Mechanism Design with Predictions for Facility Location Games with Candidate Locations

Mechanism Design with Predictions for Facility Location Games with Candidate Locations *

Jiazhu Fang[1,2] · Qizhi Fang[1,2] · Wenjing Liu[1,2] ·
Qingqin Nong[1,2] · Alexandros A. Voudouris[3]

Abstract We study mechanism design with predictions in the single (obnoxious) facility location games with candidate locations on the real line, which complements the existing literature on mechanism design with predictions. We first consider the single facility location games with candidate locations, where each agent prefers the facility (e.g., a school) to be located as close to her as possible. We study two social objectives: minimizing the maximum cost and the social cost, and provide deterministic, anonymous, and group strategy-proof mechanisms with predictions that achieve the best possible trade-offs between consistency and robustness, respectively. Additionally, we represent the approximation ratio as a function of the prediction error, indicating that mechanisms can achieve better performance even when predictions are not fully accurate. We also consider the single obnoxious facility location games with candidate locations, where each agent prefers the facility (e.g., a garbage transfer station) to be located as far away from her as possible. For the objective of maximizing the minimum utility, we prove that any strategy-proof mechanism with predictions is unbounded robust. For the objective of maximizing the social utility, we provide a deterministic, anonymous, and group strategy-proof mechanism with prediction that achieves the best possible trade-off between consistency and robustness.

Keywords Mechanism design with predictions · Facility location · Candidate locations · Consistency · Robustness

---

* A preliminary version of this paper appeared in *TAMC 2024* [10].

✉ Wenjing Liu
E-mail: liuwj@ouc.edu.cn

Jiazhu Fang
E-mail: fjz@stu.ouc.edu.cn

Qizhi Fang
E-mail: qfang@ouc.edu.cn

Qingqin Nong
E-mail: qqnong@ouc.edu.cn

Alexandros A. Voudouris
E-mail: alexandros.voudouris@essex.ac.uk

[1] School of Mathematical Sciences, Ocean University of China, Qingdao, 266100, Shandong Province, China

[2] Laboratory of Marine Mathematics, Ocean University of China, Qingdao, Shandong 266100, China

[3] School of Computer Science and Electronic Engineering, University of Essex, United Kingdom

1 Introduction

The design of *algorithms with imperfect predictions* has recently become a highly active research direction, laying the foundation for exploring methods that go beyond *worst-case analysis*, which is a traditional approach commonly used to analysis algorithm performance in computer science. While worst-case analysis can ensure the robustness of the algorithm, the worst-case scenario often does not occur in practice, which limits the design of higher-performing algorithms. Moreover, the lower bounds of the worst-case are often determined by uncertain information, such as unknown future information in online problems [5] and private information of agents in mechanism design problems [24]. Therefore, introducing prediction about uncertain information provides the possibility to break through the lower bounds.

In the framework of algorithmic design with imperfect predictions, it is generally assumed that machine learning methods can provide predictions about unknown information, which algorithm designers can utilize to design algorithms. The goal of such algorithms is to achieve good performance that surpasses the worst-case bound, when the predictions are accurate, and to ensure that their performance does not degrade significantly compared to the best algorithms without predictions, even when the predictions are arbitrarily inaccurate. Lykouris and Vassilvitskii [21] formally introduced the terms *consistency* and *robustness* to represent the competitive ratios of an online algorithm with respect to accurate and arbitrarily inaccurate predictions respectively. These terms have now become standard for analyzing the performance of algorithms with imperfect predictions. Much of the related work has focused on online algorithm design, and many classic online problems have been considered, such as ski rental [16, 25], caching [26], and scheduling [17, 19].

Mechanism design and online algorithm design are similar in that they both involve unknown information. Due to the private information held by strategic agents, the strategy-proof (no agent can benefit from misreporting) mechanism generally can only produce approximately optimal solutions. Introducing predictions could potentially enable us to obtain an optimal strategy-proof mechanism. Therefore, we hope to design an ideal mechanism using predictions, which can achieve optimal performance when predictions are accurate, and can capture the performance of the best-known mechanism without predictions when predictions are arbitrarily inaccurate, thus achieving the best of both worlds. However, in many cases, mechanisms may rely on predictions to achieve the optimal performance, leading to unbounded robustness when predictions are inaccurate. Alternatively, we hope that mechanisms can approach the above-mentioned ideal performance as closely as possible, achieving the best possible trade-off between consistency and robustness. The application of predictions to mechanism design was first introduced by Agrawal et al. [1], who considered the classical mechanism design problem without money – *facility location*, and subsequently, Xu and Lu [28] studied other mechanism design problems under predictions, both showing that mechanisms with predictions can achieve a good (optimal) trade-off between robustness and consistency.

In the framework of mechanism design with predictions, we study the facility location games with candidate locations [14, 27], which complements the results of mechanism design with predictions. In the classic facility location problem, facilities can be placed anywhere in the metric space. However, in reality, due to various constraints, most facilities can only be built in specific locations. For example, schools can only be built in spacious areas, while garbage transfer stations can only be built in specific locations due to wind direction and surrounding residential areas. Note that the facility location problem with candidate locations is more general than its counterpart without candidate locations, as the latter corresponds to the special case where all points in the metric space are considered as candidate locations.

1.1 Related Work

Algorithm (Mechanism) with Predictions. Early relevant work on algorithm design with predictions can be found in this survey [22]. Lykouris and Vassilvitskii [21] formally introduced consistency and robustness as two main indicators for measuring the performance of algorithms with predictions. Subsequently, a series of classic problems have been studied, such as secretary problems [8], online graph problems [2], maximum flow [23], shortest path [11] as well as k-means clustering [9] and Nash social welfare maximization [4]. Here, we mainly introduce the literature on mechanism design with predictions. Agrawal et al. [1] first introduced this framework into the field of mechanism design and studied facility location games in two-dimensional Euclidean space. For the maximum cost objective, they provided a 1-consistent and

$(1+\sqrt{2})$-robust strategy-proof mechanism, and for the social cost objective, they introduced a confidence value parameter $c \in (0,1)$ and provided a $\frac{\sqrt{2c^2+2}}{1+c}$-consistent and $\frac{\sqrt{2c^2+2}}{1-c}$-robust strategy-proof mechanism, achieving an optimal trade-off between consistency and robustness. They also provided the optimal 1-consistent and 2-robust strategy-proof mechanism for the one-dimensional case under the maximum cost objective. Xu and Lu [28] studied four classic mechanism design problems, and improved the known approximation ratio using predictions for the two facility location game on the real line. Istrate and Bonchis [18] considered the obnoxious facility location problem and achieved trade-offs between robustness and consistency on line segments, squares, circles, and trees. Subsequently, Gkatzelis et al. [15] extended the prediction framework to decentralized mechanism design in strategic settings. Balkanski et al. [3] studied classic strategic scheduling problems, employing predictions to provide a 6-consistent and $2n$-robust strategy-proof mechanism. For more related literature, please refer to the webpage of algorithms with predictions [20].

Facility Location Games. Due to the extensive literature on facility location problems, we only provide a brief review of the most relevant literature. Procaccia and Tennenholtz [24] proposed approximate mechanism design without money for facility location games, which uses the approximation ratio to measure the performance of the strategy-proof mechanism, thereby initiating a research trend in strategic facility location. Tang et al. [27] considered the single and two facility location games with candidate locations under social cost and maximum cost objective. For the single facility location game with maximum cost objective, they proposed a deterministic 3-approximate group strategy-proof (no coalition of agents can benefit from misreporting) mechanism and proved that there are no deterministic (or randomized) strategy-proof mechanism with approximation ratio better than 3 (or 2). Feldman et al. [12] studied the relationship among three types of candidate selection mechanisms: voting, ranking and location mechanisms. They proved that under social cost, the median mechanism that places the facility at the candidate location closest to the median agent is strategy-proof and 3-approximate, and no deterministic strategy-proof mechanism can perform better. Gai et al. [14] studied the single and two obnoxious facility location games with social utility objective on the real line. For the single obnoxious facility location game, they provided a deterministic 3-approximate group strategy-proof mechanism, which is also the best possible strategy-proof mechanism. For further literature on facility location, please refer to [7].

1.2 Our Contribution

Facility Location Game with Candidate Locations. For the single facility location game with candidate locations on the real line, we consider two social objectives: minimizing the maximum cost and minimizing the social cost. For the maximum cost objective, we present a deterministic, anonymous (the output of the mechanism does not depend on the identities of the agents) and group strategy-proof mechanism with predictions, which is 1-consistent and 3-robust. Based on the lower bound of 3 for any deterministic strategyproof mechanism provided by Tang et al. [27], our results achieve the best trade-off between consistency and robustness. More generally, we represent the approximation ratio of the mechanism as a function of the prediction error $\delta \geq 0$, obtaining that it is $\min\{1+\delta, 3\}$-approximate, which corresponds to the optimal mechanism when $\delta = 0$ and gradually increases linearly to 3-approximate. We also analyze its performance under the social cost objective, which is 1-consistent and $(2n-1)$-robust. By introducing the confidence value parameter γ proposed by Agrawal et al. [1] to reduce reliance on predictions, we present a deterministic, anonymous and group strategy-proof mechanism which is $\frac{3-\gamma}{1+\gamma}$-consistent and $\frac{3+\gamma}{1-\gamma}$-robust under social cost, where $\gamma \in [0,1)$ can be adjusted based on the degree of trust in the prediction accuracy. If the designer is not confident in the prediction, setting $\gamma = 0$ will yield a 3-consistent and 3-robust mechanism that matches the best performance guarantee without predictions [12]. As γ gradually increases from 0 to 1, the consistency decreases from 3 to 1, achieving the optimal performance, but at the cost of increasing the robustness. We also show that this the best possible trade-off between consistency and robustness achievable by any deterministic strategyproof mechanism. Finally, We represent the approximation ratio of the mechanism as a function of the prediction error δ, obtaining that it is $\min\{\frac{3-\gamma}{1+\gamma}+\delta, \frac{3+\gamma}{1-\gamma}\}$-approximate, which starts at $\frac{3-\gamma}{1+\gamma}$-approximate when $\delta = 0$, corresponding to the consistency guarantee, and gradually increases linearly to $\frac{3+\gamma}{1-\gamma}$-approximate, corresponding to the robustness guarantee.

