
Proportional Fairness in Clustering: A Social Choice Perspective

Leon Kellerhals

Technische Universität Clausthal
leon.kellerhals@tu-clausthal.de

Jannik Peters

National University of Singapore
peters@nus.edu.sg

Abstract

We study the proportional clustering problem of [Chen et al.](#) (ICML’19) and relate it to the area of multiwinner voting in computational social choice. We show that any clustering satisfying a weak proportionality notion of [Brill and Peters](#) (EC’23) simultaneously obtains the best known approximations to the proportional fairness notion of [Chen et al.](#), but also to individual fairness ([Jung et al.](#), FORC’20) and the “core” ([Li et al.](#), ICML’21). In fact, we show that any approximation to proportional fairness is also an approximation to individual fairness and vice versa. Finally, we also study stronger notions of proportional representation, in which deviations do not only happen to single, but multiple candidate centers, and show that stronger proportionality notions of [Brill and Peters](#) imply approximations to these stronger guarantees.

1 Fair clustering

Fair decision-making is a crucial research area in artificial intelligence and machine learning. To ensure fairness, a plethora of different fairness notions, algorithms and settings have been introduced, studied, and implemented. One area in which fairness has been applied extensively is (*centroid*) *clustering*: We are given a set of n data points which we want to partition into k clusters by choosing k “centers” and assigning each point to a center by which it is *represented well*. Fairness now comes into play when, e.g., the data points correspond to human individuals.

Fairness notions in clustering usually depend on one decision: whether one takes demographic information (such as gender, income, etc.) into account or whether one is agnostic to it. A large part of work on fair clustering has focused on incorporating such demographic information, starting with the seminal work of [Chierichetti et al.](#) [2017] who aimed to proportionally balance the number of people of a certain type in each cluster center. However, not all work on fair clustering relies on demographic information. Independently, and in different contexts, [Jung, Kannan, and Lutz](#) [2020] and [Chen, Fain, Lyu, and Munagala](#) [2019] instead tried to derive fairness notions from the instance itself. For [Jung et al.](#) this lead to their notion of *individual fairness*: Given a population of size n , with k cluster centers to be opened, every agent should be entitled to a cluster center not further away than their $\frac{n}{k}$ -th neighbor. While this is not always achievable, [Jung et al.](#) gave a simple algorithm achieving a 2-approximation to this notion. [Chen et al.](#) were motivated not by being fair towards individual members of the population (or agents), but towards groups of agents, defining their notion of *proportional fairness*: no group of size at least $\frac{n}{k}$ should be able to suggest a cluster center they all would be better off with. This notion is also not always achievable, and [Chen et al.](#) gave a simple $(1 + \sqrt{2})$ -approximation for it.

So far, the individual and proportional fairness notions (and some other related fairness notions) have existed in parallel, with similarities between the two being acknowledged but not formalized.¹ In their survey, [Dickerson et al. \[2023b\]](#) highlight this as a general issue in fair clustering: “each notion that was introduced [...] does not refer to or consider the interaction with the previously introduced fairness notions in clustering”. Moreover, they call for “other fairness notions in clustering that are also compatible with one another” and “general notions which possibly encompass existing ones”.

We follow this call and prove proportional and individual fairness, as well as a fairness notion by [Li et al. \[2021\]](#) which we will call the *transferable core*, to be tightly related to another. In an effort to encompass these three notions, we make use of proportionality axioms from *multiwinner voting*, an area in computational social choice [[Lackner and Skowron, 2022](#)]. Here, given the votes of n agents, the goal is to elect a size- k committee which fulfills some proportionality guarantee. We lift one of the simplest proportionality guarantees (JR) to work with metric distances and prove that any clustering fulfilling our guarantee also fulfills the best approximations for the three notions, all *simultaneously*. Moreover, such a clustering can be computed in polynomial time. Taking the multiwinner voting approach further, we also look at the lifted version of a stronger proportionality guarantee (PJR). This changes how points (agents) interact with cluster centers as they become represented not by one, but possibly multiple centers. While this is not standard for “vanilla” clustering, it is very fitting for more democratic settings, where the chosen “centers” end up possessing voting power to represent the agents. The resulting proportionality guarantee indeed highly relates to work by [Ebadian and Micha \[2024\]](#) who, motivated by *sortition* (the randomized selection of citizens’ panels [[Flanigan et al., 2021](#)]), introduced a generalization of the proportional fairness notion. Indeed, the multiwinner voting perspective allows us to prove better approximation guarantees for their fairness notion.

Our contributions. As our first main result, we provide a simple bridge between proportional fairness and individual fairness (see Section 2). Any approximation of the former is also an approximation of the latter. In particular, for any $\alpha, \beta \geq 1$ we show that (i) any α -approximation to proportional fairness is also an $(1 + \alpha)$ -approximation to individual fairness and (ii) any β -approximation to individual fairness is also a 2β -approximation to proportional fairness. These approximations are tight. We also prove a similar connection between proportional fairness and the transferable core. Our connections imply for instance that bi-criteria approximations that optimize k -means and, say, individual fairness [[Vakilian and Yalçiner, 2022](#), [Bateni et al., 2024](#)] also maintain approximations guarantees to the other fairness notions. Further, if one wants to show incompatibility of a different clustering notion with approximate proportional or individual fairness, it is sufficient to show this for one of the two notions, instead of creating instances for both (as done by [Dickerson et al. \[2023a\]](#)).

Secondly, in Section 3, we draw a connection to the area of multiwinner voting and reinterpret proportionality notions introduced by [Brill and Peters \[2023\]](#) to work with distance metrics; we call the resulting guarantees mJR and mPJR. Both of these are efficiently computable when the space of possible centers is finite. Remarkably, with simple proofs, we are able to show that any clustering satisfying mJR achieves the *best known* approximations to individual and proportional fairness notions and the transferable core. For the transferable core, we even improve upon the bound derived by [Li et al. \[2021\]](#). Finally, motivated by settings such as sortition and multiwinner voting in which agents do not only care about their closest cluster center but are represented by multiple centers, we show that a strong core stability guarantee (introduced by [Ebadian and Micha \[2024\]](#)) can be achieved by any clustering satisfying mPJR. We also deal with the case in which the center candidate space is unbounded (e.g., in Euclidean clustering settings), in which the above-mentioned algorithms can become intractable. Here, we show that satisfying the proportionality guarantees only for the set of agents is sufficient to obtain constant-factor approximations to proportional fairness and the core stability guarantee by [Ebadian and Micha \[2024\]](#).

Lastly, in Section 4, we focus on sortition: Here, the set of agents and cluster candidates is equal and each agent must be chosen with equal probability. Employing techniques from the above results, we are able to give a simpler proof achieving a better approximation guarantee for the core notion by [Ebadian and Micha \[2024\]](#).

