
Approximate Group Fairness for Clustering

Bo Li¹ Lijun Li² Ankang Sun³ Chenhao Wang^{4*} Yingfan Wang⁵

Abstract

We incorporate group fairness into the algorithmic centroid clustering problem, where k centers are to be located to serve n agents distributed in a metric space. We refine the notion of proportional fairness proposed in [Chen et al., ICML 2019] as *core fairness*, and k -clustering is in the core if no coalition containing at least n/k agents can strictly decrease their total distance by deviating to a new center together. Our solution concept is motivated by the situation where agents are able to coordinate and utilities are transferable. A string of existence, hardness and approximability results is provided. Particularly, we propose two dimensions to relax core requirements: one is on the degree of distance improvement, and the other is on the size of deviating coalition. For both relaxations and their combination, we study the extent to which relaxed core fairness can be satisfied in metric spaces including line, tree and general metric space, and design approximation algorithms accordingly.

1. Introduction

Motivated by various real-world machine learning algorithm deployments where the data points are real human beings who should be treated unbiasedly, fairness is increasingly concerned. Most traditional algorithms mainly focus on the efficiency or profit, and thus fail to ensure fairness for individual point or collection of points. Accordingly, the past several years have seen considerable efforts in developing fair learning algorithms (Chierichetti et al., 2017; Chen et al., 2019; Backurs et al., 2019; Bera et al., 2019).

Following (Chen et al., 2019), we revisit the group fairness

¹Department of Computing, The Hong Kong Polytechnic University, Hong Kong, China ²School of Mathematical Sciences, Ocean University of China, Qingdao, China ³Warwick Business School, University of Warwick, United Kingdom ⁴University of Nebraska-Lincoln, United States ⁵Department of Computer Science, Duke University, United States. Correspondence to: Chenhao Wang <chenhwang4-c@my.cityu.edu.hk>.

in unsupervised learning – specifically, centroid clustering. A canonical clustering problem is described as: given a metric space \mathcal{X} with distance measure $d : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^+ \cup \{0\}$, a multiset $\mathcal{N} \subseteq \mathcal{X}$ of n data points (in this work, each data point is an *agent*), a set $\mathcal{M} \subseteq \mathcal{X}$ of possible centers and a positive integer k , the task is to find a subset $Y \subseteq \mathcal{M}$ of k cluster centers and assign each data point to its closest center in Y . The commonly studied objective is to make data points to be as close to their assigned centers as possible. Standard algorithms, such as k -means and k -medians, solve the clustering problem by satisfying a global criterion, where individual or group-wise happiness has not been taken into consideration. It has been noted that the globally efficient solutions are less preferred, especially when the application scenario is about public resources allocation (Conitzer et al., 2017; Fain et al., 2018). We consider the following facility location problem proposed in (Chen et al., 2019; Micha and Shah, 2020).

Example 1. *Suppose the government plans to build $k = 11$ identical parks to serve the residents, where every resident’s cost (e.g. gasoline) is proportional to the distance between her home and the closest park. There is a dense urban center with a population of 10,000 and 10 small suburbs, each of which has a population of 100. Suppose the suburbs are close to each other (e.g. 10km) compared with the distance between them and the urban center (e.g. 500km).*

Accordingly, by k -means or k -medians algorithms, the government will build 1 park at the urban center, and 10 parks for each small suburb. It is not hard to see such a plan is not fair: a single park is used to serve 10,000 people in the urban area, but each small suburb of 100 people has its own a park. This intuition is formalized by Moulin as the *principle of equal entitlement* (Moulin, 2003): A solution is fair if it respects the entitlements of groups of agents, namely, every subset of agents that is of sufficiently large size is entitled to choose a center for themselves. Such group-wise fairness is also referred as *core*, which has been extensively studied in game theory (Deng and Papadimitriou, 1994; Chalkiadakis et al., 2011), social choice (McKelvey, 1986; Feldman and Serrano, 2006) and fair division (Abdulkadiroğlu and Sönmez, 1998; Fain et al., 2018).

Chen et al. (2019) first formally and theoretically studied group fairness in the clustering problem, and proposed

the proportionality, which is essentially *core with non-transferable utilities*. Informally, a k -clustering is proportional if there is no coalition of agents with size at least n/k such that by deviating together to a new center each member of this coalition gets strictly better off. We note that proportionality overlooks the situation where the agents in a coalition can facilitate internal monetary transfers between themselves if they can file a justified claim to build a park at a new location that is better for the coalition overall.

In this work, we refine proportionality by *core with transferable utilities* or *core* for short. Formally, the cost of agent $i \in \mathcal{N}$ induced by a cluster center $y \in \mathcal{M}$ is $d(i, y)$, i.e., the distance between i and y ; and the cost induced by a k -clustering Y is $d(i, Y) \triangleq \min_{y \in Y} d(i, y)$, i.e., the minimum distance from i to any cluster center in Y .

Definition 2 (Core). For any k -clustering Y , a group of agents S with $|S| \geq \frac{n}{k}$ is called a blocking coalition, if there is a new center $y' \in \mathcal{M} \setminus Y$ such that by deviating to y' together, the total distance of S can be strictly decreased, i.e., $\sum_{i \in S} d(i, y') < \sum_{i \in S} d(i, Y)$. A k -clustering is called in the core or a core clustering if there is no blocking coalition.

It is not hard to verify that core fairness is stronger than proportionality (Chen et al., 2019) in the sense that a core clustering must be proportional, but not vice versa. Particularly, in Example 1, although a proportional clustering builds 10 parks for the city center and 1 park for all the suburbs, the one for suburbs can be arbitrarily built at any suburb’s location. However, a core clustering will select a more preferred location for this park to serve the 10 suburbs by minimizing their total distance. Finally, it can be shown that traditional learning algorithms (such as k -means) can be arbitrarily bad with respect to core fairness.

1.1. Main Contribution

As core clusterings are not guaranteed to exist for all instances, in this work, we provide two relaxation dimensions to weaken the corresponding requirements. The first dimension is in parallel with (Chen et al., 2019; Micha and Shah, 2020), where a valid blocking coalition should have large distance improvement. In the second dimension, different from their works, we study the relaxation when the size of a valid blocking coalition is required to be large enough. We formalize the two relaxations in the following definition.

Definition 3 (Approximate Core). For $\alpha \geq 1$ and $\beta \geq 1$, we say a k -clustering Y is in the (α, β) -core or an (α, β) -core clustering if there is no $S \subseteq \mathcal{N}$ with $|S| \geq \alpha \cdot \frac{n}{k}$ and $y' \in \mathcal{M} \setminus Y$ such that $\beta \cdot \sum_{i \in S} d(i, y') < \sum_{i \in S} d(i, Y)$.

We investigate to what extent these two relaxations can be satisfied by finding the smallest α and β such that (α, β) -core is nonempty. The relaxation on the size of blocking coalitions is regarded as α -dimension and that on the dis-

tance improvement as β -dimension. We consider both general metric space and special cases, including real line and discrete tree. Line and tree are two widely studied metric spaces, which is partly because they extensively exist in the real world. For example, line can be used to model the situations where people want to set proper temperatures for classrooms or schedule meeting times (Feldman and Wilf, 2013), and trees can be used to model referral or query networks (Kleinberg and Raghavan, 2005; Babai et al., 2012). As argued in (Micha and Shah, 2020), though objectives like truthfulness (Alon et al., 2010) and social welfare maximization (Feldman and Wilf, 2013) have received significant attention, fairness is largely overlooked.

Table 1. Our results for $(1, \beta)$ -core.