Obnoxious Facility Location Game with Candidate Locations. For the single obnoxious facility location game with candidate locations on the real line, we also consider two social objectives: maximizing the minimum utility and maximizing the social utility. For the minimum utility objective, we show that any strategy-proof mechanism with predictions is unbounded robust. We mainly focus on the social utility objective. By regarding the prediction as the location chosen by γn virtual voters and adopting a voting mechanism for the leftmost and rightmost candidate points, we obtain a deterministic, anonymous and group strategy-proof mechanism which is $\frac{3-\gamma}{1+\gamma}$-consistent and $\frac{3+\gamma}{1-\gamma}$-robust; again, this is best possible over all deterministic strategyproof mechanisms. Since, in this case, the predictions of the mechanism can only result in either accurate or inaccurate outcomes, there is no need to analyze the mechanism's performance through prediction error when the predictions are not fully accurate.

1.3 Paper Organizations

In Section 2, we formally define our model and introduce necessary notations and definitions. In Section 3, two group strategy-proof mechanisms with predictions for the single facility location game with candidate locations are presented. In Section 4, a group strategy-proof mechanism with predictions for the single obnoxious facility location game with candidate locations is provided. Section 5 summarizes this paper and discusses some open problems.

2 Preliminaries

In the single (obnoxious) facility location game with candidate locations, let $N = \{1, \ldots, n\}$ be the set of agents located on the real line, and each agent $i \in N$ has a private location $x_i \in \mathbb{R}$. We use $\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{R}^n$ to denote the *location profile* of the n agents. Let $M \subseteq \mathbb{R}$ be a compact set representing the set of candidate locations for the facility. Let $a, b \in M$ be the leftmost candidate point and the rightmost candidate point, respectively. The distance between any two points $x, y \in \mathbb{R}$ is $d(x, y) = |x - y|$. Denote an instance by $I(\mathbf{x}, M)$ or simply by I without confusion. A *(deterministic) mechanism* is a function $f : \mathbb{R}^n \to M$ which maps a given location profile $\mathbf{x}$ to a facility location. A mechanism f is *anonymous*, if its outcome does not depend on identities of the agents, i.e., for every location profile $\mathbf{x}$ and every permutation of agents $\pi : N \to N$, $f(x_1, \ldots, x_n) = f(x_{\pi(1)}, \ldots, x_{\pi(n)})$. Here, we focus on anonymous mechanisms and assume without loss of generality that $x_1 \le \ldots \le x_n$ for every location profile $\mathbf{x}$.

Single facility location game with candidate locations. In the single facility location setting, each agent prefers the facility to be located as close to her as possible. Given an instance I and a facility location $y \in M$, the distance from the facility to agent $i \in N$ is considered as her cost, denoted by $\text{cost}(x_i, y) = d(x_i, y)$. We study two standard social objectives: minimizing the maximum cost and minimizing the social cost. The *maximum cost* of a facility location $y \in M$ with respect to the profile $\mathbf{x} \in \mathbb{R}^n$ is the maximum distance to all agents, denoted by

$$MC(\mathbf{x}, y) = \max_{i \in N} \text{cost}(x_i, y) = \max_{i \in N} d(x_i, y),$$

and the *social cost* of y with respect to $\mathbf{x}$ is the total distance to all agents, denoted by

$$SC(\mathbf{x}, y) = \sum_{i \in N} \text{cost}(x_i, y) = \sum_{i \in N} d(x_i, y).$$

Single obnoxious facility location game with candidate locations. In the single obnoxious facility location setting, each agent prefers the facility to be located as far away from her as possible, and the distance from the facility to agent $i \in N$ is interpreted as her utility, denoted by $\text{utility}(x_i, y) = d(x_i, y)$. We also consider two social objectives: maximizing the minimum utility and maximizing the social utility. The *minimum utility* of a facility location $y \in M$ with respect to the profile $\mathbf{x} \in \mathbb{R}^n$ is the minimum distance to all agents, denoted by

$$MU(\mathbf{x}, y) = \min_{i \in N} \text{utility}(x_i, y) = \min_{i \in N} d(x_i, y),$$

and the *social utility* of y with respect to $\mathbf{x}$ is the total distance to all agents, denoted by

$$SU(\mathbf{x}, y) = \sum_{i \in N} \text{utility}(x_i, y) = \sum_{i \in N} d(x_i, y).$$

The following definitions are only explicitly given in the single facility location setting and they are analogous in the obnoxious setting.

Strategyproofness. In the strategic setting, each agent may be incentivized to misreport to reduce her cost. A mechanism f is *strategy-proof* (SP) if no agent can benefit from misreporting, regardless of the reports of the others, that is, for every $i \in N$,

$$\text{cost}\left(x_i, f(x_i, \mathbf{x}_{-i})\right) \leq \text{cost}\left(x_i, f(x'_i, \mathbf{x}_{-i})\right),$$

for every $x_i, x'_i \in \mathbb{R}$ and $\mathbf{x}_{-i} = (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n) \in \mathbb{R}^{n-1}$.

A mechanism f is *group strategy-proof* (GSP), if no coalition of agents can decrease each members cost from misreporting, regardless of the reports of the other agents, that is, for every $G \subseteq N$, every $\mathbf{x}_G = (x_i)_{i \in G}, \mathbf{x}'_G = (x'_i)_{i \in G} \in \mathbb{R}^{|G|}$ and $\mathbf{x}_{-G} = (x_i)_{i \notin G} \in \mathbb{R}^{|N \setminus G|}$, there exists $i \in N$ such that

$$\text{cost}\left(x_i, f(\mathbf{x}_G, \mathbf{x}_{-G})\right) \leq \text{cost}\left(x_i, f(\mathbf{x}'_G, \mathbf{x}_{-G})\right).$$

Approximation ratio. Given a social objective function C to be minimized and an instance $I(\mathbf{x}, M)$, denote by $o(I)$ and $OPT(I)$ the optimal facility location and the optimum value respectively, or simply by o and OPT when I is clear from the context. A mechanism f achieves an *approximation ratio* of $\alpha \geq 1$ if for every instance $I(\mathbf{x}, M)$,

$$C(\mathbf{x}, f(\mathbf{x})) \leq \alpha \cdot C(\mathbf{x}, o) = \alpha \cdot OPT(I).$$

The framework of mechanism design with predictions. In the problem of mechanism design with predictions, the private location for each agent is predicted. When the predicted location profile $\hat{\mathbf{x}} = (\hat{x}_1, \ldots, \hat{x}_n)$ is obtained, it is easy to compute the optimal location $\hat{o}$ for $\hat{\mathbf{x}}$. Therefore, we directly design mechanisms using the prediction $\hat{o}$, corresponding to the optimal location o for $\mathbf{x}$. A (deterministic) *mechanism with predictions* is a function $f : \mathbb{R}^n \times M \to M$ which maps a given location profile $\mathbf{x}$ and a predicted optimal location $\hat{o}$ to a facility location. Similarly, we stipulate that the mechanism with predictions, denoted as $f(\mathbf{x}, \hat{o})$, needs to be (G)SP, and use consistency and robustness to measure its performance. Given a social objective function C to be minimized, a mechanism f is *α-consistent* if it achieves an approximation ratio of α when the prediction is accurate ($\hat{o} = o$), i.e., for every instance $I(\mathbf{x}, M)$,

$$C(\mathbf{x}, f(\mathbf{x}, o)) \leq \alpha \cdot C(\mathbf{x}, o) = \alpha \cdot OPT(I).$$

A mechanism f is *β-robust* if it achieves an approximation ratio of β even when the prediction is arbitrarily inaccurate, i.e., for every instance $I(\mathbf{x}, M)$,

$$\max_{\hat{o}} C(\mathbf{x}, f(\mathbf{x}, \hat{o})) \leq \beta \cdot C(\mathbf{x}, o) = \beta \cdot OPT(I).$$

To achieve a smooth trade-off between consistency and robustness, we represent the mechanism's approximation ratio as a function of the *prediction error* to observe the performance of mechanisms when the prediction is not fully accurate. Given an instance $I(\mathbf{x}, M)$, we define the *prediction error* as $\delta(\hat{o}, I) = \frac{d(\hat{o}, o)}{OPT(I)}$ when the social objective function is minimizing the maximum cost, representing the distance between the predicted optimal value $\hat{o}$ and the true optimal value o, normalized by the optimal value $OPT(I)$. For the social cost objective, define the *prediction error* as $\delta(\hat{o}, I) = \frac{d(\hat{o}, o)}{OPT(I)/n}$, representing the distance between the predicted optimal value $\hat{o}$ and the true optimal value o, normalized by $OPT(I)/n$.

A mechanism f is called *$\alpha(\delta)$-approximate*, if given an upper bound δ on the prediction error, it achieves an approximation ratio of $\alpha(\delta)$ when the prediction error of every instance does not exceed δ, i.e., for every instance $I(\mathbf{x}, M)$,

$$\max_{I, \hat{o}:\ \delta(\hat{o}, I) \leq \delta} C(\mathbf{x}, f(\mathbf{x}, \hat{o})) \leq \alpha(\delta) \cdot C(\mathbf{x}, o) = \alpha(\delta) \cdot OPT(I).$$

It can be observed that when $\delta = 0$, the approximation ratio of the mechanism corresponds to the consistency guarantee. As $\delta \to \infty$, it corresponds to the robustness guarantee. For $\delta \in (0, \infty)$, if $\alpha(\delta)$, as a function of δ, does not grow too rapidly, we can still achieve better approximation guarantees compared to the case without predictions.

