



PDF Download

3744252.pdf

28 December 2025

Total Citations: 0

Total Downloads: 172

Latest updates: <https://dl.acm.org/doi/10.1145/3744252>

RESEARCH-ARTICLE

Hybrid k-Clustering: Blending k-Median and k-Center

FEDOR V. FOMIN, University of Bergen, Bergen, Hordaland, Norway

PETR A GOLOVACH, University of Bergen, Bergen, Hordaland, Norway

TANMAY INAMDAR, Indian Institute of Technology Jodhpur, Jodhpur, RJ, India

SAKET SAURABH, University of Bergen, Bergen, Hordaland, Norway

MEIRAV ZEHAVI, Ben-Gurion University of the Negev, Beer Sheba, Southern District, Israel

Open Access Support provided by:

Indian Institute of Technology Jodhpur

University of Bergen

Ben-Gurion University of the Negev

Published: 28 November 2025

Online AM: 09 June 2025

Accepted: 12 May 2025

Revised: 16 April 2025

Received: 27 September 2024

[Citation in BibTeX format](#)

Hybrid k-Clustering: Blending k-Median and k-Center

FEDOR FOMIN, Department of Informatics, University of Bergen, Bergen, Norway

PETR GOLOVACH, Department of Informatics, University of Bergen, Bergen, Norway

TANMAY INAMDAR, Department of Computer Science and Engineering, IIT Jodhpur, Jodhpur, India

SAKET SAURABH, Institute of Mathematical Sciences, Chennai, India and Department of Informatics, University of Bergen, Bergen, Norway

MEIRAV ZEHAVI, Ben-Gurion University of the Negev, Beer-Sheva, Israel

We propose a novel clustering model encompassing two well-known clustering models: k -center clustering and k -median clustering. In the HYBRID k -CLUSTERING problem, given a set P of points in \mathbb{R}^d , an integer k , and a non-negative real r , our objective is to position k closed balls of radius r to minimize the sum of distances from points not covered by the balls to their closest balls. Equivalently, we seek an optimal L_1 -fitting of a union of k balls of radius r to a set of points in the Euclidean space. When $r = 0$, this corresponds to k -MEDIAN; when the minimum sum is zero, indicating complete coverage of all points, it is k -CENTER.

Our primary result is a bicriteria approximation algorithm that, for a given $\varepsilon > 0$, produces a hybrid k -clustering with balls of radius $(1 + \varepsilon)r$. This algorithm achieves a cost at most $1 + \varepsilon$ of the optimum, and it operates in time $2^{(kd/\varepsilon)^{O(1)}} \cdot n^{O(1)}$. Notably, considering the established lower bounds on k -center and k -median, our bicriteria approximation stands as the best possible result for HYBRID k -CLUSTERING.¹

CCS Concepts: • Theory of computation → Facility location and clustering; Fixed parameter tractability;

Additional Key Words and Phrases: clustering, k -center, k -median, Euclidean space, fpt approximation

ACM Reference Format:

Fedor Fomin, Petr Golovach, Tanmay Inamdar, Saket Saurabh, and Meirav Zehavi. 2025. Hybrid k -Clustering: Blending k -Median and k -Center. *ACM Trans. Comput. Theory* 17, 4, Article 23 (November 2025), 19 pages. <https://doi.org/10.1145/3744252>

¹A short version of this article was published in APPROX 2024.[18]

Fedor V. Fomin and Petr A. Golovach acknowledge support from the Research Council of Norway via the project BWCA (grant no. 314528). Tanmay Inamdar acknowledges support from IIT Jodhpur via the Research Initiation Grant I/RIG/TNI/20240072. Saket Saurabh acknowledges support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 819416); and he also acknowledges the support of Swarnajayanti Fellowship grant DST/SJF/MSA-01/2017-18. Meirav Zehavi acknowledges support from the European Research Council (ERC) grant no. 101039913 (PARAPATH).

Authors' Contact Information: Fedor Fomin, Department of Informatics, University of Bergen, Bergen, Vestland, Norway; e-mail: fedor.fomin@uib.no; Petr Golovach, Department of Informatics, University of Bergen, Bergen, Vestland, Norway; e-mail: petr.golovach@uib.no; Tanmay Inamdar, Department of Computer Science and Engineering, IIT Jodhpur, Jodhpur, India; e-mail: taninamdar@gmail.com; Saket Saurabh, Institute of Mathematical Sciences, Chennai, Tamilnadu, India and Department of Informatics, University of Bergen, Bergen, Hordaland, Norway; e-mail: saket@imsc.res.in; Meirav Zehavi, Ben-Gurion University of the Negev, Beer-Sheva, Israel; e-mail: meiravze@bgu.ac.il.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 1942-3454/2025/11-ART23

<https://doi.org/10.1145/3744252>

1 Introduction

Suppose we want to install a set of k **access points (APs)** at certain locations to provide wireless internet (Wi-Fi) coverage to a group of people belonging to a certain area. Each AP is capable of providing Wi-Fi within a circular-shaped region (i.e., a *disk*) of fixed radius r , and it may not be possible to cover the entire region with k such disks. Thus, after placing k APs, some people may be *outliers*, that lie outside any of the k disks and do not receive Wi-Fi coverage. We can model this scenario as the classical k -CENTER WITH OUTLIERS problem, which is a crude model since it only cares about the *number* of outliers. However, our scenario is more nuanced. All people that lie within any of the k disks of radius r already receive Wi-Fi, whereas a person lying outside all of the k disks must travel to the boundary of the nearest disk in order to receive coverage. Naturally, we would like to minimize the total distance traveled by people. Motivated by this and several other problems in computational geometry/clustering, we consider the following clustering problem, which encompasses two fundamental variants of clustering: k -CENTER and k -MEDIAN. Given a set P of points in some metric space and integer k and real $r \geq 0$, our objective is to position k closed balls of radius r in a way that minimizes the sum of distances from points uncovered by the balls to their closest balls. In Figure 1, we provide an example of such clustering with $k = 2$ and $r = 2$.

To define the new clustering formally, we need some definitions. We consider Euclidean inputs, i.e., all points belong to \mathbb{R}^d for some $d \geq 1$ and the distance function $\text{dist}(\cdot, \cdot)$ is given by the Euclidean (ℓ_2) distance. For a point $p \in P$ and a finite set of points $Q \subset \mathbb{R}^d$, we define $\text{dist}(p, Q) := \min_{q \in Q} \text{dist}(p, q)$. Further, for $x, y \in P$, and a real $r \geq 0$, we define the shorthand $\text{dist}_r(x, y) := \max\{\text{dist}(x, y) - r, 0\}$.

HYBRID k -CLUSTERING

Input. A set $P \subset \mathbb{R}^d$ of n points, an integer $k \geq 1$, and a real $r \geq 0$.

Task. Find a set $F \subset \mathbb{R}^d$ of size at most k , that minimizes:

$$\text{cost}_r(P, F) := \sum_{p \in P} \text{dist}_r(p, F) \quad (1)$$

We denote an instance of HYBRID k -CLUSTERING as $\mathcal{I} = (P, k, r, d)$, where d denotes the dimension. When $r = 0$, the optimal cost of HYBRID k -CLUSTERING equals the optimal k -MEDIAN clustering cost of the instance. Thus in this case, HYBRID k -CLUSTERING reduces to k -MEDIAN. However, when $r > 0$, $\text{dist}_r(\cdot, \cdot)$ does not form a metric, and hence we cannot simply reduce the problem to k -MEDIAN. On the other hand, the minimum value r that guarantees the cost of HYBRID k -CLUSTERING to be zero is equal to the optimal k -CENTER value. In this sense, HYBRID k -CLUSTERING reduces to k -CENTER.

1.1 Our Result and Techniques

The main result of this article is a bicriteria approximation algorithm for HYBRID k -CLUSTERING. An α -approximation to an instance $\mathcal{I} = (P, k, r, \text{dist})$ is a subset $F \subset \mathbb{R}^d$ of size k with $\text{cost}_r(P, F) \leq \alpha \cdot \text{OPT}_r$, where $\text{OPT}_r := \text{cost}_r(P, F^*)$ denotes the cost of an optimal solution $F^* \subset \mathbb{R}^d$ of size at most k . Furthermore, an (α, β) -bicriteria approximation is a solution $F \subset \mathbb{R}^d$ with $\text{cost}_{\beta r}(P, F) \leq \alpha \cdot \text{OPT}_r$. Here, $\text{cost}_{\beta r}(P, F) = \sum_{p \in P} \text{dist}_{\beta r}(p, F)$.

Consider the special case of $r = r^*$, where r^* is the optimal radius for k -CENTER. Then, $\text{OPT}_{r^*} = 0$. Therefore, a $(\alpha, 1)$ -bicriteria approximation would return a solution of $\text{cost} \alpha \cdot \text{OPT}_{r^*} = 0$ using radius $1 \cdot r^*$, i.e., an optimal solution for k -CENTER. On the other hand, a $(1, \beta)$ -bicriteria approximation, for the special case of $r = 0$, would return an optimal-cost solution using the radius of $\beta r = 0$.

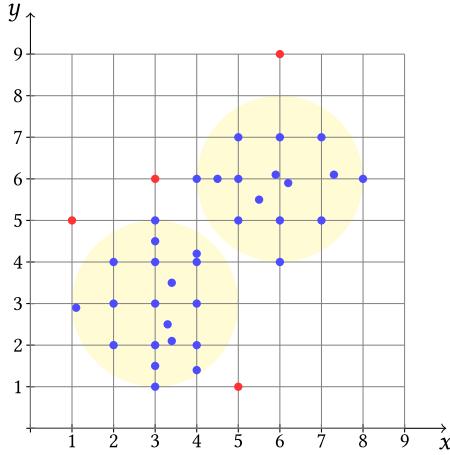


Fig. 1. Two disks of radius 2 cover all except four points that are colored red. The total sum of distances from these points to the yellow disks is $2(1 + \sqrt{8} - 2)$.

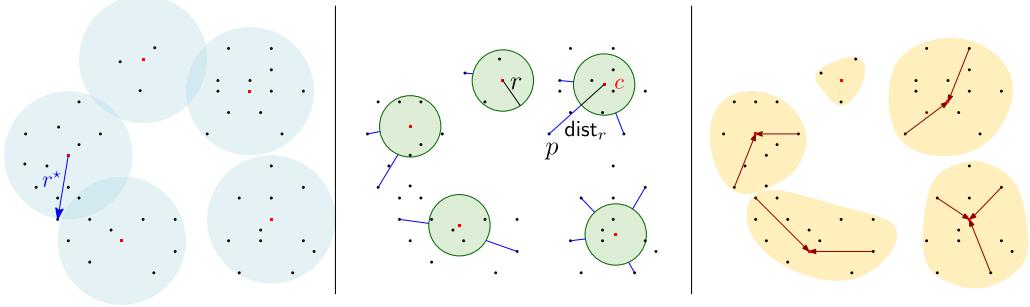


Fig. 2. Left: k -CENTER clustering, a special case of HYBRID k -CLUSTERING with $r = r^*$. All points are covered by k balls of radius r^* and $\text{OPT}_{r^*} = 0$. Right: k -MEDIAN clustering, a special case of HYBRID k -CLUSTERING with $r = 0$, and every point contributes its distance to the closest center (some are shown as brown arrows). Middle: A general instance of HYBRID k -CLUSTERING lies somewhere *in between* the two cases, where points outside radius- r balls contribute the distance to the boundary (shown in blue).

