Bayesian Strategy-Proof Facility Location via Robust Estimation

Manolis Zampetakis
UC Berkeley

Fred Zhang
UC Berkeley

Abstract

A seminal work by Moulin (1980) shows that the median voting scheme fully characterizes (deterministic) strategy-proof facility location mechanism for single-peaked preferences. In this simple setting, median also achieves the optimal social cost. In d dimensions, strategy-proof mechanism is characterized by coordinate-wise median, which is known to have a large $\sqrt{d}$ approximation ratio of the social cost in the Euclidean space, whereas the socially optimal mechanism fails at being strategy-proof. In light of the negative results in the classic, worst-case setting, we initiate the study of Bayesian mechanism design for strategy-proof facility location for multi-dimensional Euclidean preferences, where the agents' preferences are drawn from a distribution. We approach the problem via connections to algorithmic high-dimensional robust statistics. Specially, our contributions are the following:

(i) We provide a general reduction from any robust estimation scheme to Bayesian approximately strategy-proof mechanism. This leads to new strategy-proof mechanisms for Gaussian and bounded moment distributions, by leveraging recent advances in robust statistics.

(ii) We show that the Lugosi-Mendelson median arising from heavy-tailed statistics can be used to obtain Bayesian approximately strategy-proof single-facility mechanism with asymptotically optimal social cost, under mild distributional assumptions.

(iii) We provide Bayesian approximately strategy-proof multi-facility mechanisms for Gaussian mixture distributions with nearly optimal social cost.

Proceedings of the 26th International Conference on Artificial Intelligence and Statistics (AISTATS) 2023, Valencia, Spain. PMLR: Volume 206. Copyright 2023 by the author(s).

1 INTRODUCTION

Facility location with strategic agents is a fundamental problem in Mechanism Design without money (Moulin, 1980; Procaccia and Tennenholtz, 2013; Feldman et al., 2016). In the facility location game of n agents, each agent reports a location, and the mechanism chooses a set of m facility locations. The cost of an agent is the distance from their true location to the closest facility location. The agents are strategic and seek to minimize their individual cost. We study the problem of designing strategy-proof mechanisms, where an agent cannot be better off by misreporting their location. In addition, we aim to design mechanisms with minimal social cost, the sum of individual costs over all agents.

Identifying truthful mechanisms with small social cost turns out to be a challenging problem, even in simple settings. In a seminal work, Moulin (1980) first provides a complete characterization of strategy-proof single-facility mechanisms on the real line, known as the (generalized) median voter schemes. As a special case, the result of Moulin implies that taking the median of the location profile is strategy-proof. Border and Jordan (1983) extend the result to multi-dimensional Euclidean space and shows that it is strategy-proof to apply the median voter schemes in each dimension separately. This natural coordinate-wise median approach, however, does not lead to socially desirable outcome. Indeed, Meir (2019) shows that it achieves only a $\sqrt{d}$ approximation factor for the social cost, in d-dimensional Euclidean space. Such approximation quality can be unacceptable in high dimensional settings. On the other hand, the socially optimal mechanism—taking the geometric median that minimizes the total distance—violates the strategy-proof property (Goel and Hann-Caruthers, 2020). Hence, the negative results suggest a dilemma that single-facility mechanisms cannot attain strategy-proofness and optimal social cost simultaneously. Fundamental work has also been done for placing a single facility in other metric spaces as well (Feldman and Wilf, 2013).

The problem of placing multiple facilities is significantly harder as the existing literature suggests. Procaccia and Tennenholtz (2013) introduce the framework of mechanism design without money and initiates the study of m-facility strategy-proof mechanisms. For $m = 2$ and the line metric, they show that placing one facility at the leftmost peak and the other at the rightmost is strategy-proof, but the scheme only achieves an $(n-2)$-approximation for the social cost. Fotakis and Tzamos (2014) further prove that this approximation is tight for any deterministic mechanism. Though the approximation factor can be improved to 4 using a randomized PROPORTION mechanism, it is observed that PROPORTION fails to be strategy-proof for more than two facilities (Lu et al., 2010). Moreover, the approximation ratio becomes unbounded for $m > 2$ facilities and anonymous, strategy-proof mechanisms (Fotakis and Tzamos, 2014).

In light of the negative results in the classic, worst-case settings, we initiate the study of *Bayesian strategy-proof mechanisms* for facility location in multi-dimensions, where the location profile is drawn from a distribution. Our work seeks to break the tension between incentive-compatibility and social cost, for both single- and multi-facility mechanisms. In particular, we ask:

Can we achieve strong social cost guarantees and maintain strategy-proofness simultaneously in Bayesian facility location mechanisms design?

1.1 Our Contributions

We provide a mostly positive answer to the main question. Our approach is inspired by an intimate connection between strategy-proof mechanisms and *robust statistics*—the study of statistics in presence of data outliers—regarding the use of medians. On one hand, median voting schemes is the only class of deterministic strategy-proof mechanisms on the real line (Moulin, 1980). On the other, one-dimensional median is known to be resilient to outliers Huber (1973) and commonly used as a robust location estimator.

Our work can be seen as an extension of this connection to high dimensions and exploit high-dimensional median constructions. Recent literature on algorithmic robust statistics proposed new high-dimensional medians for outlier-robust parameter estimation. In particular, we first study the performance of *Lugosi-Mendelson (LM) median*, a notion of high-dimensional median for mean estimation with heavy-tailed data, due to Lugosi and Mendelson (2019a). Concretely, a *LM r-median* is a point close to the median of the data under any one-dimensional projection, up an additive factor of r; see Definition 2.7 for a formal definition. We consider applying LM median for single-facility mechanism.

Single-facility mechanism. One naïve hope may be that the LM median is a strategy-proof single-facility mechanism for multi-dimensional Euclidean preferences, even in classic, non-Bayesian setting.[1] Unfortunately, as a negative result, we prove that this is not true. Indeed, Peters et. al. show that any strategy-proof, Pareto-optimal, and anonymous mechanism must be a coordinate-wise median scheme. As LM median satisfies Pareto-optimality and anonymity, it must fail at being truthful. We formalize this argument in Theorem 3.1.

Given this result, we move on to study (approximately) strategy-proof mechanisms in Bayesian settings. We say that a mechanism is Bayesian (ε, k)-group strategy-proof if no agent in a group of k can be better off by ε in expectation, by collectively misreporting their true preferences. To achieve Bayesian strategy-proofness, we give a general connection to the robust statistics. We say that an algorithm is (r, k)-robust if its output does not change by r, in Euclidean norm, if k input points are adversarial and others are i.i.d. Roughly speaking, we show as a general reduction:

Theorem 1.1 (Informal; see Theorem 4.5)**.** *Any (r, k)-robust algorithm can be used as a Bayesian $(O(r), O(k))$-group strategy-proof mechanism.*

This allows us to exploit a wide range of results in the recent literature on algorithmic robust statistics (Diakonikolas and Kane, 2019), and to immediately obtain approximately strategy-proof mechanisms for various distributions, including Gaussian and bounded-moment distributions.