The goal is to design (G)SP mechanisms that can capture an optimal solution when the prediction is accurate and matches the best-known results without predictions when the prediction is arbitrarily inaccurate. Furthermore, we also consider the mechanism's approximation guarantees when the prediction is not fully accurate.

3 Single Facility Location Game

In this section, we study the single facility location game with candidate locations on the real line, where each agent wants the facility (e.g., a school) to be built as close to her as possible, and the distance from the facility to her is considered as her cost. We focus on two social objectives: minimizing the maximum cost and minimizing the social cost.

3.1 Maximum Cost

For the objective of minimizing the maximum cost, Tang et al. [27] proposed a simple deterministic 3-approximate GSP mechanism without predictions that chooses the candidate location closest to the leftmost agent, and proved that it is also the best deterministic SP mechanism without predictions. In this subsection, we first present a deterministic, anonymous and GSP mechanism with predictions that is 1-consistent, capturing an optimal solution when the prediction is accurate, and 3-robust, matching the best approximation guarantee without predictions, thereby achieving the best of both worlds.

Mechanism 1

Input: A reported location profile $\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{R}^n$; Candidate locations $M \subseteq \mathbb{R}$; A prediction $\hat{o} \in M$.
Output: A facility location $y \in M$.
if $\hat{o} \in [x_1, x_n]$ **then**
return $y = \hat{o}$.
else if $\hat{o} < x_1$ **then**
return $y \in \arg\min_{g \in M} |x_1 - g|$, breaking ties in any deterministic way.
else
return $y \in \arg\min_{g \in M} |x_n - g|$, breaking ties in any deterministic way.
end if

Given the predicted optimal location $\hat{o}$ and the reported location profile $\mathbf{x} = (x_1, \ldots, x_n)$, we determine the facility location of by comparing the location relationship between $\hat{o}$ and $\mathbf{x}$, as illustrated by Mechanism 1. Specifically, when $\hat{o}$ falls within the range $[x_1, x_n]$, indicating that the prediction is not significantly inaccurate, and we directly output location $\hat{o}$. When $\hat{o}$ falls outside the range $[x_1, x_n]$, signifying a highly inaccurate prediction, we output the candidate location closest to either x_1 or x_n based on the location of $\hat{o}$, breaking ties in any deterministic way. In the following, we analyze the performance of the Mechanism 1.

Lemma 1 *Mechanism 1 is anonymous and GSP.*

Proof Mechanism 1 is clearly anonymous, because its outcome depends only on the reported locations and not on the identities of the agents.

We will now prove that for any group of agents G, there exists at least one member who cannot benefit from G's misreporting. First consider the case where the prediction $\hat{o} \in [x_1, x_n]$ when all the agents in G truthfully report their locations, implying that Mechanism 1 outputs $y = \hat{o}$. Note that if G's misreporting cannot change the output of Mechanism 1, then all agents in G cannot benefit from misreporting. If G's misreporting changes the output of Mechanism 1 to $y' < y$, G must contain all the agents with locations greater than or equal to $\hat{o}$, and they must misreport their locations less than $\hat{o}$. Obviously, these agents cannot benefit from misreporting. The proof for $y' > y$ is similar and omitted.

Now consider the case where $\hat{o} < x_1$ when all the agents in G truthfully report their locations, implying that Mechanism 1 outputs $y \in \arg\min_{g\in M} |x_1 - g|$. If G contains any agent located at x_1, then she cannot benefit from any misreporting by G, since she has already obtained the minimum possible cost. Otherwise, if G's misreporting can change the output of Mechanism 1 to y', G must contain some agent i with location $x_i > x_1$ but misreporting $x_i' < x_1$, which results in $y' \leq y$ and she cannot benefit. The proof for the case of $\hat{o} > x_n$ is similar and omitted. □

We further show that mechanism 1 is 1-consistent and 3-robust. According to the lower bound of 3 for any deterministic SP mechanism provided by Tang et al. [27], it is evident that the trade-off between consistency and robustness in our results is optimal.

Theorem 1 *For the single facility location game on the real line, Mechanism 1 is 1-consistent and 3-robust under the maximum cost objective.*

Proof We use f to denote Mechanism 1. For any instance $I(\mathbf{x}, M)$, let $y = f(\mathbf{x}, \hat{o})$. Denote by g_1, g_n the candidate location closest to x_1 and x_n respectively, that is, $g_1 \in \arg\min_{g\in M} |x_1 - g|$ and $g_n \in \arg\min_{g\in M} |x_n - g|$. Let m_x be the midpoint of x_1 and x_n, then the truly optimal location o for the maximum cost objective is the candidate location closest to m_x. Denote $L = d(x_1, x_n)$.

Consistency: When $\hat{o} = o \in [x_1, x_n]$, Mechanism 1 returns the optimal solution. We only need to consider the cases of $\hat{o} < x_1$ and $\hat{o} > x_n$. Assume w.l.o.g. that $\hat{o} = o < x_1$. In this case, the candidate location closest to x_1 is also the candidate location closest to m_x, hence Mechanism 1 returns the optimal solution as well.

Robustness: When $[x_1, x_n] \cap M = \emptyset$, both o and y are either g_1 or g_n. Assuming w.l.o.g. that $o = g_1$, we have

$$d(g_1, x_1) \leq d(g_n, x_n) \text{ and } OPT(I) = d(g_1, x_n).$$

If $y = g_1$, the optimal solution is obtained; otherwise, we have

$$MC(\mathbf{x}, y) = d(g_n, x_1) \leq d(g_n, g_1) = d(g_1, x_n) + d(x_n, g_n) \leq 2 \cdot d(g_1, x_n) = 2 \cdot OPT(I),$$

where the last inequality follows from the definition of g_n.

When $[x_1, x_n] \cap M \neq \emptyset$, we have

$$\frac{L}{2} \leq OPT(I) \leq L \text{ and } x_1 - L \leq y \leq x_n + L.$$

By symmetry, we only need to consider the case of $x_1 - L \leq y \leq m_x$.

Case 1: $x_1 - \frac{L}{2} \leq y \leq m_x$. In this case, we have

$$MC(\mathbf{x}, y) = d(x_n, y) \leq \frac{3}{2}L \leq 3 \cdot OPT(I).$$

Case 2: $x_1 - L \leq y < x_1 - \frac{L}{2}$. Here, $MC(\mathbf{x}, y) = d(x_n, y)$ and $OPT(I) \geq d(x_1, o)$. Thus, we have

$$MC(\mathbf{x}, y) \leq d(x_n, x_1) + d(x_1, y) \leq L + d(x_1, o) \leq L + OPT(I) \leq 3 \cdot OPT(I),$$

where the second inequality holds since y is the candidate location closest to x_1 by Mechanism 1. □

In order to evaluate the performance of Mechanism 1 when the prediction is not fully accurate, we represent the approximation ratio as a function of the prediction error δ, and obtain that Mechanism 1 is $\min\{1 + \delta, 3\}$-approximate.

Theorem 2 *For the single facility location game on the real line, Mechanism 1 achieves a* $\min\{1+\delta, 3\}$*-approximation under the maximum cost objective, where* δ *is an upper bound on the prediction error.*

Proof We use f to denote Mechanism 1. For any instance $I(\mathbf{x}, M)$, let $g(x_1)$ and $g(x_n)$ be the candidate locations closest to x_1 and x_n, respectively. Let $y = f(\mathbf{x}, \hat{o})$ and $y' = f(\mathbf{x}, \tilde{o})$ be the output locations of f when the predictions are $\hat{o}$ and $\tilde{o}$, respectively. We claim that $d(y, y') \leq d(\hat{o}, \tilde{o})$. Without loss of generality, we assume $\hat{o} \leq \tilde{o}$. We analyze the different locations of $\hat{o}$ and $\tilde{o}$.

– If $\hat{o} \leq \tilde{o} < x_1$ or $x_n < \hat{o} \leq \tilde{o}$, then under two different predictions, the output locations of f are the same, i.e., $d(y, y') = 0 \leq d(\hat{o}, \tilde{o})$.
– If $\hat{o} < x_1$ and $\tilde{o} \geq x_1$, then $\hat{o} \leq y = g(x_1) \leq \tilde{o}$, and $y \leq y' \leq \tilde{o}$. Therefore, $d(y, y') \leq d(\hat{o}, \tilde{o})$.
– If $\hat{o} \in [x_1, x_n]$ and $\tilde{o} > x_n$, then $y = \hat{o}$, and $\hat{o} \leq y' = g(x_n) \leq \tilde{o}$. Therefore, $d(y, y') \leq d(\hat{o}, \tilde{o})$.
– If $\hat{o}, \tilde{o} \in [x_1, x_n]$, then $y = \hat{o}$, and $y' = \tilde{o}$. Therefore, $d(y, y') = d(\hat{o}, \tilde{o})$.

