# Proportionally Representative Clustering* 

Haris Aziz<br>UNSW Sydney<br>Sydney, Australia<br>Barton E. Lee<br>ETH Zürich, Switzerland<br>Sean Morota Chu<br>UNSW Sydney<br>Sydney, Australia<br>Jeremy Vollen<br>UNSW Sydney<br>Sydney, Australia


#### Abstract

In recent years, there has been a surge in effort to formalize notions of fairness in machine learning. We focus on clustering - one of the fundamental tasks in unsupervised machine learning. We propose a new axiom "proportional representation fairness" (PRF) that is designed for clustering problems where the selection of centroids reflects the distribution of data points and how tightly they are clustered together. Our fairness concept is not satisfied by existing fair clustering algorithms. We design efficient algorithms to achieve PRF both for unconstrained and discrete clustering problems. Our algorithm for the unconstrained setting is also the first known polynomial-time approximation algorithm for the well-studied Proportional Fairness (PF) axiom [17]. Our algorithm for the discrete setting also matches the best known approximation factor for PF.


## 1. Introduction

As artificial intelligence and machine learning serve as the 'new electricity' of systems, it is becoming increasingly important to ensure that AI systems embody societal values such as privacy, transparency, and fairness. Fairness is a growing concern, as AI systems are used to make decisions that can critically affect our daily lives, finances, and careers [20]. The impulse to tackle the issue of fairness in AI is prominently reflected in several governmental policy documents [see, e.g. 21, 3, 27], the emergence of dedicated conferences on Ethical AI, and the formation of AI ethics boards by various big tech companies. Over the past years, several developments have taken place in the theory and application of fairness in supervised learning. One prominent approach is to use information derived from labelled data in supervised learning to formalize some notions of equity within and across (pre-specified and) protected attributes [see, e.g. 25]. In contrast, in several models of unsupervised learning, protected attributes or their labelling are unknown. Fairness in unsupervised machine learning is becoming important with the growth of data collection both in public (e.g., via IoT devices) and private sectors (e.g., map services and social media). Unfettered and irresponsible use of such data has the danger of increasing inequity [40.

We focus on fairness in clustering, which is one of the most widely applied tasks in unsupervised machine learning. A (centroid selection) clustering problem has a metric space $\mathscr{X}$ with a distance measure $d$ :

[^0]$\mathscr{X} \times \mathscr{X} \rightarrow \mathbb{R}^{+} \cup\{0\}$, a multiset $\mathscr{N} \subseteq \mathscr{X}$ of $n$ data points (agents), a set $\mathscr{M}$ of candidate centers, and positive integer $k \leq n$. In many clustering problems, each data point corresponds to an individual's attribute profile; thus, we refer to the data points as agents. The goal is to find a size- $k$ subset of centers, $Y \subseteq \mathscr{M}:|Y|=k$. Such centroid selection problems can be widely applied to various applications including recommender systems and data-compression. The problem also captures social planning scenarios where $k$ facility locations need to be chosen to serve a set of agents.

In such a (centroid selection) clustering problem, there are two versions of the problem. In the discrete version, the set of candidate centers $\mathscr{M}$ is a finite subset of $\mathscr{X}$. In the unconstrained or continuous setting, $\mathscr{M}=\mathscr{X}$. We first focus on the unconstrained setting as done by Micha and Shah [38]. Later we discuss how our concepts and ideas extend to the discrete setting. For a location $i$ and set of locations $S \subseteq \mathscr{M}$, we will denote $\min _{c \in S} d(i, c)$ as $d(i, S)$. Standard clustering solutions include $k$-center, $k$-medians, and $k$-means $\underbrace{1}$ These standard objectives can be viewed as achieving some form of global welfare but are not well-suited for proportional representational and fairness concerns. The problems of computing outcomes optimising these objectives are also NP-hard [see, e.g. 36, 5].

We seek a suitable concept for clustering that captures general fairness principles such as nondiscrimination, equality of opportunity, and equality of outcome. In many applications, the data points correspond to real individuals or their positions or opinions. Such individuals may expect the clustering outcome to capture natural fairness requirements. Previous axiomatic studies of clustering have measured fair representation by the distance from each data point to the nearest centroid. We initiate a proportional representation perspective on fairness in clustering. In capturing this desiderata axiomatically, we require a property in which representation guarantees depend on the number (or proportion) of data points and how "tightly" they are clustered together.

Our perspective is motivated by applications in which groups of individuals (to which the data points correspond) may expect to receive centroid representation proportional to their size. For instance, the centroids may be citizens selected to make representative decisions on behalf of the public, as is done in a jury or a sortition panel. Similarly, computing a proportionally fair clustering is also meaningful in facility location, where the location of the facility should depend on the density and the number of people it is serving. More generally, our axioms and algorithms are well-motivated in any setting in which one may want to take a representative sample from datapoints between which a distance metric is well-defined. In this paper, we consider the following fundamental research question. What is a meaningful concept of proportional fairness in centroid clustering? Is such a concept guaranteed to be achievable?

Contributions.. In contrast to traditional optimization objectives in clustering and centroid selection, we take an axiomatic fairness perspective. We first identify a potential shortcoming of some recently introduced concepts that aim to capture proportional fairness for clustering; in particular, we show how the concepts do not capture an intuitive and minimal requirement of proportional representation called unanimous proportionality. We then propose a new axiom called proportional representative fairness (PRF), which is our central conceptual contribution.

We propose PRF for both the unconstrained and discrete clustering setting. We discuss how PRF overcomes some of the drawbacks of the previously introduced concepts. In particular, it implies the minimal requirement of unanimous proportionality. PRF has several other desirable features: a PRF outcome is guaranteed to exist, the axiom does not require a preset specification of protected groups and it is robust to outliers and multiplicative scaling. We show that none of the traditional clustering algorithms or new algorithms for fair clustering satisfy PRF.

We design polynomial-time algorithms to achieve PRF both in the unconstrained and discrete settings $2^{2}$ We prove that our algorithm for the unconstrained setting also gives 3 -approximate Proportional Fairness $(P F)$ 17], and hence is the first known polynomial-time approximation algorithm for PF on general metric spaces. We also prove that our algorithm for the discrete setting matches the best known approximation

[^1]|  | Unconstrained |  |  | Discrete |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | $\rho$-PF | PRF | Complex. | $\rho$-PF | PRF | Complex. |
| Our Algorithms | 3 | $\checkmark$ | in P | $\approx 2.4$ | $\checkmark$ | in P |
| Greedy Capture | $\approx 2.4$ | - | NP-hard | $\approx 2.4$ | - | in P |

Table 1: Theoretical comparison of our algorithms (Algorithm 1 for the unconstrained setting and Algorithm 2 for the discrete setting) with Greedy Capture, which gives the best known approximate PF guarantees. Unconstrained and discrete Greedy Capture results come from Micha and Shah [38] and Chen et al. [17], respectively. Note that Micha and Shah 38 provide a $(2+\epsilon)$-PF PTAS, but this result does not hold for general metric spaces.
of Proportional Fairness (PF), thereby providing no compromise on PF approximation but additionally satisfying PRF. We summarize our theoretical contributions in Table 1 via a comparison of our algorithms with the best guarantees previously known.

Finally, we experimentally compare our algorithm with two standard algorithms and show that it performs especially well when the mean squared distance of the agents is calculated based on their distance from not only the closest center but also the next $j$-closest centers (for large $j$ ).

## 2. Related Work

Fairness in machine learning and computer systems is a vast topic [32]-for an overview, see Chouldechova and Roth [20] or the survey by Mehrabi et al. [37]. We focus on the fundamental learning problem of clustering, which is a well-studied problem in computer science, statistics, and operations research [see, e.g. 28]. Our approach is axiomatic, focused on the concept of proportional fairness (or representation), and does not rely on the use of a pre-specified or protected set of attributes. Thus, the most closely related papers are those of Chen et al. [17], Micha and Shah [38], Li et al. [35], and Jung et al. [33], which we review below.

Chen et al. [17] proposed the concept of proportional fairness (PF) for the discrete setting, which is based on the entitlement of groups of agents of large enough size. They reasoned that PF is a desirable fairness concept but it is not guaranteed to be achievable for every instance. On the other hand, they showed that outcomes satisfying reasonable approximations of PF are guaranteed to exist and can be computed in polynomial time via the 'Greedy Capture' algorithm.

