Clustering What Matters: Optimal Approximation for Clustering with Outliers *

Akanksha Agrawal[1], Tanmay Inamdar[2], Saket Saurabh[2, 3], Jie Xue[4]

[1] Indian Institute of Technology Madras, Chennai, India.
[2] University of Bergen, Bergen, Norway
[3] The Institute of Mathematical Sciences, HBNI, Chennai, India
[4] New York University Shanghai, China
akanksha@cse.iitm.ac.in, Tanmay.Inamdar@uib.no, saket@imsc.res.in, jiexue@nyu.edu

Abstract

Clustering with outliers is one of the most fundamental problems in Computer Science. Given a set X of n points and two integers k and m, the clustering with outliers aims to exclude m points from X, and partition the remaining points into k clusters that minimizes a certain cost function. In this paper, we give a general approach for solving clustering with outliers, which results in a fixed-parameter tractable (FPT) algorithm in k and m (i.e., an algorithm with running time of the form $f(k, m) \cdot n^{O(1)}$ for some function f), that almost matches the approximation ratio for its outlier-free counterpart.

As a corollary, we obtain FPT approximation algorithms with optimal approximation ratios for k-MEDIAN and k-MEANS with outliers in general metrics. We also exhibit more applications of our approach to other variants of the problem that impose additional constraints on the clustering, such as fairness or matroid constraints.

Introduction

Clustering is a family of problems that aims to group a given set of objects in a meaningful way—the exact "meaning" may vary based on the application. These are fundamental problems in Computer Science with applications ranging across multiple fields like pattern recognition, machine learning, computational biology, bioinformatics and social science. Thus, these problems have been a subject of extensive studies in the field of Algorithm Design (and its subfields), see for instance, the surveys on this topic (and references therein) (Xu and Tian 2015; Rokach 2009; Blömer et al. 2016).

Two of the central clustering problems are k-MEDIAN and k-MEANS. In the standard k-MEDIAN problem, we are given a set X of n points, and an integer k, and the goal is to find a set $C^* \subseteq X$ of at most k *centers*, such that the following cost function is minimized over all subsets C of size at most k.

$$\mathsf{cost}(X, C) := \sum_{p \in X} \min_{c \in C} \mathsf{d}(p, c)$$

In k-MEANS, the objective function instead contains the sum of squares of distances.

Often real world data are contaminated with a small amount of noise and these noises can substantially change the clusters that we obtain using the underlying algorithm. To circumvent the issue created by such noises, there are several studies of clustering problems with outliers, see for instance, (Chen 2008; Krishnaswamy, Li, and Sandeep 2018; Goyal, Jaiswal, and Kumar 2020; Feng et al. 2019; Friggstad et al. 2019; Almanza et al. 2022).

In outlier extension of the k-MEDIAN problem, which we call k-MEDIANOUT, we are also given an additional integer $m \geq 0$ that denotes the number of *outliers* that we are allowed to drop. We want to find a set C of at most k centers, and a set $Y \subseteq X$ of at most m outliers, such that $\mathsf{cost}(X \setminus Y, C) := \sum_{p \in X \setminus Y} \min_{c \in C} \mathsf{d}(p, c)$ is minimized over all (Y, C) satisfying the requirements. Observe that the cost of clustering for k-MEDIANOUT equals the sum of distances of each point to its nearest center, after excluding a set of m points from consideration [1]. We remark that in a similar spirit we can define the outlier version of the k-MEANS problem, which we call k-MEANSOUT.

In this paper, we will focus on approximation algorithms. An algorithm is said to have an approximation ratio of α, if it is guaranteed to return a solution of cost no greater than α times the optimal cost, while satisfying all other conditions. That is, the solution must contain at most k centers, and drop m outliers. If the algorithm is randomized, then it must return such a solution with *high probability*, i.e., probability at least $1 - n^{-c}$ for some $c \geq 1$.

For a fixed set C of centers, the set of m outliers is automatically defined, namely the set of m points that are farthest from C (breaking ties arbitrarily). Thus, an optimal

*The authors are listed in the alphabetical order of the last names.
A full version of the paper containing missing proofs and other details can be found is available on arXiv:2212.00696 (Agrawal et al. 2022).

[1]In the technical section, we consider a more general formulation of k-MEDIAN, where the set of candidate centers may be different from the set X of points to be clustered.

clustering for k-MEDIANOUT, just like k-MEDIAN, can be found in $n^{O(k)}$ time by enumerating all center sets. On the other hand, we can enumerate all $n^{O(m)}$ subsets of outliers, and reduce the problem directly to k-MEDIAN. Other than these straightforward observations, there are several non-trivial approximations known for k-MEDIANOUT, which we discuss in a subsequent paragraph.

Our Results. In this work, we describe a general framework that reduces a *clustering with outliers* problem (such as k-MEDIANOUT or k-MEANSOUT) to its outlier-free counterpart in an approximation-preserving fashion. More specifically, given an instance $\mathcal{I}$ of k-MEDIANOUT, our reduction runs in time $f(k, m, \epsilon) \cdot n^{O(1)}$, and produces multiple instances of k-MEDIAN, such that a β-approximation for *at least* one of the produced instances of k-MEDIAN implies a $(\beta + \epsilon)$-approximation for the original instance $\mathcal{I}$ of k-MEDIANOUT. This is the main result of our paper.

Our framework does not rely on the specific properties of the underlying metric space. Thus, for special metrics, such as Euclidean spaces, or shortest-path metrics induced by sparse graph classes, for which FPT $(1 + \epsilon)$-approximations are known for k-MEDIAN, our framework implies matching approximation for k-MEDIANOUT. Finally, our framework is quite versatile in that one can extend it to obtain approximation-preserving FPT reductions for related *clustering with outliers* problems, such as k-MEANSOUT, and clustering problems with *fair outliers* (such as (Bandyapadhyay et al. 2019; Jia, Sheth, and Svensson 2020)), and MATROID MEDIAN WITH OUTLIERS. We conclude by giving a partial list of the corollaries of our reduction framework. The running time of each algorithm is $f(k, m, \epsilon) \cdot n^{O(1)}$ for some function f that depends on the problem and the setting. Next to each result, we also cite the result that we use as a black box to solve the outlier-free clustering problem.

- $(1 + 2/e + \epsilon) \approx (1.74 + \epsilon)$-approximation (resp. $1 + 8/e + \epsilon$)-approximation) for k-MEDIANOUT (resp. k-MEANSOUT) in general metrics (Cohen-Addad, Saulpic, and Schwiegelshohn 2021). These approximations are tight even for $m = 0$, under a reasonable complexity theoretic hypothesis, as shown in the same paper.
- $(1 + \epsilon)$-approximation for k-MEDIANOUT and k-MEANSOUT in (i) metric spaces of constant doubling dimensions, which includes Euclidean spaces of constant dimension, (ii) metrics induced by graphs of bounded treewidth, and (iii) metrics induced by graphs that exclude a fixed graph as a minor (such as planar graphs). (Cohen-Addad, Saulpic, and Schwiegelshohn 2021).
- $(2 + \epsilon)$-approximation for MATROID MEDIAN WITH OUTLIERS in general metrics, where k refers to the *rank* of the matroid. (Cohen-Addad et al. 2019)
- $(1 + 2/e + \epsilon)$-approximation for COLORFUL k-MEDIAN in general metrics, where m denotes the total number of outliers across all color classes (Cohen-Addad et al. 2019). The preceding two problems are orthogonal generalizations of k-MEDIANOUT, and are formally defined in Section .

