

Analysis and Optimization of Multi-Dimensional Percentile Mechanisms

Xin Sui

University of Toronto
Dept. of Computer Science
xsui@cs.toronto.edu

Craig Boutilier

University of Toronto
Dept. of Computer Science
cebly@cs.toronto.edu

Tuomas Sandholm

Carnegie Mellon University
Computer Science Department
sandholm@cs.cmu.edu

Abstract

We consider the mechanism design problem for agents with single-peaked preferences over multi-dimensional domains when multiple alternatives can be chosen. Facility location and committee selection are classic embodiments of this problem. We propose a class of *percentile mechanisms*, a form of generalized median mechanisms, that are strategy-proof, and derive worst-case approximation ratios for social cost and maximum load for L_1 and L_2 cost models. More importantly, we propose a sample-based framework for optimizing the choice of percentiles relative to any prior distribution over preferences, while maintaining strategy-proofness. Our empirical investigations, using social cost and maximum load as objectives, demonstrate the viability of this approach and the value of such optimized mechanisms *vis-à-vis* mechanisms derived through worst-case analysis.

1 Introduction

Mechanism design deals with design of protocols to elicit the preferences of self-interested agents so as to achieve a certain social objective. An important property of mechanisms is *strategy-proofness*, which requires that agents have no incentive to misreport their preferences to the mechanism. While payments are often used to ensure that mechanisms are strategy-proof [Vickrey, 1961; Clarke, 1971; Groves, 1973], in many settings payments are infeasible and restrictions on preferences are required. The simple but elegant class of *single-peaked preferences* is one such example: roughly speaking, each agent has a single, most-preferred point in the alternative space, with alternatives becoming less preferred the further they are from that point. In such settings, choosing a single alternative can be accomplished in a strategy-proof fashion using the well-known *median mechanism* [Black, 1948] and its generalizations [Moulin, 1980; Barberà, 2010]. Such models are used frequently for political choice, facility location, and other problems. They also have potential applications in areas such as in the design of a family of products, customer segmentation, and related tasks, as we discuss below.

Unfortunately, such mechanisms are efficient (e.g., w.r.t. social cost) only in very limited circumstances. Furthermore, allowing the choice of multiple alternatives (e.g., mul-

iple facilities) generally causes even these limited guarantees to evaporate. In response, research has begun to address the question of *approximate mechanism design without money* [Procaccia and Tennenholtz, 2009], which focuses on the design of strategy-proof mechanisms for problems such as multi-facility location that are approximately efficient (i.e., have good approximation ratios) [Lu *et al.*, 2010; Fotakis and Tzamos, 2010]. This work provides some positive results, but is generally restricted to settings involving two facilities (or adopts other restrictions) and L_2 (Euclidean) preferences.

In this paper, we propose *percentile mechanisms*, a type of generalized median mechanism [Barberà *et al.*, 1993; Barberà, 2010], but we address a more general class of problems. Specifically: (a) we consider problems involving selection of *multiple* alternatives (e.g., multi-facilities) in a multi-dimensional alternative space; (b) we address both social cost and maximum load as performance metrics; and (c) we analyze our mechanisms relative to L_1 (Manhattan) and L_2 (Euclidean) preferences.

Our first contribution is the analysis of the approximation ratios of percentile mechanisms under various assumptions. The performance guarantees of such mechanisms under worst-case assumptions are quite discouraging (like previous results). Indeed, designing mechanisms that have worst-case guarantees may lead to poor performance in practice. Our second contribution is the development of a sample-based *empirical framework for optimizing percentile mechanisms* relative to a known preference distribution. In most realistic applications, such as facility location, product design, and many others, the designer will have *some* knowledge of the preferences of participating agents. Assuming this takes the form of a distribution, we use profiles sampled from this distribution to optimize percentiles while maintaining strategy-proofness. Our empirical results demonstrate that, by exploiting probabilistic domain knowledge, we obtain strategy-proof mechanisms that outperform mechanisms designed to guard against worst-case profiles. Our framework can be viewed as a form of *automated mechanism design (AMD)*, which advocates the use of preference (or type) distributions to optimize mechanisms [Conitzer and Sandholm, 2002; Sandholm, 2003].

2 Preliminaries

In this section, we introduce our model along with required concepts, notation, and motivation, and then briefly discuss a selection of related work.

2.1 The Social Choice Problem

In social choice, the goal is selection of an *outcome* $o \in O$, where each agent $i \in N = \{1, 2, \dots, n\}$ has a preference over O . Preferences are represented by a (weak) total order over O , or by a *utility function*. We focus on the m -dimensional, q -*facility location problem* (or (m, q) -problems): we must choose q points or *locations* in an m -dimensional space \mathbb{R}^m (or some bounded subspace thereof) at which to place facilities. Outcomes are thus *location vectors* of the form $\mathbf{x} = (x_1, \dots, x_q)$, with $x_j \in \mathbb{R}^m$ (for $j \leq q$). Each agent i has a type t_i fixing the *cost* $c_i(\mathbf{x}, t_i)$ associated with any location $\mathbf{x} \in \mathbb{R}^m$. Given outcome \mathbf{x} , i uses the location with least cost, hence $c_i(\mathbf{x}, t_i) = \min_{j \leq q} c_i(x_j, t_i)$.

Facility location can be interpreted literally, and naturally models the placement of q facilities (e.g., warehouses, public facilities, etc.) in some geographic space where agents use the least cost (or “closest”) facility. However, many other social choice problems fit within this class. Voting is one example [Black, 1948; Barberà, 2010]: political candidates can be ordered along several dimensions (e.g., stance on environment, health care, fiscal policy)—voters have preferences over points in this space—and one must elect q candidates to a legislative body. In product design, a vendor may launch a family of q related products, each described by an m -dimensional feature vector, with consumer preferences over these options leading them to select their most preferred. This also can serve as a form of customer segmentation.