In Theorem 1, we have already proven that f is 3-robust, which implies that the worst-case approximation ratio of f is at most 3. Now, we only need to prove that the approximation ratio of f is at most $1 + \delta$ when the prediction error of every instance is at most δ. Let $\tilde{o} = o$, and based on the definition of δ, we can obtain

$$d(f(\mathbf{x}, \hat{o}), f(\mathbf{x}, o)) \leq d(\hat{o}, o) = \delta \cdot OPT(I).$$

Furthermore, since we have already proven in Theorem 1 that f is 1-consistent, this implies that when the prediction is accurate, i.e., $\hat{o} = o$, f outputs the optimal location, i.e., $f(\mathbf{x}, o) = o$. Therefore,

$$\begin{aligned}
MC(\mathbf{x}, f(\mathbf{x}, \hat{o})) &= \max_{i \in N} d(x_i, f(\mathbf{x}, \hat{o})) \\
&\leq \max_{i \in N} (d(x_i, f(\mathbf{x}, o)) + d(f(\mathbf{x}, \hat{o}), f(\mathbf{x}, o))) \\
&\leq \max_{i \in N} d(x_i, o) + \delta \cdot OPT(I) \\
&= OPT(I) + \delta \cdot OPT(I) \\
&= (1 + \delta) \cdot OPT(I),
\end{aligned}$$

where the first inequality holds due to the triangle inequality. □

Mechanism 1 yields a smooth trade-off between consistency and robustness guarantees for the single facility location game with candidate locations under the maximum cost objective. It can be observed that it starts at 1-approximate when $\delta = 0$, corresponding to the consistency guarantee, and gradually increases linearly to 3-approximate, corresponding to the robustness guarantee. Furthermore, it indicates that when the prediction is not fully accurate, Mechanism 1 can provide better approximation guarantee than mechanisms without predictions.

3.2 Social Cost

In this subsection, we focus on the objective of minimizing the social cost. For the single facility location game with candidate locations on the real line, Feldman et al. [12] proved that the median mechanism without predictions, which returns the candidate location closest to the median agent, is deterministic, SP and 3-approximate, and no deterministic SP mechanism without predictions can achieve an approximation ratio better than 3. Can we break through this lower bound by introducing predictions, as under the maximum cost objective? We first analyze the performance of Mechanism 1 under social cost.

Theorem 3 *For the single facility location game on the real line, Mechanism 1 is 1-consistent and $(2n-1)$-robust under the social cost objective.*

Proof We use f to denote Mechanism 1. For any instance $I(\mathbf{x}, M)$, denote $y = f(\mathbf{x}, \hat{o})$ and $L = d(x_1, x_n)$. Obviously, it holds that $OPT \geq L$.

Consistency: The proof of 1-consistency here is exactly the same as that in Theorem 1 and is omitted.

Robustness: Assume w.l.o.g that $y < o$. Denote by S_1 and S_2 the sets of the agents whose locations are to the left and right of y respectively, that is, $S_1 = \{i : x_i \leq y\}$ and $S_2 = \{i : x_i > y\}$. Now we analyze the following cases.

Case 1: When $y < x_1$, we have $|S_1| = 0$ and $|S_2| = n$. The social cost of Mechanism 1 is

$$\begin{aligned} SC(\mathbf{x}, y) &= \sum_{i \in N} d(x_i, y) \\ &\leq \sum_{i \in S_2} d(x_i, x_1) + \sum_{i \in S_2} d(x_1, y) \\ &\leq \sum_{i \in S_2 \setminus \{1\}} d(x_n, x_1) + \sum_{i \in S_2} d(x_1, y) \\ &\leq (n-1) \cdot L + \sum_{i \in S_2} d(x_1, o) \\ &\leq (2n-1) \cdot OPT, \end{aligned}$$

where the third inequality holds since y is the candidate location closest to x_1 and the last inequality is due to $OPT \geq d(x_1, o)$.

Case 2: When $y \geq x_1$, we have $|S_1| \geq 1$ and $|S_2| \leq n - 1$. The social cost of Mechanism 1 is

$$SC(\mathbf{x}, y) = \sum_{i \in S_1} d(x_i, y) + \sum_{i \in S_2} d(x_i, y).$$

Denote the first term as Δ_1 and the second as Δ_2.

The optimal social cost is

$$OPT = \sum_{i \in S_1} d(x_i, o) + \sum_{i \in S_2} d(x_i, o).$$

Denote the first term as Σ_1 and the second as Σ_2.

For any agent $i \in S_1$, we have $x_i \leq y < o$, thus $\Delta_1 \leq \Sigma_1$. As to Δ_2, we have

$$\begin{aligned} \Delta_2 &= \sum_{i \in S_2} d(x_i, y) \\ &\leq \sum_{i \in S_2} d(x_i, o) + \sum_{i \in S_2} d(o, y) \\ &= \Sigma_2 + \sum_{i \in S_2} d(o, y) \\ &\leq \Sigma_2 + (n-1) \cdot \sum_{i \in S_1} d(o, y) \\ &\leq \Sigma_2 + (n-1) \cdot \sum_{i \in S_1} d(o, x_i) \\ &= \Sigma_2 + (n-1) \cdot \Delta_1, \end{aligned}$$

where the second inequality holds due to $|S_2| \leq (n-1) \cdot |S_1|$.

Summing up, we have

$$SC(\mathbf{x}, y) = \Delta_1 + \Delta_2 \leq \Sigma_1 + \Sigma_2 + (n-1) \cdot \Delta_1 \leq n \cdot \Sigma_1 + \Sigma_2 \leq n \cdot OPT.$$

□

Next we give an example to show that the analysis in Theorem 2 for the robustness of Mechanism 1 is tight.

Example 1 Consider an instance $I(\mathbf{x}, M)$ with $x_1 = 0, x_2 = \ldots = x_n = m > 0$ and $M = \{-m + \epsilon, m\}$, where $\epsilon > 0$ is a sufficiently small value. For this instance, $o = m$ and $OPT = m$. If $\hat{o} = -m + \epsilon$, then the outcome of Mechanism 1 is $-m + \epsilon$, and the social cost is

$$SC(\mathbf{x}, -m + \epsilon) = m - \epsilon + (n-1)(2m - \epsilon) = (2n-1)m - n\epsilon.$$

Then,

$$\frac{SC(\mathbf{x}, -m + \epsilon)}{OPT} = \frac{(2n-1)m - n\epsilon}{m} \to 2n - 1,$$

when ϵ tends to 0.

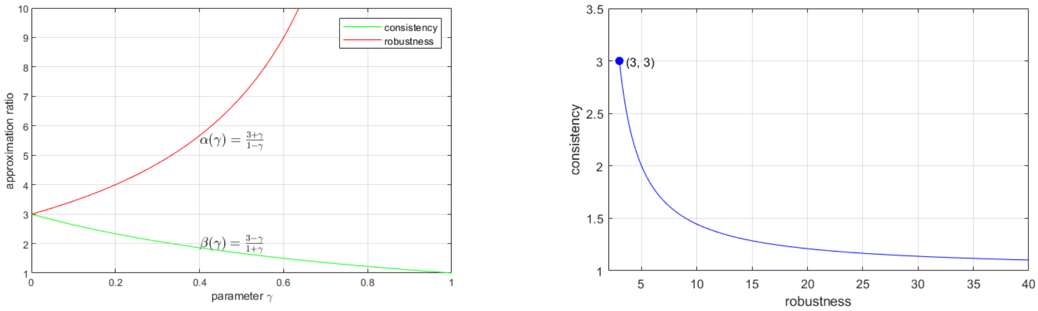

Fig. 1 (**a**): consistency and robustness as functions of the parameter γ. (**b**): the trade-off between consistency and robustness.

Notice that Mechanism 1 provides a poor robustness guarantee under social cost, as it heavily relies on the prediction to achieve 1-consistency. This can be clearly reflected in the equivalent statement of Mechanism 1 as follows: first create an $(n-1)$-dimensional virtual location profile $\mathbf{x}^v = (x_1^v, \dots, x_{n-1}^v)$ with $x_i^v = \hat{o}$ for $i \in \{1, \dots, n-1\}$, and then select the candidate location which is closest to the median of $(\mathbf{x}, \mathbf{x}^v)$.

To verify the equivalence, note that if $\hat{o} \in [x_1, x_n]$, then the candidate location closest to the median of $(\mathbf{x}, \mathbf{x}^v)$ is exactly $\hat{o}$. If $\hat{o} < x_1$ (or $\hat{o} > x_n$), the candidate locations closest to the median of $(\mathbf{x}, \mathbf{x}^v)$ are $\arg\min_{g \in M} |x_1 - g|$ (or $\arg\min_{g \in M} |x_n - g|$). Obviously, the output of Mechanism 1, which is the candidate location closest to the median of $(\mathbf{x}, \mathbf{x}^v)$, is largely biased towards the prediction $\hat{o}$. Therefore, when the prediction is inaccurate, the robustness of Mechanism 1 is compromised.

Mechanism 2

Input: A reported location profile $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$; Candidate locations $M \subseteq \mathbb{R}$; A prediction $\hat{o} \in M$; $\gamma \in [0, 1)$ such that γn is an integer.
Output: A facility location $y \in M$.
Create a γn-dimensional virtual location profile $\mathbf{x}^v = (x_1^v, \dots, x_{\gamma n}^v)$, where $x_i^v = \hat{o}$ for any $i = 1, \dots, \gamma n$.

Let x_m be the median of $(\mathbf{x}, \mathbf{x}^v)$, breaking ties in any deterministic way.
return $y \in \arg\min_{g \in M} |x_m - g|$, breaking ties in any deterministic way.

For the above reason, we introduce Mechanism 2, which reduces its reliance on the prediction. The high-level idea of Mechanism 2 is to regulate the reliance on the prediction by introducing a confidence value parameter $\gamma \in [0, 1)$, proposed by Agrawal et al. [1], such that γn is an integer. Specifically, we first create γn virtual agents located at $\hat{o}$. Subsequently, employing the idea of the median mechanism, we select the median agent among all agents, including virtual agents and real agents, and return the candidate location closest to the median agent, as illustrated by Mechanism 2. Note that we transform the degree of reliance on the prediction $\hat{o}$ into the number of virtual agents γn located at $\hat{o}$. An increasing γn implies that the output location of the mechanism gradually approaches the prediction $\hat{o}$. When $\gamma n = n - 1$, Mechanism 2 is equivalent to Mechanism 1.