Micha and Shah [38] analyzed the proportional fairness axiom in the unconstrained setting and presented similar results to those obtained by Chen et al. [17. They showed that the natural adaptation of the Greedy Capture algorithm finds an approximately PF outcome in Euclidean space. On the other hand, they showed that checking whether a PF outcome exists is NP-hard and that the clustering returned by the 'greedy capture' algorithm cannot be computed in polynomial time unless $\mathrm{P}=\mathrm{NP}$. Li et al. 35 consider the same discrete clustering setting as Chen et al. 17] and propose a stronger version of proportional fairness called core fairness. They present results on lower and upper approximation bounds for various restrictions of the discrete setting.

Jung et al. [33] consider the discrete clustering problem in which the set of agents and candidate centers coincide. For this setting, they propose a fairness concept that has a similar goal as that of Chen et al. [17. In subsequent work, it has been referred to as individual fairness [48, 15]. Individual fairness requires that, for each agent $i$, there is a selected center that is at most $r(i)$ from $i$, where $r(i)$ is the smallest radius around $i$ that includes $n / k$ agents.

There are also several concepts and algorithms in clustering where the protected groups are pre-specified based on protected attributes [1, 2, 12, 11, 34, 18, 19, 26]. In contrast, we are seeking fairness properties that apply to arbitrarily defined groups of all possible sizes. In our proposed axiom, these groups are endogenously determined from data points; in some real-world settings, this is a legal requirement since it can be forbidden by law to have decision-making processes use certain attributes as inputs.

Fairness and representation are important concerns in economic design [39. Issues related to proportional representation and fairness have been considered in many settings, including apportionment (46, 14, participatory budgeting [9, 41, portioning [7, 4], probabilistic voting [13], and multi-winner voting [29, 8, 10, 23, 24, 47. This literature has largely focused on settings where agents have ordinal preferences over candidates. Our fairness (or proportional representation) axioms take into account the additional distance-based cardinal information available when the candidates belong to a general metric space. Hence, our axioms formalize
a notion of proportional representation for $m$-dimensional Euclidean space for $m>2$. The standard fairness axioms such as proportionality of solid coalitions for ordinal preferences [8] are too weak when applied to our spatial context. Naively applying them to our spatial problem also results in a reduction from our unconstrained problem to a multiwinner voting problem with an infinite number of candidates.

The clustering problem especially has natural connections with multi-winner voting [29, 30, 43]. Agents and centers in the clustering context correspond to voters and candidates in multi-winner voting. Each agent's (voter's) preference over centers (candidates) can then be derived using the distance metric with closer centers being more preferred. With this view, multi-winner voting rules imply clustering solutions. For example, the Chamberlin-Courant rule takes agents' distances from candidates using Borda count and then implements the $k$-median objective. There is also work on proportional representation in muti-winner voting and more general settings with cardinal utilities 41 but the axioms do not necessarily and suitably capture distance-based concerns. Relatedly, clustering can also be viewed as a facility location problem 31, where multiple facilities (centers) are to be located in a metric space; for a recent overview, see the survey by Chan et al. [16].

## 3. Towards an Axiom Capturing Proportional Representation

Proportional fairness for clustering has been considered in a series of recent papers. Chen et al. 17] first proposed a proportional fairness axiom that requires that there is no set of agents of size at least $n / k$ that can find an unselected center that is "better" (i.e., located closer than the closest selected center) for each of them.

Definition 1 (Proportional Fairness [17]). $X \subseteq \mathscr{M}$ with $|X|=k$ satisfies proportional fairness if $\forall S \subseteq \mathscr{N}$ with $|S| \geq\left\lceil\frac{n}{k}\right\rceil$ and for all $c \in \mathscr{M}$, there exists $i \in S$ with $d(i, c) \geq d(i, X)$.

The idea of proportional fairness was further strengthened to core fairness that requires that there is no set of agents of size at least $n / k$ that can demand a more preferred location for a center [35]. One rationale for proportional fairness is explained in the following example, which was first presented by Micha and Shah [38] and then reused by Li et al. [35]. We reuse the same example. Consider a set of data points/agents. The agents are divided into 11 subsets of clusters each of which is densely clustered. One cluster of agents has size 10,000 . The other 10 clusters have sizes 100 each. The big cluster is very far from the smaller clusters. The small clusters are relatively near to each other. Micha and Shah [38] and Li et al. [35] are of the view that any set of centers that satisfies a reasonable notion of proportional fairness would place 10 centers for the big cluster and 1 center serving the small clusters. Next, we point out that proportional fairness and core fairness do not necessarily satisfy this minimal requirement, which we formalize via the unanimous proportionality axiom below.

Definition 2 (Unanimous Proportionality (UP)). $X \subseteq \mathscr{M}$ with $|X|=k$ satisfies unanimous proportionality if $\forall S \subseteq \mathscr{N}$ with $|S| \geq \ell\left\lceil\frac{n}{k}\right\rceil$ and each agent in $S$ has the same location $x \in \mathscr{X}$, then $\ell$ possible locations closest to $x$ are selected.

UP captures a principle which Chen et al. [17] highlighted as desirable - in their words, one which requires that "any subset of at least $r\lceil n / k\rceil$ individuals is entitled to choose $r$ centers. ${ }^{3}$ The following example shows that proportional fairness [17] does not imply Unanimous Proportionality. Similarly, it can be shown that the fairness concept proposed by Jung et al. [33] does not imply unanimous fairness.

Example 1 (Proportional fairness and core fairness). Suppose agents are located in the $[0,1]$ interval and $k=11$. 10,000 agents are at point 0 and 1000 agents are at point 1. UP requires that 10 centers are at or just around point 0 and 1 center is at point 1. However, placing 1 center at point 0 and 10 centers at point 1 satisfies the proportional fairness concept of Chen et al. 17] and also the core fairness concept of Li et al. [35].

[^2]The example above shows that there is a need to formalize new concepts in order to capture proportional representative fairness. The property should imply unanimous proportionality but should also be meaningful if no two agents/data points completely coincide in their location. One reason why the existing fairness concepts do not capture UP is that they make an implicit assumption that an agent only cares about the distance from the nearest center and not on how many centers are nearby. As discussed in the introduction, this assumption is not necessarily accurate in contexts such as data abstraction, sortition, and political representation.

In addition to capturing proportional representation in a robust sense, we would also like our axiom (and algorithms) to guarantee the existing proportional fairness concepts, at least in an approximate fashion. Since a PF outcome may not exist, Chen et al. [17] and Micha and Shah [38] studied an approximate notion of proportional fairness in the discrete and unconstrained settings, respectively.
Definition 3 ( $\rho$-approximate Proportional Fairness). $X \subseteq \mathscr{M}$ with $|X|=k$ satisfies $\rho$-approximate PF (or $\rho-P F)$ if $\forall S \subseteq N$ with $|S| \geq\left\lceil\frac{n}{k}\right\rceil$ and for all $c \in \mathscr{M}$, there exists $i \in S$ with $\rho \cdot d(i, c) \geq d(i, X)$.

We seek to devise an axiom which is compatible with approximate PF , and subsequently to design algorithms which satisfy our axiom while also providing good approximate guarantees with respect to PF. We close the section with an example presented by Micha and Shah 38 that shows that a PF outcome may not exist. We will use the same example to illustrate our new fairness concept in the next section.
Example 2. Consider Figure 1. For the 6 agents and $k=3$, it has been shown by Micha and Shah [38] that a PF outcome may not exist.

| 1 |  | 4 |
| :--- | :--- | :--- | :--- |
| 2 | 3 |  | | 6 | 5 |
| :--- | :--- |

Figure 1: An example instance with 6 agents and $k=3$ for which a PF outcome does not exist.

## 4. Fairness for the Unconstrained Setting

We propose a new concept Proportionally Representative Fairness (PRF) that captures proportional representation fairness concerns. The intuition behind the concept is based on two natural principles: (1) a set of agents that is large enough deserves to have a proportional number of centers 'near' it, and (2) the requirement of proximity of the nearby centers to a subset of agents should also depend on how densely together the particular subset of agents are. We use these two principles to propose a new fairness concept.

Definition 4 (Proportionally Representative Fairness (PRF) for Unconstrained Clustering). An outcome $X \subseteq \mathscr{M}$ with $|X|=k$ satisfies Proportionally Representative Fairness (PRF) if the following holds. For any set of agents $S \subseteq \mathscr{N}$ of size at least $|S| \geq \ell n / k$ such that the maximum distance between pairs of agents in $S$ is $y$, the following holds: $\mid c \in X: \exists i \in S$ s.t $d(i, c) \leq y \mid \geq \ell$.

PRF has the following features: it implies unanimous proportionality; it is oblivious to and does not depend on the specification of protected groups or sensitive attributes; it is robust to outliers in the data, since the fairness requirements are for groups of points that are sizable enough; and it is scale invariant to multiplication of distances (i.e., PRF outcomes are preserved if all the distances are multiplied by a constant).

