

Approximate Group Fairness for Clustering¹

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

¹Department of Computing, The Hong Kong Polytechnic University

²School of Mathematical Sciences, Ocean University of China

³Warwick Business School, University of Warwick

⁴Department of Computer Science and Engineering, University of Nebraska-Lincoln

⁵Department of Computer Science, Duke University

* Corresponding Author: chenhwang4-c@my.cityu.edu.hk

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*. A 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. We also conduct experiments on synthetic and real-world data to examine the performance of our algorithms.

1 Introduction

Fairness is increasingly concerned in many areas including designing machine learning algorithms. This is partly motivated by various real-world applications where data points of the algorithms are real human beings who should be treated unbiasedly. However, many traditional algorithms fail to ensure fairness for individual point or collection of points, which may undermine normal running of the whole society. Accordingly, the past several years have seen considerable efforts in developing fair learning algorithms [16, 14, 6, 9].

Following [14], we revisit the group fairness 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 [17, 20]. We consider the following facility location problem proposed in [14, 33].

¹This work is supported by The Hong Kong Polytechnic University (Grant NO. P0034420).

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 park closest to her. 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 these suburbs are relatively 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* [34]: 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 [19, 13], social choice [31, 21] and fair division [1, 20].

Chen et al. [14] 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 [14] 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. We present such an instance in Supplementary Material.

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 [14, 33], 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 distance 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 [22], and trees can be used to model referral or query networks [28, 5]. As argued in [33], though objectives like truthfulness [2] and social welfare maximization [22] have received significant attention, fairness is largely overlooked.

	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)

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

$(1, \beta)$ -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 [14], 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.

	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)

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

$(\alpha, 1)$ -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 complement 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 $\alpha \leq \min\{k, \max\{\frac{n}{4}, \frac{k}{2}\}\}$, 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, [16] 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, [16] 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 [24, 10] and running time [26, 36, 6]. On the other hand, [9, 35] and [12] extended the setting in [16] 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 [39, 38, 7]. We refer readers to the survey [32] 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 [37, 8, 29]. The second one is to consider the fairness of arbitrarily formed groups of agents [4, 11, 18, 25]. 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 [17, 15, 20, 30]. 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, whose proof and the proof for Proposition 6 are in Supplementary Material.

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

When $k = 1$, the single location in \mathcal{M} that minimizes the total distance to all agents is in the core. When $k \geq n - 1$, a k -clustering that places as many centers as possible on the agent locations is in the core, because for any $\lceil \frac{n}{k} \rceil = 2$ agents, at least one of them is a center and thus they have no incentive to deviate. When $k \geq \frac{n}{2}$ and the metric space is a connected graph $G(V, E)$ with $V = \mathcal{X} = \mathcal{N} = \mathcal{M}$, by contrast with the result in [33] where a proportional k -clustering always exists and can be computed efficiently, in our setting a core clustering is not guaranteed.

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.*

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 exists a line instance with empty $(1, o(\sqrt{n}))$ -core.*

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 $r = \min\{\lambda k, n\}$.
 - 6: 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 $N_1 = \{i \in S \mid x_i \leq y_j\}$, $N_2 = \{i \in S \mid y_j < x_i < y_{j+1}\}$, and $N_3 = \{i \in S \mid x_i \geq y_{j+1}\}$. Note that the algorithm guarantees $|N_2| \leq \lambda - 1$. Then we have

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

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

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. We only state the following theorem and defer the detailed discussion to Supplementary Material.

Theorem 9. *For any tree instance, a $(1, O(\sqrt{n}))$ -core clustering can be computed efficiently. 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 2, 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 [27], and It has been proved [14, 33] 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 2 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  $\exists x \in X$  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:      $Y \leftarrow Y \cup \{x\}$ 
9:      $N \leftarrow N \setminus B(x, \delta)$ 
10:  end while
11: 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.

When k , the number 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 $|V| = n$.

Theorems 11 and 12 are proved in Supplementary Material.

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. Due to space limit, in this section, all results except Theorem 13 are proved in Supplementary Material.

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 $iK + 1$ and $\{i_{C+1}, \dots, i_{2C-1}\}$ are located at $iK + 2$.

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 $x \geq i_C + 1$. 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

Again, for tree spaces, we only state our main results and defer the detailed discussion to Supplementary Material.

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 $\alpha \geq 1$, 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, which is proved in Supplementary Material.

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

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 one 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,

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

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 $\beta < r = r_1 + r_2 \leq \max\{4, \frac{2}{\alpha+1} + 3\} = \beta$. \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.

To improve the performance, we refine ALG_g in Algorithm 3, 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 [3]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.




Figure 1: The clustering centers when two algorithms cluster Gaussian dataset for $k = 10$.

Algorithm 3 $ALG_g^+(\text{obj})$ 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**
-

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 for k from 8 to 17. (2) Mopsi locations in clustering benchmark datasets [23] (real-world): a

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

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.