That is, such an algorithm would optimally solve k -MEDIAN. In Figure 2, we give an example of clustering the same point-set using k -CENTER, k -MEDIAN, and Hybrid HYBRID k -CLUSTERING. Combining these observations with the established lower bounds from the literature for k -MEDIAN and in Euclidean spaces, implies the following bounds for HYBRID k -CLUSTERING.

PROPOSITION 1.1. *The following holds for HYBRID k -CLUSTERING even when the input is from \mathbb{R}^2 .*

- *For any $\alpha \geq 1$, there exists no FPT in k algorithm that returns an $(\alpha, 1)$ -approximation, unless $\text{FPT} = \text{W}[1]$ [31].*
- *For any finite $\beta \geq 1$, there exists no polynomial-time algorithm that returns a $(1, \beta)$ -approximation unless $\text{P} = \text{NP}$ [33].*

Further, assuming the Exponential-Time Hypothesis (ETH), if the input is from \mathbb{R}^d with $d \geq 4$, then there exists no $n^{o(k)}$ time algorithm that returns a $(1, \beta)$ -approximation, for any finite $\beta \geq 1$ [11].

Given these results, a natural question arises: *Can we achieve a $(1 + \varepsilon, 1 + \varepsilon)$ -approximation for HYBRID k -CLUSTERING, running in time $f(k, \varepsilon) \cdot n^{O(1)}$, particularly in low-dimensional Euclidean spaces?* Our main theorem answers this question.

THEOREM 1.2. *Let $0 < \varepsilon < 1$. There exists a randomized algorithm that, given an instance of HYBRID k -CLUSTERING in \mathbb{R}^d , runs in time $2^{(\frac{kd}{\varepsilon})^{O(1)}} \cdot n^{O(1)}$, and returns a $(1 + \varepsilon, 1 + \varepsilon)$ -approximation with probability at least a positive constant.*

This randomized algorithm and the proof of correctness are described in Section 2. Here we discuss some of the main ideas. Recall that our objective, as the problem name suggests, is a “hybrid” of k -CENTER and k -MEDIAN. In our preprocessing steps, we first handle the inputs that behave *almost like* either of the two problems. Suppose we (approximately) know the optimal value of HYBRID k -CLUSTERING for the given set of points P , called OPT_r . First, in Lemma 2.1, if $r > \text{OPT}_r$, then we show that an approximate solution can be found using techniques used for approximating k -CENTER. Specifically, for each of the k centers in the optimal solution, we find a “nearby” center within distance εr via overlaying a fine grid in the space. Thus, we can assume that $r \leq \text{OPT}_r$. Next, we consider the case when r is *too small* compared to OPT_r , namely, when $r < \frac{\varepsilon \text{OPT}_r}{n}$, and show that in this case, the input behaves like k -MEDIAN – an approximate k -MEDIAN solution is also an approximation for HYBRID k -CLUSTERING (Lemma 2.2). In this manner, we preprocess to handle inputs that resemble k -CENTER and k -MEDIAN, we obtain a relation between r and OPT_r , which can be used to discretize the distances, which can be used to bound the aspect ratio (i.e., the ratio of maximum to minimum positive distance) (Lemma 2.3).

After the preprocessing step, we obtain inputs that are not *immediately reducible* to k -CENTER/MEDIAN. To handle such inputs, we design an intricate recursive algorithm that, at each step, tries to simultaneously handle parts (i.e., clusters) of the input that can be handled by either of the two techniques. This algorithm is inspired by the sampling approach of Kumar, Sabharwal, and Sen [28, 29] (also Jaiswal, Kumar, and Sen [27]). In this approach, one first takes a large enough sample that can be used to pin down the location of the largest cluster center. Then, one removes enough points from the vicinity of this center, so that the next largest cluster becomes dominant, and hence a subsequent sample contains sufficiently many points from the second cluster, and so on.

However, our scenario is more intricate and challenging for several reasons due to the peculiar nature of the objective. Nevertheless, in principle, one can classify each cluster as either being more *1-center-like*, or more *1-median like* (see Figure 3 for an illustration). In a *1-center-like* cluster, a large fraction of points lie within a ball of radius $O(r/\varepsilon)$. On the other hand, in a *1-median-like* cluster, a vast majority of points lie outside the $O(r/\varepsilon)$ -radius ball. Note that any such point loses very little due to the “ $-r$ ” term in the clustering cost, i.e., its dist_r and dist_s values are approximately equal. Hypothetically, if we knew the partition of the input points into k clusters, then we could use this classification to handle each type of cluster separately – an almost-optimal center of a *1-center-like* cluster can be found using a grid, whereas one can use an approximation for 1-median (as a black box) to handle a *1-median-like* cluster. However, the actual clusters are obviously unknown to the algorithm. Hence, the algorithm has to carefully navigate between the two types of clusters based on the random sample obtained, and must simultaneously handle both scenarios using branching (i.e., recursion). The analysis of the algorithm is also much more involved due to the various cases in which the distinction between two types of clusters is murkier. Nevertheless, we are able to show that the algorithm returns a $(1 + \varepsilon, 1 + \varepsilon)$ -approximation in time $2^{(kd/\varepsilon)^{O(1)}} \cdot n^{O(1)}$ with good probability. Note that we incur an exponential dependence on the dimension d due to “grid-arguments” used to handle *1-center-like* clusters, unlike the approach of [28]. However, such dependence seems unavoidable using our approach.

1.2 Related Problems

Euclidean Clustering. An extensive body of literature exists on approximation algorithms for k -CENTER and k -MEDIAN in the Euclidean space. For k -MEDIAN in \mathbb{R}^d , **Polynomial-Time**

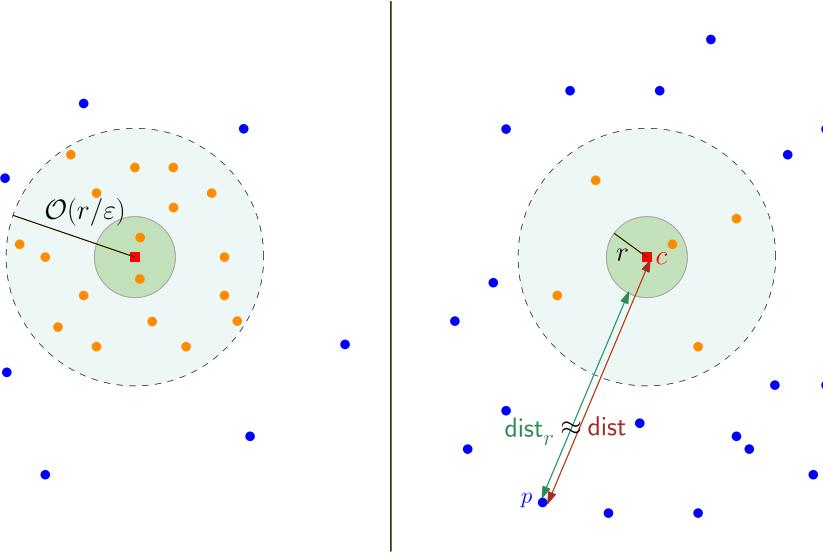


Fig. 3. Example of two different types of clusters. In each figure, we show the cluster center in red, a ball of radius r around the center in green, and a larger ball of radius $O(r/\epsilon)$ in cyan with a dashed outline. Left: A 1-center-like cluster. Note that a large chunk of points lies within the radius $O(r/\epsilon)$ ball around the center. Right: A 1-median-like cluster. Note that most of the points lie outside the $O(r/\epsilon)$ radius ball around c , and for any such point, e.g., p that is outside the $O(r/\epsilon)$ radius ball, $\text{dist}_r(p, c) \approx \text{dist}(p, c)$.

Approximation Schemes (PTASes) with a running time of $n^{f(\epsilon, d)}$ have been developed, leveraging local search techniques [13, 19]. Additionally, various **Fixed-Parameter Tractable Approximation Schemes (FPT-AS)** with a running time of $f(k, d, \epsilon) \cdot n^{O(1)}$ are known for this problem [12, 14, 27, 29]. The dependence on dimension d can be eliminated through dimensionality reduction techniques [10, 30].

For k -CENTER, an FPT-AS was introduced by Agarwal and Procopiuc in [3], with a runtime of $O(n \log k) + (k/\epsilon)^{O(dk^{1-1/d})}$ in \mathbb{R}^d . Subsequent work by Badoiu, Har-Peled, and Indyk [4] improved the running time to $2^{O(k \log k)/\epsilon^2} \cdot dn$.

Common generalizations of k -Center and k -Median. In [34], Tamir introduces a common generalization of the two clustering problems, namely, ℓ -CENTRUM. In this problem, one ignores ℓ closest points from the cost. Notably, k -MEDIAN ignores 0 points, and k -CENTER ignores all but one point. While this problem greatly resembles HYBRID k -CLUSTERING, it is not immediate whether one can give a (approximation-preserving) reduction from HYBRID k -CLUSTERING to ℓ -CENTRUM. Nevertheless, some of the techniques used in the literature for obtaining approximations for the latter problem, e.g., [8], can be used to obtain bicriteria approximations for HYBRID k -CLUSTERING, as we observe later in the conclusion. ORDERED k -MEDIAN is a further generalization of ℓ -CENTRUM, where the solution is evaluated by first sorting the client connection costs and then multiplying them with a predefined non-increasing weight vector (higher connection costs are taken with larger weights) [7]. Finally, the latter problem has been further generalized to the objectives that contain a general norm of the distance vector of the clients' connection costs [1]. Approximation algorithms for these and other generalizations have been developed in [1, 5, 7, 8, 26].

k -center clustering with outliers. In k -CENTER clustering, we are given sets P (clients) and \mathbb{F} (facilities) of points. Given an integer k , the task is to identify k centers $F \subseteq \mathbb{F}$ minimizing the

maximum distance of any point in P from its closest center. A popular variant of k -center is a formulation that considers outliers. For a selected parameter x , up to x points are allowed not to be allocated to any center. A plethora of approximation algorithms for this problem, and the related problems of covering points by disks and minimum enclosing balls with outliers, exist in the literature [4, 9, 15–17, 21, 24, 32]. HYBRID k -CLUSTERING could be seen as a variant of k -CENTER with outliers, where we focus on the sum of distances to outliers rather than their numbers.