It is natural to assume agent preferences are *single-peaked* in settings such as those above. Intuitively, this means the agent has a single “ideal” location, and its cost for any chosen location increases as it “moves away from” this ideal. Formally, we need only a strict ordering on alternatives, rather than a distance metric, to define a *betweenness relation*.

Definition 1 [Barberà et al., 1993] *An agent i ’s preference on m -dimensional space \mathbb{R}^m is single-peaked if there exists a most preferred alternative $\tau(t_i)$ such that, $\forall \alpha, \beta \in \mathbb{R}^m$ satisfying $\|\tau(t_i) - \beta\|_1 = \|\tau(t_i) - \alpha\|_1 + \|\alpha - \beta\|_1$, we have $c_i(\alpha, t_i) \leq c_i(\beta, t_i)$, where $\|\cdot\|_1$ is the L_1 -norm.*

Single-peaked preferences require that if a point α lies within the “bounding box” of $\tau(t_i)$ and β , then α is at least as preferred as β . Intuitively, as we move farther away from i ’s ideal location $\tau(t_i)$ we can reach α via some path before we reach β . Note that this does not restrict i ’s relative preference for α and β if neither lies within the other’s bounding box with respect to the ideal point $\tau(t_i)$.

An agent’s ideal location $\tau(t_i)$ does not *fully* determine its preferences, even when single-peaked. Despite this, we will equate an agent’s type t_i with its ideal location (for reasons that become clear below). However, within the class of single-peaked preferences, we can adopt specific cost functions that *are* fully determined by the ideal location t_i . Often

distance metrics are used, and we consider both L_1 (Manhattan) and L_2 (Euclidean) distances below. Specifically, we define distance-based cost functions for i as follows:

$$c_i^p(\mathbf{x}, t_i) = \min_{j \leq q} \|t_i - x_j\|_p \quad (1)$$

where $p \in \{1, 2\}$ reflects either L_1 or L_2 distance from i ’s nearest facility. We use $x^p[i; \mathbf{x}]$ to denote i ’s closest facility in the location vector \mathbf{x} under the L_p -norm.

The aim in facility location is to select a set of q facilities that minimize some social objective. One natural objective is to minimize *social cost (SC)* given type profile \mathbf{t} , where social cost (relative to some norm p) is given by:

$$SC_p(\mathbf{x}, \mathbf{t}) = \sum_i c_i^p(\mathbf{x}, t_i) \quad (2)$$

Alternatively, we could try to balance the *load* by ensuring no facility is used by too many agents. Defining load on facility j given outcome \mathbf{x} and type profile \mathbf{t} as $l_j^p(\mathbf{x}, \mathbf{t}) = |\{i | x^p[i; \mathbf{x}] = j\}|$, we minimize the *maximum load (ML)*:

$$ML_p(\mathbf{x}, \mathbf{t}) = \max_j l_j^p(\mathbf{x}, \mathbf{t}). \quad (3)$$

This objective makes sense, for instance, when a product designer launches a family of q new products, consumers purchase the product closest to their ideal product, but costs are minimized by balancing production; or when facility management costs increase superlinearly with load. Many other fundamental social objectives, such as fairness (e.g., maximum agent distance), and combinations thereof can be adopted depending on one’s design goals.

2.2 Mechanisms

We now consider *direct mechanisms*, where agents reveal their types and an outcome is chosen based on the revealed types to maximize some social objective. In facility location with single-peaked preferences, agents declare their ideal locations: hence a *mechanism* is a function f that maps a declared type profile \mathbf{t} to an outcome $f(\mathbf{t}) \in (\mathbb{R}^m)^q$ (i.e., q m -dimensional alternatives).

A mechanism f is *strategy-proof* (or truthful) if:¹

$$c_i(f(t_i, \mathbf{t}_{-i}), t_i) \leq c_i(f(t'_i, \mathbf{t}_{-i}), t_i), \quad \forall i, t_i, t'_i, \mathbf{t}_{-i}$$

In other words, f is strategy-proof if no agent can obtain a better outcome by misreporting its true type (ideal location). *Group strategy-proofness* is defined similarly, but requires that no group of agents $S \subseteq N$ can misreport their types, in a coordinated fashion, so that all agents in S gain. That is, for all \mathbf{t} , for all S , and for all \mathbf{t}'_S , there is some $i \in S$ such that $c_i(f(\mathbf{t}), t_i) \leq c_i(f(\mathbf{t}'_S, \mathbf{t}_{-S}), t_i)$.

While the ideal is to design strategy-proof mechanisms that achieve some social objective (e.g., minimizing social cost), this is not always feasible. In $(1, 1)$ -facility location problems, if agent preferences are single-peaked, the *median mechanism*, which selects the median of all reported ideal locations, is in fact (group) strategy-proof [Black, 1948;

¹We use *strategy-proof* to refer to dominant strategy incentive compatibility; (individual) rationality is assured in our settings.

Moulin, 1980], and minimizes social cost if agent preferences are all determined by a suitable distance metric. However, when one moves to even just two facilities, strategy-proofness and efficiency are incompatible [Procaccia and Tennenholtz, 2009]. Procaccia and Tennenholtz [2009] propose *approximate mechanisms* to handle such situations: mechanisms that are strategy-proof, and come as close as possible to achieving the social objective. Formally:

Definition 2 A mechanism f has an approximation ratio ε w.r.t. social objective C if:

$$C(f(\mathbf{t}), \mathbf{t}) \leq \varepsilon \cdot \min_{\mathbf{x}} C(\mathbf{x}, \mathbf{t}).$$

Such mechanisms as ε -optimal w.r.t. objective C (or ε -efficient if C is social cost/welfare). When minimizing social cost, we assume the number of agents is greater than the number of facilities (otherwise, we can trivially locate facilities at each agent’s ideal to obtain a (group) strategy-proof, efficient mechanism). Also we focus on *non-imposing* mechanisms: once facilities are selected, agents are free to choose their favorite (otherwise, one can trivially minimize ML by assigning agents to facilities in an arbitrary balanced way).