As an interesting special case, we again consider LM median, originally used for mean estimation under heavy-tailed distributions. Lugosi and Mendelson show that with probability at least $1-2^{-\Omega(k)}$, a LM r-median exists for n random vectors i.i.d. from a distribution with variance bounded by $O(1)$ at every direction, for $r = O(\sqrt{d/n} + \sqrt{k/n})$. Further, an r-median $\widehat{\mu}$ satisfies that $\|\widehat{\mu} - \mu\| \leq r$, where μ is the mean of the distribution. Note that for fixed k and d, the bound tends to 0 as $n \to \infty$. Moreover, subsequent work (Lei et al., 2020; Depersin and Lecué, 2019) prove that this guarantee holds even when $k/100$ of the input points are arbitrary, and they give polynomial-time algorithms for finding a LM r-median. This shows that the LM r-median is $(O(r), O(k))$-robust, around mean μ. The property allows us to apply the reduction (Theorem 1.1) and obtain Bayesian approximately strategy-proof mechanism as a result.

On the other hand, another objective in facility location mechanism design is to achieve socially desirable outcome. For (centrally) symmetric distributions, we show that the LM r-median achieves asymptotically optimal social cost. Combining this with the strategy-proof property, we get:

Theorem 1.2 (Informal; see Theorem 5.2)**.** *For any symmetric distribution with bounded covariance, the LM r-median (1) obtains asymptotically optimal social cost, as $n \to \infty$, and (2) is $(\varepsilon, O(k))$-group strategy-proof with $\varepsilon \to 0$ as $n \to \infty$.*

[1] Here, we consider a slight variant of the original LM median, in order to guarantee existence and uniqueness. We give a natural definition of such "unique LM median" (see Definition 2.8)

Multiple facility mechanism. We also study multi-facility mechanisms, focusing on Gaussian and Gaussian mixtures. We consider the practical case of a small number of facilities. The intuition here is simple. For high-dimensional Gaussian, the location profile is sufficiently well spread-out in all directions. Thus, we do not expect that a constant number of facilities would dramatically reduce the social cost than an optimal 1-facility mechanism. We formally prove such key lemma in section 7 for spherical Gaussian and its mixture. The lemma implies that placing just 1 facility at the mean of the Gaussian (component) is socially near optimal. On the other hand, Bakshi et al. (2020) recently established that Gaussian and Gaussian mixtures can be learned robustly. Again by exploiting our reduction to robust statistics, we obtain approximately Bayesian group strategy-proof mechanisms in these settings. See Theorem D.2 and Theorem 7.2 for formal statements.

1.2 Related Work

Strategy-proof facility location. Motivated by social choice theory, classic literature on facility location was mostly focused on single facility. Moulin Moulin (1980) proves a full characterization of any strategy-proof deterministic mechanism under single-peaked preferences. The result was later extended to multidimensional outcome space by Border and Jordan (1983); Barberà et al. (1993). The modern era of strategy-proof facility location started off with Procaccia and Tennenholtz (2013), which studies multi-facility mechanisms. This leads to a sequence of followup work Alon et al. (2010); Lu et al. (2009, 2010); Sui et al. (2013); Fotakis and Tzamos (2014). See Chan et al. (2021) for a recent survey of the area.

For single dimensional preferences, Caragiannis et al. (2016) studies univariate facility location problem. They show that in the Baeysian setting, the generalized median schemes still fully characterize the strategy-proof mechanisms. Within this class, they provide estimators that achieve small error. As generalized median mechanisms suffer high social cost in high dimensions, this motivates us to relax the exact strategy-proofness requirement.

More closely related to our paper, there has been recent work on strategy-proof facility location in multi-dimensions, using other notions of median. In particular, El-Mhamdi et al. (2021) provides an analysis for geometric median and Goel and Hann-Caruthers (2020) focuses on the cost of coordinate-wise median under other social objectives. Neither considers the Lugosi-Mendelson median or shows a general connection to robust statistics. Walsh Walsh (2020) studies the problem in Manhattan space. Finally, we mention that Caragiannis et al. (2016) also considers distributional settings, though only in one dimension.

Robust statistics. Robust statistics is a classic area, dating back to the the 1960s Tukey (1960); Huber (1964). There has been a recent surge of interests in designing computationally efficient estimation schemes in high dimensions, for classic problem such as mean estimation. The work of Lai et al. (2016); Diakonikolas et al. (2021) first gave polynomial time algorithms for statistically near optimal robust mean estimation under Gaussian and bounded covariance distributions. We leverage their results to provide approximate Bayesian group strategy-proof mechanisms when the location profile is drawn i.i.d. from such distributions. These results have since been extended and improved (Balakrishnan et al., 2017; Diakonikolas et al., 2017, 2018a, 2019a; Cheng et al., 2019; Hopkins and Li, 2018; Dong et al., 2019). See Diakonikolas and Kane (2019) for a recent survey.

A closely related area is estimation under heavy-tailed distributions, that is, distributions with weak concentration properties. For bounded second moment distributions, there has been a long line of work on mean estimation at sub-gaussian error rate (Minsker, 2015; Devroye et al., 2016; Hsu and Sabato, 2016; Joly et al., 2017; Lugosi and Mendelson, 2021, 2019b,a; Lee and Valiant, 2022). Most relevant to us, Lugosi and Mendelson (2019a) proposes a notion of high-dimensional median that acts as statistically optimal estimator in this setting. The estimator admits efficient algorithms (Hopkins, 2020; Cherapanamjeri et al., 2019; Lei et al., 2020; Depersin and Lecué, 2019). We will use some of its properties to design facility location mechanisms. See Lugosi and Mendelson (2019c) for a survey of the area in general. Finally, we mention that Hopkins et al. (2020) shows the LM median is equivalent of the filter-based estimators arising from robust statistics literature. However, the latter is not explicitly defined as a generalization of median to high dimensions.

2 PRELIMINARIES

For a set of reals $K \subseteq \mathbb{R}$, let med(K) be its (left) median. Let $N = \{1, 2, \cdots, n\}$ be a set of agents who are located in a metric space (M, d). In this paper, we focus on M being a (multi-dimensional) real space, and the metric d the standard Euclidean metric. We use $X = (x_1, x_2, \cdots, x_n) \in M^n$ to denote the true location profile of the n agents. A k-facility mechanism is a function $f : M^n \to M^k$ that maps the agents' profile to k facility locations. Given a set of facilities, each agent has a cost $c\left(f(X), x_i\right) = \min_{y \in f(X)} d\left(x_i, y\right)$. We define the *social cost* of a facility profile (with respect to X) as the sum of individual costs $\sum_{i \in N} c(f(X), x_i)$.

For any $X \in M^n$ and $K \subseteq N$, let x_K denote $\{x_i : i \in K\}$ and x_{-K} its complement. We drop the bracket when K is a singleton set.

2.1 Mechanism Design

We first recall the standard definition of (group) strategy-proofness and its variants.

Definition 2.1 (strategy-proofness)**.** *We say that a mechanism f is strategy-proof if for all $i \in N$, $x_i, x_i' \in M$ and $x_{-i} \in M^{n-1}$,*

$$c\left(f\left(x_i, x_{-i}\right), x_i\right) \leq c\left(f\left(x_i', x_{-i}\right), x_i\right).$$

In the Bayesian setting, the locations are drawn from a distribution $\mathcal{D}$ over M, and we consider the expected individual costs.

