

A Fast Spectral Algorithm for Mean Estimation with Sub-Gaussian Rates

Zhixian Lei

Harvard University

ZHIXIANLEI@G.HARVARD.EDU

Kyle Luh

Harvard University

KYLELUH@GMAIL.COM

Prayaag Venkat

Harvard University

PVENKAT@G.HARVARD.EDU

Fred Zhang

UC Berkeley

Z0@BERKELEY.EDU

Editors: Jacob Abernethy and Shivani Agarwal

Abstract

We study the algorithmic problem of estimating the mean of a heavy-tailed random vector in \mathbb{R}^d , given n i.i.d. samples. The goal is to design an efficient estimator that attains the optimal sub-gaussian error bound, only assuming that the random vector has bounded mean and covariance. Polynomial-time solutions to this problem are known but have high runtime due to their use of semi-definite programming (SDP). Moreover, conceptually, it remains open whether convex relaxation is truly necessary for this problem.

In this work, we show that it is possible to go beyond SDP and achieve better computational efficiency. In particular, we provide a *spectral* algorithm that achieves the optimal statistical performance and runs in time $\tilde{O}(n^2d)$, improving upon the previous fastest runtime $\tilde{O}(n^{3.5} + n^2d)$ by Cherapanamjeri *et al.* (COLT '19). Our algorithm is spectral in that it only requires (approximate) eigenvector computations, which can be implemented very efficiently by, for example, power iteration or the Lanczos method.

At the core of our algorithm is a novel connection between the furthest hyperplane problem introduced by Karnin *et al.* (COLT '12) and a structural lemma on heavy-tailed distributions by Lugosi and Mendelson (Ann. Stat. '19). This allows us to iteratively reduce the estimation error at a geometric rate using only the information derived from the top singular vector of the data matrix, leading to a significantly faster running time.

Keywords: High-dimensional statistics, mean estimation, robust statistics, spectral algorithm

1. Introduction

Estimating the mean of a multivariate distribution from samples is among the most fundamental statistical problems. Surprisingly, it was only recently that a line of works in the statistics literature culminated in an estimator achieving the optimal statistical error under minimal assumptions (Lugosi and Mendelson (2019b)). However, from an algorithmic point of view, computation of this estimator appears to be intractable. On the other hand, fast estimators, such as the empirical average, tend to achieve sub-optimal statistical performance. The following question remains open:

Can we provide simple, fast algorithm that computes a statistically optimal mean estimator in high dimensions, under minimal assumptions?

In this paper, we make progress towards this goal, under the classic setting where only finite mean and covariance are assumed. Formally, our problem is defined as follows. Given n i.i.d. copies $\mathbf{X}_1, \dots, \mathbf{X}_n$ of a random vector $\mathbf{X} \in \mathbb{R}^d$ with bounded mean $\boldsymbol{\mu} = \mathbb{E} \mathbf{X}$ and covariance $\boldsymbol{\Sigma} = \mathbb{E}(\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})^T$, compute an estimate $\hat{\boldsymbol{\mu}} = \hat{\boldsymbol{\mu}}(\mathbf{X}_1, \dots, \mathbf{X}_n)$ of the mean $\boldsymbol{\mu}$. Our goal is to show that for any failure probability $\delta \in (0, 1]$,

$$\Pr(\|\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}\| > r_\delta) \leq \delta,$$

for as small a radius r_δ as possible. Moreover, we would like to compute $\hat{\boldsymbol{\mu}}$ efficiently. The naïve estimator is simply the empirical mean $\bar{\boldsymbol{\mu}} = \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i$. It is well known that among all estimators, the empirical mean minimizes mean squared error. However, if we instead use the size of the deviations to quantify the quality of the estimator, the empirical mean is only optimal for sub-gaussian random variables (Catoni (2012)). When $\mathbf{X} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ we have with probability at least $1 - \delta$,

$$\|\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}\| \leq \sqrt{\frac{\text{Tr}(\boldsymbol{\Sigma})}{n}} + \sqrt{\frac{2\|\boldsymbol{\Sigma}\| \log(1/\delta)}{n}} \quad (1.1)$$

An estimator that achieves above is said to have *sub-gaussian performance* or *sub-gaussian rate*.

In practical settings, assuming that the samples obey a Gaussian distribution may be unrealistic. In an effort to design robust estimators, it is natural to study the mean estimation problem under very weak assumptions on the data. A recent line of works (Catoni (2012); Minsker (2015); Devroye et al. (2016); Joly et al. (2017); Lugosi and Mendelson (2019a,b); Lerasle et al. (2019)) study the mean estimation problem when the samples obey a heavy-tailed distribution.

For heavy-tailed distributions the performance of the empirical mean is abysmal. If we only assume that \mathbf{X} has finite mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$, then by Chebyshev’s inequality, the empirical mean only achieves error of order $\sqrt{\text{Tr}(\boldsymbol{\Sigma})/\delta n}$, which is worse than the sub-gaussian rate in two ways. First, its dependence on $\frac{1}{\delta}$ is exponentially worse. Second, the $\text{Tr}(\boldsymbol{\Sigma})$ term, which may grow with the dimension d , is multiplied the dimension-independent term $\sqrt{1/\delta n}$, whereas in the Gaussian case, the two are separate.

Median-of-means paradigm Surprisingly, recent work has shown that it is possible to improve on the performance of the empirical mean using the *median-of-means* approach. For $d = 1$, the following construction, originally due to Nemirovsky and Yudin (1983); Jerrum et al. (1986); Alon et al. (1999), achieves sub-gaussian performance:

- (i) First, bucket the data into $k = \lceil 10 \log(1/\delta) \rceil$ disjoint groups and compute their means Z_1, Z_2, \dots, Z_k .
- (ii) Then, output the median $\hat{\boldsymbol{\mu}}$ of Z_1, Z_2, \dots, Z_k .

A long line of work has followed this paradigm and generalized it to higher dimensions (Catoni (2012); Devroye et al. (2016); Joly et al. (2017); Lugosi and Mendelson (2019a,b)). The key challenge is to correctly define a notion of median for a collection of points in \mathbb{R}^d . Minsker (2015) considered $\hat{\boldsymbol{\mu}}_{GM}$ defined to be the *geometric median* of the bucket means Z_1, \dots, Z_k . For some constant c_{GM} , with probability at least $1 - \delta$, it satisfies

$$\|\hat{\boldsymbol{\mu}}_{GM} - \boldsymbol{\mu}\| \leq c_{GM} \sqrt{\frac{\text{Tr} \boldsymbol{\Sigma} \cdot \log(1/\delta)}{n}}. \quad (1.2)$$

This achieves the correct dependence on δ , but the dimension dependent and independent terms are still not separated. Following this work, [Lugosi and Mendelson \(2019a\)](#) described a tournament-based estimator, which finally achieved the optimal sub-gaussian radius. The idea is to consider *all 1-dimensional projections* of the bucket means and try to find an estimate that is close to the median of the means of all projections. This construction has been further simplified by [Hopkins \(2018\)](#). Formally, it was shown that the following estimator achieves the optimal, sub-gaussian error:

$$\hat{\boldsymbol{\mu}}_{LM} = \arg \min_{\boldsymbol{x} \in \mathbb{R}^d} \max_{\boldsymbol{u} \in \mathbb{S}^{d-1}} \left| \text{median} \{ \langle \boldsymbol{Z}_i, \boldsymbol{u} \rangle \}_{i=1}^k - \langle \boldsymbol{x}, \boldsymbol{u} \rangle \right|. \quad (1.3)$$

Clearly, searching over all directions in \mathbb{S}^{d-1} requires exponential time. The key question, therefore, is whether one can achieve both computational and statistical efficiency simultaneously.

Computational considerations A priori, it is unclear that the Lugosi-Mendelson estimator can be computed in polynomial time as a direct approach involves solving an intractable optimization problem. Moreover, the Lugosi-Mendelson analysis seems to suggest that estimation in the heavy-tailed model is conceptually harder than under (adversarial) corruptions. In the latter, each sample can be classified as either an inlier or an outlier. In the heavy-tailed setting, Lugosi-Mendelson shows that there is a majority of the bucket means that cluster around the true mean along any projection. However, a given sample may be an inlier by being close to the mean when projected onto one direction, but an outlier when projected onto another. In other words, the *set* of inliers may change from one direction to another.

Surprisingly, a recent line of works have established the polynomial-time computability of Lugosi-Mendelson estimator. [Hopkins \(2018\)](#) formulates $\hat{\boldsymbol{\mu}}_{LM}$ as the solution of a low-degree polynomial optimization problem and showed that using the Sum-of-Squares SDP hierarchy to relax this problem yields a sub-gaussian estimator. While the run-time of this algorithm is polynomial, it involves solving a large SDP. Soon after, [Cherapanamjeri et al. \(2019a\)](#) provided an iterative method in which each iteration involves solving a smaller, explicit SDP, leading to a run-time of $\tilde{O}(n^{3.5} + n^2d)$ ¹. Even more recently, a concurrent and independent work by [Lecué and Depersin \(2019\)](#) gave an estimator with sub-gaussian performance that can be computed in time $\tilde{O}(n^2d)$. The construction is inspired by a near-linear time algorithm for robust mean estimation under adversarial corruptions due to [Cheng et al. \(2019\)](#). The algorithm requires solving (covering) SDPs.

We note, however, that a common technique in these algorithms is SDP, which tends to be impractical for large sample sizes and in high dimensions. In contrast, our algorithm only requires approximate eigenvector computations. For a problem as fundamental as mean estimation, it is desirable to obtain simple and ideally practical solutions. A key conceptual message of our work is that SDP is indeed unnecessary and can be replaced by simple spectral techniques.

Our result In this work, we demonstrate for the first time that mean estimation with sub-gaussian rates can be achieved efficiently *without* SDP. The runtime of the algorithm matches the independent work of [Lecué and Depersin \(2019\)](#). In addition, our algorithm enjoys robustness against (additive) corruptions, where the number of adversarial points is a small fraction of k .

It is known that there exists an information-theoretic requirement for achieving such rates—that is, $\delta \geq 2^{-O(n)}$ ([Devroye et al. \(2016\)](#)). Under this assumption, we give an efficient spectral algorithm.

1. Throughout we use $\tilde{O}(\cdot)$ to hide polylogarithmic factors (in n , d and $\log(1/\delta)$).

Theorem 1 *Let $\delta \geq Ae^{-n}$ for a constant A and $k = \lceil 3600 \log(1/\delta) \rceil$. Given n points $\mathcal{G} \cup \mathcal{B}$, where \mathcal{G} are i.i.d. samples from a distribution over \mathbb{R}^d with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$ and \mathcal{B} a set of arbitrary points with $|\mathcal{B}| \leq k/200$, there is an efficient algorithm that outputs an estimate $\hat{\boldsymbol{\mu}} \in \mathbb{R}^d$ such that with probability at least $1 - \delta$,*

$$\|\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}\| \leq C \left(\sqrt{\frac{\text{Tr}(\boldsymbol{\Sigma})}{n}} + \sqrt{\frac{\|\boldsymbol{\Sigma}\| \log(1/\delta)}{n}} \right),$$

for a constant C . Furthermore, the algorithm runs in time $O(nd + k^2 d \text{polylog}(k, d))$.

The algorithm is iterative. Each iteration only requires an (approximate) eigenvector computation, which can be implemented in nearly linear time by power iteration or the Lanczos algorithm. We believe that our algorithm can be fairly practical.

Other related work Recently, Prasad et al. (2019) established a formal connection between the Huber contamination model and the heavy-tailed model we study in this paper. They leverage this connection to use an existing $\tilde{O}(nd^2)$ -time mean estimation algorithm of Diakonikolas et al. (2019) to design estimators for the heavy-tailed model. Under moment assumptions, their estimator achieves performance better than geometric median (1.2), yet worse than sub-gaussian.