Next, let us illustrate an example where proportional fairness is not achievable but one of the most natural clustering outcomes satisfies PRF.
Example 3 (Requirements of PRF). We revisit Example 2 to illustrate the requirements that PRF imposes. For each set $S$ of size at least $\ell\left\lceil\frac{n}{k}\right\rceil=2 \ell, P R F$ requires $\ell$ centers in the relevant neighborhood of agents in $S$ (see Figure 2). Therefore, for agent 1 and 2, PRF requires that one selected candidate center should be in one of the circles around 1 or 2. For the set of agents $\{1,2,3\}$, PRF also requires that one candidate center should be in one of the circles around 1, 2 or 3. A natural solution is to have one center located between agents 1, 2, and 3, another center located between agents 4, 5, and 6, and the final center located somewhere not too far from all agents (e.g., between the two groups of agents). Such solutions are intuitively reasonable and fair and, furthermore, satisfy PRF; however, as mentioned in Example 2, this instance does not admit a proportionally fair solution.


Figure 2: Some of the requirements of PRF for instance in Example 2

In addition to capturing natural outcomes when PF outcomes may not exist, PRF actually guarantees a $\rho$-PF committee for $\rho \approx 3.6$ in the unconstrained setting on general metric spaces.

Proposition 1. For metric spaces, if an outcome $X \subseteq \mathscr{M}$ with $|X|=k$ satisfies PRF for Unconstrained Clustering, then $X$ is $\frac{3+\sqrt{17}}{2}$-approximate $P F$, and there exists an instance for which this bound is tight.

Next, we highlight that the well-known $k$-means or $k$-median solutions do not necessarily satisfy PRF. Example 4 is adapted from an example by Chen et al. 17.

Example 4 ( $k$-means does not satisfy PRF). Consider Figure 3 in which there are $n / 3$ agents uniformly distributed on the perimeter of each of the three circles. If $k=3$, a $k$-means or $k$-median solution will place one center by the small circles and two centers in the big circle. In contrast, PRF requires that each of the circles gets its own centroid. It can also be shown that $k$-center does not satisfy PRF.


Figure 3: The k-means solution may not satisfy PRF.

On face value, it is not clear whether an outcome satisfying PRF exists as similar concepts such as PF cannot always be guaranteed [38. Secondly, even if a PRF outcome is guaranteed to exist, the complex constraints seem computationally challenging. PRF is a property that requires conditions on an exponential number of subsets of agents. For each subset of agents, it enforces representational constraints pertaining to an infinite number of neighborhood distances. Perhaps surprisingly, we show that not only is a PRF outcome guaranteed to exist, it can be computed in polynomial time.

The algorithm intuitively works as follows. Firstly, instead of considering infinite possible centers, we restrict our attention to the $n$ candidates centers that coincide with each agent's location. Each agent is given an initial weight of 1 . This weight dynamically decreases over the course of the algorithm. The algorithm can be viewed as gradually expanding the neighborhoods of each of the agents ${ }^{4}$ Instead of continuously expanding the neighborhood, we only consider at most $n^{2}$ possible values of neighborhood radii. If there exists some location in the candidate set such that it is in the intersection of the neighborhoods of a group of agents with total weight of at least $n / k$, we select one such candidate, and reduce the agents' collective weight by $n / k$. The process is continued until $k$ centers are chosen. The algorithm is formally specified as Algorithm 1. The following is the main result of the present section and summarizes the theoretical guarantees of Algorithm 1

Theorem 1. Algorithm 1 terminates in polynomial time $O\left(n^{4} k^{2}\right)$ and returns a set of $k$ centers which satisfy PRF and 3-approximate PF.

We will now formally prove two key lemmas used to prove Theorem 1
Lemma 1. Algorithm 1 returns an outcome that satisfies PRF.