Definition 2.2 (Bayesian strategy-proofness)**.** *We say that a mechanism f is Bayesian strategy-proof if for all $i \in N$, $x_i, x_i' \in M$,*

$$\begin{aligned}&\mathop{\mathbb{E}}_{x_{-i}\sim\mathcal{D}^{n-1}} c\left(f\left(x_i, x_{-i}\right), x_i\right) \\ &\qquad \leq \mathop{\mathbb{E}}_{x_{-i}\sim\mathcal{D}^{n-1}} c\left(f\left(x_i', x_{-i}\right), x_i\right).\end{aligned}$$

The most general notion that we will deal with in the paper is approximate Bayesian group strategy-proofness. We also consider a variant, where incentive-compatibility holds with high probability rather than in expectation,

Definition 2.3 (Bayesian (ε, k)-group strategy-proofness)**.** *We say that a mechanism f is Bayesian (ε, k)-group strategy-proof if for all $K \subseteq N$ with $|K| = k$, $x_K, x_K' \in M^k$, $i \in K$*

$$\begin{aligned}&\mathop{\mathbb{E}}_{x_{-K}\sim\mathcal{D}^{n-k}} c\left(f\left(x_K, x_{-K}\right), x_i\right) \\ &\qquad \leq \mathop{\mathbb{E}}_{x_{-K}\sim\mathcal{D}^{n-k}} c\left(f\left(x_K', x_{-K}\right), x_i\right) + \varepsilon.\end{aligned}$$

Definition 2.4 (Bayesian (ε, δ, k)-group strategy-proofness)**.** *We say that a mechanism f is Bayesian (ε, δ, k)-group strategy-proof if, with probability at least $1 - \delta$ over $x_{-K} \sim \mathcal{D}^{n-k}$, for all $K \subseteq N$ with $|K| = k$, $x_K, x_K' \in M^k$, $i \in K$*

$$c\left(f\left(x_K, x_{-K}\right), x_i\right) \leq c\left(f\left(x_K', x_{-K}\right), x_i\right) + \varepsilon.$$

We also define Pareto-optimality and anonymity.

Definition 2.5 (Pareto-optimal)**.** *A mechanism f is Pareto optimal if for no $X \in M^n$, there is a set of facilities $F \in M^k$ with $c(F, X_i) \leq c(f(X), X_i)$ for all $i \in n$ such that at least one of the inequalities is strict.*

Definition 2.6 (anonymous)**.** *A mechanism is anonymous if $f(\sigma(X)) = f(X)$ for any profile $X \in M^n$ and any permutation $\sigma : M^n \to M^n$.*

2.2 High Dimensional Medians

Lugosi-Mendelson median. For any unit-norm $v \in \mathbb{R}^d$ and $X \subseteq \mathbb{R}^d$, let $X_v = \{\langle x_i, v\rangle\}$, and for $r \in \mathbb{R}_{\geq 0}$ define

$$\mathsf{slab}_r(v; X) = \{x \in \mathbb{R}^d : |\langle x, v\rangle - \mathrm{med}(X_v)| \leq r\},$$

to be the slab around the median under the 1d projection. Informally, the Lugosi-Mendelson (LM) median Lugosi and Mendelson (2019a) is any point close to the median of X under every 1d projection, up to an additive factor of r. (It can be seen as a relaxation of the classic Tukey median (Tukey, 1975), which is NP-hard to compute.) Note that it may not exist (for small r) and it may *not* be unique either.

Definition 2.7 (Lugosi-Mendelson median)**.** *For a point set $X \subseteq \mathbb{R}^d$, $y \in \mathbb{R}^d$ and radius $r \geq 0$, we say that y is a Lugosi-Mendelson r-median for X if*

$$y \in \bigcap_{v:\|v\|_2=1} \mathsf{slab}_r(v; X).$$

Observe that one can always take r to be the diameter of X to guarantee existence. On the other hand, to force uniqueness, we will also choose the minimum r so that the intersection of the slabs are non-empty. This yields the following definition.

Definition 2.8 (unique Lugosi-Mendelson median)**.** *For a point set $X \subseteq \mathbb{R}^d$, let $r \geq 0$ be the minimum value such that $\bigcap_{v:\|v\|_2=1} \mathsf{slab}_r(v; X)$ is non-empty. We say that $y \in \mathbb{R}^d$ is the unique Lugosi-Mendelson median if*

$$y = \bigcap_{v:\|v\|_2=1} \mathsf{slab}_r(v; X).$$

Geometric median. For $X \subseteq \mathbb{R}^d$, the geometric median is defined to be a point that minimizes the social cost. That is, we say that y is the geometric median if

$$y \in \arg\min \sum_{i\in N} d(x_i, y). \tag{2.1}$$

3 LUGOSI-MENDELSON MEDIAN IS NOT EXACTLY STRATEGY-PROOF

We now prove a negative result, showing that the unique Lugosi-Mendelson median (Definition 2.8) is not exactly strategy-proof (in the classic, non-Bayesian setting).

Theorem 3.1 (Non-trufulness of LM median)**.** *The unique Lugosi-Mendelson median is not a strategy-proof mechanism for Euclidean preferences in $\mathbb{R}^2$.*

We give a proof sketch by a picture and delay the technical details to Appendix A. A classic result by Peters et al. (1992) shows that in this setting a mechanism is Pareto-optimal, anonymous, and strategy-proof if and only if it is a coordinate-wise median. For 3 points on a plane, their coordinate-wise medians can be fully characterized algebraically. On the the other hand, we make a geometric observation that for 3 points forming a triangle in $\mathbb{R}^2$, the LM median is simply the incenter of the triangle. However, the incenter may not be any of the coordinate-wise medians, even for, say, an equilateral triangle. It is easy to observe that the incenter is Pareto-optimal and anonymous. This implies, by the result of Peters et al. (1992), that LM median cannot be strategyproof.

4 APPROXIMATE BAYESIAN STRATEGY-PROOFNESS FROM ROBUSTNESS

We now turn to approximate strategy-proof mechanisms in Bayesian setting, given that the unique Lugosi-Mendelon median does not achieve exact strategy-proofness (Section 3) and yet coordinate-wise median can be a bad approximation for the social cost (up to a factor of $\sqrt{d}$) in the classic setting (Meir, 2019). We provide two results, an approximate strategy-proof analysis for the Lugosi-Mendelson r-median (Definition 2.7) and a general reduction from strategy-proofness to robust estimation.

4.1 Strategy-Proofness of Lugosi-Mendelson Median

We show that for an appropriate choice of r, a Lugosi-Mendelson r-median is approximate group strategy-proof in a fairly general Bayesian setting. The theorem can be seen as an example of obtaining group strategy-proofness from robustness. Subsequently, we will describe such a general reduction in Section 4.2.

We focus on a Bayesian setting with n agents, whose location profile $x_1, x_2, \cdots, x_n$ are drawn i.i.d. from a distribution $\mathcal{D}$ over $\mathbb{R}^d$. We assume that $\mathcal{D}$ has bounded support (in particular, $\|X\|_2 \leq \beta$ for any X in the support of $\mathcal{D}$) and has unknown mean μ and a covariance Σ. In practical applications, we expect that real-world distributions to have bounded support. For example, agents may claim their desired physical location of a facility, or they may vote in a budget allocation setting, and these values are bounded.

Under the assumptions, we consider the single-facility mechanism described by Algorithm 1. The mechanism applies the Lugosi-Mendelson construction, which is based on median-of-means, in a straightforward way.

Lugosi and Mendelson (2019a) shows that with the choice of r given by (4.1), a LM r-median exists with probability at least $1 - e^{-\Omega(k)}$:

Theorem 4.1 (Lugosi and Mendelson (2019a)). *Let $X_1, X_2, \cdots, X_n$ be i.i.d. from a distribution over $\mathbb{R}^d$ with mean μ and (finite) covariance Σ. Let $\{Z_i\}_{i=1}^k$ be the bucket means of k disjoint buckets that partition $X_1, \cdots, X_n$. Then with probability at least $1 - e^{-\Omega(k)}$, an LM r-median exists, for r defined as in* (4.1).