Figure 1 (left) gives an overview over our results and our achieved approximation guarantees. Proofs of some results are deferred to a full version of this manuscript [[Kellerhals and Peters, 2023](#)].

¹For instance, in a recent tutorial on fair clustering [[Awasthi et al., 2022](#)], the two notions were treated as separate unconnected paradigms.




Figure 1: *Left*: An overview over connections between and bounds on fairness notions, i.e., α -proportional fairness (α -PF), β -individual fairness (β -IF), the (γ, α) -transferable core ((γ, α) -TC), and the α - q -core. See Sections 2 and 3 for the corresponding definitions and results. If $A \rightarrow B$, then algorithm A produces outcomes satisfying B. If $B \rightarrow C$, then any outcome satisfying B also satisfies C. If C takes a parameter α , then the label specifies the parameter that can be satisfied (for the transferable core, the result holds for all $\gamma > 1$). *Right*: The metric space for the examples used throughout the paper. Edges without labels have length 1, the distance between any two points is given by the length of the shortest path between them.

Related work. *Individual fairness* was introduced by Jung et al. [2020]. Since then, follow-up work mainly focused on bi-criteria approximation guarantees [Mahabadi and Vakilian, 2020, Negahbani and Chakrabarty, 2021, Vakilian and Yalçın, 2022, Chhaya et al., 2022, Bateni et al., 2024]. Additionally, Han et al. [2023] studied individual fairness for clustering with outliers and Sternbach and Cohen [2023] incorporated demographic information into individual fairness. The individual fairness notion was also carried over to the setting of approval-based multiwinner voting [Brill et al., 2024]. We mention that the name “individual fairness” is also used for other (unrelated) fairness notions [e.g. Kar et al., 2023, Chakrabarti et al., 2022].

Proportional fairness was first studied by Chen et al. [2019]. Micha and Shah [2020] showed that the GREEDY CAPTURE algorithm by Chen et al. achieves better approximation guarantees in certain metric spaces (including the Euclidean space with the 2-norm) and studied its complexity. Li et al. [2021] introduced notions inspired by Chen et al., which are related to the transferable core concept from algorithmic game theory. Aziz et al. [2024] introduced proportionality axioms and rules directly inspired from social choice theory to proportional clustering. Among other things, they showed that every outcome satisfying DPRF (see Section 3.1) achieves an $(1 + \sqrt{2})$ -approximation to proportional fairness. Further connections to social choice or relations between the above fairness notions of Jung et al. [2020] or Li et al. [2021] remain unexplored, though.

Ebadian and Micha [2024] study proportionality in the setting of *sortition* (see e.g., Flanigan et al. [2021]), proposing a generalization of proportional fairness and a refined variant of GREEDY CAPTURE. This variant and its proportionality were used by Caragiannis et al. [2024a] to construct panels whose decisions align with that of the underlying population. The most recent work directly related to ours was created independently and in parallel to ours by Kalaycı et al. [2024]. They study proportional fairness and the transferable core in an incomplete information setting and show that just knowing the order of the distances to between agents and center candidates suffices to achieve a 5.71-approximation to proportional fairness.

Caragiannis et al. [2024b] study proportional fairness in a *non-centroid* based fair clustering setting, in which points are not assigned to cluster centers. Instead they are grouped into clusters and derive utility based on the other agents in their cluster. For this setting, Caragiannis et al. [2024b] also build on the proportional fairness framework studied in this work and also take inspiration from concepts from multiwinner voting: in their case, they adopt the *FJR* axiom of Peters et al. [2021]. The setting of non-centroid clustering is closely related to the study of *hedonic games*, with the difference being that in hedonic games, the number of clusters is not pre-determined. See for instance, Fanelli et al. [2021], Demeulemeester and Peters [2023], and Fioravanti et al. [2023] for recent works on approximate core stability in hedonic games. Ahmadi et al. [2022], Aamand et al. [2023], and Mosenzon and Vakilian [2024] further studied a notion of individual stability for clustering, in which individual agents should not be able to deviate from their clusters. This, however, is unrelated to group stability as studied in this work.

Multiwinner voting is the branch of computational social choice theory dealing with selecting multiple instead of just one candidate as a winner. A main branch herein focuses on *proportionality*. While much of the literature on proportionality, starting with Aziz et al. [2017], focuses on approval preferences (see Lackner and Skowron [2022] for a recent book on this topic), proportionality notions also exist for ordinal preferences [Dummett, 1984]. These notions were recently strengthened by Aziz and Lee [2020, 2021] and Brill and Peters [2023], with the latter forming the basis for the proportionality axioms we discuss in this paper. We are further closely related to the works of Caragiannis et al. [2022] and Ebadian et al. [2022] who studied the representation of a given committee by investigating the distances of agents to their q -closest committee member.

Model and notation. Let $(\mathcal{X}, d)$ be a (pseudo)-metric space with a distance function $d: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ satisfying $d(i, i) = 0$, $d(i, j) = d(j, i)$ and $d(i, j) + d(j, k) \geq d(i, k)$. Let $i \in \mathcal{X}$ be a point. For $r \in \mathbb{R}$, define $B(i, r) = \{j \in \mathcal{X} : d(i, j) \leq r\}$ to be the ball of radius r around i . For $W \subseteq \mathcal{X}$, let $d(i, W) = \min_{c \in W} d(i, c)$. For $q \leq |W|$, $d^q(i, W)$ is distance to the q -th closest point in W to i . Note that $d^1(i, W) = d(i, W)$ and that $d^q(i, W) \leq d(i, j) + d^q(j, W)$ for $i, j \in N$.

Throughout the paper, we are given a set of agents $N = [n]$ and a (possibly infinite) set of candidates (facilities) C , both of which lie in a metric space $(\mathcal{X}, d)$, and a number $k \in \mathbb{N}^+$. A *clustering* or *outcome* is a subset $W \subseteq C$ of at most k candidates. The elements $c \in W$ are called *centers*. Our examples use the (weighted) graph metric in which the points are the vertices of a graph with edge lengths, and the distance between two points is the length of a shortest path between them.

2 Relations between proportional fairness notions

In this section, we prove the relations between proportional fairness [Chen et al., 2019], individual fairness [Jung et al., 2020], and the transferable core [Li et al., 2021]. We first define the notions.

The idea of *proportional fairness* is the following: If there is a candidate c such that at least $\frac{n}{k}$ agents are closer to c by a factor α than to their closest cluster center in the outcome W , then we say that the agents will *deviate* to c . If there is no such candidate, the outcome satisfies α -proportional fairness.