Our Techniques. Our reduction is inspired from the following seemingly simple observation that relates k-MEDIANOUT and k-MEDIAN. Let $\mathcal{I}$ be an instance of k-MEDIANOUT, where we want to find a set C of k centers, such that the sum of distances of all except at most m points to the nearest center in C is minimized. By treating the outliers in an optimal solution for $\mathcal{I}$ as *virtual centers*, one obtains a solution for $(k+m)$-MEDIAN *without outliers* whose cost is at most the optimal cost of $\mathcal{I}$. In other words, the optimal cost of an appropriately defined instance $\widetilde{\mathcal{I}}$ of $(k+m)$-MEDIAN is a *lower bound* on the optimal cost of $\mathcal{I}$. Since k-MEDIAN is a well-studied problem, at this point, one would hope that it is sufficient to restrict the attention to $\widetilde{\mathcal{I}}$. That is, if we obtain a solution (i.e., a set of $k + m$ centers) for $\widetilde{\mathcal{I}}$, can then be modified to obtain a solution (i.e., a set of k centers and m outliers) for $\mathcal{I}$. However, it is unclear whether one can do such a modification without blowing up the cost for $\mathcal{I}$. Nevertheless, this connection between $\widetilde{\mathcal{I}}$ and $\mathcal{I}$ turns out to be useful, but we need several new ideas to exploit it.

As in before, we start with a constant approximation for $\widetilde{\mathcal{I}}$, and perform a sampling similar to (Chen 2009) to obtain a weighted set of points. This set is obtained by dividing the set of points connected to each center in the approximate solution into *concentric rings*, such that the "error" introduced in the cost by treating all points in the ring as identical is negligible. Then, we sample $O((k+m) \log n/\epsilon)$ points from each ring, and give each point an appropriate weight. We then prove a crucial concentration bound (cf. Lemma 1), which informally speaking relates the connection cost of original set of points in a ring, and the corresponding weighted sample. In particular, for *any set of k centers*, with good probability, the difference between the original and the weighted costs is "small", *even after excluding at most m outliers from both sets*. Intuitively speaking, this concentration bound holds because the sample size is large enough compared to both k and m. Then, by taking the union of all such samples, we obtain a weighted set S of $O(((k + m) \log n/\epsilon)^2)$ points that preserves the connection cost to any set of k centers, even after excluding m outliers with at least a constant probability. Then, we enumerate all sets Y of size m from S, and solve the resulting k-MEDIAN instance induced on $S \setminus Y$. Finally, we argue that at least one of the resulting instances $\mathcal{I}'$ will have the property that, a β-approximation for $\mathcal{I}'$ implies a $(\beta + \epsilon)$-approximation for $\mathcal{I}$.

Related Work. The first constant approximation for k-MEDIANOUT was given by (Chen 2008) for some large constant. More recently, (Krishnaswamy, Li, and Sandeep 2018; Gupta, Moseley, and Zhou 2021) gave constant approximations based on iterative LP rounding technique, and the 6.994-approximation by the latter is currently the best known approximation. These approximation algorithms run in polynomial time in n. (Krishnaswamy, Li, and Sandeep 2018) also give the best known polynomial approximations for related problems of k-MEANSOUT and MATROID MEDIAN.

Now we turn to FPT approximations, which is also the setting for our results. To the best of our knowledge, there are three works in this setting, (Feng et al. 2019; Goyal, Jaiswal, and Kumar 2020; Statman, Rozenberg, and Feldman 2020). The idea of relating k-MEDIAN WITH m OUTLIERS to $(k+m)$-MEDIAN that we discuss above is also present in these works. Even though it is not stated explicitly, the approach of Statman et al. (Statman, Rozenberg, and Feldman 2020) can be used to obtain FPT approximations in general metrics; albeit with a worse approximation ratio. However, by using additional properties of Euclidean k-MEDIANOUT/k-MEANSOUT (where one is allowed to place centers anywhere in $\mathbb{R}^d$) their approach yields a $(1+\epsilon)$-approximation in FPT time. The best FPT approximations in general metrics, to the best of our knowledge, are $3+\epsilon$ for k-MEDIANOUT by Goyal et al. (Goyal, Jaiswal, and Kumar 2020), and $6+\epsilon$ for k-MEANSOUT by Feng et al. (Feng et al. 2019). Thus, our FPT approximation algorithms with ratio $1+\frac{2}{e}+\epsilon$ for k-MEDIANOUT, and $1+\frac{8}{e}+\epsilon$ for k-MEANSOUT improve on these results. Furthermore, our result is essentially an *approximation-preserving reduction* from k-MEDIANOUT to k-MEDIAN in the same kind of metric, which automatically yields improved approximations in some special metrics as discussed earlier.

On the lower bound side, (Guha and Khuller 1999) showed it is NP-hard to approximate k-MEDIAN (and thus k-MEDIANOUT) within a factor $1+2/e-\epsilon$ for any $\epsilon > 0$. Recently, (Cohen-Addad et al. 2019) strengthened this result under a reasonable complexity-theory hypothesis, and showed that an $(1+2/e-\epsilon)$-approximation algorithm must take at least $n^{k^{g(\epsilon)}}$ time for some function $g()$.

Bicriteria approximations relax the strict requirement of using at most k centers, or dropping at most m outliers, in order to give improved approximation ratios, or efficiency (or both). For k-MEDIANOUT, (Charikar et al. 2001) gave a $4(1+1/\epsilon)$-approximation, while dropping $m(1+\epsilon)$ outliers. (Gupta et al. 2017) gave a constant approximation based on local search for k-MEANSOUT that drops $O(km\log(n\Delta))$ outliers, where Δ is the diameter of the set of points. (Friggstad, Rezapour, and Salavatipour 2019) gave a $(25+\epsilon)$-approximation that uses $k(1+\epsilon)$ centers but only drops m outliers. In Euclidean spaces, they also give a $(1+\epsilon)$-approximation that returns a solution with $k(1+\epsilon)$ centers.

Preliminaries

Basic notions. Let (Γ, d) be a metric space, where Γ is a finite set of points, and $\mathsf{d}: \Gamma\times\Gamma \to \mathbb{R}$ is a distance function satisfying symmetry and triangle inequality. For any finite set $S \subseteq \Gamma$ and a point $p \in \Gamma$, we let $\mathsf{d}(p,S) := \min_{s\in S}\mathsf{d}(p,S)$, and let $\mathsf{diam}(S) := \max_{x,y\in S}\mathsf{d}(x,y)$. For two non-empty sets $S, C \subseteq \Gamma$, let $\mathsf{d}(S,C) = \min_{p\in S}\mathsf{d}(p,S) = \min_{p\in S}\min_{c\in C}\mathsf{d}(p,c)$. For a point $p \in \Gamma, r \geq 0$, and a set $C \subseteq \Gamma$, let $B_C(p,r) = \{q \in C : \mathsf{d}(p,c) \leq r\}$. Let T be a finite (multi)set of n real numbers, for some positive integer n, and let $1 \leq m \leq n$. Then, we use the notation $\mathsf{sum}_{\sim m}(T)$ to denote the sum of $n-m$ smallest values in T (including repetitions in case of a multi-set).