Moreover, our argument for the strategy-proofness requires an algorithmic and robust version of this result.

Algorithm 1: Mechanism LM-MEDIAN for a Single Facility Location

```
Input: Agents' profile $X = (x_1, x_2, \ldots, x_n) \in \mathbb{R}^d$,
       parameter $k$
Output: A facility location $y \in \mathbb{R}^d$.

Partition the input locations $X$ into $k$ (artbitrary)
  disjoint buckets $B_1, B_2, \cdots, B_k$ of equal size.
Compute the bucket means $Z_i = \frac{1}{|B_i|} \sum_{j \in B_i} x_i$.
Return the Lugosi-Mendelson $r$-median $\widehat{\mu}$ of $\{Z_i\}_{i=1}^k$
  with
      $r = 120 \cdot \left( \sqrt{\frac{\operatorname{Tr} \Sigma}{n}} + \sqrt{\frac{\|\Sigma\| k}{n}} \right).$     (4.1)
If the algorithm fails to find a Lugosi-Mendelson
  $r$-median, return any point in the convex hull of $X$.
```

Lemma 4.2 (Depersin and Lecué (2019); Lei et al. (2020); see also Cherapanamjeri et al. (2019); Hopkins (2020)). *Let X be a partition of $\mathcal{I} \cup \mathcal{O}$, where (i) points in $\mathcal{O}$ are arbitrary such that $|\mathcal{O}| \leq k/300$ and (ii) the points in $\mathcal{I}$ are drawn i.i.d. from a distribution $\mathcal{D}$ (with mean μ and covariance Σ). Let r be defined in* (4.1). *Then there is a polynomial time algorithm that computes the Lugosi-Mendelson r-median $\widehat{\mu}$ of the k bucket means, with probability at least $1 - e^{-\Omega(k)}$ over the input data.*

Moreover, $\|\mu - \widehat{\mu}\| \leq 100r$.[2]

Our theorem shows that Algorithm 1 achieves Bayesian $(O(r), O(k))$-group strategy-proofness, for certain bound of β. The intuition is that the above lemma guarantees that the Lugosi-Mendelson median is a robust estimator against k outliers. It stays close to the true mean under perturbations. In particular, any group of k agents' misreporting their location would not improve their cost by $O(r)$, since the output would not change by more than a factor of $O(r)$.

Theorem 4.3 (approximate Bayesian strategy-proofness of LM median). *Let the location profile X be i.i.d. from a distribution with covariance Σ and (unknown) mean μ and support bounded by β. For any $\beta \leq e^k \cdot O(r)$, Algorithm 1 achieves Bayesian $(O(r), O(k))$-group strategy-proofness.*

Note that for constant $\operatorname{Tr} \Sigma = O(1)$ and $n \ll \operatorname{poly}(d)$, we can take $k = C \log d$ for a sufficiently large C, and our theorem gives a Bayesian $(O(r), O(\log d))$-group strategy-proof mechanism for any bounded distribution whose support is bounded by $\operatorname{poly}(d)$.

Proof. Let $\Theta \subseteq \mathbb{R}^d$ denote the support of $\mathcal{D}$ and $K \subseteq N$ be a subset of $\lfloor k/300 \rfloor$ agents. Let $X, X' \in M^n$ be two location profiles such that $x_{-K} = x'_{-K}$ are drawn from $\mathcal{D}^{n-k}$

[2] We remark that, as in prior work, we make no efforts in optimizing any of the constants here.

and $x_K, x'_K \in \Theta^k$ are both arbitrary. Then by construction, X, X' can be partitioned into the outlier set $\mathcal{O} = K$ and inlier set $\mathcal{I} = N \setminus K$. Let $\widehat{\mu}, \widehat{\mu}'$ be the output of LM-MEDIAN(X, k), LM-MEDIAN(X', k), respectively.

Now applying Lemma 4.2 and triangle inequality, we have that with probability at least $1 - e^{-\Omega(k)}$, $\|\widehat{\mu} - \widehat{\mu}'\| \leq 200r$. In this case, we obtain that for all $i \in K$

$$\begin{aligned} &c(\text{LM-MEDIAN}(X, k), x_i) \\ &\quad \leq c(\text{LM-MEDIAN}(X', k), x_i) + 200r. \end{aligned}$$

On the other hand, if this guarantee fails, which occurs with probability at most $e^{-\Omega(k)}$, we get that for all $i \in K$

$$\begin{aligned} &c(\text{LM-MEDIAN}(X, k), x_i) \\ &\quad \leq c(\text{LM-MEDIAN}(X', k), x_i) + \beta, \end{aligned}$$

since the algorithm would output some point in the convex hull of X. Finally, combining the facts above and taking expectations completes the proof. □

We observe that one can drop the boundedness assumption and obtain Bayesian (ε, δ, k)-group strategy-proofness instead:

Theorem 4.4 (approximate Bayesian strategy-proofness of LM median)**.** *Let the location profile X be i.i.d. from a distribution with covariance Σ and (unknown) mean μ. Then Algorithm 1 achieves Bayesian $(O(r), e^{-\Omega(k)}, O(k))$-group strategy-proofness.*

The proof of the theorem is analogous of that of Theorem 4.3.

4.2 Strategy-Proofness from Robust Estimation

We now generalize the analysis of Lugosi-Mendelson median to any robust estimator. Suppose we have an access to a robust algorithm $\mathcal{A}$ whose output does not change by a factor of $g(k)$ (in ℓ_2 norm), against k outliers. Then we can immediately obtain a Bayesian $(2g(k), k)$-group strategy-proof mechanism by feeding the location profile into $\mathcal{A}$ and return its output.

For simplicity, we focus on robust mean estimators. We remark that the result applies to general normed space beyond Euclidean, though in the literature most robust mean estimators attain ℓ_2 error guarantees.

Theorem 4.5 (robust estimator to Bayesian strategy-proof mechanism)**.** *Let X be a partition of $\mathcal{I} \cup \mathcal{O}$, where (i) points in $\mathcal{O}$ are arbitrary such that $|\mathcal{O}| \leq k$ and (ii) the points in $\mathcal{I}$ are drawn i.i.d. from a distribution $\mathcal{D}$ over M with unknown mean μ. Suppose there is an algorithm $\mathcal{A}$ that given input X, outputs $\widehat{\mu}$ such that $\|\widehat{\mu} - \mu\| \leq g(k)$. Then there is a Bayesian $(2g(k), k)$-group strategy-proof single-facility mechanism.*

The proof is similar to that of Theorem 4.3 and can be found in Appendix B.

Many estimation algorithm provided by the recent flurry of work on robust statistics are randomized procedures. They typically succeed with high probability. Hence, we also consider (ε, δ, k)-group strategy-proofness (Definition 2.4). The proof of the following theorem is analogous to that of Theorem 4.5, and we omit the details.

