

Mechanism Design for Facility Location Problems: A Survey

Hau Chan¹, Aris Filos-Ratsikas², Bo Li³, Minming Li⁴ and Chenhao Wang^{1*}

¹Department of Computer Science and Engineering, University of Nebraska-Lincoln

²Department of Computer Science, University of Liverpool

³Department of Computing, The Hong Kong Polytechnic University

⁴Department of Computer Science, City University of Hong Kong

hchan3@unl.edu, aris.filos-ratsikas@liverpool.ac.uk, comp-bo.li@polyu.edu.hk,
minming.li@cityu.edu.hk, cwang55@unl.edu

Abstract

The study of *approximate mechanism design* for facility location problems has been in the center of research at the intersection of artificial intelligence and economics for the last decades, largely due to its practical importance in various domains, such as social planning and clustering. At a high level, the goal is to design mechanisms to select a set of locations on which to build a set of *facilities*, aiming to optimize some social objective and ensure desirable properties based on the preferences of *strategic agents*, who might have incentives to misreport their private information such as their locations. This paper presents a comprehensive survey of the significant progress that has been made since the introduction of the problem, highlighting the different variants and methodologies, as well as the most interesting directions for future research.

1 Introduction

The study of facility location problems arises from combinatorial optimization, and aims at computing the optimal placement of facilities to minimize transportation costs for servicing customers [Hochbaum, 1982]. Besides locating actual facilities, the problem has also found a wide range of applications in other fields such as healthcare [Ahmadi-Javid *et al.*, 2017], clustering [Friedman *et al.*, 2001], and even problems that are not of a geographical nature, ranging from as simple as choosing the temperature for a classroom, to more advanced ones, like selecting a committee to represent people with different political views. Because of its practical importance, variants of facility location problems have attracted significant attention for a long time from different fields, such as operations research, theoretical computer science, economics and computational game theory.

In the last two research areas mentioned above, facility location problems are studied under a different light: strategic customers (referred to as *agents*) may have incentives to misreport their private information, such as their preferred locations, in order to manipulate the final outcome in their favor. A *strategyproof mechanism* is a predefined rule which

maps the preferences of the agents into appropriately chosen locations of the facilities, in a way that does not incentivize the agents to engage in such strategic considerations. Strategyproof mechanisms are certainly desirable, but they come at a cost; to ensure the truthful behavior of the participants, they often need to return sub-optimal solutions. One of the main investigations of the field of *computational social choice theory* [Brandt *et al.*, 2016] is precisely to quantify this loss in efficiency due to the requirement for strategyproofness.

In fact, more than a decade ago, Procaccia and Tennenholtz [2009] used one of the most basic variants of facility location problems to put forward the agenda of *approximate mechanism design without money*, which advocates the study of strategyproof mechanisms for various optimization problems through the lens of the *approximation ratio*, a celebrated notion in the field of theoretical computer science and approximation algorithms. The major difference between this agenda and prior work in algorithms is that here, the need for approximation does not come from computational limitations, but primarily due to the need for strategyproofness. The setting studied by Procaccia and Tennenholtz [2009] is that of a single facility to be built on a real line, and in which the agents' preferences are given by their ideal locations (their "peaks"), together with linear cost functions that are increasing at a fixed rate as one moves away from the peak.

A key characteristic of this setting, which makes it particularly amenable to the design of strategyproof mechanisms, is that the preferences of the agents are *single-peaked*.¹ This type of preferences, originally studied by Black [1948], is known to escape classic impossibility results and allows for winners of pairwise majority votes (referred to in the literature as "Condorcet winners"). Single-peaked preferences were popularized by the result of Moulin [1980], who identified all the possible strategyproof mechanisms for this setting by means of a *characterization*. Among those mechanisms, we highlight the *Median-point mechanism* below, phrased in appropriate terminology for facility location problems.

Mechanism (Median-point mechanism). Place the single facility at the median of the peaks of the agents.

It is not hard to see that this mechanism is strategyproof; an agent can only affect the position of the median by reporting

¹In fact, these are specific types of single-peaked preferences called *1-Euclidean*, see [Hotelling, 1990].

*Corresponding author

her peak to be on the “opposite” side of the median, but in that case the median can only “move away” from her peak. As a matter of fact, the Median-point mechanism is simply a majority outcome, something that was already observed by Black [1948] in his original paper.

Since the introduction of the agenda by Procaccia and Tennenholtz [2009], research in this area has flourished. For facility location problems in particular, a very fruitful line of work in the major venues in artificial intelligence and economics and computation has been concerned with several variants and generalizations of the main setting, often migrating from the “convenience” of single-peakedness and thus giving rise to much more challenging problems. The ultimate goal of this survey is to provide a comprehensive overview of the achievements of this literature on designing strategyproof mechanisms for these different variants. The only other survey on strategyproof facility location problems that we are aware of is that of [Cheng and Zhou, 2015], which is however mainly focused on the more classic settings. While we do highlight some of the earlier results on the problem as well, we put extra emphasis on the results of the last half decade, which has witnessed remarkable achievements.

1.1 The Classic Setting

We formally introduce the classic facility location setting in this section, which was studied in most of the earlier works on the problem; various extensions are presented in later sections. Let $N = \{1, \dots, n\}$ be a set of agents who are located in a metric space (M, d) , where M is a set of all possible positions and $d : M \times M \rightarrow \mathbb{R}$ is a distance function that satisfies the triangle inequality. We use $\mathbf{x} = (x_1, \dots, x_n) \in M^n$ to denote the location profile of the n agents. In a k -facility location setting, where the goal is to place k facilities, a (deterministic) *mechanism* is a function $f : M^n \rightarrow M^k$ that maps the agents’ locations \mathbf{x} to k locations of the facilities. A randomized mechanism outputs a probability distribution over all subsets of M with size k . Given a location profile $f(\mathbf{x})$ of k facilities, each agent $i \in N$ has a cost $c(f(\mathbf{x}), x_i)$. Usually, the cost is defined as the distance to the nearest facility, i.e., $c(f(\mathbf{x}), x_i) = \min_{y \in f(\mathbf{x})} d(x_i, y)$. When $f(\mathbf{x})$ is a distribution, the cost is defined as the expected distance. We remark that when there is only one facility ($k = 1$) and the metric space is the real line, we recover the original setting of [Procaccia and Tennenholtz, 2009].