The k-median problem. In the k-MEDIAN problem, an instance is a triple $\mathcal{I} = (X, F, k)$, where X and F are finite sets of points in some metric space (Γ, d), and $k \geq 1$ is an integer. The points in X are called *clients*, and the points in F are called *facilities* or *centers*. The task is to find a subset $C \subseteq F$ of size at most k that minimizes the cost function

$$\mathsf{cost}(X,C) := \sum_{p\in X}\mathsf{d}(p,C).$$

The *size* of an instance $\mathcal{I} = (X,F,k)$ is defined as $|\mathcal{I}| = |X\cup F|$, which we denote by n.

k-median with outliers. The input to k-MEDIANOUT contains an additional integer $0 \leq m \leq n$, and thus an instance is given by a 4-tuple $\mathcal{I} = (X,F,k,m)$. Let $C \subseteq F$ be a set of facilities. We define $\mathsf{cost}_m(X,C) := \mathsf{sum}_{\sim m}\{\mathsf{cost}(p,C) : p\in X\}$, i.e., the sum of $n-m$ smallest distances of points in X to the set of centers C. The goal is to find a set of centers C minimizing $\mathsf{cost}_m(X,C)$ over all sets $C \subseteq F$ of size at most k. Given a set $C \subseteq F$ of centers, we denote the corresponding solution by (Y,C), where $Y \subseteq X$ is a set of m outlier points in X with largest distances realizing $\mathsf{cost}_m(X,C)$. Given an instance $\mathcal{I}$ of k-MEDIANOUT, we use $\mathsf{OPT}(\mathcal{I})$ to denote the value of an optimal solution to $\mathcal{I}$.

Weighted sets and random samples. During the course of the algorithm, we will often deal with *weighted sets* of points. Here, $S \subseteq X$ is a weighted set, with each point $p \in S$ having integer weight $w(p) \geq 0$. For any set $C \subseteq F$ and $1 \leq m \leq |S|$, define $\mathsf{wcost}_m(S,C) := \mathsf{sum}_{\sim m}\{d(p,C)\cdot w(p) : p\in S\}$. A *random sample* of a finite set S refers to a random subset of S. Throughout this paper, random samples are always generated by picking points *uniformly* and *independently*.

k-Median with Outliers

In this section, we give our FPT reduction from k-MEDIANOUT to the standard k-MEDIAN problem. Formally, we shall prove the following theorem.

Theorem 1. *Suppose there exists a β-approximation algorithm for k-MEDIAN with running time $T(n,k)$, and a τ-approximation algorithm for $k+m$-MEDIAN with polynomial running time, where β and τ are constants. Then there exists a $(\beta+\epsilon)$-approximation algorithm for k-MEDIANOUT with running time $\left(\frac{k+m}{\epsilon}\right)^{O(m)}\cdot T(n,k)\cdot n^{O(1)}$, where n is the instance size and m is the number of outliers.*

Combining the above theorem with the known $(1+2/e+\epsilon)$-approximation k-median algorithm (Cohen-Addad et al. 2019) that runs in $(k/\epsilon)^{O(k)}\cdot n^{O(1)}$ time, we directly have the following result.

Corollary 1. There exists a $(1+2/e+\epsilon)$-approximation algorithm for k-MEDIANOUT with running time $\left(\frac{k+m}{\epsilon}\right)^{O(m)}\cdot\left(\frac{k}{\epsilon}\right)^{O(k)}\cdot n^{O(1)}$, where n is the instance size and m is the number of outliers.

The rest of this section is dedicated to proving Theorem 1. Let $\mathcal{I} = (X, F, k, m)$ be an instance of k-MEDIANOUT. We define a $(k+m)$-MEDIAN instance $\mathcal{I}' = (X, F \cup X, k+m)$, where in addition to the original set of facilities, there is a facility co-located with each client. We have the following observation.

Observation 1. $\mathsf{OPT}(\mathcal{I}') \leq \mathsf{OPT}(\mathcal{I})$, i.e., the value of an optimal solution to $\mathcal{I}'$ is a lower bound on the value of an optimal solution to $\mathcal{I}$.

Proof. Let (Y^*, C^*) be an optimal solution to $\mathcal{I}$ realizing the value $\mathsf{OPT}(\mathcal{I})$. We define a solution (Y', C') for $\mathcal{I}'$ as follows: let $Y' = X$, and $C' = C^* \cup Y^*$. That is, the set of centers C' is obtained by adding a facility co-located with each outlier point from Y^*. Now we argue about the costs. Since $C^* \subseteq C'$, for each point $p \in Y^*$, $\mathsf{d}(p, C') \leq \mathsf{d}(p, C^*)$. On the other hand, for each $q \in X \setminus Y^*$, $\mathsf{d}(q, C') = 0$, since there is a co-located center in C^*. This implies that $\mathsf{cost}_0(X, C') \leq \mathsf{cost}_m(X, C)$. Since the solution (Y', C') is feasible for the instance $\mathcal{I}'$, it follows that $\mathsf{OPT}(\mathcal{I}')$ is no larger than the cost $\mathsf{cost}_0(X, C')$. ☐

Now, we use τ-approximation algorithm guaranteed by the theorem, for the instance $\mathcal{I}'$, and obtain a set of at most $k' \leq k + m$ centers A such that $\mathsf{cost}_0(X, A) \leq \tau \cdot \mathsf{OPT}(\mathcal{I}') \leq \tau \cdot \mathsf{OPT}(\mathcal{I})$. By assumption, running this algorithm takes polynomial time. Let $R = \frac{\mathsf{cost}_0(X,A)}{\tau n}$ be a lower bound on average radius, and $\phi = \lceil \log(\tau n) \rceil$. For each center $c_i \in A$, let $X_i \subseteq X$ denote the set of points whose closest center in A is c_i. By arbitrarily breaking ties, we can assume that the sets X_i are disjoint, i.e., $\{X_i\}_{1 \leq i \leq k'}$ forms a partition of X. Now we further partition each X_i into smaller groups such that the points in each group have similar distances to c_i. Specifically, we define

$$X_{i,j} := \begin{cases} B_{X_i}(c_i, R) & \text{if } j = 0, \\ B_{X_i}(c_i, 2^j R) \setminus B_{X_i}(c_i, 2^{j-1} R) & \text{if } j \geq 1. \end{cases}$$

Let $s = \frac{c\tau^2}{\epsilon^2}(m + k \ln n + \ln(1/\lambda))$, for some large enough constant c. We define a weighted set of points $S_{i,j} \subseteq X_{i,j}$ as follows. If $|X_{i,j}| \leq s$, then we say $X_{i,j}$ is *small*. In this case, define $S_{i,j} := X_{i,j}$ and let the weight $w_{i,j}$ of each point $p \in S_{i,j}$ be 1. Otherwise, $|X_{i,j}| > s$ and we say that $X_{i,j}$ is *large*. In this case, we take a random sample $S_{i,j} \subseteq X_{i,j}$ of size s. We set the weight of every point in $S_{i,j}$ to be $w_{i,j} = |X_{i,j}|/|S_{i,j}|$. For convenience, assume the weights $w_{i,j}$ to be integers [2]. Finally, let $S = \bigcup_{i,j} S_{i,j}$. The set S can be thought of as an ϵ-coreset for the k-MEDIANOUT instance $\mathcal{I}$. Even though we do not define this notion formally, the key properties of S will be proven in Lemma 2 and 3. Thus, we will often informally refer to S as a *coreset*.

