Facility Location With Approval Preferences: Strategyproofness and Fairness

Edith Elkind
University of Oxford
United Kingdom
elkind@cs.ox.ac.uk

Minming Li
City University of Hong Kong
Hong Kong
miming.li@cityu.edu.hk

Houyu Zhou
City University of Hong Kong
Hong Kong
houyuzhou2-c@my.cityu.edu.hk

ABSTRACT

We develop a formal model of multiwinner facility location with approval preferences in one dimension: there is a set of facilities, a set of potential locations, and the goal is to build k facilities at these locations. Agents have approval preferences over 'facility, location' pairs, and may misreport their preferences if they can benefit from doing so. We consider both unit-demand agents and agents with additive demands, and the social objectives of coverage and utilitarian welfare. We ask whether these social objectives can be satisfied in a computationally efficient and strategyproof way. We also initiate the study of proportional representation in the context of facility location. We show that the axiom of justified representation, which is used to capture proportionality in multiwinner voting with approval preferences, is not well-suited for the facility location setting, and provide a relaxation of this axiom that can handle incompatibilities and may be of broader interest.

KEYWORDS

facility location; multiwinner voting; justified representation

ACM Reference Format:
Edith Elkind, Minming Li, and Houyu Zhou. 2022. Facility Location With Approval Preferences: Strategyproofness and Fairness. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022*, IFAAMAS, 9 pages.

1 INTRODUCTION

Consider a residential community that is planning to offer additional amenities to its residents. Several new facilities have been proposed: a playground, an outdoor gym, a fountain, and a secure bicycle parking. There are a few potential locations for these facilities, and a decision has been made to build exactly two facilities. Now, each resident's preferences over these new facilities may depend on where they will be built: a bicycle parking is only useful if it is very close to one's house, and a playground should be within a toddler's walking distance, whereas one may benefit from a fountain even if it is a bit further away. On the other hand, for those who do not have children, a playground is of no use no matter how close it is. A simple way to model this setting is to assume that each resident has an *approval radius* for each facility: that is, for each resident i and facility f_j there is a value r_{ij} such that i benefits from f_j (i.e., derives a unit of utility) if and only if the distance between the location of i and the location of f_j is at most r_{ij}. How should a central planner decide which facilities to build, given the agents' reports about their preferences, so as to come up with a plan that benefits the residents, and encourages truthful reporting?

This scenario shares important features with two settings that are well-studied in algorithmic game theory: facility location and multiwinner voting. Indeed, in facility location problems, the goal is typically to place one or more facilities so as to serve a network of agents. There are many variants of this basic model, and in recent years this research area has received a lot of attention from the multiagent systems and algorithmic game theory communities [6]. However, in the facility location problem it is typically assumed that all of the facilities will be built, so the only decision to be made is where to place them (see, however, [7]). In contrast, in multiwinner voting [11], the goal is to select k candidates from a given set, based on voters' ballots. While this task is similar to ours, in that we, too, need to select a subset of facilities, we then also need to decide where to build them. Perhaps a more useful way to conceptualize our problem as a multiwinner voting problem would be to think of pairs 'facility, location' as candidates. However, the resulting instance of multiwinner voting has some unusual features: for instance, some pairs of candidates are incompatible, and if the set of potential locations is infinite, then so is the set of candidates. Thus, neither the facility location problem nor the multiwinner voting problem fully capture the challenge that we face.

One may be tempted to separate our decision-making into two steps: first, we decide which facilities to build (using existing multiwinner voting procedures), and then we decide how to place them (using mechanisms developed for the facility location problem). However, it is easy to see that this approach may fail to produce a good allocation: the agents' preferences over facilities depend on where these facilities will be placed, so if we ask the agents to vote for a set of facilities, we may then be unable to choose locations for the selected facilities so as to benefit the voters who supported them in the first round. We face a similar problem if we first decide on locations for the facilities, and then select which facilities are to be placed there. Thus, we need an integrated procedure that can properly account for agents' preferences over the possible solutions. Our desiderata for such a procedure are as follows:

- *Solution Quality.* We are interested in optimizing some measure of collective agent utility, such as *coverage* (the number of agents who obtain positive utility in a given solution) or *social welfare* (the sum of agents' utilities).
- *Incentives.* We would like the agents to truthfully report their locations and their approval radii for all facilities, i.e., we are interested in designing mechanisms that are *strategyproof.*
- *Computational complexity.* For our procedures to be practically useful, they should be implementable by polynomial-time algorithms.

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online.

- *Fairness.* All agents and groups of agents should be treated fairly. In particular, if a large group of agents can be served by placing a single facility at a particular location, this group should not remain unrepresented.

1.1 Our Contribution

We develop a formal model of multiwinner facility location with approval preferences in one dimension. Our model is flexible enough to capture both (1) settings where facilities can be co-located and (2) settings where there can be at most one facility at each location. Further, it can handle both (a) the case where there is a finite set of possible locations and (b) the case where each facility can be placed anywhere on the line. We consider two types of agents: unit-demand agents, who are satisfied as long as one of their approved options is selected, and additive agents, whose utility scales linearly with the number of approved options. From the social planner's perspective, we consider two natural objectives, namely, coverage and social welfare maximization.

We show that, in general, maximizing coverage is NP-hard, but identify a few natural special cases where this problem is polynomial-time solvable. We also propose a mechanism for maximizing coverage that is strategyproof for unit-demand agents. For additive agents, the social welfare can be maximized in polynomial time. Moreover, if co-location is allowed, a welfare-maximizing outcome can be implemented in a strategyproof manner. In contrast, if co-location is not allowed, no mechanism that maximizes the social welfare is strategyproof. Further, if the social planner's goal and the agents' objectives are not aligned, in the sense that the planner wants to maximize coverage, while the agents are additive and each agent would like to maximize her own utility, strategyproof mechanisms do not exist either. We then focus on the study of fairness. As our starting point, we use the axiom of *justified representation* from the multiwinner voting literature [2]. We establish that a natural adaptation of this axiom to the facility location setting cannot be satisfied, because of incompatibilities between different options. Moreover, it is NP-hard to decide whether it can be satisfied in a given instance of our problem. We then propose a more general representation axiom, which is suitable for a setting with incompatibilities, and show that it can always be satisfied.

1.2 Related Work

In what follows, we briefly discuss relevant work on multiwinner voting and multiple facility location problems.

Multiwinner voting. In multiwinner voting, voters report their preferences over a set of candidates (which may be rankings of candidates or approval ballots, indicating which candidates they find acceptable), and the goal is to choose a fixed-size set of candidates, known as a *committee* [8, 11]. In Section 2 we show that our model generalizes multiwinner voting with approval ballots. An important concern in our context is strategyproofness, and Peters [18] and, subsequently, Lackner and Skowron [15] consider the problem of designing strategyproof multiwinner voting rules.