In the strategic setting, the location x_i of each agent $i \in N$ is private information. For any $S \subseteq N$ and $i \in N$, we use \mathbf{x}_S to denote \mathbf{x} projected to S and \mathbf{x}_{-i} to $N \setminus \{i\}$. Therefore, mechanism f is *strategyproof* if for all $\mathbf{x} \in M^n$,

$$c(f(\mathbf{x}), x_i) \leq c(f(x_i, \mathbf{x}_{-i}), x_i) \text{ for all } i \in N, x'_i \in M.$$

In simple words, a strategyproof mechanism requires that no agent can benefit from misreporting, regardless of the reported positions of the others agents. A stronger notion is *group strategyproofness*, which requires that for any location profile \mathbf{x} and any coalition $S \subseteq N$, there is no joint deviation \mathbf{x}'_S of the agents in S such that all the agents in S benefit.

As we mentioned earlier, the goal of approximate mechanism design without money is to optimize some objective under the requirement that the employed mechanisms are strat-

egyproof. The two most widely studied objectives in the literature of the problem are the social cost and the maximum cost. Formally, given agent profile \mathbf{x} and facility locations \mathbf{y} , the *social cost* is defined as

$$SC(\mathbf{y}, \mathbf{x}) = \sum_{i \in N} c_i(\mathbf{y}, x_i)$$

and the *maximum cost* is defined as

$$MC(\mathbf{y}, \mathbf{x}) = \max_{i \in N} c_i(\mathbf{y}, x_i).$$

In economic terms, the social cost is analogous to the *utilitarian social welfare*, whereas the maximum cost is analogous to the *egalitarian social welfare*.

The objective of a strategyproof mechanism f is to compute locations \mathbf{y} such that $SC(\mathbf{y}, \mathbf{x})$ or $MC(\mathbf{y}, \mathbf{x})$ is minimized. A mechanism f has an approximation ratio $\alpha \geq 1$ for the social cost objective, if for any location profile $\mathbf{x} \in M^n$,

$$SC(f(\mathbf{x}), \mathbf{x}) \leq \alpha \cdot \min_{\mathbf{y} \in M^k} SC(\mathbf{y}, \mathbf{x}).$$

The approximation ratio is defined similarly for the maximum cost objective, or any of the other objectives that we will discuss in later sections.

Challenges. Strategyproofness is a rather stringent and fragile property, which makes the design of good strategyproof mechanisms particularly challenging, since robustness to strategic behavior has to be guaranteed for *every possible instance of the problem*. To demonstrate this, consider the single-facility on the real line setting studied by Procaccia and Tennenholtz [2009]. It is not hard to observe that the Median-point mechanism is in fact optimal for the social cost objective; for any other choice, the cost of more than half of the agents is increased by the same amount that the cost of the remaining (less than half) of the agents is decreased by. For the maximum cost objective however, the optimal choice, which is to place the facility at the midpoint of the two extreme reported locations, is no longer strategyproof; this is due to the fact e.g., the agent at the leftmost position would be incentivized to report a position further to the left, forcing the midpoint to be closer to her real location. The Median-point mechanism in this case only achieves a 2-approximation and in fact, this is the best possible among any strategyproof mechanism for the problem [Procaccia and Tennenholtz, 2009]. In more general metric spaces and for more facilities, the situation becomes ever more challenging.

One way to circumvent such barriers is to use randomization, and consider a slightly weaker notion of strategyproofness, often referred to as *truthfulness in expectation*; this notion stipulates that an agent will not decrease her expected cost by misreporting, even if her cost might in fact decrease in some realizations of randomness. An example of such a mechanism proposed by Procaccia and Tennenholtz [2009] for their setting is the following one.

Mechanism (Left-right-middle (LRM) mechanism). Select the peak of the leftmost agent with probability $1/4$, the peak of the rightmost agent with probability $1/4$ and the midpoint of the interval defined by those peaks with probability $1/2$.

It was shown in [Procaccia and Tennenholtz, 2009] that this mechanism achieves an approximation ratio of $3/2$ for the maximum cost objective, a notable improvement of the 2-approximation barrier for deterministic mechanisms. In more general settings, the benefits of randomization are far more pronounced, improving the best possible approximations from linear functions in the number of agents to small constants [Lu *et al.*, 2010].

The ultimate goal of mechanism design for facility location is to be applied to practical scenarios of social planning, by providing strategyproof mechanisms with guarantees on social efficiency. While the classic setting described above is a crucial starting point in these investigations, the plethora of different variants that have appeared over the years take further steps towards applicability in real-world scenarios.

Roadmap. We first introduce the results on the classic setting of strategyproof facility location in Section 2, and we proceed to discuss more recent variants in the next sections. Section 3 is devoted to settings with more advanced agents’ preferences, whereas Section 4 is concerned with mechanisms that either have enhanced capabilities, or need to be robust against stronger types of manipulations. In Section 5 we discuss settings with more complicated constraints on the facilities, and in Section 6, we discuss several further variants of the problem along different axes. In Section 7, we identify interesting avenues for future work on the problem.

2 The Classic Facility Location Setting: Methods and Results

We begin our exposition with the earlier works on the problem, which we refer to as the “classic” setting.

2.1 Locating a Single Facility

The foundations of locating a single facility on the real line are based on the work of [Moulin, 1980]. As we explained earlier, the Median-point mechanism is strategyproof (in fact, group strategyproof) and optimal for the social cost objective, and is 2-approximate for the maximum cost, which is the best possible among strategyproof mechanisms [Procaccia and Tennenholtz, 2009]. In terms of randomized mechanisms, the LRM mechanism is $3/2$ -approximate, which is also best possible among randomized strategyproof mechanisms, as proven in the same work. In that sense, the *approximate* mechanism design problem for a single facility on the line was settled in [Procaccia and Tennenholtz, 2009].

The single-facility problem for more general metrics was studied by Alon *et al.* [2010], who designed strategyproof mechanisms for circles and general graphs. A highlight of their results is that randomization can be used to escape a rather strong impossibility result, namely that in circle networks, the approximation ratio of any deterministic mechanism for the social cost is $\Omega(n)$. In particular, they employed the following well-known randomized strategyproof mechanism and showed that it achieves an approximation ratio of $2 - 2/n$ on general graphs.