Proposition 1. *We have $|S| = O(((k + m) \log n/\epsilon)^2)$ if λ is a constant.*

[2]For a large $X_{i,j}$, the quantity $\frac{|X_{i,j}|}{|S_{i,j}|}$ may not be an integer. We ignore this technicality for now, and discuss in the full version how to modify the construction slightly to ensure that the weights are integral.

Proof. For any $p \in X$, $\mathsf{d}(p, A) \leq \mathsf{cost}_0(X, A) = \tau n \cdot R \leq 2^\phi R$. Therefore, for any $c_i \in A$ and $j > \phi$, $X_{i,j'} = \emptyset$, and $X_i = \bigcup_{j=0}^{\phi} X_{i,j}$. It follows that the number of non-empty sets $X_{i,j}$ is at most $|A| \cdot (1 + \log(\tau n)) = O((k+m) \log n)$, since $|A| \leq k + m$ and τ is a constant. For each non-empty $X_{i,j}$, $|S_{i,j}| \leq 2s = O((m + k \log n)/\epsilon^2)$, if λ is a constant. Since $S = \bigcup_{i,j} S_{i,j}$, the claimed bound follows. ☐

Proposition 2. *(Chen 2009; Haussler 1992) Let $M \geq 0$ and η be fixed constants, and let $h(\cdot)$ be a function defined on a set V such that $\eta \leq h(p) \leq \eta + M$ for all $p \in V$. Let $U \subseteq V$ be a random sample of size s, and $\delta > 0$ be a parameter. If $s \geq \frac{M^2}{2\delta^2} \ln(2/\lambda)$, then*

$$\Pr\left[\left|\frac{h(V)}{|V|} - \frac{h(U)}{|U|}\right| \geq \delta\right] \leq \lambda,$$

where $h(U) := \sum_{u \in U} h(u)$, and $h(V) := \sum_{v \in V} h(v)$.

Lemma 1. Let (Γ, d) be a metric space, $V \subseteq \Gamma$ be a finite set of points, $\lambda', \xi > 0$, $q \geq 0$ be parameters, and define $s' = \frac{4}{\xi^2}\left(q + \ln \frac{2}{\lambda'}\right)$. Suppose $U \subseteq V$ is a random sample of size s'. Then for any fixed finite set $C \subseteq F$ with probability at least $1 - \lambda'$ it holds that for any $0 \leq t \leq q$,

$$|\mathsf{cost}_t(V, C) - \mathsf{wcost}_{t'}(U, C)| \leq \xi |V| (\mathsf{diam}(V) + \mathsf{d}(V, C)),$$

where $t' = \lfloor t|U|/|V| \rfloor$ and $w(u) = |V|/|U|$ for all $u \in U$.

Proof. Throughout the proof, we fix the set $C \subseteq F$ and $0 \leq t \leq q$ as in the statement of the lemma. Next, we define the following notation. For each $v \in V$, let $h(v) := d(v, C)$, and let $h(V) := \sum_{v \in V} h(v)$, and $h(U) := \sum_{u \in U} h(u)$. Analogously, let $h'(V) := \mathsf{cost}_t(V, C)$, and $h'(U) := \mathsf{cost}_{t'}(U, C)$. Let $\eta(V) := \min_{v \in V} \mathsf{d}(v, C)$, and $\eta(U) := \min_{u \in U} \mathsf{d}(u, C)$.

By applying Proposition 2 with $\eta = \eta(V)$, $M = \mathsf{diam}(V)$ and $\delta = \xi M/2$, we know with probability at most λ', the following event happens.

$$\left|\frac{\sum_{v \in V} \mathsf{d}(v, C)}{|V|} - \frac{\sum_{u \in U} \mathsf{d}(u, C)}{|U|}\right| \geq \frac{\xi}{2} \mathsf{diam}(V).$$

As $h(V) = \sum_{v \in V} \mathsf{d}(v, C)$ and $h(U) = \sum_{u \in U} \mathsf{d}(u, C)$, with probability at least $1 - \lambda'$, we have that

$$\left|\frac{h(V)}{|V|} - \frac{h(U)}{|U|}\right| \leq \frac{\xi}{2} \mathsf{diam}(V). \tag{1}$$

Claim 1. *Let $\Delta = \frac{h'(V)}{|V|} - \frac{h'(U)}{|U|}$. If Equation 1 holds, then we have*

$$-\xi(\mathsf{diam}(V) + \mathsf{d}(V, C)) \leq \Delta \leq \xi \mathsf{diam}(V).$$

The proof of the above claim involves a lot of calculations, so we defer it to the full version. Assuming its correctness, the inequality in Claim 1 implies

$$\left|h'(V) - h'(U) \cdot \frac{|V|}{|U|}\right| \leq \xi |V| \cdot (\mathsf{diam}(V) + \mathsf{d}(V, C)).$$

Since Equation 1 holds with probability at least $1 - \lambda'$, the above inequality also holds with probability at least $1 - \lambda'$.

The preceding inequality is equivalent to the one in the lemma, because $h'(V) = \mathsf{cost}_t(V, C)$, and $h'(U) \cdot \frac{|V|}{|U|} = \frac{|V|}{|U|} \cdot \mathsf{cost}_{t'}(U, C) = \mathsf{wcost}_{t'}(U, C)$. Finally, notice that Claim 1 holds when the h' function is defined with respect to any choice of $t \in \{0, 1, \dots, q\}$. Therefore, with probability at least $1 - \lambda'$, the inequality in the lemma holds for every $t \in \{0, 1, \dots, q\}$, which completes the proof. □

Next, we show the following observation, whose proof is identical to an analogous proof in (Chen 2008).

Observation 2. The following inequalities hold.

- $\sum_{i,j} |X_{i,j}| 2^j R \le 3 \cdot \mathsf{cost}_0(X, A) \le 3\tau \cdot \mathsf{OPT}(\mathcal{I})$.
- $\sum_{i,j} |X_{i,j}| \mathsf{diam}(X_{i,j}) \le 6 \cdot \mathsf{cost}_0(X, A) \le 6\tau \cdot \mathsf{OPT}(\mathcal{I})$.