Shape fitting. A natural problem arising in machine learning, statistics, data-mining, and many other fields is to fit a shape γ to a set of points P in \mathbb{R}^d . Har-Peled in [23] introduces the following formalization of this problem. For a family of shapes \mathcal{F} (points, lines, hyperplanes, spheres, etc.) we seek for a shape $\gamma \in \mathcal{F}$ with the best fit to P . The typical criteria for measuring how well a shape γ fits a set of points P could be the maximum distance between a point of P and its nearest point on γ (L_∞ -fitting), sum of the distances from P to γ (L_1 -fitting) or the sum of the squares of the distances (L_2 -fitting). In this setting, HYBRID k -CLUSTERING is the problem of L_1 -fitting to a shape from \mathcal{F} , where \mathcal{F} is the family of shapes defined by unions (not necessarily disjoint) of k balls in \mathbb{R}^d . Some relevant work in this direction includes [2, 23, 25, 35].

2 Bicriteria FPT Approximation Scheme in Euclidean Spaces

We first set up some notation and define an important subroutine. For $Y \subset \mathbb{R}^d$, and $y \in Y$, let $\text{cl}(y, Y) \subseteq P$ denote the subset of points of P , whose closest point in Y is y . Ties are broken arbitrarily. Note that $\{\text{cl}(y, Y) : y \in Y\}$ forms a partition of P .

Let $p \in \mathbb{R}^d$ be a point and $\lambda \geq 0$, let $B(p, \lambda) = \{q \in \mathbb{R}^d : \text{dist}(p, q) \leq \lambda\}$ denote the ball of radius λ centered at p . For $0 \leq \tau \leq \lambda$, let $\text{Grid}(p, \lambda, \tau)$ be the outcome of the following procedure: we place a grid of sidelength τ/\sqrt{d} (of arbitrary offset). From each grid cell L that (partially) intersects with $B(p, \lambda)$ (i.e., L contains a point q with $\text{dist}(p, q) \leq \lambda$), we pick an arbitrary point from L and add it to the set $\text{Grid}(p, \lambda, \tau)$. Note that $\text{Grid}(p, \lambda, \tau)$ can be computed in time proportional to the size of the output and d . We have the following observations that follow from simple geometric arguments.

OBSERVATION 1.

- (1) $|\text{Grid}(p, \lambda, \tau)| \leq O((2\sqrt{d}\lambda/\tau)^d)$, where d is the dimension.
- (2) For any $q \in B(p, \lambda)$, there exists some $q' \in \text{Grid}(p, \lambda, \tau)$ such that $\text{dist}(q, q') \leq \tau$.

2.1 Preprocessing

Suppose we know an estimate of OPT_r up to a constant factor – this can be done by an exponential search or by first finding a bicriteria (constant) approximation. For simplicity of exposition, we assume that we know OPT_r exactly.

Step 1. Obtaining $\text{OPT}_r \geq r \geq \frac{\varepsilon \text{OPT}_r}{2n}$.

First, in the following lemma, we handle *k*-center-like instances, which we can handle using “grid arguments”. If this is not applicable, we obtain that $r \leq \text{OPT}_r$.

LEMMA 2.1. *If $r > \text{OPT}_r$, then in time $(\frac{d}{\varepsilon})^{O(dk)} \cdot n^{O(1)}$ one can find a set $F \subset \mathbb{R}^d$ of size k , such that, $\text{cost}_{(1+\varepsilon)r}(P, F) \leq (1 + \varepsilon)\text{OPT}_r$.*

PROOF. Consider an optimal solution F^* such that $\text{cost}_r(P, F^*) = \text{OPT}_r$. We assume that F^* is *minimal*, i.e., no strict subset of F^* also yields the optimal cost. It follows that every point $p \in P$ is within a distance of at most $2r$ from some $c \in F^*$, and each $c \in F^*$ has some $p \in P$ within distance $2r$. Thus, P can be covered using at most k balls of radius $2r$. Using a polynomial-time 2-approximation for k -CENTER (e.g. [22]), we find a set F such that $P \subseteq \bigcup_{c \in F} B(c, 4r)$. This implies that $F^* \subseteq \bigcup_{c \in F} B(c, 6r)$.

Let $R = \bigcup_{c \in F} \text{Grid}(c, 6r, \varepsilon r)$. By Observation 1, $|R| \leq k(\frac{\sqrt{d}}{\varepsilon})^{O(d)}$. Now, we iterate over all subsets $F' \subseteq R$ of size at most k , and look at the subset F' minimizing $\text{cost}_r(P, F')$. The number of such subsets is at most $k \binom{|R|}{k} \leq k(\frac{\varepsilon |R|}{k})^k \leq (\frac{d}{\varepsilon})^{O(dk)}$. Also, by Observation 1, there exists a subset $F \subseteq R$ of size at most k , such that for each $c^* \in F^*$, there exists some $c' \in F$ such that $\text{dist}(c^*, c') \leq \varepsilon r$. It is easy to see that F satisfies the claimed properties. \square

In the following lemma, we handle *k-median-like* instances, where r is very small compared to OPT_r . We directly reduce such instances to *k-MEDIAN* (where $r = 0$). If this is not applicable, then we obtain that r is not “too small” compared to OPT_r .

LEMMA 2.2. *Let $0 < \varepsilon < 1$. Suppose for an instance \mathcal{I} , $\text{OPT}_r \geq \frac{2nr}{\varepsilon}$. Then, $\text{OPT}_0 \leq (1 + \varepsilon/2) \cdot \text{OPT}_r \leq (1 + \varepsilon/2) \cdot \text{OPT}_0$. Furthermore, if $F \subset \mathbb{R}^d$ satisfies that $\text{cost}_0(P, F) \leq (1 + \varepsilon/3) \cdot \text{OPT}_0$. Then, $\text{cost}_r(P, F) \leq (1 + \varepsilon) \cdot \text{OPT}_r$. Such a set F can be found in time $2^{(k/\varepsilon)^{O(1)}} \cdot nd$.*

PROOF. For any $F' \subset \mathbb{R}^d$, it holds that $\text{cost}_0(P, F') \leq \text{cost}_r(P, F') + nr$. In particular, let $F^* \subset \mathbb{R}^d$ be an optimal solution of size k , i.e., $\text{cost}_r(P, F^*) = \text{OPT}_r$. Then, it holds that $\text{cost}_0(P, F^*) \leq \text{cost}_r(P, F^*) + n \cdot \frac{\varepsilon \text{OPT}_r}{2n} = (1 + \varepsilon/2) \cdot \text{OPT}_r$. Since $\text{cost}_0(P, F^*) \geq \text{OPT}_0$, since F^* is a feasible solution of size k for $r = 0$. This shows the first inequality.

Now, let $F \subset \mathbb{R}^d$ be a $(1 + \varepsilon/3)$ -approximate solution of size k for cost_0 , as in the statement of the lemma, i.e., $\text{cost}_0(P, F) \leq (1 + \varepsilon/3) \cdot \text{OPT}_0$. Then, the first inequality implies that $\text{cost}_0(P, F) \leq (1 + \varepsilon/3) \cdot (1 + \varepsilon/2) \cdot \text{OPT}_r \leq (1 + \varepsilon) \cdot \text{OPT}_r$. We can use a $(1 + \varepsilon/3)$ -approximation algorithm (e.g., [28, 29]) for *k-MEDIAN* to find such a solution. \square

For a given input P , we try the procedures from Lemma 2.1 and 2.2 and keep them as candidate solutions. However, if P does not satisfy the conditions required to apply these lemmas, then we must have that $\frac{\varepsilon \text{OPT}_r}{2n} \leq r \leq \text{OPT}_r$. In this case, we use the next step before proceeding to the main algorithm.

Step 2. Bounding the aspect ratio.

In this step, we suitably discretize the distances in order to bound the aspect ratio of the metric (i.e., the maximum ratio of inter-point distances) by $\mathcal{O}(\frac{n^2}{\varepsilon})$. This procedure preserves the cost of an optimal solution up to a factor of $1 + \varepsilon$.

LEMMA 2.3. *Let P be a set of points satisfying $\frac{\varepsilon \text{OPT}_r}{2n} \leq r \leq \text{OPT}_r$. Then, in polynomial time we can obtain another (multi)set of points P' such that, for any solution $F \subset \mathbb{R}^d$, $\text{cost}_r(P', F) \in (1 \pm \varepsilon) \cdot \text{cost}_r(P, F)$, and $\frac{\max_{p, q \in P'} \text{dist}(p, q)}{\min_{p, q \in P': \text{dist}(p, q) \neq 0} \text{dist}(p, q)} \leq \frac{4n^2}{\varepsilon}$.*

PROOF. Imagine an auxiliary graph $G = (P, E)$, where $pq \in E$ iff $\text{dist}(p, q) \leq 2(\text{OPT}_r + r) \leq 4\text{OPT}_r$. If G has more than one connected component, then, note that two points belonging to different connected components cannot belong to the same optimal cluster. Hence, we can solve the problem separately on the points belonging to different connected components and combine the solutions by a simple dynamic programming.

Thus, we can handle each connected component of G separately. In any connected component, the maximum distance between any two points is at most $n \cdot 4\text{OPT}_r$. Now, we place a grid of sidelength $\frac{\varepsilon \text{OPT}_r}{\sqrt{dn}}$ and move each point $p \in P$ to the center of the grid. It can be easily shown that this process does not change the cost of an optimal solution by more than a $(1 + \varepsilon)$ factor. Thus, the smallest non-zero inter-point distance is now at least $\frac{\varepsilon \text{OPT}_r}{n}$. Thus, the aspect ratio of the instance is bounded by $\frac{4n^2}{\varepsilon}$. \square

ALGORITHM 1: HybridClustering(F', k, m)

$F' \subseteq \mathbb{R}^d$ is a subset of centers of size at most $k - m$ added to the solution so far
 $\beta = \frac{1}{\delta^{\epsilon'}}$ as required in Proposition 2.4 and $\beta' := \beta \cdot \frac{150k}{\delta^3}$.

```

1: if  $m = 0$  then
2:   return  $F'$ 
3: end if
4:  $R \leftarrow \bigcup_{c' \in F'} \text{Grid}(c', 16r, \delta r)$ 
5: for each  $q$  of the form  $2^j$  in the range  $[8r, \text{dist}_{\max}]$  do
6:    $P_q := P \setminus (\bigcup_{c' \in F'} B(c', q))$ 
7:   Let  $S_q$  be a sample of size  $\beta'$  chosen uniformly at random from  $P_q$ 
8:    $R \leftarrow R \cup \bigcup_{p \in S_q} \text{Grid}(p, \frac{8r}{\delta}, \delta r)$ 
9:   for each  $S \subseteq S_q$  of size  $\beta$  do
10:     $c' \leftarrow \text{ApproxSolutionOnSample}(S, \delta/8)$                                  $\triangleright$  Algorithm from Proposition 2.4
11:     $R \leftarrow R \cup \{c'\}$ 
12:   end for
13: end for
14: for each  $c \in R \setminus F'$  do
15:   Call HybridClustering( $F' \cup \{c\}, k, m - 1$ )
16: end for
17: Call HybridClustering( $F', k, m - 1$ )
18: return solution  $\tilde{F}$  minimizing  $\text{cost}_{r'}(P, \tilde{F})$  over recursive calls made in lines 15 and 17

```

Bounding the aspect ratio by $O(\frac{n^2}{\epsilon})$ means that an exponential search over distances has at most $\log_2(\frac{n^2}{\epsilon}) = O(\frac{\log(n)}{\epsilon})$ levels, which will be useful in our main algorithm. By slightly abusing the notation, we continue to use P for referring to the discretized (multi)set P' returned by Lemma 2.3. If there are any co-located points in P , we will treat them as separate points, and hence use set terminology instead of multiset terminology.