Mechanism. (Random Dictatorship (RD)) Select an agent uniformly at random, and place the facility at her peak.

Meir [2019] showed that for three agents on a circle, improvements over the RD mechanism are possible. Feldman and Wilf [2013] proposed a family of generalized mechanisms on trees (the “parameterized boomerang mechanisms”) which generalize both the RD and the Median-point mechanism. Dokow *et al.* [2012] presented similar characterizations to that of [Schummer and Vohra, 2002] for discrete lines and circles. Recently, Tang *et al.* [2020a] provided a characterization of deterministic and randomized group strategyproof unanimous mechanisms in convex spaces (generalizing and extending some classic results of [Bordes *et al.*, 1990]), and showed approximation ratio bounds that can be obtained as a consequence of their characterization. Filimonov and Meir [2021] gave a full characterization of onto and strategyproof mechanisms in discrete trees.

2.2 Locating Two Facilities

For two facilities, most of the related work is focused on the line metric. This setting was in fact also studied by Procaccia and Tennenholtz [2009] in their original paper. In particular, they showed that the very simple mechanism that places one facility on the leftmost peak and the other on the rightmost peak is group strategyproof and achieves an $(n - 2)$ -approximation for the social cost. They also provided a lower bound of $3/2$ for any strategyproof mechanism; this was later improved to almost 2 [Lu *et al.*, 2009] and to an asymptotically tight bound of $\Omega(n)$ [Lu *et al.*, 2010]. Fotakis and Tzamos [2014] proved a tight lower bound of $n - 2$, thus settling the social cost objective for this setting. These works also considered randomized strategyproof mechanisms, with the best known approximation ratio being 4, given by the *Proportional Mechanism* of [Lu *et al.*, 2010]. The same approximation ratio in fact applies to general metrics as well.

For the maximum cost objective, Procaccia and Tennenholtz [2009] showed that their “two extremes” mechanism achieves a best possible 2-approximation, and showed how randomization can be used to achieve improved bounds.

2.3 Beyond Two Facilities

Escoffier *et al.* [2011] designed strategyproof mechanisms to locate $n - 1$ facilities in both general metric spaces and trees, providing linear upper bounds and constant lower bounds. For the real line metric, Fotakis and Tzamos [2016] presented a randomized strategyproof mechanism for any number of facilities k , the *Equal Cost* mechanism, which achieves approximation ratios of n and 2 for the social cost and the maximum cost respectively. Interestingly, the mechanism achieves these bounds even in the more general case where the cost functions are concave. For the social cost objective, Fotakis and Tzamos [2014] showed that *anonymous* deterministic strategyproof mechanisms for $k \leq 3$ cannot achieve any bounded ratio, even if the metric is the real line and there are 3 facilities and 4 agents; this includes the *Percentile mechanisms* studied in [Sui *et al.*, 2013]. Walsh [2020b] studied the k -facility location problem in 2-dimensional Euclidean and Manhattan spaces and provided impossibility results in terms of axiomatic properties.

3 More Involved Preferences

A common characteristic of the earlier works in facility location is that they study settings in which agents have a most-preferred point, and the cost functions are given as distances from that point. The second half of the past decade however witnessed the emergence of a plethora of results on different models, in which the preferences of the agents are more involved, giving rise to new challenges in terms of the design of good strategyproof mechanisms. In this section, we highlight the most prominent of those variants.

3.1 Obnoxious Facility Location

The location of an *obnoxious* (in the sense of “unwanted”) facility is in fact a rather classic problem in the literature on algorithms and optimization (e.g., see [Church and Garfinkel, 1978]). In this setting each agent would rather have the facility (e.g., a waste facility) built as far away as possible from her specified location (e.g., the location of her house); and the distance from that location to the facility is interpreted as her *utility*. In the context of approximate mechanism design, this setting was studied most notably by Cheng *et al.* [2013], who proposed strategyproof mechanisms for networks, which are based on a *majority vote* between two prespecified locations. For the line metric and different objectives, the problem was studied by Ye *et al.* [2015]. Ibara and Nagamochi [2012] provided characterization results for (group) strategyproof mechanisms for some special cases of the problem.

3.2 Heterogeneous Facility Location

Building upon the ideas of the obnoxious facility literature, a follow-up line of work, (e.g., [Feigenbaum and Sethuraman, 2015; Zou and Li, 2015; Serafino and Ventre, 2014; Serafino and Ventre, 2015]) considered models in which facilities can be both *desirable* and *undesirable*, from the perspective of *different* agents. In terms of strategyproofness, here it makes sense to distinguish between the cases where the *preferences* (desirable vs undesirable), or the locations of the agents are private information, or both. While these settings appear under many different names in the literature, we will use the term *heterogeneous facility location* to classify them under the same umbrella; we highlight the major variants within this part of the literature.

Dual preferences. One of the first heterogeneous facility location models studied in this literature was the *dual preference model*, proposed independently by [Zou and Li, 2015] and [Feigenbaum and Sethuraman, 2015]. In this setting each agent finds some facility either desirable or undesirable, but different agents have potentially different opinions. For the case of known locations, Zou and Li [2015] proved that it is possible to design an optimal strategyproof mechanism for the *social welfare* (the equivalent of the social cost for settings with utilities), whereas when both preferences and locations are unknown, they proposed a 3-approximate deterministic group strategyproof mechanism. In fact, the results of [Ibara and Nagamochi, 2012] for the obnoxious facility location (which is a special case of the dual preference setting) imply that this is actually best possible among strategyproof

mechanisms; this was also proven independently by Feigenbaum and Sethuraman [2015], who further provided lower bounds for randomized mechanisms. Kyropoulou *et al.* [2019] studied a related setting in which the facility can be located on a constrained feasible region of the Euclidean plane.

Optional preferences. A somewhat different model was considered first by Serafino and Ventre [2014; 2015], that of *optional preferences*. In that setting, there are two facilities to be built, each agent is “interested” in either facility or both, and her cost is equal to the sum of distances from her position to the locations of the facilities she is interested in (the *sum* function). The setting was later extended by Yuan *et al.* [2016] who considered the *min* and *max* functions instead, and they proved upper and lower bounds on the approximation ratio of deterministic strategyproof mechanisms for different objective functions. The upper bound for the min function was later improved [Li *et al.*, 2019a]. Recently, Deligkas *et al.* [2021] studied a utility version: A mechanism takes as input the positions and the preferences of the agents, and chooses to locate a single facility. Each agent has a utility equal to one minus her distance to the facility located if she is interested in it, and zero otherwise. Under three different settings depending on the level of agent-related information that is public or private, they designed strategyproof mechanisms that achieve a good approximation for maximizing the total utility.