[^3]```
Input: metric space \(\mathscr{X}\) with a distance measure \(d: \mathscr{X} \times \mathscr{X} \rightarrow \mathbb{R}^{+} \cup\{0\}\), a finite multiset \(\mathscr{N} \subseteq \mathscr{X}\) of \(n\) data
points (agents), and positive integer \(k\).
Output: A multiset of \(k\) centers
\(w_{i} \longleftarrow 1\) for each \(i \in \mathscr{N}\)
Consider the set \(D=\left\{d\left(i, i^{\prime}\right) \mid i, i^{\prime} \in N\right\}\). Order the entries in \(D\) as \(d_{1} \leq d_{2} \leq \cdots d_{|D|}\).
\(\mathscr{M} \longleftarrow \mathscr{N}\)
\(j \longleftarrow 1\)
\(W \longleftarrow \emptyset\)
while \(|W|<k\) do
    \(C^{*}=\left\{c \in \mathscr{M} \mid \sum_{i \in \mathscr{N} \mid d(i, c) \leq d_{j}} w_{i} \geq n / k\right\}\)
    if \(C^{*}=\emptyset\) then
        \(j \longleftarrow j+1\)
    else
        Select some candidate \(c^{*}\) from \(C^{*}\) such that \(c^{*}=\arg \max _{c^{\prime} \in C^{*}} \sum_{i \in \mathscr{N} \mid d\left(i, c^{\prime}\right) \leq d_{j}} w_{i}\) :
        \(W \longleftarrow W \cup\left\{c^{*}\right\} ; \mathscr{M} \longleftarrow \mathscr{M} \backslash\left\{c^{*}\right\}\)
        \(N^{\prime} \longleftarrow\left\{i \in \mathscr{N}: d\left(i, c^{*}\right) \leq d_{j}\right\}\)
        Modify the weights of voters in \(N^{\prime}\) so the total weight of voters in \(N^{\prime}\), i.e., \(\sum_{i \in N^{\prime}} w_{i}\), decreases by exactly
        \(n / k\).
Return \(W\)
```

Algorithm 1: Algorithm for Unconstrained Clustering

Proof. Suppose for contradiction that the algorithm finds a set of $k$ centers $W$ that violates PRF. In that case, there exist a set of agents $S \subseteq \mathscr{N}$ of size at least $|S| \geq \ell n / k$ such that the maximum distance between pairs of agents in $S$ is $y$ but the number of locations that are within $y$ of some agent in $S$ is at most $\ell-1$. In that case, consider the snapshot of the algorithm when neighborhood distance $y=\max _{i, j \in S} d(i, j)$ is considered. At this point, $\mid c \in X: \exists i \in S$ s.t $d(i, c) \leq y \mid \leq \ell-1$. Since agents in $\mathscr{N}$ have only used their weight towards selecting locations within $y$ up to this point, it follows that $\sum_{i \in S} w_{i} \geq \ell \frac{n}{k}-(\ell-1) \frac{n}{k}=\frac{n}{k}$. Therefore, the agents in $S$ still have total weight at least $\frac{n}{k}$ to select one more location that is a distance of at most $y$ from some agent in $S$. Hence, at this stage, when the neighborhood distance is $y$, the algorithm would have selected at least one more center within distance $y$ than in the outcome $W$.

Lemma 1 implies guaranteed existence of an outcome satisfying our axiom, which was one of our goals in devising PRF. In conjunction with Proposition 1, the lemma also gives us that the outcome returned by Algorithm 1 is $\frac{1}{2}(3+\sqrt{17})$-PF. This is noteworthy as it establishes Algorithm 1 as the first known polynomialtime approximation algorithm for PF on general metric spaces in the unconstrained setting. As we will show next, Algorithm 1 actually provides the stronger guarantee of 3-PF.

Lemma 2. Algorithm 1 finds an outcome that satisfies 3 -approximate $P F$.
Proof. For an instance of unconstrained clustering, let $X$ be the outcome returned by Algorithm 1 . Fix $S \subseteq N, c \in \mathscr{M}$ with $|S| \geq\left\lceil\frac{n}{k}\right\rceil$. We will show that there exists an agent $i \in S$ and center $x \in X$ such that $d(i, x) \leq 3 d(i, c)$. We assume that $S \cap X=\emptyset$, since otherwise $d(i, x)=0$ for some agent and center, and we are done.

At the start of the algorithm, the collective weight of agents in $S$ is $|S|$. When the algorithm terminates, the collective weight of $S$ is zero. Thus, in some iteration, the collective weight of $S$ decreases for the first time. We denote the centroid selected in this iteration by $x^{*}$ and denote an arbitrary agent whose weight is decreased in this iteration by $i^{*}$. Let $y$ be the maximum distance between $i^{*}$ and any other agent in $S$.

We first point out that $y \geq d\left(i^{*}, x^{*}\right)$. To see this, suppose for contradiction $y<d\left(i^{*}, x^{*}\right)$. Since $y$ and $d\left(i^{*}, x^{*}\right)$ are both distances between pairs of agents in $N, y=d_{j}$ and $d\left(i^{*}, x^{*}\right)=d_{j^{\prime}}$ for some $j$ and $j^{\prime}$. Furthermore, $j<j^{\prime}$, since $D$ is in increasing order and by assumption that $y<d\left(i^{*}, x^{*}\right)$. Thus, when the algorithm considers neighborhoods of $d_{j}=y$, the collective weight of $S$ has not yet decreased, and we have that

$$
\sum_{i \in \mathscr{N} \mid d\left(i, i^{*}\right) \leq y} w_{i} \geq|S| \geq\left\lceil\frac{n}{k}\right\rceil \geq \frac{n}{k}
$$

Therefore, before transitioning to the next element in $D$, the algorithm will either add $i^{*}$ to $X$ or the weight of some agent in $S$ will decrease. The first possibility contradicts $S \cap X=\emptyset$. For the second possibility, by
assumption, it must be that $x^{*}$ was added to $X$. But, since $d\left(i^{*}, x^{*}\right)>y, i^{*} \notin N^{\prime}$ and $w_{i^{*}}$ does not decrease, which gives us a contradiction.

We divide the remainder of the argument into cases. First, consider the case in which $d\left(i^{*}, c\right) \geq \frac{y}{3}$. It follows that $d\left(i^{*}, x^{*}\right) \leq y \leq 3 d\left(i^{*}, c\right)$. If, instead, $d\left(i^{*}, c\right)<\frac{y}{3}$, then consider the agent $j \in S \backslash\{i\}$ such that $d\left(i^{*}, j\right)=y$. By the triangle inequality, we have that

$$
\begin{aligned}
d(j, c) & \geq y-d\left(i^{*}, c\right) \geq \frac{2 y}{3} \\
& \geq \frac{1}{3}\left(d\left(j, i^{*}\right)+d\left(i^{*}, x^{*}\right)\right) \geq \frac{1}{3} d\left(j, x^{*}\right) .
\end{aligned}
$$

## 5. Fairness for the Discrete Setting

We presented PRF as a desirable concept for the unconstrained setting. It is not immediate how to adapt the concept to the discrete setting. Applying the unconstrained PRF definition to the discrete setting leads to the issues of non-existence. On the other hand, some natural attempts to account for the coarseness of the candidate space, $\mathscr{M}$, can lead to a concept that is too weak. For example, one alternative is to require that an outcome $X$ be such that if, for any set of agents $S \subseteq \mathscr{N}$ of size at least $|S| \geq \ell n / k$ such that the maximum distance between pairs of agents in $S$ is $y$, the following holds: $\mid c \in X: \exists i \in S$ s.t $d(i, c) \leq y \mid \geq$ $\min \left\{\ell,\left|\cup_{j \in S} B_{y}(j) \cap \mathscr{M}\right|\right\}$, where $B_{r}(i)$ denote a ball of radius $r$ around agent $i$. Alternatively, we could replace the right-hand side of the final inequality with $\min \left\{\ell,\left|\cap_{i \in S} B_{y}(i) \cap \mathscr{M}\right|\right\}$. Both of these versions are too weak. To see this, suppose $k=1$, all agents are located at 0 , and $\mathscr{M}=(\ldots,-2,-1,1,2,3, \ldots)$, then neither version places any restriction on the facility location and do not even imply UP.

To resolve these issues, we need to take into account that the nearest candidate locations may be very far from a subset of agents. We adapt our PRF concept for the discrete setting via a careful adaptation of the PRF concept for the unconstrained setting.

Definition 5 (Proportionally Representative Fairness (PRF) for Discrete Clustering). For any set of agents $S \subseteq \mathscr{N}$ of size at least $|S| \geq \ell n / k$, if there are $\ell^{\prime} \leq \ell$ candidates from $\mathscr{M}$ that are at distance at most $y$ from all agents in $S$, the following holds: $\mid c \in X: \exists i \in S$ s.t $d(i, c) \leq y \mid \geq \ell^{\prime}$.

We begin by highlighting the fact that, in the discrete setting, PRF implies $(1+\sqrt{2})$-approximate PF, the closest PF approximation known to exist in the discrete setting [17]. Our proof follows similarly to Theorem 1 from Chen et al. [17].

Proposition 2. For metric spaces, if an outcome $X \subseteq \mathscr{M}$ with $|X|=k$ satisfies PRF for Discrete Clustering, then $X$ is $(1+\sqrt{2})$-approximate $P F$.

For the discrete setting, we propose a new algorithm (Algorithm 2). Algorithm 2 is similar to Algorithm 1 , which was designed for the unconstrained setting. In fact, Algorithm 1 can be viewed as first setting $\mathscr{M}$ to the multiset of candidate locations corresponding to agents in $\mathscr{N}$ and then running Algorithm 2, Similar to Algorithm 1, Algorithm 2 terminates in polynomial time and returns $k$ centers and its output satisfies PRF.

Theorem 2. Algorithm 1 terminates in polynomial time $O\left(|\mathscr{M}|^{2} n^{2}\right)$ and returns a set of $k$ centers which satisfy PRF and $(1+\sqrt{2})$-approximate $P F$.

Notably, by satisfying PRF, Algorithm 2 matches the best known PF approximation factor in the discrete setting. As the following remark highlights, however, the existing algorithms do not capture PRF.

Remark 1. Chen et al. 17] present Greedy Capture. An equivalent algorithm, called $A L G_{g}$, is also presented by Li et al. [35]. However, these algorithms do not satisfy PRF. In fact, for given $k$, these algorithms may fail to output $k$ candidate locations ${ }^{5}$

[^4]```
Input: metric space \(\mathscr{X}\) with a distance measure \(d: \mathscr{X} \times \mathscr{X} \rightarrow \mathbb{R}^{+} \cup\{0\}\), a finite multiset \(\mathscr{N} \subseteq \mathscr{X}\) of \(n\) data
points (agents), a finite set of candidate locations \(\mathscr{M}\), and positive integer \(k\).
Output: A multiset of \(k\) centers
\(w_{i} \longleftarrow 1\) for each \(i \in \mathscr{N}\)
Consider the set \(D=\{d(i, c) \mid i \in \mathscr{N}, c \in \mathscr{M}\}\). Order the entries in \(D\) as \(d_{1} \leq d_{2} \leq \cdots \leq d_{|D|}\).
\(j \longleftarrow 1\)
\(W \longleftarrow \emptyset\)
while \(|W|<k\) do
    \(C^{*}=\left\{c \in \mathscr{M} \mid \sum_{i \in \mathscr{N} \mid d(i, c) \leq d_{j}} w_{i} \geq n / k\right\}\)
    if \(C^{*}=\emptyset\) then
        \(j \longleftarrow j+1\)
    else
        Select some candidate \(c^{*}\) from \(C^{*}\) such that \(c^{*}=\arg \max _{c^{\prime} \in C^{*}} \sum_{i \in \mathscr{N} \mid d\left(i, c^{\prime}\right) \leq d_{j}} w_{i}\) :
        \(W \longleftarrow W \cup\left\{c^{*}\right\} ; \mathscr{M} \longleftarrow \mathscr{M} \backslash\left\{c^{*}\right\}\)
        \(N^{\prime} \longleftarrow\left\{i \in \mathscr{N}: d\left(i, c^{*}\right) \leq d_{j}\right\}\)
        Modify the weights of voters in \(N^{\prime}\) so the total weight of voters in \(N^{\prime}\), i.e., \(\sum_{i \in N^{\prime}} w_{i}\), decreases by exactly
        \(n / k\).
Return \(W\)
```

Algorithm 2: Algorithm for Discrete Clustering

|  | MSD to 1 |  |  | MSD to $k / 2$ |  |  | MSD to $k$ |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Dataset | $k++$ Average MSD | Alg 1 <br> Difference in MSD relative to $k++$ | $A L G_{g}$ <br> Difference in MSD <br> relative to $k++$ | $k++$ Average MSD | Alg 1 <br> Difference in MSD <br> relative to $k++$ | $A L G_{g}$ <br> Difference in MSD <br> relative to $k++$ | $k++$ Average MSD | Alg 1 <br> Difference in MSD <br> relative to $k++$ | $A L G_{g}$ <br> Difference in MSD relative to $k++$ |
| Wholesale | 34,977,392.00 | 664\% | 112\% | 11,668,107,198.00 | -10\% | -12\% | 100,307,739,343.00 | -63\% | -55\% |
| HCV | 811.60 | 519\% | 429\% | 582,360.00 | -4\% | -6\% | 4,858,449.00 | -70\% | -62\% |
| Buddy | 782.20 | 8\% | 5\% | 181,700.00 | -8\% | -7\% | 784,360.00 | -13\% | -9\% |
| Seeds | 0.54 | 23\% | 8\% | 228.68 | -1\% | -1\% | 1,449.94 | -4\% | -1\% |

Table 2: Performance of clustering algorithms with respect to MSD to $1, \mathrm{k} / 2$, and k centroids.