The justified representation (JR) axiom and its variants (such as proportional and extended justified representation) aim to capture a basic notion of fairness in multiwinner voting with approval ballots [2, 3, 20]. In this setting every instance admits a committee that satisfies this axiom, so the major research focus is on determining which of the existing voting rules provide JR. In contrast, we prove that, in our setting, a committee that provides JR need not exist.

Facility location problem. In the facility location problem, the goal is to understand how to locate one or more facilities in a metric space (typically, a line). Moulin [17] first characterizes the strategyproof mechanisms in the facility location problem where agents have single-peaked preferences. Procaccia and Tennenholtz [19] study how to approximately optimize certain social objectives in a strategyproof manner; for extensions, see [13, 22]. The problem of placing multiple facilities has been considered by several groups of authors [4, 14, 16, 21]. It is typically assumed that all facilities will be built (the only exception we are aware of is the work of Deligkas et al. [7]; however, they focus on the case of selecting one facility out of two), while we are interested in the case where only a subset of facilities will be built. Furthermore, most of the works study the facility location problem from a perspective of strategyproof mechanism design, while in our setting even the basic optimization problem turns out to be challenging; moreover, we are also interested in fairness. For further discussion, we point the reader to the survey of Chan et al. [6].

2 PRELIMINARIES

There is a set $F = \{f_1, \dots, f_m\}$ of m *facilities* and a set Y of *potential locations*. We will primarily focus on the *discrete* model, where $Y = \{y_1, \dots, y_s\}$, with $y_1, \dots, y_s \in \mathbb{R}$; when considering computational problems, we assume that all real parameters, such as $y_1, \dots, y_s$, are rational numbers represented as fractions, with numerators and denominators given in binary. In Section 5, we will also consider the *continuous* model, where $Y = \mathbb{R}$. The goal is to build k of the facilities in F, where $1 \le k \le m$. Facilities can only be built at locations in Y. We distinguish between the *setting with co-location*, where several facilities can be placed at the same point, and the *setting without co-location*, where this is not allowed.

There is a set of agents $N = [n]$. Each agent $i \in N$ is described by her *profile* $p_i = \langle x_i, \mathbf{r}_i \rangle$, where $x_i \in \mathbb{R}$ is her *location* and $\mathbf{r}_i = (r_{i1}, \dots, r_{im}) \in (\mathbb{R}_{\ge 0} \cup \{-\infty\})^m$ is a vector of *approval radii*: agent i derives utility 1 from facility f_j if f_j is located at distance at most r_{ij} from i, and utility 0 otherwise. Note that $r_{ij} = 0$ means that agent i benefits from f_j if and only if f_j is placed at x_i, and $r_{ij} = -\infty$ if agent i receives no benefit from f_j no matter where f_j is located.

Let $C = F \times Y$; we will refer to elements of C as *candidates*. We will denote an element (f_j, y_ℓ) of C by $c_{j\ell}$. The preferences of each agent $i \in N$ can then be described by her *approval ballot* A_i, i.e., the set of candidates that she approves: an agent i *approves* candidate (f_j, y) if and only if $|x_i - y| \le r_{ij}$. Our problem, then, can be phrased as selecting k candidates from C, where selecting a candidate (f_j, y) means building facility f_j at location y. We note that we cannot build the same facility at two distinct locations, i.e., we cannot select both (f_j, y) and (f_j, y') if $y \ne y'$. To capture this, we introduce the notion of *compatibility* between candidates: if co-location is allowed, we say that two distinct candidates (f, y) and (f', y') are *compatible* if $f \ne f'$, and if co-location is not allowed, we additionally require that $y \ne y'$.

A *committee* is a set of k candidates $W \subseteq C$ such that all candidates in W are pairwise compatible. We will now describe how agents evaluate committees. Specifically, we consider *unit-demand agents*, which are satisfied by W as long as they approve at least one candidate in W, and *additive agents*, whose utility is the number of candidates in W that they approve. Formally, for a unit-demand agent i with approval ballot A_i, her *utility* from a committee W is given by $u_i(W) = 1$ if $A_i \cap W \neq \emptyset$ and $u_i(W) = 0$ otherwise, whereas for an additive agent i with approval ballot A_i her *utility* from a committee W is given by $u_i(W) = |A_i \cap W|$.

The *(utilitarian) social welfare* of a committee W is the sum of agents' utilities, i.e., $\sum_{i \in N} u_i(W)$. The *coverage* of W is the number of agents who derive positive utility from W, i.e., $|\{i \in N : u_i(W) > 0\}|$. From the designer's perspective, a natural objective is to maximize the social welfare or coverage. For unit-demand agents these two objectives coincide, but for additive agents this is not the case.

Example 1. *Suppose there are* 5 *agents located at* 0, 3 *agents located at* 2 *and* 1 *agent located at* 10. *There is a set of three facilities* $F = \{f_1, f_2, f_3\}$. *For each agent, their approval radius for* f_1 *and* f_2 *is* 1. *For two of the agents at* 0 *and two of the agents at* 2, *their approval radius for* f_3 *is* 3. *All other approval radii are* $-\infty$. *Let* $k = 2$.

Suppose that $Y = \{1, 2, 9\}$. *To maximize coverage, we can place* f_1 *at* 1 *and* f_2 *at* 9, *to cover all agents. Now, suppose the agents are additive and we want to maximize the social welfare. Then the answer depends on whether co-location is allowed. If yes, we can place* f_1 *and* f_2 *at* 1, *so the social welfare is* $2 \cdot 8 = 16$. *If not, we can place* f_1 *at* 1 *and* f_3 *at* 2, *so the social welfare is* $8 + 4 = 12$.

Multiwinner Voting. Our setting can be viewed as a generalization of the well-studied setting of multiwinner voting with approval ballots. Formally, an instance of the multiwinner voting problem is a set of candidates C, a set of voters $\mathcal{N}$, a list of approval ballots $(\mathcal{A}_i)_{i \in \mathcal{N}}$, where $\mathcal{A}_i \subseteq C$ for each $i \in \mathcal{N}$, and an integer $\kappa \leq |C|$; the goal is to output a subset of C of size exactly κ. We can transform any such instance $\langle C, \mathcal{N}, (\mathcal{A}_i)_{i \in \mathcal{N}}, \kappa \rangle$ into an instance of our problem as follows. Assume for convenience that $C = \{1, \dots, m\}$. We set $N = \mathcal{N}$, $F = \{f_1, \dots, f_m\}$, and $Y = \{1, \dots, \kappa\}$. For each $i \in N$ we let $x_i = 0$, and set $r_{ij} = \kappa$ if $j \in \mathcal{A}_i$ and $r_{ij} = -\infty$ otherwise. Finally, we set $k = \kappa$. The resulting instance has the property that, once we decide to build a facility, its location does not matter, so the decision boils down to selecting k facilities to build, which is equivalent to selecting κ candidates in $\langle C, \mathcal{N}, (\mathcal{A}_i)_{i \in \mathcal{N}}, \kappa \rangle$. We will use this observation to derive some of our complexity results.