2.3 Related Work

Black [1948] first proposed the median mechanism for $(1, 1)$ -facility location, showing it to be strategy-proof for single-peaked preferences. Moulin’s [1980] generalized median scheme allows *phantom peaks*, and is the unique class of (anonymous) strategy-proof mechanisms for such preferences. Barberà *et al.* [1993] later generalized this class further using *coalitional systems*, and provided a characterization result for $(m, 1)$ -problems. These m -dimensional *generalized median* schemes select a (single) location by choosing its coordinates in each dimension independently (in a “median-like” fashion).

Some work has considered strategy-proof mechanisms with even more restricted preferences and domain assumptions, for example: m -dimensional, *separable star-shaped* preferences (including quadratic preferences) [Border and Jordan, 1983]; symmetric, single-peaked preferences (of which L_1 and L_2 are instances) [2011]; and location on a graph (e.g., a network) relative to an extended notion of single-peakedness [Schummer and Vohra, 2002; Dokow *et al.*, 2012].

Recent attention has been focused on algorithmic aspects and approximation in strategy-proof facility location when agents have L_2 preferences. Procaccia and Tennenholtz [2009] study the one-dimensional problem, providing upper and lower bounds on the approximation ratio for social cost. Of interest here is their deterministic *left-right mechanism*, which is $(n - 1)$ -efficient for $(1, 2)$ -problems. Fotakis and Tzamos [2012] characterize the class of all deterministic strategy-proof mechanisms with bounded approximation ratios for $(1, 2)$ -problems on the real line (and show that no such mechanisms exist for $(1, q)$ -problems, $q \geq 3$). Lu *et al.* [2010] define the (randomized) *proportional mechanism* with an approximation ratio of 4 for general distance metrics, but it cannot be applied with more than two facilities. Fotakis and Tzamos [2010] show that a *winner-imposing* variant of the proportional mechanism is strategy-proof for any




Figure 1: The $(0.25, 0.75)$ -percentile mechanism ($n = 9$).

number of facilities, with an approximation ratio of $4q$. Escoffier *et al.* [2011] define the first mechanism for general multi-dimensional location problems, a randomized mechanism with approximation ratio $n/2$, but only in the very restrictive setting where the number of agents is exactly one more than the number of facilities.

Work on load balancing games is somewhat related, but differs in that cost functions reflect the externalities agents impose on one another (by sharing a facility or some other resource). Considerable research has developed price of anarchy [Koutsoupias and Papadimitriou, 1999; Berthold, 2007] and related results. However, externalities give those models a very different character than ours.

3 Percentile Mechanisms

In this section, we introduce the class of *percentile mechanisms*, a special case of *generalized median mechanisms* [Barberà *et al.*, 1993; Barberà, 2010].

3.1 One-dimensional Percentile Mechanisms

We begin with one-dimensional facility location problems to develop intuitions. We wish to place q facilities, with each agent i having a single ideal location t_i and single-peaked preferences. Without loss of generality, we rename the agents so their ideal locations are ordered: $t_1 \leq t_2 \leq \dots \leq t_n$. A *percentile mechanism* is specified by a vector $\mathbf{p} = (p_1, p_2, \dots, p_q)$, where $0 \leq p_1 \leq p_2 \leq \dots \leq p_q \leq 1$: the \mathbf{p} -percentile mechanism locates the j th facility at the p_j th percentile of the reported ideal locations. In other words, the j th location is placed at $x_j = t_{i_j}$, where $i_j = \lfloor (n - 1) \cdot p_j \rfloor + 1$.² Intuitively, we can decompose the mechanism into q independent rules, each locating one facility.

Example 1 We illustrate the $(0.25, 0.75)$ -percentile mechanism for a two-facility problem with $n = 9$ agents in Fig. 1. Ordering reported locations so that $t_1 \leq \dots \leq t_9$, the mechanism locates the first facility at $x_1 = t_3$ (since $\lfloor 8 \cdot 0.25 \rfloor + 1 = 3$) and the second at $x_2 = t_7$.

The following is an important property of the mechanism:

Theorem 1 The \mathbf{p} -percentile mechanism is (group) strategy-proof for any \mathbf{p} .

Proof sketch: We prove the theorem for the case of $q = 2$ (the proof for $q > 2$ is similar).

Let $S \subseteq N$ be a coalition of agents, $\mathbf{x} = (x_1, x_2)$ be the location vector if agents truthfully report their ideals, and $\mathbf{x}' = (x'_1, x'_2)$ be the location vector if agents in S jointly deviate from their peaks. In addition, let $\Delta_1 = x_1 - x'_1$ and

²We could equivalently use order statistics; but the percentile formulation removes dependence on the number of the agents in the mechanism’s specification.

$\Delta_2 = x'_2 - x_2$. We show that if either Δ_1 or Δ_2 is strictly greater or strictly less than 0, some agent in S is worse off in the percentile mechanism, which is sufficient to establish (group) strategy-proofness. There are four cases to consider:

I. $\Delta_1 \geq 0$ and $\Delta_2 \geq 0$. We can ignore the case where both Δ_1 and Δ_2 are 0, since no agent in S gains by misreporting if neither facility moves. Assume, w.l.o.g., that $\Delta_1 > 0$ and $\Delta_2 \geq 0$. Recall that x_1 is the p_1 th percentile among all reported peaks. Hence $\Delta_1 > 0$ implies that some agent $i \in S$, with $t_i \geq x_1$, reports a new ideal to the left of x_1 . Agent i 's cost is now:

$$\begin{aligned} c_i(\mathbf{x}', t_i) &= \min\{t_i - x'_1, x'_2 - t_i\} \\ &\geq \min\{t_i - x_1, x_2 - t_i\} = c_i(x, t_i) \end{aligned}$$

II. $\Delta_1 \geq 0$ and $\Delta_2 < 0$. In this case, there must be an $i \in S$, with $t_i \geq x_2$, that reports a new ideal to the left of x_2 ; it's cost is:

$$c_i(\mathbf{x}', t_i) = t_i - x'_2 > t_i - x_2 = c_i(x, t_i)$$

III. $\Delta_1 < 0$ and $\Delta_2 \geq 0$. This case is completely symmetric to case II.

IV. $\Delta_1 < 0$ and $\Delta_2 < 0$. The case is similar to case II: There must be an $i \in S$ whose ideal is to the right of x_2 but misreports to the left of x_2 , increasing its cost.