We prove that Mechanism 2 with some constant $\gamma \in [0, 1)$, is a deterministic, anonymous and GSP mechanism, and it is $\frac{3-\gamma}{1+\gamma}$-consistent and $\frac{3+\gamma}{1-\gamma}$-robust. During the implementation of the mechanism, the mechanism designer can choose the value of γ based on her confidence in the prediction $\hat{o}$. If the designer is not confident in the prediction $\hat{o}$, setting $\gamma = 0$ will yield a 3-consistent and 3-robust mechanism that matches the best performance guarantee without predictions [12]. As γ gradually increases from 0 to 1, the consistency decreases from 3 to 1, achieving the optimal performance, but at the cost of increasing the robustness (see Fig.1).

Theorem 4 *For the single facility location game on the real line, Mechanism 2 with some constant $\gamma \in [0, 1)$ is deterministic, anonymous and GSP, and it is $\frac{3-\gamma}{1+\gamma}$-consistent and $\frac{3+\gamma}{1-\gamma}$-robust under the social cost objective.*

Proof Mechanism 2 is clearly anonymous, because its outcome only depends on the set of reported locations and not on the identities of the agents.

GSP: We now prove that for any group of agents G, there exists at least one member who cannot benefit from G's misreporting. Assuming all the agents in G truthfully report their locations, Mechanism 2 outputs y. If G contains any agent located at x_m, then she cannot benefit from any misreporting by G, since she has already obtained the minimum possible cost. Otherwise, first consider the case where G contains both agents with locations less than x_m and agents with locations greater than x_m. Note that if G's misreporting cannot change the output of Mechanism 2, then all agents in G cannot benefit from misreporting. If G's misreporting changes the output of Mechanism 2 to $y' > y$, G must contain some agent i with location $x_i < x_m$ but misreporting $x'_i > x_m$. Clearly, she cannot benefit from misreporting. The proof for $y' < y$ is similar and omitted.

Now consider the case where G only contains agents with locations less than x_m. If G's misreporting cannot change the output of Mechanism 2, then all agents in G cannot benefit from misreporting. If G's misreporting changes the output of Mechanism 2 to y', G must contain some agent i with location $x_i < x_m$ but misreporting $x'_i > x_m$, which results in $y' \geq y$ and she cannot benefit. The proof for the case where G only contains agents with locations larger than x_m is similar and omitted.

We use f to denote Mechanism 2. For any instance $I(\mathbf{x}, M)$, let $y = f(\mathbf{x}, \hat{o})$ be the outcome of Mechanism 2 and o be the optimal facility location. For each agent $i \in N$, denote by g_i the candidate location closest to x_i, that is, $g_i \in \arg\min_{g \in M} |x_i - g|$. We denote by S_1 and S_2 the sets of the agents whose closest candidate locations are to the left and right of y, respectively, that is $S_1 = \{i : g_i \leq y\}$ and $S_2 = \{i : g_i > y\}$. Then,

The social cost of Mechanism 2 is

$$SC(\mathbf{x}, y) = \sum_{i \in S_1} d(x_i, y) + \sum_{i \in S_2} d(x_i, y).$$

Denote the first term as Δ_1 and the second as Δ_2.

The optimal social cost is

$$OPT = \sum_{i \in S_1} d(x_i, o) + \sum_{i \in S_2} d(x_i, o).$$

Denote the first term as Σ_1 and the second as Σ_2.

Consistency: Assume w.l.o.g that $y < o$. In this case, $x_m < o$, since y is the candidate location closest to x_m. Since x_m is the median of $(\mathbf{x}, \mathbf{x}^v)$, then the total number of agents and virtual points whose locations are not greater than x_m is at least $\frac{(1+\gamma)n}{2}$. When the prediction is accurate, all the γn virtual points of $\mathbf{x}^v$ are located at o, i.e., $\hat{o} = o > x_m$. Therefore, the number of agents whose locations are not greater than x_m is at least $\frac{(1+\gamma)n}{2}$. Clearly, for these agents, the candidate locations closest to them cannot be greater than y. Hence, S_1 must contains at least these $\frac{(1+\gamma)n}{2}$ agents. Then, we have

$$|S_1| \geq \frac{(1+\gamma)n}{2} \text{ and } |S_2| \leq \frac{(1-\gamma)n}{2}.$$

For any agent $i \in S_1$, if $g_i < y$, then $x_i < y < o$; otherwise, by the definition of g_i, we have $d(x_i, y) \leq d(x_i, o)$. Therefore, $\Delta_1 \leq \Sigma_1$. As to Δ_2, we have

$$\begin{aligned}
\Delta_2 &= \sum_{i \in S_2} d(x_i, y) \\
&\leq \sum_{i \in S_2} d(x_i, o) + \sum_{i \in S_2} d(o, y) \\
&= \Sigma_2 + \sum_{i \in S_2} d(o, y) \\
&\leq \Sigma_2 + \frac{1-\gamma}{1+\gamma} \cdot \sum_{i \in S_1} d(o, y) \\
&\leq \Sigma_2 + \frac{1-\gamma}{1+\gamma} \cdot \left(\sum_{i \in S_1} d(o, x_i) + \sum_{i \in S_1} d(x_i, y)\right) \\
&= \Sigma_2 + \frac{1-\gamma}{1+\gamma} \cdot (\Delta_1 + \Sigma_1),
\end{aligned}$$

where the second inequality holds due to $|S_2| \leq \frac{1-\gamma}{1+\gamma} \cdot |S_1|$.

Summing up, we have

$$SC(\mathbf{x}, y) = \Delta_1 + \Delta_2 \leq \Sigma_1 + \Sigma_2 + \frac{1-\gamma}{1+\gamma}(\Delta_1 + \Sigma_1) \leq \frac{3-\gamma}{1+\gamma}\Sigma_1 + \Sigma_2 \leq \frac{3-\gamma}{1+\gamma}OPT.$$

Robustness: Without loss of generality, also assume that $y < o$. When the prediction is arbitrarily inaccurate, the total number of agents and virtual points whose locations are not greater than x_m is at least $\frac{(1+\gamma)n}{2}$, since x_m is the median of $(\mathbf{x}, \mathbf{x}^v)$. Since the number of these virtual points is at most γn, the number of agents whose locations are not greater than x_m is at least $\frac{(1-\gamma)n}{2}$. Clearly, for these $\frac{(1-\gamma)n}{2}$ agents, the candidate locations closest to them cannot be greater than y. Hence, S_1 must contains at least these $\frac{(1-\gamma)n}{2}$ agents. Then, we have

$$|S_1| \geq \frac{(1-\gamma)n}{2} \text{ and } |S_2| \leq \frac{(1+\gamma)n}{2}.$$

Obviously, $\Delta_1 \leq \Sigma_1$ still holds. As to Δ_2, we have

$$\begin{aligned}
\Delta_2 &= \sum_{i \in S_2} d(x_i, y) \\
&\leq \sum_{i \in S_2} d(x_i, o) + \sum_{i \in S_2} d(o, y) \\
&= \Sigma_2 + \sum_{i \in S_2} d(o, y) \\
&\leq \Sigma_2 + \frac{1+\gamma}{1-\gamma} \cdot \sum_{i \in S_1} d(o, y) \\
&\leq \Sigma_2 + \frac{1+\gamma}{1-\gamma} \cdot (\sum_{i \in S_1} d(o, x_i) + \sum_{i \in S_1} d(x_i, y)) \\
&= \Sigma_2 + \frac{1+\gamma}{1-\gamma} \cdot (\Delta_1 + \Sigma_1)
\end{aligned}$$

where the second inequality holds due to $|S_2| \leq \frac{1+\gamma}{1-\gamma} \cdot |S_1|$.

Summing up, we have

$$SC(\mathbf{x}, y) = \Delta_1 + \Delta_2 \leq \Sigma_1 + \Sigma_2 + \frac{1+\gamma}{1-\gamma}(\Delta_1 + \Sigma_1) \leq \frac{3+\gamma}{1-\gamma}\Sigma_1 + \Sigma_2 \leq \frac{3+\gamma}{1-\gamma}OPT.$$

□

Below, we prove that the trade-off between consistency and robustness achieved by Mechanism 2 is optimal.

Theorem 5 *For the single facility location game on the real line, the robustness of any deterministic SP mechanism with consistency strictly smaller than $\frac{3-\gamma}{1+\gamma}$ is at least $\frac{3+\gamma}{1-\gamma} - \delta$ for any $\gamma \in [0, 1)$ and $\delta > 0$, under the social cost objective function.*

Proof Consider an instance I with two candidate locations at ε and 2, and n agents such that $\frac{(1+\gamma)}{2}n - 1$ of them are positioned at 1 while the remaining $\frac{(1-\gamma)}{2}n + 1$ are positioned at 2, where $\varepsilon \in (0, 1)$. The social costs of the two possible locations 0 and 2 in this instance are

$$\begin{aligned}
SC(\mathbf{x}, \varepsilon) &= \left(\frac{1+\gamma}{2}n - 1\right)(1-\varepsilon) + \left(\frac{1-\gamma}{2}n + 1\right)(2-\varepsilon) \\
&= \frac{3-\gamma}{2}n + 1 - \varepsilon
\end{aligned}$$

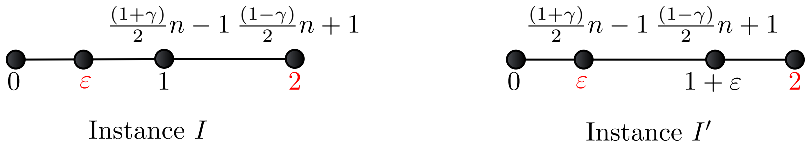

Fig. 2 The two instances of Theorem 5 and Theorem 9 to prove the lower bound.

and

$$SC(\mathbf{x}, 2) = \frac{1+\gamma}{2}n - 1.$$

Observe that the optimal solution is to place the facility at 2. So, when given a correct prediction about the optimal solution, if the mechanism chooses ε, it would achieve consistency $\frac{\frac{3-\gamma}{2}n+1-\varepsilon}{\frac{1+\gamma}{2}n-1} > \frac{3-\gamma}{1+\gamma}$, for any $\varepsilon < 1$. Since the mechanism is assumed to achieve a consistency ratio strictly smaller than that much, it must output the optimal solution when given I as input, thus it places the facility at 2.