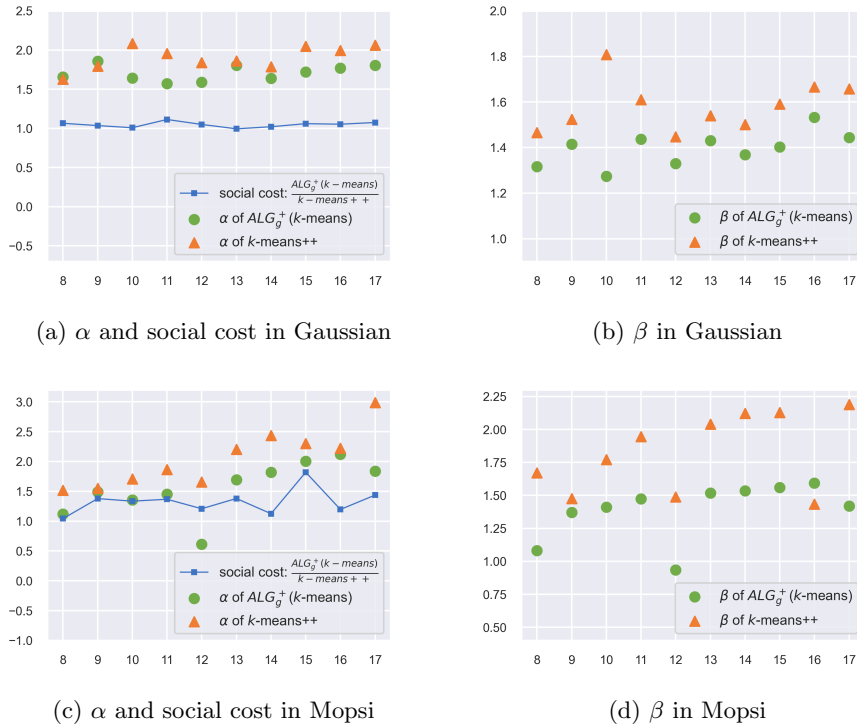


Figure 2: For a range of $k = 8, \dots, 17$ (horizontal axis), (a) and (b) (resp. (c) and (d)) compare the fairness and efficiency in Gaussian dataset (resp. Mopsi locations).

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 2a, 2c 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 2b, 2d 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. More experiments can be found in Supplementary Material. To conclude, empirically, our algorithm ensures better core fairness for the agents than classic clustering algorithms, and in the meanwhile, it has descent efficiency guarantee 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.

References

- [1] Atila Abdulkadiroğlu and Tayfun Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.
- [2] Noga Alon, Michal Feldman, Ariel D. Procaccia, and Moshe Tennenholtz. Strategyproof approximation of the minimax on networks. *Math. Oper. Res.*, 35(3):513–526, 2010.
- [3] David Arthur and Sergei Vassilvitskii. k-means++: the advantages of careful seeding. In *SODA*, pages 1027–1035. SIAM, 2007.
- [4] Haris Aziz and Simon Rey. Almost group envy-free allocation of indivisible goods and chores. In *IJCAI*, pages 39–45. ijcai.org, 2020.
- [5] Moshe Babaioff, Shahar Dobzinski, Sigal Oren, and Aviv Zohar. On bitcoin and red balloons. In *EC*, pages 56–73. ACM, 2012.
- [6] Arturs Backurs, Piotr Indyk, Krzysztof Onak, Baruch Schieber, Ali Vakilian, and Tal Wagner. Scalable fair clustering. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 405–413. PMLR, 2019.
- [7] Maria-Florina Balcan, Travis Dick, Ritesh Noothigattu, and Ariel D. Procaccia. Envy-free classification. In *NeurIPS*, pages 1238–1248, 2019.
- [8] Nawal Benabbou, Mithun Chakraborty, Edith Elkind, and Yair Zick. Fairness towards groups of agents in the allocation of indivisible items. In *IJCAI*, pages 95–101. ijcai.org, 2019.
- [9] Suman Kalyan Bera, Deeparnab Chakrabarty, Nicolas Flores, and Maryam Negahbani. Fair algorithms for clustering. In *NeurIPS*, pages 4955–4966, 2019.
- [10] Ioana Oriana Bercea, Martin Groß, Samir Khuller, Aounon Kumar, Clemens Rösner, Daniel R. Schmidt, and Melanie Schmidt. On the cost of essentially fair clusterings. In *APPROX-RANDOM*, volume 145 of *LIPICs*, pages 18:1–18:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [11] Marcus Berliant, William Thomson, and Karl Dunz. On the fair division of a heterogeneous commodity. *Journal of Mathematical Economics*, 21(3):201–216, 1992.
- [12] Vladimir Braverman, Shaofeng H.-C. Jiang, Robert Krauthgamer, and Xuan Wu. Core-sets for ordered weighted clustering. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 744–753. PMLR, 2019.
- [13] Georgios Chalkiadakis, Edith Elkind, and Michael J. Wooldridge. *Computational Aspects of Cooperative Game Theory*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2011.
- [14] Xingyu Chen, Brandon Fain, Liang Lyu, and Kamesh Munagala. Proportionally fair clustering. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 1032–1041. PMLR, 2019.
- [15] Yu Cheng, Zhihao Jiang, Kamesh Munagala, and Kangning Wang. Group fairness in committee selection. In *EC*, pages 263–279. ACM, 2019.
- [16] Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, and Sergei Vassilvitskii. Fair clustering through fairlets. In *NIPS*, pages 5029–5037, 2017.

- [17] Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 629–646, 2017.
- [18] Vincent Conitzer, Rupert Freeman, Nisarg Shah, and Jennifer Wortman Vaughan. Group fairness for the allocation of indivisible goods. In *AAAI*, pages 1853–1860. AAAI Press, 2019.
- [19] Xiaotie Deng and Christos H. Papadimitriou. On the complexity of cooperative solution concepts. *Math. Oper. Res.*, 19(2):257–266, 1994.
- [20] Brandon Fain, Kamesh Munagala, and Nisarg Shah. Fair allocation of indivisible public goods. In *EC*, pages 575–592. ACM, 2018.
- [21] Allan M Feldman and Roberto Serrano. *Welfare Economics and Social Choice Theory*. Springer Science & Business Media, 2006.
- [22] Michal Feldman and Yoav Wilf. Strategyproof facility location and the least squares objective. In *EC*, pages 873–890. ACM, 2013.
- [23] Pasi Fränti and Sami Sieranoja. K-means properties on six clustering benchmark datasets. *Appl. Intell.*, 48(12):4743–4759, 2018.
- [24] Elfarouk Harb and Ho Shan Lam. KFC: A scalable approximation algorithm for k -center fair clustering. In *NeurIPS*, 2020.
- [25] Safwan Hossain, Andjela Mladenovic, and Nisarg Shah. Designing fairly fair classifiers via economic fairness notions. In *WWW*, pages 1559–1569. ACM / IW3C2, 2020.
- [26] Lingxiao Huang, Shaofeng H.-C. Jiang, and Nisheeth K. Vishnoi. Coresets for clustering with fairness constraints. In *NeurIPS*, pages 7587–7598, 2019.
- [27] Kamal Jain and Vijay V. Vazirani. Primal-dual approximation algorithms for metric facility location and k -median problems. In *FOCS*, pages 2–13. IEEE Computer Society, 1999.
- [28] Jon M. Kleinberg and Prabhakar Raghavan. Query incentive networks. In *FOCS*, pages 132–141. IEEE Computer Society, 2005.
- [29] Maria Kyropoulou, Warut Suksompong, and Alexandros A. Voudouris. Almost envy-freeness in group resource allocation. *Theor. Comput. Sci.*, 841:110–123, 2020.
- [30] Bo Li and Yingkai Li. Fair resource sharing and dorm assignment. In *AAMAS*, pages 708–716. International Foundation for Autonomous Agents and Multiagent Systems, 2020.
- [31] Richard D McKelvey. Covering, dominance, and institution-free properties of social choice. *American Journal of Political Science*, pages 283–314, 1986.
- [32] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *CoRR*, abs/1908.09635, 2019.
- [33] Evi Micha and Nisarg Shah. Proportionally fair clustering revisited. In *ICALP*, volume 168 of *LIPICs*, pages 85:1–85:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