This establishes (group) strategy-proofness. ■

Since any percentile mechanism is strategy-proof for any class of single-peaked preferences, it prevents strategic manipulation even when applied to specific cost/preference functions. Unfortunately, percentile mechanisms can give rise to poor approximation ratios when we consider specific cost functions, specifically, L_1 or L_2 costs.

Theorem 2 *Let agents have L_1 or L_2 preferences. Let $\mathbf{p} = (p_1, p_2, \dots, p_q)$ define a percentile mechanism M . If $q \geq 3$, the approximation ratio of M w.r.t. social cost is unbounded. The approximation ratio w.r.t. maximum load is $q \cdot z$, where $z = \max_{1 \leq j \leq q} (p_{j+1} - p_{j-1})$ (let $p_0 = 0$ and $p_{q+1} = 1$).*

Proof sketch: We first show the approximation ratio is unbounded when the objective is social cost minimization. The intuition is that for any percentile vector \mathbf{p} , there is a type profile for which optimal social cost is arbitrarily small, while the mechanism-induced social cost is constant. We prove this for the case of $q = 3$ (the proof can easily be extended to $q > 3$).

Let $q = 3$, and assume that each agent's types is one of only four possible ideal locations, 0, δ , 2 and 3. For any percentile vector $\mathbf{p} = (p_1, p_2, p_3)$, consider a type profile $(\tau(t_1), \tau(t_2), \dots, \tau(t_n))$ where $\tau(t_1) = \dots = \tau(t_{i_1}) = 0$ and $\tau(t_{i_1+1}) = \dots = \tau(t_{i_2}) = \delta$, with $\lfloor (n-1) \cdot p_1 \rfloor + 1 \leq i_1 < \lfloor (n-1) \cdot p_2 \rfloor + 1 \leq i_2$. Given these reports, the \mathbf{p} -percentile mechanism locates the first two facilities at locations 0 and δ . In addition, let n_1, n_2, n_3 and n_4 be the number of agents whose ideal locations are 0, δ , 2 and 3, respectively. When δ is small enough, the mechanism incurs a social cost of n_3 (if the third facility is located at 3) or n_4 (if the third facility is located at 2). However, the optimal location of the three facilities for this profile is 0, 2 and 3, which has optimal social cost of $n_2 \cdot \delta$. Thus the approximation ratio is unbounded.

For maximum load, assume a percentile vector $\mathbf{p} = (p_1, \dots, p_q)$, and the induced location vector $\mathbf{x} = (x_1, \dots, x_q)$. For each $1 \leq j \leq q$, the number of agents using facility x_j is at most $l_j(t, f(t)) = n \cdot (p_{j+1} - p_{j-1})$; this occurs when each agent with a peak in (x_{j-1}, x_{j+1}) is closest to x_j , in which case maximum load is $ML(\mathbf{t}, f(\mathbf{t})) = n \cdot z$, where $z = \max_{1 \leq j \leq q} (p_{j+1} - p_{j-1})$. However, optimal maximum load, which is $\lceil n/q \rceil$, occurs using a location vector such that each facility is evenly loaded. So the approximation ratio is $\frac{n \cdot z}{\lceil n/q \rceil} \leq q \cdot z$. ■

Notice that the theorem does not hold for social cost with $q = 2$ facilities: the *left-right mechanism*, or $(0, 1)$ -percentile mechanism, has a bounded approximation ratio of $n - 1$ for social cost [Procaccia and Tennenholtz, 2009]. Indeed, the $(0, 1)$ -percentile mechanism is the *only* mechanism within the percentile family that has a bounded approximation ratio, and the only anonymous, deterministic mechanism with a bounded approximation ratio for $(1, 2)$ -problems.³ For $q \geq 3$, not only does no percentile mechanism have bounded approximation ratio, it has recently been shown that no deterministic mechanism has bounded approximation ratio [Fotakis and Tzamos, 2012]. This gives further motivation for the use of probabilistic priors to optimize percentiles for average-case rather than worst-case performance (see Sec. 4).

With respect to maximum load, it is natural to ask which percentile vector \mathbf{p} minimizes z in Thm. 2. We can show that the percentile mechanism that “evenly distributes” facilities is approximately optimal, and that it has the smallest approximation ratio within the family.

Proposition 1 *Let agents have L_1 or L_2 preferences. If q is odd, then the percentile mechanism with $p_j = \frac{j}{q+1}$, $\forall 1 \leq j \leq q$, is $\frac{2q}{q+1}$ -optimal w.r.t. maximum load. If q is even, then the percentile mechanism with $p_j = p_{j+1} = \frac{j+1}{q+2}$, $\forall j = 2j' - 1, 1 \leq j' \leq q/2$, is $\frac{2q}{q+2}$ -optimal w.r.t. maximum load.⁴ In each case, the mechanism has the smallest approximation ratio within the percentile family.*

All proofs omitted here can be found in a longer version of this paper⁵

3.2 Multi-dimensional Percentile Mechanisms

As discussed above, many social choice problems can be interpreted as “facility location” problems when viewed as choice in a higher dimensional space, such as selection of political/committee representatives, product design, and the like. We now analyze a generalization of the percentile mechanism to multi-dimensional spaces.

³The characterization results of Fotakis and Tzamos [2012] show that the only deterministic mechanisms for $(1, 2)$ -problems are the $(0, 1)$ -percentile mechanism and dictatorial mechanisms.

⁴For even q , the mechanism is partially imposing. We locate two facilities at each selected location, and balance the agents choosing any location; they are indifferent to the “imposed” assignment, so it isn't truly imposing mechanisms (we don't remove choice from the agents [Fotakis and Tzamos, 2010]). We use this for convenience; there are strictly non-imposing mechanisms with the same ratio.