After the two preprocessing steps, we now proceed to the description of the main algorithm.

2.2 Main Algorithm

Our goal is to prove Theorem 1.2, that is, to design a randomized bicriteria FPT approximation for HYBRID k -CLUSTERING. We define some parameters. Let $\delta := \frac{\epsilon}{10k} < \frac{1}{2}$, $\delta' := \frac{\delta}{3}$ and $r' := (1 + \delta')r$.

Algorithm 1 is a recursive algorithm, and is called HybridClustering. It takes three parameters F' , k , and m . $F' \subset \mathbb{R}^d$ is a subset of centers added to the solution so far and has size $k - m$. Further, k is the total size of the solution, and m is an upper bound on the remaining solution (since we have already added $k - m$ centers). At a high (and imprecise) level, the goal of each recursive step is to find an *approximate replacement* for each center in an unknown optimal solution.

In line 2, we check whether $m = 0$, i.e., whether we have used our budget of k centers, and if so, we return the same set F' of centers built through the recursive process. Otherwise (line 4 onward), we assume that $m > 0$, i.e., we are yet to add a set of centers. Throughout this process (line 4 to 13), we will build a set R consisting of candidate centers, at least one of which will be an approximate replacement of an *unseen* center (i.e., one whose approximate replacement has not already been found) from an optimal solution. Finally, in line 15, we will make a recursive call by adding each candidate to the current solution F' . Now we discuss how we build the set R .

First, in line 4, for each center $c' \in F'$ added so far, we add a set of “nearby” centers by placing a grid. This handles the case when an unseen optimal center is close to one of the already chosen centers in F' . Next, in the outer **for** loop (line 5 to 13), we handle the case when all new optimal centers are relatively far from the already chosen centers. In this **for** loop, we iterate over a range

of values for the parameter q via exponential search. Parameter q tries to approximate half of the minimum distance between the already chosen and new optimal centers. Thus, for the “correct” value of q , the set of points C_q lying “far” from the centers of F' (line 6), leaves all of the m unseen optimal clusters untouched. At this point, we aim to use a sample of faraway size (chosen in line 7), to find an approximate replacement for one of these m unseen centers. We do this by using the sample in two different ways, to handle two different situations. First, if our sample happens to contain a point “nearby” an unseen center, say c^* , then the points chosen from the fine grid in line 8 will find such an approximate replacement for c^* . Otherwise, the idea is that, if we have removed a significant fraction of points from the “seen” clusters in line 6, by virtue of being close to F' , then the sample contains sufficiently many (i.e., at least β) points from the largest unseen cluster, say C^* , with reasonable probability, and these points can be used to find an approximate replacement of the cluster center (using Proposition 2.4). However, *a priori* we do not know which subset of the sample comes from C^* . Therefore, we iterate over all subsets of size β in the inner **for** loop (lines 9 to 12) to find such a subset of size β that comes entirely from C^* and use a known subroutine, called `ApproxSolutionOnSample`, to find an approximate replacement. Finally, in 15, we make a recursive call by adding each center from $R \setminus F'$, and in line 17, we make a recursive call by not adding any new center (to handle a particular case). In line 18, we return the minimum-cost solution found over all recursive calls. This completes the description of the algorithm.

2.3 Analysis

The crux of the analysis is to establish that Algorithm 1 satisfies the following invariant.

Invariant. Let $0 \leq m \leq k$. Suppose for the given F' of size at most $k - m$, there exists some $F = F' \uplus F_o \subset \mathbb{R}^d$, such that

$$(1) |F_o| \leq m, \text{ and}$$

$$(2)$$

$$\sum_{c \in F'} \text{cost}_{r'}(\text{cl}(c, F), c) + \sum_{c \in F_o} \text{cost}_r(\text{cl}(c, F), c) \leq (1 + \delta)^{k-m} \cdot \text{OPT}_r \quad (2)$$

Then, there exists some $0 < \alpha < 1$ (that does not depend on m), such that with probability at least α^m , the algorithm returns a solution $\tilde{F} \subset \mathbb{R}^d$, such that

$$(1) |\tilde{F}| \leq k,$$

$$(2) F' \subseteq \tilde{F}, \text{ and}$$

$$(3)$$

$$\sum_{c \in \tilde{F}} \text{cost}_{r'}(\text{cl}(c, F), c) \leq (1 + \delta)^k \cdot \text{OPT}_r. \quad (3)$$

Proof of Correctness. The proof is by induction on m . For the base case, consider $m = 0$. In the base case (Algorithm 1), we return the same $F' = F$ with probability one. In this case, the invariant tells us that $\text{cost}_r(P, F) \leq (1 + \delta)^k \cdot \text{OPT}_r$, which is what we need to prove. Now we assume that the claim is true for some $m - 1 \geq 0$ and we prove it for m by considering different cases.

Easy case: $F = F'$. This is a much simpler case since we have already found the desired set. In this case, any solution \tilde{F} returned by a recursive call always contains $F = F'$ as a subset. Then, in this case, we have that:

$$\text{cost}_{r'}(P, \tilde{F}) \leq \text{cost}_{r'}(P, F) \leq (1 + \delta)^{k-m} \cdot \text{OPT}_r \leq (1 + \delta)^k \cdot \text{OPT}_r.$$

Here, the first inequality follows from the assumption that $\tilde{F} \supseteq F = F'$, and the second inequality follows from (2) of the invariant. Note that we do not need to rely on the induction here.

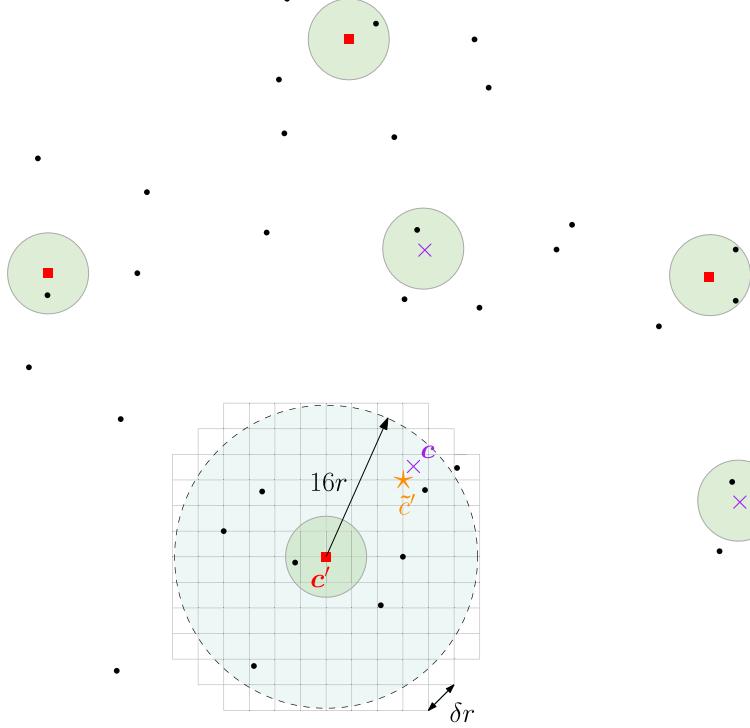


Fig. 4. Illustration for Case 1. Centers in F' are shown as red squares and unseen centers of $F \setminus F'$ are shown as purple crosses. c is the closest center to F' and $\text{dist}(c, c') \leq 16r$. Then, a nearby center \tilde{c}' can be found using a δr grid.

Main case: $F' \subsetneq F$. This is the case where we are yet to discover some subset (namely, $F \setminus F'$) of centers. We will analyze this case by considering different scenarios based on the inter-center distances, as well as their relative sizes.

First, since $F' \subsetneq F$, there exists some $c \in F \setminus F'$. Now, let $c \in F \setminus F'$ and $c' \in F'$ be the pair of centers with the smallest distance, i.e., (c, c') is a pair realizing $\min_{c_1 \in F \setminus F', c_2 \in F'} \text{dist}(c_1, c_2)$. Now we consider different cases depending on $\text{dist}(c, c')$, namely the closest distance between an already chosen center $c' \in F'$, and an “unseen center” $c \in F \setminus F'$.

Case 1. Nearby center: $\text{dist}(c, c') \leq 16r$. In this case, via Observation 1, we conclude that there exists some $\tilde{c} \in \text{Grid}(c', 16r, \delta r)$ with $\text{dist}(\tilde{c}, c) \leq \delta r$. Let $\tilde{F} = F' \cup \{\tilde{c}\}$. Then, the proof follows from the following claim (see Figure 4).

CLAIM 1. *Let $c_1 \in F_o$ and let $\tilde{c}_1 \in \mathbb{R}^d$ be such that $\text{dist}(c_1, \tilde{c}_1) \leq \delta'r$. Then, with probability at least α^{m-1} , $\text{HybridClustering}(F \cup \tilde{c}_1, k, m-1)$ returns a solution \tilde{F} that satisfies the required properties.*

PROOF. Consider c_1, \tilde{c}_1 as defined in the statement. Define $F_{\text{new}} := F'_{\text{new}} \uplus F'_o$, where $F'_{\text{new}} := F' \cup \{\tilde{c}_1\}$ and $F'_o := F_o \setminus \{c_1\}$. First, we show the following inequality:

$$\begin{aligned} \text{cost}_{r'}(\text{cl}(c_1, F), \tilde{c}_1) + \sum_{c \in F'_{\text{new}} \setminus \{\tilde{c}_1\}} \text{cost}_{r'}(\text{cl}(c, F), c) + \sum_{c \in F'_o} \text{cost}_r(\text{cl}(c, F_{\text{new}}), c) \\ \leq \sum_{c \in F'} \text{cost}_{r'}(\text{cl}(c, F), c) + \sum_{c \in F_o} \text{cost}_r(\text{cl}(c, F), c) \quad (4) \end{aligned}$$

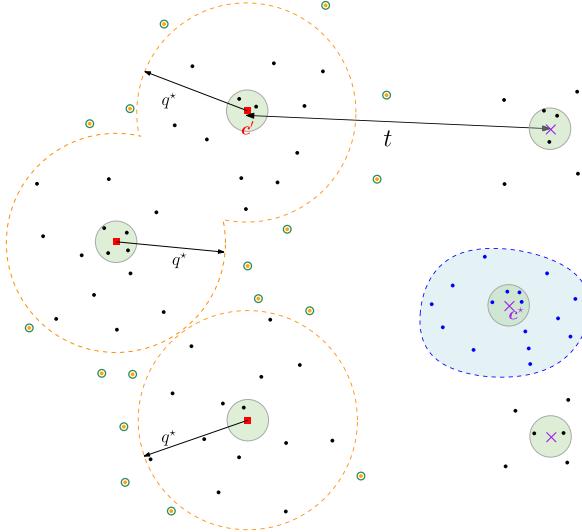


Fig. 5. Illustration for Case 2. Centers of F' are shown as red squares and unseen centers of $F \setminus F'$ are shown as purple crosses. Balls of radius q^* around F' are shown in dashed orange. P' are the points lying outside these balls. Among the points of P' , D is the set of points belonging to clusters around F' , and shown as green-orange filled dots. Finally, the cluster around c^* is the largest unseen cluster (marked in dashed blue shape), L . We analyze different cases depending on the relative sizes of L and D .