- [34] Hervé Moulin. *Fair division and collective welfare*. MIT Press, 2003.
- [35] Clemens Rösner and Melanie Schmidt. Privacy preserving clustering with constraints. In *ICALP*, volume 107 of *LIPICs*, pages 96:1–96:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- [36] Melanie Schmidt, Chris Schwiegelshohn, and Christian Sohler. Fair coresets and streaming algorithms for fair k -means clustering. *CoRR*, abs/1812.10854, 2018.
- [37] Erel Segal-Halevi and Warut Suksompong. Democratic fair allocation of indivisible goods. *Artif. Intell.*, 277, 2019.
- [38] Berk Ustun, Yang Liu, and David C. Parkes. Fairness without harm: decoupled classifiers with preference guarantees. In *ICML*, volume 97 of *Proceedings of Machine Learning Research*, pages 6373–6382. PMLR, 2019.
- [39] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, Krishna P. Gummadi, and Adrian Weller. From parity to preference-based notions of fairness in classification. In *NIPS*, pages 229–239, 2017.

Supplementary Material

The Supplementary Material is organized as follows. In Section A, we present an instance to show that the classic clustering algorithm can be arbitrarily bad with respect to the core fairness. In Section B, we give a lower bound to show that a $(1, \frac{14}{13} - \epsilon)$ -core or a $(\frac{28}{25} - \epsilon, 1)$ -core is not guaranteed to exist for any $\epsilon > 0$, even the space is an unweighted tree and each vertex lies exactly one agent. In Section C, we present all omitted proofs in the main part of this work. In Section D, we show all experiment results.

A Classic Algorithms Can Be Arbitrarily Unfair

Regarding core fairness, we note that the traditional learning algorithms (e.g., Lloyd-style algorithms for k -means and k -medians) can be arbitrarily unfair in the worst case. Take k -medians as an example. Recall that k -medians algorithm returns the cluster set $Y \in [\mathcal{M}]^k$ minimizing the total distance. Consider the following special setting when $\mathcal{X} = \mathcal{M} = \mathbb{R}$ and $k = 1$. For this case, k -medians algorithm degenerates to *selecting the median point of* \mathcal{N} . That is, given n locations of agents with $a_1 \leq a_2 \leq \dots \leq a_n$, the algorithm outputs $Y = \{a_{\lceil \frac{n}{2} \rceil}\}$. Similarly, when k is arbitrary, k -medians algorithm partitions \mathcal{N} into k subsets and within each subset, the median point of it is selected as the center.

It is not hard to verify that when $k = 1$, k -medians algorithm outputs a core clustering. However, when $k \geq 2$, k -medians algorithm can be arbitrarily unfair. Consider k groups (with different sizes) of points with a total size n , where each group is far from the others. Then k -medians algorithm will set one cluster center for each group (at its median point) to minimize the total distance. Assume there is a group of $2 \lceil \frac{n}{k} \rceil + 1$ points, where $\lceil \frac{n}{k} \rceil$ of these points are located at 0, a single point is at 1, and the remaining $\lceil \frac{n}{k} \rceil$ are at 2. Accordingly, k -medians algorithm puts the cluster center for this group of points at 1 to minimize the total 1-norm distance. However, all the $\lceil \frac{n}{k} \rceil$ points located at 0 can deviate to a new cluster center at 0 that decreases their total distance to 0; and thus

$$\beta = \frac{\lceil \frac{n}{k} \rceil}{0} = \infty,$$

which means k -medians algorithm returns an arbitrarily unfair clustering in the β -dimension relaxation.

B Lower Bound for Simple Trees

Theorem 20. *There exists an instance in an unweighted tree such that $(1, \beta)$ -core and $(\alpha, 1)$ -core for all $\beta < \frac{14}{13}$ and $\alpha < \frac{28}{25}$ are both empty, even if each vertex lies exactly one agent.*

Proof. Consider the following instance constructed on the tree with unit weights. As shown in Figure 3, there are $n = 50$ vertices distributed on 7 branches and at each vertex lies an agent. Let $k = 7$ be the number of cluster centers.

We first prove that $(1, \beta)$ -core is empty for all $\beta < \frac{14}{13}$. For clarity, in the figure we number the vertices and group them into 3 types (sets): A , B and C , each of which contains 7 parts of vertices with identical structure on each branch denoted by $A = \bigcup_{i=1}^7 A_i$, $B = \bigcup_{i=1}^7 B_i$ and $C = \bigcup_{i=1}^7 C_i$. Regarding deviation, by Lemma 5, we only need to consider a potential blocking coalition with size $\lceil \frac{n}{k} \rceil = 8$.

Case 1: There exists at least one part of vertices in A without a center located within. W.l.o.g., suppose vertices in A_1 have no center located on them. Considering the group of

agents $S = \{v_1, v_2, \dots, v_8\}$, their total distance from the center assigned to them in this case is at least $\sum_{i \in S} d(i, v_1) = 25$. However, they can form a blocking coalition to deviate together to a new center v_5 , which can decrease their total distance to the lowest(optimal): $\sum_{i \in S} d(i, v_5) = 13$. Therefore, for $\beta < \frac{25}{13}$, all the clustering solutions in Case 1 are not a $(1, \beta)$ -core clustering.