⁵See: <http://www.cs.toronto.edu/~cebly/papers.html>.

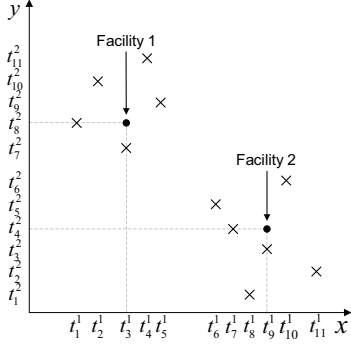


Figure 2: A percentile mechanism for the $(2, 2)$ -problem ($n = 11$).

As above, we assume that agents have single-peaked preferences (see Defn. 1). Reported types t_i are now points in \mathbb{R}^m . For any type profile \mathbf{t} , let $t_1^k \leq t_2^k \leq \dots \leq t_n^k$ be the *ordered projection* of \mathbf{t} in the k th dimension (for $k \leq m$). In other words, we simply order the reported coordinates in each dimension independently. An m -dimensional *percentile mechanism* is specified by a $q \times m$ matrix $\mathbf{P} = (\mathbf{p}_1; \mathbf{p}_2; \dots; \mathbf{p}_q)$, where each $\mathbf{p}_j \in [0, 1]^m$ is an m -vector in the unit cube, with $\mathbf{p}_j = (p_j^1, p_j^2, \dots, p_j^m)$. Given a reported profile \mathbf{t} , the \mathbf{P} -percentile mechanism locates the j th facility by selecting, for each dimension $k \leq m$, the p_j^k th percentile of the ordered projection of \mathbf{t} in the k th dimension as the coordinate of facility j in that dimension. Formally:

$$x_j = (t_{\lfloor (n-1) \cdot p_j^1 \rfloor + 1}^1, t_{\lfloor (n-1) \cdot p_j^2 \rfloor + 1}^2, \dots, t_{\lfloor (n-1) \cdot p_j^m \rfloor + 1}^m).$$

Example 2 Fig. 2 illustrates a 2-D, two facility problem with 11 agents. With $\mathbf{P} = (0.2, 0.7; 0.8, 0.3)$, the \mathbf{P} -percentile mechanism locates the first facility at the x -coordinate of t_3 (since $\lfloor 10 \cdot 0.2 \rfloor + 1 = 3$) and at the y -coordinate of t_8 ; and the second facility is placed at the x -coordinate of t_9 and the y -coordinate of t_4 . Notice facilities need not be located at the ideal point of any agent.

The following result says that the m -dimensional percentile mechanism is strategy-proof.

Theorem 3 *The m -dimensional \mathbf{P} -percentile mechanism is strategy-proof for any \mathbf{P} .*

Strategy-proofness can be easily verified. The mechanism is not group strategy-proof: in a two-dimensional model, two agents i and j can collude to misreport their preferences such that i 's misreport benefits j in one dimension, and j 's misreport benefits i in the other, making both better off.

The following results generalize the corresponding one-dimensional results above.

Theorem 4 *Let agents have L_1 or L_2 preferences, and \mathbf{P} define a percentile mechanism M for an (m, q) -facility location problem with $m > 1$. The approximation ratio of M is unbounded w.r.t. social cost. The approximation ratio of M is $q \cdot z$ w.r.t. maximum load, where $z = \prod_{k=1}^m \max_{1 \leq j \leq q} (p_{j+1}^k - p_{j-1}^k)$ (let $p_0^k = 1$ and $p_{q+1}^k = 1$).*

Notice that this result differs from the one-dimensional case, where the $(0, 1)$ -percentile (i.e., left-right) mechanism has a

Distribution		$q = 2$	$q = 3$	$q = 4$
D_u	SC	(0.25, 0.75)	(0.16, 0.5, 0.84)	(0.12, 0.37, 0.63, 0.88)
	ML	(0.49, 0.50)	(0.33, 0.35, 0.98)	(0.25, 0.26, 0.74, 0.75)
D_g	SC	(0.25, 0.75)	(0.15, 0.5, 0.85)	(0.1, 0.35, 0.65, 0.9)
	ML	(0.49, 0.50)	(0.33, 0.35, 0.9)	(0.25, 0.26, 0.74, 0.75)
D_{gm}	SC	(0.17, 0.68)	(0.16, 0.59, 0.93)	(0.12, 0.37, 0.68, 0.94)
	ML	(0.49, 0.50)	(0.14, 0.65, 0.66)	(0.17, 0.34, 0.73, 0.74)

Table 1: Optimal percentiles for different distributions, objectives, and numbers of facilities.

bounded approximation ratio for social cost. When $m > 1$, no percentile mechanism has this property—this holds because the mechanism may place no facility at the ideal location of any agent. As above, however, we can optimize the percentiles for maximum load, when $q = \tilde{q}^m$ for some \tilde{q} by exploiting Prop. 1 in each dimension:

Proposition 2 *Let $q = \tilde{q}^m$. If \tilde{q} is odd, the mechanism that locates one facility at each $\frac{1}{\tilde{q}+1}$ th percentile in each dimension is $\left(\frac{2\tilde{q}}{\tilde{q}+1}\right)^m$ -optimal w.r.t. maximum load. If \tilde{q} is even, the mechanism that locates two facilities at each $\frac{2}{\tilde{q}+2}$ th percentile in each dimension is $\left(\frac{2\tilde{q}}{\tilde{q}+2}\right)^m$ -optimal w.r.t. maximum load. Moreover, these are the smallest approximation ratios within the family of percentile mechanisms.*

4 Optimizing Percentile Mechanisms

We have seen that percentile mechanisms are strategy-proof for general (m, q) -facility location problems, and can offer bounded approximation ratios for L_1 and L_2 preferences (though only under certain conditions for social cost). Unfortunately, these guarantees require optimizing the choice of percentiles w.r.t. worst-case profiles, which can lead to poor performance in practice. For example, in a $(1, 2)$ -problem, decent approximation guarantees for social cost require using the $(0, 1)$ -percentile mechanism; but if agent preferences are uniformly distributed in one dimension, this will perform quite poorly. Intuitively, the $(0.25, 0.75)$ -percentile mechanism should have lower expected social cost due to its “probabilistically suitable” placement of two facilities, each for use by half of the agents.

