From Recommendation Systems to Facility Location Games

Omer Ben-Porat, Gregory Goren, Itay Rosenberg, Moshe Tennenholtz

Technion - Israel Institute of Technology

Haifa 32000 Israel

{omerbp@campus, gregory.goren@campus, itayrose@campus, moshet@ie}.technion.ac.il

Abstract

Recommendation systems are extremely popular tools for matching users and contents. However, when content providers are strategic, the basic principle of matching users to the closest content, where both users and contents are modeled as points in some semantic space, may yield low social welfare. This is due to the fact that content providers are strategic and optimize their offered content to be recommended to as many users as possible. Motivated by modern applications, we propose the widely studied framework of facility location games to study recommendation systems with strategic content providers. Our conceptual contribution is the introduction of a *mediator* to facility location models, in the pursuit of better social welfare. We aim at designing mediators that a) induce a game with high social welfare in equilibrium, and b) intervene as little as possible. In service of the latter, we introduce the notion of *intervention cost*, which quantifies how much damage a mediator may cause to the social welfare when an off-equilibrium profile is adopted. As a case study in high-welfare low-intervention mediator design, we consider the one-dimensional segment as the user domain. We propose a mediator that implements the socially optimal strategy profile as the unique equilibrium profile, and show a tight bound on its intervention cost. Ultimately, we consider some extensions, and highlight open questions for the general agenda.

1 Introduction

Publishing, blogging, and content creation are fundamental to data science. Indeed, many of the AI-based recommendation systems aim at matching users with data created by content providers. While at first glance this task might be viewed as purely computational, a major topic that should be tackled is the participants' *incentives*.

To illustrate such incentives, consider the following example inspired by Hotelling's seminal work (Hotelling 1929). A blogging recommendation system recommends users with relevant blogs. Two players (i.e., publishers/blog owners) write one blog each. For simplicity, assume every blog is represented as a point along the $[0, 1]$ segment, e.g., the mix between news and articles in the blog. Thus, each player selects a point in that segment. A set of users, where user preferences are uniformly distributed along the same segment, approach the recommendation system to have a recommendation for one blog. The social cost of the users is defined as the sum of absolute distances between preferences and recommended contents. As Hotelling demonstrates, the only pure Nash equilibrium (PNE hereinafter) is obtained when both players select the center of the segment, producing identical content. The social cost of the PNE profile is far from optimum, and users effectively suffer from strategic behavior of the players.

The difference between this example and the one considered by Hotelling in his seminal work is that in modern applications, users are not explicitly exposed to players' content, but rather the recommendation system serves as a *mediator* for matching users with content. However, despite their commercial success, current recommendation systems basically (very efficiently) implement the above principle: users are matched with the "closest" content. A prominent example is the vector space model for document retrieval in response to a user query (Salton, Wong, and Yang 1975). Indeed, the recommendation system is just a tool for having the users "choose" the offered content most similar to their taste.

1.1 The Connection to Facility Location Games

Facility location games (Anderson, De Palma, and Thisse 1992; Brenner 2010; Fournier and Scarsini 2014), which are extensively studied in economics, operations research and computer science, portray recommendation systems with strategic content providers incredibly well, due to the nature of recommendation systems described above. This is true since in many (and perhaps even most) recommendation systems, the computational task of matching users with content is carried out by modeling both into a joint metric space. In the vast majority of facility location games literature, it is assumed that users are attracted to their nearest facility[1]. In the example above, this amounts to each user selecting the closest blog to her preference.

Given a strategy profile, i.e., the selection of each player for a particular mix of news and articles in her blog, the best (in terms of social cost) recommendation system matches each user with the nearest blog to her preference. However,

---

[1]Some exceptions appear in Subsection 1.4.

the way the recommendation system operates affects the strategic behavior of players, and more precisely redefines player payoffs. As a result, it may induce a game with a different set of PNE, potentially with improved social welfare (i.e., the social cost with negative sign). This calls for the design of a mediator that takes into account strategic behavior of players, aiming at achieving low social cost in equilibrium.

1.2 Our Agenda: High Welfare Low Intervention Mediator Design

Our conceptual contribution is the introduction of a mediator to facility location models, in the pursuit of better outcomes for recommendation systems with strategic content providers. We aim at designing mediators that a) induce a game with high social welfare in equilibrium, and b) intervene as little as possible.

To better describe the latter, let NIME be the No Intervention Mediator, namely the default mediator that displays each user with its nearest content. A question in this regard is whether a mediator that intervenes with the process can do better than NIME, in terms of social welfare. Clearly, one can design a mediator that *dictates* each player which content to select, thereby forcing the players to any desired strategy profile. Consider the DICT mediator that operates as follows: it commands each player to play a particular strategy, and if she disobeys DICT directs no user to her content. Indeed, by adopting DICT the mediator designer is guaranteed to have any specific strategy profile as the unique PNE of the induced recommendation game. However, this is a huge and potential harmful intervention, and doing so may lead to poor performance in case that specific profile is not materialized. Crucially, this can happen even when the players are rational, but some of the assumptions the mediator holds regarding the players are violated, e.g., players have constraints on their strategy space, slightly different payoff functions, they are able to form coalitions and so on.

A major aspect of our approach is the quantification of *intervention cost*. Given the above intuition, we define the intervention cost as the maximal increase in user social cost incurred when the players use arbitrary strategies (i.e., select arbitrary contents). Formally, let $\mathcal{M}$ denote a mediator, S be the set of all strategy profiles, $s \in S$ be a strategy profile, and let $\mathrm{SC}(\mathcal{M}, s)$ denote the social cost under $\mathcal{M}$ and s. The intervention cost of a mediator $\mathcal{M}$ is defined as $\sup_{s \in S} \{\mathrm{SC}(\mathcal{M}, s) - \mathrm{SC}(\mathrm{NIME}, s)\}$. The term inside the supremum is the increase in social cost compelled by $\mathcal{M}$ under a profile s, which is non-negative (due to the definition of NIME). The intervention cost is therefore the potential increase in social cost, which might be caused under all possible (even off-equilibrium) profiles.