In addition, algorithmic robust statistics has gained much attention in the theoretical computer science community in recent years. A large body of works have studied the mean estimation problem with *adversarially* corrupted samples, with the focus on providing efficient algorithms (Diakonikolas et al. (2019); Lai et al. (2016); Cheng et al. (2019); Dong et al. (2019)). For a more complete survey, see Diakonikolas and Kane (2019)

Going beyond mean estimation, there has been a recent spate of works on other statistical problems under heavy-tailed distributions. We refer the readers to Lugosi and Mendelson (2019c) for a survey.

Technical overview Our main algorithm builds upon the iterative approach of Cherapanamjeri et al. (2019a). For simplicity, assume there is no adversarial point. At a high level, for each iteration t , the algorithm will maintain a current guess \mathbf{x}_t of the true mean. To update, Cherapanamjeri et al. study the inner maximization of $\hat{\boldsymbol{\mu}}_{LM}$ (1.3) with $\mathbf{x} = \mathbf{x}_t$. They showed that under Lugosi-Mendelson structural condition, the problem is essentially equivalent of following program, which we call $\mathcal{M}(\mathbf{x}_t, \mathbf{Z})$:

$$\begin{aligned} & \max \quad \theta \\ & \text{subject to} \quad b_i \langle \mathbf{Z}_i - \mathbf{x}_t, \mathbf{u} \rangle \geq b_i \theta \text{ for } i = 1, \dots, k \\ & \quad \sum_{i=1}^k b_i \geq 0.95k \\ & \quad \mathbf{b} \in \{0, 1\}^k, \mathbf{u} \in \mathbb{S}^{d-1}. \end{aligned}$$

It can be shown that an optimal solution $\mathbf{u} \in \mathbb{S}^{d-1}$ will align with the unit vector in the direction of $\boldsymbol{\mu} - \mathbf{x}_t$, and θ approximates $\|\boldsymbol{\mu} - \mathbf{x}_t\|$. Hence, one can perform the update $\mathbf{x}_{t+1} \leftarrow \mathbf{x}_t + \gamma \theta \mathbf{u}$, for some appropriate constant γ , to geometrically decrease the distance of \mathbf{x}_t to $\boldsymbol{\mu}$.

In this work, we start by drawing a connection between the above program and the furthest hyperplane problem (FHP) of Karnin et al. (2012). This allows us to avoid the SDP approach in

Cherapanamjeri et al. (2019a). The problem can be formulated as the following:

$$\begin{aligned} \max \quad & \theta && \text{(FHP)} \\ \text{subject to} \quad & |\langle \mathbf{Z}_i - \mathbf{x}_t, \mathbf{u} \rangle| \geq \theta \text{ for } i = 1, \dots, k && (1.4) \\ & \mathbf{u} \in \mathbb{S}^{d-1}. \end{aligned}$$

In the original formulation due to Karnin *et al.*, the goal is to find a maximum margin linear classifier for a collection of points, where the margin is *two-sided*. Notice that any feasible solution to $\mathcal{M}(\mathbf{x}_t, \mathbf{Z})$ satisfies at least $0.95k$ constraints of FHP as well. For an arbitrary dataset, the two-sided margin requirement indeed provides a relaxation. One technical observation of this work is that it is *not* a significant one, for the random data we care about—if a major fraction of the constraint (1.4) are satisfied, then most constraints of $\mathcal{M}(\mathbf{x}_t, \mathbf{Z})$ are satisfied as well.

Unfortunately, the algorithm of Karnin *et al.* cannot directly apply, as it only works under a strong promise that there exists a feasible solution that satisfies *all* of the constraints (1.4). In our setting, there may not be such a feasible solution; we can only guarantee that there exists a unit vector (namely, the one in the direction of $\boldsymbol{\mu} - \mathbf{x}_t$) that satisfies most of constraints with large margin.

Our main contribution is to provide an algorithm that works even under this weak promise. We now briefly review the algorithm of Karnin *et al.*, show why it fails for our purpose, and explain how we address the issues that arise. Suppose that there exists a unit vector \mathbf{u}^* and θ^* which are feasible for the FHP problem. Then, averaging the constraints tells us that

$$\frac{1}{k} \sum_{i=1}^k \langle \mathbf{Z}_i, \mathbf{u}^* \rangle^2 \geq \theta^{*2}.$$

Hence, if we define \mathbf{u} to be the top right singular vector of the matrix \mathbf{A} whose rows are \mathbf{Z}_i , then

$$\|\mathbf{A}\mathbf{u}\|_2^2 = \sum_{i=1}^k \langle \mathbf{Z}_i, \mathbf{u} \rangle^2 \geq \sum_{i=1}^k \langle \mathbf{Z}_i, \mathbf{u}^* \rangle^2 \geq k\theta^{*2},$$

so \mathbf{u} satisfies the constraints in (FHP) *on average*. However, the distribution of the quantities $\langle \mathbf{Z}_i, \mathbf{u} \rangle^2$ may be extremely skewed, so that \mathbf{u} only satisfies a few of the constraints with large margin. If this happens, however, we can downweight those constraints which are satisfied by \mathbf{u} with large slack to encourage it to satisfy more constraints. This reweighting procedure is repeated several times, and at the end we use a simple rounding scheme to yield a single output vector with the desired properties from all the repetitions. In particular, this weighting scheme is essentially the same as the classic *multiplicative weights update* (MWU) method (Arora et al. (2012)) for regret minimization, as we show in Appendix I.

If we are only guaranteed that \mathbf{u}^* satisfies most, but not all, of the constraints, then the inequality $\sum_{i=1}^k \langle \mathbf{Z}_i, \mathbf{u}^* \rangle^2 \geq k\theta^{*2}$ may no longer hold when the points \mathbf{Z}_i get re-weighted and the algorithm of Karnin *et al.* cannot be guaranteed to converge. To illustrate this point, consider the following extreme case. Suppose that after the first iteration, the algorithm finds the vector \mathbf{u}^* as the top right singular vector of \mathbf{A} . In the re-weighting procedure, the constraints i for which $\langle \mathbf{Z}_i, \mathbf{u}^* \rangle^2 \geq \theta^{*2}$ may be down-weighted significantly, whereas the remaining constraints may be unaffected. This may result in most of the weight being concentrated on the constraints i where $\langle \mathbf{Z}_i, \mathbf{u}^* \rangle^2 \ll \theta^{*2}$. In the second iteration, we have no guarantee of the behavior of the top singular vector of the re-weighted matrix because all the weight is concentrated on a small set consisting of these “bad” constraints.

To address this scenario, our key technical idea is to project the weights onto the set of *smooth distributions* after each update. Informally, the notion of smooth distribution enforces that no point can take too much probability mass—say, more than $4/k$. This prevents the weights from ever being concentrated on too small a subset and allows us to guarantee that $\sum_{i=1}^k \langle \mathbf{Z}_i, \mathbf{u}^* \rangle^2 \geq k\theta^{*2}$ still holds approximately. Moreover, the appropriate notion of projection here is that of a Bregman projection. Leveraging our earlier MWU interpretation of the algorithm (Section I), we apply a classic regret bound for MWU under Bregman projection (Arora et al. (2012)), and this yields the same guarantee of the original algorithm. Finally, we remark that the projection can be computed quickly. Combining all these ideas together, we manage to bypass the barrier of having bad points, under the much weaker assumption on \mathbf{u}^* .

Organization The remainder of this article is organized as follows. In Section 2, we set up the notations and specify assumptions on the data. In Section 3, we explain the high level approach based on an iterative descent procedure from Cherapanamjeri et al. (2019a). The procedure requires us to approximately maximize a (non-convex) objective, and we discuss its properties in Section 4. Section 5 contains the main technical innovations of this work, where we design and analyze a faster algorithm for the aforementioned optimization problem.

2. Preliminaries and Assumptions

In the following, we use $r_\delta = \sqrt{\text{Tr}(\boldsymbol{\Sigma})/n} + \sqrt{\|\boldsymbol{\Sigma}\| \log(1/\delta)/n}$ to denote the optimal, sub-gaussian error rate and $k = \lceil 3200 \log(8/\delta) \rceil$. The input data $\{\mathbf{X}_i\}_{i=1}^n$ consist of \mathcal{G} , a set of i.i.d. points, and \mathcal{B} , a set of adversarial points, with $|\mathcal{B}| \leq k/200$. Our algorithm preprocesses the data \mathbf{X}_i into the bucket means $\mathbf{Z}_1, \mathbf{Z}_2, \dots, \mathbf{Z}_{2k} \in \mathbb{R}^d$.² Let \mathcal{B}_j be the set of \mathbf{X}_i in bucket j . We say that a bucket mean \mathbf{Z}_j is *contaminated* if \mathcal{B}_j contains an adversarial $X_i \in \mathcal{B}$ and *uncontaminated* otherwise. Note that the number of contaminated bucket means is at most $k/200$.

Our argument is built on the Lugosi-Mendelson condition. It states that under any one-dimensional projection, most of the (uncontaminated) bucket means are close to the true mean, by an additive factor of $O(r_\delta)$. Throughout, we pessimistically assume all contaminated bucket means do not satisfy this property (under any projection) and condition on the following event.

Assumption 1 (Lugosi-Mendelson condition) *Under the setting above, for all unit \mathbf{v} , we have*

$$|\{i : \langle \mathbf{v}, \mathbf{Z}_i \rangle - \langle \mathbf{v}, \boldsymbol{\mu} \rangle \geq 600r_\delta\}| \leq 0.05k.$$

Lemma 2 (Lugosi and Mendelson (2019b)) *Assumption 1 holds with probability at least $1 - \delta/8$.*

3. Descent Procedure

At a high level, our algorithm builds upon the iterative descent paradigm of Cherapanamjeri et al. (2019a). It maintains a sequence of estimates and updates via distance and gradient estimate.

Definition 3 (distance estimate) *We say that d_t is a distance estimate (with respect to \mathbf{x}_t) if*

- (i) *when $\|\boldsymbol{\mu} - \mathbf{x}_t\| \leq 14000r_\delta$, we have $d_t \leq 28000r_\delta$; and*

2. We assume δ is such that $k \leq n/2$; as we mentioned in the introduction, this is information-theoretically necessary, up to a constant (Devroye et al. (2016)).

(ii) when $\|\boldsymbol{\mu} - \mathbf{x}_t\| > 14000r_\delta$, we have

$$\frac{1}{21}\|\boldsymbol{\mu} - \mathbf{x}_t\| \leq d_t \leq 2\|\boldsymbol{\mu} - \mathbf{x}_t\| \quad (3.1)$$

Definition 4 (gradient estimate) We say that \mathbf{g}_t is a gradient estimate (with respect to \mathbf{x}_t) if

$$\left\langle \mathbf{g}_t, \frac{\boldsymbol{\mu} - \mathbf{x}_t}{\|\boldsymbol{\mu} - \mathbf{x}_t\|} \right\rangle \geq \frac{1}{200} \quad (3.2)$$

whenever $\|\boldsymbol{\mu} - \mathbf{x}_t\| > 14000r_\delta$.

1. **Input:** Buckets means $\mathbf{Z}_1, \dots, \mathbf{Z}_k \in \mathbb{R}^d$, initial estimate \mathbf{x}_0 , iteration count T_{des} , and step size η .
2. For $t = 1, \dots, T_{\text{des}}$:
 - (a) Compute $d_t = \text{DISTEST}(\mathbf{Z}', \mathbf{x}_t)$.
 - (b) Compute $\mathbf{g}_t = \text{GRADEST}(\mathbf{Z}', \mathbf{x}_t)$.
 - (c) Update $\mathbf{x}_{t+1} = \mathbf{x}_t + \eta d_t \mathbf{g}_t$.
3. **Output:** \mathbf{x}_{t^*} , where $t^* = \arg \min_t d_t$.