Proof. For any $p \in X_{i,j}$, it holds that $2^j R \le \max\{2\mathsf{d}(p, A), R\} \le 2\mathsf{d}(p, A) + R$. Therefore,

$$\begin{aligned}
\sum_{i,j} |X_{i,j}| \cdot 2^j R &\le \sum_{i,j} \sum_{p \in X_{i,j}} 2^j R \\
&\le \sum_{i,j} \sum_{p \in X_{i,j}} 2\mathsf{d}(p, A) + R \\
&= 2 \sum_{p \in X} \mathsf{d}(p, A) + |X| \cdot R \\
&\le 2 \cdot \mathsf{cost}_0(X, A) + nR \\
&\le 3 \cdot \mathsf{cost}_0(X, A) \quad \text{(By definition of } R\text{)} \\
&\le 3\tau \mathsf{OPT}(\mathcal{I}') \le 3\tau \mathsf{OPT}(\mathcal{I}).
\end{aligned}$$

We also obtain the second item by observing that $\mathsf{diam}(X_{i,j}) \le 2 \cdot 2^j \cdot R$. □

Next, we show that the following lemma, which informally states that the union of the sets of sampled points approximately preserve the cost of clustering w.r.t. *any* set of at most k centers, *even after* excluding at most m outliers overall.

Lemma 2. The following event happens with probability at least $1 - \lambda/2$:
For all sets $C \subseteq F$ of size at most k, and for all sets of non-negative integers $\{m_{i,j}\}_{i,j}$ such that $\sum_{i,j} m_{i,j} \le m$,

$$\left| \sum_{i,j} \mathsf{cost}_{m_{i,j}}(X_{i,j}, C) - \sum_{i,j} \mathsf{wcost}_{t_{i,j}}(S_{i,j}, C) \right| \le \epsilon \cdot \sum_{i,j} \mathsf{cost}_{m_{i,j}}(X_{i,j}, C) \qquad (2)$$

where $t_{i,j} = \lfloor m_{i,j} / w_{i,j} \rfloor$.

Proof. Fix an arbitrary set $C \subseteq F$ of at most k centers, and the integers $\{m_{i,j}\}_{i,j}$ such that $\sum_{i,j} m_{i,j} \le m$ as in the statement of the lemma. For each $i = 1, \dots, |A|$, and $0 \le j \le \phi$, we invoke Lemma 1 by setting $V = X_{i,j}$, and $U = S_{i,j}$, $\xi = \frac{\epsilon}{8\tau}$, $\lambda' = n^{-k} \lambda / (4(k+m)(1+\phi))$, and $q = m$. This implies that, the following inequality holds with probability at least $1 - \lambda'$ for each set $X_{i,j}$, and the corresponding $m_{i,j} \le m$,

$$\begin{aligned}
&\left| \mathsf{cost}_{m_{i,j}}(X_{i,j}, C) - \mathsf{wcost}_{t_{i,j}}(S_{i,j}, C) \right| \\
&\le \frac{\epsilon}{8\tau} |X_{i,j}| (\mathsf{diam}(X_{i,j}) + \mathsf{d}(X_{i,j}, C)) \qquad (3)
\end{aligned}$$

Note that the sample size required in order for this inequality to hold is

$$\begin{aligned}
s' &= \left\lceil \frac{4}{\xi^2} \left(m + \ln\left(\frac{2}{\lambda'} \right) \right) \right\rceil \\
&= \left\lceil 4 \left(\frac{8\tau}{\epsilon} \right)^2 \cdot \left(m + \ln\left(\frac{8n^k (k+m)(1+\phi)}{\lambda} \right) \right) \right\rceil \le s.
\end{aligned}$$

For any i, j, if $X_{i,j} < s$ (i.e., $X_{i,j}$ is *small*), then the sample $S_{i,j}$ is equal to $X_{i,j}$, and each point in $S_{i,j}$ has weight equal to 1. This implies that $\mathsf{cost}_{m_{i,j}}(X_{i,j}, C) = \mathsf{wcost}_{t_{i,j}}(S_{i,j}, C)$ for all such $X_{i,j}$, and their contribution to the right hand side of inequality (2) is zero. Thus, it suffices to restrict the sum on the right hand side of (2) over *large* sets $X_{i,j}$'s. Let $\mathcal{L}$ consist of all pairs (i, j) such that $X_{i,j}$ is large. We have the following claim.

Claim 2. $\sum_{(i,j) \in \mathcal{L}} |X_{i,j}| \mathsf{d}(X_{i,j}, C) \le 2\mathsf{cost}_m(X, C)$.

Proof. Let Y denote the farthest m points in X from the set of centers C. Now, fix $(i, j) \in \mathcal{L}$ and let $q_{i,j} := |X_{i,j} \cap Y| \le m$ denote the number of outliers in $X_{i,j}$. Since $|X_{i,j}| \ge 2m \ge 2q_{i,j}$, the set $X_{i,j} \setminus Y$ is non-empty, and all points $X_{i,j} \setminus Y$ contribute towards $\mathsf{cost}_m(X, C)$. That is,

$$\sum_{(i,j) \in \mathcal{L}} \sum_{p \in X_{i,j} \setminus Y} \mathsf{d}(p, C) \le \mathsf{cost}_m(X, C) \qquad (4)$$

For any $p \in X_{i,j} \setminus Y$, $\mathsf{d}(X_{i,j}, C) \le \mathsf{d}(p, C)$ from the definition. Therefore,

$$\begin{aligned}
&\sum_{(i,j) \in \mathcal{L}} |X_{i,j}| \cdot \mathsf{d}(X_{i,j}, C) \\
&\le \sum_{(i,j) \in \mathcal{L}} 2|X_{i,j} \setminus Y| \cdot \mathsf{d}(X_{i,j}, C) \\
&\le 2 \cdot \sum_{(i,j) \in \mathcal{L}} \sum_{p \in X_{i,j} \setminus Y} \mathsf{d}(p, C) \\
&\le 2 \cdot \mathsf{cost}_m(X, C)
\end{aligned}$$

Here, to see the second inequality, see that $|X_{i,j}| \ge 2q_{i,j}$, which implies that $|X_{i,j}| - q_{i,j} \le 2(|X_{i,j}| - q_{i,j})$. The last inequality follows from (4). □

Thus, by revisiting (3) and (2), we get:

$$\begin{aligned}
&\sum_{(i,j) \in \mathcal{L}} \left| \mathsf{cost}_{m_{i,j}}(X_{i,j}, C) - \mathsf{wcost}_{t_{i,j}}(S_{i,j}, C) \right| \\
&\le \frac{\epsilon}{8\tau} \sum_{(i,j) \in \mathcal{L}} |X_{i,j}| (\mathsf{diam}(X_{i,j}) + \mathsf{d}(X_{i,j}, C)) \\
&\qquad \text{(From (3))} \\
&\le \frac{\epsilon}{8\tau} \cdot (6\tau \cdot \mathsf{OPT}(\mathcal{I}) + 2\mathsf{cost}_m(X, C)) \\
&\qquad \text{(From Obs. 2 and Claim 2)} \\
&= \frac{\epsilon}{8\tau} (8\tau \cdot \mathsf{cost}_m(X, C)) = \epsilon \cdot \mathsf{cost}_m(X, C)
\end{aligned}$$

Where, in the last inequality, since C is an arbitrary set of at most k centers, $\mathsf{OPT}(\mathcal{I}) \le \mathsf{cost}_m(X, C)$. Note that the preceding inequality holds for a fixed set C of centers with probability at least $1 - |A| \cdot (1+\phi)\lambda' = 1 - n^{-k}\lambda/2$, which follows from taking the union bound over all sets $X_{i,j}$, $1 \le i \le |A| \le k+m$, and $0 \le j \le \phi$.