Evidently, NIME has a zero intervention cost and high social cost, and serves as a benchmark for low intervention cost. On the other hand, DICT has high intervention cost but low social cost (when dictating the socially optimal strategy profile), and it serves as a benchmark for low social cost. The main challenge in our agenda is the design of a mediator with low intervention cost *and* low social cost in equilibrium.

1.3 Our Results

This paper presents a case study in high welfare low intervention mediator design. The mathematical model we adopt in service of the above is based on pure location Hotelling games (Hotelling 1929), with the $[0, 1]$ segment. Section 2 presents a formal mathematical model for the setting, stating former widely-known results for pure Hotelling games (or equivalently, recommendation games with NIME as a mediator) on the segment with uniform user distribution. In addition, in Subsection 2.3 we define the intervention cost, and bound the intervention cost of DICT.

The main technical results of the paper appear in Section 3. We introduce the Limited Intervention Mediator, LIME. We provide the intuition behind it, as well as a thorough example to illustrate how it operates. We then prove that the game induced under LIME possesses a unique PNE, which attains the minimal social cost. Then, we establish upper and lower bounds on its intervention cost, which are almost tight. We show that not only is its intervention cost lower than that of DICT for every n, it is also $O(\frac{1}{n})$. Since DICT and LIME have the same social cost under (the unique) PNE profile, the results given in this section provide a highly encouraging answer to the challenge given above.

We subsequently study three natural extensions of the basic setting. First, we discuss neutral mediators. Informally, a mediator is neutral if when two players swap their strategies, the mediator's direction also swaps. We show an impossibility result for neutral mediators that aim at optimizing the social cost in the two-player case. Second, we design a mediator with a configurable intervention cost. This mediator is important where e.g. one seeks to design a mediator that minimizes social cost but is penalized for intervention (similarly to regularization in machine learning models). We show that for some cases of n, it induces Pareto optimal solutions for the setting. Third, we consider non-uniform user distributions. We propose the General Limited Intervention Mediator, which depends on the distribution quantiles solely. We then prove that it always induces a game with a unique PNE. This becomes even more striking when one recalls that PNE may not exist in facility location games (with no mediator) (Osborne 1993), and are generally hard to characterize (see, e.g., (Shilony 1981; Ewerhart 2015)). We bound its intervention cost (for the uniform distribution), and show that its intervention cost is lower than that of DICT.

From an algorithmic perspective, we deal with one-dimensional problems. This may sound disappointing, but recall that we intentionally focus on relatively simple, structured problems, and that this domain is extremely well-studied in the facility location literature. Due to space constraints, the proofs are deferred to the appendix.

1.4 Related Work

The notion of a mediator in a game-theoretic setting was first proposed by Aumann in his seminal work on Correlated Equilibrium (Aumann 1974). In the setting Aumann considers, a mediator may send signals to the players, where the signal is designed to drive the players to achieve better payoffs. The work of Shoham and Tennenholtz (1995) grants

stronger capabilities to the mediator, by setting constraints on participants' behavior. A different type of mediator, considered by Monderer and Tennenholtz (2004), is allowed to change player payoffs. As in our work, the mediator is designed to "force" players into playing some desired subset of their strategy sets with minimal interference with the payoff functions. Unlike the above work, where the mediator intervenes with the outcome of the game to improve the social welfare of the (strategic) players, in the model we study the mediator intervenes in order to improve the social welfare of the (non-strategic) users, as we describe next.

The work of Ben-Basat, Tennenholtz, and Kurland (2017) introduces a formal model of adversarial information retrieval. Each author (a player) has several strategies, where each strategy corresponds to a document she can publish. Each pair of document-query (where a document is a player's strategy and a query represents a user) has a score, termed "quality", which measures the extent to which the document is relevant to the query. The mediator (i.e., search engine) provides each user the document with the highest quality relative to her query, among those selected under the corresponding strategy profile. In their model, every author seeks to maximize a function that takes into account the number of queries for which her document is the most relevant (representing users directed to that document by the mediator) along with the document-query quality.

As Ben-Basat, Tennenholtz, and Kurland show, displaying each user with the document with the highest quality (i.e., no intervention) may lead to deteriorated content and low user utility. They argue against no intervention, and claim that introducing randomization into the mediator leads to an overall user utility that transcends that of no intervention at all. Nevertheless, they neither a) provide a systematic approach for doing so under PNE profiles; nor b) show that a PNE always exists. Recently, Ben-Porat, Rosenberg, and Tennenholtz (2019) claim in favor of no intervention in a similar model, and show that no intervention policy leads to convergence of any better-response dynamics; thus, a PNE is guaranteed to exist. Importantly, Ben-Basat, Tennenholtz, and Kurland (2017) as well as Ben-Porat, Rosenberg, and Tennenholtz (2019) do not consider the task of mediator design for better outcomes.