Next, consider a sequence of instances in which each of the agents located at 1 are one-by-one moved to ε. Due to strategyproofness, the mechanism must place the facility at 2 in each such instance as otherwise, if the solution is ε, the agent that moves from 1 to ε would decrease their cost from 1 to $1-\varepsilon$. So, in the last instance of this sequence, we have $\frac{(1+\gamma)}{2}n - 1$ agents at ε and $\frac{(1-\gamma)n}{2} + 1$ agents at 2. Similarly, we now move one-by-one the agents at 2 to $1+\varepsilon$. The mechanism must still place the facility at 2 in all these instance as otherwise an agent would prefer to move back from $1+\varepsilon$ to 2. So, we finally have an instance I' in which there are $\frac{(1+\gamma)n}{2} - 1$ agents at ε and $\frac{(1-\gamma)n}{2} + 1$ agents at $1+\varepsilon$, and the mechanism places the facility at 2 (see Fig.2). We now have

$$SC(\mathbf{x}, \varepsilon) = \frac{1-\gamma}{2}n + 1$$

and

$$\begin{aligned} SC(\mathbf{x}, 2) &= \left(\frac{1+\gamma}{2}n - 1\right)(2-\varepsilon) + \left(\frac{1-\gamma}{2}n + 1\right)(1-\varepsilon) \\ &= \frac{3+\gamma}{2}n - 1 - \varepsilon. \end{aligned}$$

Hence, the robustness of the mechanism is

$$\frac{\frac{3+\gamma}{2}n - 1 - \varepsilon}{\frac{1-\gamma}{2}n + 1} = \frac{3+\gamma}{1-\gamma} - \delta,$$

where δ approaches 0 as n approaches infinity and ε approaches 0. □

We now analyze the performance of Mechanism 2 when the prediction is not fully correct via the prediction error.

Theorem 6 *For the single facility location game on the real line, Mechanism 2 with some constant $\gamma \in [0, 1)$ achieves a $\min\{\frac{3-\gamma}{1+\gamma} + \delta, \frac{3+\gamma}{1-\gamma}\}$-approximation under the social cost objective, where δ is an upper bound on the prediction error.*

Proof We use f to denote Mechanism 2. For any instance $I(\mathbf{x}, M)$, let $y = f(\mathbf{x}, \hat{o})$ and $y' = f(\mathbf{x}, \tilde{o})$ be the output locations of f when the predictions are $\hat{o}$ and $\tilde{o}$, respectively. We claim that $d(y, y') \le d(\hat{o}, \tilde{o})$. Without loss of generality, we assume $\hat{o} \le \tilde{o}$. We analyze different locations of y.

- If $y < \hat{o}$ or $y > \tilde{o}$, then when γn virtual locations change from $\hat{o}$ to $\tilde{o}$, the median location remains unchanged. Therefore, when the prediction is $\tilde{o}$, f still outputs location y, i.e., $y' = y$. Thus, $d(y, y') = 0 \le d(\hat{o}, \tilde{o})$.
- If $\hat{o} \le y \le \tilde{o}$, then the total number of agents with locations not exceeding $\tilde{o}$ and virtual locations not exceeding $\tilde{o}$ is at least $\lceil \frac{n+\gamma n}{2} \rceil$. Therefore, when the prediction is $\tilde{o}$, the output location of f will not exceed $\tilde{o}$, i.e., $y \le y' \le \tilde{o}$. Thus, $d(y, y') \le d(\hat{o}, \tilde{o})$.

In Theorem 4, we have already proven that f is $\frac{3+\gamma}{1-\gamma}$-robust, which implies that the worst-case approximation ratio of f is at most $\frac{3+\gamma}{1-\gamma}$. Now, we only need to prove that the approximation ratio of f is at most $\frac{3-\gamma}{1+\gamma}+\delta$ when the prediction error of every instance is at most δ. Let $\tilde{o}=o$, and based on the definition of δ, we can obtain

$$d(f(\mathbf{x},\hat{o}),f(\mathbf{x},o))\leq d(\hat{o},o)=\delta\cdot\frac{OPT(I)}{n}.$$

Furthermore, since we have already proven in Theorem 4 that f is $\frac{3-\gamma}{1+\gamma}$-consistent, this implies that when the prediction is accurate, i.e., $\hat{o}=o$, $SC(\mathbf{x},f(\mathbf{x},o))\leq\frac{3-\gamma}{1+\gamma}\cdot OPT(I)$ holds. Therefore,

$$\begin{aligned}SC(\mathbf{x},f(\mathbf{x},\hat{o}))&=\sum_{i\in N}d(x_i,f(\mathbf{x},\hat{o}))\\&\leq\sum_{i\in N}(d(x_i,f(\mathbf{x},o))+d(f(\mathbf{x},\hat{o}),f(\mathbf{x},o)))\\&\leq SC(\mathbf{x},f(\mathbf{x},o))+\delta\cdot OPT(I)\\&\leq\frac{3-\gamma}{1+\gamma}\cdot OPT(I)+\delta\cdot OPT(I)\\&=(\frac{3-\gamma}{1+\gamma}+\delta)\cdot OPT(I),\end{aligned}$$

where the first inequality holds due to the triangle inequality. □

For any $\gamma\in[0,1)$, Mechanism 2 yields a smooth trade-off between consistency and robustness guarantees for the single facility location game with candidate locations under the social cost objective. It starts at $\frac{3-\gamma}{1+\gamma}$-approximate when $\delta=0$, corresponding to consistency guarantee, and gradually increases linearly to $\frac{3+\gamma}{1-\gamma}$-approximate, corresponding to robustness guarantee.

4 Single Obnoxious Facility Location Game

In this section, we study the single obnoxious facility location game with candidate locations on the real line, where each agent wants the facility (e.g., a garbage dump) to be built as far away from her as possible, and the distance from her location to the facility can be interpreted as her utility. We consider two social objectives: maximizing the minimum utility and maximizing the social utility.

4.1 Minimum Utility

For the objective of maximizing the minimum utility, Tang et al. [27] showed that all SP mechanisms without predictions have unbounded approximation ratios, which implies that any SP mechanism with predictions is unbounded robust. For completeness, we also give a specific description.

Theorem 7 *For the single obnoxious facility location game on the real line, any SP mechanism with predictions has an unbounded robustness under the minimum utility objective.*

Proof Assuming that f is an SP mechanism with β-robustness, consider an instance $I(\mathbf{x},M)$ with $\mathbf{x}=(0,2+\epsilon)$ and $M=\{1,2\}$, where $0<\epsilon<1/\beta$. For this instance, $o(I)=1$ and $OPT(I)=1$. If f outputs 2, then the minimum utility of f is ϵ and its approximation ratio is no less than $1/\epsilon>\beta$, which is a contradiction. Thus, for any prediction $\hat{o}\in M, f(\mathbf{x},\hat{o})=1$ and agent 1' utility under f is 1. Consider another instance $I'(\mathbf{x}',M)$ with $\mathbf{x}'=(1-\epsilon^2,2+\epsilon)$. For this new instance, $o(I')=2$ and $OPT(I')=\epsilon$. Note that for I', f must output 2, since if f outputs 1, then the minimum utility of f is ϵ^2 and its approximation ratio is no less than $\epsilon/\epsilon^2=1/\epsilon>\beta$. Above all, agent 1 at location 0 can benefit herself by misreporting $1-\epsilon^2$, which contradicts f's strategy-proofness. □

4.2 Social Utility

We now focus on the objective of maximizing the social utility. For the single obnoxious facility location game with candidate locations on the real line, Gai et al. [14] proposed a deterministic 3-approximate GSP mechanism without predictions and showed that no deterministic GSP mechanism without predictions can achieve an approximation ratio better than 3.

Obviously, the optimal solution to the problem is either the leftmost candidate location a or the rightmost one b. Therefore, the predicted optimal solution $\hat{o}$ is also in $\{a, b\}$. We introduce the reliance parameter $\gamma \in [0, 1)$ just like in Section 3.2 and adopt the voting idea to determine the facility location. Specifically, there are $(1+\gamma)n$ voters voting for two candidate locations $\{a, b\}$, where each agent $i \in N$ supports the candidate location farther from her, and the other γn virtual voters support $\hat{o}$. Then the facility location is determined by the majority rule. Note that we transform the degree of reliance on the prediction $\hat{o}$ into the number of virtual voters. Clearly, a larger value of γn implies a greater probability that the mechanism outputs the prediction $\hat{o}$. For this mechanism (denoted as Mechanism 3), we have the following conclusion.

Mechanism 3

Input: A reported location profile $\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{R}^n$; Candidate locations $M \subseteq \mathbb{R}$; A prediction $\hat{o} \in \{a, b\}$; $\gamma \in [0, 1)$.
Output: A facility location $y \in \{a, b\}$.
Let n_1 be the number of agents with $x_i \leq \frac{a+b}{2}$, and n_2 be the number of agents with $x_i > \frac{a+b}{2}$.
if $(\hat{o} = b \;\&\&\; n_1 + \gamma n > n_2) \| (\hat{o} = a \;\&\&\; n_1 > n_2 + \gamma n)$ **then**
return $y = b$.
else
return $y = a$.
end if

Theorem 8 *For the single obnoxious facility location game on the real line, Mechanism 3 is deterministic, anonymous, GSP, and it is $\frac{3-\gamma}{1+\gamma}$-consistent and $\frac{3+\gamma}{1-\gamma}$-robust under the social utility objective.*

Proof Mechanism 3 is clearly anonymous, because its outcome only depends on the set of reported locations and not on the identities of the agents.