We construct an assignment of clients to the centers in F_{new} , where we may not assign a client to its closest center. To construct this assignment, we consider different cases. Firstly, for any point $p \in \text{cl}(c_1, F)$, it holds that $\text{dist}(p, \tilde{c}_1) \leq \text{dist}(p, c_1) + \text{dist}(c_1, \tilde{c}_1) \leq \text{dist}(p, c_1) + \delta'r$, which implies that $\text{cost}_{r'}(\text{cl}(\tilde{c}_1, P)) \leq \text{cost}_r(\text{cl}(c_1, P))$. Next, for $c \in F' \cup F_o \setminus \{c_1\}$, we assign all points $p \in \text{cl}(c, F)$ to c . The contribution of all such points is the same as the right-hand side of (4). Finally, we assign all points in $\text{cl}(c_1, F)$ to \tilde{c}_1 . By the choice of \tilde{c}_1 , $\text{cost}_{r'}(\text{cl}(c_1, F), \tilde{c}_1) \leq \text{cost}_r(\text{cl}(c_1, F), c_1)$, which is the contribution of such points on the right-hand side. Since the cost on the left-hand side is no larger than the cost of the assignment thus constructed, it shows (4).

Note that the right-hand side of (4) is at most $(1 + \delta)^{k-m} \cdot \text{OPT}_r$ due to the invariant, and hence F_{new} satisfies the properties required to apply the inductive hypothesis for $m - 1$. This implies that with probability at least α^{m-1} the recursive call $\text{Recursive}(F' \cup \{\tilde{c}_1\}, k, m - 1)$ returns a solution \tilde{F} satisfying $\text{cost}_{r'}(P, \tilde{F}) \leq (1 + \delta)^k \cdot \text{OPT}_r$. \square

Case 2. Faraway center: $16r < \text{dist}(c, c^*) \leq \text{dist}_{\max}$. Let $t = \text{dist}(c, c^*)$ and q^* be the largest power of 2 that is at most $t/2$. Consider $P_{q^*} = P \setminus (\bigcup_{c_1 \in F' \setminus F'} B(c_1, q^*))$. Let $c^* \in F \setminus F'$ denote the center of the maximum-size cluster, i.e., $c^* = \arg \max_{c_1 \in F \setminus F'} |\text{cl}(c_1, F)|$, and $L := \text{cl}(c^*, F)$ denote the largest cluster. Finally, let $D := \bigcup_{c_{\text{old}} \in F'} \text{cl}(c_{\text{old}}, F) \cap P_{q^*}$ denote the set of clients that are distant from the respective centers in F' . Let us summarize some consequences of these definitions in the following observation (its proof is essentially discussed above). Also see Figure 5.

OBSERVATION 2.

- (1) $P_{q^*} = D \uplus \biguplus_{c_1 \in F \setminus F'} \text{cl}(c_1, F) \cap P_{q^*}$.
- (2) In particular, $\text{cl}(c, F), \text{cl}(c^*, F) \subseteq P_{q^*}$.

We consider different sub-cases based on the relative sizes of L and D .

Case 2.1. New cluster is tiny: $|L| \leq \delta^2/4 \cdot |D|$. Let $N := \text{cl}(c, F)$. Note that the definition of c^* , combined with the case assumption, implies that $|N| \leq |L| \leq \frac{\delta^2}{4} \cdot |D|$. We summarize a few technical consequences of these definitions in the following claim.

CLAIM 2.

$$\text{cost}_r(N, c') \leq \delta \cdot \text{cost}_r(D, F) + (1 + \delta) \cdot \text{cost}_r(N, c). \quad (5)$$

PROOF. Note that for each $p \in D$, $\text{dist}(p, F) \geq \frac{t}{4}$. Thus, $\text{dist}_r(p, F) \geq \frac{t}{4} - r \geq \frac{3t}{16}$, where the last inequality follows from the case assumption, namely $t > 16r$. Thus, each point $p \in D$ contributes at least $\frac{3t}{16}$ to $\text{cost}_r(P, F)$, and their total contribution to $\text{cost}_r(P, F)$ is

$$\text{cost}_r(D, F) := \sum_{p \in D} \text{dist}_r(p, F) \geq |D| \cdot \frac{3t}{16}. \quad (6)$$

Now we upper bound the cost of assigning points of N to c' . To this end, we partition $N = N_{\text{near}} \uplus N_{\text{far}}$, where $N_{\text{near}} := \{p \in N : \text{dist}(p, c) \leq 2r/\delta\}$ and $N_{\text{far}} := \{p \in N : \text{dist}(p, c) > 2r/\delta\}$. Note that, for each $p \in N_{\text{far}}$, $\text{dist}_r(p, c) = \text{dist}(p, c) - r \geq \text{dist}(p, c) - \frac{\delta}{2} \cdot \text{dist}(p, c)$, which implies that, for each $p \in N_{\text{far}}$,

$$\text{dist}(p, c) \leq \left(\frac{1}{1 - \frac{\delta}{2}} \right) \cdot \text{dist}_r(p, c) \leq (1 + \delta) \cdot \text{dist}_r(p, c). \quad (7)$$

Now consider,

$$\begin{aligned} \text{cost}_r(N, c') &\leq \sum_{p \in N} \text{dist}(p, c') && \text{(Since } \text{dist}_r(\cdot, \cdot) \leq \text{dist}(\cdot, \cdot) \text{)} \\ &\leq \sum_{p \in N} \text{dist}(p, c) + \text{dist}(c, c') && \text{(Triangle inequality)} \\ &= |N| \cdot \text{dist}(c, c') + \sum_{p \in N_{\text{near}}} \text{dist}(p, c) + \sum_{p \in N_{\text{far}}} \text{dist}(p, c) \\ &\leq \frac{\delta^2}{4} \cdot |D| \cdot t + \sum_{p \in N_{\text{near}}} \frac{2r}{\delta} + \sum_{p \in N_{\text{far}}} (1 + \delta) \cdot \text{dist}_r(p, c) && \text{(From case assumption and (7))} \\ &\leq \frac{\delta^2}{4} \cdot |D| \cdot t + \frac{2r}{\delta} \cdot |N_{\text{near}}| + (1 + \delta) \cdot \text{cost}_r(N, c) \\ &\leq \frac{\delta^2}{4} \cdot |D| \cdot t + \delta \cdot |D| \cdot \frac{t}{32} + (1 + \delta) \cdot \text{cost}_r(N, c) \\ &&& \text{(Since } |N_{\text{near}}| \leq |N| \leq \frac{\delta^2}{4} |D| \text{ and } t > 16r \text{)} \\ &\leq |D| \cdot \frac{3t}{16} \cdot \delta \cdot \left(\frac{4\delta}{3} + \frac{1}{3} \right) + (1 + \delta) \cdot \text{cost}_r(N, c) \\ &\leq \delta \cdot \text{cost}_r(D, F) + (1 + \delta) \cdot \text{cost}_r(N, c). \end{aligned} \quad (8)$$

where the last inequality follows from (6) and $\delta < 1/2$. \square

Thus, consider the solution $F \setminus \{c\}$. To upper bound $\text{cost}_r(P, F \setminus \{c\})$, we assign all points in $N = \text{cl}(c, F)$ to c' . The cost of this solution can be upper bounded as follows:

$$\begin{aligned} \text{cost}_r(P, F \setminus \{c\}) &\leq \text{cost}_r(P, F) - \text{cost}_r(N, c) + \text{cost}_r(N, c') \\ &\leq \text{cost}_r(P, F) + \delta \cdot \text{cost}_r(D, F) + \delta \cdot \text{cost}_r(N, c) && \text{(From (5))} \\ &\leq \text{cost}_r(P, F) + \delta \cdot \text{cost}_r(P, F) && \text{(Since } D \uplus N \subseteq P \text{)} \\ &\leq (1 + \delta)^{k-m+1} \cdot \text{OPT}_r. \end{aligned} \quad (9)$$

where the last inequality follows from the invariant. Further, observe that $c \in F \setminus F'$, which implies that $|(F \setminus \{c\}) \setminus F'| \leq m - 1$. This, combined with (9), shows that the solution $F \setminus \{c\} = F' \uplus (F_o \setminus \{c\})$ satisfies the conditions of the invariant for $m - 1$. Then, by using inductive hypothesis, $\text{HybridClustering}(F, k, m - 1)$, with probability at least $\alpha^{m-1} \geq \alpha^m$, returns a solution \tilde{F} such that (a) $|\tilde{F}| \leq k$, (b) $F \subseteq \tilde{F}$, and (c) $\text{cost}_{(1+\delta)r'}(F, \tilde{F}) \leq (1 + \delta)^k \cdot \text{OPT}_r$, completing the induction.

Case 2.2. New cluster is large enough: $|L| > \delta^2/4 \cdot |D|$. That is, the largest “untouched” cluster is at least a $\delta^2/4$ fraction of the remaining points. Since $L = \text{cl}(c^*, F)$ is the largest “untouched” cluster, $|L| \geq |\text{cl}(c_1, F)|$ for all $c_1 \in F \setminus F'$. Then, by Observation 2, we have that

$$|P_{q^*}| = |D| + \sum_{c_1 \in F \setminus F'} |\text{cl}(c_1, F)| \leq |D| + |F \setminus F'| \cdot |L| \leq \frac{4}{\delta^2} \cdot |L| + k \cdot |L| \leq \frac{5k}{\delta^2} \cdot |L|.$$

In other words, $|L| \geq \frac{\delta^2}{5k} \cdot |P_{q^*}|$.