Mechanism Design. We aim to design mechanisms that take as input the set of facilities, the set of potential locations, and agents' locations and approval radii, and output a size-k committee. We would like these mechanisms to optimize (exactly or approximately) the designer's objectives, as well as to enjoy additional desirable properties. We focus on three such properties: (1) computational tractability, (2) strategyproofness and (3) fairness. We will formally define strategyproofness now, and postpone the discussion of fairness till Section 6. For readability, we assume that the set of facilities and possible locations as well as the committee size are fixed, so that the input to the mechanism is the list of agents' profiles.

Definition 1. *(Strategyproofness) A facility location mechanism* M *is* strategyproof *if for every agent* $i \in N$, *every list of agents' profiles* $(p_i)_{i \in N}$, *and every profile* p'_i, *we have*

$$u_i(M(p_1, \dots, p_i, \dots, p_n)) \geq u_i(M(p_1, \dots, p'_i, \dots, p_n)).$$

We will be interested in mechanisms that aim to maximize some quantity, such as social welfare or coverage. However, there are often multiple optimal outcomes, so it is important to specify how our mechanisms break ties. To illustrate, we will now consider two mechanisms for $k = 1$ that both maximize coverage, but break ties differently. Note that for $k = 1$ there is no difference between unit-demand agents and additive agents, and maximizing coverage is equivalent to maximizing the social welfare.

Theorem 1. *Let* $k = 1$. *Consider a mechanism that selects a candidate with the highest number of approvals, with ties broken lexicographically: if* (f_j, y) *and* $(f_{j'}, y')$ *both have the maximum number of approvals, then* (f_j, y) *is chosen over* $(f_{j'}, y')$ *if* $j < j'$ *or* $j = j'$, $y < y'$. *The resulting mechanism is strategyproof.*

Proof. Consider an agent i. If i has utility 1 when reporting truthfully, she has no incentive to misreport. Now, suppose that i's utility is 0, and, by misreporting, i changes the outcome from (f, y) to (f', y') and increases her utility to 1. Then i approves (f', y') and does not approve (f, y). Hence, by misreporting, i cannot decrease the number of approvals received by (f, y) or increase the number of approvals received by (f', y'). But then she cannot change the outcome from (f, y) to (f', y'), a contradiction. □

We will now present a different coverage-maximizing mechanism for $k = 1$ that is not strategyproof.

Example 2. *Consider the mechanism for* $k = 1$ *that selects a candidate with the highest number of approvals, but, to break ties, it favors facilities with a larger sum of approval radii. Specifically, given a set of tied candidates* $C' \subseteq F \times Y$, *it computes the set of facilities* $F' = \{f_j : (f_j, y) \in C' \text{ for some } y \in Y\}$, *sets* $R(f_j) = \sum_{i \in N : r_{ij} > 0} r_{ij}$ *for each* $f_j \in F'$, *chooses* $f \in \arg\max_{f_j \in F'} R(f_j)$ *(breaking ties according to a fixed order over F), picks the smallest* $y \in Y$ *such that* $(f, y) \in C'$, *and finally outputs* (f, y). *By construction, this mechanism maximizes coverage. However, it is manipulable.*

To see this, consider an instance with two agents and two facilities, where $Y = \{0, 7\}$. *Agent 1 is located at* 0 *and has approval radius* 0 *for* f_1 *and* $-\infty$ *for* f_2. *Agent 2 is located at* 8 *and has approval radius* 1 *for* f_2 *and* $-\infty$ *for* f_1. *Then* $(f_1, 0)$ *and* $(f_2, 7)$ *receive one approval each, and our tie-breaking rule sets* $F' = \{f_1, f_2\}$, *establishes that* $R(f_2) > R(f_1)$, *and outputs* $(f_2, 7)$. *However, agent 1 can misreport her location as* 2 *and her approval radius for* f_1 *as* 2, *in which case the tie-breaking rule establishes that* $R(f_1) > R(f_2)$ *and outputs* $(f_1, 0)$.

In what follows, we focus on the case $k \geq 2$ and pay close attention to tie-breaking rules.

3 UNIT-DEMAND AGENTS

In this section, we focus on unit-demand agents. Note that for such agents maximizing the utilitarian social welfare is equivalent to maximizing coverage.

3.1 Complexity

We start by formally defining the coverage maximization problem.

> An instance of FL-Coverage is given by a set of agents N, a set of facilities F, a set of locations Y, a list of agent profiles $(p_i)_{i \in N}$, a committee size k and a parameter λ. It is a yes-instance if there is a committee

of size k such that at least λ agents approve at least one candidate in the committee, and a no-instance otherwise.

In the multiwinner setting, the problem of computing a coverage-maximizing committee corresponds to computing the outcome of a well-known multiwinner voting rule, namely, the Chamberlin–Courant rule with approval preferences (CCAV). This rules selects a committee so as to maximize the number of voters who approve at least one candidate in the committee. An instance of the associated CCAV-Score problem is a tuple $\langle C, \mathcal{N}, (\mathcal{A}_i)_{i\in\mathcal{N}}, \kappa\rangle$ and a positive integer λ; it is a yes-instance if there exists a committee $\mathcal{W} \subseteq C$, $|\mathcal{W}| = \kappa$, such that $|\{i \in \mathcal{N} : \mathcal{A}_i \cap \mathcal{W} \neq \varnothing\}| \geq \lambda$, and a no-instance otherwise. This problem is known to be equivalent to MaxCoverage and hence NP-complete [12]. We will now show that this implies NP-hardness of FL-Coverage.

Theorem 2. *FL-Coverage is NP-complete both with and without co-location.*

Proof. To see that FL-Coverage is in NP, note that, given a set of candidates W, we can check if all candidates in W are pairwise compatible, and compute the resulting coverage. To prove NP-hardness, consider the mapping from an instance of multiwinner voting to an instance of our problem described in Section 2. A committee that provides a coverage of at least λ in the resulting instance of the facility location problem (with or without co-location) corresponds to a committee that satisfies at least λ voters. Thus, this mapping is a reduction from CCAV-Score to FL-Coverage. □

The proof of Theorem 2 implies that FL-Coverage remains hard if each agent has the same approval radius for all facilities she approves (i.e., there is a value R such that $r_{ij} \in \{-\infty, R\}$ for all $i \in N$ and $f_j \in F$), and even if all agents are co-located. We will now discuss several settings where this hardness result can be circumvented.