Case 2: There is one center in every part of A type vertices. We further discuss two subcases.

Subcase 2.1: There is exact one center in every part of B type vertices. Considering a potential blocking coalition of 8 agents $S = \{v_1, v_2, v_9, v_{16}, v_{23}, v_{30}, v_{37}, v_{44}\}$ located at the center part of this tree, their total distance from the center assigned to them is at least $d(v_2, v_4) \cdot 7 + d(v_1, v_4) = 17$. They can deviate to a new center v_1 , which can decrease their total distance to $\sum_{i \in S} d(i, v_1) = 7$. Therefore, for $\beta < \frac{17}{7}$, all the clustering solutions in Subcase 2.1 are not a $(1, \beta)$ -core clustering.

Subcase 2.2: There exists one part of B without centers located in it, that is, there exists one part of C having a center located in it. W.l.o.g., suppose C_1 has a center located within while B_1 does not have one. Consider the group of agents $S = \{v_1, v_2, \dots, v_8\}$. Their total distance from the centers assigned to them is at least $\sum_{i \in S} \min\{d(i, v_3), d(i, v_9)\} = 14$. However, they can form a blocking coalition to deviate together to a new center v_5 , which can decrease their total distance to $\sum_{i \in S} d(i, v_5) = 13$. Hence, for $\beta < \frac{14}{13}$, all the clustering solutions in Subcase 2.2 are not a $(1, \beta)$ -core clustering.

Therefore, in this instance, for all $\beta < \frac{14}{13}$, $(1, \beta)$ -core is empty.

When $\beta = 1$ and $\alpha < \frac{56}{50} = \frac{28}{25}$, for the above instance there always exists a blocking coalition of size $\lceil \frac{\alpha n}{k} \rceil = 8$. Therefore, $(\alpha, 1)$ -core is empty for all $\alpha < \frac{28}{25}$. \square




Figure 3: A lower bound for simple trees

C Missing Proofs

C.1 Proof of Lemma 5

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

C.2 Proof of Proposition 6

Proof. We consider the graph space induced by a complete graph $G = (V, E)$ with n vertices, where at each vertex lies an agent, and $V = \mathcal{X} = \mathcal{M}$. 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 be 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

REMARK. The second claim in Proposition 6 can be stated more generally. For every $k \in [2, n - 2]$, the existence of a $(1, \beta)$ -core clustering with $\beta < \frac{\lceil \frac{n}{k} \rceil}{\lceil \frac{n}{k} \rceil - 1} = 1 + \frac{1}{\lceil \frac{n}{k} \rceil - 1}$ is not guaranteed. Considering the graph space induced by a complete graph defined in the above proof, for any k -clustering Y , there is a group of $\lceil \frac{n}{k} \rceil$ vertices having no center located at their locations. They form a blocking coalition (w.r.t. $(1, \beta)$ -core) to deviate to a new center located at any one of them, such that their total distance to new center is $\lceil \frac{n}{k} \rceil - 1$, while that to Y is $\lceil \frac{n}{k} \rceil$.

C.3 $(1, \beta)$ -Core in Tree and Proof of Theorem 9

We consider $(1, \beta)$ -core in tree spaces. For the proportionality fairness, Micha and Shah [33] 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 $|\text{ST}(x)|$ be the number of agents contained in the subtree $\text{ST}(x)$. 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 $\text{ALG}_t(\lambda)$ with a tuneable parameter λ , which controls the threshold value for locating a center. When $\lambda = \lceil \frac{n}{k} \rceil$, $\text{ALG}_t(\lambda)$ is equivalent to PFCT.

Algorithm 4 $ALG_t(\lambda)$ 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: $G^{\ell-1} \leftarrow G^\ell$
 - 6: **for** every $x \in V$ with level $(x) = \ell$ and $|\text{ST}(x)| \geq \lambda$ **do**
 - 7: **if** $|Y| < k$ **then**
 - 8: $Y \leftarrow Y \cup \{x\}$
 - 9: $G^{\ell-1} \leftarrow G^{\ell-1} \setminus \text{ST}(x)$
 - 10: **end if**
 - 11: **end for**
 - 12: **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 $\lambda - 1$. Actually, $ALG_t(\lambda)$ is an extension of $ALG_l(\lambda)$ in the tree. They share the similar proofs for the upper-bound.

Theorem (Restatement of 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. In addition, there exists an instance in the tree with $k = \frac{n}{k+1}$, such that the $(1, o(\sqrt{n}))$ -core is empty.*

Proof. Define $\lambda = \lceil \frac{n}{k} \rceil$ if $k = \Omega(\sqrt{n})$, and $\lambda = \lceil \frac{n}{k+1} \rceil$ otherwise. The output of the algorithm $ALG_t(\lambda)$ is location profile $Y = (y_1, \dots, y_k)$. Suppose for contradiction that there is a blocking coalition $S \subseteq \mathcal{N}$ of agents with $|S| \geq \lceil \frac{n}{k} \rceil$ and a deviating center $y' \in V \setminus Y$. Let $d(u, v)$ denote the distance between nodes u and v , i.e., the length of the unique path from u to v . Define a partition (N_1, N_2) of group S (as Figure 4 shows): N_2 consists of the agents in S who are in the component containing y' , where the components are obtained by removing all centers in Y from the tree; N_1 consists of the remaining group members. Then we have

$$\sum_{i \in S} d(i, Y) \leq \sum_{i \in N_1} d(i, Y) + \sum_{i \in N_2} (d(i, y') + d(y', Y)), \quad (3)$$

$$\sum_{i \in S} d(i, y') \geq \sum_{i \in N_1} (d(i, Y) + d(Y, y')) + \sum_{i \in N_2} d(i, y'). \quad (4)$$