Algorithm 3.1: Main algorithm—DESCENT

Suppose we initialize the estimate with coordinate-wise median-of-means which achieves an error rate $\sqrt{\|\boldsymbol{\Sigma}\|kd/n}$ (Lemma 16). The following lemma states that if DISTEST and GRADEST provide distance and gradient estimate, then the algorithm DESCENT succeeds in logarithmic iterations. The lemma has essentially appeared in Cherapanamjeri et al. (2019a), albeit with a general initialization and a different set of constants. We give a proof in Appendix C for completeness.

Lemma 5 (convergence rate; see Cherapanamjeri et al. (2019a)) Assume that for all $t \leq T_{\text{des}}$, d_t is a distance estimate and \mathbf{g}_t is a gradient estimate (with respect to \mathbf{x}_t). Suppose $\|\boldsymbol{\mu} - \mathbf{x}_0\| \leq O\left(\sqrt{\|\boldsymbol{\Sigma}\|kd/n}\right)$. Then the output of Algorithm 3.1 DESCENT instantiated with $T_{\text{des}} = \Theta(\log d)$ and $\eta = 1/8000$ satisfies $\|\mathbf{x}_{t^*} - \boldsymbol{\mu}\| \leq O(r_\delta)$.

4. Inner Maximization and its Two-Sided Relaxation

Cherapanamjeri et al. (2019a) obtains gradient and distance estimates by solving the *inner maximization* problem of the Lugosi-Mendelson estimator, denoted by $\mathcal{M}(\mathbf{x}, \mathbf{Z})$:

$$\begin{aligned} & \max \quad \theta \\ & \text{subject to} \quad b_i \langle \mathbf{Z}_i - \mathbf{x}, \mathbf{w} \rangle \geq b_i \theta \text{ for } i = 1, \dots, k \\ & \quad \quad \quad \sum_{i=1}^k b_i \geq 0.95k \end{aligned}$$

$$\mathbf{b} \in \{0, 1\}^k, \mathbf{w} \in \mathbb{S}^{d-1}.$$

We also denote its feasibility version for a fixed θ by $\mathcal{M}(\theta, \mathbf{x}, \mathbf{Z})$. Note that the constraint of $\mathcal{M}(\mathbf{x}, \mathbf{Z})$ dictates that 0.95 fraction of the data must lie on *one* side of the hyperplane \mathbf{w} with a margin θ . As discussed in the introduction, we relax it by allowing a two-sided margin: $\mathcal{M}_2(\mathbf{x}, \mathbf{Z})$.

$$\begin{aligned} & \max \quad \theta \\ & \text{subject to} \quad b_i |\langle \mathbf{Z}_i - \mathbf{x}, \mathbf{w} \rangle| \geq b_i \theta \text{ for } i = 1, \dots, k \\ & \quad \quad \quad \sum_{i=1}^k b_i \geq 0.95k \\ & \quad \quad \quad \mathbf{b} \in \{0, 1\}^k, \mathbf{w} \in \mathbb{S}^{D-1}. \end{aligned}$$

One technical observation here is that under the Lugosi-Mendelson condition, this relaxation is insignificant. Indeed, approximately solving the problem suffices for gradient and distance estimates.

Lemma 6 *Let θ^* be the optimal value of $\mathcal{M}(\mathbf{x}, \mathbf{Z})$ and \mathbf{w} be a unit vector such that for at least $k/8$ of the \mathbf{Z}_i , we have $|\langle \mathbf{w}, \mathbf{Z}_i - \mathbf{x} \rangle| \geq \theta$, where $\theta = 0.1\theta^*$. We have that (i) θ is a distance estimate and (ii) either \mathbf{w} or $-\mathbf{w}$ is a gradient estimate.*

We give a proof in Appendix D. The intuition here is simple. If $\|\mathbf{x} - \boldsymbol{\mu}\| \ll r_\delta$, then the Lugosi-Mendelson condition ensures at most $0.05k$ points are far from \mathbf{x} by $O(r_\delta)$ (under any projection), so $\theta = O(r_\delta)$. On the other hand, if $\|\mathbf{x} - \boldsymbol{\mu}\| \gg r_\delta$, along the gradient direction, a majority of data lie only on *one* side of the hyperplane, the side that contains the true mean, so the two-sided constraint does not make a difference.

5. Approximating the Inner Maximization

We now give an algorithm that efficiently computes a approximate solution to the relaxation of the inner maximization. This will provide gradient and distance estimates for each iteration of the main DESCENT algorithm (Algorithm 3.1).

The run-time of the algorithm is proportional $1/\theta^2$. For technical reasons, we need to ensure that $\|\mathbf{Z}_i - \mathbf{x}\| \leq 1$ for all i . However, naively scaling all the data would decrease θ , thereby blowing up the running time. Hence, as a preprocessing step, we prune out a small fraction of points $\mathbf{Z}_i - \mathbf{x}$ with large norm before scaling.

5.1. Pruning and scaling

The preprocessing step (Algorithm 5.1) will be executed *only once* in the algorithm. After the pruning step and an appropriate scaling, we may assume the following structures on the data.

Assumption 2 *Given a current estimate \mathbf{x} , the pruned dataset $\mathbf{Z} \in \mathbb{R}^{k' \times d}$ of size k' , let $\mathbf{Z}'_i = \frac{1}{B}(\mathbf{Z}_i - \mathbf{x})$, where $B = \max_i \|\mathbf{Z}'_i - \mathbf{x}\|$. We assume (i) $\|\mathbf{Z}'_i\| \leq 1$; (ii) $k' \geq 0.9k$; and (iii) there exists $\theta = \Omega(1/\sqrt{d})$ and a unit vector \mathbf{w} such that for at least $0.8k$ points $|\langle \mathbf{Z}'_i, \mathbf{w} \rangle| \geq \theta$.*

We analyze the subroutine and prove the following lemma in Appendix E.

1. **Input:** Dataset $\mathbf{Z}_1, \mathbf{Z}_2, \dots, \mathbf{Z}_k \in \mathbb{R}^d$, initial estimate \mathbf{x}_0
2. Compute the distances $d_i = \|\mathbf{Z}_i - \mathbf{x}_0\|$.
3. Sort the points by d_i in decreasing order.
4. Remove the top 1/10 fraction of them. Let $\mathbf{Z}_1, \dots, \mathbf{Z}_{k'}$ be the remaining data.
5. **Output:** $\mathbf{Z}_1, \dots, \mathbf{Z}_{k'}$

Algorithm 5.1: PRUNE

Lemma 7 *With probability at least $1 - \delta/8$, [Assumption 2](#) holds for any \mathbf{x} such that $\|\mathbf{x} - \boldsymbol{\mu}\| \leq O\left(\sqrt{\|\boldsymbol{\Sigma}\|kd/n}\right)$ and $\|\mathbf{x} - \boldsymbol{\mu}\| \geq \Omega(r_\delta)$.*

In the remainder of the section, given a current estimate \mathbf{x} , we work with the pruned and scaled data, centered at \mathbf{x} , which we call $\mathbf{Z}' \in \mathbb{R}^{k' \times d}$.

We will aim at proving the following lemma, under [Assumption 2](#).

Lemma 8 (key lemma) *Assume [Assumption 2](#). Let $\delta \in (0, 1)$ and $T_{des} = \Theta(\log d)$. Suppose that there exists $\mathbf{w}^* \in \mathbb{S}^{d-1}$ which satisfies $|\langle \mathbf{Z}'_i, \mathbf{w}^* \rangle| \geq \theta^*$ for $0.8k$ points in $\{\mathbf{Z}'_i\}$. Then there is an algorithm APPROXBREGMAN which, with probability at least $1 - \delta/4T_{des}$, outputs $\mathbf{w} \in \mathbb{S}^{d-1}$ such that for at least 0.45 fraction of the points \mathbf{Z}'_i , it holds that $|\langle \mathbf{Z}'_i, \mathbf{w} \rangle| \geq 0.1\theta^*$.*

Further, APPROXBREGMAN runs in time $\tilde{O}(k^2d)$.

5.2. Approximation via Bregman Projection

In this section, we give the main algorithm for approximating \mathcal{M}_2 . Suppose (by binary search) that we know the optimal margin θ in [Lemma 8](#). The goal is to find a unit vector \mathbf{w} such that a constant fraction of \mathbf{Z}'_i has margin $|\langle \mathbf{Z}'_i, \mathbf{w} \rangle| \geq \theta$. The intuition is that we can start by computing the top singular vector of \mathbf{Z}' . Then the margin would be large on average: certain points may overly satisfy the margin demand while other may under-satisfy it. Hence, we would downweight those data points that achieve large margin and compute the top singular vector of the weighted matrix again.

However, it may stop making progress if it puts too much weight on the points that do not satisfy the margin bound. In this section, we show how to prevent this scenario from occurring. The key idea is that at every iteration, we “smooth” the weight vector τ_t so that we can guarantee progress is being made. We will formulate our algorithm in the well-studied regret-minimization framework and appeal to existing machinery [Arora et al. \(2012\)](#) to derive the desired approximation guarantees.

First, we define what type of distribution we would like τ_t to be.

Definition 9 (Smooth distributions) *The set of smooth distributions on $[k']$ is defined to be*

$$\mathcal{K} = \left\{ p \in \Delta_{k'} : p(i) \leq \frac{4}{k'} \text{ for every } i \in [k'] \right\},$$

where $\Delta_{k'}$ is the set of probability distributions on $[k']$,

$$\Delta_{k'} = \left\{ p : [k'] \rightarrow [0, 1] : \sum_{i \in [k']} p(i) = 1 \right\}.$$

In the course of the algorithm, after updating τ_t as in the previous section, it may no longer be smooth. Hence, we will replace it by the closest smooth weight vector (under KL divergence). The following fact confirms that finding this closest smooth weight vector can be done quickly.

Fact 10 (Barak et al. (2009)) For any $p \in \Delta_k$ with support size at least $k'/2$, computing

$$\Pi_{\mathcal{K}}(p) = \arg \min_{q \in \mathcal{K}} KL(p||q)$$

can be done in $\tilde{O}(k')$ time, where $KL(\cdot||\cdot)$ denotes the Kullback-Leibler divergence.

Remark 11 In our algorithm, we will only compute Bregman projections of distributions of support size at least $k'/2$. This is because neither our reweighting procedure nor the actual projection algorithm of Barak et al. (2009) sets any coordinates to 0 and the initial weight is uniform.

1. **Input:** Buckets means $\mathbf{Z}' \in \mathbb{R}^{k' \times d}$, margin θ , iteration count $T \in \mathbb{N}$
2. Initialize weights: $\tau_1 = \frac{1}{k'}(1, \dots, 1) \in \mathbb{R}^{k'}$.
3. For $t = 1, \dots, T$, repeat:
 - (a) Let \mathbf{A}_t be the $k' \times d$ matrix whose i th row is $\sqrt{\tau_t(i)}(\mathbf{Z}'_i)$ and \mathbf{w}_t be its approximate top right singular vector .
 - (b) Set $\sigma_t(i) = |\langle \mathbf{Z}'_i, \mathbf{w}_t \rangle|$.
 - (c) Reweight: If $\|\mathbf{A}_t \mathbf{w}_t\|_2^2 \geq \frac{\theta^2}{10}$, then $\tau_{t+1}(i) = \tau_t(i) (1 - \sigma_t(i)^2/2)$ for $i \in [k']$. Otherwise, do not change the weights.
 - (d) Normalize: Let $Z = \sum_{i \in [k']} \tau_{t+1}(i)$ and redefine $\tau_{t+1} \leftarrow \frac{1}{Z} \tau_{t+1}$.
 - (e) Compute the Bregman projection: $\tau_{t+1} \leftarrow \Pi_{\mathcal{K}}(\tau_{t+1})$.
4. **Output:** $\mathbf{w} \leftarrow \text{ROUND}(\mathbf{Z}', \{\mathbf{w}_j\}_{j=1}^T, \theta)$ (or report FAIL if ROUND fails).