	Line	Tree	Metric Space
Upper Bound	$O(\sqrt{n})$ (Thm 8)	$O(\sqrt{n})$ (Thm 9)	$2\lceil \frac{n}{k} \rceil + 1$ (Thm 10)
Lower Bound	$\Omega(\sqrt{n})$ (Thm 7)	$\Omega(\sqrt{n})$ (Thm 9)	$\Omega(\sqrt{n})$ (Thm 7)

$(1, \beta)$ -Core. We first, in Section 3, study the relaxation in β -dimension, where α is fixed to be 1, i.e., $(1, \beta)$ -core. Our results are summarized in Table 1. Different to the study of proportionality in (Chen et al., 2019), where a constant approximation can be guaranteed for any metric space, we show that for core fairness, the existence of $(1, o(\sqrt{n}))$ -core clustering is not guaranteed, even in a real line. On the other hand, when the metric space is a real line or a tree, we present efficient algorithms that always output a $(1, O(\sqrt{n}))$ -core clustering, and thus we get the optimal approximation algorithm by relaxing the distance requirement solely. With respect to general metric space, we show that a greedy algorithm ensures $O(\frac{n}{k})$ -approximation. Beyond the study for arbitrary number k , when $k \geq \frac{n}{2}$, we show that for any metric space, there is a polynomial time algorithm which returns a $(1, 2)$ -core clustering, whereas determining the existence of $(1, 2 - \epsilon)$ -core clustering is NP-complete.

Table 2. Our results for $(\alpha, 1)$ -core.

	Line	Tree	Metric Space
Upper Bound	2 (Thm 14)	2 (Thm 15)	k (Thm 16)
Lower Bound	2 (Thm 13)	2 (Thm 15)	$\min\{k, \max\{\frac{k}{2}, \frac{n}{4}\}\}$ (Thm 16)

$(\alpha, 1)$ -Core. Next, in Section 4, we study the relaxation in α -dimension, where β is fixed to be 1, i.e., $(\alpha, 1)$ -core. Our results are summarized in Table 2. Different to the relaxation in β -dimension, we prove a $(2, 1)$ -core clustering is guaranteed to exist in line and tree spaces. We comple-

ment this result with a line instance where $(2 - \epsilon, 1)$ -core is empty for any $\epsilon > 0$. Thus our algorithms are optimal. For general metric space, we observe that a trivial upper-bound for α is k , which can be guaranteed by placing the k centers such that the total distance of all agents is minimized. We also complement this observation with a lower-bound instance where $(\alpha, 1)$ -core is empty for any $\alpha \leq \min\{k, \max\{\frac{k}{2}, \frac{n}{4}\}\}$, and thus our algorithm is tight up to a constant factor. Finally, we end this section by proving that determining the existence of $(\alpha, 1)$ -core clustering for any constant $\alpha \geq 1$ in general metric space is NP-complete.

(α, β) -Core In Section 5, we integrate previous results and study the case when both dimensions can be relaxed. Intuitively, sacrificing the approximation ratio in one dimension should be able to improve that in the other. We prove this intuition affirmatively by quantifying the tradeoff between the two relaxation dimensions. Specifically,

- for line or tree space and any $\alpha \in (1, 2]$, $(\alpha, \frac{1}{\alpha-1})$ -core is always non-empty (Thm 18);
- for general metric space and $\alpha > 1$, $(\alpha, \max\{4, \frac{2}{\alpha-1} + 3\})$ -core is always non-empty (Thm 19).

We want to highlight the significance of the above two results, especially for the general metric space, which is regarded as the major theoretical contribution of the current work. The results in Sections 3 and 4 imply that, in the general metric space, if $\alpha = 1$ is not relaxed, the best possible approximation ratio for β is $\Theta(\sqrt{n})$; on the other hand, if $\beta = 1$ is not relaxed, the best possible approximation ratio for α is $\max\{\frac{k}{2}, \frac{n}{4}\}$. However, our results in this section show that if we sacrifice a small constant on the approximation for one dimension, we can guarantee constant approximation for the other dimension. For example, by relaxing α to 2, $(2, 5)$ -core is always non-empty, and by relaxing β to 4, $(3, 4)$ -core is always non-empty.

Experiments Finally, in Section 6, we conduct experiments to examine the performance of our algorithms. We note that our algorithms have good theoretical guarantees in the worst case, but they may not find the fairest clustering for every instance. Accordingly, we first propose a two-stage algorithm to refine the clusters and then use synthetic and real-world data sets to show how much it outperforms classic ones regarding core fairness. Actually, the second stage of our algorithm provides us an interface to balance fairness and social efficiency. As shown by the experiments, our solution does not sacrifice much efficiency.

1.2. Other Related Works

Fairness Study in Machine Learning. The necessity of incorporating fairness into machine learning algorithms has

been well recognized in the recent decade. Various fairness concepts have been proposed based on different principles and to adapt to different machine learning tasks. For example, in another seminal work, Chierichetti et al. (2017) also studied fairness issue under clustering context but defined fairness as preserving equal representation for each protected class in every cluster based on disparate impact in US law. To satisfy their fairness requirement, Chierichetti et al. (2017) designed a two-step algorithm which first decomposes data points into fairlets and then runs classical clustering algorithms on those fairlets. On one hand, the algorithm is improved by a number of subsequent works in the sense of approximation ratios (Harb and Lam, 2020; Bercea et al., 2019) and running time (Huang et al., 2019; Schmidt et al., 2018; Backurs et al., 2019). On the other hand, Bera et al. (2019); Rösner and Schmidt (2018) and Braverman et al. (2019) extended the setting in (Chierichetti et al., 2017) to allow overlap in protected groups, consider multiple protected features or address fairness towards remote data points. Under the context of classification, fairness is captured by envy-freeness in (Zafar et al., 2017; Ustun et al., 2019; Balcan et al., 2019). We refer readers to the survey by Mehrabi et al. (2019) for a detailed discussion on fairness study in various machine learning contexts.

Fair Resource Allocation. Group fairness is recently considered in the resource allocation field when resources are allocated among groups of agents. Based on how groups are formed, the models can be classified into two categories. The first one is when agents are partitioned into fixed groups and the fairness is only concerned with the pre-existing groups (Segal-Halevi and Suksompong, 2019; Benabbou et al., 2019; Kyropoulou et al., 2020). The second one is to consider the fairness of arbitrarily formed groups of agents (Aziz and Rey, 2020; Berliant et al., 1992; Conitzer et al., 2019; Hossain et al., 2020). Our problem falls under the umbrella of the second category. However, the items in our work are public and fairness is defined for collective utilities. Fair clustering problem is also related to public resource allocation, where resources can be shared by agents. For public resources, one popular research agenda is to impose (combinatorial) constraints on the allocations (Conitzer et al., 2017; Cheng et al., 2019; Fain et al., 2018; Li and Li, 2020). However, in our setting, all centers will be built without constraints.

2. Preliminaries

Recall that $\mathcal{N} \subseteq \mathcal{X}$ is a multiset of n agents in a metric space (\mathcal{X}, d) , where for any $\{x_1, x_2, x_3\} \subseteq \mathcal{X}$, $d(x_1, x_2) + d(x_2, x_3) \geq d(x_1, x_3)$. Note that we allow repetitions in \mathcal{N} which means multiple agents can be located at the same position. We refer to agents and their locations interchangeably, when no confusion arises. $\mathcal{M} \subseteq \mathcal{X}$ is the set of feasible

locations for cluster centers. Our task is to find $Y \in \mathcal{M}^k$ to place the k centers such that the clustering is in the (approximate) core. We first present an example where an exact core clustering does not exist, and illustrate the two relaxation dimensions so that an approximate core clustering exists.