In the Appendix, we consider two natural and stronger notions of PRF in the discrete setting (that also apply to the continuous setting). We will show that the versions that we consider do not guarantee the existence of a fair outcome in all settings.

## 6. Experiment

We now apply Algorithm 1 to four real-world datasets. We consider the discrete domain in which $\mathscr{M}=\mathscr{N}$ as most clustering data sets only have the data points. Algorithm 1 is guaranteed to satisfy PRF; therefore, we are interested in analyzing the performance of the algorithm with respect to other objectives. A prominent objective for clustering algorithms is to minimize the Mean Squared Distance (MSD) to the closest center: the average distance between each datapoint and its closest center. We consider this performance measure as the number of centers ranges from $k=1$ to 100 .

In addition, to MSD to the closest center, we analyze two other related measures. MSD to the closest $k / 2$ centers and MSD to the closest $k$ centers. Intuitively, MSD to the closest $k / 2$ centers is the average distance between each datapoint and its $k / 2$-closest centers (MSD to the closest $k$ centers is similar). These measures capture - to varying extents - the idea that it may be desirable to have datapoints located close to more than just one center; this idea is at the core of the proportional representation concept. To benchmark Algorithm 1 1 s performance, we also implement two other algorithms: the well-known $k$-means++ algorithm ${ }^{6}$ of Arthur and Vassilvitskii [6] and $\mathrm{ALG}_{g}$ of Li et al. [35]. For all of our MSD measures, a smaller value is more desirable.

Datasets. The four datasets that we analyze are from the UCI Machine Learning Repository [22]. We

[^5]summarize these below ${ }^{7}$ HCV dataset (contains information about patients with Hepatitis C and blood donors.) ; Seeds dataset (describes measurements of geometrical properties of kernels belonging to three different varieties of wheat.); Buddy-move dataset (consists of user information from user reviews published in www.holidayiq.com about point of interests in South India); and Wholesale dataset (consists of client information for a wholesale distributor. It includes annual spending on various product categories.)

Results. Table 2 summarizes our results. We illustrate our analysis in figures in the Appendix. Across all the datasets, Algorithm 11 s performance relative to $k$-means improves when moving from MSD to closest center to MSD to closest $k / 2$ centers and MSD to closest $k$ centers. A similar pattern is observed for $A L G_{g}$, which may be expected and is reassuring since - like Algorithm $1-A L G_{g}$ is motivated by the idea of proportional fairness. For two of the datasets (Wholesale and HCV), Algorithm 1 offers substantially better performance compared to $k$-means with respect to MSD to closest $k$ centers. In these cases, Algorithm 1 also outperforms $A L G_{g}$. For MSD to the closest $k / 2$ centers, $A L G_{g}$ and Algorithm 1 either outperform or perform similar to $k$-means. In contrast, and as expected, the $k$-means outperforms both algorithms with respect to the MSD to closest center. With this measure, $A L G_{g}$ typically outperforms Algorithm 1 . The experiments illustrate that other well-known algorithms produce different outcomes compared to our Algorithm 1, and this difference can be quantified via an intuitive PRF-related metric, i.e., minimizing the MSD to multiple centers.

## 7. Discussion

We revisited a classical AI problem with a novel solution concept that builds on ideas from social choice and fair division. We proposed a natural fairness concept called PRF for clustering. PRF is guaranteed to exist and has desirable features: it implies unanimous proportionality, does not require the specification of protected groups, and is robust to outliers and multiplicative scaling. Even if our only focus is on the PF axiom, we make significant progress on efficient algorithms for approximate PF for general metric spaces. We did not focus on strategic aspects. Unfortunately, PRF and PF are incompatible in general with the incentive of an agent to report their location truthfully (see Appendix). A natural next step is to identify stronger versions of PRF that still guarantee existence. It will be interesting to consider refinements of the PRF algorithm that make more clever choices when selecting candidates with sufficient weighted support. Another direction for future research is to understand how much impact the constraint of PRF imposes on standard optimization objectives such as for $k$-means, $k$-median, or $k$-center.