Algorithm 5.2: Approximate inner maximization via Bregman projection—APPROXBREGMAN

Since Algorithm 5.2 is the MWU method with Bregman projections onto the set \mathcal{K} , we will apply the following regret guarantee.³

3. To be more precise, the iterations t in which $\|\mathbf{A}_t \mathbf{w}_t\|_2^2 \geq \frac{\theta^2}{10}$ behave according to the MWU method. Whenever $\|\mathbf{A}_t \mathbf{w}_t\|_2^2 < \frac{\theta^2}{10}$, the algorithm does not update the weights, which has no effect on the other iterations.

Theorem 12 (Theorem 2.4 of Arora et al. (2012)) *Suppose that for $\sigma_t^2(i) \in [0, 1]$ for all $i \in [k']$ and $t \in [T]$. Then after T iterations of Algorithm 5.2, for any $p \in \mathcal{K}$, it holds that:*

$$\sum_{t=1}^T \langle \tau_t, \sigma_t^2 \rangle \leq \frac{3}{2} \sum_{t=1}^T \langle p, \sigma_t^2 \rangle + 2KL(p || \tau_1).$$

Finally, we comment that we cannot naïvely apply the power method for the singular vector computation. The power method has failure probability of $1/10$, whereas our algorithm should fail with probability at most $\delta = O(\exp(-k))$ that is exponentially low. However, we note that the algorithm computes the top singular vectors of a sequence of matrices $\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_T$. Observe that as long as $T = \Omega(\log(1/\delta)) = \Omega(k)$, with probability at least $1 - \delta/8$, the power method will succeed for $0.9T$ of the matrices. We will show that this many successes suffice to guarantee correctness of our algorithm.

We first prove the following lemma, a requirement for the rounding algorithm to succeed.

Lemma 13 (regret analysis) *After $T = O\left(\max\left(\frac{\log k'}{\theta^2}, \log(T_{des}/\delta)\right)\right)$ iterations of Algorithm 5.2, for all but a $1/4$ fraction of $i \in [k']$:*

$$\sum_{t=1}^T \langle \mathbf{Z}'_i, \mathbf{w}_t \rangle^2 \geq 100 \log k'.$$

Proof Let $S = \{i \in [k'] : |\langle \mathbf{Z}'_i, \mathbf{w}^* \rangle| \geq \theta\}$ be the set of constraints satisfied by the unit vector \mathbf{w}^* whose existence is guaranteed in the hypothesis of Lemma 8. By assumption, we have that $|S| \geq 0.8k'$. We simply calculate each of the terms in Theorem 12.

First, let $\mathcal{I} = \{t \in [T] : \mathbf{w}_t \text{ is a } 1/2\text{-approximate top singular vector of } \mathbf{A}_t\}$. Then we have for any $t \in \mathcal{I}$:

$$\begin{aligned} \langle \tau_t, \sigma_t^2 \rangle &= \sum_{i=1}^{k'} \tau_t(i) \langle \mathbf{Z}'_i, \mathbf{w}_t \rangle^2 && \text{(by definition)} \\ &\geq \frac{1}{2} \sum_{i=1}^{k'} \tau_t(i) \langle \mathbf{Z}'_i, \mathbf{w}^* \rangle^2 && \text{(because } \mathbf{w}_t \text{ is an approximate top eigenvector)} \\ &\geq \frac{1}{2} \sum_{i \in S} \tau_t(i) \langle \mathbf{Z}'_i, \mathbf{w}^* \rangle^2 \\ &\geq \frac{1}{2} \sum_{i \in S} \tau_t(i) \theta^2 && \text{(by definition of } S) \\ &\geq \frac{1}{2} \cdot \frac{1}{5} \theta^2 = \frac{\theta^2}{10} && \text{(because } |S| \geq 0.8k' \text{ and } \tau_t \in \mathcal{K}). \end{aligned}$$

Summing this inequality over $t \in [T]$, we have that

$$\sum_{t=1}^T \langle \tau_t, \sigma_t^2 \rangle \geq \sum_{t \in \mathcal{I}} \langle \tau_t, \sigma_t^2 \rangle \geq \frac{|\mathcal{I}|}{10} \theta^2.$$

By Chernoff-Hoeffding bound combined with the guarantee of power iteration ([Fact 17](#)), as long as $T = \Omega(\log(T_{\text{des}}/\delta))$, then with probability at least $1 - \frac{\delta}{8T_{\text{des}}}$, for at least $\frac{4}{5}T$ iterations, it will be the case that \mathbf{w}_t is an approximate top singular vector. In other words, $|\mathcal{I}| \geq \frac{4}{5}T$, so that we have:

$$\sum_{t=1}^T \langle \boldsymbol{\tau}_t, \boldsymbol{\sigma}_t^2 \rangle \geq \frac{2T}{25} \theta^2.$$

Next, note that if we choose $\mathbf{p} = \mathbf{e}_i$, then

$$\sum_{t=1}^T \langle \mathbf{p}, \boldsymbol{\sigma}_t^2 \rangle = \sum_{t=1}^T \langle \mathbf{Z}'_i, \mathbf{w}_t \rangle^2.$$

Because $\boldsymbol{\tau}_1$ is uniform, the relative entropy term in [Theorem 12](#) is at most $\log k'$. Let's pretend for a moment that $\mathbf{e}_i \in \mathcal{K}$ (it is not). Then after plugging in the above calculations to [Theorem 12](#) and rearranging, we have that for every $i \in [k']$

$$\sum_{t=1}^T \langle \mathbf{Z}'_i, \mathbf{w}_t \rangle^2 \geq \frac{2T}{25} \theta^2 - 2 \log k' \geq 100 \log k',$$

by setting $T \geq \frac{10^5 \log k'}{\theta^2}$. This gives the bound claimed in the statement of the lemma, but it remains to fix the invalid assumption that $\mathbf{e}_i \in \mathcal{K}$. To do so, we will construct, for most $i \in [k']$, another distribution $\mathbf{p}' \in \mathcal{K}$ such that

$$\sum_{t=1}^T \langle \mathbf{e}_i, \boldsymbol{\sigma}_t^2 \rangle \geq \sum_{t=1}^T \langle \mathbf{p}', \boldsymbol{\sigma}_t^2 \rangle.$$

Combining this with $\sum_{t=1}^T \langle \mathbf{p}', \boldsymbol{\sigma}_t^2 \rangle \geq 100 \log k'$ gives the desired lower bound, for most i . Write $\boldsymbol{\alpha} = \sum_{t=1}^T \boldsymbol{\sigma}_t^2$, and without loss of generality assume that

$$\boldsymbol{\alpha}_1 \geq \boldsymbol{\alpha}_2 \geq \dots \geq \boldsymbol{\alpha}_{k'}.$$

For $i = 1, \dots, 4k'/5$, take \mathbf{p}' to be uniform on those $j \in [k']$ such that $\boldsymbol{\alpha}_i \geq \boldsymbol{\alpha}_j$ (there are at least $k'/5$ such j). By construction, we have that $\langle \boldsymbol{\alpha}, \mathbf{e}_i \rangle \geq \langle \boldsymbol{\alpha}, \mathbf{p}' \rangle$. Finally, observe that $\mathbf{p}' \in \mathcal{K}$ because \mathbf{p}' is uniform on a set of size at least $k'/5$. \blacksquare

Observe that the APPROXBREGMAN produces a sequence of vectors by the end. [Karnin et al. \(2012\)](#) provides a rounding algorithm that combines them into one with the desired margin bound. We describe the algorithm and prove the following lemma in [Appendix F](#).

Lemma 14 *The algorithm ROUND (Algorithm [F.1](#)) outputs \mathbf{w} that satisfies $|\langle \mathbf{Z}'_i, \mathbf{w} \rangle| \geq 0.1\theta$ for $0.45k$ of the points, with probability at least $1 - \delta/4T_{\text{des}}$.*

Finally, we can prove the key lemma using APPROXBREGMAN ([Appendix G](#)) and put everything together in [Appendix H](#).

Acknowledgments

The authors would like to thank Boaz Barak and Jelani Nelson for helpful conversations. In particular, we would like to thank Boaz for directing us to the paper (Barak et al. (2009)).

Zhixian Lei and Prayaag Venkat have been supported by NSF awards CCF 1565264 and CNS 1618026, and the Simons Foundation. Kyle Luh has been partially supported by NSF postdoctoral fellowship DMS-1702533. Prayaag Venkat has also been supported by an NSF Graduate Fellowship under grant DGE1745303.

References

- Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. *J. Comput. System Sci.*, 58(1, part 2):137–147, 1999. ISSN 0022-0000. doi: 10.1006/jcss.1997.1545.
- Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012. doi: 10.4086/toc.2012.v008a006.
- Boaz Barak, Moritz Hardt, and Satyen Kale. The uniform hardcore lemma via approximate Bregman projections. In *ACM-SIAM Symposium on Discrete Algorithms (SODA '09)*, pages 1193–1200, 2009. URL <http://dl.acm.org/citation.cfm?id=1496770.1496899>.
- Avrim Blum, John Hopcroft, and Ravindran Kannan. *Foundations of Data Science*. Cambridge University Press, 2019.
- Olivier Catoni. Challenging the empirical mean and empirical variance: a deviation study. *Ann. Inst. Henri Poincaré Probab. Stat.*, 48(4):1148–1185, 2012. ISSN 0246-0203. doi: 10.1214/11-AIHP454.
- Yu Cheng, Ilias Diakonikolas, and Rong Ge. High-dimensional robust mean estimation in nearly-linear time. In *ACM-SIAM Symposium on Discrete Algorithms (SODA '19)*, pages 2755–2771, 2019. doi: 10.1137/1.9781611975482.171.
- Yeshwanth Cherapanamjeri, Nicolas Flammarion, and Peter L Bartlett. Fast mean estimation with sub-gaussian rates. In *Conference on Learning Theory (COLT '19)*, pages 786–806, 2019a. URL <https://arxiv.org/abs/1902.01998>.
- Yeshwanth Cherapanamjeri, Samuel B. Hopkins, Tarun Kathuria, Prasad Raghavendra, and Nilesh Tripuraneni. Algorithms for heavy-tailed statistics: Regression, covariance estimation, and beyond. *arXiv preprint arXiv:1912.11071*, 2019b.
- Luc Devroye, Matthieu Lerasle, Gabor Lugosi, and Roberto I Oliveira. Sub-gaussian mean estimators. *Annals of Statistics*, 44(6):2695–2725, 2016. doi: 10.1214/16-AOS1440.
- Ilias Diakonikolas and Daniel M. Kane. Recent advances in algorithmic high-dimensional robust statistics. *arXiv preprint arXiv:1911.05911*, 2019. URL <https://arxiv.org/abs/1911.05911>.