The work of Ben-Porat and Tennenholtz (2018a) is the first to consider mediator design in recommendation systems with strategic content providers, in a mathematical model that extends that of Ben-Basat, Tennenholtz, and Kurland (2017) and Ben-Porat, Rosenberg, and Tennenholtz (2019). They highlight several fairness-related properties that a mediator should arguably satisfy, along with the requirement of PNE existence. They show that the no-intervention mediator satisfies the fairness-related properties, but may lead to a game without PNEs. Then, they design a mediator that is based on the Shapley value (Shapley 1952), prove it satisfies the fairness properties and the game it induces always possesses a PNE. On the other hand, Ben-Porat and Tennenholtz (2018a) take an axiomatic approach, and do not address user utility as one of the axioms. Moreover, they show that the user utility under their proposed mediator can be arbitrarily low. In contrast, this paper stems from a user utility opti-

mization point of view, and so are the solution concepts it proposes.

Several extensions of pure location Hotelling games have been suggested recently, assuming non-deterministic user behavior (Feldman, Fiat, and Obraztsova 2016; Shen and Wang 2017; Ben-Porat and Tennenholtz 2017). More precisely, users do not select their nearest facility, but rather have a *reaction function*, mapping every strategy profile to a distribution over player indices, possibly skipping all options. The work of Ben-Porat and Tennenholtz (Ben-Porat and Tennenholtz 2017) associates users with a reaction function motivated by decision theory literature, and shows that the induced facility location game always possesses a PNE, regardless of the metric space in which users are embedded. Shen and Wang (2017) show a result of similar flavor for a different reaction function. Interestingly, every reaction function can be implemented as a mediator. Unlike this line of research, in this paper (as in the information retrieval setting) the mediator is required to direct users to facilities w.p. 1, i.e., skipping all facilities is not an option. We elaborate on the latter in Section 5.

Finally, a different line of research in the algorithmic game theory literature (Nisan et al. 2007) is the study of facility location, in the context of approximate mechanism design without money (Procaccia and Tennenholtz 2009). That literature deals with the case where only one entity dictates the place of a facility (or several facilities), while user preferences are their private information and are strategically reported (see, e.g., (Schummer and Vohra 2007)). In contrast, the setting we consider is a complete information setting.

2 Mathematical Formulation

In this section we introduce our formal mathematical model. The non-cooperative game we consider is formally defined as follows:

1. A continuous density function g over the unit interval $[0, 1]$, representing user distribution.

2. A set of players $[n] = \{1, 2, \ldots, n\}$, where the pure[2] strategy set of player $i \in [n]$ is $S_i = [0, 1]$. It will sometimes be convenient to say that player i has/owns a *facility* in s_i if the strategy she adopts is s_i. The set of all strategy profiles is denoted by $S \stackrel{\text{def}}{=} \prod_{i=1}^{n} S_i$.

3. A mediator $\mathcal{M}$ is a mapping from the set of strategy profiles and users to the set of distributions over player indices, i.e., $\mathcal{M} : S \times [0, 1] \to \Delta([n])$. Given a pure strategy profile $\boldsymbol{s} = (s_1, \ldots, s_n) \in S$, a user $t \in [0, 1]$ and a player $i \in [n]$, we denote by $\mathcal{M}(\boldsymbol{s}, t)_i$ the probability that $\mathcal{M}$ will send user t to player i under the strategy profile $\boldsymbol{s}$.

4. Given a pure strategy profile $\boldsymbol{s} \in S$, the payoff of each player i is the proportion of users $\mathcal{M}$ directs to her facility. Namely, $\pi_i(\boldsymbol{s}) = \int_0^1 \mathcal{M}(\boldsymbol{s}, t)_i \cdot g(t) dt$. Each player aims at maximizing her payoff.

According to these assumptions, every game is fully described by the number of players, the mediator, and the user

[2]In this paper we discuss pure strategy profiles only.

Number of players	Social Cost (former results)			Intervention Cost (Sections 2 and 3)	
	Optimal	NIME, best PNE	NIME, worst PNE	$\mathcal{I}_n(\mathsf{DICT})$	$\mathcal{I}_n(\mathsf{LIME})$
$n = 2$	$1/8$	$1/4$	$1/4$	$1/4$	$3/16$
$n = 3$	$1/12$	no PNE exists		≥ 0.278	$\in (0.236, 0.278)$
$n \geq 4$	$1/4n$	$1/4(n-2)$	$1/4\lceil n/2 \rceil$	$\geq 1/2 - 3/4n + 1/4n^2$	$\in \left(\frac{2n-4}{n^2}, \frac{2n-3.5}{n^2} \right)$

Table 1: A summary of former results known for pure location Hotelling games, and the results of this paper. The social cost is given for three scenarios: the optimal social cost, for the profile $\boldsymbol{o}^n$ (see Subsection 2.2); the best PNE in terms of social cost under NIME; and the worst PNE in terms of social cost under NIME. Under both DICT and LIME, the induced game possesses a unique PNE, with social cost of $1/4n$, namely the optimal social cost. For these mediators, the table reports the bounds on the intervention cost (see Subsection 2.3) as obtained in Sections 2 and 3.

distribution function, i.e., $G(n, \mathcal{M}, g)$. In addition, unless stated otherwise, g is the uniform distribution; hence, for convenience, we shall use the notation $G(n, \mathcal{M})$ to describe the n-player game induced by selecting $\mathcal{M}$ as a mediator.

Equilibrium For a vector $\boldsymbol{v} = (v_1, \ldots v_n)$, we denote by $\boldsymbol{v}_{-i} \overset{\text{def}}{=} (v_1, \ldots v_{i-1}, v_{i+1}, \ldots v_n)$ the vector that does not contain the i-th coordinate of $\boldsymbol{v}$.

