*}$ as outliers and cover the remaining points with a ball of diameter $L^* + 3 \cdot 2^{j^*-1}z$, where

$$L^* := \sum_{j=1}^{j^*-1} 2^j z/\varepsilon - z/\varepsilon + 2\ell^* \cdot 2^{j^*-1}z = (2^{j^*} - 3)z/\varepsilon + \ell^* 2^{j^*}z < (2^{j^*+1})z/\varepsilon,$$

where the last inequality uses that $\ell^* \leq 1/\varepsilon$. Hence, $\text{OPT}_{k,z}(P(t_{p^*}^-)) = (L^* + 3 \cdot 2^{j^*-1}z)/2$.

At time $t_{p^*}^+$ the point p^* expires. Hence, at time $t_{p^*}^+$ the subgroup $G_{i^*}^{j^*,\ell^*}$ has only z points, which can all be designated as outliers. Hence, C_{i^*} minus the outliers can be covered by a ball of diameter $L^* + 2 \cdot 2^{j^*-1}z$ and so $\text{OPT}_{k,z}(P(t_{p^*}^+)) = (L^* + 2 \cdot 2^{j^*-1}z)/2$.

Since the algorithm must report a valid solution at time $t_{p^*}^-$ and it must report the same solution at time $t_{p^*}^+$, the approximation ratio at time $t_{p^*}^+$ is at least

$$\frac{\text{OPT}_{k,z}(P(t_{p^*}^-))}{\text{OPT}_{k,z}(P(t_{p^*}^+))} \geq \frac{L^* + 3 \cdot 2^{j^*-1}z}{L^* + 2 \cdot 2^{j^*-1}z} = 1 + \frac{2^{j^*-1}z}{L^* + 2^{j^*}z} > 1 + \frac{2^{j^*-1}z}{2^{j^*+2}z/\varepsilon} = 1 + \varepsilon/8$$

We conclude that any algorithm that gives a $(1 + \varepsilon')$ -approximation must store at least $k(z+1)(s-1)/(8\varepsilon')$ objects in the worst case.

The spread of the point set used in the construction, including the subgroup G_{new} , is⁶

$$\sigma = kL + (k-1)2L + 2^{j^*}z < 3kL < 3k \cdot 2^{s+1}z/\varepsilon.$$

⁶ In our definition of spread, we are allowed to reuse points, since $\Delta_{\min} = 0$ is defined as the minimum distance between distinct points in the stream. We can also avoid reusing points in our lower-bound scenario, without asymptotically influencing the spread of the points.

If the spread is not too small, namely when $\sigma \geq (3kz/\varepsilon)^2$, we have $s \geq \frac{1}{2} \log \sigma - 1$. We obtain the following theorem.

► **Theorem 6.** *Let ALG be a $(1 + \varepsilon)$ -approximation algorithm for the k -center problem with z outliers in \mathbb{R}^1 , with $z \geq 1$, that works in the model described above. For any spread $\sigma \geq (3kz/\varepsilon)^2$, there is a problem instance of spread σ that requires ALG to store $\Omega((kz \log \sigma)/\varepsilon)$ points.*

References

- 1 Pankaj K. Agarwal, Sariel Har-Peled, and Hai Yu. Robust shape fitting via peeling and grating coresets. *Discret. Comput. Geom.*, 39(1-3):38–58, 2008. doi:10.1007/s00454-007-9013-2.
- 2 Pankaj K. Agarwal and R. Sharathkumar. Streaming algorithms for extent problems in high dimensions. *Algorithmica*, 72(1):83–98, 2015. doi:10.1007/s00453-013-9846-4.
- 3 Matteo Ceccarello, Andrea Pietracaprina, and Geppino Pucci. Solving k -center clustering (with outliers) in mapreduce and streaming, almost as accurately as sequentially. *Proc. VLDB Endow.*, 12(7):766–778, 2019. doi:10.14778/3317315.3317319.
- 4 Timothy M. Chan and Vinayak Pathak. Streaming and dynamic algorithms for minimum enclosing balls in high dimensions. *Comput. Geom.*, 47(2):240–247, 2014. doi:10.1016/j.comgeo.2013.05.007.
- 5 Timothy M. Chan and Bashir S. Sadjad. Geometric optimization problems over sliding windows. *Int. J. Comput. Geom. Appl.*, 16(2-3):145–158, 2006. doi:10.1142/S0218195906001975.
- 6 Moses Charikar, Chandra Chekuri, Tomás Feder, and Rajeev Motwani. Incremental clustering and dynamic information retrieval. *SIAM J. Comput.*, 33(6):1417–1440, 2004. doi:10.1137/S0097539702418498.
- 7 Moses Charikar, Samir Khuller, David M Mount, and Giri Narasimhan. Algorithms for facility location problems with outliers. In *Proc. 12th Annual ACM-SIAM symposium on Discrete Algorithms (SODA)*, pages 642–651. Society for Industrial and Applied Mathematics, 2001.
- 8 Moses Charikar, Liadan O’Callaghan, and Rina Panigrahy. Better streaming algorithms for clustering problems. In *Proc. 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 30–39, 2003. doi:10.1145/780542.780548.
- 9 Vincent Cohen-Addad, Chris Schwiegelshohn, and Christian Sohler. Diameter and k -center in sliding windows. In *Proc. 43rd International Colloquium on Automata, Languages, and Programming, (ICALP)*, volume 55 of *LIPICs*, pages 19:1–19:12, 2016. doi:10.4230/LIPICs.ICALP.2016.19.
- 10 Hu Ding, Haikuo Yu, and Zixiu Wang. Greedy strategy works for k -center clustering with outliers and coresets construction. In *Proc. 27th Annual European Symposium on Algorithms (ESA)*, volume 144 of *LIPICs*, pages 40:1–40:16, 2019. doi:10.4230/LIPICs.ESA.2019.40.
- 11 Joan Feigenbaum, Sampath Kannan, and Jian Zhang. Computing diameter in the streaming and sliding-window models. *Algorithmica*, 41(1):25–41, 2005. doi:10.1007/s00453-004-1105-2.
- 12 Behnam Hatami and Hamid Zarrabi-Zadeh. A streaming algorithm for 2-center with outliers in high dimensions. *Comput. Geom.*, 60:26–36, 2017. doi:10.1016/j.comgeo.2016.07.002.
- 13 Richard Matthew McCutchen and Samir Khuller. Streaming algorithms for k -center clustering with outliers and with anonymity. In *Proc. 11th and 12th International Workshop on Approximation, Randomization and Combinatorial Optimization (APPROX and RANDOM)*, volume 5171 of *Lecture Notes in Computer Science*, pages 165–178, 2008. doi:10.1007/978-3-540-85363-3_14.
- 14 Hamid Zarrabi-Zadeh. Core-preserving algorithms. In *Proceedings of the 20th Annual Canadian Conference on Computational Geometry (CCCG)*, 2008.
- 15 Hamid Zarrabi-Zadeh and Asish Mukhopadhyay. Streaming 1-center with outliers in high dimensions. In *Proc. 21st Annual Canadian Conference on Computational Geometry (CCCG)*, pages 83–86, 2009. URL: http://cccg.ca/proceedings/2009/cccg09_22.pdf.