Definition 1. For $\alpha \geq 1$ an outcome W satisfies α -proportional fairness, if there is no group $N' \subseteq N$ of agents with $|N'| \geq \frac{n}{k}$ and $c \notin W$ such that $\alpha \cdot d(i, c) < d(i, W)$ for all $i \in N'$.

While $(2 - \varepsilon)$ -proportional fair outcomes need not exist (for any $\varepsilon > 0$), $(1 + \sqrt{2})$ -proportional fair outcomes can be computed for any metric space [Chen et al., 2019, Micha and Shah, 2020].

To define *individual fairness*, denote by $r_{N,k}(i)$ be the radius of the smallest ball around an agent $i \in N$ that encloses at least $\frac{n}{k}$ agents, i.e., $r_{N,k}(i) = \min\{r \in \mathbb{R} : |B(i, r) \cap N| \geq \frac{n}{k}\}$. We drop the subscripts N and k if clear from context. For this definition to properly work, we additionally need the assumption that $N \subseteq C$, i.e., any agent can be chosen as center. Otherwise, a secluded group of agents without any possible cluster centers around them would never be able to get a center close to them in the outcome. Indeed, this is a plausible restriction in metric clustering, as oftentimes the centers may be picked from the (infinite) set of points in the metric space.

Definition 2. For an instance with $N \subseteq C$, for $\beta \geq 1$ an outcome W satisfies β -individual fairness if $d(i, W) \leq \beta r_{N,k}(i)$ for all $i \in N$.

It is known that an outcome satisfying 2-individual fairness always exists, while there are instances with no $(2 - \varepsilon)$ -individually fair outcome [Jung et al., 2020].

The *transferable core*² notion is based on the concept of transferable utilities from game theory. Comparing to proportional fairness, the notion considers the average utility for each group.

Definition 3. For $\gamma, \alpha \geq 1$, an outcome W is in the (γ, α) -transferable core if there is no group of agents $N' \subseteq N$ and candidate $c \notin W$ with $|N'| \geq \gamma \frac{n}{k}$ and $\alpha \sum_{i \in N'} d(i, c) < \sum_{i \in N'} d(i, W)$.

It is known that for any $\gamma > 1$ there are outcomes in the $(\gamma, \max(4, \frac{3\gamma-1}{\gamma-1}))$ -transferable core while there need not be outcomes in the $(\gamma, \min(1, \frac{1}{\gamma-1}))$ -transferable core [Li et al., 2021].

²We remark that Li et al. [2021] call this notion just “core”, we rename it to avoid confusion with the core notion of Ebadian and Micha [2024].

Example 1. Consider the instance depicted in Figure 1 (right) with $k = 5$ and the associated graph distance metric. Assume that cluster centers can only be placed on the depicted agents. We have $\frac{n}{k} = 2$; thus any two agents are able to deviate to another center. The outcome $W = \{1, 2, 3, 6, 9\}$ satisfies 1-proportional fairness: The agents 1, $\dots$, 4 have distance 0 to a center, while every remaining agent has distance at most 1 to a center.

To see the difference between proportional fairness, individual fairness, and the transferable core, consider the same instance with $k = 4$, so $\frac{n}{k} = 2.5$. Here, the outcome $W = \{1, 2, 6, 7\}$ satisfies 1-proportional fairness, however it does not satisfy 1-individual fairness. Agent 8 could look at their 2 closest neighbors, 5 and 9, both at a distance of 1. However, the distance of 8 to the outcome is 2. Observe that W also is not in the $(1, 1)$ -transferable core. Here, for the group $N' = \{8, 9, 10\}$ and candidate $c = 9$, we have $\sum_{i \in N'} d(i, c) = 2 < \sum_{i \in N'} d(i, W) = 4$. $\diamond$

2.1 Proportional and individual fairness

We first show that proportional and individual fairness are the same up to a factor of at most 2.

Theorem 1. *Let $\alpha, \beta \geq 1$. If $N \subseteq C$, then an outcome that satisfies α -proportional fairness also satisfies $(1 + \alpha)$ -individual fairness, and an outcome that satisfies β -individual fairness also satisfies 2β -proportional fairness. If $N = C$, then an outcome that satisfies β -individual fairness also satisfies $(1 + \beta)$ -proportional fairness.*

Proof. Let $W \subseteq C$ be an outcome satisfying α -proportional fairness, $j \in N$ be any agent, and $N_j = \{i \in N : d(i, j) \leq r(j)\}$. As $N \subseteq C$, there is an $i \in N_j$ with $d(i, W) \leq \alpha d(i, j)$; otherwise the coalition N_j deviates to candidate j . Thus, by the triangle inequality, $d(j, W) \leq d(i, j) + d(i, W) \leq (1 + \alpha)d(i, j) \leq (1 + \alpha)r(j)$, and hence W satisfies $(1 + \alpha)$ -individual fairness.

Now suppose the outcome W satisfies β -individual fairness. Let $N' \subseteq N$ with $|N'| \geq \frac{n}{k}$ and $c \notin W$ be an unchosen candidate. Take $i^* \in N'$ to be the agent in N' furthest away from c . If $N \subseteq C$, then the radius $r(i^*)$ containing $\lceil \frac{n}{k} \rceil$ agents is at most as large as the most distant agent in N' , i.e., there is an $i' \in N'$ with $r(i^*) \leq d(i^*, i') \leq d(i^*, c) + d(c, i')$. Then $d(i^*, W) \leq \beta r(i^*) \leq \beta(d(i^*, c) + d(c, i')) \leq 2\beta d(i^*, c)$. If $N = C$, then, since $|N'| \geq \frac{n}{k}$, we have $r(c) \leq d(c, i^*)$; thus $d(c, W) \leq \beta d(c, i^*)$. Therefore, $d(i^*, W) \leq d(i^*, c) + d(c, W) \leq (1 + \beta)d(i^*, c)$, and thus W also satisfies $(1 + \beta)$ -proportional fairness. $\square$

Indeed, we also show that all three provided bounds are tight.

Theorem 2. *For every $\alpha, \beta \geq 1$ and $\varepsilon > 0$, there are instances with $N = C$ for which there exists (1) an outcome which satisfies α -proportional fairness, but not $(1 + \alpha - \varepsilon)$ -individual fairness, and (2) an outcome which satisfies β -individual fairness, but not $(1 + \beta - \varepsilon)$ -proportional fairness. Moreover, there are instances with $N \subseteq C$ for which there exists (3) an outcome which satisfies β -individual fairness, but not $(2\beta - \varepsilon)$ -proportional fairness.*

2.2 Proportional fairness and the transferable core

It is easy to see that the $(1, \alpha)$ -transferable core implies α -proportional fairness. For $\gamma > 1$ however, the (γ, α) -transferable core does not imply any meaningful proportional fairness approximation (consider $\frac{n}{k}$ agents on one point and $(\gamma - 1)\frac{n}{k}$ agents “far” away). Hence, we focus on the other direction and show that a proportional fairness approximation implies one to the transferable core.