Because S is a blocking coalition, we have $r := \frac{\sum_{i \in S} d(i, Y)}{\sum_{i \in S} d(i, y')} > 1$. It follows from (3) and (4) that

$$r \leq \frac{\sum_{i \in N_2} d(y', Y)}{\sum_{i \in N_1} d(y', Y)} = \frac{|N_2|}{|N_1|}. \quad (5)$$

When $k = \Omega(\sqrt{n})$ and $\lambda = \lceil \frac{n}{k} \rceil$, the algorithm guarantees that $|N_2| \leq \lceil \frac{n}{k} \rceil - 1$, and thus $|N_1| = |S| - |N_2| \geq 1$, implying that $r \leq \lceil \frac{n}{k} \rceil - 1$ and S cannot be a blocking coalition w.r.t. $(1, \lceil \frac{n}{k} \rceil - 1)$ -core. Therefore, Y is a $(1, \lceil \frac{n}{k} \rceil - 1)$ -core clustering.

When $k = o(\sqrt{n})$ and $\lambda = \lceil \frac{n}{k+1} \rceil$, the algorithm guarantees that $|N_2| \leq \lceil \frac{n}{k+1} \rceil - 1$, and thus $|N_1| = |S| - |N_2| \geq \lceil \frac{n}{k} \rceil - \lceil \frac{n}{k+1} \rceil + 1$. A simple computation shows that $r \leq k$, indicating that S cannot be a blocking coalition. So Y is a $(1, k)$ -core clustering.




Figure 4: The partition (N_1, N_2) of group S , where solid points indicate centers in Y .

For the lower bound, we note that the instance constructed in the line in the proof of Theorem 7 can be adapted to the graph space induced by a path graph, where $V = \{1, 2, \dots, k\}$ is the vertex set of the path, and each vertex accommodates $k+1$ agents. Using the similar analysis, it can be shown that the $(1, o(\sqrt{n}))$ -core for this path is empty. \square

C.4 Proof of Theorem 10

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

$$\sum_{i \in S} d(i, Y) \leq \sum_{i \in S} d(i, y^*) \leq \sum_{i \in S} d(i, y') + \sum_{i \in S} d(y^*, y'). \quad (6)$$

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 (6), it follows that

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

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

C.5 Proof of Theorem 11

Proof. Construct an n -vertex complete graph by setting the vertices to be the locations of n agents, and the weight of each edge to be the distance of the two endpoints. Consider

the following algorithm: run Prim's algorithm to find a minimum spanning tree (MST) over the complete graph, and then find a vertex cover of size k for the MST, which can be done efficiently by selecting every second vertex. The algorithm locates k center on the locations in this vertex cover. Note that every two adjacent agents in the MST have at least one center located on themselves. Now we show this solution (denoted by Y) is a $(1, 2)$ -core clustering. We refer to agents and vertices interchangeably. The time complexity comes from Prim's algorithm.

First, by Lemma 5, we only need to consider groups consisting of 2 agents, denoted by $\{u, v\}$. Second, if one group member has a center on his/her location, then it cannot be a blocking coalition, because the group's total distance to any new center is at least $d(u, v)$ by the metric, while the total distance to Y is at most $d(u, v)$. So we only need to consider group $\{u, v\}$ without center on u 's nor v 's location. By the definition of vertex cover, edge (u, v) is not in the MST. Recall that Prim's algorithm operates by building the tree one vertex at a time, from an arbitrary starting vertex (say r), at each step adding the cheapest possible connection from the tree to another vertex. Assume u is added earlier than v , and the vertex set at the time of adding u is W . Assume v is added via an edge (y, v) . By Prim's algorithm, it must be $d(u, v) \geq d(y, v)$, otherwise v is added by edge (u, v) , a contradiction. We discuss the following two cases.

(a) If the unique r - v path in the MST passes through u , denote by $(r, \dots, u, z, \dots, v)$. It is easy to see that $d(u, z) \leq d(u, v)$, otherwise edge (u, v) is added into the MST. Thus we have $2d(u, v) \geq d(y, v) + d(u, z)$. Note that both agents y and z have facilities on their locations. The total distance of group $\{u, v\}$ to Y is at most $d(y, v) + d(u, z)$, and that to any deviating center is at least $d(u, v)$. It indicates that $\{u, v\}$ cannot form a blocking coalition w.r.t. $\beta = 2$. Hence, Y is a $(1, 2)$ -core clustering.

(b) If the unique r - v path in the MST does not pass through u , denote by $(r, \dots, w, z, \dots, v)$, and assume $w \in W$, $z \notin W$. It is easy to see that $d(u, W) \leq d(z, W)$, otherwise z is added into W , a contradiction. Because z is added earlier than v , we have $d(z, W) \leq d(w, z) \leq d(u, v)$. Thus we have $2d(u, v) \geq d(y, v) + d(z, W) \geq d(y, v) + d(u, W)$. Because both the agent y and the agent incident to u in W have centers on their locations, the total distance of group $\{u, v\}$ to Y is at most $d(y, v) + d(u, W)$, while that to any deviating center is at least $d(u, v)$. So $\{u, v\}$ cannot form a blocking coalition w.r.t. $\beta = 2$, and Y is a $(1, 2)$ -core clustering. \square

C.6 Proof of Theorem 12

Proof. We give a reduction from the minimum vertex cover problem (MVC). Let $(G' = (V', E'), k')$ with $|V'| = n'$ be an arbitrary instance of the MVC, which asks whether a vertex cover of G' with size k' exists. We construct an instance I of our clustering problem with $n = 4n'$ agents and $k = 2n' + k'$ centers as follows. In instance I , the space is induced by a graph $G = (V, E)$ with $n' + 1$ connected components, one being graph G' and other n' components being triangles K_3 . We assume w.l.o.g. that every two components are connected by an edge with infinite weight, and all other edges have unit weight. Define the distance between any two vertices to be the length of the shortest path between them. At each vertex of G lies an agent, so there are $4n'$ agents in total. We show that the answer of (G', k') is "yes" if and only if I has a $(1, 2 - \epsilon)$ -core clustering. By Lemma 5, we only need to consider a group of agents of size $\lceil \frac{n}{k} \rceil = 2$ as a possible blocking coalition.