Theorem 4.6 (robust estimator to Bayesian strategy-proof mechanism, high probability)**.** *Let X be a partition of $\mathcal{I} \cup \mathcal{O}$, where (i) points in $\mathcal{O}$ are arbitrary such that $|\mathcal{O}| \leq k$ and (ii) the points in $\mathcal{I}$ are drawn i.i.d. from a distribution $\mathcal{D}$ over M with unknown mean μ. Suppose there is an algorithm $\mathcal{A}$ that given input X, outputs $\widehat{\mu}$ such that $\|\widehat{\mu} - \mu\| \leq g(k)$ with probability at least $1 - \delta$. Then there is a Bayesian $(2g(k), \delta, k)$-group strategy-proof single-facility mechanism.*

Let $\alpha = k/n$. We name a few computationally efficient robust estimation schemes that achieve the conditions of Theorem 4.6.

- Algorithms for robust estimation of multidimensional Gaussian $\mathcal{N}(\mu, \Sigma)$, with error $g(k) = O\left(\alpha\sqrt{\log 1/\alpha}\right)$ (Lai et al., 2016; Diakonikolas et al., 2019b, 2018b).
- Algorithms for robust estimation of bounded second moment distributions, with error $g(k) = O(\sqrt{\alpha})$ (Lai et al., 2016; Diakonikolas et al., 2019b).

All the algorithms run in time polynomial in $d, n, 1/\alpha$ and $\log(1/\delta)$. For the two canonical settings above, near-linear time algorithms also exist (Cheng et al., 2019; Dong et al., 2019; Hopkins et al., 2020).

5 SOCIAL COST ANALYSIS FOR LUGOSI-MENDELSON MEDIAN

In addition to incentive-compatibility, another key objective in designing facility location mechanism is to achieve small social cost. In this section, we continue to study the high dimensional Bayesian setting. We further assume that the input location profile is drawn from a centrally symmetric distribution.

Definition 5.1 (central symmetry)**.** *For a distribution $\mathcal{D}$ with mean μ and density $p(\cdot)$, we say that it is centrally symmetric (or symmetric for short) if $p(x - \mu) = p(x - \mu)$ for any x in the support of $\mathcal{D}$.*

Recall that the geometric median obtains the optimal social cost. We first observe that the geometric median converges to the mean μ, under the symmetry condition.

Lemma 5.1. *Let $X \sim \mathcal{D}^n$. Then for any symmetric distribution $\mathcal{D}$ over $\mathbb{R}^d$ with mean μ, the geometric median of X converges to μ as $n \to \infty$.*

We delay the proof of the lemma to Appendix C.

Now we observe that the LM r-median, if it exists, converges to the mean. It follows that the LM median obtains asymptotically optimal social cost.

Theorem 5.2 (social cost analysis for LM median)**.** *Suppose that the LM r median exists for the k bucket means. For fixed k and any Σ such that $\operatorname{Tr}\Sigma = o(n)$, the mechanism* LM-MEDIAN *(Algorithm 1) achieves asymptotically optimal social cost as $n \to \infty$.*

Proof. For fixed k and Σ such that $\operatorname{Tr}\Sigma = o(n)$, we note that $r \to 0$ as $n \to \infty$. Let $\widehat{\mu}$ be the LM r-median, the output of Algorithm 1. Then we have that $\|\widehat{\mu} - \mu\| \to 0$ as $n \to \infty$. Applying Lemma 5.1 shows that $\widehat{\mu}$ converges to the geometric median of X. □

Finally, we observe that our argument implies that the robust mean estimator yields an approximate strategy-proof mechanism, as long as the distribution has its geometric median converging to the mean. Symmetry is one condition that ensures this, but is not necessary. The general question of what distributions have this property is purely statistical in nature and remains an interesting direction for future work.

6 MULTIPLE FACILITIES FOR A SINGLE GAUSSIAN

We now turn to multi-facility mechanisms in Bayesian setting, focusing on a single multi-dimensional Gaussian.

We first show that having 2 facilities offers negligible improvement on the expected social cost, when the location profile are drawn from normal $\mathcal{N}(\mu, I_d)$. Therefore, it is an nearly optimal choice to set up facilities at the mean.

Lemma 6.1. *Let*

$$C = \min_{w,z} \mathop{\mathbb{E}}_{X \sim \mathcal{N}(\mu, I_d)} \min\{\|X - w\|, \|X - z\|\}$$

be the minimum expected social cost for any 2-facility mechanism. Then we have

$$C \geq (1 - o(1))\sqrt{d}.$$

Proof. Let

$$(w, z) \in \arg\min_{w,z} \mathop{\mathbb{E}}_{X \sim \mathcal{N}(\mu, I_d)} \min\{\|X - w\|, \|X - z\|\}$$

be the optimal 2-facility locations, in terms of the expected social cost. Assume without loss of generality that $w_i = z_i = 0$ for any $i > 2$. Furthermore, by rotation symmetry of the standard Gaussian distribution, we can also assume that $w_1 = -z_1, w_2 = z_2 = 0$. Now we can write

$$C = \min_{w_1} \mathop{\mathbb{E}}_{Y \sim \chi^2_{d-1}} \mathop{\mathbb{E}}_{X \sim \mathcal{N}(\mu_1, 1)} \min\left\{\sqrt{(X - w_1)^2 + Y}, \sqrt{(X - z_1)^2 + Y}\right\},$$

where χ^2_{d-1} denotes the Chi-squared distribution with $d-1$ degrees of freedom. By standard concentration of Chi-squared random variable (Laurent and Massart, 2000; Wainwright, 2019), we have for any $t > 0$

$$\Pr_{Y \sim \chi^2_{d-1}} \left(Y \geq d - 1 - 2\sqrt{(d-1)t}\right) \geq 1 - e^{-t}.$$

By taking $t = 100 \log d$, we conclude that $C \geq (1 - o(1))\sqrt{d}$. □

On the other hand, we note that taking the mean as a single-facility achieves a expected social cost of at most $\sqrt{d}$, since $\mathbb{E}_{X \sim \mathcal{N}(0, I_d)} \|X\|_2 \leq \sqrt{d}$. Given the robust estimation procedure for mean of a Gaussian, we have:

Theorem 6.2. *Let $d, n > 0$, $k \leq n/10$, and $\alpha = k/n$. There exists a Bayesian $(O(\alpha\sqrt{\log 1/\alpha}), \delta, k)$-group strategy-proof 2-facility mechanism for standard Gaussian location profile in $\mathbb{R}^d$ over n agents. Furthermore,*

- *the mechanism achieves $1 + o(1)$ approximation of the expecpted social cost;*
- *the mechanism runs in time polynomial in $n, d, 1/\alpha$ and $\log(1/\delta)$.*

Proof. To ensure strategy-proofness and runtime, we appeal to our general reduction Theorem 4.6 and any known estimation procedure for mean of Gaussian (e.g., Diakonikolas et al. (2018b)). Lemma 6.1 provides the social cost guarantee when setting both facilities at the mean. □

We also consider general m-facility mechanisms and give analogous results. See Appendix D.

7 MULTIPLE FACILITIES FOR A MIXTURE OF GAUSSIAN

We consider m-facility mechanisms for a mixture of m spherical Gaussians in $\mathbb{R}^d$. Our result is a natural one: setting the m facilities at the mean of each component of the mixture achieves nearly optimal expected social cost (for constant m). This allows us to apply known robust estimation procedures for Gaussian mixture and appeal to the general reduction (Theorem 4.6) to obtain approximate Bayesian group strategy-proof mechanisms.

For $\mu \in \mathbb{R}^{d \times m}$ and $w, \sigma \in \mathbb{R}^m$, let $\mathcal{N}(m, w, \sigma, \mu) = \sum_{i=1}^m w_i \mathcal{N}(\mu_i, \sigma_i^2 I_d)$ denote a mixture of m spherical

Gaussians in $\mathbb{R}^d$, where the ith component has weight w_i, mean μ_i and covariance $\sigma_i^2 I$.