Heterogeneous preferences for k facilities. Anastasiadis and Deligkas [2018] considered a setting with k facilities such that, for each facility, an agent prefers to be either close to the facility, far from it, or is indifferent between the two choices. In a sense, their setting merges the dual preference and the optional preference models described above, and extends them to more facilities. They studied mainly the egalitarian welfare (the utility equivalent of the maximum cost objective), but also obtained some results on the social welfare as a byproduct.

Fractional preferences. Another generalization of the optional preference model was proposed in [Fong *et al.*, 2018], that of the *fractional preference model*. In this setting, the preference of each agent for the facility is a number between 0 and 1 (rather than either 0 or 1, as in the case of optional preferences), with these numbers summing up to 1. Among other results, the authors show that when the locations are known, it is possible to achieve the maximum social welfare by a strategyproof mechanism, when the preferences are known, a 2 approximation is possible and when the locations and the preferences are private, there is a deterministic 4-approximate mechanism for the problem.

4 More Involved Incentives

In all the aforementioned works, the ability of the agents to misreport is limited to either their locations or their preferences. In this section we discuss some representative variants where the space of possible manipulations is extended, as well as variants in which mechanisms are given additional power to deal with the agents’ incentives.

Multiple locations. Besides the basic setting, Procaccia and Tennenholtz [2009] in their seminal paper also considered a

natural extension where each agent controls *multiple locations*, and her cost is either the total distance or the maximum distance from her locations to the facility. Hossain *et al.* [2020] further assumed that each agent may control multiple locations with different degrees of importance, and designed mechanisms which elicit the locations of the agents, as well as different levels of information about their importance. In their setting, in addition to misreporting their locations, agents may also *hide* some of their locations, if that serves their purposes. Yan and Chen [2020] enhanced the space of possible manipulations even further, assuming that the agents can *replicate* some of their locations as well, and fully characterized all mechanisms that are anonymous, efficient, and robust to manipulation with respect to this richer space of strategic behavior. Babaioff *et al.* [2016] studied three-stage mechanisms by introducing *strategic mediators*, where each agent is associated with exactly one mediator, and the cost of a mediator is the total cost of the agents that she represents. In this setting, the agents first strategically report their locations to the mediators, and then the mediators strategically report them to the mechanism.

False-name manipulations. Another type of manipulation which has been studied in the literature of mechanism design for *anonymous* environments is that of *false-name manipulation* (also known as *sybil-proofness*, e.g., see [Babaioff *et al.*, 2012]), where an agent can report more than once by creating and using fake identifiers (e.g., different email addresses or repeatedly logging onto an online service). A mechanism that is robust against standard and false-name manipulations is called *false-name-proof* (FNP), and is thus *stronger* than a simply strategyproof mechanism. False-name manipulations in facility location were first considered by Todo *et al.* [2011], who provided a characterization of FNP mechanisms for single-facility location on the real line, and proved that the mechanism that places the facility at the location of the leftmost agent achieves the best possible approximation for both the social cost and the maximum cost. Sonoda *et al.* [2016] studied FNP mechanisms for locating two facilities and proved that the “two extreme” mechanism is the best possible deterministic FNP mechanism on a line, while randomized mechanisms can achieve improved approximations. Nehama *et al.* [2019] defined a family of graphs that is called *ZV-line graphs*, and proved that there is a general facility location mechanism for these graphs that is strategyproof, false-name-proof and Pareto optimal. Other works that studied this type of manipulation in the context of facility location are [Ono *et al.*, 2017; Todo *et al.*, 2020].

Mechanism design with verification. Motivated by settings in which strategyproofness is not enough to guarantee small approximation ratios, Fotakis and Tzamos [2010] and Nisim *et al.* [2012] first considered *winner-imposing mechanisms*, which are capable of penalizing possibly deceitful agents by restricting how these agents interact with the outcome, e.g., by requiring agents to “use” the facility that is closer to their *reported* location. Fotakis and Tzamos [2010] proved that a winner-imposing extension of the Proportional Mechanism of [Lu *et al.*, 2010] is strategyproof and achieves

a $4k$ -approximation for the k -Facility Location game on the real line. In general, this line of research fits under the umbrella of mechanism design with verification (e.g., see [Caragiannis *et al.*, 2012]).

5 Constrained Facilities

Besides the different preference structures that were studied in previous models, the motivation from real-world applications drove researchers to consider additional constraints on the facilities rather than the agents, which might have to do with their capacities (e.g., in terms of physical storage or service time) or their possible locations. We highlight the major associated results below.

Limited locations. In the standard setting, the location of the facility is assumed to be some point in a continuum of possible points. In reality however, the “allowed” points for the placement of the facility could be limited - e.g., a bus stop needs to be located on the main road. Feldman *et al.* [2016a] studied the single-facility location setting in the context of voting embedded in a metric space, a topic very much related to that of *distortion in metric social choice* [Anshelevich *et al.*, 2018]. Among other results, for the line metric, they provided a deterministic 3-approximation mechanism, which places the facility at the candidate location that is closest to the median agent, as well as a randomized 2-approximation mechanism for the social cost, together with matching lower bounds. Tang *et al.* [2020b] further considered the maximum cost objective and the case of two-facility location. Walsh [2020a] considered a related setting in which each facility can be placed in a finite set of feasible subintervals.

Distance constraints. Zou and Li [2015] initiated the study of settings with two facilities, in which the placement of one of them imposes constraints on the placement of the other. In particular, they considered the *maximum distance requirement*, namely that the distance between the two facilities cannot exceed a certain threshold, for the *two-opposite-facility problem*; in this setting, there is one desired and one undesired facility, and the utility of each agent is the distance to undesired one minus the distance to the desired one. Later, Tang *et al.* [2018] and Chen *et al.* [2021] relaxed the maximum distance constraint by instead imposing a *penalty*, which is linear in the degree of violation. Xu *et al.* [2021] studied the analogous *minimum distance requirement* for heterogeneous two-facility locations, where the cost of an agent is the sum of her distances to the two facilities and Xu *et al.* [2020] studied the case with an *exact* distance requirement.