[^6]
## References

[1] Abbasi, M., Bhaskara, A., Venkatasubramanian, S., 2021. Fair clustering via equitable group representations. In: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. FAccT '21. Association for Computing Machinery, New York, NY, USA, pp. 504-514.
[2] Ahmadian, S., Epasto, A., Kumar, R., Mahdian, M., 2019. Clustering without over-representation. In: Proceedings of the 25 th ACM SIGKDD International Conference on Knowledge Discovery \& Data Mining. ACM, pp. 267-275.
[3] AI-Committee, 2019. The national artificial intelligence research and development strategic plan: 2019 update. Tech. rep., National Science and Technology Council, Washington, D.C.
[4] Airiau, S., Aziz, H., Caragiannis, I., Kruger, J., Lang, J., Peters, D., 2019. Portioning using ordinal preferences: Fairness and efficiency. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI). IJCAI.
[5] Aloise, D., Deshpande, A., Hansen, P., Popat, P., 2009. Np-hardness of euclidean sum-of-squares clustering. Machine Learning 75 (2), 245-248.
[6] Arthur, D., Vassilvitskii, S., 2007. k-means++: the advantages of careful seeding. In: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, New Orleans, Louisiana, USA, January 7-9, 2007. SIAM, pp. 1027-1035.
[7] Aziz, H., Bogomolnaia, A., Moulin, H., 2019. Fair mixing: the case of dichotomous preferences. In: Proceedings of the 20th ACM Conference on Electronic Commerce (ACM-EC). pp. 753-781.
[8] Aziz, H., Lee, B., 2020. The expanding approvals rule: Improving proportional representation and monotonicity. Social Choice and Welfare 54 (1), 1-45.
[9] Aziz, H., Lee, B. E., 2021. Proportionally representative participatory budgeting with ordinal preferences. In: Proceedings of the 35 th AAAI Conference on Artificial Intelligence (AAAI).
[10] Aziz, H., Lee, B. E., 2022. A characterization of proportionally representative committees. Games and Economic Behavior 133, 248-255.
[11] Backurs, A., Indyk, P., Onak, K., Schieber, B., Vakilian, A., Wagner, T., 2019. Scalable fair clustering. In: Proceedings of the 36th International Conference on Machine Learning, ICML. Vol. 97 of Proceedings of Machine Learning Research. PMLR, pp. 405-413.
[12] Bera, S. K., Chakrabarty, D., Flores, N., Negahbani, M., 2019. Fair algorithms for clustering. In: Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019,. pp. 4955-4966.
[13] Brandt, F., 2017. Rolling the dice: Recent results in probabilistic social choice. In: Endriss, U. (Ed.), Trends in Computational Social Choice. AI Access, Ch. 1, pp. 3-26.
[14] Brill, M., Gölz, P., Peters, D., Schmidt-Kraepelin, U., Wilker, K., 2020. Approval-based apportionment. In: Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI).
[15] Chakrabarty, D., Negahbani, M., 2021. Better algorithms for individually fair k-clustering. CoRR abs/2106.12150.
[16] Chan, H., Filos-Ratsikas, A., Li, B., Li, M., Wang, C., 2021. Mechanism design for facility location problems: A survey. In: Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI).
[17] Chen, X., Fain, B., Lyu, L., Munagala, K., 2019. Proportionally fair clustering. In: Proceedings of the 36th International Conference on Machine Learning, ICML. Vol. 97 of Proceedings of Machine Learning Research. PMLR, pp. 1032-1041.
[18] Chierichetti, F., Kumar, R., Lattanzi, S., Vassilvitskii, S., 2017. Fair clustering through fairlets. In: Proceedings of the 31st Annual Conference on Neural Information Processing Systems (NeurIPS). pp. 5029-5037.
[19] Chiplunkar, A., Kale, S., Ramamoorthy, S. N., 2020. How to solve fair k-center in massive data models. In: Proceedings of the 37th International Conference on Machine Learning, ICML 2020. Vol. 119. PMLR, pp. 1877-1886.
[20] Chouldechova, A., Roth, A., 2020. A snapshot of the frontiers of fairness in machine learning. Communications of the ACM 63 (5), 82-89.
[21] Congressional-Research-Service, 2021. Artificial intelligence: Background, selected issues, and policy considerations. Tech. rep., Congressional Research Service Reports.
[22] Dua, D., Graff, C., 2017. UCI machine learning repository. University of California, Irvine, School of Information and Computer Sciences.
URL http://archive.ics.uci.edu/ml
[23] Dummett, M., 1984. Voting Procedures. Oxford University Press.
[24] Dummett, M., 1997. Principles of Electoral Reform. Oxford University Press.
[25] Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R. S., 2012. Fairness through awareness. In: Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012. ACM, pp. 214-226.
[26] Esmaeili, S. A., Brubach, B., Srinivasan, A., Dickerson, J. P., 2021. Fair clustering under a bounded cost. In: Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS).
[27] European Commission, Feb. 2020. White paper on artificial intelligence: a european approach to excellence and trust. White Paper $\operatorname{COM}(2020) 65$ final, European Commission, Brussels.
[28] Everitt, B. S., Landau, S., Leese, M., 2009. Cluster Analysis, 4th Edition. Wiley Publishing.
[29] Faliszewski, P., Skowron, P., Slinko, A., Talmon, N., 2017. Multiwinner voting: A new challenge for social choice theory. In: Endriss, U. (Ed.), Trends in Computational Social Choice. Ch. 2.
[30] Faliszewski, P., Slinko, A., Stahl, K., Talmon, N., 2018. Achieving fully proportional representation by clustering voters. Journal of Heuristics 24 (5), 725-756.
[31] Farahani, R. Z., Hekmatfar, M., 2009. Facility location: concepts, models, algorithms and case studies. Springer Science \& Business Media.
[32] Friedman, B., Nissenbaum, H., Jul. 1996. Bias in computer systems. ACM Trans. Inf. Syst. 14 (3), 330-347.
[33] Jung, C., Kannan, S., Lutz, N., 2020. Service in your neighborhood: Fairness in center location. In: 1st Symposium on Foundations of Responsible Computing, FORC 2020, June 1-3, 2020, Harvard University, Cambridge, MA, USA (virtual conference). Vol. 156 of LIPIcs. pp. 5:1-5:15.
[34] Kleindessner, M., Samadi, S., Awasthi, P., Morgenstern, J., 2019. Guarantees for spectral clustering with fairness constraints. In: Proceedings of the 36th International Conference on Machine Learning (ICML). Vol. 97. PMLR, pp. 3458-3467.
[35] Li, B., Li, L., Sun, A., Wang, C., Wang, Y., 2021. Approximate group fairness for clustering. In: Proceedings of the 38th International Conference on Machine Learning. Vol. 139 of Proceedings of Machine Learning Research. PMLR, pp. 6381-6391.
[36] Megiddo, N., Supowit, K. J., 1984. On the complexity of some common geometric location problems. SIAM Journal on Computing 13 (1), 182-196.
[37] Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., Galstyan, A., 2021. A survey on bias and fairness in machine learning. ACM Computing Surveys 54 (6).
[38] Micha, E., Shah, N., 2020. Proportionally fair clustering revisited. In: 47th International Colloquium on Automata, Languages, and Programming, ICALP 2020. Vol. 168. pp. 85:1-85:16.
[39] Moulin, H., 2003. Fair Division and Collective Welfare. The MIT Press.
[40] O'Neil, C., 2016. Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Crown Publishing Group, USA.
[41] Peters, D., Pierczyński, G., Skowron, P., 2021. Proportional participatory budgeting with additive utilities. In: Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS). pp. 12726-12737.
[42] Peters, H., van der Stel, H., Storcken, T., 1992. Pareto optimality, anonymity, and strategy-proofness in location problems. International Journal of Game Theory 21 (3), 221-235.
[43] Pierczynski, G., Skowron, P., 2021. Core-stable committees under restricted domains. CoRR.
[44] Procaccia, A. D., 2008. Towards a theory of incentives in machine learning. SIGecom Exchanges 7 (2).
[45] Procaccia, A. D., Tennenholtz, M., 2013. Approximate mechanism design without money. ACM Transactions on Economics and Computation 1 (4).
[46] Pukelsheim, F., 2014. Proportional Representation: Apportionment Methods and Their Applications. Springer.
[47] Sánchez-Fernández, L., Elkind, E., Lackner, M., Fernández, N., Fisteus, J. A., Basanta Val, P., Skowron, P., 2017. Proportional justified representation. In: Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI). AAAI Press.
[48] Vakilian, A., Yalçiner, M., 2021. Improved approximation algorithms for individually fair clustering. CoRR abs/2106.14043.

## Appendix A. Omitted proofs

Proof of Proposition 1. Let $X$ be an outcome which satisfies PRF for Unconstrained Clustering. Fix $S \subseteq$ $N, c \in \mathscr{M}$ with $|S| \geq\left\lceil\frac{n}{k}\right\rceil$.

Let $y$ be the maximum distance between pairs of agents in $S$, i.e., $y=\max _{i, j \in S} d(i, j)$, and let $i_{1}$ and $i_{2}$ be two agents in $S$ that maximize this distance. We assume that $y>0$; otherwise, if $y=0$, then every agent in $S$ is located at the same point, and there exists $x \in X$ such that $d(i, x)=0 \leq d(i, c)$ for all $i \in S$. Denote by $i^{*}$ the agent in $S$ for which the distance to $c$ is maximized. We have that $d\left(i^{*}, c\right) \geq \max \left(d\left(i_{1}, c\right), d\left(i_{2}, c\right)\right) \geq \frac{1}{2}\left(d\left(i_{1}, c\right)+d\left(i_{2}, c\right)\right) \geq \frac{y}{2}$.

Next, since $|S| \geq \frac{n}{k}$ and every pair of agents in $S$ are a distance of at most $y$ from each other, it follows from $X$ satisfying Unconstrained PRF that there exist $x \in X, i \in S$ such that $d(i, x) \leq y$. Using this along with the triangle inequality, we can establish the following upper bound: $d\left(i^{*}, x\right) \leq d\left(i^{*}, c\right)+d(c, x) \leq$ $d\left(i^{*}, c\right)+d(i, c)+y$. From these observations, the minimum PF approximation ratio obtained by $i$ or $i^{*}$ becomes

$$
\begin{aligned}
& \min \left(\frac{d(i, x)}{d(i, c)}, \frac{d\left(i^{*}, x\right)}{d\left(i^{*}, c\right)}\right) \\
& \leq \min \left(\frac{y}{d(i, c)}, \frac{d\left(i^{*}, c\right)+d(i, c)+y}{d\left(i^{*}, c\right)}\right) \\
& \leq \min \left(\frac{y}{d(i, c)}, 3+\frac{2 \cdot d(i, c)}{y}\right) \\
& \leq \max _{z \geq 0}(\min (z, 3+2 / z))=\frac{3+\sqrt{17}}{2}
\end{aligned}
$$

Note that the final inequality above holds since replacing $\frac{y}{d(i, c)}$ by the non-negative number $z$, which maximizes the expression, can only weakly increase the expression. It can be easily verified that this expression is maximized when $z=\frac{3+\sqrt{17}}{2}$, hence the final transition. Thus, for any $c \in \mathscr{M}$, there exists $i \in S$ such that $d(i, X) \leq \frac{3+\sqrt{17}}{2} d(i, c)$.

Proof of Theorem 11. We prove by induction that if the number of candidates selected is less than $k$, then there is a way to select at least one more center. If the number of candidates selected is less than $k$, there is still aggregate weight of at least $n / k$ on the set of all agents. These agents can get a candidate selected from their neighborhoods if the neighborhood radius is large enough. In particular, for $d_{|D|}=\max _{i, j \in N} d(i, j)$, some unselected candidate from $M$ can be selected.

Note that for a given radius $d_{j}$, at most $k$ candidates can be selected. If no more candidates can be selected for a given $d_{j}$, the next radius $d_{j+1}$ is considered. There are $O\left(n^{2}\right)$ different distances that are to be considered. Next, we bound the running time.

There are $\left(n^{2}\right)$ different distances that are to be considered. These distances are sorted in $O\left(n^{2} \log \left(n^{2}\right)\right)$ time. For each distance that we consider, we check for each of the $k n$ locations whether they get sufficient support of $n / k$. It takes time $k n^{2}$ to check if some location has support $n / k$ for the current neighborhood distance. If no location has enough support, we move to the next distance. If some location has enough support, we need to select one less location. Therefore, we need to check whether some location has enough support at most $\max \left(k, n^{2} k\right)$ times. Therefore the running time is $O\left(n^{2} \log \left(n^{2}\right)\right)+O\left(\left(k n^{2}\right)\left(n^{2} k\right)\right)=O\left(n^{4} k^{2}\right)$.