Theorem 3. *An outcome satisfying α -proportional fairness is in the $(\gamma, \frac{\gamma(\alpha+1)}{\gamma-1})$ -transferable core for any $\alpha \geq 1$ and $\gamma > 1$.*

Proof. Let $W \subseteq C$ satisfy α -proportional fairness. Let $N' \subseteq N$ be a group of agents of size $n' \geq \gamma \frac{n}{k}$, $c \notin W$, and shorten $\eta = \lceil \frac{n}{k} \rceil$. Further, assume the agents $N' = \{i_1, \dots, i_{n'}\}$ are ordered by their increasing distance to c , i.e., $d(i_j, c) \leq d(i_{j+1}, c)$ for every $j \in [n' - 1]$. Let $J_0 = \{i_1, \dots, i_\eta\}$ and $j_0 \in J_0$ such that $d(j_0, W) \leq \alpha d(j_0, c)$; such an agent must exist due to α -proportional fairness. Next, for $\lambda = 1, \dots, n' - \eta$, we inductively define $J_\lambda = \{i_1, \dots, i_{\eta+\lambda}\} \setminus \{j_0, \dots, j_{\lambda-1}\}$, and choose $j_\lambda \in J_\lambda$ such that $d(j_\lambda, W) \leq \alpha d(j_\lambda, c) \leq \alpha d(i_{\eta+\lambda}, c)$ (note that $|J_\lambda| = \eta$). Thus,

$$\sum_{\lambda=0}^{n'-\eta} d(j_\lambda, W) \leq \alpha \sum_{z=\eta}^{n'} d(i_z, c) \leq \alpha \sum_{i \in N'} d(i, c). \quad (1)$$

Next, for each $i \in N'' = N' \setminus \{j_0, \dots, j_{n'-\eta}\}$, we can bound the distance to W as follows:

$$d(i, W) \leq d(i, c) + d(c, j_0) + d(j_0, W) \leq d(i, c) + (1 + \alpha)d(i_\eta, c). \quad (2)$$

Note that $d(i_\eta, c) \leq \frac{1}{n'-|N''|} \sum_{z=\eta}^{n'} d(i_z, c)$ as each of the summands is at least $d(i_\eta, c)$. Thus,

$$\sum_{i \in N''} d(i, W) \leq \sum_{i \in N''} (d(i, c) + (1 + \alpha)d(i_\eta, c)) \leq (1 + \alpha) \frac{|N''|}{n'-|N''|} \sum_{z=\eta}^{n'} d(i_z, c) + \sum_{i \in N''} d(i, c). \quad (3)$$

As $n' \geq \gamma \frac{n}{k}$ and $|N''| = \eta - 1 \leq \frac{n}{k}$, we have $\frac{|N''|}{n'-|N''|} \leq \frac{1}{\gamma-1}$. In all, $\sum_{i \in N'} d(i, W)$ is the sum of (1) and (3), which is at most

$$\left(\frac{\alpha+1}{\gamma-1} + \alpha + 1\right) \sum_{i \in N'} d(i, c) = \frac{\gamma(\alpha+1)}{\gamma-1} \sum_{i \in N'} d(i, c). \quad (4)$$

Hence, W is in the $\left(\gamma, \frac{\gamma(\alpha+1)}{\gamma-1}\right)$ -transferable core. $\square$

We remark that the denominator $\frac{1}{\gamma-1}$ in α is inevitable. This is because for $\gamma \leq 2$, the $\left(\gamma, \frac{1}{\gamma-1}\right)$ -transferable core may be non-empty [Li et al., 2021, Theorem 18]. We complement the above upper bound with an asymptotically tight lower bound.

Theorem 4. For any $\alpha \geq 1$, $\gamma > 1$, and $\varepsilon > 0$ there exists an instance in which an α -proportional fair outcome is not in the $\left(\gamma, \frac{\gamma\alpha+1}{\gamma-1} - \varepsilon\right)$ -transferable core.

3 Fairness notions for multiwinner voting axioms

In this section we show a connection between the research on computational social choice, specifically approval-based multiwinner voting (also known as approval-based committee (ABC) voting) and the fairness notions for clustering. We will first give a primer on ABC voting and introduce our *metric JR axioms*. We then focus on two of those axioms and show that (1) they are satisfied by existing, simple algorithms, and (2) they imply the best known approximation guarantees to proportional and individual fairness, the transferable core, and the q -core (see Definition 5 below). For the latter two notions, we are even able to improve upon the best currently known approximation guarantees. Finally, we will focus on the case when the candidate set is infinitely large (i.e., when we are in the Euclidean space and every point is a candidate): In this setting, the above algorithms become hard to compute. We combine two approaches to maneuver around this hardness and again match upon the best known approximation guarantees for proportional fairness and the q -core.

3.1 Metric JR axioms

In ABC voting [Lackner and Skowron, 2022], we are given a set N of voters (or agents), a set C of candidates, and a committee size k . For each voter $i \in N$ we are given a subset $A_i \subseteq C$ of candidates they approve. For such preferences, we call a set $N' \subseteq N$ of voters ℓ -large if $|N'| \geq \ell \frac{n}{k}$, and ℓ -cohesive if $|\bigcap_{i \in N'} A_i| \geq \ell$. We say that a committee satisfies

- JR if for every 1-cohesive and 1-large group N' there exists an $i \in N'$ with $|A_i \cap W| \geq 1$;
- PJR if for every $\ell \in [k]$ and ℓ -cohesive and ℓ -large group N' it holds that $|\bigcup_{i \in N'} A_i \cap W| \geq \ell$,

and remark that there are many further proportionality axioms [Lackner and Skowron, 2022]. Here, JR is short for *justified representation*. To define our *metric JR axioms* for voters and candidates in a distance metric, we follow Brill and Peters [2023] (who lifted these axioms for weak ordinal preferences and called them rank-II) and generalize their notions to look at each distance separately.

Definition 4 (Metric JR axioms). Let $(\mathcal{X}, d)$ be a distance metric, let $N, C \subseteq \mathcal{X}$. Let Π be a proportionality axiom. An outcome W satisfies $m\Pi$ (short for metric) if for all $y \in \mathbb{R}_{\geq 0}$, for the ABC voting instance in which each $i \in N$ has the approval set $B(i, y) \cap C$, the outcome W satisfies Π .

For example, an outcome W satisfies mJR if for every $y \in \mathbb{R}$ and for every group $N' \subseteq N$ of at least $\frac{n}{k}$ agents whose ball of radius y all contain a common candidate ($|\bigcap_{i \in N'} B(i, y) \cap C| \geq 1$), there exists an agent $i \in N'$ whose ball of radius y contains a center $c \in W$.