Example 4. Consider a complete graph $K_4 = (V, E)$ with 4 vertices and the distance between any two vertices is 1. Let $\mathcal{X} = \mathcal{M} = V$. Suppose that at each vertex lies an agent (i.e., $\mathcal{N} = V$), and we want to locate $k = 2$ centers to cluster these $n = 4$ agents. First, we note that the (exact) core is empty for this instance, because for any 2-clustering $Y \in \mathcal{M}^2$, the remaining $n/k = 2$ agents (say u, v) in $V \setminus Y$ can form a blocking coalition and deviate to a new center $v \in V \setminus Y$ such that $d(u, v) + d(v, v) = 1 < 2 = d(u, Y) + d(v, Y)$. Second, for all $\beta \geq 2$, every solution $Y \in \mathcal{M}^2$ is a $(1, \beta)$ -core clustering, because for any group $S \subseteq V$ with $|S| \geq 2$ and any possible deviating center $v \in V$, we have $\beta \sum_{i \in S} d(i, v) \geq 2 \geq \sum_{i \in S} d(i, Y)$. Finally, for all $\alpha > 1$, every solution $Y \in \mathcal{M}^2$ is an $(\alpha, 1)$ -core clustering, because for any group $S \subseteq V$ with $|S| \geq \alpha \cdot \frac{n}{k} > 2$ and any possible deviating center $v \in V$, we have $\sum_{i \in S} d(i, v) \geq 2 \geq \sum_{i \in S} d(i, Y)$.

Motivated by the above example, our task is to compute the smallest β or α such that $(1, \beta)$ -core or $(\alpha, 1)$ -core is non-empty. Besides the general metric space, we are also interested in two special spaces, namely, real line \mathbb{R} and graph spaces such as tree.

Line. $\mathcal{X} = \mathcal{M} = \mathbb{R}$ and the distance d is Euclidean, i.e., the agents and centers can be at anywhere in the real line.

Graph Space. Let $G = (V, E)$ be an undirected tree graph. The edges in E may have specified length. In the graph space induced by G , $\mathcal{X} = \mathcal{M} = V$, and the distance d between two vertices is the length of the shortest path. A tree space is induced by a tree graph. Note that line is a continuous space where every point is feasible for cluster centers, while the graph space is discrete and centers can only be placed on its vertex set.

The following Lemma 5 enables us to only focus on the blocking coalitions with minimum possible size.

Lemma 5. Given $\alpha, \beta \geq 1$ and a solution $Y \in [\mathcal{M}]^k$, if there is a group S of size $|S| > \lceil \frac{\alpha n}{k} \rceil$ such that $\beta \sum_{i \in S} d(i, y') < \sum_{i \in S} d(i, Y)$ for some $y' \in \mathcal{M}$, then there is a group $S' \subseteq S$ of size $|S'| = \lceil \frac{\alpha n}{k} \rceil$ such that $\beta \sum_{i \in S'} d(i, y') < \sum_{i \in S'} d(i, Y)$.

Proof. If S and $y' \in \mathcal{M}$ satisfy $|S| > \lceil \frac{\alpha n}{k} \rceil$ and $\frac{\sum_{i \in S} d(i, Y)}{\sum_{i \in S} d(i, y')} > \beta$, then there must exist an agent $w \in S$ subject to

$$\frac{d(w, Y)}{d(w, y')} \leq \frac{\sum_{i \in S} d(i, Y)}{\sum_{i \in S} d(i, y')}.$$

Let $S' = S \setminus \{w\}$. Then we have

$$\beta < \frac{\sum_{i \in S} d(i, Y)}{\sum_{i \in S} d(i, y')} \leq \frac{\sum_{i \in S'} d(i, Y)}{\sum_{i \in S'} d(i, y')}.$$

Repeating this process of removing one agent until $|S'| = \lceil \frac{\alpha n}{k} \rceil$, we establish the proof. \square

When $k = 1$ or $n - 1$, the solution that minimizes the total distance to all agents is in the core. When $k \geq \frac{n}{2}$ and the space is a connected graph $G(V, E)$ with $V = \mathcal{X} = \mathcal{N} = \mathcal{M}$, by contrast with the result in (Micha and Shah, 2020) where a proportional k -clustering always exists and can be computed efficiently, in our setting a core clustering is not guaranteed. We state a stronger result as follows.

Proposition 6. When $k = 1$ or $k \geq n - 1$, the core is always non-empty. When $\frac{n}{2} \leq k \leq n - 2$, for any $0 < \epsilon \leq 1$, the existence of a $(1, 2 - \epsilon)$ -core clustering is not guaranteed.

Proof. For the first argument, let

$$Y^* = \arg \min_{Y \in \mathcal{M}^k} \sum_{i \in N} d(i, Y), \text{ for } k = 1 \text{ or } n - 1.$$

If $k = 1$, by definition it is obvious that Y^* is in the core as the ground agent set N is the unique possible deviating coalition. For $k = n - 1$, a deviating coalition should contain 2 agents. Suppose for contradiction that i and j can decrease their cost by deviating to $y' \notin Y^*$. Then we can construct a new clustering Y' : (1) $y' \in Y'$ and thus $d(i, Y^*) + d(j, Y^*) > d(i, Y') + d(j, Y')$; (2) each of the other $n - 2$ centers is at the optimal position for one of the remaining $n - 2$ agents in $N \setminus \{i, j\}$ and thus $d(l, Y') \leq d(l, Y^*)$ for $l \in N \setminus \{i, j\}$. Therefore,

$$\sum_{i \in N} d(i, Y^*) > \sum_{i \in N} d(i, Y'),$$

which is a contradiction with the definition of Y^* .

For the second argument, we consider the graph space induced by a complete graph $G = (V, E)$ with n vertices, where at each vertex lies an agent, and $V = \mathcal{X} = \mathcal{N} = \mathcal{M}$. The distance between any two vertices is 1. When $\frac{n}{2} \leq k \leq n - 2$, the minimum size of a possible blocking coalition is $\lceil \frac{n}{k} \rceil = 2$. For every k -clustering Y , there must exist two agents without center at their locations. Then they form a blocking coalition (w.r.t. $(1, 2 - \epsilon)$ -core) with a deviating center on one of them: their total distance to the deviating center is 1, while their total distance to Y is 2. Hence, a $(1, 2 - \epsilon)$ -core clustering does not exist. \square

3. $(1, \beta)$ -Core

In this section, we study the relaxation on distance improvement, and show to what extent a $(1, \beta)$ -core is non-empty.

3.1. Line

First, a $(1, o(\sqrt{n}))$ -core clustering is not guaranteed to exist.

Theorem 7. *There is a line instance such that the $(1, o(\sqrt{n}))$ -core is empty.*

Proof. Consider an instance building in the real line and a set of integer points $V = \{1, 2, \dots, k+1\}$ in this line, each of which accommodates k agents. A k -clustering is required to serve all $n = k(k+1)$ agents by k centers.

Suppose for contradiction that Y is a $(1, o(\sqrt{n}))$ -core clustering. Since there are $k+1$ agents' locations and k centers to be located, there must be a point $j \in V$ so that no center is located in the interval $(j - \frac{1}{2}, j + \frac{1}{2})$, i.e., $(j - \frac{1}{2}, j + \frac{1}{2}) \cap Y = \emptyset$. Assume w.l.o.g. that $j \neq k+1$ by symmetry. Consider a group S consisting of k agents located at j and one agent located at $j+1$. Its size is $|S| = k+1 = \frac{n}{k}$, which entitles itself to choose a center. Because every agent at j has a distance at least $\frac{1}{2}$ from its location to solution Y , the total distance of this group is $\sum_{i \in S} d(i, Y) \geq \frac{1}{2} \cdot k$. However, if they deviate to a new center $y' = j$, their total distance changes into $\sum_{i \in S} d(i, y') = 1$. Then we have $\frac{\sum_{i \in S} d(i, Y)}{\sum_{i \in S} d(i, y')} \geq \frac{k}{2} > o(\sqrt{n})$, which is a contradiction. Hence, the $(1, o(\sqrt{n}))$ -core is empty. \square

Algorithm 1 $ALG_l(\lambda)$ for Line.