We consider a framework for empirical optimization of percentiles within the family of percentile mechanisms that admits much better performance in practice. As in automated mechanism design [Conitzer and Sandholm, 2002; Sandholm, 2003], we assume a prior distribution D over agent preference profiles. One will often assume a prior model D (e.g., learned from observation) that renders individual agent preferences independent given that model, but this is not a requirement for our method. In many settings, such as facility location or product design, such distributional information will readily be available. We sample preference profiles from this distribution, and use them to optimize percentiles to ensure the best expected performance w.r.t. our social objective.

Unlike classic AMD, we restrict ourselves to the specific family of percentile mechanisms. While this limits




Figure 3: Comparison of optimized percentile mechanism and optimal value ($q = 3$).

the space of mechanisms, we do this for several reasons. First, it provides a much more compact mechanism parameterization over which to optimize than in typical AMD settings.⁶ Second, since the resulting mechanism is “automatically” strategy-proof, no matter which percentiles are chosen, the optimization need not account for incentive constraints. Third, unlike Bayesian optimization—in other words, methods that choose optimal facility placement relative to the prior with *no elicitation* of ideal locations—optimized percentile mechanisms are *responsive* to the specific preferences of the agents. (We empirically compare percentile mechanisms to Bayesian optimization below.)

Let agent type profiles $\mathbf{t} = (t_1, t_2, \dots, t_n)$ be drawn from distribution D . Given a \mathbf{P} -percentile mechanism, let $f_{\mathbf{P}}(\mathbf{t})$ denote the chosen locations when the agent type profile is \mathbf{t} . The goal is to select \mathbf{P} to minimize the expected social cost or maximum load:

$$\min_{\mathbf{P}} \mathbb{E}_D [SC_p(f_{\mathbf{P}}(\mathbf{t}), \mathbf{t})]; \text{ or } \min_{\mathbf{P}} \mathbb{E}_D [ML_p(f_{\mathbf{P}}(\mathbf{t}), \mathbf{t})]$$

Naturally, other objectives can be modelled in this way too.

Given Y sampled preference profiles, we optimize percentile selection relative to the Y sampled profiles. In our experiments below, we use simple exhaustive optimization for this purpose. Specifically, we consider all possible values for the percentile matrix \mathbf{P} . For each, we compute the average social cost (maximum load) over Y sample profiles, and select the one with minimum objective value. This is feasible for problems of the size we consider.

For larger problems, one can formulate the minimization problem as a mixed integer program (MIP) for L_1 costs, or a mixed integer quadratically constrained program (MIQCP) for L_2 costs, and use standard optimization tools, e.g., CPLEX, to solve the problem. Relaxed formulations require $O(n^m)$ variables however, rendering them intractable for problems with large numbers of agents. We also experimented with gradient and coordinate ascent algorithms from random starting points (i.e., \mathbf{P} -matrices) on all of the problems described below. These worked extremely well: no run of either algorithm on the problems below converged to a solution more than 2% from optimal (on avg. within 0.5% of optimal); and with 100 random restarts, both methods found the optimal solution in every instance (and did so quickly, in times ranging from 0.88–1.97 sec.). Further algorithmic developments remain a key focus of research.

In the following experiments, we consider problems with $n = 101$ agents, with agent preferences drawn independently

⁶AMD has been explored in parameterized mechanisms, e.g., in combinatorial auctions [Likhodedov and Sandholm, 2004; 2005].

Distr.	1D						2D		4D	
	$n = 101$			$n = 21$			$n = 101$	21	$n = 101$	21
	$q = 2$	3	4	$q = 2$	3	4	$q = 3$		$q = 2$	
D_u	2.2	3.0	3.8	9.7	18.5	24.6	1.4	7.4	1.0	6.2
D_g	1.4	2.3	3.1	11.6	19.7	27.9	1.5	5.4	0.9	2.9
D_{gm}	2.2	1.7	3.8	8.2	11.9	21	1.2	6.2	0.9	3.6

Table 2: Percentage improvement in social cost of optimized percentile mechanism vs. Bayesian optimization.

from three classes of distributions: uniform D_u , Gaussian D_g and mixture of Gaussians D_{gm} with 2 or 3 components. Each distribution reflects rather different assumptions about agent preferences: that they are spread evenly (D_u); that they are biased toward one specific location (D_g); or that they partitioned into 2 or 3 loose clusters (D_{gm}). In all cases, $T = 500$ sampled profiles are used for optimization. We examine results for both social cost and maximum load.

One-dimensional mechanisms

We begin with simple one-dimensional problems with $q = 2, 3$ or 4. Table 1 shows the percentiles resulting from our optimization for both SC and ML under each of the three distributions.⁷ For example, when agent ideal locations are uniformly distributed, the (0.25, 0.75)-percentile mechanism minimizes the expected social cost for two facilities. This is expected, since the uniform (and Gaussian) distribution partitions agents into two groups of roughly equal size, and facilities should be located at the median positions of each group.

The performance of the optimized percentile mechanisms is extremely good. Fig. 3 compares the expected social cost and maximum load of our mechanisms with those given by *optimal placement* of facilities (results for $q = 3$ are shown, but others are similar). Recognize however that optimal placement is not realizable with any strategy-proof mechanism. Despite this, optimized percentile mechanisms perform nearly as well, in expectation, as optimal placement in all three cases. Contrast this with the performance of the mechanisms with provable approximation ratios. When $q = 2$, the (0, 1)-percentile mechanism has an average social cost of 242.4, 340.9 and 523.2 for D_u , D_g and D_{gm} , respectively; but the social cost of our mechanisms are only 123.7, 76.5, and 165.1, respectively. When $q = 3$, the (0.25, 0.5, 0.75)-percentile mechanism has the best approximation ratio for ML (see Prop. 1). Its average maximum loads are 39.5, 38.7 and 38.3, which are close to (but not as good as) the loads of the optimized percentile mechanisms (36.5, 36.5, and 36.2).