Figure 2: Metric space for some of the examples. Edges without labels have length 1.

We want to point out that mJR is *significantly* weaker than mPJR. Indeed, to satisfy mPJR, an outcome W may need to contain several candidates c such that $d(i, c) > d(i, W)$ for all agents $i \in N$, i.e., c is no-one’s “first choice” among W ; mJR does not have this property. This makes mJR the more sensible of the two axioms for “vanilla” clustering, in which one only cares about the closest center to each agent. mPJR however, is a natural axiomatic choice for settings such as sortition or even social choice in general: Here, agents may benefit from having more than a single representative. We provide some intuition for mJR and mPJR and this property in the example below.

Example 2. To see the differences between the proportionality axioms, consider the instance depicted in Figure 1 (right). First consider instance (a) on the left with $N = C = \{1, \dots, 10\}$, $k = 4$, and the outcome $W = \{1, 2, 3, 6\}$. Here, $\frac{n}{k} = 2.5$. First, we note that this outcome does not satisfy 1-proportional fairness: The agents 8, 9, 10 are closer to 9 than they are to the closest winner in W . It does however satisfy mJR: Among every group of at least three agents that have a common candidate within distance y , there is one agent that has a cluster center $w \in W$ within distance y . For example, 8, 9, 10 have candidate 9 at distance 1, and the distance of 9 to the closest center is also 1. This outcome does not satisfy mPJR though, since the group 5, $\dots$, 10 would deserve at least two candidates within distance 1 in W . An outcome satisfying mPJR is $W = \{1, 2, 3, 9\}$. For $y = 0$, only the group $\{1, \dots, 4\}$ shares a candidate, but also have a center at distance 0.

If $k = 5$, then, to satisfy mPJR, an outcome must contain at least two of the four candidates. But there are outcomes satisfying mJR that contain only one of $1, \dots, 4$. This property of mPJR makes it suited for settings in which agents may want to be represented by multiple candidates, e.g., in political settings, in which the candidates end up possessing voting power to represent the agents. $\diamond$

Independently of Brill and Peters [2023], Aziz et al. [2024] introduced two notions they call *Proportionally Representative Fairness*. The first notion is called “discrete” (DPRF), and the second is called “unconstrained” (UPRF). Indeed, DPRF is equivalent to mPJR. UPRF was introduced to tackle the case when the candidate space is unbounded. We discuss how it relates to mPJR and the other fairness notions in the full version [Kellerhals and Peters, 2023].

Aziz et al. [2024] show that an outcome satisfying DPRF (mPJR) also fulfills $(1 + \sqrt{2})$ -proportional fairness. We show hereafter that this already holds for the (much weaker) mJR axiom.

3.2 Fairness bounds for mJR outcomes

We now prove the approximation guarantees implied by mJR. We remark that the bound for the transferable core below improves upon the analysis of Li et al. [2021]. The proof is deferred to the full version [Kellerhals and Peters, 2023].

Theorem 5. *Let W be an outcome satisfying mJR. Then it also satisfies $(1 + \sqrt{2})$ -proportional fairness, 2-individual fairness, and is in the $\left(\gamma, \frac{2\gamma}{\gamma-1}\right)$ -transferable core for any $\gamma > 1$.*

If the candidate space is finite, then an outcome satisfying mJR can be computed in polynomial time by the GREEDY CAPTURE algorithm [Chen et al., 2019, Micha and Shah, 2020, Li et al., 2021]. We briefly recall its procedure:

GREEDY CAPTURE starts off with an empty clustering W . It maintains a radius δ (initially $\delta = 0$) and smoothly increases δ . If there is a candidate c such that at least $\frac{n}{k}$ agents have distance at most δ to c , it adds c to W and deletes the $\frac{n}{k}$ agents. If an agent has distance at most δ to a candidate in W , then it is deleted as well. This is continued until all agents are deleted.

Example 3. Consider the instance in Figure 2. Here, with $k = 4$, GREEDY CAPTURE, would first open one of $1, \dots, 4$ with $\delta = 0$ and remove all agents from $1, \dots, 4$. Then for $\delta = 1$ it could either open 6 or 9, removing all adjacent agents to it. Then there are two agents remaining, which would

be assigned to either 6 or 9 for $\delta = 2$. Thus, in this instance, GREEDY CAPTURE only opens two clusters. $\diamond$

Proposition 6. *Any outcome returned by GREEDY CAPTURE satisfies mJR.*

3.3 Fairness bounds for mPJR outcomes

Recall that mPJR is equivalent to the DPRF notion by Aziz et al. [2024]. To satisfy their notion, they designed a generalization of the expanding approvals rule from multiwinner voting [Aziz and Lee, 2020, 2021] (in which the agents' preferences over the candidates are ordinal) to the setting of proportional clustering. They refer to this generalization as SPATIAL EXPANDING APPROVALS. As Aziz et al. show, SPATIAL EXPANDING APPROVALS can be computed in polynomial time for finite candidate spaces.

In general, SPATIAL EXPANDING APPROVALS behaves similarly to GREEDY CAPTURE. It also starts off with an empty clustering W and a radius $\delta = 0$ as well and additionally gives each agent a budget $b_i = \frac{k}{n}$. It then smoothly increases the radius δ . When there is a candidate $c \notin W$ for which the agents at a distance of at most δ have a budget of at least 1, it decreases the budget of these agents collectively by exactly 1 and adds c to W .

Example 4. Consider the instance in Figure 2. Here, with $k = 4$, SPATIAL EXPANDING APPROVALS would give each agent a budget of $\frac{4}{10} = \frac{2}{5}$. For $\delta = 0$, it will open a cluster from $1, \dots, 4$ and decrease their budgets by exactly 1, for instance it could set the budget of 1 and 2 to 0 and of 3 to $\frac{1}{5}$. Then for $\delta = 1$ it could again open 6 and 9, for instance by removing the budget of 6 and 9 to $\frac{1}{5}$ and of 5, 7, 8, 10 to zero. The remaining budget is exactly 1, which would be spent for $\delta = 10$ on 5. Thus, one possible final outcome is $\{1, 5, 6, 9\}$. $\diamond$

Remarkably, as shown by Aziz et al. [2024], it does not matter in which way the budget is subtracted and which candidate meeting the budget is selected; the outcomes of the algorithm will satisfy mPJR in any case.

Proposition 7. *Any outcome returned by SPATIAL EXPANDING APPROVALS satisfies mPJR.*

We now turn to fairness measures implied by mPJR. As any outcome satisfying mPJR also fulfills mJR, the results in Theorem 5 also hold for mPJR. Indeed, mPJR is stricter in the sense that larger groups must also be represented justly by a proportional number of candidates: an α percentage of the population should roughly be close to an α percentage of the centers.