Input: Agents $\mathbf{x} = \{x_1, \dots, x_n\}$, number $k \in \mathbb{N}^+$

Output: k -clustering $Y = \{y_1, \dots, y_k\}$

- 1: Rename the agents such that $x_1 \leq \dots \leq x_n$.
 - 2: **for** $i = 1, 2, \dots, k-1$ **do**
 - 3: Locate a center at $y_i = x_{\lambda i}$.
 - 4: **end for**
 - 5: Let $r = \min\{\lambda k, n\}$, and locate a center at $y_k = x_r$.
-

Next, we present our algorithm ALG_l , as shown in Algorithm 1, which matches the lower-bound in Theorem 7. Roughly, ALG_l has a tuneable parameter λ , and guarantees that the number of agents between any two contiguous centers to be no more than $\lambda - 1$. By selecting the optimal λ depending on k , we can obtain the tight approximation.

Theorem 8. *For any line instance, a $(1, O(\sqrt{n}))$ -core clustering can be found in linear time. Specifically, $ALG_l(\lceil \frac{n}{k} \rceil)$ gives a $(1, \lceil \frac{n}{k} \rceil - 1)$ -core clustering if $k = \Omega(\sqrt{n})$, and $ALG_l(\lceil \frac{n}{k+1} \rceil)$ gives a $(1, k)$ -core clustering if $k = o(\sqrt{n})$.*

Proof. Let $\mathbf{x} = \{x_1, \dots, x_n\}$ be the locations of agents. When $k = \Omega(\sqrt{n})$ and $k = o(\sqrt{n})$, we define $\lambda = \lceil \frac{n}{k} \rceil$ and $\lambda = \lceil \frac{n}{k+1} \rceil$ respectively, and implement $ALG_l(\lambda)$. The output is a k -clustering $Y = \{y_1, \dots, y_k\}$. Now we prove this solution is a $(1, \lceil \frac{n}{k} \rceil - 1)$ -core clustering when $k = \Omega(\sqrt{n})$,

and a $(1, k)$ -core clustering when $k = o(\sqrt{n})$. Therefore, it is always a $(1, O(\sqrt{n}))$ -core clustering.

Suppose for contradiction that there is a blocking coalition $S \subseteq \mathcal{N}$ of agents with $|S| \geq \frac{n}{k}$ and a deviating center $y' \in \mathbb{R} \setminus Y$ satisfying $r := \frac{\sum_{i \in S} d(i, Y)}{\sum_{i \in S} d(i, y')} > \lceil \frac{n}{k} \rceil - 1$ (resp. $r > k$) when $k = \Omega(\sqrt{n})$ (resp. $k = o(\sqrt{n})$). Set two virtual points $y_0 = -\infty$ and $y_{k+1} = +\infty$. Assume w.l.o.g. $y' \in (y_j, y_{j+1})$ for some $j = 0, \dots, k$ and $d(y_j, y') \leq d(y_{j+1}, y')$. Let (N_1, N_2, N_3) be a partition of S with $N_1 = \{i \in S | x_i \leq y_j\}$, $N_2 = \{i \in S | y_j < x_i < y_{j+1}\}$, and $N_3 = \{i \in S | x_i \geq y_{j+1}\}$. Note that the algorithm guarantees $|N_2| \leq \lambda - 1$. Then we have

$$\begin{aligned} \sum_{i \in S} d(i, Y) &\leq \sum_{i \in N_1} d(i, y_j) + \sum_{i \in N_2} (d(i, y') + d(y', y_j)) \\ &\quad + \sum_{i \in N_3} d(i, y_{j+1}), \end{aligned} \quad (1)$$

$$\begin{aligned} \sum_{i \in S} d(i, y') &= \sum_{i \in N_1} (d(i, y_j) + d(y_j, y')) + \sum_{i \in N_2} d(i, y') \\ &\quad + \sum_{i \in N_3} (d(i, y_{j+1}) + d(y_{j+1}, y')). \end{aligned} \quad (2)$$

Combining Equations (1) and (2), it follows that

$$r \leq \frac{\sum_{i \in N_2} d(y', y_j)}{\sum_{i \in N_1} d(y_j, y') + \sum_{i \in N_3} d(y_{j+1}, y')} \leq \frac{|N_2|}{|N_1 \cup N_3|},$$

where the second inequality is due to the assumption $d(y_j, y') \leq d(y_{j+1}, y')$.

When $k = \Omega(\sqrt{n})$ and $\lambda = \lceil \frac{n}{k} \rceil$, we have $|N_2| \leq \lceil \frac{n}{k} \rceil - 1$, and $|N_1 \cup N_3| = |S| - |N_2| \geq 1$. It indicates that $r \leq \lceil \frac{n}{k} \rceil - 1$ which is a contradiction. So Y is a $(1, \lceil \frac{n}{k} \rceil - 1)$ -core clustering.

When $k = o(\sqrt{n})$ and $\lambda = \lceil \frac{n}{k+1} \rceil$, we have $|N_2| \leq \lceil \frac{n}{k+1} \rceil - 1$, and $|N_1 \cup N_3| = |S| - |N_2| \geq \lceil \frac{n}{k} \rceil - \lceil \frac{n}{k+1} \rceil + 1$. By a simple calculation we have $r \leq k$ as a contradiction. So Y is a $(1, k)$ -core clustering. \square

3.2. Tree

Both the lower- and upper-bound results for line space can be extended to trees. For the proportionality fairness, Micha and Shah (2020) proposed an algorithm ‘‘Proportionally Fair Clustering for Trees (PFCT)’’ as follows, which always returns a proportional solution for trees. A rooted tree (G, r) is obtained by rooting the tree at an arbitrary node r . Let $\text{level}(x)$ denote the height of node x relative to the root r (with $\text{level}(r) = 1$), and $\text{ST}(x)$ denote the subtree rooted at node x (i.e. the set of nodes v with $\text{level}(v) \leq \text{level}(x)$ and

the unique path from v to r passes by x). Let $|ST(x)|$ be the number of agents contained in the subtree $ST(x)$. PFCT traverses all the nodes from the highest level ones (i.e. the leaves), locates a center on the node whose subtree contains at least $\lceil \frac{n}{k} \rceil$ agents, and then deletes this subtree. At the end, agents are assigned to the closest center in the output. We adapt PFCT into the following $ALG_t(\lambda)$ with a tuneable parameter λ , which controls the threshold value for locating a center. When $\lambda = \lceil \frac{n}{k} \rceil$, $ALG_t(\lambda)$ is equivalent to PFCT.

Algorithm 2 $ALG_t(\lambda)$ for Tree.

Input: A tree $G = (V, E)$, n agents and integer k

Output: k -clustering set Y

```

1: Let  $r$  the root of tree  $G$  and  $d$  be the height.
2:  $Y \leftarrow \emptyset$ 
3:  $G^d \leftarrow G$ 
4: for  $l = d$  to 1 do
5:    $G^{l-1} \leftarrow G^l$ 
6:   for every  $x \in V$  with level  $(x) = l$  and  $|ST(x)| \geq \lambda$  do
7:     if  $|Y| < k$  then
8:        $Y \leftarrow Y \cup \{x\}$  and  $G^{\ell-1} \leftarrow G^{\ell-1} \setminus ST(x)$ 
9:     end if
10:  end for
11: end for
    
```

Note that it always has $|Y| \leq k$, and if $|Y| < k$, we can build $k - |Y|$ more centers arbitrarily. When $\lambda = \lceil \frac{n}{k} \rceil$ or $\lceil \frac{n}{k+1} \rceil$, we observe that, after removing all centers in Y from the tree, the number of agents in every component is at most $\lambda - 1$. It can be observed that $ALG_t(\lambda)$ is an extension of $ALG_l(\lambda)$ in the tree, and we have the following theorem.