Lemma 7.1. *Let m be a positive integer, $\mu \in \mathbb{R}^{d\times m}$ and $w, \sigma \in \mathbb{R}^m$ such that $w_i \geq 0$ and $\sum_i w_i = 1$. Let $C_m = \min_\alpha \mathbb{E}_{X\sim\mathcal{N}(m,w,\sigma,\mu)} \min_{i\in[m]} \|\alpha_i - X\|_2$, where the outer min is over a set of m facility locations $\{\alpha_i\}_{i=1}^m \in \mathbb{R}^d$. Then we have*

$$C_m \geq (1-o(1))\sqrt{d-m}\sum_{i=1}^m w_i\sigma_i. \qquad (7.1)$$

Proof. By using definition of the Gaussian mixture and pushing the outer minimum over the facilities inside, we can rewrite

$$\begin{aligned} C_m &= \min_\alpha \mathop{\mathbb{E}}_{X\sim\mathcal{N}(m,w,\sigma,\mu)} \min_{i\in[m]} \|\alpha_i - X\|_2 \\ &= \min_\alpha \sum_{i=1}^m w_i \mathop{\mathbb{E}}_{X\sim\mathcal{N}(\mu_i,\sigma_i^2)} \min_{i\in[m]} \|\alpha_i - X\|_2 \\ &\geq \sum_{i=1}^m w_i \min_\alpha \mathop{\mathbb{E}}_{X\sim\mathcal{N}(\mu_i,\sigma_i^2)} \min_{i\in[m]} \|\alpha_i - X\|_2 \\ &\geq (1-o(1))\sqrt{d-m}\sum_{i=1}^m w_i\sigma_i, \end{aligned}$$

where in the last step we used Lemma D.1. □

On the other hand, observe that setting m facilities at μ_i gets a social cost of $(1+o(1))\sqrt{d}\sum_{i=1}^m w_i\sigma_i$. Lemma 7.1 shows that this is essentially optimal for small m. It remains to achieve strategy-proofness.

For $m = O(1)$, Bakshi et al. (2020) gives a polynomial-time algorithm for robust estimation of Gaussian mixture, under a constant fraction of outliers. Applying their result (Theorem 1.6 of Bakshi et al. (2020)) and the general reduction Definition 2.4, we get:

Theorem 7.2 (multi-facility mechanism for Gaussian mixture)**.** *Let $d, n > 0$, $m = O(1)$, $k < n/10$, and $\alpha = k/n$. There exists a Bayesian $(O(poly(1/\alpha)), \delta, k)$-group strategy-proof m-facility mechanism for a mixture of m spherical Gaussian location profile in $\mathbb{R}^d$ over n agents. Furthermore,*

- *the mechanism achieves $1 + o(1)$ approximation of the expected social cost; and*
- *the mechanism runs in time polynomial in $n, d, 1/\alpha$ and $\log(1/\delta)$.*

8 CONCLUSION

Motivated by the classic results under worst-case assumptions, we initiate the study of strategy-proof facility mechanisms in Bayesian settings. Our work provides a wide class of approximately strategy-proof mechanisms. The key conceptual contribution is an explicit connection between strategy-proof facilities and algorithmic high-dimensional statistics. We show that robust estimators yield, in an almost black-box fashion, approximately strategy-proof mechanisms. We also provide social cost guarantees of mechanisms arising from this connection.

We conclude by pointing out three future directions.

- First, recall that we assume that the agents' true locations are identically distributed. We acknowledge that this may be somewhat restrictive for certain applications. An exciting future direction is to relax this condition and identify strategy-proof mechanisms when each agent's location is drawn from a distinct distribution.
- Second, our social cost analysis works under a natural though not fully general assumption of symmetry. We observe, however, that the proof holds if the distribution has its geometric median converging to the mean. It is an interesting theoretical question to understand what distributions have this property.
- Third, there has been a spate of work on connecting private and robust statistics (Alabi et al., 2022; Hopkins et al., 2022; Georgiev and Hopkins, 2022). At a high level, bridging strategy-proof estimation with these areas would a conceptually intriguing direction.

Acknowledgments

We thank the anonymous reviewers of AISTATS 2023 for helpful comments.

MZ was supported by the Army Research Office (ARO) under contract W911NF-17-1-0304 as part of the collaboration between US DOD, UK MOD and UK Engineering and Physical Research Council (EPSRC) under the Multidisciplinary University Research Initiative (MURI). FZ was supported by NSF award CCF-1951384.

References

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

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

Michal Feldman, Amos Fiat, and Iddan Golomb. On voting and facility location. In *Proceedings of the 2016 ACM Conference on Economics and Computation (EC)*, 2016.

Kim C Border and James S Jordan. Straightforward elections, unanimity and phantom voters. *The Review of Economic Studies*, 50(1):153–170, 1983.

Reshef Meir. Strategyproof facility location for three agents on a circle. In *International Symposium on Algorithmic Game Theory (SAGT)*, pages 18–33, 2019.

Sumit Goel and Wade Hann-Caruthers. Coordinate-wise median: Not bad, not bad, pretty good. *arXiv preprint arXiv:2007.00903*, 2020.

Michal Feldman and Yoav Wilf. Strategyproof facility location and the least squares objective. In *Proceedings of the fourteenth ACM conference on Electronic Commerce (EC)*, 2013.

Dimitris Fotakis and Christos Tzamos. On the power of deterministic mechanisms for facility location games. *ACM Transactions on Economics and Computation (TEAC)*, 2 (4):1–37, 2014.

Pinyan Lu, Xiaorui Sun, Yajun Wang, and Zeyuan Allen Zhu. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*, 2010.

Peter J Huber. Robust regression: asymptotics, conjectures and monte carlo. *The Annals of Statistics*, pages 799–821, 1973.

Gábor Lugosi and Shahar Mendelson. Sub-gaussian estimators of the mean of a random vector. *Annals of Statistics*, 47(2):783–794, 2019a.

Ilias Diakonikolas and Daniel M Kane. Recent advances in algorithmic high-dimensional robust statistics. *arXiv preprint arXiv:1911.05911*, 2019.

Zhixian Lei, Kyle Luh, Prayaag Venkat, and Fred Zhang. A fast spectral algorithm for mean estimation with subgaussian rates. In *Conference on Learning Theory (COLT)*, 2020.

Jules Depersin and Guillaume Lecué. Robust subgaussian estimation of a mean vector in nearly linear time. *arXiv:1906.03058*, 2019.

Ainesh Bakshi, Ilias Diakonikolas, He Jia, Daniel M Kane, Pravesh K Kothari, and Santosh S Vempala. Robustly learning mixtures of k arbitrary gaussians. *arXiv preprint arXiv:2012.02119*, 2020.

Salvador Barberà, Faruk Gul, and Ennio Stacchetti. Generalized median voter schemes and committees. *Journal of Economic Theory*, 61(2):262–289, 1993.

Noga Alon, Michal Feldman, Ariel D Procaccia, and Moshe Tennenholtz. Strategyproof approximation of the minimax on networks. *Mathematics of Operations Research*, 35(3):513–526, 2010.

Pinyan Lu, Yajun Wang, and Yuan Zhou. Tighter bounds for facility games. In *International Workshop on Internet and Network Economics (WINE)*, 2009.