Capacitated facilities. The study of *capacitated facility location games* was initiated by Aziz *et al.* [2020b]. They studied a single-facility location problem with capacity constraints on the real line and provided a complete characterization of strategyproof mechanisms, quite similar to that of [Moulin, 1980]. In a follow-up work, Aziz *et al.* [2020a] considered a more general setting with k facilities and different capacities for each facility. Among other results, the authors proposed the *Extended Endpoint Mechanism* for the case of two facilities, which achieves a tight approximation ratio matching the lower bound proven in [Aziz *et al.*, 2020b].

6 Further Variants

In this section, we present some further interesting variants that do not quite fit any of the categories presented earlier.

Other social objectives. Besides the social cost and the maximum cost, the literature has also been concerned with other objectives. Feigenbaum *et al.* [2017] considered the minimization of the L_p -norm of costs for the real line metric, whereas Feldman and Wilf [2013] studied the objective of minimizing the sum of squares of the costs. In a markedly different direction, Cai *et al.* [2016] studied the *minimax envy* objective for single-facility location, where the envy of an agent with respect to another agent is simply their difference in distance from the facility, and the goal is to minimize the maximum envy over all the agents; we remark that contrary to most of the present results in this survey, the approximations of [Cai *et al.*, 2016] are additive. In a similar vein, Ding *et al.* [2020] and Liu *et al.* [2021] studied the *envy ratio*, which is defined as the maximum over the ratios between any two agents' utilities. Mei *et al.* [2019] introduced a *happiness* factor within $[0, 1]$ to measure the difference between the best possible facility location for an agent and the one selected by the mechanism, in both settings of standard and obnoxious facility games. In this regime, the goal of the agent is to maximize her happiness factor, and the mechanism aims to maximize the sum of those factors.

Double-peaked preferences. Filos-Ratsikas *et al.* [2015] studied the single-facility location problem with double-peaked preferences, where each agent has *two peaks* on the real line, and the cost increases as one moves away from either peak. The authors provided approximation ratio upper and lower bounds for deterministic and randomized strategyproof mechanisms, as well as an axiomatic characterization of group strategyproof mechanisms that satisfy some standard neutrality properties.

Externalities. Li *et al.* [2019b] considered settings where agents mutually influence each other, i.e., when one agent's utility positively or negatively affect some other agents'.

Weighted agents. Zhang and Li [2014] introduced *weights* to the agents, and showed that when weights are private information, the best that a deterministic strategyproof mechanism can do is to simply ignore them.

Distributed facility location. Recently, Filos-Ratsikas and Voudouris [2020] introduced a setting in which the facility location is selected as part of a distributed process: first, agents within groups (or *districts*) decide on a representative location and then the mechanism, oblivious to the actual locations of the agents, decides on a location from the set of representatives. The authors proved that the best possible strategyproof mechanism for this setting has an approximation ratio of 3.

Approximation and variance. Procaccia *et al.* [2018] considered the topic of reducing the variance of randomized strategyproof mechanisms while maintaining good approximation guarantees, by using the single-facility location problem as the application domain.

Additive approximations. Instead of multiplicative approximation ration, Golomb and Tzamos [2017] studied the worst-case additive approximation, and presented tight bounds for

mechanisms locating a single facility in both deterministic and randomized cases. They also argued that this measure is better suited for some situations.

Shuttle facility games. Fukui *et al.* [2020] introduced the shuttle facility games, where each agent's location is actually an interval in the line representing her commuting route. They designed a corresponding group strategyproof mechanism which also achieves the optimal social welfare.

Mechanisms with payments. Archer and Tardos [2001] studied a conceptually similar but fundamentally different facility location problem, a monetary model. Given a set of facilities holders and a set of customers, the facilities holders are strategic players and are required to report their private cost for them to be built. Once receiving the reports, the government needs to select a subset of facilities to build, and monetarily compensate these facilities with payment to guarantee truthful reports. They proven that any algorithm that solves the uncapacitated facility location problem optimally admits a truthful mechanism. Li *et al.* [2020] further studied this monetary model with a sharp budget constraint: the total payment of a mechanism is below a given budget. Chen *et al.* [2019] considered a dual-role setting where every agent plays both roles of facility holder and customer.

Automated mechanism design. There are also some works that use machine learning and deep learning approaches to design automated mechanisms; for example, see [Golowich *et al.*, 2018; Narasimhan *et al.*, 2016].

7 Future Research Directions

Many open problems related to all the different facility location variants were already implicit in our exposition in previous sections; we conclude this survey by identifying some settings which we believe deserve some special attention as part of future work.

Dynamic facility location. Most of the work on the problem assumes a static model, in which the agents' locations are fixed over time. Exceptions to this are two fairly recent works [De Keijzer and Wojtczak, 2018; Fotakis *et al.*, 2021] who considered multi-stage facility *reallocation* problems on the real line, where the facilities are being moved between stages based on the locations reported by the agents and the goal is to minimize the social cost by keeping the reallocation cost low as well. Another exception is the work of Wada *et al.* [2018], who investigated settings with variable and dynamic populations, affecting the choices of the designer. We believe that there is ample space for future work here, with more dynamic settings potentially amenable to approximate mechanism design.

Incomplete information. Another interesting avenue is to explore *stochastic* settings, where the information about the agents' locations and preferences contains uncertainty, or is based on distributional assumptions. Along those lines, Caragiannis *et al.* [2016], motivated by machine-learning applications, considered single-facility location when agents' locations are independently and identically drawn from an unknown distribution. In a more "possibilistic" approach, Menon and Larson [2019] considered a setting in which for

each agent there is an interval on the line representing all her possible locations, and they explore robust mechanisms that perform well with respect to all the possible unknown true preferred locations of the agents within those intervals. These works have only scratched the surface of the effects of incomplete information on the design of mechanisms for the problem, and more work in this direction is needed.

Incentives beyond strategyproofness. While strategyproofness is very desirable, it certainly impairs any attempt at good approximations, especially in settings where there are prohibitive lower bounds (e.g., the $n - 2$ lower bound on the social cost for the two-facility location case). It is interesting to see if relaxing strategyproofness to an ε -approximate version, in which agents have *very limited* incentive to misreport could “open up” the design space for much improved mechanisms. For example, Sui and Boutilier [2015a] provided several possibility/impossibility results for the ε -approximate (group) strategy-proofness in both constrained and unconstrained problems. Oomine *et al.* [2017] studied the trade-off between relaxation on group strategy-proofness and social benefit gain for obnoxious facilities.