A strategy profile $\boldsymbol{s} \in S$ is called a *pure Nash equilibrium* (PNE hereinafter) if for every $i \in [n]$ and every $s_i' \in S_i$ it holds that $\pi_i(s_i', \boldsymbol{s}_{-i}) \leq \pi_i(\boldsymbol{s})$. We say that player i has a *beneficial deviation* under a strategy profile $\boldsymbol{s} = (s_i, \boldsymbol{s}_{-i})$ if there exists $s_i' \in S_i$ such that $\pi_i(s_i', \boldsymbol{s}_{-i}) > \pi_i(\boldsymbol{s})$.

2.1 Social cost

We remark that any mediator always directs a user to one of the facilities selected by the players; thus, the player payoffs sum to one under every strategy profile $\boldsymbol{s}$, i.e., $\sum_{i=1}^n \pi_i(\boldsymbol{s}) = 1$. Consequently, we view the social cost of the users as the public welfare. In the blogging example given above, we consider the distance between a user's preferences and the actual attribute of the blog as the extent to which she is satisfied with that blog. Clearly, this perspective is identical to the one considered by Hotelling (1929), albeit the motivation of Hotelling is the physical world. Motivated by transportation cost, Hotelling defines the cost a customer (i.e., a user) suffers as the distance he has to travel to reach his nearest facility. This notion is naturally extended to the mediated setting.

Definition 1 (Social cost). *Given a strategy profile $\boldsymbol{s}$ and a mediator $\mathcal{M}$, the social cost is the sum of distances users must travel to reach their recommended facility. Formally, $\mathrm{SC} : S \to \mathbb{R}_{\geq 0}$ is defined by*

$$\mathrm{SC}(\mathcal{M}, \boldsymbol{s}) = \int_0^1 \sum_{i=1}^n \mathcal{M}(\boldsymbol{s}, t)_i \cdot |s_i - t| g(t) dt.$$

2.2 No Intervention Mediator

We denote by NIME the No-Intervention Mediator. Namely, NIME sends every user to his nearest facility, breaking ties uniformly. Formally,

$$\mathsf{NIME}(\boldsymbol{s}, t)_i = \frac{\mathbb{1}_{|s_i - t| \leq \min_{i'} |s_{i'} - t|}}{\sum_{j=1}^n \mathbb{1}_{|s_j - t| \leq \min_{i'} |s_{i'} - t|}}.$$

By implementing NIME we recover the non-mediated version of the setting, where under every strategy profile each

user is attracted to his nearest facility, as is in pure location Hotelling games. Given n, we denote by $\boldsymbol{o}^n = (o_1^n, \ldots, o_n^n)$, where $o_i^n \overset{\text{def}}{=} \frac{2i-1}{2n}$, the sequence of n-*socially optimal locations*. These are optimal in the following sense:

Claim 1 (e.g. (Ben-Porat and Tennenholtz 2018b)). *It holds that $\mathrm{SC}(\mathsf{NIME}, \boldsymbol{o}^n) = \inf_{\boldsymbol{s} \in S} \mathrm{SC}(\mathsf{NIME}, \boldsymbol{s}) = \frac{1}{4n}$.*

It turns out that this is the unique socially optimal profile, up to renaming the players. Moreover, for any fixed strategy profile, it is apparent that employing NIME results in the lowest possible social cost w.r.t. that particular profile. Namely,

Claim 2. *For every strategy profile $\boldsymbol{s}$ and every mediator $\mathcal{M}$, it holds that $\mathrm{SC}(\mathsf{NIME}, \boldsymbol{s}) \leq \mathrm{SC}(\mathcal{M}, \boldsymbol{s})$.*

However, as we show later, the set of PNE depends on the mediator; hence, a PNE under NIME may not be in equilibrium under another mediator, and vice versa. Since we care about the social cost in equilibrium, this turns out to be crucial. Pure location Hotelling games (or equivalently $G(n, \mathsf{NIME})$ for any n) have been studied extensively in the past decade, and its equilibrium structure is widely known (see, e.g., (Eaton and Lipsey 1975; Fournier and Scarsini 2014)). For completeness, we state the following well-known results for $G(n, \mathsf{NIME})$:

- The only PNE for $n = 2$ is $(1/2, 1/2)$.

- There is no PNE for $n = 3$, although a mixed NE exists (Shaked 1982).

- For $n \in \{4, 5\}$ a unique PNE exists (up to renaming the players).

- For $n \geq 6$, there are infinite PNEs.

For an elaborated discussion of these results, the reader is referred to (Eaton and Lipsey 1975). Table 1 summarizes the social cost for NIME. Notice the social cost under the worst PNE, which is a factor of two of the optimal feasible social cost.

2.3 Intervention Cost

Indeed, a mediator that implements the n-socially optimal locations as a PNE can greatly decrease the social cost. By intervening with the market and using a basic punishing mechanism, one can drive the players to play any arbitrary profile, by intervention that makes that profile the only PNE.

We denote by DICT the mediator[3] that dictates the strategy profile o^n, the n-socially optimal locations. Namely, DICT operates as follows: given a strategy profile s, let $A = \{i : s_i = o_i^n\}$. If A is not empty, DICT(s) directs every user t to its nearest facility from those of the players in A, i.e., DICT$(s,t)_i = \mathbb{1}_{i \in A} \cdot$ NIME$(s \cap A, t)_i$. Otherwise, if A is empty, it directs every user to a player selected at random. It is apparent that

Claim 3. *Consider* $G(n, \text{DICT})$ *for any* $n \geq 2$. *The unique PNE is* o^n.

However, sometimes the strategy profile being materialized is not the equilibrium profile. As elaborated above, a mediator intervening with the system to decrease social cost may hence have implicit negative effects. It is therefore highly desired to design a mediator that not only drives the players to a "good" equilibrium (in terms of social cost), but also does not intervene "that much". To quantify the extent to which a mediator intervenes with the natural market, we introduce the following measure.