GSP: We now prove that for any group of agents G, there exists at least one member who cannot benefit from G's misreporting. Now consider the case where Mechanism 3 outputs $y = b$ when all the agents in G truthfully report their locations. If G does not contain any agent with location less than or equal to $\frac{a+b}{2}$, then all agents in G cannot benefit from misreporting, since any G's misreporting cannot change the outcome of Mechanism 3. Othewise, if G's misreporting can change the outcome of Mechanism 3 to y', G must contain some agent i with location $x_i \leq \frac{a+b}{2}$ but misreporting $x'_i > \frac{a+b}{2}$, which results in $y' = a$ and she cannot benefit. The proof for the case of $y = a$ is similar and omitted.

We use f to denote Mechanism 3. For any instance $I(\mathbf{x}, M)$, let $y = f(\mathbf{x}, \hat{o})$ be the outcome of Mechanism 3, o be the optimal facility location, and OPT be the optimal social utility. We denote by S_1 and S_2 the sets of agents whose locations are to the left and right of $\frac{a+b}{2}$, respectively, that is $S_1 = \{i : x_i \leq \frac{a+b}{2}\}$ and $S_2 = \{i : x_i > \frac{a+b}{2}\}$.

Consistency: Assume w.l.o.g that $\hat{o} = o = a$. If $y = a$, then the consistency is 1. Consider the case of $y = b$. By Mechanism 3, it holds $n_1 > n_2 + \gamma n$. Then, we have

$$n_1 > \frac{(1+\gamma)n}{2} \text{ and } n_2 < \frac{(1-\gamma)n}{2},$$

that is, $n_2 < \frac{1-\gamma}{1+\gamma} \cdot n_1$. In this case, the social utility of Mechanism 3 is

$$SU(\mathbf{x}, y) = \sum_{i \in N} d(x_i, y) = \sum_{i \in S_1} d(x_i, y) + \sum_{i \in S_2} d(x_i, y) \geq n_1 \cdot d(\frac{a+b}{2}, y).$$

Define $D_1 = SU(\mathbf{x}, y) - n_1 \cdot d(\frac{a+b}{2}, y)$, then

$$\begin{aligned} D_1 &= \sum_{i \in S_1} (d(x_i, y) - d(\frac{a+b}{2}, y)) + \sum_{i \in S_2} d(x_i, y) \\ &= \sum_{i \in S_1} d(x_i, \frac{a+b}{2}) + \sum_{i \in S_2} d(x_i, y). \end{aligned}$$

Consider a new instance $I'(\mathbf{x}', M)$ where there are n_1 agents located at $\frac{a+b}{2}$ and n_2 agents located at y. For this instance, the optimal facility location is o, and the optimal social utility is

$$\begin{aligned} OPT(I') &= \sum_{i \in S_1} d(x_i, o) + \sum_{i \in S_2} d(x_i, o) \\ &= n_1 \cdot d(\frac{a+b}{2}, o) + n_2 \cdot d(y, o) \\ &= n_1 \cdot d(\frac{a+b}{2}, y) + 2n_2 \cdot d(\frac{a+b}{2}, y) \\ &< (2 \cdot \frac{1-\gamma}{1+\gamma} + 1) \cdot n_1 \cdot d(\frac{a+b}{2}, y) \\ &= \frac{3-\gamma}{1+\gamma} \cdot n_1 \cdot d(\frac{a+b}{2}, y), \end{aligned}$$

where the inequality holds due to $n_2 < \frac{1-\gamma}{1+\gamma} \cdot n_1$.

Define D_2 as the difference between OPT and $OPT(I')$, then

$$\begin{aligned} D_2 &= OPT - OPT(I') \\ &= \sum_{i \in S_1} (d(x_i, o) - d(\frac{a+b}{2}, o)) + \sum_{i \in S_2} (d(x_i, o) - d(y, o)). \end{aligned}$$

Therefore,

$$\begin{aligned} D_1 - D_2 = &\sum_{i \in S_1} (d(x_i, \frac{a+b}{2}) - d(x_i, o) + d(\frac{a+b}{2}, o)) \\ &+ \sum_{i \in S_2} (d(x_i, y) - d(x_i, o) + d(y, o)) \geq 0, \end{aligned}$$

where $D_1 - D_2 \geq 0$ holds due to the triangle inequality. Hence,

$$\begin{aligned} \frac{OPT}{SU(\mathbf{x}, y)} &\leq \frac{OPT - D_1}{SU(\mathbf{x}, y) - D_1} \\ &\leq \frac{OPT - D_2}{SU(\mathbf{x}, y) - D_1} \\ &\leq \frac{OPT(I')}{n_1 \cdot d(\frac{a+b}{2}, y)} \\ &< \frac{\frac{3-\gamma}{1+\gamma} \cdot n_1 \cdot d(\frac{a+b}{2}, y)}{n_1 \cdot d(\frac{a+b}{2}, y)} \\ &= \frac{3-\gamma}{1+\gamma}. \end{aligned}$$

Robustness: Assume w.l.o.g that $o = a$ and $y = b$. Then, the social utility of Mechanism 3 is

$$SU(\mathbf{x}, y) = \sum_{i \in N} d(x_i, y) = \sum_{i \in S_1} d(x_i, y) + \sum_{i \in S_2} d(x_i, y) \geq n_1 \cdot d(\frac{a+b}{2}, y).$$

Define $D_1 = SU(\mathbf{x}, y) - n_1 \cdot d(\frac{a+b}{2}, y)$, then

$$\begin{aligned} D_1 &= \sum_{i \in S_1} (d(x_i, y) - d(\frac{a+b}{2}, y)) + \sum_{i \in S_2} d(x_i, y) \\ &= \sum_{i \in S_1} d(x_i, \frac{a+b}{2}) + \sum_{i \in S_2} d(x_i, y). \end{aligned}$$

Consider a new instance $I'(\mathbf{x}', M)$ where there are n_1 agents located at $\frac{a+b}{2}$ and n_2 agents located at y. For this instance, the optimal facility location is o, and the optimal social utility is

$$\begin{aligned} OPT(I') &= \sum_{i \in S_1} d(x_i, o) + \sum_{i \in S_2} d(x_i, o) \\ &= n_1 \cdot d(\frac{a+b}{2}, o) + n_2 \cdot d(y, o) \\ &= n_1 \cdot d(\frac{a+b}{2}, y) + 2n_2 \cdot d(\frac{a+b}{2}, y). \end{aligned}$$

When the prediction is arbitrarily inaccurate, by Mechanism 3, it holds $n_1 > \frac{(1-\gamma)n}{2}$ and $n_2 < \frac{(1+\gamma)n}{2}$, that is, $n_2 < \frac{1+\gamma}{1-\gamma} \cdot n_1$. Thus, we have

$$OPT(I') < (2 \cdot \frac{1+\gamma}{1-\gamma} + 1) \cdot n_1 \cdot d(\frac{a+b}{2}, y) = \frac{3+\gamma}{1-\gamma} \cdot n_1 \cdot d(\frac{a+b}{2}, y).$$

Define D_2 as the difference between OPT and $OPT(I')$, then

$$\begin{aligned} D_2 &= OPT - OPT(I') \\ &= \sum_{i \in S_1} (d(x_i, o) - d(\frac{a+b}{2}, o)) + \sum_{i \in S_2} (d(x_i, o) - d(y, o)). \end{aligned}$$

Therefore,

$$\begin{aligned} D_1 - D_2 = &\sum_{i \in S_1} (d(x_i, \frac{a+b}{2}) - d(x_i, o) + d(\frac{a+b}{2}, o)) \\ &+ \sum_{i \in S_2} (d(x_i, y) - d(x_i, o) + d(y, o)) \geq 0, \end{aligned}$$

where $D_1 - D_2 \geq 0$ holds because of the triangle inequality. Hence,

$$\begin{aligned} \frac{OPT}{SU(\mathbf{x}, y)} &\leq \frac{OPT - D_1}{SU(\mathbf{x}, y) - D_1} \\ &\leq \frac{OPT - D_2}{SU(\mathbf{x}, y) - D_1} \\ &\leq \frac{OPT(I')}{n_1 \cdot d(\frac{a+b}{2}, y)} \\ &< \frac{\frac{3+\gamma}{1-\gamma} \cdot n_1 \cdot d(\frac{a+b}{2}, y)}{n_1 \cdot d(\frac{a+b}{2}, y)} \\ &= \frac{3+\gamma}{1-\gamma}. \end{aligned}$$

□

Below, we prove that the trade-off between consistency and robustness achieved by Mechanism 3 is optimal.

Theorem 9 *For the single obnoxious facility location game on the real line, the robustness of any deterministic SP mechanism with consistency strictly smaller than $\frac{3-\gamma}{1+\gamma}$ is at least $\frac{3+\gamma}{1-\gamma} - \delta$ for any $\gamma \in [0, 1)$ and $\delta > 0$, under the social utility objective.*

Proof Consider an instance I with two candidate locations at ε and 2, and n agents such that $\frac{(1+\gamma)}{2}n-1$ of them are positioned at 1 while the remaining $\frac{(1-\gamma)}{2}n+1$ are positioned at 2, where $\varepsilon \in (0,1)$. The social utilities of the two possible locations 0 and 2 in this instance are

$$\begin{aligned} SU(\mathbf{x},\varepsilon) &= \left(\frac{1+\gamma}{2}n-1\right)(1-\varepsilon)+\left(\frac{1-\gamma}{2}n+1\right)(2-\varepsilon) \\ &= \frac{3-\gamma}{2}n+1-\varepsilon \end{aligned}$$

and

$$SU(\mathbf{x},2) = \frac{1+\gamma}{2}n-1.$$

Observe that the optimal solution is to place the facility at ε. So, when given a correct prediction about the optimal solution, if the mechanism chooses 2, it would achieve consistency $\frac{\frac{3-\gamma}{2}n+1-\varepsilon}{\frac{1+\gamma}{2}n-1} > \frac{3-\gamma}{1+\gamma}$, for any $\varepsilon < 1$. Since the mechanism is assumed to achieve a consistency ratio strictly smaller than that much, it must output the optimal solution when given I as input, thus it places the facility at ε.