First, maximizing coverage is easy if all facilities are of the same 'type': e.g., the goal is to build k identical playgrounds, and the only decision to be made is where to locate them.

Proposition 3. *Suppose that all facilities are identical, i.e., for each agent $i \in N$ there is a value r_i such that $r_{ij} = r_i$ for all $f_j \in F$. Then FL-Coverage is in P, both with and without co-location.*

Proof. We associate each agent with an interval of length $2r_i$ centered at x_i. We can then think of our problem as voting over *locations*: we would like to select k locations to 'stab' as many agent intervals as possible. This problem then reduces to computing a winning committee under the Chamberlin–Courant rule when voters have *candidate interval* preferences, i.e., the candidates can be placed on a line so that each voter approves a contiguous segment of candidates. This problem is known to be polynomial-time solvable by dynamic programming [5, 9]. Note that there is no benefit to co-locating several facilities, so the setting with co-location is equivalent to the setting without co-location. □

Next, suppose that agents can only benefit from the facilities if these facilities are placed at their exact location: e.g., each location is a neighborhood, and the neighborhoods are so far from each other that it is impractical to use a facility if it is located in a different neighborhood. Then, maximizing coverage is easy if co-location is not allowed, but the latter condition is essential.

Proposition 4. *Suppose that each agent has approval radius $-\infty$ or 0 for each facility. Then FL-Coverage is in P if co-location is not allowed, but remains NP-hard if co-location is allowed.*

Proof. In the setting without co-location, we reduce FL-Coverage to a bipartite matching problem. We construct a complete bipartite graph with parts F and Y, where the weight of the edge $\{f_j, y_\ell\}$ is set to $w_{j\ell} = |\{i \in N : x_i = y_\ell, r_{ij} = 0\}|$, i.e., the number of agents who approve facility f_j if they are co-located with it. Then FL-Coverage reduces to finding a maximum-weight matching of size k in this graph. This problem, in turn, can be reduced to min-cost maximum flow problem [1]. Specifically, we create a graph with vertex set $F \cup Y \cup \{x, x', z\}$, which contains the edges of the original bipartite graph as well as $\{x', x\}$, $\{x, f_j\}$ for each f_j in F, and $\{y_\ell, z\}$ for each y_ℓ in Y. All edges except $\{x', x\}$ have unit capacities, while $\{x', x\}$ has capacity k; this ensures that the maximum flow in this graph is of size k, and an integer flow of this size corresponds to a size-k matching in the original graph. Further, we set the cost of the edge $\{f_j, y_\ell\}$ to $Z - w_{j\ell}$, where $Z = \max\{w_{j\ell} : f_j \in F, y_\ell \in Y\} + 1$; this ensures that all costs are positive, and an integer minimum-cost flow of size k corresponds to a maximum-weight matching of size k. It remains to observe that an integer min-cost max flow can be computed in polynomial time [1].

However, this approach no longer works for settings with co-location: indeed, in this case, we can again reduce CCAV-Score to our problem (e.g., we can place all agents at 0, and say that an agent has approval radius 0 for a facility if and only if she approves the respective candidate in the multiwinner instance). □

We also get an easiness result if each agent cares about at most one facility.

Proposition 5. *Suppose that each agent has a non-negative approval radius for at most one facility. Then FL-Coverage is in P both with and without co-location.*

Proof. If co-location is allowed, there is no dependence among facilities, so we can find an optimal location for each facility independently: for each facility f we consider a sub-instance of our problem that consists of agents whose approval radius for f is non-negative, and select a location that satisfies the maximum number of such agents. In the setting without co-location, we can reduce FL-Coverage to finding a maximum-weight k-matching, as in the proof of Proposition 4: the weight of the edge $\{f_j, y_\ell\}$ is set to the number of agents who approve $c_{j\ell}$ (note that, in a matching, each agent contributes to the weight of at most one edge). □

3.2 Strategyproofness

For the unit-demand setting, we can extend the strategyproofness result for $k = 1$ (Theorem 1) to all values of k.

Theorem 6. *Suppose all agents are unit-demand. Then the mechanism that maximizes coverage and breaks ties lexicographically is strategyproof.*

[1]For an alternative proof, see Question 121353 on MathOverflow.

Proof. Consider a winning committee W and an agent $i \in N$. If $u_i(W) = 1$, then i has no incentive to misreport. Now, suppose that $u_i(W) = 0$. By misreporting her profile, agent i cannot decrease the number of agents covered by W. Further, if $u_i(W') = 1$, then i cannot increase the number of agents covered by W' by changing her report. Hence, i cannot change the outcome from W to W'. □

Remark 1. *An alternative proof of Theorem 6 can be obtained by observing that the unit-demand setting can be conceptualized as single-winner approval voting setting, with each agent implicitly indicating which committees they approve. As approval voting with deterministic tie-breaking is strategyproof in the single-winner setting (see, e.g., [10]), the result follows.*

While the mechanism described in Theorem 6 is strategyproof and provides optimal coverage, computing its output is NP-hard. One may ask if there exists a strategyproof mechanism that achieves approximately optimal coverage and runs in polynomial time. If we make a mild assumption that each agent approves at least one candidate in C, then the dictatorship mechanism (where we ask one agent to choose the winning committee) provides a $1/n$ approximation to the optimal coverage, is strategyproof and runs in polynomial time. We can also obtain a $1/k$ approximation by asking the agents to report their approvals, choosing a single candidate with the maximum number of approvals, and selecting the remaining $k - 1$ candidates in a way that is independent of the agents' reports. This mechanism (together with lexicographic tie-breaking) is strategyproof (because it reduces to single-winner approval voting) and polynomial-time computable. However, designing a polynomial-time strategyproof mechanism that provides a constant-factor approximation to optimal coverage remains an open problem.

4 ADDITIVE AGENTS: SOCIAL WELFARE

In this section, we consider the setting where the agents have additive utilities and the goal is to maximize the social welfare.

4.1 Complexity

We first show that the associated computational problem is in P, both with or without co-location.

Proposition 7. *We can compute an outcome that maximizes the social welfare in polynomial time, both with and without co-location.*

Proof. If co-location is allowed, we can place each facility independently, choosing its location so that as many agents as possible approve placing that facility in that location.

If co-location is not allowed, we create a complete bipartite graph with parts F and Y; the weight of the edge $\{f_j, y_\ell\}$ is equal to the number of agents who approve $c_{j\ell}$. Then a maximum weight size-k matching in this graph corresponds to a committee of size k that maximizes the social welfare. Such a matching can be found in polynomial time, as argued in the proof of Proposition 4. □

4.2 Strategyproofness

When considering strategyproofness, an important consideration is whether facilities can be co-located. Indeed, in the scenario with co-location, we can optimize the social welfare in a strategyproof fashion, because we can place each facility independently. However, if co-location is not allowed, the resulting dependencies provide an incentive to act strategically.