Definition 2 (intervention cost). *The intervention cost of a mediator* $\mathcal{M}$ *is the maximum difference between the social cost of* $\mathcal{M}$ *and that of* NIME *(i.e., the No-Intervention Mediator), when the maximum is taken over all strategy profiles. Formally,*

$$\mathcal{I}_n(\mathcal{M}) \overset{def}{=} \sup_{s=(s_1, \dots s_n) \in S} \{ \text{SC}(\mathcal{M}, s) - \text{SC}(\text{NIME}, s) \},$$

where the subscript n *emphasizes the dependency on the number of players.*

Notice that SC$(\mathcal{M}, s) -$ SC(NIME, s) is the sum of distances $\mathcal{M}$ compels the users to travel beyond the minimal distance they can travel under s, which is non-negative. The intervention cost is the supremum of this term over all, possibly off-equilibrium profiles. Indeed, the intervention cost captures the measure of intervention by the mediator. By definition of NIME, it has an intervention cost of zero, i.e., $\mathcal{I}_n(\text{NIME}) = 0$ for every n. In addition, the intervention cost as defined above also demonstrates the great amount of intervention employed by DICT.

To illustrate it, consider the consequences of a strategy profile other than o^n being materialized under DICT. Take for example the two-player game, $G(2, \text{DICT})$, and the strategy profile $s = (s_1, s_2) = (0, 1)$. This profile represents the case where the offered contents are varied to the extreme. Both players defy their commands; thus, DICT sends every user to a facility selected uniformly at random, and every user travels a distance of $1/2$ in expectation. Summing over all users we get, SC$(\text{DICT}, s) = 1/2$. The social cost under no mediation is SC$(\text{NIME}, s) = 1/4$; hence, the intervention cost of DICT is lower bounded by SC$(\text{DICT}s) -$ SC$(\text{NIME}, s) = 1/4$. In fact, this profile can be extended to any n-player game by splitting the players evenly on the segment's endpoints, showing that $\mathcal{I}_n(\text{DICT}) \geq 1/4$ for any n. However, using a different construction we show a much tighter bound that depends on n. More precisely,

---

[3]One can think of other dictatorship mediators as well, by varying the punishment in case of disobedience.

Algorithm 1: Limited Intervention Mediator

Input: A strategy profile s and a user t
Output: A location in s
1 **if** $\exists i \in [n-1]$ *such that* $t \in \left(\frac{2i-1}{2n}, \frac{2i+1}{2n} \right)$ **then**
2     $l \leftarrow s \cap \left[0, \frac{2i-1}{2n} \right], r \leftarrow s \cap \left[\frac{2i+1}{2n}, 1 \right]$
3   **if** $l \neq \varnothing$ *and* $r \neq \varnothing$ **then**
4   **return** NIME$(l \cup r, t)$
5   **else if** $l \neq \varnothing$ *or* $r \neq \varnothing$ **then**
6   w.p. $1 - \epsilon$ **return** NIME$(l \cup r, t)$, otherwise **return** $random(s, t)$
7   **else**   // all facilities are inside $\left(\frac{2i-1}{2n}, \frac{2i+1}{2n} \right)$
8   **return** NIME(s, t)
9 **else**   // t is outside $\cup_{i \in [n-1]} \left(\frac{2i-1}{2n}, \frac{2i+1}{2n} \right)$
10   **return** NIME(s, t)
11 **Function** $random(s, t)$:
12   **return** an element from s uniformly at random.

Lemma 1. *For any* $n \geq 3$, *it holds that* $\mathcal{I}_n(\text{DICT}) \geq \frac{1}{2} - \frac{3}{4n} + \frac{1}{4n^2}$.

Due to its high intervention cost (among other properties) using DICT as a mediator may not be the best solution. In the next section we devise a mediator that implements o^n as the unique equilibrium, but unlike DICT has a substantially low intervention cost.

3 Limited Intervention Mediator (LIME)

In this section we take advantage of the game structure to devise a mediator that encourages the players to select the n-socially optimal locations o^n, but intervenes substantially less than DICT. This is done by exploiting the equilibria structure under no mediation. According to the characterization of PNE profiles (Eaton and Lipsey 1975), under NIME a profile can be in equilibrium only if its peripheral facilities (i.e., leftmost and rightmost facilities) are paired. In addition, if beneficial deviations do not exist, the proportion of users coming from the left/right of any facility cannot be greater than the total proportion of users served by any other facility.

Consider the Limited Intervention Mediator described in Algorithm 1, and referred to as LIME hereinafter for abbreviation. The intuition behind LIME is the following: between every two locations that belong to o^n, we construct a *potentially intervened interval* (PII). The users in every PII are not directed to a facility located in the same PII, but rather to the closest facility outside of it. In addition, if a PII does not have a facility located from its left or from its right (exclusive or), the users in that PII are sent to a random player w.p. ϵ, where $\epsilon > 0$ is an arbitrarily small constant.

To facilitate understanding of the LIME mediator, we provide the following example, which is further illustrated in Figure 1.a.