Next, consider a sequence of instances in which each of the agents located at 1 are one-by-one moved to ε. Due to strategyproofness, the mechanism must place the facility at ε in each such instance as otherwise, if the solution is 2, the agent that moves from 1 to ε would increase their utility from $1-\varepsilon$ to 1. So, in the last instance of this sequence, we have $\frac{(1+\gamma)}{2}n-1$ agents at ε and $\frac{(1-\gamma)n}{2}+1$ agents at 2. Similarly, we now move one-by-one the agents at 2 to $1+\varepsilon$. The mechanism must still place the facility at ε in all these instance as otherwise an agent would prefer to move back from $1+\varepsilon$ to 2. So, we finally have an instance I' in which there are $\frac{(1+\gamma)n}{2}-1$ agents at ε and $\frac{(1-\gamma)n}{2}+1$ agents at $1+\varepsilon$, and the mechanism places the facility at ε (see Fig.2). We now have

$$SU(\mathbf{x},\varepsilon) = \frac{1-\gamma}{2}n+1$$

and

$$\begin{aligned} SU(\mathbf{x},2) &= \left(\frac{1+\gamma}{2}n-1\right)(2-\varepsilon)+\left(\frac{1-\gamma}{2}n+1\right)(1-\varepsilon) \\ &= \frac{3+\gamma}{2}n-1-\varepsilon. \end{aligned}$$

Hence, the robustness of the mechanism is

$$\frac{\frac{3+\gamma}{2}n-1-\varepsilon}{\frac{1-\gamma}{2}n+1} = \frac{3+\gamma}{1-\gamma}-\delta,$$

where δ approaches 0 as n approaches infinity and ε approaches 0. □

While implementing Mechanism 3, the mechanism designer can choose the value of γ based on her confidence in the prediction $\hat{o}$. If the designer is not confident in the prediction $\hat{o}$, setting $\gamma = 0$ will yield a 3-consistent and 3-robust mechanism that matches the best performance guarantee without predictions [14]. As γ gradually increases from 0 to 1, the consistency decreases from 3 to 1, achieving the optimal performance, but at the cost of increasing the robustness (see Fig.1).

Here, we need not to analyze the performance of Mechanism 3 when predictions are not fully accurate via the prediction error, since the prediction $\hat{o}$ is either o or $\{a,b\}\backslash o$.

5 Conclusions and Open Problems

In this paper, we studied the single (obnoxious) facility location game with candidate locations on the real line under the framework of mechanism design with predictions. We first considered the single facility location game. Under the maximum cost objective, we proposed a deterministic, anonymous, and GSP mechanism, and proved that it is 1-consistent and 3-robust, achieving the best of both worlds. We also analyzed its performance under the social cost objective, but it provides poor robustness guarantee. By reducing its reliance on the prediction, we proposed another deterministic, anonymous, and GSP mechanism, and proved that it is $\frac{3-\gamma}{1+\gamma}$-consistent and $\frac{3+\gamma}{1-\gamma}$-robust, achieving the best possible trade-off between consistency and robustness. We respectively represented the approximation ratios of both mechanisms as functions of the prediction error, demonstrating that the mechanisms can transition from consistency guarantees to robustness guarantees in a linearly gradual manner. We also considered the single obnoxious facility location game. For the minimum utility objective, we found that any SP mechanism with predictions has an unbounded robustness. Therefore, we focused on the social utility objective and provided a deterministic, anonymous, and GSP mechanism, which is $\frac{3-\gamma}{1+\gamma}$-consistent and $\frac{3+\gamma}{1-\gamma}$-robust, achieving the best possible trade-off between consistency and robustness. However, many open questions remain in this direction:

1. For the case of locating two facilities, can we provide a GSP mechanism with predictions that achieves a good trade-off between consistency and robustness? We have attempted to do so, but the results are unsatisfactory.
2. Can we extend the line to other metric spaces, such as trees or cycles, and use predictions to achieve good results? We might also consider facility location games in other settings, such as under different preference models (e.g., double-peaked preferences [13]) or other social objectives (e.g., minimax envy [6]), but this is clearly not an easy task.
3. Randomized mechanisms often achieve better approximation ratios than the deterministic. Can we design randomized mechanisms with predictions to achieve better performance?

Acknowledgements This research was supported in part by the National Natural Science Foundation of China (12201590, 12171444, 11971447).

Conflict of interest

The authors declare that they have no conflict of interest.

References

1. Agrawal, P., Balkanski, E., Gkatzelis, V., Ou, T., Tan, X.: Learning-augmented mechanism design: Leveraging predictions for facility location. In: Proceedings of the 23rd ACM Conference on Economics and Computation, pp. 497–528 (2022)
2. Azar, Y., Panigrahi, D., Touitou, N.: Online graph algorithms with predictions. In: Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 35–66 (2022)
3. Balkanski, E., Gkatzelis, V., Tan, X.: Strategyproof scheduling with predictions. arXiv preprint arXiv:2209.04058 (2022)
4. Banerjee, S., Gkatzelis, V., Gorokh, A., Jin, B.: Online nash social welfare maximization with predictions. In: Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1–19 (2022)
5. Borodin, A., El-Yaniv, R.: Online computation and competitive analysis. Cambridge University Press, Cambridge (2005)
6. Cai, Q., Filos-Ratsikas, A., Tang, P.: Facility location with minimax envy. In: Proceedings of the 25th International Joint Conference on Artificial Intelligence, pp. 137—143 (2016)
7. Chan, H., Filos-Ratsikas, A., Li, B., Li, M., Wang, C.: Mechanism design for facility location problems: A survey. In: Proceedings of the 30th International Joint Conference on Artificial Intelligence, pp. 4356–4365 (2021)

8. Dütting, P., Lattanzi, S., Paes Leme, R., Vassilvitskii, S.: Secretaries with advice. In: Proceedings of the 22nd ACM Conference on Economics and Computation, pp. 409–429 (2021)
9. Ergun, J.C., Feng, Z., Silwal, S., Woodruff, D.P., Zhou, S.: Learning-augmented k-means clustering. arXiv preprint arXiv:2110.14094 (2021)
10. Fang, J., Fang, Q., Liu, W., Nong, Q.: Mechanism design with predictions for facility location games with candidate locations. In: X. Chen, B. Li (eds.) Theory and Applications of Models of Computation, pp. 38–49. Springer Nature Singapore, Singapore (2024)
11. Feijen, W., Schäfer, G.: Dijkstras algorithm with predictions to solve the single-source many-targets shortest-path problem. arXiv preprint arXiv:2112.11927 (2021)
12. Feldman, M., Fiat, A., Golomb, I.: On voting and facility location. In: Proceedings of the 2016 ACM Conference on Economics and Computation, pp. 269–286 (2016)
13. Filos-Ratsikas, A., Li, M., Zhang, J., Zhang, Q.: Facility location with double-peaked preferences. Auton. Agent. Mulit-Ag **31**, 1209–1235 (2017)
14. Gai, L., Liang, M., Wang, C.: Obnoxious facility location games with candidate locations. In: Q. Ni, W. Wu (eds.) AAIM 2022, LNCS, pp. 96–105. Springer, Cham (2022)
15. Gkatzelis, V., Kollias, K., Sgouritsa, A., Tan, X.: Improved price of anarchy via predictions. In: Proceedings of the 23rd ACM Conference on Economics and Computation, pp. 529–557 (2022)
16. Gollapudi, S., Panigrahi, D.: Online algorithms for rent-or-buy with expert advice. In: K. Chaudhuri, R. Salakhutdinov (eds.) ICML 2019, vol. 97, pp. 2319–2327. Long Beach, California, USA (2019)
17. Im, S., Kumar, R., Montazer Qaem, M., Purohit, M.: Non-clairvoyant scheduling with predictions. In: Proceedings of the 33rd ACM Symposium on Parallelism in Algorithms and Architectures, pp. 285–294 (2021)
18. Istrate, G., Bonchis, C.: Mechanism design with predictions for obnoxious facility location. arXiv preprint arXiv:2212.09521 (2022)
19. Lattanzi, S., Lavastida, T., Moseley, B., Vassilvitskii, S.: Online scheduling via learned weights. In: Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, pp. 1859–1877 (2020)
20. Lindermayr, A., Megow, N.: Algorithms with predictions. https://algorithms-with-predictions.github.io. 31 May 2024
21. Lykouris, T., Vassilvitskii, S.: Competitive caching with machine learned advice. J. ACM **68**(4), 1–25 (2021)
22. Mitzenmacher, M., Vassilvitskii, S.: Algorithms with predictions. Commum. ACM **65**(7), 33–35 (2022)
23. Polak, A., Zub, M.: Learning-augmented maximum flow. arXiv preprint arXiv:2207.12911 (2022)
24. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. ACM Trans. Econ. Comput. **1**(4), 1–26 (2013)
25. Purohit, M., Svitkina, Z., Kumar, R.: Improving online algorithms via ml predictions. In: S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (eds.) NeurIPS 2018, vol. 31, pp. 9661–9670. Montréal, Canada (2018)
26. Rohatgi, D.: Near-optimal bounds for online caching with machine learned advice. In: Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1834–1845 (2020)
27. Tang, Z., Wang, C., Zhang, M., Zhao, Y.: Strategyproof facility location with limited locations. J. Oper. Res. Soc. China pp. 1–15 (2022)
28. Xu, C., Lu, P.: Mechanism design with predictions. In: Proceedings of the 31st International Joint Conference on Artificial Intelligence, pp. 571–577 (2022)