In the next claim, we summarize some properties of the sample S chosen in line 7 of the algorithm, in the current case, i.e., when $|L| \geq \frac{\delta^2}{5k} |C_{q^*}|$.

CLAIM 3. *Consider the iteration of the **for** loop of Algorithm 1, when $q = q^*$, and the corresponding sample S_q obtained in Algorithm 1. The following statements hold.*

- (1) *With probability at least 1/2, S_q contains at least β points of L .*
- (2) *$S_q \cap L$ has the same distribution as selecting $|S_q \cap L|$ points uniformly at random from L .*
- (3) *Let $L' \subseteq L$ be an arbitrary subset of size at least $\frac{\delta}{10} |L|$. Then, with probability at least 1/2, S_q contains at least 1 point from L , i.e., $S_q \cap L' \neq \emptyset$.*

PROOF. Recall that $\beta' = \beta \cdot \frac{150k}{\delta^3}$ and $|L| \geq \frac{\delta^2}{5k} |C_{q^*}|$. So, the first item follows, say, via Markov’s inequality.² The second item is an easy consequence of conditional distributions. The proof of the third item is analogous to the first item, combined with the bound on $|L'|$. \square

Now we condition on the event that $|S_{q^*} \cap L| \geq \beta$, which, by Claim 3 happens with probability at least $\frac{1}{2}$. Then, let $S' \subseteq S \cap L$ be such a subset of size β . Let $L = L_{\text{near}} \uplus L_{\text{far}}$, where $L_{\text{near}} = \{p \in L : \text{dist}(p, c^*) \leq \frac{8r}{\delta}\}$ and $L_{\text{far}} = \{p \in L : \text{dist}(p, c^*) > \frac{8r}{\delta}\}$. We consider different cases depending on the relative sizes of L_{near} and L_{far} . In the first case below (2.2.1), when $|L_{\text{near}}|$ is not very tiny compared to $|L_{\text{far}}|$, we show that our sample contains at least one point from L_{near} with good probability, and hence a δr grid around that point will contain an approximate center. In the complementary case (2.2.2), $|L_{\text{near}}|$ is very tiny compared to $|L_{\text{far}}|$, and in this case, we argue that, instead of finding an approximate HYBRID 1-MEDIAN, we can focus on finding an approximate 1-MEDIAN, which can be found using the sample. Now we formally analyze each of these cases.

Case 2.2.1. $|L_{\text{near}}| > \frac{\delta}{8} \cdot |L_{\text{far}}|$. In this case, letting $L' \leftarrow L_{\text{near}}$ in Claim 3, we infer that with at least probability 1/2, $S' \cap L_{\text{near}} \neq \emptyset$. We condition on this event. Then, since $\text{dist}(p, c^*) \leq \frac{8r}{\delta}$, it follows that there exists a $\tilde{c}^* \in \text{Grid}(p, \frac{8r}{\delta}, \delta r)$, such that $\text{dist}(\tilde{c}^*, c^*) \leq \delta r$. Since we branch on each point in $\bigcup_{p' \in S} \text{Grid}(p, \frac{8r}{\delta}, \delta r)$, we will branch on \tilde{c}^* in particular. Then, by Claim 1, $\text{HybridClustering}(F' \cup \{\tilde{c}^*\}, k, m - 1)$ returns a solution \tilde{F} , with probability at least $1/2 \cdot \alpha^{m-1} \geq \alpha^m$.

Case 2.2.2. $|L_{\text{near}}| \leq \frac{\delta}{8} \cdot |L_{\text{far}}|$. We prove two claims, namely Claim 4, and Claim 5. The latter essentially reduces the problem to finding an approximate solution to 1-median on L . Intuitively speaking, this follows from the following two reasons: (1) As we show in (10), a similar statement holds for the points in L_{far} . This essentially follows from the fact that, since each point of L_{far} has distance at least $\frac{8r}{\delta}$ to c , the subtraction of r from their distances has little effect on the cost, and

²In fact, a closer inspection reveals that the probability is much closer to 1, but “at least 1/2” suffices for our purpose.

(2) Due to the case assumption, the points of L_{far} vastly outnumber the points of L_{near} . Hence, the preceding claim also translates to the points of $L = L_{\text{far}} \cup L_{\text{near}}$, at a further small approximation error.

$$\text{CLAIM 4. } \sum_{p \in L} \text{dist}_r(p, c^*) \leq \sum_{p \in L} \text{dist}(p, c^*) \leq \left(1 + \frac{3\delta}{4}\right) \cdot \sum_{p \in L} \text{dist}_r(p, c^*).$$

PROOF. First, consider,

$$\sum_{p \in L_{\text{far}}} \text{dist}_r(p, c^*) = \sum_{p \in L_{\text{far}}} \text{dist}(p, c^*) - r \geq \sum_{p \in L_{\text{far}}} \text{dist}(p, c^*) - \frac{\delta}{8} \cdot \text{dist}(p, c^*).$$

(Since $\text{dist}(p, c^*) \geq \frac{8r}{\delta} > r$)

Then, the inequality between the first and the last term can be rewritten as

$$\sum_{p \in L_{\text{far}}} \text{dist}(p, c^*) \leq \frac{1}{1 - \delta/8} \cdot \sum_{p \in L_{\text{far}}} \text{dist}_r(p, c^*) \leq \left(1 + \frac{\delta}{4}\right) \cdot \sum_{p \in L_{\text{far}}} \text{dist}_r(p, c^*). \quad (10)$$

The following inequality will be used later to show that the contribution of points of L_{near} is negligible to the overall cost.

$$\begin{aligned} \sum_{p \in L_{\text{far}}} \text{dist}(p, c^*) &\geq \sum_{p \in L_{\text{far}}} \text{dist}_r(p, c^*) \geq \sum_{p \in L_{\text{far}}} \left(1 - \frac{\delta}{8}\right) \cdot \text{dist}(p, c^*) && \text{(From (10))} \\ &\geq \left(1 - \frac{\delta}{8}\right) \cdot \frac{8r}{\delta} \cdot |L_{\text{far}}| && \text{(Definition of } L_{\text{far}}\text{)} \\ &\geq \frac{1}{2} \cdot \frac{8r}{\delta} \cdot \frac{8}{\delta} \cdot |L_{\text{near}}| && \text{(Case assumption: } |L_{\text{far}}| \geq \frac{8}{\delta} \cdot |L_{\text{near}}|\text{)} \\ &= \frac{4}{\delta} \cdot \frac{8r}{\delta} \cdot |L_{\text{near}}| \\ &\geq \frac{4}{\delta} \cdot \sum_{p \in L_{\text{near}}} \text{dist}_r(p, c^*). \end{aligned} \quad (11)$$

where the last inequality follows from the definition of L_{near} .

The next sequence of inequalities shows a bound similar to (10), but when the sum is taken over all points of L (instead of only the points of L_{far} , as in Equation (10)).

$$\begin{aligned} \sum_{p \in L} \text{dist}_r(p, c^*) &\leq \sum_{p \in L} \text{dist}(p, c^*) = \sum_{p \in L_{\text{near}}} \text{dist}_r(p, c^*) + \sum_{p \in L_{\text{far}}} \text{dist}_r(p, c^*) \\ &= \left(1 + \frac{\delta}{4}\right) \cdot \sum_{p \in L_{\text{far}}} \text{dist}(p, c^*) && \text{(From (11))} \\ &\leq \left(1 + \frac{\delta}{4}\right) \cdot \left(1 + \frac{\delta}{4}\right) \cdot \sum_{p \in L_{\text{far}}} \text{dist}_r(p, c^*) && \text{(From (10))} \\ &\leq \left(1 + \frac{3\delta}{4}\right) \cdot \sum_{p \in L} \text{dist}_r(p, c^*). \end{aligned} \quad (12)$$

Here, the last inequality follows from (i) $(1 + \frac{\delta}{4}) \cdot (1 + \frac{\delta}{4}) \leq 1 + \frac{3\delta}{4}$ and (ii) $L_{\text{far}} \subseteq L$. This completes the proof of the claim. \square

Using this claim, we prove the following claim, which shows that it is sufficient to find an approximate 1-median solution for L , which will also be a good approximation for HYBRID 1-MEDIAN for L . To this end, let $\tilde{c}^* \in \mathbb{R}^d$ denote the optimal 1-median for L .

CLAIM 5. Let c_1 be a $(1 + \frac{\delta}{8})$ -approximation for 1-MEDIAN for L , i.e., $\sum_{p \in L} \text{dist}(p, c_1) \leq (1 + \frac{\delta}{8}) \cdot \sum_{p \in L} \text{dist}(p, \tilde{c}^*)$. Then, it is also a $(1 + \delta)$ -approximation for HYBRID 1-MEDIAN for L , i.e.,

$$\sum_{p \in L} \text{dist}_r(p, c_1) \leq (1 + \delta) \cdot \sum_{p \in L} \text{dist}_r(p, c^*) \quad (13)$$

PROOF. Let $c_1, \tilde{c}^* \in \mathbb{R}^d$ as defined above. Then,

$$\begin{aligned} \sum_{p \in L} \text{dist}_r(p, c_1) &\leq \sum_{p \in L} \text{dist}(p, c_1) \leq \left(1 + \frac{\delta}{8}\right) \cdot \sum_{p \in L} \text{dist}(p, \tilde{c}^*) && \text{(By definition of } c_1) \\ &\leq \left(1 + \frac{\delta}{8}\right) \cdot \sum_{p \in L} \text{dist}(p, c^*) \\ &\quad \text{(Since } \tilde{c}^* \in \mathbb{R}^d \text{ is an optimal 1-median and } c_1 \in \mathbb{R}^d \text{ is a feasible median)} \\ &\leq \left(1 + \frac{\delta}{8}\right) \cdot \left(1 + \frac{3\delta}{4}\right) \cdot \sum_{p \in L} \text{dist}_r(p, c^*) && \text{(From Claim 4)} \\ &\leq (1 + \delta) \cdot \sum_{p \in L} \text{dist}_r(p, c^*). \end{aligned}$$

□

Thus, now the task reduces to finding a $1 + \frac{\delta}{8}$ -approximate 1-MEDIAN solution for L . To this end, we have the following result from [28, 29].

PROPOSITION 2.4 ([28, 29]). Let $X \subset \mathbb{R}^d$ be a set of n points and $0 < \delta < 1$. Let $S \subseteq X$ be a uniform sample chosen from X of size $\beta = (\frac{1}{\delta})^{c'}$. Then, there exists an algorithm that runs in time $2^{O(1/\delta^c)}d$, and with probability at least α' , returns a $(1 + \delta)$ -approximate 1-median for X . Here, c, c' are absolute constants independent of the dimension d , and $0 < \alpha' < 1$ is a constant.