Theorem 8. *For additive agents, there is a polynomial-time computable mechanism that maximizes the social welfare and is strategyproof if co-location is allowed.*

Proof. Suppose that co-location is allowed. We can efficiently compute the *score* of each facility f_j as $\sigma(f_j) = \max_{y \in Y} |\{i \in N : |x_i - y| \le r_{ij}\}|$, i.e., the maximum number of agents who approve a candidate (f_j, y), where the maximum is taken over all $y \in Y$. We can then select k facilities with the highest scores, breaking ties in favor of lower-indexed facilities. We place each selected facility f_j at a point where it is approved by $\sigma(f_j)$ agents; among all such points, we pick the leftmost point. The resulting mechanism maximizes the social welfare; we claim that it is strategyproof.

Let G be the set of facilities selected by our mechanism, and let W be the respective committee. Consider an agent i, and suppose that i can benefit from misreporting a profile $p_i' \ne p_i$, so that the mechanism selects a set of facilities G', where W' is the respective committee. We will say that i *downvotes* (respectively, *upvotes*) f if the score of f decreases (respectively, increases) when i reports p_i' instead of p_i. If $G' = G$, then i's misreport only changes the placement of the selected facilities. Suppose a facility f was moved from location y to a location y'. This move only improves i's utility if i approves (f, y') and disapproves (f, y). But in this case i cannot misreport her preferences so as to move f from y to y'.

We now consider the case $G' \ne G$. We claim that there is a matching between $G \setminus G'$ and $G' \setminus G$ such that if $a \in G \setminus G'$ is matched to $b \in G' \setminus G$ then i downvotes a or upvotes b. To see this, imagine that i proceeds in two stages: first, she downvotes some set of facilities A and then she upvotes some set of facilities B. After each stage, form a directed matching, with arcs pointing from the facilities that move out of the set of top k facilities (where the facilities are ordered by score, with ties broken lexicographically) to those that move into this set. Let these matchings be $\mathcal{M}$ and $\mathcal{M}'$, respectively. To create the final matching $\mathcal{M}^*$, consider a graph on vertex set F that has $\mathcal{M} \cup \mathcal{M}'$ as its set of arcs. One can check that this graph is a collection of pairwise disjoint directed arcs and two-arc paths; we replace each such path $a \to b \to c$ with an arc $a \to c$. The resulting directed matching pairs facilities in $G \setminus G'$ with those in $G' \setminus G$ and has the desired property.

Now, consider candidate $(f, y) \in W$ with f in $G \setminus G'$, such that f is downvoted. This means that i approves (f, y). On the other hand, consider a candidate $(f', y') \in W'$ with $f' \in G' \setminus G$ such that f' is upvoted. This means that i does not approve (f', y'). Thus, no edge of our matching makes a positive contribution to the utility of i, and hence i does not benefit from the manipulation.

We note that the continuous version of this problem (see Section 5) can be seen as a facility location problem with a threshold, where each agent's utility is 0 or 1 [23], and the optimal mechanism has been shown to be strategyproof in this setting. □

The tie-breaking rule plays an important role in our proof.

Example 3. *Suppose there are two additive agents located at* 0 *and four facilities* a, b, c, d, *and the set* Y *contains* 0. *Agent 1's approval radius for* a, c, *and* d *is* 0, *and her approval radius for* b *is* $-\infty$. *Agent 2's approval radius for* b *is* 0, *and her approval radius for other facilities* *is $-\infty$. Suppose $k = 2$, and co-location is allowed. Note that all selected facilities will be placed at 0. Choosing any pair of facilities maximizes the social welfare. Suppose the tie-breaking rule selects $\{a, b\}$ in this case. Further, suppose that agent 1 misreports her approval radius for a as $-\infty$. Any size-2 subset of $\{b, c, d\}$ would then maximize the social welfare. If this tie is broken in favor of $\{c, d\}$, then agent 1 has an incentive to misreport.*

Next, we consider settings without co-location. It turns out that, for additive agents, no welfare-maximizing mechanism is strategyproof if co-location is not allowed.

Proposition 9. *For additive agents, if co-location is not allowed, no welfare-maximizing mechanism is strategyproof. This holds even if each agent's approval radius for each facility is $-\infty$, 0 or 1, and even if agents cannot misreport their locations.*

Proof. We construct an instance with $F = \{f_1, f_2\}$, $Y = \{1, 2, 3\}$ and $k = 2$. There are 10 agents: 2 agents at 1, 6 agents at 2, and 2 agents at 3. For all agents, their approval radius for each facility is $-\infty$ or 0. We specify the profiles of agents 1 and 2; all other agents approve a single candidate, so we will simply list their approval ballots. We have $p_1 = \langle 1, (0, 0)\rangle$, $p_2 = \langle 3, (0, 0)\rangle$. The remaining agents at 1 and 3 approve $(f_1, 1)$ and $(f_2, 3)$ respectively. The 6 candidates at 2 are split into two equal groups: one group approves $(f_1, 2)$, and the other group approves $(f_2, 2)$.

If co-location is not allowed, there are two optimal solutions, with social welfare 5 each: $\{(f_1, 1), (f_2, 2)\}$ and $\{(f_1, 2), (f_2, 3)\}$. If the first solution is selected, the utility of agent 2 is 0. She can benefit by changing her approval radius for f_1 to 1. Then the social welfare of the second solution becomes 6, so any welfare-maximizing mechanism has to select it, and the utility of agent 2 becomes 1. Symmetrically, if in the original instance the second solution is selected, agent 1 can benefit by changing his approval radius for f_2 to 1, so that the mechanism is forced to select the first solution. □

4.3 Coverage

We will also briefly discuss what happens when the designer aims to cover as many agents as possible, but the agents themselves are additive. That is, we ask whether a coverage-maximizing mechanism can be strategyproof for additive agents. Perhaps unsurprisingly, the answer is 'no'.

Theorem 10. *For additive agents, no coverage-maximizing mechanism is strategyproof, both with and without co-location.*

Proof. Let $F = \{f_1, f_2\}$, $Y = \{1, 2\}$, $k = 2$. We have four agents with the following locations and approval radii: $p_1 = \langle 1, (1, 0)\rangle$, $p_2 = \langle 2, (1, 0)\rangle$, $p_3 = \langle 1, (0, 1)\rangle$, $p_4 = \langle 2, (0, 1)\rangle$.