Another option would be to not require any form of strategyproofness and allow the agents to engage in a *strategic game*; these games could be then studied using known techniques for analyzing *Nash equilibria*, and their performance via the celebrated notion of the *Price of Anarchy* [Koutsoupias and Papadimitriou, 1999]. For example, Hotelling games are proposed by Hotelling [1990] to model the situations where companies compete for attracting customers by strategically locating their facilities, and the customers prefer closer facilities. Recent works have been focusing on characterizing the equilibria and bounding the corresponding Price of Anarchy (see [Feldman *et al.*, 2016b; Ben-Porat and Tennenholtz, 2019]). In another dimension, Sui and Boutilier [2015b] study the computational problem of group manipulation: given a mechanism, how can the agents form collusion to manipulate the outcome, in order to minimize the social cost? When the mechanism is “Quantile”, they solve this problem for single-facility and multi-dimensional spaces by formulating as cone programs. The group manipulation for other types of mechanisms is also worthy being studied.

8 Concluding Remarks

From the abundance of different variants and interesting results that we have presented in this survey, it is evident that facility location problems are one of the most fundamental and most important examples of approximate mechanism design, and that it is at the very center of a vibrant literature in the major venues in artificial intelligence. And yet, despite the marvelous achievements that the literature on the problem has witnessed during the last decade, the future is nothing short of exciting. We believe that this survey will serve as a solid point of reference for current and future researchers in their quest to understand the major challenges, identify the most intriguing questions, and ultimately advance the state of the art even further.

Acknowledgements

The authors thank anonymous reviewers for the constructive comments. Bo Li is partially supported by The Hong Kong Polytechnic University (Grant No. P0034420). Minming Li is partially supported by NSFC (Grant No. 11771365).

References

- [Ahmadi-Javid *et al.*, 2017] A. Ahmadi-Javid, P. Seyedi, and S. S. Syam. A survey of healthcare facility location. *Comput. Oper. Res.*, 79:223–263, 2017.
- [Alon *et al.*, 2010] N. Alon, M. Feldman, A. D. Procaccia, and M. Tennenholtz. Strategyproof approximation of the minimax on networks. *Math. Oper. Res.*, 35(3):513–526, 2010.
- [Anastasiadis and Deligkas, 2018] E. Anastasiadis and A. Deligkas. Heterogeneous facility location games. In *AAMAS*, pages 623–631, 2018.
- [Anshelevich *et al.*, 2018] E. Anshelevich, O. Bhardwaj, E. Elkind, J. Postl, and P. Skowron. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264:27–51, 2018.
- [Archer and Tardos, 2001] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *FOCS*, pages 482–491, 2001.
- [Aziz *et al.*, 2020a] H. Aziz, H. Chan, B. Lee, B. Li, and T. Walsh. Facility location problem with capacity constraints: Algorithmic and mechanism design perspectives. In *AAAI*, pages 1806–1813, 2020.
- [Aziz *et al.*, 2020b] H. Aziz, H. Chan, B. E Lee, and D. C. Parkes. The capacity constrained facility location problem. *Games Econ. Behav.*, 124:478–490, 2020.
- [Babaiouff *et al.*, 2012] M. Babaiouff, S. Dobzinski, S. Oren, and A. Zohar. On bitcoin and red balloons. In *Proceedings of the 13th ACM conference on electronic commerce*, pages 56–73, 2012.
- [Babaiouff *et al.*, 2016] M. Babaiouff, M. Feldman, and M. Tennenholtz. Mechanism design with strategic mediators. *ACM Trans. Econ. Comput.*, 4(2):1–48, 2016.
- [Ben-Porat and Tennenholtz, 2019] O. Ben-Porat and M. Tennenholtz. Multiunit facility location games. *Math. Oper. Res.*, 44(3):865–889, 2019.
- [Black, 1948] D. Black. On the rationale of group decision-making. *J. Polit. Econ.*, 56(1):23–34, 1948.
- [Bordes *et al.*, 1990] G. Bordes, G. Laffond, and M. Le Breton. Strategy-proofness issues in some economic and political domains. *Unpublished Manuscript, University of Bordeaux*, 1990.
- [Brandt *et al.*, 2016] F. Brandt, V. Conitzer, U. Endriss, J.e Lang, and A. D. Procaccia. *Handbook of computational social choice*. Cambridge University Press, 2016.
- [Cai *et al.*, 2016] Q. Cai, A. Filos-Ratsikas, and P. Tang. Facility location with minimax envy. In *IJCAI*, pages 137–143, 2016.