Example 1. Consider $G(4, \text{LIME})$, and the strategy profile $s = (1/16, 4/16, 10/16, 12/16)$. User t_1, located in $0.5/16$, is directed by LIME to player 1's facility in $1/16$, since t_1 is outside the PIIs (Line 10). User t_2, who is located in $3/16$, is
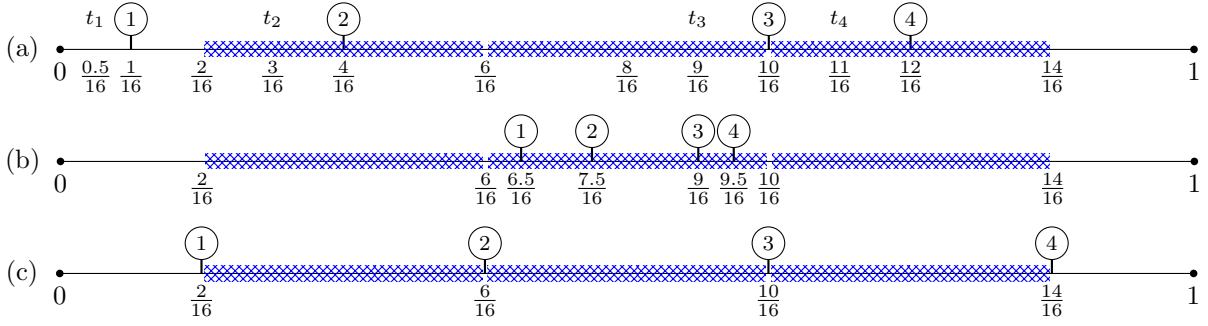
Figure 1: Three strategy profiles for $G(4, \mathsf{LIME})$. The blue (dotted) segments represent potentially intervened intervals (PIIs), and the numbered circles are the locations selected by the corresponding players. Sub-figure (a) visualizes Example 1. Sub-figure (b) exemplifies a case where all players locate their facilities in the PII $(6/16, 10/16)$. In this case, LIME directs every user outside the PIIs to his nearest facility; every user in $(6/16, 10/16)$ is also directed to her nearest facility (Line 8 of Algorithm 1); and every user inside $(2/16, 6/16)$ or $(10/16, 14/16)$ is directed to his nearest facility w.p. $1-\epsilon$, and with the remaining probability to a facility selected uniformly at random (Line 6). The unique PNE of this game (due to Theorem 1) is demonstrated in Sub-figure (c).

inside the PII $(2/16, 6/16)$. Notice that this PII has a facility from its left and a facility from the right. Namely, according to the if condition in Line 3, $\boldsymbol{r} = \boldsymbol{s} \cap [6/16, 1] \neq \varnothing$, and $\boldsymbol{l} = \boldsymbol{s} \cap [0, 2/16] \neq \varnothing$; thus, t_2 is directed to his nearest facility outside $(2/16, 6/16)$, which is $1/16$ (Line 4). User t_3 is inside the PII $(6/16, 10/16)$, and is therefore directed to his nearest facility outside this PII, player 3's facility in $10/16$ (PIIs are open intervals, and this facility lies in the exterior of $(6/16, 10/16)$). User t_4 is inside $(10/16, 14/16)$, and this PII does not have a facility from its right (i.e., $\boldsymbol{r} = \boldsymbol{s} \cap [14/16, 1] = \varnothing$); thus, w.p. $1 - \epsilon$ LIME directs him to the facility in $10/16$, and with the remaining probability he is directed to a facility selected uniformly at random (Line 6).

The only event not covered by Example 1 is the case where all the facilities are located in the same PII, as illustrated in Figure 1.b.

3.1 Pure Nash Equilibrium

We now show that LIME is carefully constructed to mitigate players' incentives. In particular, we show that $\boldsymbol{o}^n$ is the unique PNE of $G(n, \mathsf{LIME})$. First, we show that under any PNE, players are not encouraged to locate their facilities in PIIs.

Lemma 2. *Consider* $G(n, \mathsf{LIME})$ *for any* $n \geq 2$. *If* $\boldsymbol{s}$ *is an equilibrium profile, then* $s_j \notin \left(\frac{2i-1}{2n}, \frac{2i+1}{2n} \right)$ *for every* $j \in [n]$ *and* $i \in [n-1]$.

Next, we leverage Lemma 2 to show that there is a unique equilibrium under LIME, which is composed of the n-socially optimal locations, i.e., $\boldsymbol{o}^n$. The case of $n = 2$ is an interesting exception and is therefore discussed separately in Subsection 4.1.

Theorem 1. *Consider* $G(n, \mathsf{LIME})$ *for any* $n \geq 3$. *The unique PNE (up to renaming the players) is* $\boldsymbol{o}^n$.

See Figure 1.c for illustration. Importantly, Theorem 1 holds for any $\epsilon < 1/2$, and implies that $\boldsymbol{o}^n$ is an *exact* PNE of

$G(n, \mathsf{LIME})$ and not an approximated PNE. In addition, under the profile $\boldsymbol{o}^n$ all the facilities are outside the PIIs; hence, every user is directed to his nearest facility. Consequently,

Corollary 1. *Consider* $G(n, \mathsf{LIME})$ *for any* $n \geq 3$. *The unique PNE,* $\boldsymbol{o}^n$, *attains the optimal social cost,* $\mathrm{SC}(\mathsf{LIME}, \boldsymbol{o}^n) = \frac{1}{4n}$.

Therefore, the social cost of LIME under the PNE profile matches the optimal (see Table 1).

3.2 Intervention Cost

Having demonstrated that LIME obtains the optimal social cost in the (unique) equilibrium, we now claim that its intervention cost is low. By definition of the LIME mediator, it suffices to consider users inside PIIs only, as users outside PIIs are directed to their nearest facility; hence, they do not contribute to the intervention cost. Moreover, for simplicity, we take ϵ to be arbitrarily small (see Line 6), and neglect it in our analysis. By doing so we make the bounds more meaningful, and do not harm the results obtained in the previous sub-section.

For $n = 2$, i.e., the two-player game, the intervention cost of LIME can be determined precisely.

Proposition 1. *It holds that* $\mathcal{I}_2(\mathsf{LIME}) = \frac{3}{16}$.