- Ilias Diakonikolas, Gautam Kamath, Daniel M 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, 2019. doi: 10.1137/17M1126680.
- Yihe Dong, Samuel B Hopkins, and Jerry Li. Quantum entropy scoring for fast robust mean estimation and improved outlier detection. *arXiv preprint arXiv:1906.11366*, 2019. URL <https://arxiv.org/abs/1906.11366>.
- Samuel B Hopkins. Sub-gaussian mean estimation in polynomial time. *arXiv preprint arXiv:1809.07425*, 2018. URL <https://arxiv.org/abs/1809.07425>.
- Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoret. Comput. Sci.*, 43(2-3):169–188, 1986. ISSN 0304-3975. doi: 10.1016/0304-3975(86)90174-X.
- 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. doi: 10.1214/17-EJS1228.
- Zohar Karnin, Edo Liberty, Shachar Lovett, Roy Schwartz, and Omri Weinstein. Unsupervised SVMs: On the complexity of the furthest hyperplane problem. In *Conference on Learning Theory (COLT '12)*, pages 1–17, 2012. URL <http://proceedings.mlr.press/v23/karnin12.html>.
- Kevin A Lai, Anup B Rao, and Santosh Vempala. Agnostic estimation of mean and covariance. In *Symposium on Foundations of Computer Science (FOCS '16)*, pages 665–674, 2016. doi: 10.1109/FOCS.2016.76.
- Guillaume Lecué and Jules Depersin. Robust subgaussian estimation of a mean vector in nearly linear time. *arXiv preprint arXiv:1906.03058*, 2019. URL <https://arxiv.org/abs/1906.03058>.
- Matthieu Lerasle, Zoltan Szabo, Timothée Mathieu, and Guillaume Lecue. Monk outlier-robust mean embedding estimation by median-of-means. In *International Conference on Machine Learning*, pages 3782–3793, 2019.
- Gábor Lugosi and Shahar Mendelson. Near-optimal mean estimators with respect to general norms. *Probability Theory and Related Fields*, 2019a. ISSN 1432-2064. doi: 10.1007/s00440-019-00906-4.
- Gábor Lugosi and Shahar Mendelson. Sub-Gaussian estimators of the mean of a random vector. *Annals of Statistics*, 47(2):783–794, 2019b. ISSN 0090-5364. doi: 10.1214/17-AOS1639.
- Gábor Lugosi and Shahar Mendelson. Mean estimation and regression under heavy-tailed distributions: A survey. *Foundations of Computational Mathematics*, Aug 2019c. doi: 10.1007/s10208-019-09427-x.
- Stanislav Minsker. Geometric median and robust estimation in banach spaces. *Bernoulli*, 21(4): 2308–2335, 2015. doi: 10.3150/14-BEJ645.

A. S. Nemirovsky and D. B. and Yudin. *Problem complexity and method efficiency in optimization*. A Wiley-Interscience Publication. John Wiley & Sons, Inc., New York, 1983. ISBN 0-471-10345-4. Translated from the Russian and with a preface by E. R. Dawson, Wiley-Interscience Series in Discrete Mathematics.

Masashi Okamoto. Some inequalities relating to the partial sum of binomial probabilities. *Annals of the institute of Statistical Mathematics*, 10(1):29–35, 1959.

Adarsh Prasad, Sivaraman Balakrishnan, and Pradeep Ravikumar. A unified approach to robust mean estimation. *arXiv preprint arXiv:1907.00927*, 2019. URL <https://arxiv.org/abs/1907.00927>.

Appendix A. Main algorithm description

1. **Input:** Dataset \mathbf{Z}' and current estimate \mathbf{x}_t
2. $\mathbf{Z}'_i \leftarrow (\mathbf{Z}'_i - \mathbf{x}_t) / B$ by scaling each point by $B = \max_i \|\mathbf{Z}'_i - \mathbf{x}_t\|$.
3. $\theta \leftarrow$ the largest margin θ such that $\text{APPROXBREGMAN}(\mathbf{Z}', \theta, T)$ does not FAIL, where $T = O(\log k / \theta^2)$.
4. **Output:** $\hat{d} = \frac{B}{10}\theta$.

Algorithm A.1: Distance estimation—DISTEST

1. **Input:** Dataset \mathbf{Z}' and current estimate \mathbf{x}_t
2. $\mathbf{Z}'_i \leftarrow (\mathbf{Z}'_i - \mathbf{x}_t) / B$ by scaling each point by $B = \max_i \|\mathbf{Z}'_i - \mathbf{x}_t\|$.
3. $\theta \leftarrow$ the largest margin θ such that $\text{APPROXBREGMAN}(\mathbf{Z}', \theta, T)$ does not FAIL, where $T = O(\log k / \theta^2)$.
4. $\hat{\mathbf{g}} \leftarrow \text{APPROXBREGMAN}(\mathbf{Z}', \theta, T)$
5. If $\langle \hat{\mathbf{g}}, \mathbf{Z}'_i \rangle \geq 0.1\theta$ for at least $0.5k$ of the \mathbf{Z}'_i , **output** $\hat{\mathbf{g}}$; otherwise, **output** $-\hat{\mathbf{g}}$.

Algorithm A.2: Gradient estimation—GRADEST

Appendix B. Technical facts

We formally state the statistical guarantee of empirical average and coordinate-wise median-of-means. The former is an application of the Chebyshev's inequality. The latter is folklore but can follow easily from the Lugosi-Mendelson condition by considering the projections onto standard basis vectors.

Lemma 15 (empirical mean) *Let $\delta \in (0, 1)$. Given n i.i.d. copies $\mathbf{X}_1, \dots, \mathbf{X}_n$ of a random vector $\mathbf{X} \in \mathbb{R}^d$ with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$, let $\bar{\boldsymbol{\mu}} = \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i$. Then with probability at least $1 - \delta$,*

$$\|\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}\| \leq \sqrt{\frac{\text{Tr}(\boldsymbol{\Sigma})}{\delta n}}.$$

Lemma 16 (coordinate-wise median-of-means) *Assume the Lugosi-Mendelson condition (Assumption 1). Let $\{\mathbf{Z}'_i\}_{i=1}^k$ be the bucket means from n points (with at most $k/200$ contaminated) and $\hat{\boldsymbol{\mu}}$ be their coordinate-wise median-of-means. Then with probability at least $1 - \delta/8$,*

$$\|\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}\| \leq 600\sqrt{dr_\delta} \lesssim \sqrt{\frac{d\|\boldsymbol{\Sigma}\| \log(1/\delta)}{n}}.$$

1. **Input:** Dataset $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n \in \mathbb{R}^d$
2. Let $k = 3600 \log(1/\delta)$. Divide the data into $2k$ groups.
3. Compute the bucket mean of each group: $\mathbf{Z}_1, \mathbf{Z}_2, \dots, \mathbf{Z}_{2k} \in \mathbb{R}^d$.
4. Compute the coordinate-wise median-of-means of the second half of bucket means:

$$\mathbf{x}_0 \leftarrow \text{MEDIANOFMEANS}(\{\mathbf{Z}_{k+1}, \dots, \mathbf{Z}_{2k}\}).$$
5. Prune the first half of bucket means, where \mathbf{Z} is the data matrix of $\{\mathbf{Z}_i\}_{i=1}^k$:

$$\mathbf{Z}' \leftarrow \text{PRUNE}(\mathbf{Z}, \mathbf{x}_0).$$
6. $T_{\text{des}} \leftarrow \Theta(\log d), \eta \leftarrow 1/8000$
7. Run the main descent procedure: $\hat{\boldsymbol{\mu}} \leftarrow \text{DESCENT}(\mathbf{Z}', \mathbf{x}_0, T_{\text{des}}, \eta)$, using DISTEST and GRADEST as above.
8. **Output:** $\hat{\boldsymbol{\mu}}$

Algorithm A.3: Final algorithm

Our algorithm requires computing an approximation of the top (right) singular vector of a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$. The classic power method is efficient for this task.

Fact 17 (power iteration; see Theorem 3.1 of Blum et al. (2019)) *Let $\lambda(\mathbf{A}) = \max_{\mathbf{x} \in \mathbb{S}^{n-1}} \|\mathbf{A}\mathbf{x}\|_2^2$. With probability at least $9/10$, the power method (with random initialization) outputs a unit vector \mathbf{w} such that $\|\mathbf{A}\mathbf{w}\|_2^2 \geq \frac{\lambda(\mathbf{A})}{2}$ in $O(\log n)$ iterations. Moreover, each iteration can be performed in $O(mn)$ time.*

The following is a standard bound on binomial tail.

Lemma 18 (Okamoto (1959)) *Let $H(n, p)$ be a binomial random variable. Then*

$$\Pr(H(n, p) \geq 2np) \leq \exp(-np/3).$$

Appendix C. Omitted proofs from Section 3

Proof [Proof of Lemma 5] First, suppose that in some iteration t it holds that $\|\boldsymbol{\mu} - \mathbf{x}_t\| \leq 14000r_\delta$. Then

$$\frac{1}{21} \|\boldsymbol{\mu} - \mathbf{x}_{t^*}\| \leq d_{t^*} \leq d_t \leq 2\|\boldsymbol{\mu} - \mathbf{x}_t\| \leq 28000r_\delta,$$

so that we may conclude $\|\boldsymbol{\mu} - \mathbf{x}_{t^*}\| \leq 588000r_\delta$. Second, suppose that in all iterations t it holds that $\|\boldsymbol{\mu} - \mathbf{x}_t\| > 14000r_\delta$. Then by the update rule with $\eta = 1/8000$,

$$\|\mathbf{x}_{t+1} - \boldsymbol{\mu}\|^2 = \|\mathbf{x}_t - \boldsymbol{\mu}\|^2 + 2\eta d_t \langle \mathbf{x}_t - \boldsymbol{\mu}, \mathbf{g}_t \rangle + \eta^2 d_t^2 \|\mathbf{g}_t\|^2$$

$$\begin{aligned}
 &\leq \|\mathbf{x}_t - \boldsymbol{\mu}\|^2 - \frac{1}{800000} d_t \|\boldsymbol{\mu} - \mathbf{x}_t\| + \frac{1}{16000000} \|\mathbf{x}_t - \boldsymbol{\mu}\|^2 \\
 &\leq \|\mathbf{x}_t - \boldsymbol{\mu}\|^2 - \frac{1}{1680000} \|\boldsymbol{\mu} - \mathbf{x}_t\|^2 + \frac{1}{16000000} \|\mathbf{x}_t - \boldsymbol{\mu}\|^2 \\
 &= \left(1 - \frac{179}{336000000}\right) \|\mathbf{x}_t - \boldsymbol{\mu}\|^2
 \end{aligned}$$

Hence, the error bound drops at a geometric rate. The conclusion follows since $\|\boldsymbol{\mu} - \mathbf{x}_0\| \leq O(\sqrt{kd\|\boldsymbol{\Sigma}\|/n}) \leq O(\sqrt{dr_\delta})$. \blacksquare

Appendix D. Omitted proof from Section 4

First recall that Cherapanamjeri et al. (2019a) showed that the optimal solution to $\mathcal{M}(\mathbf{x}_t, \mathbf{Z})$ satisfies the property that θ is a valid distance estimate (Definition 3) and \mathbf{w} a gradient estimate (Definition 4).

Lemma 19 (Lemma 1 of Cherapanamjeri et al. (2019a)) *For all $t = 1, 2, \dots, T$, let $d_t = \theta^*$ be the optimal value of $\mathcal{M}(\mathbf{x}_t, \mathbf{Z})$. Then $|d_t - \|\boldsymbol{\mu} - \mathbf{x}_t\|| \leq 600r_\delta$, so d_t is a distance estimate with respect to \mathbf{x}_t .*

Lemma 20 (Lemma 2 of Cherapanamjeri et al. (2019a)) *For all $t = 1, 2, \dots, T$, let $(\theta^*, \mathbf{b}^*, \mathbf{w}^*)$ be the optimal solution of $\mathcal{M}(\mathbf{x}_t, \mathbf{Z})$. Then \mathbf{g}_t is a distance estimate with respect to \mathbf{x}_t .*

We now start by proving a generic claim that any reasonably good *bicriteria approximation* of $\mathcal{M}(\mathbf{x}_t, \mathbf{Z})$ suffices to provide gradient and distance estimates.