If G' has a vertex cover $W \subseteq V'$ of size k' , then we have a k -clustering of G (denoted by $Y \subseteq V$) which locates k' centers on vertices in W , and locates two centers on vertices of each triangle. Note that no possible blocking coalition can have two agents in different components. Clearly, any two agents in a triangle cannot form a blocking coalition, as their distance to Y is at most 1. Thus it suffices to consider a group $S \subseteq V'$ consisting of two

agents in G' . If no member of S is in the vertex cover W , then there is no edge between the two members, and the total distance of S to any new center is at least 2. Since the total distance to Y is 2, group S cannot form a blocking coalition. If there is some member in W , then the total distance to Y is at most 1, implying that no deviation can occur. Hence, there is no blocking coalition, indicating that our solution is a core clustering (and thus a $(1, 2 - \epsilon)$ -core clustering).

If I has a $(1, 2 - \epsilon)$ -core clustering Y , it must locate at least two centers on each triangle (otherwise two agents in a triangle form a blocking coalition), and thus locate at most k' centers on G' . For any two adjacent agents $u, v \in V'$ in G' , if neither of them has a center located on themselves, they can deviate to a new center u such that

$$d(u, u) + d(v, u) = 1 < \frac{2}{(2 - \epsilon)} \leq \frac{d(u, Y) + d(v, Y)}{2 - \epsilon},$$

implying that they can form a blocking coalition. Therefore, there is at least one center between any two adjacent agents in G' , and thus the center locations in G' is a vertex cover. \square

C.7 Proof of Theorem 14

Proof. For any instance in the line, we show that $ALG_l(\lceil \frac{n}{k} \rceil)$ outputs a $(2, 1)$ -core clustering. Let $\mathbf{x} = \{x_1, \dots, x_n\}$ be the locations of agents. Recall that the algorithm $ALG_l(\lceil \frac{n}{k} \rceil)$ locates a center at $y_i = x_{\lceil \frac{n}{k} \rceil \cdot i}$ for $i = 1, \dots, k - 1$, and $y_k = x_n$. The output is centers $Y = \{y_1, \dots, y_k\}$. Now we prove that Y is a $(2, 1)$ -core clustering.

Consider an arbitrary group $S \subseteq \mathcal{N}$ of agents with $|S| \geq \frac{2n}{k}$, and an arbitrary point $y' \in \mathbb{R} \setminus Y$. 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 $j = 0, \dots, k$.

Let $N_1 = \{i \in S | y_j < x_i < y_{j+1}\}$ be the set of members who are located in the interval (y_j, y_{j+1}) , and $N_2 = S \setminus N_1$ be the set of members who are outside. By the algorithm, we have $|N_1| \leq \lceil \frac{n}{k} \rceil - 1$, and $|N_2| = |S| - |N_1| \geq |N_1|$. We construct pairs

$$\{(v_1, u_1), (v_2, u_2), \dots, (v_{|N_1|}, u_{|N_1|})\}$$

between N_1 and N_2 , so that $v_p \in N_1, u_p \in N_2$ and $v_p \neq v_q, u_p \neq u_q$, for any different $p, q \leq |N_1|$. Then it is easy to see that for any pair of agents, their total distance to the potential deviating center y' cannot be smaller than that to the solution Y , i.e.,

$$d(v_p, y') + d(u_p, y') \geq d(v_p, Y) + d(u_p, Y)$$

for any $p \leq |N_1|$. Further, for those agents in N_2 but not in any pair (if any), they also have a smaller distance from solution Y than that from y' . Therefore, S cannot be a blocking coalition deviating to y' . By the arbitrariness of S and y' , Y is a $(2, 1)$ -core clustering. \square

C.8 Proof of Theorem 15

Proof. For any instance in the tree, we show that $ALG_t(\lceil \frac{n}{k} \rceil)$ outputs a $(2, 1)$ -core clustering. Let Y be the solution output by $ALG_t(\lceil \frac{n}{k} \rceil)$. Suppose for contradiction that there is a blocking coalition $S \subseteq \mathcal{N}$ with $|S| \geq \frac{2n}{k}$ and a deviating center y' . Define a partition (N_1, N_2) of group S : N_2 consists of the agents in S who are in the component containing y' , where the components are obtained by removing all centers in Y from the tree; N_1 consists of the remaining group members.

Recall that $ALG_t(\lceil \frac{n}{k} \rceil)$ guarantees

$$|N_2| \leq \lceil \frac{n}{k} \rceil - 1$$

and thus

$$|N_1| \geq |S| - (\lceil \frac{n}{k} \rceil - 1) \geq \lceil \frac{n}{k} \rceil + 1.$$

By Equation (5), we have

$$\frac{\sum_{i \in S} d(i, Y)}{\sum_{i \in S} d(i, y')} \leq \frac{|N_2|}{|N_1|} \leq \frac{\lceil \frac{n}{k} \rceil - 1}{\lceil \frac{n}{k} \rceil + 1} < 1.$$

This contradicts the definition of a blocking coalition.

For the lower bound, we note that the instance constructed in the line in the proof of Theorem 13 can be adapted to the graph space induced by a path graph. Using the similar analysis, it can be shown that the $(2 - \epsilon, 1)$ -core for this path is empty. \square

C.9 Proof of Theorem 16

Proof. Consider the graph space induced by a complete graph $K_n = (V, E)$ with $n = 2k$ vertices, where at each vertex lies an agent, and the distance between any two vertices is 1. We show that, for $\alpha \leq \frac{k}{2} = \frac{n}{4}$, the $(\alpha, 1)$ -core is empty. Let $Y \subseteq V$ with $|Y| = k$ be an arbitrary k -clustering. Consider the group $V \setminus Y$ of agents with size $|V \setminus Y| = k \geq \alpha \cdot \frac{n}{k}$. The total distance of this group to Y is k , as each member has a distance 1 to Y . However, they may deviate to a new center $v \in V \setminus Y$, such that the total distance becomes $\sum_{i \in V \setminus Y} d(i, v) = k - 1 < k$. So they form a blocking coalition. This completes the proof. \square

C.10 Proof of Theorem 17

Proof. We give a reduction from the minimum vertex cover problem. Let $(G' = (V', E'), k')$ be an arbitrary instance of the MVC with $|V'| = n' \geq 2k'$ vertices, which asks whether a vertex cover of G' with size k' exists. Define an integer $x := n' \lceil \alpha \rceil + 2 \lceil 2\alpha \rceil - 2 - 2k'$.