We now consider games with $n \geq 3$ players, which require a more delicate treatment. First, we lower bound the intervention cost of LIME.

Lemma 3. *For any* $n \geq 3$, *it holds that* $\mathcal{I}_n(\mathsf{LIME}) \geq \frac{2n-4}{n^2}$.

Next, we move on to upper bounding the intervention cost of LIME.

Theorem 2. *For any* $n \geq 3$, *it holds that* $\mathcal{I}_n(\mathsf{LIME}) \leq \frac{2n-3.5}{n^2}$.

Notice that the upper bound almost matches the lower bound. We summarize the intervention cost of LIME in Table 1. Observe that not only does LIME intervene less than DICT, but also its intervention cost diminishes in a $1/n$ scale, similarly to the optimal social cost (which is $1/4n$).

4 Extensions

Beyond the main analysis of the paper given in the previous sections, we find it important to examine the three following extensions.

4.1 Neutral Mediators

Dilemmas other than intervention cost may arise when implementing a mediator. To illustrate such dilemmas and intertwine our recommendation games with computational social choice (Brandt et al. 2016), we analyze the two-player game. As Theorem 1 shows, o^n is the unique PNE for $n \geq 3$. In $G(2, \mathsf{LIME})$, however, LIME does not induce a game with a unique PNE.

Proposition 2. *Consider $G(2, \mathsf{LIME})$. A strategy profile (s_1, s_2) is a PNE if and only if $(s_1, s_2) \in \left\{ \frac{1}{4}, \frac{3}{4} \right\} \times \left\{ \frac{1}{4}, \frac{3}{4} \right\}$.*

Notice that the selection of the socially optimal locations remains a PNE, similarly to greater values of n, but this PNE is not *unique*. To clarify the intuition behind the latter phenomena, consider the profile $(1/4, 1/4)$. According to Line 6 in Algorithm 1, LIME directs all users in the segment $(1/4, 3/4)$ to a random player w.p. ϵ. Under this profile, random direction collides with directing these users to their nearest facility; hence, the randomness LIME employs for breaking symmetry is not enough and other tools should be applied. In fact, this problem can be solved by adding the following (rather complicated) tie-breaker: under $(1/4, 1/4)$ break ties in favor of player 1, and under $(3/4, 3/4)$ break ties in favor of player 2. When augmented to LIME, this tie-breaker induces a game with the socially optimal locations as the unique PNE.

Indeed, our notion of intervention cost is not affected by such tie-breaking intervention. However, the mediator should treat players *fairly*, in some sense of the word. This fact is meaningful with respect to the generality of our agenda, since it introduces other considerations that involve the players rather than the users, borrowing ideas from social choice theory. The most immediate characterization of fairness in our setting is through the notion of *neutrality*. Formally, a mediator $\mathcal{M}$ is neutral if for every $i, j \in [n]$ and $s = (s_{-\{i,j\}}, s_i, s_j)$ it holds that $\pi_i(s) = \pi_j(s_{-\{i,j\}}, s_j, s_i)$. Noticeably, the DICT mediator is not neutral, while the LIME mediator and all other mediators discussed in the paper and the appendix are neutral. The inability of LIME to handle the two-player game then becomes apparent: it turns out that for every neutral mediator that induces a game with a PNE, there is a PNE when both players select the same location.

Proposition 3. *There is no neutral mediator such that*

- *(x, y) is a PNE for some $x, y \in [0, 1]$ such that $x \neq y$, and*

- *for every $z \in [0, 1]$, (z, z) is not a PNE.*

So, no neutral mediator implementing the socially optimal locations as a unique PNE exists for $n = 2$, while for $n \geq 3$ we show LIME does precisely that under low intervention cost.

Algorithm 2: GLIME

Input: A strategy profile s, a user t
Output: A location in s

1   $a_i \leftarrow q_{2i-1/2n}$ for $i \in [n-1]$
2   **if** $\exists i \in [n]$ *such that* $t \in (a_i, a_{i+1})$ **then**
3      $l \leftarrow s \cap [0, a_i], r \leftarrow s \cap [a_{i+1}, 1]$
4   **if** $l \neq \varnothing$ *and* $r \neq \varnothing$ **then**
5   w.p. $\frac{1}{2}$ **return** $\mathsf{NIME}(l, t)$, otherwise **return** $\mathsf{NIME}(r, t)$
6   **else if** $l \neq \varnothing$ *or* $r \neq \varnothing$ **then**
7   w.p. $1 - \epsilon$ **return** $\mathsf{NIME}(l \cup r, t)$, otherwise **return** $random(s, t)$
8   **else**
9   **return** $\mathsf{NIME}(s, t)$
10 **else**
11   **return** $\mathsf{NIME}(s, t)$

4.2 Configurable Intervention Cost

In some situations it might be beneficial to control the amount of intervention. A mediator may offer a tradeoff between social cost in equilibrium and intervention cost. Namely, we devise a parameterized mediator, where its parameter determines a desired intervention cost. The mediator attains the desired intervention cost but suffers some increase in its social cost in equilibrium (compared to the optimal one). For some values of n we show decisive positive results for the applicability of such a mediator in creating desired tradeoffs. This extension is deferred to the appendix.