We combine the properties of the sample S_{q^*} proved in Claim 3 along with the previous proposition, to complete the proof. To this end, note that the first item of Claim 3 implies that, with probability at least $1/2$, S_{q^*} contains at least $\beta = \frac{1}{\delta^c}$ points from L . Then, we use the algorithm of Proposition 2.4, that returns with probability at least α' , a $(1 + \frac{\delta}{8})$ -approximate 1-median $c_1 \in \mathbb{R}^d$ for L . It follows that

$$\begin{aligned} &\sum_{c \in F'} \text{cost}_r(\text{cl}(c, F), c) + \sum_{c \in F_o \setminus \{c^*\}} \text{cost}_r(\text{cl}(c, F), c) + \sum_{p \in L} \text{dist}_r(p, c_1) \\ &\leq \sum_{c \in F'} \text{cost}_{r'}(\text{cl}(c, F), c) + \sum_{c \in F_o} \text{cost}_r(\text{cl}(c, F), c) - \text{cost}_r(\text{cl}(c^*, F), c^*) + \sum_{p \in L} \text{dist}_r(p, c_1) \\ &\leq \sum_{c \in F'} \text{cost}_{r'}(\text{cl}(c, F), c) + \sum_{c \in F_o} \text{cost}_r(\text{cl}(c, F), c) + \delta \cdot \sum_{c \in F_o} \text{cost}_r(\text{cl}(c, F), c) && \text{(From Claim 5)} \\ &\leq (1 + \delta) \cdot \left(\sum_{c \in F'} \text{cost}_{r'}(\text{cl}(c, F), c) + \sum_{c \in F_o} \text{cost}_r(\text{cl}(c, F), c) \right) \\ &\leq (1 + \delta)^{k-m+1} \cdot \text{OPT}_r. \end{aligned} \quad (14)$$

Then, by induction hypothesis, $\text{HybridClustering}(F' \cup \{c_1\}, k, m - 1)$, with probability at least α^{m-1} , returns a solution \tilde{F} such that $\text{cost}_{r'}(P, \tilde{F}) \leq (1 + \delta)^k \cdot \text{OPT}_r$. The overall probability of this event is at least $\frac{1}{2} \cdot \alpha' \cdot \alpha^{m-1} = \alpha^m$, completing the induction.

This finishes the case analysis, and thus we have established the invariant using induction. Using the invariant, we can show the following key lemma.

LEMMA 2.5. *HybridClustering(\emptyset, k, k) returns a $(1+\varepsilon, 1+\varepsilon)$ -bicriteria approximation solution to the given instance of HYBRID k -CLUSTERING with probability at least α^k for some constant $0 < \alpha < 1$.*

PROOF. We first show the following:

$$\text{CLAIM 6. } |R| \leq \left(\frac{k\sqrt{d}}{\delta}\right)^{O(d)} + \frac{k \log n}{\delta^{O(1)}} \cdot \left(\left(\frac{\sqrt{d}}{\delta}\right)^{O(d)} + \binom{\beta'}{\beta}\right) \leq (\log n) \cdot 2^{(kd/\delta)^{O(1)}}.$$

PROOF. First, in algorithm 1, we add the points returned by $\text{Grid}(c', 16r, \delta r)$ for each $c' \in F'$. The number of such points is $\left(\frac{16r\sqrt{d}}{\delta r}\right)^d = \left(\frac{\sqrt{d}}{\delta}\right)^{O(d)}$. Next, there are at most $\frac{\log n}{\delta^{O(1)}}$ values for q (this follows from the second preprocessing step, cf. Lemma 2.3), corresponding to each iteration of the for loop. In each iteration, we take a sample S of size $\beta' = O\left(\frac{150k\beta}{\delta^3}\right)$. Then, for each $p \in S$, we add to R the points of $\text{Grid}(p, 8r/\delta, \delta r)$, and the number of such points is at most $\left(\frac{8\sqrt{d}}{\delta^2}\right)^d = \left(\frac{\sqrt{d}}{\delta^2}\right)^{O(d)}$. In addition, we iterate over each subset $S' \subseteq S$ of size β , and the number of such subsets is $\binom{\beta'}{\beta} \leq \left(\frac{e\beta'}{\beta}\right)^\beta = \left(\frac{k}{\delta^3 \cdot \beta}\right)^\beta = \left(\frac{k}{\delta^3}\right)^{(1/\delta)^{O(1)}} = k^{1/\delta^{O(1)}}$. Thus, overall, the size of R is bounded by $(\log n) \cdot 2^{(kd/\delta)^{O(1)}}$. \square

To bound the running time of the algorithm, let $T(m)$ denote an upper bound on $\text{HybridClustering}(F', k, m)$ for any $F' \subset \mathbb{R}^d$. Note that we make a recursive call on each point in R . Further by Proposition 2.4, the time taken to compute a center in algorithm 1 is at most $2^{(1/\delta)^{O(1)}}$; and this algorithm is used in each of the at most $\frac{k \log n}{(1/\delta)^{O(1)}} \cdot \binom{\beta'}{\beta} \leq (\log n) \cdot k^{(1/\delta)^{O(1)}}$. Thus, $T(m)$ can be bounded by the following recurrence:

$$\begin{aligned} T(m) &\leq |R| \cdot T(m-1) + (\log n) \cdot k^{(1/\delta)^{O(1)}} \cdot n^{O(1)} \\ &\leq (\log n) \cdot 2^{\left(\frac{kd}{\delta}\right)^{O(1)}} \cdot T(m-1) + k^{(1/\delta)^{O(1)}} \cdot n^{O(1)}. \end{aligned}$$

It can be shown that this recurrence solves to $T(m) \leq 2^{\left(\frac{kd}{\delta}\right)^{O(1)}} \cdot n^{O(1)}$ – here we use the standard argument that $(\log n)^t \leq t^{O(t)} \cdot n^{O(1)}$.

Finally, note that our first call to the recursive algorithm is $\text{HybridClustering}(F' = \emptyset, k, k)$. At this point, the precondition of the invariant is satisfying setting $F_o \leftarrow F^*$, an optimal solution satisfying $\text{cost}_r(P, F^*) = \text{OPT}_r$. Then, the correctness of the invariant implies that, with probability at least α^k , the algorithm returns a solution \tilde{F} of size at most k , that is a $(1+\varepsilon, 1+\varepsilon)$ -bicriteria approximation – here we use that $\delta = \frac{\varepsilon}{10k}$, which implies that $(1+\delta)^k \leq (1+\varepsilon)$. \square

We now conclude with the following theorem, which is restated for convenience.

THEOREM 1.2. *Let $0 < \varepsilon < 1$. There exists a randomized algorithm that, given an instance of HYBRID k -CLUSTERING in \mathbb{R}^d , runs in time $2^{\left(\frac{kd}{\varepsilon}\right)^{O(1)}} \cdot n^{O(1)}$, and returns a $(1+\varepsilon, 1+\varepsilon)$ -approximation with probability at least a positive constant.*

PROOF. From Lemma 2.5, the success probability of the algorithm is α^k for some constant $\alpha > 0$. Thus, we need to repeat the algorithm α^{-k} times to boost the probability to at least a positive constant, which gets absorbed in the $2^{\left(\frac{kd}{\varepsilon}\right)^{O(1)}}$ factor. \square

A Hybrid of k -CENTER and k -MEANS. We note that an almost identical algorithm also implies a $(1+\varepsilon, 1+\varepsilon)$ bicriteria approximation for an analogous generalization of k -CENTER and k -MEANS. In this problem, the objective of (1) is replaced by the following: $\text{cost}_r(C, F) := \sum_{p \in C} \text{dist}_r(p, C)^2$. Let us refer to this problem as HYBRID $(k, 2)$ -CLUSTERING – the “2” in the name refers to the squares of the distance-thresholds that feature in the objective. Most of the analysis can be adapted to deal with the squares of the distances, by appropriately changing the sizes and distance-thresholds. The only significant change is that instead of Proposition 2.4, one needs to use an algorithm that

computes an approximate 1-MEANS solution given a large enough uniform sample of the cluster – such an algorithm can also be found in [29]. Then, one obtains the following theorem.

THEOREM 2.6. *Let $0 < \varepsilon < 1$. There exists a randomized algorithm that, given an instance of HYBRID $(k, 2)$ -CLUSTERING in \mathbb{R}^d , runs in time $2^{(\frac{kd}{\varepsilon})^{O(1)}} \cdot n^{O(1)}$ and returns a $(1+\varepsilon, 1+\varepsilon)$ -approximation with probability at least a positive constant.*

More generally, one can also define HYBRID (k, z) -CLUSTERING analogously, where the threshold-distances feature the z th power of distances. Again, our approach easily extends to this problem, modulo a version of Proposition 2.4 for the vanilla (k, z) -CLUSTERING in Euclidean spaces. To the best of our knowledge, such an algorithm is not explicitly known in the literature; however, it may be possible to obtain such an algorithm using the approach of [27, 29].

3 Conclusion and Future Directions

In this article, we proposed a novel clustering objective and defined a new problem, called HYBRID k -CLUSTERING, that generalizes both k -MEDIAN and k -CENTER. For d -dimensional Euclidean inputs, we designed a randomized $(1 + \varepsilon, 1 + \varepsilon)$ -bicriteria approximation scheme for HYBRID k -CLUSTERING running in time $2^{(kd/\varepsilon)^{O(1)}}$, for any $\varepsilon > 0$. Further, essentially the same algorithm also generalizes for a hybrid objective of k -CENTER and k -MEANS. We remind that improving either of the two $(1 + \varepsilon)$ factors to 1 would imply an exact FPT algorithm for k -CENTER/MEDIAN(/MEANS) in Euclidean spaces, which is unlikely to exist.

In the conference version of this article [18], we asked whether obtaining an approximation scheme with an improved FPT dependence on the dimension d – or ideally, *without* an FPT dependence on d – is possible, similar to the approach of [29] for k -MEDIAN/ k -MEANS. We also asked whether designing coresets for HYBRID k -CLUSTERING is possible. Subsequently, both of these problems were answered in the affirmative in the work of Gadekar and Inamdar [20]. Firstly, they designed dimension-free FPT approximation schemes for HYBRID k -CLUSTERING, and indeed generalizations of the problem to so-called *norm objectives* that also capture (k, z) -clustering variant, as defined above. In fact, their results hold in a much broader class of metric spaces introduced in [1], namely, metric spaces of bounded algorithmic scatter dimension. This class of metrics includes Euclidean spaces of arbitrary dimension, metric spaces of bounded doubling dimension, shortest path metrics in graphs excluding a fixed minor, and bounded treewidth graphs [1, 6]. Secondly, they designed coresets for HYBRID k -CLUSTERING in metric spaces of bounded doubling dimension.

In light of our results, and the subsequent improvement by [20], HYBRID k -CLUSTERING remains greatly unexplored in general metric spaces. Here, a primal-dual algorithm from [8] can be adapted to obtain an (α, β) -bicriteria approximation in polynomial time, for some constants α and β . Indeed, a quick examination of the proof of [8] suggests that $(18, 6)$ -bicriteria approximation easily follows by mimicking their approach step-by-step; with further improvements possible with a more careful analysis. Exploring the best possible constants in the bicriteria approximation in general metric spaces would be an interesting avenue for future research.