Since F has at most n^k subsets of size at most k, the statement of the lemma follows from taking a union bound. $\square$

Now we are ready to prove Theorem 1. We enumerate every subset $T \subseteq S$ of size at most m. For each T, we compute a β-approximation solution for the (weighted) k-median instance $(S\backslash T, F, k)$. Theorem 1 only assumes the existence of a β-approximation algorithm for unweighted k-median, which cannot be applied to weighted point sets. However, we can transform $S\backslash T$ to an equivalent unweighted sets R, which contains, for each $x \in S\backslash T$, $w(x)$ copies of (unweighted) x, where $w(x)$ is the weight of x in $S\backslash T$. It is clear that $\mathsf{wcost}(S\backslash T, C) = \mathsf{cost}(R, C)$ for all $C \subseteq F$. Thus, we can apply the β-approximation k-MEDIAN algorithm on (R, F, k) to compute a center set $C \subseteq F$ of size k such that $\mathsf{wcost}(T, C) \le \beta \cdot \mathsf{wcost}(T, C')$ for any $C' \subseteq F$ of size k. We do this for all $T \subseteq S$ of size at most m. Let $\mathcal{C}$ denote the set of all center sets C computed. We pick a center set $C^* \subseteq \mathcal{C}$ that minimizes $\mathsf{cost}_m(X, C^*)$, and return (Y^*, C^*) as the solution where $Y^* \subseteq X$ consists of the m points in X farthest to the center set C^*.

Lemma 3. With probability at least $1 - \frac{\lambda}{2}$, for all $C \subseteq F$ of size k we have

$$\mathsf{cost}_m(X, C^*) \le \frac{1+\epsilon}{1-\epsilon} \cdot \beta \mathsf{cost}_m(X, C).$$

Proof. The statement in Lemma 2 holds with probability at least $1 - \lambda/2$. Thus, it suffices to assume the statement in Lemma 2, and show $\mathsf{cost}_m(X, C^*) \le (1+\epsilon)^2\beta \cdot \mathsf{cost}_m(X, C)$ for any $C \subseteq F$ of size k. Fix a subset $C \subseteq F$ of size k. Let $Y \subseteq X$ consist of the m points in X farthest to C, and define $m_{i,j} = |Y \cap X_{i,j}|$. Set $t_{i,j} = \lfloor m_{i,j}/w_{i,j} \rfloor$. Note that $\mathsf{cost}_m(X, C) = \sum_{i,j} \mathsf{cost}_{m_{i,j}}(X_{i,j}, C)$. Furthermore, by Lemma 2, we have

$$\begin{aligned}\sum_{i,j} \mathsf{wcost}_{t_{i,j}}(S_{i,j}, C) &\le (1+\epsilon) \cdot \sum_{i,j} \mathsf{cost}_{m_{i,j}}(X_{i,j}, C) \\ &= (1+\epsilon) \cdot \mathsf{cost}_m(X, C). \end{aligned} \tag{5}$$

Now let $T_{i,j} \subseteq S_{i,j}$ consist of the $t_{i,j}$ points in $S_{i,j}$ farthest to C, and define $T = \bigcup_{i,j} T_{i,j}$. Since $|T| \le m$, T is considered by our algorithm and thus there exists a center set $C' \in \mathcal{C}$ that is a β-approximation solution for the (weighted) k-median instance $(S\backslash T, F, k)$. We have

$$\begin{aligned}\mathsf{wcost}(S\backslash T, C') &\le \beta \cdot \mathsf{wcost}(S\backslash T, C) \\ &= \beta \sum_{i,j} \mathsf{wcost}_{t_{i,j}}(S_{i,j}, C). \end{aligned} \tag{6}$$

Note that $\mathsf{wcost}(S\backslash T, C') \ge \sum_{i,j} \mathsf{wcost}_{t_{i,j}}(S_{i,j}, C')$. Furthermore, by applying Lemma 2 again, we have $\sum_{i,j} \mathsf{wcost}_{t_{i,j}}(S_{i,j}, C') \ge (1-\epsilon) \cdot \sum_{i,j} \mathsf{cost}_{m_{i,j}}(X_{i,j}, C')$. It then follows that

$$\begin{aligned}(1-\epsilon) \cdot \mathsf{cost}_m(X_{i,j}, C') &\le (1-\epsilon) \cdot \sum_{i,j} \mathsf{cost}_{m_{i,j}}(X_{i,j}, C') \\ &\le \mathsf{wcost}(S\backslash T, C'). \end{aligned} \tag{7}$$

Finally, we have $\mathsf{cost}_m(X, C^*) \le \mathsf{cost}_m(X, C')$ by the construction of C^*. Combining this with (5), (6), and (7), we have $\mathsf{cost}_m(X, C^*) \le \frac{1+\epsilon}{1-\epsilon} \cdot \beta \mathsf{cost}_m(X, C)$, which completes the proof. $\square$

By choosing $\lambda > 0$ to be a sufficiently small constant, and by appropriately rescaling ϵ [3], the above lemma shows that our algorithm outputs a $(\beta+\epsilon)$-approximation solution with a constant probability. By repeating the algorithm a logarithmic number of rounds, we can guarantee the algorithm succeeds with high probability. The number of subsets $T \subseteq S$ of size at most m is bounded by $|S|^{O(m)}$, which is $((k+m)\log n/\epsilon)^{O(m)}$ by Proposition 1. Note that $(\log n)^{O(m)} \le \max\{m^{O(m)}, n^{O(1)}\}$. Thus, the number of subsets $T \subseteq S$ of size at most m is bounded by $f(k, m, \epsilon) \cdot n^{O(1)}$, where $f(k, m, \epsilon) = \left(\frac{k+m}{\epsilon}\right)^{O(m)}$. Thus, we need to call the β-approximation k-MEDIAN algorithm $f(k, m, \epsilon) \cdot n^{O(1)}$ times, which takes $f(k, m, \epsilon) n^{O(1)} \cdot T(n, k)$ time overall. The first call of the algorithm for obtaining a τ-approximation to the $(k+m)$-MEDIAN instance takes polynomial time. Besides this, the other parts of our algorithm can all be implemented in polynomial time. This completes the proof of Theorem 1.

Extensions

k-Means with Outliers

This is similar to k-MEDIANOUT, except that the cost function is the sum of *squares* of distances of all except m outlier points to a set of k facilities. This generalizes the well-known k-MEANS problem. Here, the main obstacle is that the squares of distances do not satisfy triangle inequality, and thus it does not form a metric. However, they satisfy a *relaxed* version of triangle inequality (i.e., $\mathsf{d}(p,q)^2 \le 2(\mathsf{d}(p,r)^2 + \mathsf{d}(r,q)^2)$). This technicality makes the arguments tedious, nevertheless, we can follow the same approach as for k-MEDIANOUT, to obtain optimal FPT approximation schemes. Our technique implies an optimal $(1 + 8/e + \epsilon)$-approximation for k-MEANSOUT (using the result of (Cohen-Addad et al. 2019) as a black-box), improving upon polynomial-time 53.002-approximation from (Krishnaswamy, Li, and Sandeep 2018), and $(9+\epsilon)$-approximation from (Goyal, Jaiswal, and Kumar 2020) in time FPT in k, m and ϵ.