- [Caragiannis *et al.*, 2012] I. Caragiannis, E. Elkind, M. Szegedy, and L. Yu. Mechanism design: from partial to probabilistic verification. In *EC*, pages 266–283, 2012.
- [Caragiannis *et al.*, 2016] I. Caragiannis, A. D. Procaccia, and N. Shah. Truthful univariate estimators. In *ICML*, pages 127–135, 2016.
- [Chen *et al.*, 2018] X. Chen, X. Hu, X. Jia, M. Li, Z. Tang, and C. Wang. Mechanism design for two-opposite-facility location games with penalties on distance. In *SAGT*, pages 256–260, 2018.
- [Chen *et al.*, 2019] X. Chen, M. Li, C.J. Wang, C Wang, and Y. Zhao. Truthful mechanisms for location games of dual-role facilities. In *AAMAS*, pages 1470–1478, 2019.
- [Chen *et al.*, 2021] X. Chen, X. Hu, Z. Tang, and C. Wang. Tight efficiency lower bounds for strategy-proof mechanisms in two-opposite-facility location game. *Inf. Process. Lett.*, 168:106098, 2021.
- [Cheng and Zhou, 2015] Y. Cheng and S. Zhou. A survey on approximation mechanism design without money for facility games. In *Advances in global optimization*, pages 117–128. Springer, 2015.
- [Cheng *et al.*, 2013] Y. Cheng, W. Yu, and G. Zhang. Strategy-proof approximation mechanisms for an obnoxious facility game on networks. *Theor. Comput. Sci.*, 497:154–163, 2013.
- [Church and Garfinkel, 1978] R. L. Church and R. S. Garfinkel. Locating an obnoxious facility on a network. *Transp. Sci.*, 12(2):107–118, 1978.
- [De Keijzer and Wojtczak, 2018] B. De Keijzer and D. Wojtczak. Facility reallocation on the line. In *IJCAI*, pages 188–194, 2018.
- [Deligkas *et al.*, 2021] A. Deligkas, A. Filos-Ratsikas, and A. Voudouris. Heterogeneous facility location with limited resources. *arXiv preprint arXiv:2105.02712*, 2021.
- [Ding *et al.*, 2020] Y. Ding, W. Liu, X. Chen, Q. Fang, and Q. Nong. Facility location game with envy ratio. *Comput. Ind. Eng.*, 148:106710, 2020.
- [Dokow *et al.*, 2012] E. Dokow, M. Feldman, R. Meir, and I. Nehama. Mechanism design on discrete lines and cycles. In *EC*, pages 423–440, 2012.
- [Escoffier *et al.*, 2011] B. Escoffier, L. Gourvès, N. K. Thang, F. Pascual, and O. Spanjaard. Strategy-proof mechanisms for facility location games with many facilities. In *ADT*, pages 67–81, 2011.
- [Feigenbaum and Sethuraman, 2015] I. Feigenbaum and J. Sethuraman. Strategyproof mechanisms for one-dimensional hybrid and obnoxious facility location models. In *AAAI Workshops*, 2015.
- [Feigenbaum *et al.*, 2017] I. Feigenbaum, J. Sethuraman, and C. Ye. Approximately optimal mechanisms for strategyproof facility location: Minimizing l_p norm of costs. *Mathematics of Operations Research*, 42(2):434–447, 2017.
- [Feldman and Wilf, 2013] M. Feldman and Y. Wilf. Strategyproof facility location and the least squares objective. In *EC*, pages 873–890, 2013.
- [Feldman *et al.*, 2016a] M. Feldman, A. Fiat, and I. Golomb. On voting and facility location. In *EC*, pages 269–286, 2016.
- [Feldman *et al.*, 2016b] M. Feldman, A. Fiat, and S. Obraztsova. Variations on the hotelling-downs model. In *AAAI*, volume 30, 2016.
- [Filimonov and Meir, 2021] A. Filimonov and R. Meir. Strategyproof facility location mechanisms on discrete trees. In *AAMAS*, pages 510–518, 2021.
- [Filos-Ratsikas and Voudouris, 2020] A. Filos-Ratsikas and A. A. Voudouris. Approximate mechanism design for distributed facility location. *arXiv preprint arXiv:2007.06304*, 2020.
- [Filos-Ratsikas *et al.*, 2015] A. Filos-Ratsikas, M. Li, J. Zhang, and Q. Zhang. Facility location with double-peaked preferences. In *AAAI*, pages 893–899, 2015.
- [Fong *et al.*, 2018] C. K. Ken Fong, M. Li, P. Lu, T. Todo, and M. Yokoo. Facility location games with fractional preferences. In *AAAI*, pages 1039–1046, 2018.
- [Fotakis and Tzamos, 2010] D. Fotakis and C. Tzamos. Winner-imposing strategyproof mechanisms for multiple facility location games. In *WINE*, pages 234–245, 2010.
- [Fotakis and Tzamos, 2014] D. Fotakis and C. Tzamos. On the power of deterministic mechanisms for facility location games. *ACM Trans. Econ. Comput.*, 2(4):15:1–15:37, 2014.
- [Fotakis and Tzamos, 2016] D. Fotakis and C. Tzamos. Strategyproof facility location for concave cost functions. *Algorithmica*, 76(1):143–167, 2016.
- [Fotakis *et al.*, 2021] D. Fotakis, L. Kavouras, P. Kostopaniotis, P. Lazos, S. Skoulakis, and N. Zarifis. Reallocating multiple facilities on the line. *Theoretical Computer Science*, 2021.
- [Friedman *et al.*, 2001] J. Friedman, T. Hastie, and R. Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics New York, 2001.
- [Fukui *et al.*, 2020] Y. Fukui, A. Shurbevski, and H. Nagamochi. Group strategy-proof mechanisms for shuttle facility games. *Journal of Information Processing*, 28:976–986, 2020.
- [Golomb and Tzamos, 2017] I. Golomb and C. Tzamos. Truthful facility location with additive errors. *arXiv preprint arXiv:1701.00529*, 2017.
- [Golowich *et al.*, 2018] N. Golowich, H. Narasimhan, and David C. Parkes. Deep learning for multi-facility location mechanism design. In *IJCAI*, pages 261–267, 2018.
- [Hochbaum, 1982] D. S. Hochbaum. Heuristics for the fixed cost median problem. *Math. Program.*, 22(1):148–162, 1982.