References

- [1] Fateme Abbasi, Sandip Banerjee, Jaroslaw Byrka, Parinya Chalermsook, Ameet Gadekar, Kamyar Khodamoradi, Dániel Marx, Roohani Sharma, and Joachim Spoerhase. 2023. Parameterized approximation schemes for clustering with general norm objectives. In *Proceedings of the 64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023*. IEEE, 1377–1399. DOI: <https://doi.org/10.1109/FOCS57990.2023.00085>
- [2] Pankaj K. Agarwal, Sariel Har-Peled, and Hai Yu. 2008. Robust shape fitting via peeling and grating coresets. *Discret. Comput. Geom.* 39, 1–3 (2008), 38–58. DOI: <https://doi.org/10.1007/S00454-007-9013-2>
- [3] Pankaj K. Agarwal and Cecilia Magdalena Procopiuc. 2002. Exact and approximation algorithms for clustering. *Algorithmica* 33, 2 (2002), 201–226. DOI: <https://doi.org/10.1007/S00453-001-0110-Y>

- [4] Mihai Badoiu, Sariel Har-Peled, and Piotr Indyk. 2002. Approximate clustering via core-sets. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing, (STOC)*. ACM, 250–257. DOI : <https://doi.org/10.1145/509907.509947>
- [5] MohammadHossein Bateni, Vincent Cohen-Addad, Alessandro Epasto, and Silvio Lattanzi. 2024. A scalable algorithm for individually fair k -means clustering. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*. Sanjoy Dasgupta, Stephan Mandt, and Yingzhen Li (Eds.), *Proceedings of Machine Learning Research*, Vol. 238, PMLR, 3151–3159. Retrieved from <https://proceedings.mlr.press/v238/bateni24a.html>
- [6] Romain Bourneuf and Marcin Pilipczuk. 2025. Bounding ϵ -scatter dimension via metric sparsity. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025*. Yossi Azar and Debmalya Panigrahi (Eds.), SIAM, 3155–3171. DOI : <https://doi.org/10.1137/1.9781611978322.101>
- [7] Jaroslaw Byrka, Krzysztof Sornat, and Joachim Spoerhase. 2018. Constant-factor approximation for ordered k -median. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. ACM, 620–631. DOI : <https://doi.org/10.1145/3188745.3188930>
- [8] Deeparnab Chakrabarty and Chaitanya Swamy. 2018. Interpolating between k -median and k -center: Approximation algorithms for ordered k -median. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP) (LIPIcs)*, Vol. 107. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 29:1–29:14. DOI : <https://doi.org/10.4230/LIPIcs.ICALP.2018.29>
- [9] Moses Charikar, Samir Khuller, David M. Mount, and Giri Narasimhan. 2001. Algorithms for facility location problems with outliers. In *Proceedings of the 12th Annual Symposium on Discrete Algorithms (SODA)*. ACM/SIAM, 642–651. Retrieved from <http://dl.acm.org/citation.cfm?id=365411.365555>
- [10] Moses Charikar and Erik Waingarten. 2022. The Johnson-Lindenstrauss lemma for clustering and subspace approximation: From coresets to dimension reduction. arXiv:2205.00371. Retrieved from <https://arxiv.org/abs/2205.00371>
- [11] Vincent Cohen-Addad, Arnaud de Mesmay, Eva Rotenberg, and Alan Roytman. 2018. The bane of low-dimensionality clustering. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*. Artur Czumaj (Ed.), SIAM, 441–456. DOI : <https://doi.org/10.1137/1.9781611975031.30>
- [12] Vincent Cohen-Addad, Anupam Gupta, Amit Kumar, Euiwoong Lee, and Jason Li. 2019. Tight FPT approximations for k -median and k -means. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP) (LIPIcs)*, Vol. 132. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 42:1–42:14. DOI : <https://doi.org/10.4230/LIPIcs.ICALP.2019.42>
- [13] Vincent Cohen-Addad, Philip N. Klein, and Claire Mathieu. 2019. Local search yields approximation schemes for k -means and k -median in euclidean and minor-free metrics. *SIAM J. Comput.* 48, 2 (2019), 644–667. DOI : <https://doi.org/10.1137/17M112717X>
- [14] Vincent Cohen-Addad, Kasper Green Larsen, David Saulpic, and Chris Schwiegelshohn. 2022. Towards optimal lower bounds for k -median and k -means coresets. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. ACM, 1038–1051. DOI : <https://doi.org/10.1145/3519935.3519946>
- [15] Mark de Berg, Leyla Biabani, and Morteza Monemzadeh. 2023. k -center clustering with outliers in the MPC and streaming model. In *Proceedings of the IEEE International Parallel and Distributed Processing Symposium, (IPDPS)*. IEEE, 853–863. DOI : <https://doi.org/10.1109/IPDPS54959.2023.00090>
- [16] Hu Ding, Haiku Yu, and Zixiu Wang. 2019. Greedy strategy works for k -center clustering with outliers and coreset construction. In *Proceedings of the 27th Annual European Symposium on Algorithms (ESA) (LIPIcs)*, Vol. 144. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 40:1–40:16. DOI : <https://doi.org/10.4230/LIPIcs.ESA.2019.40>
- [17] Alon Efrat, Micha Sharir, and Alon Ziv. 1994. Computing the smallest k -enclosing circle and related problems. *Comput. Geom.* 4, 3 (1994), 119–136. <https://www.sciencedirect.com/science/article/pii/0925772194900035>
- [18] Fedor V. Fomin, Petr A. Golovach, Tanmay Inamdar, Saket Saurabh, and Meirav Zehavi. 2024. Hybrid k -clustering: Blending k -median and k -center. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2024*. Amit Kumar and Noga Ron-Zewi (Eds.), LIPIcs, Vol. 317, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 4:1–4:19. DOI : <https://doi.org/10.4230/LIPIcs.APPROX/RANDOM.2024.4>
- [19] Zachary Friggstad, Mohsen Rezapour, and Mohammad R. Salavatipour. 2019. Local search yields a PTAS for k -means in doubling metrics. *SIAM J. Comput.* 48, 2 (2019), 452–480. DOI : <https://doi.org/10.1137/17M1127181>
- [20] Ameet Gadekar and Tanmay Inamdar. 2025. Dimension-free parameterized approximation schemes for hybrid clustering. In *Proceedings of the 42nd International Symposium on Theoretical Aspects of Computer Science, STACS 2025*. Olaf Beyersdorff, Michal Pilipczuk, Elaine Pimentel, and Kim Thang Nguyen (Eds.), LIPIcs, Vol. 327, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 35:1–35:20. DOI : <https://doi.org/10.4230/LIPIcs.STACS.2025.35>
- [21] Rajiv Gandhi, Samir Khuller, and Aravind Srinivasan. 2004. Approximation algorithms for partial covering problems. *J. Algorithms* 53, 1 (2004), 55–84.
- [22] Moses Charikar and Erik Waingarten. 2025. The johnson-lindenstrauss lemma for clustering and subspace approximation: From coresets to dimension reduction. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12-15, 2025*. SIAM, 3172–3209. DOI : <https://doi.org/10.1137/1.9781611978322.102>

- [23] Sariel Har-Peled. 2007. How to get close to the median shape. *Comput. Geom.* 36, 1 (2007), 39–51. DOI : <https://doi.org/10.1016/J.COMGEO.2006.02.003>
- [24] Sariel Har-Peled and Soham Mazumdar. 2005. Fast algorithms for computing the smallest k-enclosing circle. *Algorithmica* 41, 3 (2005), 147–157. DOI : <https://doi.org/10.1007/S00453-004-1123-0>
- [25] Sariel Har-Peled and Yusu Wang. 2004. Shape fitting with outliers. *SIAM J. Comput.* 33, 2 (2004), 269–285. DOI : <https://doi.org/10.1137/S0097539703427963>
- [26] Martin G. Herold, Evangelos Kipouridis, and Joachim Spoerhase. 2025. Clustering to minimize cluster-aware norm objectives. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025*. Yossi Azar and Debmalya Panigrahi (Eds.), SIAM, 255–287. DOI : <https://doi.org/10.1137/1.9781611978322.8>
- [27] Ragesh Jaiswal, Amit Kumar, and Sandeep Sen. 2014. A simple D 2-sampling based PTAS for k-means and other clustering problems. *Algorithmica* 70, 1 (2014), 22–46. DOI : <https://doi.org/10.1007/S00453-013-9833-9>
- [28] Amit Kumar, Yogish Sabharwal, and Sandeep Sen. 2005. Linear time algorithms for clustering problems in any dimensions. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming (ICALP) (Lecture Notes in Computer Science)*, Vol. 3580. Springer, 1374–1385. DOI : https://doi.org/10.1007/11523468_111
- [29] Amit Kumar, Yogish Sabharwal, and Sandeep Sen. 2010. Linear-time approximation schemes for clustering problems in any dimensions. *J. ACM* 57, 2 (2010), 5:1–5:32.
- [30] Konstantin Makarychev, Yury Makarychev, and Ilya P. Razenshteyn. 2019. Performance of Johnson-Lindenstrauss transform for k-means and k-medians clustering. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. ACM, 1027–1038. DOI : <https://doi.org/10.1145/3313276.3316350>
- [31] Dániel Marx. 2005. Efficient approximation schemes for geometric problems?. In *Proceedings of the 13th Annual European Symposium on Algorithms (ESA) (Lecture Notes in Computer Science)*, Vol. 3669. Springer, 448–459. DOI : https://doi.org/10.1007/11561071_41
- [32] Jirí Matousek. 1995. On enclosing k points by a circle. *Inf. Process. Lett.* 53, 4 (1995), 217–221. DOI : [https://doi.org/10.1016/0020-0190\(94\)00190-A](https://doi.org/10.1016/0020-0190(94)00190-A)
- [33] Nimrod Megiddo and Kenneth J. Supowit. 1984. On the complexity of some common geometric location problems. *SIAM J. Comput.* 13, 1 (1984), 182–196.
- [34] Arie Tamir. 2001. The k-centrum multi-facility location problem. *Discret. Appl. Math.* 109, 3 (2001), 293–307. DOI : [https://doi.org/10.1016/S0166-218X\(00\)00253-5](https://doi.org/10.1016/S0166-218X(00)00253-5)
- [35] Hai Yu, Pankaj K. Agarwal, Raghunath Paredy, and Kasturi R. Varadarajan. 2008. Practical methods for shape fitting and kinetic data structures using coresets. *Algorithmica* 52, 3 (2008), 378–402. DOI : <https://doi.org/10.1007/S00453-007-9067-9>

Received 27 September 2024; revised 16 April 2025; accepted 12 May 2025