In fact, using our technique, we can get improved approximation guarantees for (k,z)-CLUSTERING WITH OUTLIERS, where the cost function involves z-th power of distances, where $z \ge 1$ is fixed for a problem. Note that the cases $z = 1$ and $z = 2$ correspond to k-MEDIANOUT and k-MEANSOUT respectively. We give the details for (k,z)-CLUSTERING WITH OUTLIERS in the full version.

[3] Since Lemma 3 implies a $\beta(1+O(\epsilon))$-approximation, and β is a constant, it suffices to redefine $\epsilon = \epsilon/c$ for some large enough constant c to get the desired result.

Matroid Median with Outliers

A *matroid* is a pair $\mathcal{M} = (F, \mathcal{S})$, where F is a ground set, and $\mathcal{S}$ is a collection of subsets of F with the following properties: (i) $\emptyset \in \mathcal{S}$, (ii) If $A \in \mathcal{S}$, then for every subset $B \subseteq A$, $B \in \mathcal{S}$, and (iii) For any $A, B \in \mathcal{S}$ with $|B| < |A|$, there exists an $b \in B \setminus A$ such that $B \cup \{b\} \in \mathcal{S}$. The rank of a matroid $\mathcal{M}$ is the size of the largest independent set in $\mathcal{S}$.

An instance of MATROID MEDIAN WITH OUTLIERS is given by $(X, F, \mathcal{M}, m)$, where $\mathcal{M} = (F, \mathcal{S})$ is a matroid with rank k defined over a finite ground set F, and X, F are sets of clients and facilities, belonging to a finite metric space (Γ, d). The objective is to find a set $C \subseteq F$ of facilities that minimizes $\mathsf{cost}_m(X, C)$, and $C \in \mathcal{S}$, i.e., C is an independent set in the given matroid. Note that an explicit description of a matroid of rank k may be as large as n^k. Therefore, we assume that we are given an efficient *oracle access* to the matroid $\mathcal{M}$. That is, we are provided with an algorithm $\mathcal{A}$ that, given a candidate set $S \subseteq F$, returns in time $T(\mathcal{A})$ (which is assumed to be polynomial in $|F|$), returns whether $S \in \mathcal{I}$.

We can adapt our approach to MATROID MEDIAN WITH OUTLIERS in a relatively straightforward manner. Recall that our algorithm needs to start with an instance of outlier-free problem (i.e., MATROID MEDIAN) that provides a lower bound on the optimal cost of the given instance. To this end, given an instance $\mathcal{I} = (X, F, \mathcal{M} = (F, \mathcal{S}), m)$ of MATROID MEDIAN WITH OUTLIERS, we define an instance $\mathcal{I}' = (X, F, \mathcal{M}')$ MATROID MEDIAN, where $\mathcal{M}' = (F \cup X, \mathcal{S}')$ is defined as follows. $\mathcal{S}' = \{Y \cup C : Y \subseteq X \text{ with } |Y| \leq m \text{ and } C \subseteq F \text{ with } C \in \mathcal{S}\}$. That is, each independent set of $\mathcal{M}'$ is obtained by taking the union of an independent set of facilities from $\mathcal{M}$, and a subset of X of size at most m. It is straightforward to show that $\mathcal{M}'$ is a matroid over the ground set $F \cup X$. In particular, it is the direct sum of $\mathcal{M}$ and a uniform matroid over X of rank m (i.e., where any subset of X of size at most m is independent). Using the oracle algorithm $\mathcal{A}$, we can simulate an oracle algorithm to test whether a candidate set $C \subseteq F \cup X$ is independent in $\mathcal{M}'$. Therefore, using a $(2+\epsilon)$-approximation for MATROID MEDIAN (Cohen-Addad et al. 2019) in time FPT in k and ϵ, we can find a set $A \subseteq F \cup X$ of size at most $k + m$ that we can use to construct a coreset. The details about enumeration are similar to that for k-MEDIANOUT, and are thus omitted.

Colorful k-Median

This is an orthogonal generalization of k-MEDIANOUT to ensure a certain notion of *fairness* in the solution (see (Jia, Sheth, and Svensson 2020)). Suppose the set of points X is partitioned into ℓ different colors $X_1 \uplus X_2 \uplus \ldots \uplus X_\ell$. We are also given the corresponding number of outliers $m_1, m_2, \ldots, m_\ell$. The goal is to find a set of at most facilities C to minimize the connection cost of all except at most m_t outliers from each color class X_t, i.e., we want to minimize the cost function: $\sum_{t=1}^{\ell} \mathsf{cost}_{m_t}(X_t, C)$. This follows a generalizations of the well-known k-CENTER problem introduced in (Bandyapadhyay et al. 2019) and (Anegg et al. 2020; Jia, Sheth, and Svensson 2020) , called COLORFUL k-CENTER. Similar generalization of FACILITY LOCATION has also been studied in (Chekuri et al. 2022).

Using our ideas, we can find an FPT approximation parameterized by k, $m = \sum_{t=1}^{\ell} m_t$, and ϵ. To this end, we sample sufficiently many points from each color class X_t separately, and argue that it preserves the cost appropriately. The technical details follow the same outline as that for k-Median with m Outliers. In particular, during the enumeration phase—just like that for k-MEDIANOUT—we obtain several instances of k-MEDIAN. That is, our algorithm is *color-agnostic* after constructing the coreset. Thus, we obtain a tight $(1 + 2/e + \epsilon)$-approximation for this problem. This is the first non-trivial true approximation for this problem – previous work (Gupta, Moseley, and Zhou 2021) only gives a *pseudo-approximation*, i.e., a solution with cost at most a constant times that of an optimal cost, but using slightly more than k facilities.

A Combination of Above Generalizations

Our technique also works for a combination of the aforementioned generalizations that are orthogonal to each other. To consider an extreme example, consider COLORFUL MATROID MEDIAN with ℓ different color classes (a similar version for k-CENTER objective has been recently studied by (Anegg, Koch, and Zenklusen 2022)), where we want to find a set of facilities that is independent in the given matroid, in order to minimize the sum of distances of all except m_t outlier points for each color class X_t. By using a combination of the ideas mentioned above, one can get FPT approximations for such generalizations.

Concluding Remarks

In this paper, we give a reduction from k-MEDIANOUT to k-MEDIAN that runs in time FPT in k, m, and ϵ, and preserves the approximation ratio up to an additive ϵ factor. As a consequence, we obtain improved FPT approximations for k-MEDIANOUT in general as well as special kinds of metrics, and these approximation guarantees are known to be tight. Furthermore, our technique is versatile in that it also gives improved approximations for related clustering problems, such as k-MEANSOUT, MATROID MEDIAN WITH OUTLIERS, and COLORFUL k-MEDIAN, among others.