4.3 Non-Uniform User Distribution

The results obtained for LIME do not hold for general user distribution, since players may have beneficial deviations. However, the case of non-uniform user distribution is appealing, especially given the fact that PNEs are hard to characterize (see, e.g., (Shilony 1981; Ewerhart 2015)) or may not even exist (Osborne 1993). In this subsection we devise a mediator for general user distribution, showing initial results for this setting as well. Let g be an arbitrary continuous function supported in the $[0, 1]$ segment, and denote by $G(n, \mathcal{M}, g)$ the game induced by the mediator $\mathcal{M}$, n players and g as the user distribution function. Let q_a be the a-th quantile of g. Consider the General Limited Intervention Mediator described in Algorithm 2, and referred to as GLIME hereinafter for abbreviation. GLIME operates similarly to LIME, but carries out two things differently. First, it incorporates the quantiles of g instead of those of the uniform distribution (o^n defined earlier). The other difference is given in Line 5. Instead of directing users in PIIs to their nearest facility outside that PII, it directs such users w.p. $1/2$ to the nearest facility right of the PII (if such exists), and with the remaining probability to the nearest facility left of the PII. Importantly, GLIME induces games with a unique PNE, which is highly desired.

Lemma 4. *Consider $G(n, \mathsf{GLIME}, g)$ for any $n \geq 3$ and general density function g. The unique PNE (up to renaming the players) is $s = (q_{2i-1/2n})_{i=1}^n$.*

Next, we bound the intervention cost of GLIME. Since the analysis for general distributions is highly challenging, we leave it for future work and only focus on its intervention cost under the uniform distribution.

Proposition 4. *For any $n \geq 2$, it holds that $\mathcal{I}_n(\mathsf{GLIME}) \geq \frac{1}{4} - \frac{1}{2n} + \frac{1}{2n^2}$.*

Indeed, while the intervention cost of GLIME is greater than that of LIME, it is better than that of DICT.

Lemma 5. *For every n, $\mathcal{I}_n(\mathsf{GLIME}) < \mathcal{I}_n(\mathsf{DICT})$.*

5 Discussion

Another interesting solution concept that can be employed, if the domain permits it, is to allow the mediator to direct users in some cases with w.p. less than 1. Namely, to skip the players' offers and avoid producing a recommendation. Ben-Porat and Tennenholtz (2018a) implement this idea while adopting an axiomatic approach to player fairness, and introduce a mediator that induces a PNE for games with any user space (e.g., k-dimensional space or even non-metric space). However, they do not discuss user welfare. This suggests one should seek mediators that balance user welfare and player welfare simultaneously, hopefully by not intervening too much.

Acknowledgments

References

Anderson, S. P.; De Palma, A.; and Thisse, J. F. 1992. *Discrete choice theory of product differentiation.* MIT press.

Aumann, R. J. 1974. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1(1):67–96.

Ben-Basat, R.; Tennenholtz, M.; and Kurland, O. 2017. A game theoretic analysis of the adversarial retrieval setting. *Journal of Artificial Intelligence Research* 60:1127–1164.

Ben-Porat, O., and Tennenholtz, M. 2017. Shapley facility location games. In *International Conference on Web and Internet Economics (WINE) 2017*, 58–73. Springer.

Ben-Porat, O., and Tennenholtz, M. 2018a. A game-theoretic approach to recommendation systems with strategic content providers. In *Advances in Neural Information Processing Systems (NIPS) 2018*.

Ben-Porat, O., and Tennenholtz, M. 2018b. Multi-unit facility location games. *Accepted to Mathematics of Operations Research*. Preliminary version appeared in WINE 2016.

Ben-Porat, O.; Rosenberg, I.; and Tennenholtz, M. 2019. Convergence of learning dynamics in information retrieval games. In *The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI 2019)*.

Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D. 2016. *Handbook of computational social choice.* Cambridge University Press.

Brenner, S. 2010. Location (hotelling) games and applications. *Wiley Encyclopedia of Operations Research and Management Science*.

Eaton, B. C., and Lipsey, R. G. 1975. The principle of minimum differentiation reconsidered: Some new developments in the theory of spatial competition. *The Review of Economic Studies* 42(1):27–49.

Ewerhart, C. 2015. Mixed equilibrium in a pure location game: The case of n \geq 4 firms. *The BE Journal of Theoretical Economics* 15(2):457–472.

Feldman, M.; Fiat, A.; and Obraztsova, S. 2016. Variations on the hotelling-downs model. In *Thirtieth AAAI Conference on Artificial Intelligence*.

Fournier, G., and Scarsini, M. 2014. Hotelling games on networks: efficiency of equilibria. *Available at SSRN 2423345*.

Hotelling, H. 1929. Stability in competition. The Economic Journal 39(153): 41–57, 1929.

Monderer, D., and Tennenholtz, M. 2004. k-implementation. *Journal of Artificial Intelligence Research* 21:37–62.

Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V. 2007. *Algorithmic game theory*, volume 1. Cambridge University Press Cambridge.

Osborne, M. J. 1993. Candidate positioning and entry in a political competition. *Games and Economic Behavior* 5(1):133–151.

Procaccia, A., and Tennenholtz, M. 2009. Approximate mechanism design without money. In *EC-09*.

Salton, G.; Wong, A.; and Yang, C.-S. 1975. A vector space model for automatic indexing. *Communications of the ACM* 18(11):613–620.

Schummer, J., and Vohra, R. V. 2007. Mechanism design without money. *Algorithmic Game Theory* 10:243–299.

Shaked, A. 1982. Existence and computation of mixed strategy nash equilibrium for 3-firms location problem. *The Journal of Industrial Economics* 93–96.

Shapley, L. S. 1952. A value for n-person games. Technical report, DTIC Document.

Shen, W., and Wang, Z. 2017. Hotelling-downs model with limited attraction. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, 660–668. International Foundation for Autonomous Agents and Multiagent Systems.

Shilony, Y. 1981. Hotelling's competition with general customer distributions. *Economics Letters* 8(1):39–45.

Shoham, Y., and Tennenholtz, M. 1995. On social laws for artificial agent societies: off-line design. *Artificial Intelligence* 73(1-2):231–252.