- [Hossain *et al.*, 2020] S. Hossain, E. Micha, and N. Shah. The surprising power of hiding information in facility location. In *AAAI*, pages 2168–2175, 2020.
- [Hotelling, 1990] H. Hotelling. Stability in competition. In *The collected economics articles of Harold Hotelling*, pages 50–63. Springer, 1990.
- [Ibara and Nagamochi, 2012] K. Ibara and H. Nagamochi. Characterizing mechanisms in obnoxious facility game. In *COCOA*, pages 301–311, 2012.
- [Koutsoupias and Papadimitriou, 1999] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *STACS*, pages 404–413, 1999.
- [Kyropoulou *et al.*, 2019] M. Kyropoulou, C. Ventre, and X. Zhang. Mechanism design for constrained heterogeneous facility location. In *SAGT*, pages 63–76. Springer, 2019.
- [Li *et al.*, 2019a] M. Li, P. Lu, Y. Yao, and J. Zhang. Strategyproof mechanism for two heterogeneous facilities with constant approximation ratio. In *IJCAI*, pages 238–245, 2019.
- [Li *et al.*, 2019b] M. Li, L. Mei, Y. Xu, G. Zhang, and Y. Zhao. Facility location games with externalities. In *AAMAS*, pages 1443–1451, 2019.
- [Li *et al.*, 2020] M. Li, C. Wang, and M. Zhang. Budgeted facility location games with strategic facilities. In *IJCAI*, pages 400–406, 2020.
- [Liu *et al.*, 2021] W. Liu, Y. Ding, X. Chen, Q. Fang, and Q. Nong. Multiple facility location games with envy ratio. *Theor. Comput. Sci.*, 2021.
- [Lu *et al.*, 2009] P. Lu, Y. Wang, and Y. Zhou. Tighter bounds for facility games. In *WINE*, pages 137–148, 2009.
- [Lu *et al.*, 2010] P. Lu, X. Sun, Y. Wang, and Z. A. Zhu. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *EC*, pages 315–324, 2010.
- [Mei *et al.*, 2019] L. Mei, M. Li, D. Ye, and G. Zhang. Facility location games with distinct desires. *Discrete Applied Mathematics*, 264:148–160, 2019.
- [Meir, 2019] R. Meir. Strategyproof facility location for three agents on a circle. In *SAGT*, pages 18–33, 2019.
- [Menon and Larson, 2019] V. Menon and K. Larson. Mechanism design for locating a facility under partial information. In *SAGT*, pages 49–62, 2019.
- [Moulin, 1980] H. Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.
- [Narasimhan *et al.*, 2016] H. Narasimhan, S. Agarwal, and D. C. Parkes. Automated mechanism design without money via machine learning. In *IJCAI*, pages 433–439, 2016.
- [Nehama *et al.*, 2019] I. Nehama, T. Todo, and M. Yokoo. Manipulations-resistant facility location mechanisms for z -line graphs. In *AAMAS*, pages 1452–1460, 2019.
- [Nissim *et al.*, 2012] K. Nissim, R. Smorodinsky, and M. Tennenholtz. Approximately optimal mechanism design via differential privacy. In *ITCS*, pages 203–213, 2012.
- [Ono *et al.*, 2017] T. Ono, T. Todo, and M. Yokoo. Rename and false-name manipulations in discrete facility location with optional preferences. In *PRIMA*, pages 163–179, 2017.
- [Oomine *et al.*, 2017] M. Oomine, A. Shurbevski, and H. Nagamochi. Parameterization of strategy-proof mechanisms in the obnoxious facility game. *J. Graph Algorithms Appl.*, 21(3):247–263, 2017.
- [Procaccia and Tennenholtz, 2009] A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. In *EC*, pages 177–186, 2009.
- [Procaccia *et al.*, 2018] A. D. Procaccia, D. Wajc, and H. Zhang. Approximation-variance tradeoffs in facility location games. In *AAAI*, 2018.
- [Schummer and Vohra, 2002] J. Schummer and R. V. Vohra. Strategy-proof location on a network. *J. Econ. Theory*, 104(2):405–428, 2002.
- [Serafino and Ventre, 2014] P. Serafino and C. Ventre. Heterogeneous facility location without money on the line. In *ECAI*, pages 807–812, 2014.
- [Serafino and Ventre, 2015] P. Serafino and C. Ventre. Truthful mechanisms without money for non-utilitarian heterogeneous facility location. In *AAAI*, pages 1029–1035, 2015.
- [Sonoda *et al.*, 2016] A. Sonoda, T. Todo, and M. Yokoo. False-name-proof locations of two facilities: Economic and algorithmic approaches. In *AAAI*, volume 30, 2016.
- [Sui and Boutilier, 2015a] X. Sui and C. Boutilier. Approximately strategy-proof mechanisms for (constrained) facility location. In *AAMAS*, pages 605–613, 2015.
- [Sui and Boutilier, 2015b] X. Sui and C. Boutilier. Optimal group manipulation in facility location problems. In *ADT*, pages 505–520, 2015.
- [Sui *et al.*, 2013] X. Sui, C. Boutilier, and T. Sandholm. Analysis and optimization of multi-dimensional percentile mechanisms. In *IJCAI*, pages 367–374, 2013.
- [Tang *et al.*, 2020a] P. Tang, D. Yu, and S. Zhao. Characterization of group-strategyproof mechanisms for facility location in strictly convex space. In *EC*, pages 133–157, 2020.
- [Tang *et al.*, 2020b] Z. Tang, C. Wang, M. Zhang, and Y. Zhao. Mechanism design for facility location games with candidate locations. In *COCOA*, pages 440–452. Springer, 2020.
- [Todo *et al.*, 2011] T. Todo, A. Iwasaki, and M. Yokoo. False-name-proof mechanism design without money. In *AAMAS*, pages 651–658, 2011.
- [Todo *et al.*, 2020] T. Todo, N. Okada, and M. Yokoo. False-name-proof facility location on discrete structures. In *ECAI*, pages 227–234, 2020.

- [Wada *et al.*, 2018] Y. Wada, T. Ono, T. Todo, and M. Yokoo. Facility location with variable and dynamic populations. In *AAMAS*, pages 336–344, 2018.
- [Walsh, 2020a] T. Walsh. Strategy proof mechanisms for facility location at limited locations. *arXiv preprint arXiv:2009.07982*, 2020.
- [Walsh, 2020b] T. Walsh. Strategy proof mechanisms for facility location in Euclidean and Manhattan space. *arXiv preprint arXiv:2009.07983*, 2020.
- [Xu *et al.*, 2020] X. Xu, M. Li, and L. Duan. Strategyproof mechanisms for activity scheduling. In *AAMAS*, pages 1539–1547, 2020.
- [Xu *et al.*, 2021] X. Xu, B. Li, M. Li, and L. Duan. Two-facility location games with minimum distance requirement. *J. Artif. Intell. Res.*, 70, 2021.
- [Yan and Chen, 2020] X. Yan and Y. Chen. Strategyproof facility location mechanisms with richer action spaces. *arXiv preprint arXiv:2002.07889*, 2020.
- [Ye *et al.*, 2015] D. Ye, L. Mei, and Y. Zhang. Strategy-proof mechanism for obnoxious facility location on a line. In *COCOON*, pages 45–56, 2015.
- [Yuan *et al.*, 2016] H. Yuan, K. Wang, C. K. K. Fong, Y. Zhang, and M. Li. Facility location games with optional preference. In *ECAI*, pages 1520–1527, 2016.
- [Zhang and Li, 2014] Q. Zhang and M. Li. Strategyproof mechanism design for facility location games with weighted agents on a line. *J. Comb. Optim.*, 28(4):756–773, 2014.
- [Zou and Li, 2015] S. Zou and M. Li. Facility location games with dual preference. In *AAMAS*, pages 615–623, 2015.