We also compare the performance (w.r.t. social cost only) of our optimized percentile mechanism with Bayesian optimization (see column **1D** in Table 2). Bayesian optimization performs almost as well as the optimal percentile mechanism when the number of agents is large. However, for the smaller population, eliciting ideal locations using the percentile mechanism gives much better results than the Bayesian approach. For example, when $q = 4$, the optimized percentile mechanism has an expected cost that is 3.1% better

⁷ D_u is uniform on $[0, 10]$. D_g is Gaussian $\mathcal{N}(0, 2)$. D_{gm} is a Gaussian mixture with 3 components: $\mathcal{N}(-4, 4)$ (weight 0.4), $\mathcal{N}(0, 1)$ (weight 0.45), and $\mathcal{N}(5, 2)$ (weight 0.15).




Figure 4: Optimized Percentiles for (a) **2D**: Uniform, (b) **2D**: Gaussian, (c) **2D**: Gaussian mixture, and (d) **4D**.

than the Bayesian model with $n = 101$ agents; but the performance gaps grows to 27.9% with $n = 21$ agents. In addition, we see that the agent-facility ratio also matters (i.e., when there are more facilities, the percentile mechanism tends to exhibit a greater performance gap).

These results are not surprising in this i.i.d. setting: indeed simple law-of-large-numbers arguments suggest that no elicitation of ideal points is needed at all for optimal placement given a sufficiently large population.⁸ However, our framework does not require this i.i.d. assumption—preferences can be arbitrarily correlated. In such a case, Bayesian optimization can work extremely poorly. For example, consider 1-D, 2-facility problem in which a latent variable V correlates preferences: if V is true, ideal points are drawn from a Gaussian $\mathcal{N}(\mu_1, \sigma)$; otherwise, they are drawn from $\mathcal{N}(\mu_2, \sigma)$. If each realization of V is equally likely, optimal Bayesian placement selects facilities at each of μ_1 and μ_2 . By contrast, the optimal percentile mechanism is a simple function of σ , and will place facilities around the mean of the single “true” Gaussian, greatly improving social cost.

Multi-dimensional mechanisms

We also experimented with two additional problems. **2D** is a (2, 3)-problem where agents have L_2 preferences, capturing, say, the placement of three public projects like libraries, or warehouses. **4D** is a (4, 2)-problem with L_1 preferences, which might model the selection of 2 products for launch, each with four attributes that predict consumer demand.⁹

For the problem **2D** we show the *expected* placement of facilities given the selected percentiles in Fig. 4(a)-(c), for both *SC* and *ML*, for each of the three distributions. (*Actual* facility placement will shift to match the reported type profile in each instance.) Placement for *SC* tends to be distributed appropriately, while *ML* places two facilities adjacent to one another. For **4D**, we measure performance rather than visualizing locations. Fig. 4(d) compares expected *SC* and *ML* of our optimized percentile mechanisms to those using true optimal facility placements: the percentile mechanisms are always optimal for *ML*;¹⁰ and for *SC*,

⁸Thanks to Lirong Xia for this observation.

⁹For **2D**, D_u is uniform over $[0, 10]$ in each dimension. D_g is normal with mean $\mu = [3, 2]$ and covariance $\Sigma = [2, 1]\mathbf{I}$. D_{gm} is a 2 component mixture: $\mathcal{N}([-2, -1], [2, 1]\mathbf{I})$ (weight 0.3) and $\mathcal{N}([0, 2], [1, 3]\mathbf{I})$ (weight 0.7). For **4D**, D_u is uniform over $[0, 10]$ in each dimension. D_g is $\mathcal{N}([3, 2, 1, 2], [2, 3, 4, 1]\mathbf{I})$. D_{gm} is a 2 component mixture: $\mathcal{N}([2, 1, 0, 1], [4, 6, 8, 5]\mathbf{I})$ (weight 0.4) and $\mathcal{N}([1, 2, 1, 0], [7, 4, 5, 8]\mathbf{I})$ (weight 0.6).

¹⁰This is because the mechanism locates two facilities at almost the same position, and achieves optimal maximum load. However,

placements using our optimized strategy-proof mechanisms are only 1.77%-4.66% worse than the corresponding non-strategy-proof optimal placements. This strongly suggests that percentile mechanisms, optimized using priors over preferences, are well-suited to multi-dimensional, single-peaked domains. The improvement of optimized percentile mechanisms over Bayesian optimization (see columns **2D** and **4D** in Table 2) exhibits trends similar to those in the **1D** case.

5 Conclusions and Future Research

Our *percentile mechanisms* for multi-dimensional, multi-facility location problems, are strategy-proof for single-peaked preferences, but at the same time admit considerable flexibility in optimization. While worst-case approximation ratios seem discouraging, sample-based optimization that exploits priors over preferences allows strong performance w.r.t. one’s social objectives. Indeed, our mechanisms give solutions that are, in practice, extremely close to the optimum attainable with exact knowledge of agent preferences.

This work is just a starting point for the design of optimized mechanisms for single-peaked domains, and can be extended in several ways, including: mechanisms for more specific classes of single-peaked preferences (e.g., quadratic [Border and Jordan, 1983] or symmetric single-peaked [Massó and Moreno de Barreda, 2011]); and other social objectives, including those that combine various desiderata (e.g., *SC* and *ML*), and those that trade off facility cost with benefit (e.g., new facilities may decrease social cost, but their expense must be factored in [Lu and Boutilier, 2011]). Further development of optimization methods for percentile mechanisms (e.g., our MIP or MIQCP formulations) are needed to make our approach more scalable; preliminary experiments with local search techniques are very promising in this respect. Sample complexity results—theoretical bounds on the number of sampled profiles needed by our technique to ensure near-optimal results with high probability—are also of interest. Finally, incremental (or multi-stage) mechanisms that trade off social cost, communication costs, and agent privacy [Sandholm *et al.*, 2007; Feigenbaum *et al.*, 2010; Sui and Boutilier, 2011] would be extremely valuable.