Definition 21 (bicriteria solution) *Let θ^* be the optimal value of $\mathcal{M}(\mathbf{x}, \mathbf{Z})$. We say that $(\theta, \mathbf{b}, \mathbf{w})$ is a (α, β) -bicriteria solution to $\mathcal{M}(\mathbf{x}, \mathbf{Z})$ if $\sum_i b_i \geq \alpha k$ and $b_i \langle \mathbf{Z}_i - \mathbf{x}, \mathbf{w} \rangle \geq b_i \theta$ for all i , where $\theta = \beta \theta^*$.*

Lemma 22 (distance estimate) *Let $(\theta, \mathbf{b}, \mathbf{w})$ be a $(1/10, 1/20)$ -bicriteria solution to $\mathcal{M}(\mathbf{x}_t, \mathbf{Z})$. Then $d_t = \theta$ is a distance estimate with respect to \mathbf{x}_t .*

Proof [Proof of Lemma 22] By Lemma 19, the optimal value θ^* lies in the range

$$[\|\boldsymbol{\mu} - \mathbf{x}_t\| - 600r_\delta, \|\boldsymbol{\mu} - \mathbf{x}_t\| + 600r_\delta].$$

Moreover, since $\theta^*/20 \leq \theta \leq \theta^*$, we have that

$$\frac{\|\boldsymbol{\mu} - \mathbf{x}\|}{20} - 30r_\delta \leq \theta \leq \frac{\|\boldsymbol{\mu} - \mathbf{x}\|}{20} + 30r_\delta. \quad (\text{D.1})$$

- When $\|\boldsymbol{\mu} - \mathbf{x}\| \geq 14000r_\delta$, we get from the inequality (D.1) that

$$\frac{\|\boldsymbol{\mu} - \mathbf{x}\|}{21} \leq \theta \leq \frac{\|\boldsymbol{\mu} - \mathbf{x}\|}{19}.$$

- When $\|\boldsymbol{\mu} - \mathbf{x}\| \leq 14000r_\delta$, $\theta \leq 730r_\delta < 28000r_\delta$, again by (D.1).

■

Lemma 23 (gradient estimate) *Let $(\theta, \mathbf{b}, \mathbf{w})$ be a $(1/10, 1/20)$ -bicriteria solution to $\mathcal{M}(\mathbf{x}_t, \mathbf{Z})$. Then $\mathbf{g}_t = \mathbf{w}$ is a gradient estimate with respect to \mathbf{x}_t .*

Proof [Proof of Lemma 23] Let $\mathbf{g}^* = (\boldsymbol{\mu} - \mathbf{x}_t) / \|\boldsymbol{\mu} - \mathbf{x}_t\|$ be the true gradient. We need to show that $\langle \mathbf{g}^*, \mathbf{g}_t \rangle \geq 1/20$. On the one hand, by Lemma 19, we have

$$d_t = \theta \geq \frac{1}{20} (\|\boldsymbol{\mu} - \mathbf{x}_t\| - 600r_\delta). \quad (\text{D.2})$$

On the other hand, for at least $k/10$ points, we have $\langle \mathbf{Z}_i - \mathbf{x}_t, \mathbf{g}_t \rangle \geq d_t$ and for at least $0.95k$ points, we have $\langle \mathbf{Z}_i - \boldsymbol{\mu}, \mathbf{g}_t \rangle \leq 600r_\delta$ by Assumption 1. Hence, there must be a point \mathbf{Z}_j that satisfies both inequalities, so it follows that

$$d_t \leq \langle \mathbf{Z}_j - \mathbf{x}_t, \mathbf{g}_t \rangle = \langle \mathbf{Z}_j - \boldsymbol{\mu}, \mathbf{g}_t \rangle + \langle \boldsymbol{\mu} - \mathbf{x}_t, \mathbf{g}_t \rangle \leq 600r_\delta + \|\boldsymbol{\mu} - \mathbf{x}_t\| \langle \mathbf{g}^*, \mathbf{g}_t \rangle. \quad (\text{D.3})$$

Using (D.2) and (D.3) and rearranging,

$$\langle \mathbf{g}^*, \mathbf{g}_t \rangle \geq \frac{1}{20} - \frac{630r_\delta}{\|\boldsymbol{\mu} - \mathbf{x}_t\|} \geq \frac{1}{200},$$

where we use $\|\boldsymbol{\mu} - \mathbf{x}_t\| \geq 14000r_\delta$. ■

Now we show that the optimal solution to the two-sided relaxation give distance and gradient estimate.

Lemma 24 *Let $(\theta', \mathbf{b}', \mathbf{w}')$ be an optimal solution of $\mathcal{M}_2(\mathbf{x}, \mathbf{Z})$. We have that*

- (i) *the value θ' lies in $[\|\boldsymbol{\mu} - \mathbf{x}\| - 600r_\delta, \|\boldsymbol{\mu} - \mathbf{x}\| + 600r_\delta]$; and*
- (ii) *one of the following two statements must hold, if $\|\boldsymbol{\mu} - \mathbf{x}\| \geq 14000r_\delta$:*
 - *there is a set \mathcal{C} of at least $0.9k$ points such that $\langle \mathbf{Z}_i - \mathbf{x}, \mathbf{w}' \rangle \geq \theta'$ for all $i \in \mathcal{C}$; or*
 - *there is a set \mathcal{C} of at least $0.9k$ points such that $\langle \mathbf{Z}_i - \mathbf{x}, -\mathbf{w}' \rangle \geq \theta'$ for all $i \in \mathcal{C}$.*

Proof [Proof of Lemma 24] Let θ be the optimal value of $\mathcal{M}(\mathbf{x}, \mathbf{Z})$. To prove (i), first recall that Lemma 19 states that $\theta \geq \|\boldsymbol{\mu} - \mathbf{x}\| - 600r_\delta$. Therefore, we get that $\theta' \geq \|\boldsymbol{\mu} - \mathbf{x}\| - 600r_\delta$, as $\theta' \geq \theta$. For the upper bound, assume for the sake of a contradiction that $\theta' > \|\boldsymbol{\mu} - \mathbf{x}\| + 600r_\delta$. Then one side of the hyperplane defined by \mathbf{w}' must contain at least $19/40$ fraction of points, so let's suppose without loss of generality that

$$\langle \mathbf{Z}_i - \mathbf{x}, \mathbf{w}' \rangle \geq \theta' > \|\boldsymbol{\mu} - \mathbf{x}\| + 600r_\delta \quad (\text{D.4})$$

for at least $19k/40$ \mathbf{Z}_i 's. Also, note that

$$\langle \mathbf{Z}_i - \mathbf{x}, \mathbf{w}' \rangle = \langle \mathbf{Z}_i - \boldsymbol{\mu}, \mathbf{w}' \rangle + \langle \boldsymbol{\mu} - \mathbf{x}, \mathbf{w}' \rangle \leq \|\boldsymbol{\mu} - \mathbf{x}\| + \langle \mathbf{Z}_i - \boldsymbol{\mu}, \mathbf{w}' \rangle. \quad (\text{D.5})$$

Combining (D.4) and (D.5), it follows that for at least $19k/40$ \mathbf{Z}_i 's we have

$$\langle \mathbf{Z}_i - \boldsymbol{\mu}, \mathbf{w}' \rangle > 600r_\delta. \quad (\text{D.6})$$

On the other hand, consider projections of all bucket means \mathbf{Z}_i onto \mathbf{w}' . [Assumption 1](#) implies that

$$|\{i : \langle \mathbf{w}', \mathbf{Z}_i \rangle - \langle \mathbf{w}', \boldsymbol{\mu} \rangle \geq 600r_\delta\}| \leq 0.05k.$$

This means that at most $k/20$ points satisfy $\langle \mathbf{Z}_i - \boldsymbol{\mu}, \mathbf{w}' \rangle \geq 600r_\delta$, contradicting (D.6).

To prove (ii), let $S^+ = \{i : \langle \mathbf{Z}_i - \mathbf{x}, \mathbf{w}' \rangle \geq \theta'\}$ and $S^- = \{i : \langle \mathbf{Z}_i - \mathbf{x}, -\mathbf{w}' \rangle \geq \theta'\}$. Notice that since $\|\boldsymbol{\mu} - \mathbf{x}\| \geq 14000r_\delta$, S^+ and S^- are disjoint. Now let

$$B = \{i : |\langle \mathbf{w}', \mathbf{Z}_i - \boldsymbol{\mu} \rangle| \leq 600r_\delta\} = \{i : |\langle \mathbf{w}', \mathbf{Z}_i - \mathbf{x} \rangle - \langle \mathbf{w}', \boldsymbol{\mu} - \mathbf{x} \rangle| \leq 600r_\delta\}.$$

By [Assumption 1](#), $|B| \geq 19k/20$. Consider the two cases.

- If $\langle \mathbf{w}', \boldsymbol{\mu} - \mathbf{x} \rangle \geq 0$, observe that B must intersect S^+ but not S^- . This implies that $|S^-| \leq k/20$, so $|S^+| \geq 9k/10$, since $|S^+| + |S^-| = 19k/20$ and they are disjoint.
- If $\langle \mathbf{w}', \boldsymbol{\mu} - \mathbf{x} \rangle < 0$, by the same argument, we have $|S^-| \geq 9k/10$.

■

Next, we show that approximating \mathcal{M}_2 in a bicriteria manner achieves a similar guarantee.

Lemma 25 *Let θ^* be the optimal value of $\mathcal{M}(\mathbf{x}, \mathbf{Z})$ and \mathbf{w}' be a unit vector such that for at least $k/8$ of the \mathbf{Z}_i , we have $|\langle \mathbf{w}', \mathbf{Z}_i - \mathbf{x} \rangle| \geq \theta'$, where $\theta' = 0.1\theta^*$. One of the following two statements must hold if $\|\boldsymbol{\mu} - \mathbf{x}\| \geq 14000r_\delta$.*

- there is a set \mathcal{C} of at least $0.95k$ points such that $\langle \mathbf{Z}_i - \mathbf{x}, \mathbf{w}' \rangle \geq \theta' - 600r_\delta$ for all $i \in \mathcal{C}$;
- there is a set \mathcal{C} of at least $0.95k$ points such that $\langle \mathbf{Z}_i - \mathbf{x}, -\mathbf{w}' \rangle \geq \theta' - 600r_\delta$ for all $i \in \mathcal{C}$.

Proof [Proof of [Lemma 25](#)] Let $\mathcal{C} = \{i : |\langle \mathbf{w}', \mathbf{Z}_i - \boldsymbol{\mu} \rangle| \leq 600r_\delta\}$ be the set of ‘‘good’’ points with respect to direction \mathbf{w}' . By [Assumption 1](#), $|\mathcal{C}| \geq 19k/20$. Further, let $S = \{i : |\langle \mathbf{w}', \mathbf{Z}_i - \mathbf{x} \rangle| \geq \theta'\}$, which we assume has size at least $k/8$. Thus, by pigeonhole principle, there must be a point, say \mathbf{Z}_j , that is in both sets. There are two cases.

- Suppose $\langle \mathbf{w}', \boldsymbol{\mu} - \mathbf{x} \rangle \geq 0$. Since $j \in S$ and $\theta^* \geq 13400r_\delta$ by [Lemma 24](#), we have $|\langle \mathbf{w}', \mathbf{Z}_j - \mathbf{x} \rangle| \geq 1340r_\delta$. On the other hand, since $j \in \mathcal{C}$,

$$|\langle \mathbf{w}', \mathbf{Z}_j - \boldsymbol{\mu} \rangle| = |\langle \mathbf{w}', \mathbf{Z}_j - \mathbf{x} \rangle - \langle \mathbf{w}', \boldsymbol{\mu} - \mathbf{x} \rangle| \leq 600r_\delta. \quad (\text{D.7})$$

Hence, we observe that $\langle \mathbf{w}', \mathbf{Z}_j - \mathbf{x} \rangle \geq \theta' \geq 1340r_\delta$. By definition of \mathcal{C} , all its points cluster around \mathbf{Z}_j by an additive factor of $600r_\delta$.

- Suppose $\langle \mathbf{w}', \boldsymbol{\mu} - \mathbf{x} \rangle \leq 0$. We get the second case in the claim by the same argument.

■

Finally, we are ready to prove [Lemma 6](#).

Proof [Proof of [Lemma 6](#)] Let's first check the distance estimate ([Definition 3](#)) guarantee.

- If $\|\boldsymbol{\mu} - \mathbf{x}\| \geq 14000r_\delta$, we have

$$\theta' \geq \frac{1}{10}\|\boldsymbol{\mu} - \mathbf{x}\| - 60r_\delta \geq \frac{2}{35}\|\boldsymbol{\mu} - \mathbf{x}\|,$$

since $\theta' = 0.1\theta^*$ and $\theta^* \geq \|\boldsymbol{\mu} - \mathbf{x}\| - 600r_\delta$. The upper bound of (3.1) obviously holds.