Suppose that co-location is allowed. There are four possible committees, $W_1 = \{c_{11}, c_{22}\}$, $W_2 = \{c_{12}, c_{21}\}$, $W_3 = \{c_{11}, c_{21}\}$, $W_4 = \{c_{12}, c_{22}\}$, which cover all four agents. If W_1 or W_4 is selected, agent 1 can misreport her profile as $\langle 1, (-\infty, 0)\rangle$, so that she only approves c_{21} and is therefore not covered by the chosen committee. A coverage-maximizing rule then has to select W_2 or W_3, so the utility of agent 1 improves from 1 to 2. Symmetrically, if W_2 or W_3 is selected, agent 2 can misreport her profile to change the outcome so that f_2 is placed at 2.

If co-location is not allowed, we can simplify this argument by only considering the committees W_1 and W_2. □

5 CONTINUOUS MODEL

So far, we assumed that the set of possible locations Y is finite. Alternatively, we can assume that each facility can be located anywhere on the real line. In this case, depending on whether co-location is allowed, we distinguish between two variants of the model: continuous without co-location, and continuous with co-location. Note that in the continuous setting the set of candidates $F \times Y$ is infinite.

Example 4. *Suppose that agents 1 and 2 are located at 0, agents 3 and 4 are located at 2, $F = \{f_1, f_2\}$, we have $r_{11} = r_{31} = 1$, $r_{22} = r_{42} = 1$, and all other approval radii are $-\infty$. If co-location is allowed, we can cover all agents by placing both facilities at 1. If co-location is not allowed, at most three agents can be covered: e.g., we can place f_1 at 1 and f_2 at 0, covering agents 1, 2, and 3.*

The continuous model may appear to be harder to work with than the discrete model: e.g., as the set of candidates is infinite, we cannot consider all candidates one by one, or explicitly list all size-k committees. Nevertheless, it turns out that we can extend all of our positive results to the continuous setting by discretizing Y.

Proposition 11. *Consider an instance with a set of agents N, a set of facilities F, a set of locations $Y = \mathbb{R}$, a list of agent profiles $(p_i)_{i\in N}$ and committee size k. Let*

$$Y_0 = \{x_i + r_{ij}, x_i - r_{ij} : i \in N, f_j \in F, r_{ij} \neq -\infty\}.$$

Number the points in Y_0 as $y_1, \dots y_t$, with $y_1 < \dots < y_t$. Form Y' by starting with Y_0 and, for each $q \in [t-1]$, adding k points from the interior of the interval $[y_q, y_{q+1}]$ to Y'. Then for every size-k committee $W \subseteq F \times Y$ there is a size-k committee $W' \subseteq F \times Y'$ such that for each $i \in N$ and each $\ell \in \mathbb{Z}$ it holds that if i approves ℓ candidates in W then she approves at least ℓ candidates in W'.

Proof. Suppose that in W there are ℓ facilities located strictly between y_q and y_{q+1}; note that $\ell \leq |W| = k$. We can shift these facilities to ℓ distinct points in Y' that are located strictly between y_q and y_{q+1}: by construction, an agent who approves (f, y) also approves (f, y') for each $y_q < y' < y_{q+1}$. We do not change the positions of facilities that are placed in points in Y_0. Thus, we can move all facilities to points in Y' so that no agent is negatively affected. □

If co-location is allowed, we can simplify the construction in Proposition 11, by simply using the set Y_0; in the proof, we simply shift each facility to the nearest point in Y_0. In this case, our discretization does not depend on k.

The construction in Proposition 11 enables us to extend, e.g., Theorem 1 to the continuous setting: to select a candidate with the highest number of approvals, it suffices to consider candidates in $F \times Y'$. Importantly, the size of the set Y' is polynomial in $|N|$ and $|F|$, so this mechanism still runs in polynomial time. More broadly, if a problem admits a polynomial-time algorithm in the discrete model (with or without co-location), it also admits a polynomial-time algorithm in the continuous model (with or without co-location): we reduce the continuous problem to a discrete problem by replacing $Y = \mathbb{R}$ with Y'. This argument also shows that FL-Coverage remains in NP in the continuous setting (the reader can verify that the hardness proof goes through for the continuous case as well, so FL-Coverage is NP-complete in the continuous model).

6 JUSTIFIED REPRESENTATION

In this section, we study fairness in multiwinner facility location settings, as captured by the notion of justified representtion and its variants. To simplify the presentation, we focus on the discrete model without co-location.

The notion of justified representation has been proposed by Aziz et al. [2]. We will now adapt this definition to our setting.

Definition 2. *Consider an instance $\langle F, Y, N, (p_i)_{i\in N}, k\rangle$ of the multiwinner facility location problem with $|N| = n$. We say that a committee $W \subseteq F \times Y$, $|W| = k$, provides* justified representation (JR) *for this instance if there is no set of agents N' of size at least n/k such that there is a candidate in $F \times Y$ that is approved by all agents in N', but no agent in N' approves any of the candidates in W.*

It is well-known that in the context of multiwinner voting with approval ballots, every instance admits a committee that provides justified representation [2]. However, in our setting this is not the case, because of incompatibilities between candidates.

Example 5. *Suppose $F = \{f_1, f_2\}$, $Y = \{y_1, y_2\}$, $k = 2$, and there are two agents, located at y_1 and y_2, with $r_{11} = r_{21} = 0$, $r_{12} = r_{22} = -\infty$. Then $n/k = 1$, agent 1 approves c_{11}, agent 2 approves c_{12}, but we can only build f_1 at one of the two locations.*

Further, we can show that it is NP-hard to decide whether a given instance admits a committee that provides JR.

Theorem 12. *Given an instance $\langle F, Y, N, (p_i)_{i\in N}, k\rangle$ of the multiwinner facility location problem with $|N| = n$, it is* NP*-complete to decide whether there exists a size-k committee that provides JR. The hardness result holds even if there is a value r such that each agent's approval radius for each facility is either r or $-\infty$.*

Proof sketch. The problem is clearly in NP: given a committee W, for each candidate $C\setminus W$ we can check how many agents approve that candidate, and then verify that fewer than n/k of them do not approve any candidate in W.

To prove NP-hardness, we reduce from the variant of 3-SAT in which each clause has exactly three variables. Consider an instance of 3-SAT with variables $\xi_1, \dots, \xi_q$ and clauses $K_1, \dots, K_t$. We will construct an instance of the multiwinner facility location problem with $n = 3(q + t) + 6$ agents and committee size $k = q + t + 2$ (so that $n/k = 3$) as follows.

We create facilities $h_0, \dots, h_q$, $f_0, \dots, f_t$. If $q \geq t$, we create potential locations $\{-2q, \dots, -q-1\} \cup \{0\} \cup \{1+q, \dots 2q\} \cup \{10q\}$. If $q < t$, we create potential locations $\{-2q, \dots, -q-1\} \cup \{0\} \cup \{1 + q, \dots 2q\} \cup \{10q\} \cup \{100q + 1, \dots, 100q + t\}$. Each agent has approval radius q or $-\infty$ for each facility. Note that the total number of facilities is $q + t + 2 = k$, so all facilities need to be placed, and it is just a matter of choosing locations for them.