Acknowledgements: Thanks to Lirong Xia and the anonymous reviewers for helpful comments and Alex Francois-Nienaber for evaluating our heuristic search methods. Sui and Boutilier acknowledge the support of NSERC. Sandholm was supported by the NSF under grants CCF-1101668 and IIS-0964579.

this is not always possible for three or more facilities.

References

- [Barberà *et al.*, 1993] Salvador Barberà, Faruk Gul, and Ennio Stacchetti. Generalized median voter schemes and committees. *J. Economic Theory*, 61(2):262–289, 1993.
- [Barberà, 2010] Salvador Barberà. Strategy-proof social choice. In K. J. Arrow, A. K. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 2, pp.731–832. North-Holland, Amsterdam, 2010.
- [Berthold, 2007] Vöcking Berthold. Selfish load balancing. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, pp.517–542. Cambridge University Press, 2007.
- [Black, 1948] Duncan Black. On the rationale of group decision-making. *J. Political Econ.*, 56(1):23–34, 1948.
- [Border and Jordan, 1983] Kim C. Border and J. S. Jordan. Straightforward elections, unanimity and phantom voters. *The Review of Economic Studies*, 50(1):153–170, 1983.
- [Clarke, 1971] Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.
- [Conitzer and Sandholm, 2002] Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. In *Proc. Eighteenth Conference on Uncertainty in Artificial Intelligence (UAI-02)*, pp.103–110, Edmonton, 2002.
- [Dokow *et al.*, 2012] Elad Dokow, Michal Feldman, Reshef Meir, and Ilan Nehama. Mechanism design on discrete lines and cycles. In *Proc. Thirteenth ACM Conference on Electronic Commerce (EC’12)*, pp.423–440, Valencia, Spain, 2012.
- [Escoffier *et al.*, 2011] Bruno Escoffier, Laurent Gourvès, Thang Nguyen Kim, Fanny Pascual, and Olivier Spanjaard. Strategy-proof mechanisms for facility location games with many facilities. In *Proc. Second International Conference on Algorithmic Decision Theory (ADT-11)*, pp.67–81, Piscataway, NJ, 2011.
- [Feigenbaum *et al.*, 2010] Joan Feigenbaum, Aaron D. Jagard, and Michael Schapira. Approximate privacy: Foundations and quantification (extended abstract). In *Proc. Eleventh ACM Conference on Electronic Commerce (EC’10)*, pp.167–178, Cambridge, MA, 2010.
- [Fotakis and Tzamos, 2010] Dimitris Fotakis and Christos Tzamos. Winner-imposing strategyproof mechanisms for multiple facility location games. In *Proc. Sixth International Workshop on Internet and Network Economics (WINE-10)*, pp.234–245, Stanford, CA, 2010.
- [Fotakis and Tzamos, 2012] Dimitris Fotakis and Christos Tzamos. On the power of deterministic mechanisms for facility location games. arXiv: 1207.0935, 2012.
- [Groves, 1973] Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [Koutsoupias and Papadimitriou, 1999] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Proc. 16th Symposium on Theoretical Aspects of Computer Science (STACS-99)*, pp.404–413, Trier, DE, 1999.
- [Likhodedov and Sandholm, 2004] Anton Likhodedov and Tuomas Sandholm. Methods for boosting revenue in combinatorial auctions. In *Proc. Nineteenth National Conference on Artificial Intelligence (AAAI-04)*, pp.232–237, Pittsburgh, PA, 2004.
- [Likhodedov and Sandholm, 2005] Anton Likhodedov and Tuomas Sandholm. Approximating revenue-maximizing combinatorial auctions. In *Proc. Twentieth National Conference on Artificial Intelligence (AAAI-05)*, pp.267–274, Pittsburgh, PA, 2005.
- [Lu and Boutilier, 2011] Tyler Lu and Craig Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proc. Twenty-second International Joint Conference on Artificial Intelligence (IJCAI-11)*, pp.280–286, Barcelona, 2011.
- [Lu *et al.*, 2010] Pinyan Lu, Xiaorui Sun, Yajun Wang, and Zeyuan A. Zhu. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *Proc. Eleventh ACM Conference on Electronic Commerce (EC’10)*, pp.315–324, Cambridge, MA, 2010.
- [Massó and Moreno de Barreda, 2011] Jordi Massó and Inés Moreno de Barreda. On strategy-proofness and symmetric single-peakedness. *Games and Economic Behavior*, 72(2):467–484, 2011.
- [Moulin, 1980] Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.
- [Procaccia and Tennenholtz, 2009] Ariel D. Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. In *Proc. Tenth ACM Conf. on Electronic Commerce (EC’09)*, pp.177–186, Stanford, CA, 2009.
- [Sandholm *et al.*, 2007] Tuomas Sandholm, Vincent Conitzer, and Craig Boutilier. Automated design of multistage mechanisms. In *Proc. Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)*, pp.1500–1506, Hyderabad, India, 2007.
- [Sandholm, 2003] Tuomas Sandholm. Automated mechanism design: A new application area for search algorithms. In *Proc. Intl. Conference on Principles and Practice of Constraint Programming (CP-03)*, Kinsale, Ireland, 2003.
- [Schummer and Vohra, 2002] James Schummer and Rakesh V. Vohra. Strategy-proof location on a network. *J. Economic Theory*, 104(2):405–428, 2002.
- [Sui and Boutilier, 2011] Xin Sui and Craig Boutilier. Efficiency and privacy tradeoffs in mechanism design. In *Proc. Twenty-fifth AAAI Conference on Artificial Intelligence (AAAI-11)*, pp.738–744, San Francisco, CA, 2011.
- [Vickrey, 1961] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.