This property makes mPJR fit well into metric social choice settings such as sortition. For this, Ebadian and Micha [2024] introduced a fairness notion that measures proportionality in this setting by considering for each agent not only the closest center, but the first q closest centers. In that, their notion called α - q -core naturally generalizes α -proportional fairness; the two are equal when $q = 1$.

Definition 5. For $\alpha \geq 1$ an outcome W is in the α - q -core, if there is no $\ell \in \mathbb{N}$ and no $N' \subseteq N$ with $|N'| \geq \ell \frac{n}{k}$ and set $C' \subseteq C$ with $q \leq |C'| \leq \ell$ such that $\alpha \cdot d^q(i, C') < d^q(i, W)$ for all $i \in N'$.

Example 5. Consider the instance in Figure 2 with $k = 5$ and the outcome $W = \{1, 2, 3, 6, 9\}$. As mentioned above, W satisfies 1-proportional fairness. For the 3-core however, consider the set $N' = \{5, \dots, 10\}$ deviating to $C' = \{6, 9, 10\}$. The distance of any member of N' to 6, 9, or 10 is at most 3, while the distance to the third most distant center in the outcome is at least 10. Thus, when considering the distances to the third most distant candidate in C' , every agent in N' would improve by at least a factor of $\frac{10}{3}$. Thus, W is only in the $\frac{10}{3}$ -3-core. $\diamond$

We mention in passing that we introduce similar generalizations for individual fairness and the transferable core, in which each agent is represented by q candidates instead of one. The definitions and obtained results can be found in the full version of this paper [Kellerhals and Peters, 2023].

Ebadian and Micha [2024] show that, if $N = C$ (every agent is a candidate and vice versa), for a given q , one can compute a $\frac{5+\sqrt{41}}{2}$ - q -core outcome.³ We show that mPJR (or DPRF) provides a better guarantee for the q -core, for all values of q simultaneously.

Theorem 8. *If an outcome satisfies mPJR, then, for every $q \leq k$, it is in the 5- q -core.*

³In addition, their randomized algorithm selects each agent with the same probability, a desirable property in the context of sortition, i.e., the randomized selection of citizen assemblies [Flanigan et al., 2021].

To prove the theorem, we use two lemmas. The first was first observed by Ebadian and Micha [2024, Lemma 1] and is proven here in a shorter fashion.

Lemma 9. *Let $\ell \geq q \geq 0$, let $N' \subseteq N$ be a set of agents with $|N'| \geq \ell \frac{n}{k}$, and let $C' \subseteq C$ be a set of $q \leq |C'| \leq \ell$ candidates such that $d^q(i, C') \leq d^q(i, W)$ for any $i \in N'$. Then there is a set $N'' \subseteq N$ of at least $q \frac{n}{k}$ agents and a candidate $c \in C'$ such that $d(i, c) \leq d^q(i, C')$ for all $i \in N''$.*

Proof. Assume that each agent marks each of their top q choices among C' . Then there are at least $q|C'| \frac{n}{k}$ marks on the candidates. Thus, there is one $c \in C'$ with at least $q \frac{n}{k}$ marks. $\square$

The next lemma bounds the α - q -core once we find two agents with specific bounds on their distances.

Lemma 10. *Let $\rho_1, \rho_2, \sigma_1, \sigma_2 \geq 0$ and let $W \subseteq C$ be an outcome. If for any set $N' \subseteq N$ of at least $\ell \frac{n}{k}$ agents and any candidate set $C' \subseteq C$ with $q \leq |C'| \leq \ell$ there are $i_1, i_2 \in N'$ such that $d^q(i_1, W) \leq \rho_1 d^q(i_1, C') + \rho_2 d^q(i_2, C')$ and $d^q(i_2, W) \leq \sigma_1 d^q(i_1, C') + \sigma_2 d^q(i_2, C')$, then W is in the α - q -core, where $\alpha \leq \rho_1 + \frac{1}{2} \left(\sigma_2 - \rho_1 + \sqrt{(\rho_1 - \sigma_2)^2 + 4\sigma_1\rho_2} \right)$.*

Proof of Theorem 8. Let $N' \subseteq N$ be a group of agents with $|N'| \geq \ell \frac{n}{k}$ and let $C' \subseteq C$ with $q \leq |C'| \leq \ell$ such that $d^q(i, C') \leq d^q(i, W)$ for any $i \in N'$. By Lemma 9 there is a candidate c being ranked in their top q among C' by $q \frac{n}{k}$ many agents $N'' \subseteq N'$. Out of N'' , let i_1 be the agent maximizing $d(i_1, c)$ and i_2 be any other agent in N'' . Also, let $C'' \subseteq C'$ be the set of the q candidates closest to i_2 . Then, for every $c' \in C''$ and every $i \in N''$, we have

$$d(i, c') \leq d(i, c) + d(c, i_2) + d(i_2, c') \leq d(i_1, c) + 2d^q(i_2, C') =: y.$$

In other words, $|\bigcap_{i \in N''} B(i, y) \cap C| \geq q$. Now mPJR implies that $|\bigcup_{i \in N''} B(i, y) \cap W| \geq q$. Thus, for every $i \in N''$ there is an agent $i' \in N''$ such that $d^q(i, W) \leq d(i, i') + y$. Since the distance of i_1 to any other agent $i' \in N''$ is $d(i_1, i') \leq d(i_1, c) + d(c, i_1)$, we have

$$d^q(i_1, W) \leq 2d(i_1, c) + y \leq 3d^q(i_1, C') + 2d^q(i_2, C').$$

As the distance of i_2 to any other agent $i' \in N''$ is at most $d(i', c) + d(c, i_2) \leq d(i_1, c) + d^q(i_2, C')$,

$$d^q(i_2, W) \leq d(i_1, c) + d^q(i_2, C') + y \leq 2d(i_1, c) + 3d^q(i_2, C') \leq 2d^q(i_1, C') + 3d^q(i_2, C').$$

Applying Lemma 10 with $\rho_1 = \sigma_2 = 3$ and $\rho_2 = \sigma_1 = 2$ yields the stated 5- q -core. $\square$

3.4 Dealing with unbounded candidate sets

Whenever the candidate space C is finite, it is straightforward to implement GREEDY CAPTURE and SPATIAL EXPANDING APPROVALS in polynomial time. However, as shown by Micha and Shah [2020], once C is unbounded and the metric space is only implicitly given (e.g., some distance norm over $C = \mathbb{R}^t$), computing GREEDY CAPTURE can become NP-hard. For Euclidean distances over $C = \mathbb{R}^t$, Micha and Shah [2020, Theorem 12] were nevertheless able to give an approximate version of GREEDY CAPTURE, which approximates proportional fairness up to a factor of $2 + \varepsilon$ for any $\varepsilon > 0$ in this special metric space. For general metric spaces, Micha and Shah [2020, Theorem 11] show that in an instance with $N \subseteq C$, any outcome which is α -proportionally fair when restricted to the instance with candidate set N is 2α -proportionally fair in the whole instance. Aziz et al. [2024] used a very similar approach to this and showed that running SPATIAL EXPANDING APPROVALS on the agents results in a 3-proportionally fair outcome. Combining both approaches, we show that any outcome W satisfying mJR when restricted to the instance with candidate set $N \cup W$ satisfies 3-proportional fairness in the entire instance. Thus, GREEDY CAPTURE restricted to the agents yields a 3-proportionally fair outcome. The same also applies to mPJR and the q -core.