For f_0, we create a group N^* containing $3q + 3$ agents at 0 who approve f_0 with radius q, so f_0 should be placed at 0 to satisfy JR.

For h_0, we create a group N_0 containing 3 agents at $10q$ who approve h_0 with radius q, so h_0 should be placed at $10q$ to satisfy JR.

For each clause K_j, we create a group N_j containing $n/k = 3$ agents. Suppose clause K_j contains literals ℓ_1, ℓ_2 and ℓ_3. There is one agent in N_j for each of these literals. All of these agents approve h_0 with approval radius q. Further, if ℓ_i is positive, i.e., $\ell_i = \xi_u$ for some $u \in [q]$, then the respective agent is located at $-q$, and approves h_u with radius q, and if ℓ_i is negative, i.e, $\ell_i = \neg\xi_u$ for some $u \in [q]$, then the respective agent is located at q and approves h_u with radius q. Note that N_j can be satisfied in two ways: either by placing h_0 at a location between $-2q$ and $2q$ (but this conflicts with the demand of N_0, which wants h_0 to be placed at $10q$), or by ensuring that at least one of the three agents in N_j is located within a distance q from the respective h-facility. Facilities $f_1, \dots, f_t$ receive no approvals.

One can verify that a placement of facilities that provides JR corresponds to a truth assignment, and vice versa; we omit the details due to space constraints. □

6.1 Maximal Justified Representation

We have seen that the JR axiom is inappropriate for settings with conflicts. We now put forward a new axiom, which we call *Maximal Justified Representation*, which waives representation requirements in case of unavoidable conflicts. We formulate this axiom in terms of approval ballots rather than agent profiles; note that, if the set of potential locations is finite, we can easily construct an approval ballot for each agent from her profile.

First, we need a formal definition of multiwinner approval voting with conflicts.

Definition 3. *An instance of the* multiwinner approval voting with conflicts *problem is a tuple $\langle C, N, (A_i)_{i\in N}, G, k\rangle$, where $C = \{c_1, \dots, c_m\}$ is a set of candidates, $N = \{1, \dots, n\}$ is a set of agents, $(A_i)_{i\in N}$ is a profile of approval ballots over C, G is an undirected graph with vertex set C such that an edge $\{c, c'\}$ indicates that candidates c and c' are incompatible, and k is the target committee size. A subset of candidates $W \subset C$ is called a* committee *if $|W| = k$ and all candidates in W are pairwise compatible.*

We note that, in general, it is NP-hard to decide if a given instance with conflicts admits any committees, as this requires finding an independent set of size k in the graph G. However, for instances that arise from the multiwinner facility location problem this task is easy: if co-location is allowed, it suffices to check that $k \leq |F|$, and otherwise we need to check that $k \leq \min\{|F|, |Y|\}$.

To define the axiom of maximal justified representation, we introduce the notions of *cohesive group* and *group-representing set.*

Definition 4. *(Cohesive group and group-representing set) Fix an instance of multiwinner approval voting with conflicts, i.e., a tuple $\langle C, N, (A_i)_{i\in N}, G, k\rangle$. We say that N^* is a* cohesive group *if $|N^*| \geq \frac{n}{k}$ and $\cap_{i\in N^*} A_i \neq \emptyset$. Let $\mathcal{N} = \{N^* : |N^*| \geq \frac{n}{k}, \cap_{i\in N^*} A_i \neq \emptyset\}$ be the set of cohesive groups. Then the group-representing set of candidate c_j is $\mathcal{N}_j = \{N^* \in \mathcal{N} : c_j \in \cup_{i\in N^*} A_i\}$.*

In words, $\mathcal{N}_j$ is a collection of cohesive groups that cannot 'complain' about a committee that contains c_j: each group in $\mathcal{N}_j$ includes a member that approves c_j. We are now ready to state the maximal justified representation axiom.

Definition 5. *(Maximal Justified Representation (MJR)) Fix an instance of multiwinner approval voting with conflicts, i.e., a tuple $\langle C, N, (A_i)_{i\in N}, G, k\rangle$. For each $i \in [m]$, let $\mathcal{N}_j$ be the group-representing set of candidate c_j. We say that a committee W of size k* provides maximal justified representation (MJR) *for this instance if* *there does not exist another committee W^* of size k such that*

$$\bigcup_{c_j \in W} \mathcal{N}_j \subset \bigcup_{c_j \in W^*} \mathcal{N}_j. \tag{1}$$

We say that an algorithm satisfies MJR *if for every instance it outputs a committee that provides MJR.*

That is, a committee W provides MJR if there is no other committee W' such that the set of cohesive groups represented by W' is a strict superset of those represented by W.

Of course, in the absence of conflicts MJR coincides with JR since both JR and MJR committees need to represent all cohesive groups.

Proposition 13. *For any instance of multiwinner approval voting without conflicts, a committee provides JR if and only if it provides MJR.*

However, in the presence of conflicts, there may exist a cohesive group such that no member of this group is represented by an MJR committee. For instance, in Example 5 each agent forms a cohesive group, so, in every feasible solution at least one of these groups will remain unsatisfied.

By construction, every instance of multiwinner voting with conflicts admits at least one committee that provides MJR: indeed, one can simply choose a committee that represents the maximum number of cohesive groups.

Proposition 14. *For any instance of multiwinner approval voting with conflicts, there is at least one committee that provides MJR.*

However, finding a committee that provides MJR is NP-hard even for instances of multiwinner voting with conflicts that are derived from multiwinner facility location problems: this follows from our observation about the relationship between MJR and JR and Theorem 12.

Proposition 15. *Given an instance of multiwinner approval voting with conflicts, it is* NP*-complete to find a size-k committee providing MJR.*

In the definition of MJR we focus on inclusion-maximality. Alternatively, one may be interested in committees that satisfy as many cohesive groups as possible. This approach can be captured by replacing condition (1) with

$$\left|\bigcup_{c_j \in W} \mathcal{N}_j\right| < \left|\bigcup_{c_j \in W^*} \mathcal{N}_j\right|; \tag{2}$$

we will refer to the resulting notion as *maximum justified representation (MMJR).*

Clearly, the MMJR condition is more demanding than MJR: every committee that provides MMJR provides MJR, but the converse is not true.