We construct an instance I of our problem on an edge-weighted connected graph $G = (V, E)$ as follows, with $|V| = n$ and $V = \mathcal{X} = \mathcal{M}$. G is constructed based on G' : connect each vertex $v \in V'$ with $\lceil \alpha \rceil - 1$ new vertices by unit-weighted edges (denote this graph by G''), and connect an arbitrary vertex of G'' with an x -vertex complete graph K_x by a large-weighted edge, as shown in Figure 5. There are $n := |V| = \lceil \alpha \rceil n' + x$ vertices, at each of which lies an agent, and $k := k' + x - \lceil 2\alpha \rceil + 1$ centers are to be located. All edges in G have unit weight except the one connecting G'' and K_x with a large weight, which guarantees that no center can serve agents in G'' and K_x simultaneously.




Figure 5: An illustration for G , where G'' is obtained by adding $\lceil \alpha \rceil - 1$ new vertices (denoted by blue nodes) for each vertex in G' , and G'' and K_x are connected by a large-weighted edge.

Note that $\frac{\alpha \cdot n}{k} = 2\alpha$. By Lemma 5, it suffices to consider the groups of size $\lceil 2\alpha \rceil$. We show that, G' admits a vertex cover $D \subseteq V'$ of size k' , if and only if G has an $(\alpha, 1)$ -core clustering.

If G' has a vertex cover D of size k' , we construct a k -clustering Y on G , which locates k' centers on D , and locates $k - k' = x - \lceil 2\alpha \rceil + 1$ centers on different vertices of K_x . Note that no possible blocking coalition can have agents in both K_x and G'' . For any group S of size $\lceil 2\alpha \rceil$ in K_x , at least one member has a center, and thus the total distance to Y is at most $\lceil 2\alpha \rceil - 1$. However, the total distance to any deviating center is at least $\lceil 2\alpha \rceil - 1$, implying that S cannot be a blocking coalition. On the other hand, for any group S of size $\lceil 2\alpha \rceil$ in G'' , if it is a blocking coalition, there must exist a deviating center $y' \in V' \setminus D$ (because a new vertex is a deviating center only if its adjacent vertex in V' is a deviating center). Note that by the definition of vertex cover, all neighbors of y' in V' must have a center in Y . Denote the set of agents who are located at the $\lceil \alpha \rceil - 1$ new vertices adjacent to y' or y' itself by $S' \subset S$. By deviating to y' , each agent in S' decreases his/her cost by 1, while each agent in $S \setminus S'$ increases his/her cost by 1. Since $|S'| \leq \lceil \alpha \rceil$, the total cost of S cannot be improved by deviating. So Y is in the $(\alpha, 1)$ -core.

If G' has no vertex cover of size k' , suppose for contradiction that I admits an $(\alpha, 1)$ -core clustering Y . First, it must locate at least $k - k' = x - \lceil 2\alpha \rceil + 1$ centers on K_x , as otherwise those $\lceil 2\alpha \rceil$ agents without center can form a blocking coalition. So it locates at most k' centers on G'' . Because G' has no vertex cover of size k' , there must be two vertices $u, v \in V'$ so that there is no center among u, v and the new vertices adjacent to them. These $2 + 2(\lceil \alpha \rceil - 1) = 2\lceil \alpha \rceil \geq \lceil 2\alpha \rceil$ vertices can form a blocking coalition by deviating to u , giving a contradiction. \square

C.11 Proof of Theorem 18

Proof. We first prove for lines. Let $\mathbf{x} = \{x_1, \dots, x_n\}$ be the locations of agents. Consider the algorithm $ALG_t(\lceil \frac{n}{k} \rceil)$: locate a center at $y_i = x_{\lceil \frac{n}{k} \rceil, i}$ for $i = 1, \dots, k - 1$, and locate a center at $y_k = x_n$. The output is $Y = \{y_1, \dots, y_k\}$. We show that Y is in (α, β) -core for any $\alpha > 1$ and $\beta = \max\{1, \frac{1}{\alpha - 1}\}$.

Suppose for contradiction that there is a blocking coalition $S \subseteq \mathcal{N}$ with $|S| \geq \alpha \cdot \frac{n}{k}$ and a deviating center $y' \in \mathbb{R} \setminus Y$. 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 $|N_2| \leq \lceil \frac{n}{k} \rceil - 1$.

Because S is a blocking coalition, we have $r := \frac{\sum_{i \in S} d(i, Y)}{\sum_{i \in S} d(i, y')} > \beta \geq 1$. However, as the proof of Theorem 8, we have

$$r \leq \frac{|N_2|}{|N_1 \cup N_3|} \leq \frac{\lceil \frac{n}{k} \rceil - 1}{\frac{\alpha n}{k} - \lceil \frac{n}{k} \rceil + 1} \leq \beta$$

where the second equation is because $|N_2| \leq \lceil \frac{n}{k} \rceil - 1$, and $|N_1 \cup N_3| = |S| - |N_2| \geq \frac{\alpha n}{k} - \lceil \frac{n}{k} \rceil + 1$. This is a contradiction.

For any instance in the tree, we show that $ALG_t(\lceil \frac{n}{k} \rceil)$ outputs a (α, β) -core clustering for any $\alpha > 1$. Let Y be the solution output by $ALG_t(\lceil \frac{n}{k} \rceil)$. Suppose for contradiction that there is a blocking coalition $S \subseteq \mathcal{N}$ with $|S| \geq \frac{\alpha n}{k}$ and a deviating center y' . Define a partition (N_1, N_2) of group S : N_2 consists of the agents in S who are in the component containing y' , where the components are obtained by removing all centers in Y from the tree; N_1 consists of the remaining group members.