The most natural direction is to improve the FPT running time while obtaining the tight approximation ratios. More fundamentally, perhaps, is the question whether we need an FPT dependence on the number of outliers, m; or whether it is possible to obtain approximation guarantees for k-MEDIANOUT matching that for k-MEDIAN, with a running time that is FPT in k and ϵ.

Acknowledgments

T. Inamdar is supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 819416). S. Saurabh is supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 819416) and Swarnajayanti Fellowship (No. DST/SJF/MSA01/2017-18).

References

Agrawal, A.; Inamdar, T.; Saurabh, S.; and Xue, J. 2022. Clustering What Matters: Optimal Approximation for Clustering with Outliers. arXiv:2212.00696.

Almanza, M.; Epasto, A.; Panconesi, A.; and Re, G. 2022. k-Clustering with Fair Outliers. In Candan, K. S.; Liu, H.; Akoglu, L.; Dong, X. L.; and Tang, J., eds., *WSDM '22: The Fifteenth ACM International Conference on Web Search and Data Mining, Virtual Event / Tempe, AZ, USA, February 21 - 25, 2022*, 5–15. ACM.

Anegg, G.; Angelidakis, H.; Kurpisz, A.; and Zenklusen, R. 2020. A Technique for Obtaining True Approximations for k-Center with Covering Constraints. In Bienstock, D.; and Zambelli, G., eds., *Integer Programming and Combinatorial Optimization - 21st International Conference, IPCO 2020, London, UK, June 8-10, 2020, Proceedings*, volume 12125 of *Lecture Notes in Computer Science*, 52–65. Springer.

Anegg, G.; Koch, L. V.; and Zenklusen, R. 2022. Techniques for Generalized Colorful k-Center Problems. In Chechik, S.; Navarro, G.; Rotenberg, E.; and Herman, G., eds., *30th Annual European Symposium on Algorithms, ESA 2022, September 5-9, 2022, Berlin/Potsdam, Germany*, volume 244 of *LIPIcs*, 7:1–7:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Bandyapadhyay, S.; Inamdar, T.; Pai, S.; and Varadarajan, K. R. 2019. A Constant Approximation for Colorful k-Center. In Bender, M. A.; Svensson, O.; and Herman, G., eds., *27th Annual European Symposium on Algorithms, ESA 2019, September 9-11, 2019, Munich/Garching, Germany*, volume 144 of *LIPIcs*, 12:1–12:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Blömer, J.; Lammersen, C.; Schmidt, M.; and Sohler, C. 2016. Theoretical analysis of the k-means algorithm–a survey. In *Algorithm Engineering*, 81–116. Springer.

Charikar, M.; Khuller, S.; Mount, D. M.; and Narasimhan, G. 2001. Algorithms for facility location problems with outliers. In Kosaraju, S. R., ed., *Proceedings of the Twelfth Annual Symposium on Discrete Algorithms, January 7-9, 2001, Washington, DC, USA*, 642–651. ACM/SIAM.

Chekuri, C.; Inamdar, T.; Quanrud, K.; Varadarajan, K.; and Zhang, Z. 2022. Algorithms for covering multiple submodular constraints and applications. *Journal of Combinatorial Optimization*, 1–32.

Chen, K. 2008. A constant factor approximation algorithm for *k*-median clustering with outliers. In Teng, S., ed., *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008*, 826–835. SIAM.

Chen, K. 2009. On coresets for k-median and k-means clustering in metric and euclidean spaces and their applications. *SIAM Journal on Computing*, 39(3): 923–947.

Cohen-Addad, V.; Gupta, A.; Kumar, A.; Lee, E.; and Li, J. 2019. Tight FPT Approximations for k-Median and k-Means. In Baier, C.; Chatzigiannakis, I.; Flocchini, P.; and Leonardi, S., eds., *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, volume 132 of *LIPIcs*, 42:1–42:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Cohen-Addad, V.; Saulpic, D.; and Schwiegelshohn, C. 2021. A new coreset framework for clustering. In Khuller, S.; and Williams, V. V., eds., *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, 169–182. ACM.

Feng, Q.; Zhang, Z.; Huang, Z.; Xu, J.; and Wang, J. 2019. Improved Algorithms for Clustering with Outliers. In Lu, P.; and Zhang, G., eds., *30th International Symposium on Algorithms and Computation, ISAAC 2019, December 8-11, 2019, Shanghai University of Finance and Economics, Shanghai, China*, volume 149 of *LIPIcs*, 61:1–61:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Friggstad, Z.; Khodamoradi, K.; Rezapour, M.; and Salavatipour, M. R. 2019. Approximation schemes for clustering with outliers. *ACM Transactions on Algorithms (TALG)*, 15(2): 1–26.

Friggstad, Z.; Rezapour, M.; and Salavatipour, M. R. 2019. Local Search Yields a PTAS for k-Means in Doubling Metrics. *SIAM J. Comput.*, 48(2): 452–480.

Goyal, D.; Jaiswal, R.; and Kumar, A. 2020. FPT Approximation for Constrained Metric k-Median/Means. In Cao, Y.; and Pilipczuk, M., eds., *15th International Symposium on Parameterized and Exact Computation, IPEC 2020, December 14-18, 2020, Hong Kong, China (Virtual Conference)*, volume 180 of *LIPIcs*, 14:1–14:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Guha, S.; and Khuller, S. 1999. Greedy Strikes Back: Improved Facility Location Algorithms. *J. Algorithms*, 31(1): 228–248.

Gupta, A.; Moseley, B.; and Zhou, R. 2021. Structural Iterative Rounding for Generalized k-Median Problems. In Bansal, N.; Merelli, E.; and Worrell, J., eds., *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPIcs*, 77:1–77:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Gupta, S.; Kumar, R.; Lu, K.; Moseley, B.; and Vassilvitskii, S. 2017. Local Search Methods for k-Means with Outliers. *Proc. VLDB Endow.*, 10(7): 757–768.

Haussler, D. 1992. Decision Theoretic Generalizations of the PAC Model for Neural Net and Other Learning Applications. *Inf. Comput.*, 100(1): 78–150.

Jia, X.; Sheth, K.; and Svensson, O. 2020. Fair Colorful k-Center Clustering. In Bienstock, D.; and Zambelli, G., eds., *Integer Programming and Combinatorial Optimization - 21st International Conference, IPCO 2020, London, UK, June 8-10, 2020, Proceedings*, volume 12125 of *Lecture Notes in Computer Science*, 209–222. Springer.

Krishnaswamy, R.; Li, S.; and Sandeep, S. 2018. Constant approximation for k-median and k-means with outliers via iterative rounding. In Diakonikolas, I.; Kempe, D.; and Henzinger, M., eds., *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, 646–659. ACM.

Rokach, L. 2009. A survey of clustering algorithms. In *Data mining and knowledge discovery handbook*, 269–298. Springer.

Statman, A.; Rozenberg, L.; and Feldman, D. 2020. k-Means: Outliers-Resistant Clustering+++. *Algorithms*, 13(12): 311.

Xu, D.; and Tian, Y. 2015. A comprehensive survey of clustering algorithms. *Annals of Data Science*, 2(2): 165–193.