The fact that the algorithm gives PRF and 3-PF follows immediately from Lemma 1 and Lemma 2 ,
Proof of Proposition 2. Let $X$ be an outcome which satisfies PRF for Discrete Clustering. Fix $S \subseteq N, c \in \mathscr{M}$ with $|S| \geq\left\lceil\frac{n}{k}\right\rceil$. Let $y=\max _{i \in S} d(i, c)$, the maximum distance any agent in $S$ is from $c$, and denote this agent $i^{*}$. We assume that $y>0$; otherwise, if $y=0$, then every agent in $S$ is located at $c, c \in X$ (by PRF), and $d(i, c)=d(i, X)$ for all $i \in S$.

Since $|S| \geq \frac{n}{k}$ and every agent in $S$ is a distance of at most $y$ from $c$, it follows from $X$ satisfying PRF that there exist $x \in X, i \in S$ such that $d(i, x) \leq y$. By the triangle inequality, $d(c, x) \leq d(i, c)+d(i, x)$ and thus $d\left(i^{*}, x\right) \leq y+d(i, c)+d(i, x)$. From these observations, the minimum PF approximation ratio obtained by $i$ or $i^{*}$ becomes

$$
\begin{aligned}
& \min \left(\frac{d(i, x)}{d(i, c)}, \frac{d\left(i^{*}, x\right)}{d\left(i^{*}, c\right)}\right) \\
& \leq \min \left(\frac{y}{d(i, c)}, \frac{y+d(i, c)+d(i, x)}{y}\right) \\
& \leq \min \left(\frac{y}{d(i, c)}, 2+\frac{d(i, c)}{y}\right) \\
& \leq \max _{z \geq 0}(\min (z, 2+1 / z))=1+\sqrt{2} .
\end{aligned}
$$

Note that the final inequality above holds since replacing $\frac{y}{d(i, c)}$ by the non-negative number $z$, which maximizes the expression, can only weakly increase the expression. It can be easily verified that this expression is maximized when $z=1+\sqrt{2}$, hence the final transition. Thus, for any $c \in \mathscr{M}$, there exists $i \in S$ such that $d(i, X) \leq(1+\sqrt{2}) d(i, c)$.

Proof of Theorem 2. Output size. We prove by induction that if the number of candidates selected is less than $k$, then then there is a way to select at least one one more center. If the number of candidates selected is less than $k$, there is still aggregate weight of at least $n / k$ on the set of all agents. These agents can get a candidate selected from their neighborhood if the neighborhood distance is large enough. In particular, for $d_{|D|}=\max _{i \in N, c \in \mathscr{M}} d(i, c)$, some unselected candidate from $M$ can be selected.

Running time. Note that for a given radius $d_{j}$, at most $k$ candidates can be selected. If no more candidates can be selected for a given $d_{j}$, the next distance $d_{j+1}$ is considered. Thus, here are $(n \cdot|\mathscr{M}|)$ different distances to be considered. These distances are sorted in $O(n \cdot|\mathscr{M}| \log (n \cdot|\mathscr{M}|))$ time. For each distance that we consider, we check for each of the $|\mathscr{M}|$ locations whether they get sufficient support of $n / k$. It takes time $|\mathscr{M}| \cdot n$ to check if some location has support $n / k$ for the current neighborhood distance. If no location has enough support, we move to the next distance. If some location has enough support, we need to select one less

|  | i | $j$ | l |
| :---: | :---: | :---: | :---: |
| i | 0 | $6+2 \sqrt{17}$ | $7+\sqrt{17}$ |
| j | $6+2 \sqrt{17}$ | 0 | $7+\sqrt{17}$ |
| l | $7+\sqrt{17}$ | $7+\sqrt{17}$ | 0 |
| c | $3+\sqrt{17}$ | $3+\sqrt{17}$ | 4 |
| $x$ | $13+3 \sqrt{17}$ | $13+3 \sqrt{17}$ | $6+2 \sqrt{17}$ |

Table B.3: Distances between agents $\{i, j, l\}$ and candidates $\{c, x\}$.
location. Therefore, we need to check whether some location has enough support at most max $(k, n \cdot|\mathscr{M}|)$ times. Therefore the running time is $O(n \cdot|\mathscr{M}| \log (n \cdot|\mathscr{M}|))+O((|\mathscr{M}| \cdot n)(k+n \cdot|\mathscr{M}|))=O\left(|\mathscr{M}|^{2} n\right)$.

PRF. Suppose for contradiction that the algorithm finds a set of $k$ centers $W$ that violates PRF. In that case, there exist a set of agents $S \subseteq \mathscr{N}$ of size at least $|S| \geq \ell n / k$ such that there are at least $\ell^{\prime} \leq \ell$ locations in $\mathscr{M}$ that are within distance $y$ of some agent in $S$, but the number of locations in $W$ that are within $y$ of at least some agent in $S$ is at most $\ell-1$. In that case, consider the snapshot of the algorithm when neighborhood distance $y$ is considered. At this point, the the number of locations that are within $y$ of at least some agent in $S$ is at most $\ell-1$. Since agents in $\mathscr{N}$ have only used their weight towards selecting locations within $y$ up till this point, it follows that $\sum_{i \in S} w_{i} \geq \ell \frac{n}{k}-(\ell-1) \frac{n}{k}=\frac{n}{k}$. It follows that the agents in $S$ still have total weight at least $\frac{n}{k}$ to select one more location that is at distance at most $y$ from some agent in $S$. Hence, at this stage, when the neighborhood distance is $y$, the algorithm would have selected at least one more center within distance $y$ than in the outcome $W$. Thus $W$ is not the correct output of the algorithm, a contradiction.

PF approximation. Since, as we just showed, the outcome returned by Algorithm 2 satisfies PRF, it follows from Proposition 2 that the outcome returned by Algorithm 2 satisfies $(1+\sqrt{2})$-PFAlgorithm 2 .

## Appendix B. Tight Analysis of PF Approximations

The following example shows the analysis in Proposition 1 is tight. That is, there is an instance for which a committee satisfying PRF for Unconstrained Clustering obtains no better than $\frac{3+\sqrt{17}}{2}-\mathrm{PF}$.

Example 5. Consider an instance of unconstrained clustering with $k=1$ and $\mathscr{N}=\left\{i_{1}, i_{2}, i_{3}\right\}$. Furthermore, let $c$ and $x$ be two candidate locations in $\mathscr{M}$ such that the distance between agents and candidates is as given in Table B.3.

It can be verified that all distances in Table B. 3 satisfy the triangle inequality. Note that, since $6+2 \sqrt{17}$ is the maximum distance between any pair of agents in $\mathscr{N}$ (call this $y$ ), and $d(l, x)=y$, it holds that $\{x\}$ satisfies PRF for Unconstrained Clustering. Through straightforward arithmetic, it can then be shown that $\frac{d(i, x)}{d(i, c)}=\frac{d(j, x)}{d(j, c)}=\frac{d(l, x)}{d(l, c)}=\frac{3+\sqrt{17}}{2}$.

## Appendix C. Strengthening of PRF

We consider two natural and stronger notions of PRF in the discrete setting (that also apply to the continuous setting). We will show that the versions that we consider do not guarantee the existence of a fair outcome in all settings.

The PRF concept can be increasingly strengthened in the following ways by making the requirements on the outcome stronger.

Definition 6 (Proportionally Representative Fairness (PRF)-II for Discrete Clustering). For any set of agents $S \subseteq \mathscr{N}$ of size at least $|S| \geq \ell n / k$, if there are $\ell^{\prime} \leq \ell$ candidates from $\mathscr{M}$ that are at distance at most $y$ from all agents in $S$, the following holds: there exists some $i \in S$ such that $|c \in X: d(i, c) \leq y| \geq \ell^{\prime}$.

Definition 7 (Proportionally Representative Fairness (PRF)-III for Discrete Clustering). For any set of agents $S \subseteq \mathscr{N}$ of size at least $|S| \geq \ell n / k$, if there are $\ell^{\prime} \leq \ell$ candidates from $\mathscr{M}$ that are at distance at most $y$ from all agents in $S$, the following holds: $\mid c \in X: d(i, c) \leq y$ for all $i \in S \mid \geq \ell^{\prime}$.