Recall that $ALG_t(\lceil \frac{n}{k} \rceil)$ guarantees $|N_2| \leq \lceil \frac{n}{k} \rceil - 1$ and thus $|N_1| \geq |S| - (\lceil \frac{n}{k} \rceil - 1) \geq \frac{\alpha n}{k} - \lceil \frac{n}{k} \rceil + 1$. By Equation (5), we have

$$\frac{\sum_{i \in S} d(i, Y)}{\sum_{i \in S} d(i, y')} \leq \frac{|N_2|}{|N_1|} \leq \frac{\lceil \frac{n}{k} \rceil - 1}{\frac{\alpha n}{k} - \lceil \frac{n}{k} \rceil + 1} \leq \beta.$$

This contradicts the definition of a blocking coalition. \square

D Experiment Results

Besides the experiments given in Section 6, we also investigate the k -medians objective empirically. We compared Lloyd’s algorithm for k -medians and our Algorithm 3 with k -medians objective in Line 8. The data sets consist of the 3-Gaussian dataset and real-world Mopsi locations in Joensuu (see 6), and S-sets in clustering benchmark datasets [23].

Gaussian dataset. We implement experiments with the k -medians objective for a range $k = 8, \dots, 17$, as shown in Figure 6. Our algorithm is clearly much fairer than Lloyd’s algorithm, in both α -dimension and β -dimension. Moreover, for the social cost (i.e., the sum of the distances from each agent to the nearest center), our algorithm is slightly better than Lloyd’s algorithm, because our algorithm is two-stage which takes both fairness and efficiency into consideration.

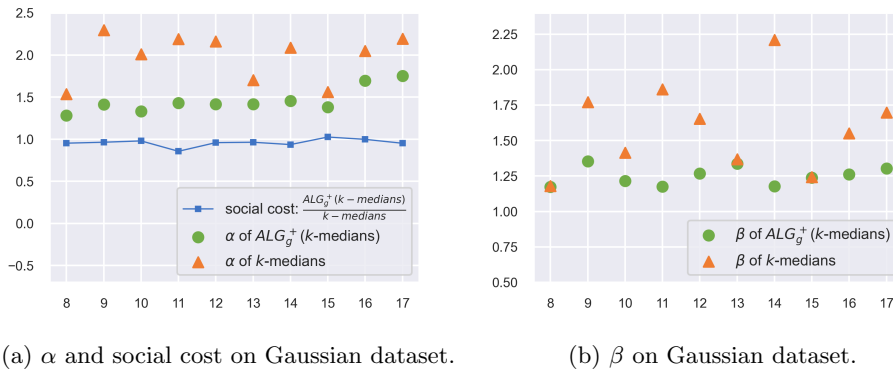
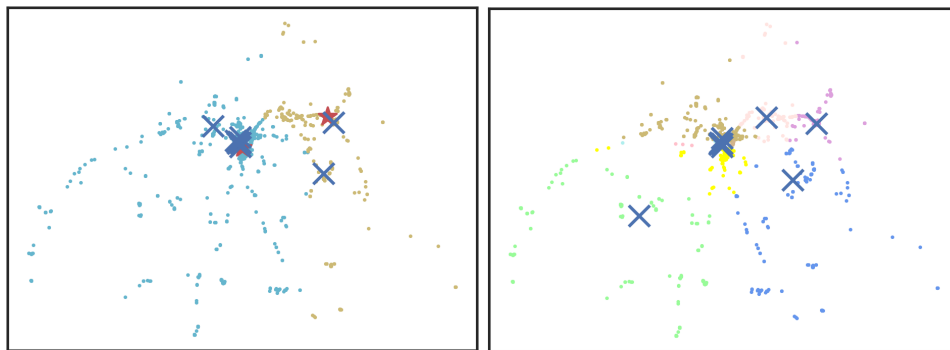


Figure 6: The comparison between Algorithm 3 and Lloyd’s algorithm with k -medians objective on Gaussian dataset.

Mopsi locations. For the dataset of Mopsi locations in Joensuu, we consider a range of $k = 8, \dots, 17$. For every such k , the clustering output by either algorithm is exactly fair, i.e., $\alpha = 1, \beta = 1$. For example, we show the results with $k = 10$ in Figure 7, where Figure 7a depicts the clustering of Algorithm 3, and Figure 7b depicts the clustering of Lloyd’s algorithm. It is easy to observe that, in both clusterings, at least 6 centers are built nearly to serve a large group of agents, which helps guarantee the core fairness. The social cost ratio in this case is 1.38.

However, we notice that the algorithms with the k -means objective cannot obtain the exact core fairness, as shown in Figure 8. The clusterings returned by Algorithm 3 is given by Figure 8a, with $\alpha = 1.49, \beta = 1.45$. The clusterings returned by k -means++ are given by Figure 8a and 8b, with $\alpha = 1.65, \beta = 1.67$. The social cost ratio is 1.42.


S-sets. It contains synthetic 2-d data with $n = 5000$ vectors and 15 Gaussian clusters with different degree of cluster overlap. Figure 9 and 10 show the experiment results on S1 set with the k -medians and k -means objectives, respectively. It can be seen that, for both objectives, our algorithm is fairer in most cases than classic algorithms, and the social cost of our algorithm is also better.



(a) Algorithm 3.

(b) Lloyd's algorithm.


Figure 7: Experiments with k -medians objective on Mopsi locations in Joensuu ($k = 10$).




(a) Algorithm 3.

(b) k -means++.

Figure 8: Experiments with k -means objective on Mopsi locations in Joensuu ($k = 10$).




(a) α and social cost on S1 dataset.




(b) β on S1 dataset.

Figure 9: Experiments with k -medians objective on S1 set.



(a) α and social cost on S1 dataset.



(b) β on S1 dataset.

Figure 10: Experiments with k -means objective on S1 set.