- If $\|\boldsymbol{\mu} - \mathbf{x}\| \leq 14000r_\delta$, we have $\theta' \leq 1460r_\delta$ by Lemma 24.

For gradient estimate, we appeal to Lemma 25 and get that if $\|\boldsymbol{\mu} - \mathbf{x}\| \geq 14000r_\delta$, then either $(\theta', \mathbf{b}', \mathbf{w}')$ or $(\theta', \mathbf{b}', -\mathbf{w}')$ is a $(19/20, 1/20)$ -bicriteria approximation of $\mathcal{M}(\mathbf{x}, \mathbf{Z})$, where \mathbf{b}' is the indicator vector of \mathcal{C} . Thus, we can apply Lemma 23, and this completes the proof. \blacksquare

Appendix E. Omitted proof from subsection 5.1

We remark that under Lugosi-Mendelson condition, the assumption $\|\boldsymbol{\mu} - \mathbf{x}_0\| \lesssim \sqrt{kd\|\boldsymbol{\Sigma}\|/n}$ can be easily achieved by initializing \mathbf{x}_0 to be the coordinate-wise median-of-means (Lemma 16) (with a failure probability at most $\delta/8$).

Lemma 26 (pruning) *Let $\beta = 600\sqrt{kd\|\boldsymbol{\Sigma}\|/n}$, and suppose $\|\boldsymbol{\mu} - \mathbf{x}_0\| \leq \beta$. Given the bucket means $\mathbf{Z} \in \mathbb{R}^{k \times d}$ such that at most $k/200$ points are contaminated, the algorithm PRUNE removes $k/10$ of the points and guarantees that with probability at least $1 - \delta/8$, among the remaining data,*

$$\max_i \|\mathbf{Z}_i - \boldsymbol{\mu}\| \leq O(\beta).$$

Further, $\text{PRUNE}(\mathbf{Z}, \mathbf{x}_0)$ can be implemented in $\tilde{O}(kd)$ time.

Proof [Proof of Lemma 26] For correctness, consider $\|\mathbf{Z}_i - \boldsymbol{\mu}\|$, and by triangle inequality,

$$\|\mathbf{Z}_i - \mathbf{x}_0\| - \|\boldsymbol{\mu} - \mathbf{x}_0\| \leq \|\mathbf{Z}_i - \boldsymbol{\mu}\| \leq \|\mathbf{Z}_i - \mathbf{x}_0\| + \|\boldsymbol{\mu} - \mathbf{x}_0\|.$$

Since $\|\boldsymbol{\mu} - \mathbf{x}_0\| \leq \beta$ by our assumption,

$$\|\mathbf{Z}_i - \mathbf{x}_0\| - \beta \leq \|\mathbf{Z}_i - \boldsymbol{\mu}\| \leq \|\mathbf{Z}_i - \mathbf{x}_0\| + \beta. \tag{E.1}$$

Let $\mathcal{S}_{\text{good}} = \{i : \|\mathbf{Z}_i - \boldsymbol{\mu}\| \leq \beta\}$ and $\mathcal{S}_{\text{bad}} = \{i : \|\mathbf{Z}_i - \boldsymbol{\mu}\| \geq 20\beta\}$. It suffices to show that with probability at least $1 - \delta/8$ all the points in \mathcal{S}_{bad} are removed. We first lower bound the number of good points. Each uncontaminated \mathbf{Z}_i is an average of $\lfloor n/k \rfloor$ i.i.d. random vectors. Applying Lemma 15 on estimation error of empirical mean, we obtain that for each uncontaminated i , with probability at least $1 - 1/1000$,

$$\|\mathbf{Z}_i - \boldsymbol{\mu}\| \leq \sqrt{1000 \cdot \text{Tr}(\boldsymbol{\Sigma})k/n} \leq \beta.$$

Therefore, each uncontaminated \mathbf{Z}_i is in $\mathcal{S}_{\text{good}}$ with probability at least $1 - 1/1000$. Let H be the number of uncontaminated points not in $\mathcal{S}_{\text{good}}$ and $p = 1/1000$. Since there are at most $k/200$ contaminated points and each uncontaminated point is independent, by a binomial tail bound (Lemma 18)

$$\Pr(H \leq 2p \cdot (199/200)k) \geq 1 - \exp(-p \cdot (199/200)k/3)$$

$$\begin{aligned} &\geq 1 - \exp(-\log(8/\delta)) \\ &= 1 - \delta/8, \end{aligned}$$

where we used $k = \lceil 3600 \log(8/\delta) \rceil$. Hence, with probability at least $1 - \delta/8$, $\mathcal{S}_{\text{good}}$ contains at least $(399/400)k$ (uncontaminated) points. We condition on this event for the rest of the proof.

Now observe that

$$\|\mathbf{Z}_i - \mathbf{x}_0\| < \|\mathbf{Z}_j - \mathbf{x}_0\| \quad \text{for each } j \in \mathcal{S}_{\text{bad}} \text{ and } i \in \mathcal{S}_{\text{good}} \quad (\text{E.2})$$

by (E.1). Suppose for a contradiction that $j \in \mathcal{S}_{\text{bad}}$ is not removed by line 4. Then it means that $d_j \leq d_i$ for $k/10$ of the \mathbf{Z}'_i 's. By pigeonhole principle, this implies $d_j \leq d_i$ for some $i \in \mathcal{S}_{\text{good}}$, since $|\mathcal{S}_{\text{good}}| \geq (399/400)k$. This contradicts condition (E.2).

Computing the distances takes $O(kd)$ time and sorting takes $O(k \log k)$ time. Thus, the algorithm PRUNE runs in time $O(kd + k \log k)$ and succeeds with probability at least $1 - \delta/8$. ■

Pruning allows us to bound the norms of the points $\mathbf{Z}_i - \mathbf{x}_t$ for each iteration t .

Corollary 27 (scaling and margin) *Suppose $\|\boldsymbol{\mu} - \mathbf{x}\| \leq O\left(\sqrt{\|\boldsymbol{\Sigma}\|kd/n}\right)$ and $\|\boldsymbol{\mu} - \mathbf{x}\| \geq \Omega(r_\delta)$. Let \mathcal{S} be the pruned dataset of size $k' \geq 9k/10$ such that $\|\mathbf{Z}_i - \boldsymbol{\mu}\| \leq O\left(\sqrt{\|\boldsymbol{\Sigma}\|kd/n}\right)$ for each $i \in \mathcal{S}$. There exists a scaling factor B , $\theta > 0$ and unit vector \mathbf{w} such that for at least $4k/5$ points in \mathcal{S} ,*

$$\left| \left\langle \frac{1}{B}(\mathbf{Z}_i - \mathbf{x}), \mathbf{w} \right\rangle \right| \geq \theta.$$

Further, we have that $1/\theta^2 = O(d)$.

Proof [Proof of Corollary 27] Let $B = \max_{i \in \mathcal{S}} \|\mathbf{Z}_i - \mathbf{x}\|$. Then B is bounded by

$$\|\mathbf{Z}_i - \mathbf{x}\| \leq \|\mathbf{Z}_i - \boldsymbol{\mu}\| + \|\boldsymbol{\mu} - \mathbf{x}\| \leq O\left(\sqrt{\|\boldsymbol{\Sigma}\|kd/n}\right). \quad (\text{E.3})$$

By Lemma 24, there exists a unit vector \mathbf{w} such that for at least $0.8k$ points in \mathcal{S} , $\langle \mathbf{Z}_i - \mathbf{x}, \mathbf{w} \rangle \geq \theta'$ and $\theta' = \Omega(r_\delta)$. Hence, we get that

$$\theta = \Omega\left(\frac{r_\delta}{B}\right) = \Omega\left(\frac{\sqrt{k\|\boldsymbol{\Sigma}\|/n} + \sqrt{\text{Tr}\boldsymbol{\Sigma}/n}}{\sqrt{\|\boldsymbol{\Sigma}\| \cdot kd/n}}\right) = \Omega\left(1/\sqrt{d}\right).$$

■

Proof [Proof of Lemma 7] The lemma follows directly from Lemma 26 and Corollary 27. ■

Appendix F. Omitted proofs from subsection 5.2

Karnin et al. (2012) provides a rounding scheme that combines the sequence of vectors produced by APPROXBREGMAN into one vector that satisfies the desired margin bound. The original routine succeeds with constant probability. We simply perform independent trials to boost the rate. We now analyze the algorithm. We cite the following lemma for the guarantee of the rounding algorithm (Algorithm F.1).

1. **Input:** Buckets means \mathbf{Z}' , unit vectors $\mathbf{w}_1, \dots, \mathbf{w}_T \in \mathbb{R}^d$, margin θ ,
2. Round to a single vector: $\mathbf{w} = \frac{\mathbf{w}'}{\|\mathbf{w}'\|}$, where $\mathbf{w}' = \sum_{t=1}^T g_t \mathbf{w}_t$ and $g_t \sim \mathcal{N}(0, 1)$, for $t = 1, \dots, T$.
3. Repeat until $|\langle \mathbf{Z}'_i, \mathbf{w} \rangle| \geq \frac{1}{10}\theta$ for at least $0.6k'$ of \mathbf{Z}'_i :
 - (a) Sample $g_t \sim \mathcal{N}(0, 1)$, for $t = 1, \dots, T$.
 - (b) Recompute $\mathbf{w} = \mathbf{w}'/\|\mathbf{w}'\|$, where $\mathbf{w}' = \sum_{t=1}^T g_t \mathbf{w}_t$.
 - (c) Report FAIL if more than $\Omega(\log(T_{\text{des}}/\delta))$ trials have been performed.
4. **Output:** \mathbf{w}

Algorithm F.1: Rounding algorithm—ROUND

Lemma 28 (Lemma 6 of Karnin et al. (2012)) *Suppose that for at least $\frac{3}{4}$ fraction of $i \in [k']$, it holds that*

$$\sum_{t=1}^T \langle \mathbf{Z}'_i, \mathbf{w}_t \rangle^2 \geq \log k'. \quad (\text{F.1})$$

Let $\mathbf{w}_1, \dots, \mathbf{w}_T$ be the unit vectors satisfying the above condition. Then with constant probability, the vector \mathbf{w} in each repetition of the step 3 of the ROUND algorithm (Algorithm F.1) satisfies $|\langle \mathbf{Z}'_i, \mathbf{w} \rangle| \geq \theta/10$ for at least a 0.45 fraction of $i \in [k']$.

Now we prove the guarantee of ROUND.

Proof [Proof of Lemma 14] By Lemma 28, it suffices to prove inequality (F.1) holds for at least a 3/4 fraction of the points. By the regret analysis (Lemma 13), the vectors $\mathbf{w}_1, \dots, \mathbf{w}_T$ produced during the iterations of Algorithm 5.2 satisfy the hypothesis of Lemma 28. Hence, the guarantee of Lemma 28 holds with constant probability. Moreover, we can test that this guarantee holds in time $O(Tk'd)$. To boost the success probability to $1 - \delta'$ (with $\delta' = \delta/4T_{\text{des}}$), ROUND algorithm performs $O(\log(1/\delta'))$ independent trials. Hence, it reports FAIL with probably at most δ' . Otherwise, by its definition, the output \mathbf{w} satisfies desired bound $|\langle \mathbf{Z}'_i, \mathbf{w} \rangle| \geq 0.1\theta$ for $0.45k$ of the points. ■

Appendix G. Full proof of key lemma

Proof [Proof of Lemma 8] The correctness follows from Lemma 14. We focus on run-time. By Assumption 2, we have that $1/\theta^2 = O(d)$. By projecting onto the subspace spanned by the bucket means, we can assume $d \leq k$. Hence, Lemma 13 implies that the iteration count is $\tilde{O}(k')$. The runtime of each iteration is bounded by the cost of computing an approximate top singular vector of a k' by d matrix via the power method, which is $\tilde{O}(k'd)$ by Fact 17. Finally, each repetition of the rounding algorithm ROUND takes time $\tilde{O}(k'd)$, and the number of trials is at most $O(\log(1/\delta'))$ by definition. Thus, the runtime of the rounding algorithm is $\tilde{O}(k^2d)$. ■

Appendix H. Full proof of main theorem

Proof [Proof of [Theorem 1](#)] Our argument assumes the following global events.