Theorem 9. *For any instance in the tree, we can find a $(1, O(\sqrt{n}))$ -core clustering efficiently. In particular, when $k = \Omega(\sqrt{n})$, $ALG_t(\lceil \frac{n}{k} \rceil)$ returns a $(1, \lceil \frac{n}{k} \rceil - 1)$ -core clustering, and when $k = o(\sqrt{n})$, $ALG_t(\lceil \frac{n}{k+1} \rceil)$ returns a $(1, k)$ -core clustering. Moreover, a $(1, o(\sqrt{n}))$ -core clustering is not guaranteed to exist.*

3.3. General Metric Space

For the general metric space, we show that a simple greedy algorithm, ALG_g , as described in Algorithm 3, has the desired theoretical guarantee. For each $x \in \mathcal{X}$, we use $B(x, \delta) = \{i \in \mathcal{N} \mid d(i, x) \leq \delta\}$ to denote the set of agents in the ball with center x and radius δ . ALG_g continuously grows the radius of each ball centered on each possible center, with the same speed. When a ball is large enough to contain at least $\lceil \frac{n}{k} \rceil$ points, we open a cluster center at that ball center. Actually, the underlying idea of expanding balls of points has been utilized for k -median problems (Jain and Vazirani, 1999), and Chen et al. (2019) and Micha and Shah (2020) have proved that ALG_g also has good

theoretical performance regarding proportional fairness. We note that ALG_g may output less than k centers, and if so, the remaining centers can be selected arbitrarily. In Section 6, we will show how to refine ALG_g beyond the worst case.

Algorithm 3 ALG_g for General Metric Space.

Input: Metric space (\mathcal{X}, d) , agents $\mathcal{N} \subseteq \mathcal{X}$, possible locations $\mathcal{M} \subseteq \mathcal{X}$, and $k \in \mathbb{N}^+$

Output: k -clustering Y

```

1:  $\delta \leftarrow 0$ ;  $Y \leftarrow \emptyset$ ;  $N \leftarrow \mathcal{N}$ .
2: while  $N \neq \emptyset$  do
3:   Smoothly increase  $\delta$ 
4:   while  $\exists x \in Y$  s.t.  $|B(x, \delta) \cap N| \geq 1$  do
5:      $N \leftarrow N \setminus B(x, \delta)$ 
6:   end while
7:   while  $\exists x \in \mathcal{M} \setminus Y$  s.t.  $|B(x, \delta) \cap N| \geq \lceil \frac{n}{k} \rceil$  do
8:      $Y \leftarrow Y \cup \{x\}$  and  $N \leftarrow N \setminus B(x, \delta)$ 
9:   end while
10: end while
    
```

Theorem 10. *For any instance in metric space (\mathcal{X}, d) , algorithm ALG_g always outputs a $(1, 2\lceil \frac{n}{k} \rceil + 1)$ -core clustering.*

Proof. Let Y be the k -clustering returned by ALG_g . By Lemma 5, it suffices to prove that, for any set of agents $S \subseteq \mathcal{N}$ with $|S| = \lceil \frac{n}{k} \rceil$, $\sum_{i \in S} d(i, Y) \leq (2\lceil \frac{n}{k} \rceil + 1) \sum_{i \in S} d(i, y')$ holds for any point $y' \in \mathcal{M} \setminus Y$. Let $y^* \in Y$ be the center closest to y' , i.e., $y^* \in \arg \min_{y \in Y} d(y, y')$. Since for any $i \in S$, $d(i, Y) \leq d(i, y^*)$, we have

$$\sum_{i \in S} d(i, Y) \leq \sum_{i \in S} d(i, y^*) \leq \sum_{i \in S} d(i, y') + \sum_{i \in S} d(y^*, y'). \quad (3)$$

We discuss two cases.

Case 1: $\max_{i \in S} d(i, y') \geq \frac{1}{2} d(y^*, y')$. If so, then we have

$$\sum_{i \in S} d(i, y') \geq \max_{i \in S} d(i, y') \geq \frac{1}{2} d(y^*, y').$$

Combining with (3), it follows that

$$\begin{aligned} \frac{\sum_{i \in S} d(i, Y)}{\sum_{i \in S} d(i, y')} &\leq 1 + \frac{\sum_{i \in S} d(y^*, y')}{\sum_{i \in S} d(i, y')} \\ &\leq 1 + \frac{\sum_{i \in S} d(y^*, y')}{d(y^*, y')/2} = 1 + 2\lceil \frac{n}{k} \rceil. \end{aligned}$$

Case 2: $\max_{i \in S} d(i, y') < \frac{1}{2} d(y^*, y')$. Let $\delta^* = \max_{i \in S} d(i, y')$, and accordingly, $S \subseteq B(y', \delta^*)$. If there exists a center $y'' \in Y$ such that $B(y'', \delta^*) \cap B(y', \delta^*) \neq \emptyset$, then we have $d(y'', y') \leq 2\delta^* < d(y^*, y')$, which however, contradicts to the definition of y^* . So when the algorithm operating radius δ as δ^* , all balls with center in

Y have an empty intersection with ball $B(y', \delta^*)$, that is, $B(y', \delta^*) \cap B(y, \delta^*) = \emptyset, \forall y \in Y$. Since $|S| = \lceil \frac{n}{k} \rceil$ and $S \subseteq B(y', \delta^*)$, point y' must be selected as a center by the algorithm, contradicting to $y' \notin Y$. Therefore, this case would never occur. \square

When the number k of cluster centers is large, we are able to get better approximations. Recall Proposition 6 that when $k \geq n - 1$, the core is always non-empty. We complement this result by the following theorem.

Theorem 11. *For any metric space with $\frac{n}{2} \leq k \leq n - 2$, there is an algorithm that computes a $(1, 2)$ -core clustering in $O(n^2 \log n)$ time.*

The $(1, 2)$ -core clustering provided above is best possible, because by Proposition 6, a $(1, 2 - \epsilon)$ -core clustering may not exist. Finally, we end this section by proving that deciding the existence of a $(1, 2 - \epsilon)$ -core clustering is hard.

Theorem 12. *For any $0 < \epsilon \leq 1$, the problem of determining the existence of a $(1, 2 - \epsilon)$ -core clustering is NP-complete, even if $k \geq \frac{n}{2}$ and it is in a graph space induced by $G = (V, E)$ with $|V| = n$.*

4. $(\alpha, 1)$ -Core

Next, we study to what extent core fairness can be achieved when only the requirement of blocking coalition's size is relaxed, i.e., finding the smallest α such that $(\alpha, 1)$ -core is non-empty.

4.1. Line

Theorem 13. *In line space, for any $\epsilon > 0$, there exists an instance such that the $(2 - \epsilon, 1)$ -core is empty.*

Proof. Consider an instance in the real line with $n = C(2C - 1)$ agents and $k = 2C - 1$ centers to be opened for some integer $C \geq \frac{3}{\epsilon}$. Let K be a sufficiently large number, say $K = C^n$. The agents are partitioned into C parts (N_1, \dots, N_C) such that for any $j \in [C]$, $N_j = \{i_1, i_2, \dots, i_{2C-1}\}$ contains $2C - 1$ agents where $\{i_1, \dots, i_{C-1}\}$ are located at jK , agent i_C is located at $jK + 1$ and $\{i_{C+1}, \dots, i_{2C-1}\}$ are located at $jK + 2$.