Xin Sui, Craig Boutilier, and Tuomas Sandholm. Analysis and optimization of multi-dimensional percentile mechanisms. In *IJCAI*, pages 367–374, 2013.

Hau Chan, Aris Filos-Ratsikas, Bo Li, Minming Li, and Chenhao Wang. Mechanism design for facility location problem: A survey. In *The 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 2021.

Ioannis Caragiannis, Ariel Procaccia, and Nisarg Shah. Truthful univariate estimators. In *International Conference on Machine Learning (ICML)*, pages 127–135, 2016.

El-Mahdi El-Mhamdi, Sadegh Farhadkhani, Rachid Guerraoui, and Lê-Nguyên Hoang. On the strategyproofness of the geometric median. *arXiv preprint arXiv:2106.02394*, 2021.

Toby Walsh. Strategy proof mechanisms for facility location in euclidean and manhattan space. *arXiv preprint arXiv:2009.07983*, 2020.

John W. Tukey. A survey of sampling from contaminated distributions. *Contributions to probability and statistics*, 2:448–485, 1960.

Peter J Huber. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964.

Kevin A Lai, Anup B Rao, and Santosh Vempala. Agnostic estimation of mean and covariance. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, 2016.

Ilias Diakonikolas, Gautam Kamath, Daniel M Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robustness meets algorithms. *Communications of the ACM*, 64(5):107–115, 2021.

Sivaraman Balakrishnan, Simon S Du, Jerry Li, and Aarti Singh. Computationally efficient robust sparse estimation in high dimensions. In *Conference on Learning Theory (COLT)*, 2017.

Ilias Diakonikolas, Gautam Kamath, Daniel M Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Being robust (in high dimensions) can be practical. In *International Conference on Machine Learning (ICML)*, 2017.

Ilias Diakonikolas, Daniel M Kane, and Alistair Stewart. List-decodable robust mean estimation and learning mixtures of spherical gaussians. In *ACM Symposium on Theory of Computing (STOC)*, 2018a.

Ilias Diakonikolas, Gautam Kamath, Daniel Kane, Jerry Li, Jacob Steinhardt, and Alistair Stewart. Sever: A robust meta-algorithm for stochastic optimization. In *International Conference on Machine Learning (ICML)*, 2019a.

Yu Cheng, Ilias Diakonikolas, and Rong Ge. High-dimensional robust mean estimation in nearly-linear time. In *Proceedings of the thirtieth annual ACM-SIAM symposium on discrete algorithms (SODA)*, pages 2755–2771, 2019.

Samuel B Hopkins and Jerry Li. Mixture models, robustness, and sum of squares proofs. In *ACM SIGACT Symposium on Theory of Computing (STOC)*, 2018.

Yihe Dong, Samuel Hopkins, and Jerry Li. Quantum entropy scoring for fast robust mean estimation and improved outlier detection. *Advances in Neural Information Processing Systems*, 32, 2019.

Stanislav Minsker. Geometric median and robust estimation in banach spaces. *Bernoulli*, 21(4):2308–2335, 2015.

Luc Devroye, Matthieu Lerasle, Gabor Lugosi, and Roberto I Oliveira. Sub-gaussian mean estimators. *The Annals of Statistics*, 44(6):2695–2725, 2016.

Daniel Hsu and Sivan Sabato. Loss minimization and parameter estimation with heavy tails. *The Journal of Machine Learning Research*, 17(1):543–582, 2016.

Emilien Joly, Gábor Lugosi, and Roberto Imbuzeiro Oliveira. On the estimation of the mean of a random vector. *Electronic Journal of Statistics*, 11(1):440–451, 2017.

Gabor Lugosi and Shahar Mendelson. Robust multivariate mean estimation: the optimality of trimmed mean. *The Annals of Statistics*, 49(1):393–410, 2021.

Gábor Lugosi and Shahar Mendelson. Near-optimal mean estimators with respect to general norms. *Probability theory and related fields*, 175(3):957–973, 2019b.

Jasper CH Lee and Paul Valiant. Optimal sub-gaussian mean estimation in very high dimensions. In *13th Innovations in Theoretical Computer Science Conference (ITCS)*, 2022.

Samuel B Hopkins. Mean estimation with sub-gaussian rates in polynomial time. *The Annals of Statistics*, 48(2): 1193–1213, 2020.

Yeshwanth Cherapanamjeri, Nicolas Flammarion, and Peter L Bartlett. Fast mean estimation with sub-gaussian rates. In *Conference on Learning Theory (COLT)*, 2019.

Gábor Lugosi and Shahar Mendelson. Mean estimation and regression under heavy-tailed distributions: A survey. *Foundations of Computational Mathematics*, 19(5):1145–1190, 2019c.

Sam Hopkins, Jerry Li, and Fred Zhang. Robust and heavy-tailed mean estimation made simple, via regret minimization. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:11902–11912, 2020.

John W Tukey. Mathematics and the picturing of data. In *Proceedings of the International Congress of Mathematicians, Vancouver, 1975*, volume 2, pages 523–531, 1975.

Hans Peters, Hans van der Stel, and Ton Storcken. Pareto optimality, anonymity, and strategy-proofness in location problems. *International Journal of Game Theory*, 21(3): 221–235, 1992.

Ilias Diakonikolas, Gautam Kamath, Daniel Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robust estimators in high-dimensions without the computational intractability. *SIAM Journal on Computing*, 48(2):742–864, 2019b.

Ilias Diakonikolas, Gautam Kamath, Daniel M Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robustly learning a gaussian: Getting optimal error, efficiently. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2018b.

Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338, 2000.

Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.

Daniel Alabi, Pravesh K Kothari, Pranay Tankala, Prayaag Venkat, and Fred Zhang. Privately estimating a gaussian: Efficient, robust and optimal. *arXiv preprint arXiv:2212.08018*, 2022.

Samuel B Hopkins, Gautam Kamath, Mahbod Majid, and Shyam Narayanan. Robustness implies privacy in statistical estimation. *arXiv preprint arXiv:2212.05015*, 2022.

Kristian Georgiev and Samuel B Hopkins. Privacy induces robustness: Information-computation gaps and sparse mean estimation. *arXiv preprint arXiv:2211.00724*, 2022.

A Proof of Theorem 3.1

We start by recalling a classic result that fully characterizes the strategy-proof facility in this setting.

Lemma A.1 (Peters et al. (1992))**.** *In the Euclidean metric space $\mathbb{R}^2$ with odd number of agents, a mechanism is Pareto-optimal, anonymous, and strategy-proof if, and only if, it is a coordinate-wise median.*

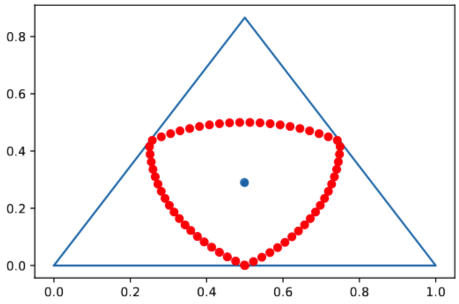

Figure 1: A proof by picture for Theorem 3.1. The red points are the coordinate-wise medians of the three vertices, and the central blue point is the incenter.

The following lemma characterizes the unique Lugosi-Mendelson median in case of three points on a plane.

Lemma A.2. *The unique Lugosi-Mendelson median of any three points forming a triangle is the incenter of the triangle.*

Proof. Recall that the incenter is the center point of the inscribed circle of the triangle. Now let $r \geq 0$ be the minimum value such that $\bigcap_{v:\|v\|_2=1} \mathsf{slab}_r(v; X)$ is non-empty, where X is the three vertex points of the triangle.