Theorem 11. *Consider an instance I with $N \subseteq C$ and an outcome W and let I' be the instance with agent set N and candidate set $N \cup W$. If W satisfies mJR in I' , then W satisfies 3-proportional fairness. If W satisfies mPJR in I' , then W is in the 4- q -core for all $q \leq k$.*

4 Stronger fairness bounds for sortition

Ebadian and Micha [2024] introduced FAIR GREEDY CAPTURE, a randomized generalization of GREEDY CAPTURE for the setting of sortition. It works in the setting in which $N = C$ and is parameterized by some parameter $q \leq k$. Like GREEDY CAPTURE it smoothly increases a radius around each agent/candidate. Once this radius contains at least $q \frac{n}{k}$ agents, it selects q of them uniformly at random and deletes in total $\lceil q \frac{n}{k} \rceil$ of these agents. Together with an adequate final sampling step, one can show that this selects each agent with a probability of exactly $\frac{k}{n}$.

Ebadian and Micha show that any clustering returned by the algorithm is in the $\frac{3+\sqrt{17}}{2}$ -1-core when parameterized by $q = 1$ and in the $\frac{5+\sqrt{41}}{2} \approx 5.7$ - q -core when parameterized by $q > 1$.⁴ We improve upon their analysis (with a simpler proof) and show that FAIR GREEDY CAPTURE satisfies a better bound for every parameter $q \leq k$.

Theorem 12. *Let $N = C$ and $q \leq k$. Then any outcome W returned by FAIR GREEDY CAPTURE parameterized by q is in the $\frac{3+\sqrt{17}}{2}$ - q -core.*

5 Conclusion and future work

In this paper, we studied proportional clustering from a social choice perspective and showed that our new *metric JR axioms* enable near-optimal approximations of fairness notions for clustering. An interesting open question, both relevant to social choice and clustering is related to a different relaxation of proportional fairness (or core fairness) introduced by Jiang et al. [2020]. Instead of bounding the factor by which the agents can improve, they bound the size of the deviating coalition (similar to the transferable core). In that sense, no group of size $\gamma \frac{n}{k}$ should exist, who could all deviate to a candidate they like more. In their work, they show that there are instances for which no solution with $\gamma < 2$ can exist while for any $\varepsilon > 0$ a solution with $\gamma = 16 + \varepsilon$ exists. Since these results only care about the relative ordering of the candidates, they also translate to clustering. Closing this bound, or improving it for certain metric spaces, seems like an interesting problem. It would be also intriguing to study the probabilistic analog of the core [Cheng et al., 2020, Jiang et al., 2020], especially if the results generalize to the q -core and if certain metric spaces admit simple algorithms to compute it.

Further, SPATIAL EXPANDING APPROVALS (Section 3.3) is more of a family of algorithms, parameterized by how candidates are selected and how budgets are deducted. Is there any way to axiomatically (or quantitatively) distinguish its different parameterizations? In the context of approval-based multi-winner voting, the *Method of Equal Shares* [Peters and Skowron, 2020] can be seen as an instantiation of SPATIAL EXPANDING APPROVALS which provides stronger proportionality guarantees than other algorithms in the family. Is something similar possible for our setting, e.g., can one go from proportionality axioms inspired by PJR to axioms inspired by the stronger EJR axiom [Aziz et al., 2017]? As shown by Brill and Peters [2023, Example 7] the straightforward extension of studying mEJR (or rank-EJR in their notation) is not possible, as outcomes satisfying mEJR may easily fail to exist. However, the metric variant of the PJR+ axiom of Brill and Peters [2023] may be of greater interest in the clustering setting. It is easy to see that SPATIAL EXPANDING APPROVALS satisfies mPJR+. Is it also possible to derive better proportionality or core approximations from mPJR+?

Naturally, our work still leaves several open questions when it comes to the approximation factors of our notions. What are the best attainable factors for proportional fairness and the q -core? Further, the questions of Jung et al. [2020] whether the bound of 2 on individual fairness can be improved for Euclidean spaces and of Micha and Shah [2020] whether for (unweighted) graph metrics with $N = C$ a 1-proportional fair clustering always exist, are still open.

References

A. Aamand, J. Chen, A. Liu, S. Silwal, P. Sukprasert, and A. Vakilian. Constant approximation for individual preference stable clustering. In *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS)*, pages 43646–43661, 2023.

⁴Recall that the algorithm receives q as a parameter in the input; thus, as opposed to our Theorem 8, the bounds only hold for one q at a time. In a recently updated version of their preprint, Ebadian and Micha [2024] modified their initial version of FGC to additionally satisfy the 6- q core for all q simultaneously.