As we only open $2C - 1$ centers, for any $(2 - \epsilon, 1)$ -core clustering, there exists N_j such that all agents in N_j are served by a single center. Since the agents in $\{i_1, \dots, i_{C-1}\}$ and $\{i_{C+1}, \dots, i_{2C-1}\}$ are symmetric with respect to i_C , assume w.l.o.g. the center is placed at $x \geq jK + 1$. Note that $2C - 3 \geq (2 - \epsilon)\lceil \frac{n}{k} \rceil$. We show that $S = \{i_1, \dots, i_{2C-3}\}$ forms a blocking coalition w.r.t. $(2 - \epsilon, 1)$ -core by deviating to a new center i_1 :

$$\sum_{i \in S} d(i, i_1) = 1 + 2(C - 3) < 2C - 4$$

$$= \sum_{i \in S} d(i, i_C) \leq \sum_{i \in S} d(i, x).$$

Hence, there does not exist a $(2 - \epsilon, 1)$ -core clustering. \square

Fortunately, by selecting a proper parameter λ , Algorithm ALG_l guarantees the tight approximation ratio.

Theorem 14. *For any line instance, $ALG_l(\lceil \frac{n}{k} \rceil)$ returns a $(2, 1)$ -core clustering.*

4.2. Tree

We continue to consider the tree spaces.

Theorem 15. *For any tree instance, we can find a $(2, 1)$ -core clustering efficiently. For any $\epsilon > 0$, a $(2 - \epsilon, 1)$ -core clustering is not guaranteed to exist in a tree.*

4.3. General Metric Space

By the definition of $(\alpha, 1)$ -core clustering, a $(k, 1)$ -core clustering always exists, because any potential blocking coalition must contain all agents and any solution containing the optimal single center is a $(k, 1)$ -core clustering. Next, we show that this trivial upper-bound is asymptotically tight.

Theorem 16. *The $(k, 1)$ -core is always non-empty. Further, the existence of an $(\alpha, 1)$ -core clustering, for any $\alpha \leq \min\{k, \max\{\frac{k}{2}, \frac{n}{4}\}\}$, is not guaranteed.*

Finally, we give a hardness result for the existence of an $(\alpha, 1)$ -core clustering.

Theorem 17. *For any given constant $\alpha \geq 1$, the problem of determining the existence of an $(\alpha, 1)$ -core clustering is NP-complete.*

5. (α, β) -Core

Finally, we study the approximate fairness by relaxing both dimensions simultaneously. Recalling the results in Sections 3.1 and 4.1, if $\alpha = 1$, the best possible approximation of β is $\Theta(\sqrt{n})$; however, if α is relaxed to 2, we are able to get the optimum in the β -dimension, i.e., $\beta = 1$. The following theorem shows the exact tradeoff between the approximations in both dimensions for line and tree spaces.

Theorem 18. *For any $\alpha > 1$, every line or tree instance has a non-empty (α, β) -core with $\beta = \max\{1, \frac{1}{\alpha-1}\}$.*

Next, we extend the above theorem to general metric space. Again our results in Sections 3.3 and 4.3 show that if α -dimension is not relaxed, the best approximation in β -dimension is $\Theta(\sqrt{n})$; and if β -dimension is not relaxed, the best approximation in α -dimension is $\max\{\frac{k}{2}, \frac{n}{4}\}$. As we will see in the following theorem, however, if we sacrifice a small constant for one dimension, we can guarantee constant approximation for the other as well.

Theorem 19. For any $\alpha > 1$, any instance in a metric space has a non-empty (α, β) -core with $\beta = \max\left\{4, \frac{2}{\alpha-1} + 3\right\}$.

Proof. Let Y be the k -clustering returned by ALG_g . By Lemma 5, it suffices to prove that, for any set of agents $S \subseteq \mathcal{N}$ with $|S| = \lceil \frac{\alpha n}{k} \rceil$,

$$\sum_{i \in S} d(i, Y) \leq \max\left\{4, \frac{2}{\alpha-1} + 3\right\} \cdot \sum_{i \in S} d(i, y')$$

holds for any point $y' \in \mathcal{M} \setminus Y$. Suppose for contradiction that there is a blocking coalition $S \subseteq \mathcal{N}$ with $|S| = \lceil \frac{\alpha n}{k} \rceil$ and a deviating center $y' \in \mathcal{M} \setminus Y$.

Let $y^* \in Y$ be the cluster center closest to y' , i.e., $y^* \in \arg \min_{y \in Y} d(y, y')$. Note that for any $i \in S$, $d(i, Y) \leq d(i, y^*)$. Because S is a blocking coalition, we have

$$r := \frac{\sum_{i \in S} d(i, Y)}{\sum_{i \in S} d(i, y')} > \beta = \max\left\{4, \frac{2}{\alpha-1} + 3\right\}.$$

Consider an open ball X centered at y' with radius $R := \frac{d(y', y^*)}{2}$, $X = \{i \in S \mid d(i, y') < R\}$. Note that $|X| \leq \lceil \frac{n}{k} \rceil - 1$, otherwise by ALG_g , y' should be selected as a center. Define

$$r_1 := \frac{\sum_{i \in X} d(i, Y)}{\sum_{i \in S} d(i, y')} \quad \text{and} \quad r_2 := \frac{\sum_{i \in S \setminus X} d(i, Y)}{\sum_{i \in S} d(i, y')}.$$

Then $r = r_1 + r_2$. For r_1 , if $r_1 \geq 1$, we have,

$$\begin{aligned} r_1 &\leq \frac{\sum_{i \in X} d(i, y') + \sum_{i \in X} d(y', y^*)}{\sum_{i \in X} d(i, y') + \sum_{i \in S \setminus X} d(i, y')} \leq \frac{\sum_{i \in X} d(y', y^*)}{\sum_{i \in S \setminus X} d(i, y')} \\ &\leq \frac{|X| \cdot 2R}{|S \setminus X| \cdot R} \leq \frac{2 \left(\lceil \frac{n}{k} \rceil - 1\right)}{\lceil \frac{\alpha n}{k} \rceil - \lceil \frac{n}{k} \rceil + 1} \leq \frac{2}{\alpha - 1}. \end{aligned}$$

Combining with the other case $r_1 < 1$, we have $r_1 \leq \max\left\{1, \frac{2}{\alpha-1}\right\}$. Then, we derive an upper-bound for r_2 .

$$r_2 \leq \frac{\sum_{i \in S \setminus X} (d(i, y') + d(y', y^*))}{\sum_{i \in S \setminus X} d(i, y')} \leq 1 + \frac{|S \setminus X| \cdot 2R}{|S \setminus X| \cdot R} = 3,$$

where the last inequality is due to $d(y', y^*) = 2R$ and $d(i, y') \geq R$ for any $i \in S \setminus X$. Combing the upper-bounds of r_1, r_2 , we have $\beta < r = r_1 + r_2 \leq \max\left\{4, \frac{2}{\alpha-1} + 3\right\} = \beta$, which is a contradiction. \square

6. Experiments

6.1. A Refined Algorithm $ALG_g^+(\text{obj})$

Before examining the performance of our algorithm, we note that though ALG_g has good guarantee in the worst-case, it

may not produce the fairest clustering in every instance. For example, in Figure 1a, we randomly partition each of 1,000 nodes into 3 Gaussian-distributed sets with probability 0.2, 0.3 and 0.5 (from left to right). We want to build $k = 10$ centers to serve the nodes; however, ALG_g only returns 4 clusters whose centers are shown by the red stars. Obviously, the extra 6 centers can significantly improve its performance, regarding both fairness and efficiency.