- (i) The Lugosi-Mendelson condition ([Assumption 1](#)) holds.
- (ii) The initial estimate \mathbf{x}_0 satisfies $\|\boldsymbol{\mu} - \mathbf{x}_0\| \leq 600\sqrt{\|\boldsymbol{\Sigma}\|kd/n}$.
- (iii) The pruning step succeeds: $\|\mathbf{Z}'_i - \boldsymbol{\mu}\| \leq O\left(\sqrt{\|\boldsymbol{\Sigma}\|kd/n}\right)$

We consider our main algorithm ([Algorithm A.3](#)) and first prove the correctness of DISTEST and GRADEST. Let \mathbf{Z}' be defined as in line 2 of DISTEST and GRADEST. [Lemma 24](#) states that there exists a margin θ^* in $[\|\boldsymbol{\mu} - \mathbf{x}\| - 600r_\delta, \|\boldsymbol{\mu} - \mathbf{x}\| + 600r_\delta]$. When $\|\boldsymbol{\mu} - \mathbf{x}_t\| \geq 14000r_\delta$, we have that for at least $0.8k$ points \mathbf{Z}'_i it holds $B \cdot |\langle \mathbf{Z}'_i, \mathbf{w}^* \rangle| \geq \theta^*$ for some unit vector \mathbf{w}^* , since the data are scaled by B . Furthermore, when the pruning step succeeds, [Assumption 1](#) holds. This allows us to apply the key lemma ([Lemma 8](#)).

- (i) For GRADEST, we use binary search in line 3 to find $\theta = \theta^*/B$. By the guarantee of [Lemma 8](#), $|\langle \mathbf{w}, \mathbf{Z}'_i \rangle| \geq \frac{\theta}{10}$ for at least $k/8$ of the \mathbf{Z}_i . It follows that $|\langle \mathbf{w}, \mathbf{Z}_i - \mathbf{x}_t \rangle| \geq \frac{B\theta}{10}$ for at least $k/8$ of the \mathbf{Z}_i . Thus, [Lemma 6](#) implies that the output \mathbf{g}_t is a gradient estimate.
- (ii) By the same argument, we apply [Lemma 6](#) and conclude that \hat{d}_t of DISTEST is a distance estimate.

Finally, we apply [Lemma 5](#) for the guarantee of DESCENT.

Next we bound several failure probabilities of the algorithm. The first three correspond to the global conditions.

- By [Lemma 2](#), the Lugosi-Mendelson condition [Assumption 1](#) fails with probability at most $\delta/8$.
- By [Lemma 16](#), the coordinate-wise median-of-means error bound fails with probability at most $\delta/8$.
- By [Lemma 7](#), the guarantee of our pruning and scaling procedure ([Assumption 2](#)) fails with probability at most $\delta/8$.
- Conditioned on above, the APPROXBREGMAN satisfies the guarantee of the key lemma ([Lemma 8](#)). The failure probability is at most $\delta/4T_{\text{des}}$ each iteration. We take union bound over all these iterations.

Overall, the failure probability of the entire algorithm ([Algorithm A.3](#)) is bounded by δ via union bound.

The runtime follows from [Lemma 8](#) which claims each iteration takes time $\tilde{O}(k^2d)$ and the fact that $T_{\text{des}} = \tilde{O}(1)$. ■

Appendix I. Interpretation of FHP algorithm Karnin et al. (2012) as regret minimization

Here, we review the bicriteria approximation algorithm of Karnin *et al.* Karnin et al. (2012) and show how it can be interpreted in the multiplicative weights update (MWU) framework for regret minimization. Given $\mathbf{Z}_1, \dots, \mathbf{Z}_k \in \mathbb{R}^d$ such that $\|\mathbf{Z}_i\| \leq 1$, we study the following *furthest hyperplane problem*:

$$\begin{aligned} &\text{Find } \mathbf{w} \in S^{d-1} \\ &\text{subject to } |\langle \mathbf{Z}_i, \mathbf{w} \rangle| \geq r \text{ for } i = 1, \dots, k, \end{aligned}$$

where we are promised that there does indeed exist a feasible solution \mathbf{w}^* . Since this problem is (provably) hard (even to approximate) we will settle for *bicriteria approximate* solutions. By this, we simply mean that we require the algorithm to output some \mathbf{w} such that $|\langle \mathbf{Z}_i, \mathbf{w} \rangle| \geq \frac{r}{10}$ for most of the $i \in [k]$. For our applications, the particular constants will not matter much, as long as they are actually constants.

See Algorithm I.1 for a formal description. First we give some intuition and then we sketch the important steps in the analysis.

1. **Input:** $\mathbf{Z}_1, \dots, \mathbf{Z}_k \in \mathbb{R}^d$ and iteration count $T \in \mathbb{N}$.
2. Initialize weights: $\tau_1 = \frac{1}{k}(1, \dots, 1) \in \mathbb{R}^k$.
3. For $t = 1, \dots, T$, repeat:
 - (a) Let \mathbf{A}_t be the $k \times d$ matrix whose i th row is $\sqrt{\tau_t(i)}\mathbf{Z}_i$ and \mathbf{w}_t be the top right unit singular vector of \mathbf{A}_t .
 - (b) Set $\sigma_t(i) = |\langle \mathbf{Z}_i, \mathbf{w}_t \rangle|$.
 - (c) Reweight: $\tau_{t+1}(i) = \tau_t(i)\eta^{-\sigma_t^2(i)}$ for $i \in [k]$ for an appropriately chosen constant η . In MWU language, σ_t^2 is the loss vector at time t .
 - (d) Normalize: Let $Z = \sum_{i \in [k]} \tau_{t+1}(i)$ and redefine $\tau_{t+1} \leftarrow \frac{1}{Z}\tau_{t+1}$.
4. **Output:** $\mathbf{w}_1, \dots, \mathbf{w}_T \in S^{d-1}$.

Algorithm I.1: Iterative MWU procedure

I.1. Intuition

Because we are promised that \mathbf{w}^* exists, averaging the constraints yields:

$$\frac{1}{k} \sum_{i=1}^k \langle \mathbf{Z}_i, \mathbf{w}^* \rangle^2 \geq r^2.$$

Note that if we define \mathbf{A}_1 as in Algorithm I.1, then the definition of singular vector tells us that:

$$\max_{\mathbf{w} \in S^{d-1}} \|\mathbf{A}_1 \mathbf{w}\|^2 = \frac{1}{k} \sum_{i=1}^k \langle \mathbf{Z}_i, \mathbf{w} \rangle^2 \geq \frac{1}{k} \sum_{i=1}^k \langle \mathbf{Z}_i, \mathbf{w}^* \rangle^2 \geq r^2.$$

Thus, \mathbf{w}_1 , the top singular vector as defined in Algorithm I.1, satisfies the constraints *on average*. It could be the case that $\langle \mathbf{Z}_4, \mathbf{w}_1 \rangle^2 \gg r^2$ but $\langle \mathbf{Z}_i, \mathbf{w} \rangle^2 \ll r^2$ for all $i \neq 4$. To fix this issue, we would simply down-weight \mathbf{Z}_4 in the next iteration, so that \mathbf{w}_2 aligns more with \mathbf{Z}_i for $i \neq 4$. We repeat this several times, with each \mathbf{w}_t improving upon \mathbf{w}_{t-1} .

At the end, the algorithm produces a collection of vectors $\mathbf{w}_1, \dots, \mathbf{w}_T$ which each satisfy a certain property. While it seems natural to just output \mathbf{w}_T as the final answer, it turns out that this will not work. Instead, we need to apply a randomized rounding procedure to extract a single vector \mathbf{w} from $\mathbf{w}_1, \dots, \mathbf{w}_T$.

I.2. Analysis

Lemma 29 *When Algorithm I.1 terminates after $T = O(\frac{\log k}{r^2})$ iterations, for every $i \in [k]$ it holds that:*

$$\sum_{t=1}^T \langle \mathbf{Z}_i, \mathbf{w}_t \rangle^2 \geq \frac{\log k}{\log \eta}.$$

Proof Algorithm I.1 is simply the MWU algorithm with the experts corresponding to the k constraints and the loss of expert i at time t being $\sigma_t^2(i)$. Using the regret guarantee from Theorem 2.1 in Arora et al. (2012) with respect to the fixed expert e_i and step size η :

$$\sum_{t=1}^T \langle \boldsymbol{\tau}_t, \boldsymbol{\sigma}_t^2 \rangle - (1 + \eta) \sum_{t=1}^T \langle \mathbf{e}_i, \boldsymbol{\sigma}_t^2 \rangle \leq \frac{\log k}{\eta}. \quad (\text{I.1})$$

Note that

$$\sum_{t=1}^T \langle \mathbf{e}_i, \boldsymbol{\sigma}_t^2 \rangle = \sum_{t=1}^T \langle \mathbf{Z}_i, \mathbf{w}_t \rangle^2$$

and

$$\begin{aligned} \sum_{t=1}^T \langle \boldsymbol{\tau}_t, \boldsymbol{\sigma}_t^2 \rangle &= \sum_{t=1}^T \sum_{i=1}^k \tau_t(i) \sigma_t^2(i) \\ &= \sum_{t=1}^T \sum_{i=1}^k \tau_t(i) \langle \mathbf{Z}_i, \mathbf{w}_t \rangle^2 && \text{(by definition of the algorithm)} \\ &\geq \sum_{t=1}^T \sum_{i=1}^k \tau_t(i) \langle \mathbf{Z}_i, \mathbf{w}^* \rangle^2 && \text{(since } \mathbf{w}^* \text{ is the top eigenvector)} \\ &\geq \sum_{t=1}^T \sum_{i=1}^k \tau_t(i) r^2 \\ &= T r^2. \end{aligned}$$

Substituting into and simplifying the regret formula and taking $\eta = 1/3$ gives the claim. \blacksquare

Given the previous lemma, we can just apply the rounding algorithm as a black-box to the output of Algorithm I.1.

Lemma 30 (Karnin et al. (2012)) *Let $\alpha \in (0, 1)$ and $\mathbf{w}_1, \dots, \mathbf{w}_T$ be unit vectors satisfying the conclusion of the previous lemma. Then with probability at least $1/147$, the output \mathbf{w} of the Rounding Algorithm F.1 satisfies $|\langle \mathbf{Z}_i, \mathbf{w} \rangle| \geq \alpha r$ for at least a $1 - 3\alpha$ fraction of $i \in [k]$.*

Appendix J. Conclusion and Discussion

In this paper, we provided a faster algorithm for estimating the mean of a heavy-tailed random vector that achieves subgaussian performance. Unlike previous algorithms, our faster running time is achieved by the use of a simple spectral method that iteratively updates the current estimate of the mean until it is sufficiently close to the true mean.

Our work suggests two natural directions for future research. First, is it possible to achieve subgaussian performance for heavy-tailed *covariance* estimation in polynomial time? Currently, the best polynomial-time covariance estimators do not achieve the optimal statistical rate (see [Lugosi and Mendelson \(2019c\)](#); [Cherapanamjeri et al. \(2019b\)](#)), while a natural generalization of the (computationally intractable) Lugosi-Mendelson estimator is known to achieve subgaussian performance. One approach would be to build on our framework; the key technical challenge is to design an efficient subroutine for producing bi-criteria approximate solutions to the natural generalization of the inner maximization problem to the covariance setting.

Another direction is to achieve a truly linear-time algorithm for the mean estimation problem. Our iterative procedure for solving the inner maximization problem take $\tilde{O}(k)$ iterations; is it possible to reduce this to a constant?