Suppose y is a unique LM median. We first consider the three one-dimensional projects along the direction of the three sides of the triangle. The median along each projection is simply one of the vertex points, and the slab centers around the corresponding side. Observe that the minimum value of r for the three slabs to meet is the radius of the inscribed circle of the triangle.

Further, we claim that this choice of r suffices for the slabs of all other projection directions to contain the incenter. Indeed, for each direction, we can associate it with a ray with the initial point being one vertex point. Observe that the (orthogonal) distance of the ray to the incenter is at most that of one of the sides to the incenter. Moreover, the slab is precisely centered around the ray, and this ensures that the slab contains the incenter. □

We show that the unique Lugosi-Mendelson median is not strategy-proof for two-dimensional Euclidean preferences. A simple proof by picture is given in Figure 1.

Theorem A.3 (Restatement of Theorem 3.1)**.** *The unique Lugosi-Mendelson median is not a strategy-proof mechanism for Euclidean preferences in $\mathbb{R}^2$.*

Proof. First, by its definition, the unique LM median is anonymous. Moreover, it is Pareto optimal, since it lies within the convex hull of the input points. Let $X \in (\mathbb{R}^2)^3$ be three points forming an equilateral triangle. By Lemma A.1, it suffices to show that it is not any coordinate-wise median of the three vertices of X, under any coordinate system of $\mathbb{R}^2$. Now for simplicity, let's take $X = \{(0,0),(1,0),(1/2,\sqrt{3}/2)\}$. For any angle $\theta \in [0,\pi)$, let

$$R(\theta) = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}$$

be the standard rotation matrix. Then any coordinate-wise median of X can be written as

$$R(\theta)^\top \cdot \Big(\mathrm{med}(0, \cos\theta, 1/2\cos\theta + \sqrt{3}\sin\theta/2), \mathrm{med}(0, -\sin\theta, -\sin\theta/2 + \sqrt{3}\cos\theta/2)\Big)^\top \tag{A.1}$$

by rotation of axes. On the other hand, the incenter of the triangle formed by X is $(1/2, 1/2\sqrt{3})$. It is a straightforward calculation to verify that this point is not a solution to (A.1). Hence, the unique LM median is not a coordinate-wise median mechanism. By Lemma A.1 and Lemma A.2, it is not strategy-proof. □

B Proof of Theorem 4.5

Proof of Theorem 4.5. Let $\Theta \subseteq \mathbb{R}^d$ denote the support of $\mathcal{D}$ and $K \subseteq N$ be a subset of k agents. Let $X, X' \in M^n$ be two location profiles such that $x_{-K} = x'_{-K}$ are drawn from $\mathcal{D}^{n-k}$ and $x_K, x'_K \in \Theta^k$ are both arbitrary. Then by construction, X, X' can be partitioned into the outlier set $\mathcal{O} = K$ and inlier set $\mathcal{I} = N \setminus K$. Let $\widehat{\mu}, \widehat{\mu}'$ be the output of $\mathcal{A}(X, k)$, $\mathcal{A}(X', k)$, respectively.

Now applying the guarantee of algorithm $\mathcal{A}$ and triangle inequality, we have that $\|\widehat{\mu} - \widehat{\mu}'\| \leq 2g(k)$. Therefore, for all $i \in K$

$$c(\mathcal{A}(X,k), x_i) \leq c(\mathcal{A}(X',k), x_i) + 2g(k). \tag{B.1}$$

This finishes the proof. □

C Proof of Lemma 5.1

Proof of Lemma 5.1. Given a set of n points in $\mathbb{R}^d$, the geometric median y minimizes the objective of

$$F(y) = \sum_{i=1}^{n} \|y - x_i\|_2. \tag{C.1}$$

First, observe by direct calculation that the gradient of $f(x) = \|x\|_2$ is simply $\nabla f(x) = x/\|x\|_2$. Applying chain rule, we get that

$$\nabla F(y) = \sum_{i=1}^{n} \frac{y - x_i}{\|y - x_i\|_2}. \tag{C.2}$$

Intuitively, this means that each data point exerts a unit force pulling the geometric median towards it. A point is the geometric median if all the unit forces cancel out; that is, $\nabla F(y) = 0$ so that the point is stable with respect to the pulls. Now suppose that x_i's are drawn i.i.d. from $\mathcal{D}$ with density $p(x)$. Then setting the gradient (C.2) to 0, we get that in population

$$\mathop{\mathbb{E}}_{X\sim\mathcal{D}}[\nabla F(y)] = \mathbb{E}\,\frac{y - X}{\|y - X\|_2} = \int_{\mathbb{R}^d} \frac{y - X}{\|y - X\|_2} p(x)\,\mathrm{d}x = 0.$$

Setting $y = \mu$ solves the equation by definition of symmetry. It follows that as $n \to \infty$, the empirical gradient (C.2) tends 0. This completes the proof. □

D General Multi-Facilities Mechanisms for a Single Gaussian

We now consider m-facility mechanism for any spherical Gaussian in d dimensions. For standard Gaussian, our result implies again that having m facilities offers no significant social cost improvement, unless m grows with the dimension.

Lemma D.1. *Let $C_m = \min \mathbb{E}_{X\sim\mathcal{N}(\mu,\sigma^2 I)} \min_{i\in[m]} \|w_i - X\|_2$, where the outer min is over a set of m facility locations $\{w_i\}_{i=1}^m \in \mathbb{R}^d$. Then we have $C_m \geq (1 - o(1))\sigma\sqrt{d-m}$.*

Proof. By considering the subspace spanned by the m facilities and using the rotation symmetry of spherical Gaussian, we have

$$C_m \geq \min \mathop{\mathbb{E}}_{Y\sim\sigma^2\chi^2_{d-m}} \mathop{\mathbb{E}}_{X\sim N(\mu_{[m]}, I_k)} \min_{i\in[m]} \sqrt{\sum_{i=1}^{m} (w_i - X_i)^2 + Y},$$

where we assume without loss of generality that the coordinates of w_i are 0, except at the first k indices, and $\mu_{[m]}$ denotes μ projected to the subspace spanned by the k facilities. By concentration of Chi-squared random variable (Laurent and Massart, 2000; Wainwright, 2019), we have for any $t > 0$

$$\Pr_{Y\sim\sigma^2\chi^2_{d-m}} \left(\frac{Y}{\sigma^2} \geq d - m - 2\sqrt{(d-m)t} \right) \geq 1 - e^{-t}.$$

By taking $t = 100 \log d$, we conclude that $C_k \geq (1 - o(1))\sigma\sqrt{d-m}$. □

The lemma shows that it is nearly optimal to simply set up all facilities at the mean. The proof of the following theorem is analogous to that of Theorem 6.2, so we omit it.

Theorem D.2 (multi-facility mechanism for Gaussian)**.** *Let $d, n > 0$, $k \leq n/10$, and $\alpha = k/n$. There exists a Bayesian $(O(\alpha\sqrt{\log 1/\alpha}), \delta, k)$-group strategy-proof m-facility mechanism for a spherical Gaussian location profile in $\mathbb{R}^d$ over n agents. Furthermore,*

- *the mechanism achieves $1 + o(1)$ approximation of the expected social cost for any constant m; and*
- *the mechanism runs in time polynomial in $n, d, 1/\alpha$ and $\log(1/\delta)$.*