Algorithm 4 $ALG_g^+(\text{obj})$ for General Metric Space.

Input: Metric space (\mathcal{X}, d) , agents $\mathcal{N} \subseteq \mathcal{X}$, possible locations $\mathcal{M} \subseteq \mathcal{X}$, and $k \in \mathbb{N}^+$

Output: k -clustering C .

- 1: Initialize $C = \emptyset$.
 - 2: Run ALG_g on $(\mathcal{M}, \mathcal{N}, k)$ and get $Y = \{y_1, \dots, y_k\}$.
 - 3: Let \mathcal{N}_i be the corresponding cluster centered at y_i .
 - 4: Rename so that $(|\mathcal{N}_1| \bmod \frac{n}{k}) \geq \dots \geq (|\mathcal{N}_{k'}| \bmod \frac{n}{k})$.
 - 5: Let $r = k - \sum_{i=1}^{k'} \lfloor \frac{|\mathcal{N}_i|}{n/k} \rfloor$.
 - 6: **for** $i = 1, \dots, k'$ **do**
 - 7: Let $k_i = \lceil \frac{|\mathcal{N}_i|}{n/k} \rceil$ if $i \leq r$; otherwise, $k_i = \lfloor \frac{|\mathcal{N}_i|}{n/k} \rfloor$.
 - 8: $\{y_{i1}, \dots, y_{ik_i}\} = \arg \min \text{obj}(\mathcal{M}, \mathcal{N}_i, k_i)$.
 - 9: $C \leftarrow C \cup \{y_{i1}, \dots, y_{ik_i}\}$.
 - 10: **end for**
-

To improve the performance, we refine ALG_g in Algorithm 4, denoted by $ALG_g^+(\text{obj})$. Roughly, we first use ALG_g to obtain a preliminary clustering Y and resultant partition $\mathcal{N} = \cup_i \mathcal{N}_i$, which guarantees the fairness in the worst case, and then we proportionally assign all centers to these clusters according to their populations, i.e., $\sum_i k_i = k$ and $k_i \propto |\mathcal{N}_i|$. Within each preliminary cluster i , the real centers are selected to optimize a social objective function $\text{obj}(\mathcal{M}, \mathcal{N}_i, k_i)$ in Line 8. For example, when obj is the k -means objective (i.e., minimizing the squared Euclidean distance), we refer our algorithm to $ALG_g^+(k\text{-means})$. Thus, obj actually provides us an interface to balance fairness and social efficiency. In the experiment shown in Figure 1a, we feed $ALG_g^+(k\text{-means})$ with $k\text{-means++}$ algorithm (Arthur and Vassilvitskii, 2007)¹, which builds centers proportionally to the populations of the three Gaussian sets (i.e., 2,3,5 centers for each). However, if we directly use $k\text{-means++}$ algorithm on \mathcal{N} , it builds 4,3,3 centers for each Gaussian cluster, as shown in Figure 1b, where the right set contains half of all points but only gets 3 centers.

6.2. Experiments

We implement experiments on two qualitatively different datasets used for clustering. (1) Gaussian dataset (synthetic): the experiment in Section 6.1 (Figure 1a and 1b) is repeated

¹ $k\text{-means++}$ algorithm is Lloyd's algorithm for $k\text{-means}$ minimization objective with a particular initialization.

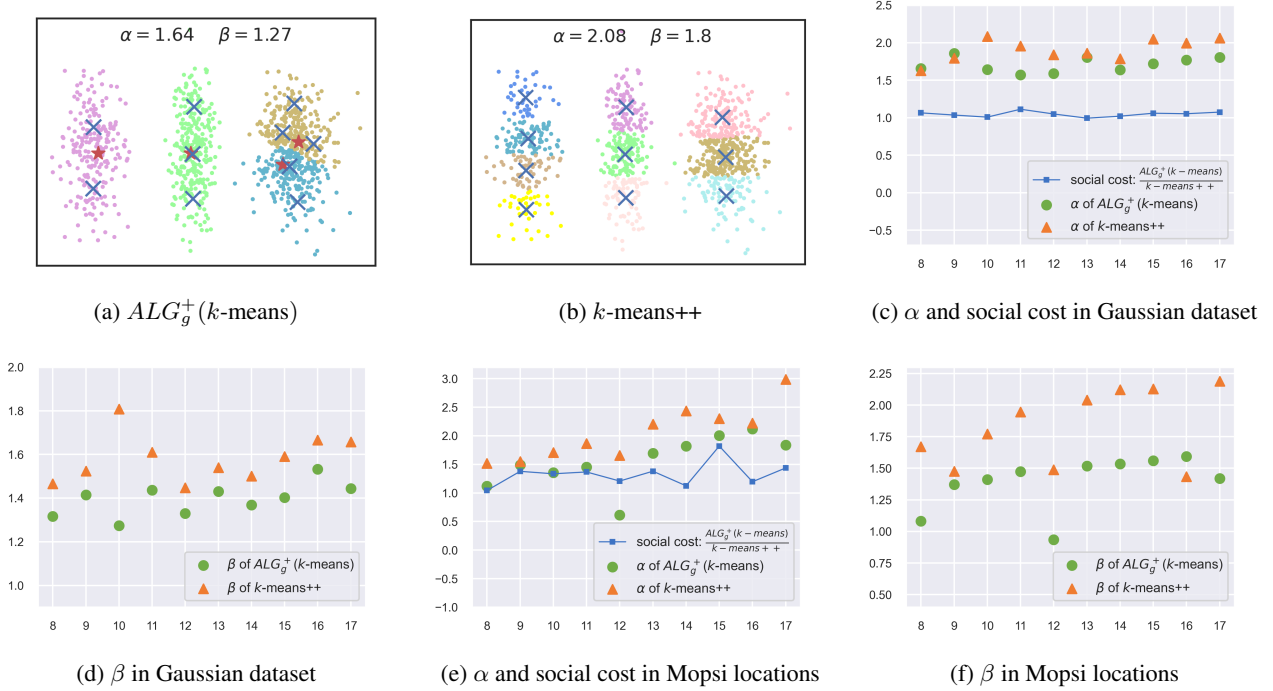


Figure 1. (a) and (b) depict the clustering centers when two algorithms cluster Gaussian dataset for $k = 10$. For a range of $k = 8, \dots, 17$ (horizontal axis), (c) and (d) (resp. (e) and (f)) compare the fairness and efficiency in Gaussian dataset (resp. Mopsi locations).

for k from 8 to 17. (2) Mopsi locations in clustering benchmark datasets (Fränti and Sieranoja, 2018) (real-world): a set of 2-D locations for $n = 6014$ users in Joensuu. Note that the second dataset concerning human beings is exactly the situation when the data points need to be treated fairly. Both datasets are in Euclidean plane.

We consider the k -means objective as social cost, and compare our algorithm $ALG_g^+(k\text{-means})$ with $k\text{-means++}$. For each dataset, we consider a range of values of k . Figure 1c, 1e show the values of α (which is the minimum value such that the output clustering is a $(\alpha, 1)$ -core), and the ratio between the social costs of the two algorithms. Figure 1d, 1f show the values of β . In terms of fairness, $ALG_g^+(k\text{-means})$ has a significantly lower α and β than $k\text{-means++}$ in most cases (though in few cases it is slightly larger). In terms of efficiency, the social cost of $ALG_g^+(k\text{-means})$ is bounded by a small constant compared with $k\text{-means++}$, and ours is even as good as $k\text{-means++}$ on Gaussian dataset. To conclude, our algorithm ensures better core fairness for the agents than classic clustering algorithms, and meanwhile, empirically has a good efficiency as well.