The following example shows that an outcome satisfying PRFII and hence PRFIII may not exist.
Example 6. Take $n=4, k=2$ with the following location profile on a unit interval $[0,1]: \boldsymbol{x}=(0,0,1,1)$. There are 4 candidate locations at $\mathscr{M}=(0,0.5,0.5,1)$. It is immediate that the 2 groups of agents at 0 and 1 must have a facility at 0 and 1. But this violates PRF-II. The set of agents $N$ (size 4) have 2 candidate locations with a distance of 0.5 of all agents; however, no agent in $N$ has 2 facility locations within a distance of 0.5 of them.

## Appendix D. $\boldsymbol{k}$-means ++

Step 1. The first centroid is selected randomly.
Step 2. Calculate the Euclidean distance between the centroid and every other data point in the dataset. The point farthest away will become our next centroid.

Step 3. Create clusters around these centroids by associating every point with its nearest centroid.
Step 4. The point which has the farthest distance from its centroid will be our next centroid.
Step 5. Repeat Steps 3 and 4 until $k$ number of centroids are located.

## Appendix E. Strategyproofness and Fairness

Most of the work on clustering assumes that the data points are truthfully reported. If data points correspond to individual agents who prefer to have a nearby center, then the issue of incentives also arises (see, e.g., Procaccia [44). It is easy to see that Algorithm 1 for the unconstrained setting coincides with the mid-point mechanism if there are two agents on a line. Since the mid-point (egalitarian mechanism) is known to be manipulable (see, e.g., Procaccia and Tennenholtz [45]), Algorithm 1 is not strategyproof even for $k=1$ (under which an agent cares about the sole facility to be as close as possible). A discretized version of the same example shows that Algorithm 2 is also not strategyproof.

Instead of understanding the strategic aspects of a particular algorithm, we next examine the tradeoff between fairness and strategyproofness. We focus on the case of $k=1$ for which the preference relation of an agent is clear: an agent wants the location to be as close as possible.

Proposition 3. Even for $k=1$ and the unconstrained setting, there exists no anonymous, PF, and strategyproof algorithm.

Proof. For $k=1$ and Euclidean space, an outcome is Pareto optimal if and only is it is in the convex hull of the agent locations 42. Next, we show that weak Pareto optimality is equivalent to Pareto optimality in this context. We show that if a location $x$ is not Pareto optimal, then it is not weakly Pareto optimal. Since $x$ is not Pareto optimal, it is strictly outside the convex hull $H$ of the agent locations. Let $y$ be the point in $H$ closest to $x$. Note that $y$ will either be a corner of $H$ or the perpendicular projection of $x$ onto an edge of $H$. Consider the line $L$ going through $y$ orthogonal to the line through $x$ and $y$. Assume without loss of generality that $L$ is vertical with $x$ to the right of $L$. Then all points of $N$ lie on $L$ or to the left of $L$. Let $i$ be an arbitrary point in $N$. If $i$ is the same point as $y$, then $d(i, y)<d(i, y)$. Now suppose $i$ is not at the same loation as $y$. Consider the triangle $(i, y, x)$. Since $x-y$ is orthogonal to $L$, it follows that the angle at point $y$ for triangle $(i, y, x)$ is at least $90^{\circ}$. The maximum angle of the other two corners of the triangle $(i, y, x)$ is strictly less than $90^{\circ}$ as the sum of the angles of a triangle is $180^{\circ}$. There is a well-known triangle inequality law that "if two angles of a triangle are unequal, then the measures of the sides opposite these angles are also unequal, and the longer side is opposite the greater angle." Hence, it follows that $d(i, y)<d(i, x)$. We have shown that $d(i, y)<d(i, y)$ for all $i \in N$. Hence $x$ is not weakly Pareto optimal.

Peters et al. 42 proved that there is no anonymous, strategyproof, and Pareto optimal mechanism for the Euclidean space with 3 dimensions. From our argument above, it follows that there is no anonymous, strategyproof, and weakly Pareto optimal mechanism for the Euclidean space with 3 dimensions. Since PF is equivalent to weak Pareto optimality for $k=1$, it follows that for $k=1$, there is no anonymous, PF , and strategyproof algorithm.

Corollary 1. Even for $k=1$, there exists no anonymous, core fair, and strategyproof algorithm.
Proposition 4. Even for $k=2$ and the discrete setting, there exists no PRF strategyproof algorithm even in one dimension.

Proof. Let $\boldsymbol{x}=(-1,1,1.2)$ and $\mathscr{M}=(0,0.9,1.3)$. Let $k=2$. PRF (for discrete setting) requires that $X=(0,0.9)$. This follows because the group of agents $S=N$ can demand 2 centers and there exists 2 centers at most $y=2.2$ for all agents locations-namely, 0 and 0.9 . The other groups of voters do not add any further demands (i.e., no more demands that are not already satisfied by considering $S=N$ ). With this outcome, agent 3 has access to 2 facilities with distances 0.3 and 1.2 , respectively.

Now consider a deviation by agent 3 to $x_{3}^{\prime}=1.4$. Now the group of agents $S=\{2,3\}$ can demand that the candidate center at 1.3 is chosen-this is because this is the only center that is within $y=0.4$ of both agents 2 and 3. There are 2 possible outcomes $X^{\prime}=(0.9,1.3)$ and $X^{\prime \prime}=(0,1.3)$. In the first case, agent 3 (with true location 1.2) has cost 0.1 and 0.3 . In the second case, agent 3 (with true location 1.2) has cost 0.1 and 1.2. In either case, the new outcome component-wise dominants the original outcome (under sincere voting). Hence, PRF is incompatible with SP in discrete setting.

## Appendix F. Experimental results



Figure F.4: Buddy dataset: MSD to closest $1, \mathrm{k} / 2$, and k centroids.


Figure F.5: Seeds dataset: MSD to closest 1, k/2, and k centroids.


Figure F.6: Wholesale dataset: MSD to closest 1, k/2, and k centroids.


Figure F.7: Buddy dataset: MSD to closest $1, \mathrm{k} / 2$, and k centroids.


Figure F.8: Seeds dataset: MSD to closest 1, k/2, and k centroids.


Figure F.9: Wholesale dataset: MSD to closest 1, k/2, and k centroids.


[^0]:    Email addresses: haris.aziz@unsw.edu.au (Haris Aziz), bartonlee@ethz.ch (Barton E. Lee), seanmorotachu@gmail.com (Sean Morota Chu), j.vollen@unsw.edu. au (Jeremy Vollen)

    * We owe thanks to John Dickerson and Bo Li for helpful comments and feedback.

[^1]:    ${ }^{1}$ The $k$-center solution outputs a subset $Y \in \arg \min _{W \subset \mathscr{M}}|W|=k \max _{i \in \mathscr{N}} d(i, W)$. The $k$-medians solution outputs a subset $Y \in \arg \min _{W \subseteq \mathscr{M},|W|=k} \sum_{i \in \mathscr{N}} d(i, W)$. The $k$-means solution outputs a subset $Y \in \arg \min _{W \subseteq \mathscr{M},|W|=k} \sum_{i \in \mathscr{N}} d(i, W)^{2}$.
    ${ }^{2}$ The results for unconstrained and discrete settings are incomparable in the following sense. The guarantee of existence of a fair outcome in the discrete setting guarantees the existence of a fair outcome in the unconstrained setting; however, a polynomial-time algorithm for computing a fair outcome in the discrete setting does not imply the same for the unconstrained setting.

[^2]:    ${ }^{3}$ In the term unanimous proportionality, unanimous refers to the condition where a set of agents have the same location, and proportionality refers to the requirements that such agents get an appropriate number of centers in proportion to the number of agents.

[^3]:    ${ }^{4}$ The idea of expanding neighborhoods is used in other algorithms in the literature, such as the Greedy Capture algorithm for the unconstrained setting studied by Micha and Shah [38]. Unlike our algorithm, it does not involve reweighting of agents. More importantly, Greedy Capture does not satisfy PRF and is not polynomial time.

[^4]:    ${ }^{5}$ To see this, let $k=3$ and let $\boldsymbol{x}=(0,0,1)$. Chen et al. s Greedy Capture algorithm and Li et al. s ALG $g$ will select candidate locations 0 and 1, but will not output a 3rd location. Of course, this issue could be rectified by arbitrarily choosing a 3rd candidate location-but then there is no guarantee that the set of locations would satisfy PRF.

[^5]:    ${ }^{6}$ I.e., Lloyd's algorithm for $k$-means minimization objective with a particular initialization. Given our focus on the discrete domain, we use a $k$-means ++ algorithms that only chooses centers from among the data points.

[^6]:    ${ }^{7}$ The datasets are licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0); for details, please refer to https://archive-beta.ics.uci.edu/ml/donation-policy