Example 6. *Suppose $F = \{f_1, f_2, f_3\}$, $Y = \{y_1, y_2, y_3\}$, $k = 3$, and there are four agents at y_1 and five agents at y_2. All agents have approval radius 0 for f_1 and approval radius $-\infty$ for f_2 and f_3. Then $n/k = 3$, and there are 5 cohesive groups formed by agents at y_1 (four groups of size 3 and one group of size 4) and 16 cohesive groups formed by agents at y_2 (ten groups of size 3, five groups of size 4 and one group of size 1). Thus, a solution that places f_1 at y_1 provides MJR but not MMJR, whereas a solution that places f_2 at y_1 provides MMJR (and hence MMR).*

The reader can verify that all of our results for MJR, such as the ones concerning existence (Proposition 14), NP-hardness (Proposition 15) and relationship with JR (Proposition 13) extend to MMJR.

7 CONCLUSION

We introduce the setting of multiwinner facility location with approval preferences, which combines facility location and multiwinner voting with approval ballots. We consider several variants of the model, which differ in whether the facilities can be co-located, whether the number of potential locations is finite or infinite, whether the agents are unit-demand or additive, and whether the social planner aims to maximize coverage or social welfare. For each setting, we explore the complexity of the associated optimization problem and whether the optimization objective can be implemented in a strategyproof way. We also propose a new representation axiom for settings where there may be incompatibilities among the candidates.

Interestingly, it turns out that in the approval-based model the continuous case can be reduced to the discrete case. In contrast, the settings with and without co-location appear to be very different, both from the computational and from the strategic perspective; as a rule, settings with co-location are easier to work with, but Proposition 4 illustrates that this is not always the case.

There are many open problems suggested by our analysis. For strategyproofness, we focused on the performance of mechanisms that implement specific optimization objectives. A natural extension of this idea is to explore the existence of (deterministic or randomized) strategyproof mechanisms that implement these objectives approximately. Another interesting direction is to explore what happens if the agents are limited in their ability to misreport, i.e., they can only misreport their locations, but not the approval radii, or vice versa (see Proposition 9 for the first step in this direction).

In our analysis of representation, we proposed a justified representation axiom that can be satisfied in any instance. However, finding the respective committee is computationally hard. It would be interesting to identify restrictions on agents' preferences under which this problem is tractable.

While our model is quite general, it could be extended further to capture additional real-life scenarios. In particular, one could specify the list of available locations for each facility, rather than a single list for all facilities: indeed, some facilities may require larger plots or special infrastructure, while others can be placed on any available plot. One can also investigate what happens if agents' preferences are expressed via ranked ballots rather than approval ballots.

ACKNOWLEDGMENTS

The work described in this paper was partially supported by a grant from Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. CityU 11205619). We would like to thank the anonymous AAMAS-22 reviewers for their very useful feedback.

REFERENCES

[1] Ravindra K Ahuja, Thomas L Magnanti, and James B Orlin. 1988. *Network flows.* Prentice Hall.

[2] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.

[3] Haris Aziz, Edith Elkind, Shenwei Huang, Martin Lackner, Luis Sánchez-Fernández, and Piotr Skowron. 2018. On the complexity of extended and proportional justified representation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI).* 902–909.

[4] Moshe Babaioff, Moran Feldman, and Moshe Tennenholtz. 2016. Mechanism design with strategic mediators. *ACM Transactions on Economics and Computation (TEAC)* 4, 2 (2016), 7:1–7:48.

[5] Nadja Betzler, Arkadii Slinko, and Johannes Uhlmann. 2013. On the computation of fully proportional representation. *Journal of Artificial Intelligence Research* 47 (2013), 475–519.

[6] Hau Chan, Aris Filos-Ratsikas, Bo Li, Minming Li, and Chenhao Wang. 2021. Mechanism design for facility location problem: A survey. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI), Survey Track.* 4356–4365.

[7] Argyrios Deligkas, Aris Filos-Ratsikas, and Alexandros A. Voudouris. 2021. Heterogeneous facility location with limited resources. *arXiv:2105.02712* (2021).

[8] Edith Elkind, Piotr Faliszewski, Piotr Skowron, and Arkadii Slinko. 2017. Properties of multiwinner voting rules. *Social Choice and Welfare* 48, 3 (2017), 599–632.

[9] Edith Elkind and Martin Lackner. 2015. Structure in dichotomous preferences. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI).* 2019–2025.

[10] Ulle Endriss. 2013. Sincerity and manipulation under approval voting. *Theory and Decision* 74, 3 (2013), 335–355.

[11] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2017. Multiwinner voting: A new challenge for social choice theory. *Trends in computational social choice* 74 (2017), 27–47.

[12] Uriel Feige. 1998. A threshold of ln n for approximating set cover. *J. ACM* 45, 4 (1998), 634–652.

[13] Michal Feldman, Amos Fiat, and Iddan Golomb. 2016. On Voting and Facility Location. In *Proceedings of the 17th ACM Conference on Economics and Computation (ACM EC).* 269–286.

[14] Dimitris Fotakis and Christos Tzamos. 2010. Winner-imposing strategyproof mechanisms for multiple facility location games. In *Proceeding of the 6th International Workshop of Internet and Network Economics (WINE).* 234–245.

[15] Martin Lackner and Piotr Skowron. 2018. Approval-Based Multi-Winner Rules and Strategic Voting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence, (IJCAI).* 340–346.

[16] Minming Li, Chenhao Wang, and Mengqi Zhang. 2020. Budgeted Facility Location Games with Strategic Facilities. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI).* 400–406.

[17] Hervé Moulin. 1980. On strategy-proofness and single peakedness. *Public Choice* 35, 4 (1980), 437–455.

[18] Dominik Peters. 2018. Proportionality and Strategyproofness in Multiwinner Elections. In *Proceedings of the 17th International Conference on Autonomous Agents and Multigent Systems (AAMAS).* 1549–1557.

[19] Ariel D. Procaccia and Moshe Tennenholtz. 2013. Approximate mechanism design without money. *ACM Transactions on Economics and Computation (TEAC)* 1, 4 (2013), 18:1–18:26.

[20] Luis Sánchez-Fernández, Edith Elkind, Martin Lackner, Norberto Fernández, Jesús Fisteus, Pablo Basanta Val, and Piotr Skowron. 2017. Proportional justified representation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI).* 670–676.

[21] Paolo Serafino and Carmine Ventre. 2016. Heterogeneous facility location without money. *Theoretical Computer Science* 636 (2016), 27–46.

[22] Zhongzheng Tang, Chenhao Wang, Mengqi Zhang, and Yingchao Zhao. 2020. Mechanism Design for Facility Location Games with Candidate Locations. In *Proceedings of the 14th International Conference on Combinatorial Optimization and Applications (COCOA).* 440–452.

[23] Qiang Zhang and Minming Li. 2014. Strategyproof Mechanism Design for Facility Location Games with Weighted Agents on a Line. *Journal of Combinatorial Optimiazation* 28, 4 (2014), 756–773.