7. Conclusion

In this work, we revisited the algorithmic fair centroid clustering problem, and proposed a novel definition for group

fairness – core. We demonstrated the extent to which an approximate core clustering is guaranteed to exist. There are many future directions that are worth exploring. For example, it will be interesting to improve the approximation bounds for other spaces such as d -dimensional Euclidean space for $d \geq 2$, and simple graphs with unit weight. In Supplementary Material, we prove that even when all agents consist of all vertices of a unit-weight tree, an exact core can still be empty. A systematic analysis of upper- and lower-bounds is left for future study.

Acknowledgement

The authors thank anonymous reviewers for their constructive comments. This work is supported by The Hong Kong Polytechnic University (Grant No. P0034420).

References

- Atila Abdulkadiroğlu and Tayfun Sönmez. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66, 3 (1998), 689–701.
- Noga Alon, Michal Feldman, Ariel D. Procaccia, and Moshe Tennenholtz. 2010. Strategyproof approximation of the minimax on networks. *Math. Oper. Res.* 35, 3 (2010), 513–526.
- David Arthur and Sergei Vassilvitskii. 2007. k-means++: the advantages of careful seeding. In *SODA*. SIAM, 1027–1035.
- Haris Aziz and Simon Rey. 2020. Almost group envy-free allocation of indivisible goods and chores. In *IJCAI*. ijcai.org, 39–45.
- Moshe Babaioff, Shahar Dobzinski, Sigal Oren, and Aviv Zohar. 2012. On bitcoin and red balloons. In *EC*. ACM, 56–73.
- Arturs Backurs, Piotr Indyk, Krzysztof Onak, Baruch Schieber, Ali Vakilian, and Tal Wagner. 2019. Scalable fair clustering. In *ICML (Proceedings of Machine Learning Research, Vol. 97)*. PMLR, 405–413.
- Maria-Florina Balcan, Travis Dick, Ritesh Noothigattu, and Ariel D. Procaccia. 2019. Envy-free classification. In *NeurIPS*. 1238–1248.
- Nawal Benabbou, Mithun Chakraborty, Edith Elkind, and Yair Zick. 2019. Fairness towards groups of agents in the allocation of indivisible items. In *IJCAI*. ijcai.org, 95–101.
- Suman Kalyan Bera, Deeparnab Chakraborty, Nicolas Flores, and Maryam Negahbani. 2019. Fair algorithms for clustering. In *NeurIPS*. 4955–4966.
- Ioana Oriana Bercea, Martin Groß, Samir Khuller, Aounon Kumar, Clemens Rösner, Daniel R. Schmidt, and Melanie Schmidt. 2019. On the cost of essentially fair clusterings. In *APPROX-RANDOM (LIPIcs, Vol. 145)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 18:1–18:22.
- Marcus Berliant, William Thomson, and Karl Dunz. 1992. On the fair division of a heterogeneous commodity. *Journal of Mathematical Economics* 21, 3 (1992), 201–216.
- Vladimir Braverman, Shaofeng H.-C. Jiang, Robert Krauthgamer, and Xuan Wu. 2019. Coresets for ordered weighted clustering. In *ICML (Proceedings of Machine Learning Research, Vol. 97)*. PMLR, 744–753.
- Georgios Chalkiadakis, Edith Elkind, and Michael J. Wooldridge. 2011. *Computational Aspects of Cooperative Game Theory*. Morgan & Claypool Publishers.
- Xingyu Chen, Brandon Fain, Liang Lyu, and Kamesh Munagala. 2019. Proportionally fair clustering. In *ICML (Proceedings of Machine Learning Research, Vol. 97)*. PMLR, 1032–1041.
- Yu Cheng, Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2019. Group fairness in committee selection. In *EC*. ACM, 263–279.
- Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, and Sergei Vassilvitskii. 2017. Fair clustering through fairlets. In *NIPS*. 5029–5037.
- Vincent Conitzer, Rupert Freeman, and Nisarg Shah. 2017. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*. 629–646.
- Vincent Conitzer, Rupert Freeman, Nisarg Shah, and Jennifer Wortman Vaughan. 2019. Group fairness for the allocation of indivisible goods. In *AAAI*. AAAI Press, 1853–1860.
- Xiaotie Deng and Christos H. Papadimitriou. 1994. On the complexity of cooperative solution concepts. *Math. Oper. Res.* 19, 2 (1994), 257–266.
- Brandon Fain, Kamesh Munagala, and Nisarg Shah. 2018. Fair allocation of indivisible public goods. In *EC*. ACM, 575–592.
- Allan M Feldman and Roberto Serrano. 2006. *Welfare Economics and Social Choice Theory*. Springer Science & Business Media.
- Michal Feldman and Yoav Wilf. 2013. Strategyproof facility location and the least squares objective. In *EC*. ACM, 873–890.
- Pasi Fränti and Sami Sieranoja. 2018. K-means properties on six clustering benchmark datasets. *Appl. Intell.* 48, 12 (2018), 4743–4759. <http://cs.uef.fi/sipu/datasets/>
- Elfarouk Harb and Ho Shan Lam. 2020. KFC: A scalable approximation algorithm for k -center fair clustering. In *NeurIPS*.
- Safwan Hossain, Andjela Mladenovic, and Nisarg Shah. 2020. Designing fairly Fair classifiers via economic fairness notions. In *WWW*. ACM / IW3C2, 1559–1569.
- Lingxiao Huang, Shaofeng H.-C. Jiang, and Nisheeth K. Vishnoi. 2019. Coresets for clustering with fairness constraints. In *NeurIPS*. 7587–7598.
- Kamal Jain and Vijay V. Vazirani. 1999. Primal-dual approximation algorithms for metric facility location and k -median problems. In *FOCS*. IEEE Computer Society, 2–13.

- Jon M. Kleinberg and Prabhakar Raghavan. 2005. Query incentive networks. In *FOCS*. IEEE Computer Society, 132–141.
- Maria Kyropoulou, Warut Suksompong, and Alexandros A. Voudouris. 2020. Almost envy-freeness in group resource allocation. *Theor. Comput. Sci.* 841 (2020), 110–123.
- Bo Li and Yingkai Li. 2020. Fair resource sharing and dorm assignment. In *AAMAS*. International Foundation for Autonomous Agents and Multiagent Systems, 708–716.
- Richard D McKelvey. 1986. Covering, dominance, and institution-free properties of social choice. *American Journal of Political Science* (1986), 283–314.
- Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2019. A survey on bias and fairness in machine learning. *CoRR* abs/1908.09635 (2019).
- Evi Micha and Nisarg Shah. 2020. Proportionally fair clustering revisited. In *ICALP (LIPIcs, Vol. 168)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 85:1–85:16.
- Hervé Moulin. 2003. *Fair division and collective welfare*. MIT Press.
- Clemens Rösner and Melanie Schmidt. 2018. Privacy preserving clustering with constraints. In *ICALP (LIPIcs, Vol. 107)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 96:1–96:14.
- Melanie Schmidt, Chris Schwiegelshohn, and Christian Sohler. 2018. Fair coresets and streaming algorithms for fair k -means clustering. *CoRR* abs/1812.10854 (2018).
- Erel Segal-Halevi and Warut Suksompong. 2019. Democratic fair allocation of indivisible goods. *Artif. Intell.* 277 (2019).
- Berk Ustun, Yang Liu, and David C. Parkes. 2019. Fairness without harm: decoupled classifiers with preference guarantees. In *ICML (Proceedings of Machine Learning Research, Vol. 97)*. PMLR, 6373–6382.
- Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, Krishna P. Gummadi, and Adrian Weller. 2017. From parity to preference-based notions of fairness in classification. In *NIPS*. 229–239.