- S. Ahmadi, P. Awasthi, S. Khuller, M. Kleindessner, J. Morgenstern, P. Sukprasert, and A. Vakilian. Individual preference stability for clustering. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, pages 162:197–246, 2022.
- P. Awasthi, B. Brubach, D. Chakrabarty, J. P. Dickerson, S. A. Esmaeili, M. Kleindessner, M. Knittel, J. Morgenstern, S. Samadi, A. Srinivasan, and L. Tsepenekas. Fairness in clustering. <https://www.fairclustering.com/>, 2022.
- H. Aziz and B. E. Lee. The expanding approvals rule: improving proportional representation and monotonicity. *Social Choice and Welfare*, 54:1–45, 2020.
- H. Aziz and B. E. Lee. Proportionally representative participatory budgeting with ordinal preferences. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5110–5118. AAAI Press, 2021.
- H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.
- H. Aziz, B. E. Lee, S. Morota Chu, and J. Vollen. Proportionally representative clustering. In *Proceedings of the 20th International Workshop on Internet and Network Economics (WINE)*, 2024. Forthcoming.
- M. Bateni, V. Cohen-Addad, A. Epasto, and S. Lattanzi. A scalable algorithm for individually fair k-means clustering. In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 3151–3159, 2024.
- M. Brill and J. Peters. Robust and verifiable proportionality axioms for multiwinner voting. In *Proceedings of the 24th ACM Conference on Economics and Computation (ACM-EC)*, page 301. ACM, 2023. Full version arXiv:2302.01989 [cs.GT].
- M. Brill, J. Israel, E. Micha, and J. Peters. Individual representation in approval-based committee voting. *Social Choice and Welfare*, 2024. Forthcoming.
- I. Caragiannis, N. Shah, and A. A. Voudouris. The metric distortion of multiwinner voting. *Artificial Intelligence*, 313:103802, 2022.
- I. Caragiannis, E. Micha, and J. Peters. Can a few decide for many? The metric distortion of sortition. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, pages 235:5660–5679. PMLR, 2024a.
- I. Caragiannis, E. Micha, and N. Shah. Proportional fairness in non-centroid clustering. In *Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS)*, 2024b. Forthcoming.
- D. Chakrabarti, J. P. Dickerson, S. A. Esmaeili, A. Srinivasan, and L. Tsepenekas. A new notion of individually fair clustering: α -equitable k-center. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 6387–6408, 2022.
- X. Chen, B. Fain, L. Lyu, and K. Munagala. Proportionally fair clustering. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 1032–1041, 2019.
- Y. Cheng, Z. Jiang, K. Munagala, and K. Wang. Group Fairness in Committee Selection. *ACM Transactions on Economics and Computation*, 8(4):23:1–23:18, 2020.
- R. Chhaya, A. Dasgupta, J. Choudhari, and S. Shit. On coresets for fair regression and individually fair clustering. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 9603–9625, 2022.
- F. Chierichetti, R. Kumar, S. Lattanzi, and S. Vassilvitskii. Fair clustering through fairlets. In *Proceedings of the 30th Conference on Neural Information Processing Systems (NeurIPS)*, pages 5029–5037, 2017.
- T. Demeulemeester and J. Peters. Relaxed core stability for hedonic games with size-dependent utilities. In *Proceedings of the 48th International Symposium on Mathematical Foundations of Computer Science (MFCS 2023)*, pages 41:1–41:14, 2023.

- J. P. Dickerson, S. A. Esmaili, J. Morgenstern, and C. J. Zhang. Doubly constrained fair clustering. In *Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS)*, pages 13267–13293, 2023a.
- J. P. Dickerson, S. A. Esmaili, J. Morgenstern, and C. J. Zhang. Fair clustering: Critique, caveats, and future directions. Neurips 2023 Workshop on Algorithmic Fairness through the Lens of Time, 2023b. Full version arXiv:2406.15960 [cs.LG].
- M. Dummett. *Voting Procedures*. Oxford University Press, 1984.
- S. Ebadian and E. Micha. Boosting sortition via proportional representation. Technical report, arXiv:2406.00913 [cs.GT], 2024.
- S. Ebadian, G. Kehne, E. Micha, A. D. Procaccia, and N. Shah. Is sortition both representative and fair? In *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS)*, pages 3431–3443, 2022.
- A. Fanelli, G. Monaco, and L. Moscardelli. Relaxed core stability in fractional hedonic games. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 182–188, 2021.
- S. Fioravanti, M. Flammini, B. Kodric, and G. Varricchio. ϵ -fractional core stability in hedonic games. In *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS)*, pages 28723–28734, 2023.
- B. Flanagan, P. Gözl, A. Gupta, B. Hennig, and A. D. Procaccia. Fair algorithms for selecting citizens’ assemblies. *Nature*, 596:548–552, 2021.
- L. Han, D. Xu, Y. Xu, and P. Yang. Approximation algorithms for the individually fair k-center with outliers. *Journal of Global Optimization*, 87(2):603–618, 2023.
- Z. Jiang, K. Munagala, and K. Wang. Approximately stable committee selection. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 463–472. ACM, 2020.
- C. Jung, S. Kannan, and N. Lutz. Service in your neighborhood: Fairness in center location. In *Proceedings of the 1st Symposium on Foundations of Responsible Computing (FORC)*, pages 5:1–5:15, 2020.
- Y. H. Kalaycı, D. Kempe, and V. Kher. Proportional representation in metric spaces and low-distortion committee selection. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 9815–9823, 2024.
- D. Kar, M. Kosan, D. Mandal, S. Medya, A. Silva, P. Dey, and S. Sanyal. Feature-based individual fairness in k-clustering. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 2772–2774, 2023.
- L. Kellerhals and J. Peters. Proportional fairness in clustering: A social choice perspective. Technical report, arXiv:2310.18162 [cs.LG], 2023.
- M. Lackner and P. Skowron. *Multi-Winner Voting with Approval Preferences*. Springer, 2022.
- B. Li, L. Li, A. Sun, C. Wang, and Y. Wang. Approximate group fairness for clustering. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, pages 6381–6391, 2021.
- S. Mahabadi and A. Vakilian. Individual fairness for k-clustering. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pages 6586–6596, 2020.
- E. Micha and N. Shah. Proportionally fair clustering revisited. In *Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 85:1–85:16, 2020.
- R. Mosenzon and A. Vakilian. Scalable algorithms for individual preference stable clustering. In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1108–1116. PMLR, 2024.

- M. Negahbani and D. Chakrabarty. Better algorithms for individually fair k -clustering. In *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS)*, pages 13340–13351, 2021.
- D. Peters and P. Skowron. Proportionality and the limits of welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation (ACM-EC)*, pages 793–794. ACM, 2020.
- D. Peters, G. Pierczyński, and P. Skowron. Proportional participatory budgeting with additive utilities. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, pages 12726–12737, 2021.
- H. Sternbach and S. Cohen. Fair facility location for socially equitable representation. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 2775–2777, 2023.
- A. Vakilian and M. Yalçiner. Improved approximation algorithms for individually fair clustering. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 8758–8779, 2022.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: All claims including claimed approximation guarantees and connections are shown via proofs

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We discuss several lower bounds and incompatibilities throughout the paper. We do not really make any assumptions (except being given a metric space), run no experiments, and we discuss computational issues in Section 3.4

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We provide proofs (or references) for all theoretical results given.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification: Our work does not contain any experimental results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: We provide no experimental results and use no data or code.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: We provide no experimental results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: We provide no experimental results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: We provide no experimental results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We do not work with human subjects, nor create datasets. We do not see any overlap with the described potential ethics violations.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: This is a purely theoretical work. We do not think that any of the examples listed below fit to be discussed here.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We do not see any potential high risk.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: We do not use any assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We do not use any assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No experiments are performed in this paper.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: No experiments are performed in this paper.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.