# List-Decodable Sparse Mean Estimation via Difference-of-Pairs Filtering 

Ilias Diakonikolas<br>UW-Madison<br>ilias@cs.wisc.edu

Daniel M. Kane<br>UC, San Diego<br>dakane@cs.ucsd.edu

Sushrut Karmalkar<br>UW-Madison<br>skarmalkar@wisc.edu

Ankit Pensia<br>UW-Madison<br>ankitp@cs.wisc.edu

Thanasis Pittas<br>UW-Madison<br>pittas@wisc.edu


#### Abstract

We study the problem of list-decodable sparse mean estimation. Specifically, for a parameter $\alpha \in(0,1 / 2)$, we are given $m$ points in $\mathbb{R}^{n},\lfloor\alpha m\rfloor$ of which are i.i.d. samples from a distribution $D$ with unknown $k$-sparse mean $\mu$. No assumptions are made on the remaining points, which form the majority of the dataset. The goal is to return a small list of candidates containing a vector $\hat{\mu}$ such that $\|\hat{\mu}-\mu\|_{2}$ is small. Prior work had studied the problem of list-decodable mean estimation in the dense setting. In this work, we develop a novel, conceptually simpler technique for listdecodable mean estimation. As the main application of our approach, we provide the first sample and computationally efficient algorithm for list-decodable sparse mean estimation. In particular, for distributions with "certifiably bounded" $t$-th moments in $k$-sparse directions and sufficiently light tails, our algorithm achieves error of $(1 / \alpha)^{O(1 / t)}$ with sample complexity $m=(k \log (n))^{O(t)} / \alpha$ and running time poly $\left(m n^{t}\right)$. For the special case of Gaussian inliers, our algorithm achieves the optimal error guarantee $\Theta(\sqrt{\log (1 / \alpha)})$ with quasi-polynomial complexity. We complement our upper bounds with nearly-matching statistical query and low-degree polynomial testing lower bounds.


## 1 Introduction

It is well-established that when a dataset is corrupted by outliers, many commonly-used estimators fail to produce reliable estimates [Tuk60, $\mathrm{ABH}^{+} 72$ ]. The field of robust statistics was developed to perform reliable statistical inference in the presence of a constant fraction of outliers, even when the data is high-dimensional [HR09, HRRS86]. Although statistical rates of high-dimensional robust estimation problems are relatively well-understood by now [DL88, Yat85, DG92, HR09, CGR16], all of the estimators developed in these works were computationally inefficient, with runtime exponential in the dimension. The goal of algorithmic robust statistics, beginning with the works of [ $\mathrm{DKK}^{+} 16$, LRV16], is to design computationally efficient algorithms for high-dimensional robust estimation tasks. We refer the reader to the survey [DK19] for an introduction to this field.

The bulk of the recent progress in algorithmic robust statistics has focused on the setting where the fraction of outliers is a small constant, and the majority of samples are inliers, see, e.g., $\left[\mathrm{DKK}^{+} 16\right.$, LRV16, KKM18]. In contrast, when the fraction of outliers outnumbers the fraction of inliers, it is generally information-theoretically impossible to output a single estimate with non-vacuous error guarantees. In such situations, we allow the algorithm to return a small list of candidates such that one of the candidates is close to the true parameter. This list-decodable learning setting was first introduced in [BBV08] and developed in [CSV17]. We define the model below.

Definition 1.1 (List-Decodable Learning). Given a parameter $0<\alpha<1 / 2$ and a distribution family $\mathcal{D}$ on $\mathbb{R}^{n}$, the algorithm specifies $m \in \mathbb{Z}_{+}$and observes a set of $m$ samples constructed as follows: First, a set $S$ of $\lfloor\alpha m\rfloor$ i.i.d. samples are drawn from an (unknown) distribution $D \in \mathcal{D}$. Then, an adversary is allowed to inspect $S$ and choose a multiset $E$ of $m-\lfloor\alpha m\rfloor$ points. The multiset $T$, defined as $T:=S \cup E$, of $m$ points is given as input to the algorithm. We say that $D$ is the distribution of inliers, the elements in $S$ are inliers, the points in $E$ are outliers, and $T$ is an $(1-\alpha)$-corrupted dataset of $S$. The goal is to output a "small" list of hypotheses $\mathcal{L}$ at least one of which is (with high probability) close to the target parameter of $D$.

The list-decodable learning setting, interesting in its own right, is closely related to several wellstudied problems. A natural example is the problem of parameter recovery from mixture models, e.g., Gaussian mixtures (see, e.g., [Das99, VW04, AK05, DS07, KK10, RV17]). List-decodable mean estimation can serve as a key step in learning mixtures, since one can treat any component of the mixture as the set of inliers (see, e.g., [CSV17, DKS18a, KS17]). In addition, list-decodable learning can be used to model data in important applications where mixture models are not sufficient, such as crowdsourcing (see, e.g., [SVC16, SKL17, MV18]) and community detection in stochastic block models (e.g., [CSV17]).

Prior work on list-decodable mean estimation has focused on the unstructured setting, where the target mean is an arbitrary dense vector (see, e.g., [CSV17, KS17, DKS18a, RY20a, CMY20, DKK20, $\left.\mathrm{DKK}^{+} 21, \mathrm{DKK}^{+} 22 \mathrm{~b}\right]$ ). Sparse models have proven to be useful in a wide range of statistical tasks, and thus understanding the statistical and computational aspects of sparse estimation is a fundamental problem (see, e.g., [EK12, HTW15, van16]). Here we study list-decodable sparse mean estimation, where the target mean vector is known to be $k$-sparse, i.e., it has at most $k$ non-zero coordinates. Given an $(1-\alpha)$-corrupted set of samples, our goal is to output, in a computationally-efficient manner, a small list of vectors containing a good approximation $\widehat{\mu}$ to the true mean $\mu$ (cf. Definition 1.1). Importantly, the goal is to achieve this with far fewer samples than in the dense setting. While the dense setting would require sample size polynomial in $n$ - the ambient dimension of the data - the goal here is to solve the problem in number of samples polynomial in $k$ and only polylogarithmic in $n$.

In this paper, we present a novel and conceptually simple technique for list-decodable mean estimation (that is applicable even in the dense setting). Combining our framework with the concentration results from [ $\mathrm{DKK}^{+} 22 \mathrm{a}$ ], we obtain the first sample and computationally efficient algorithm for list-decodable sparse mean estimation. We note that, while prior results [KS17, RY20a] can possibly be modified to incorporate the sparsity framework of $\left[\mathrm{DKK}^{+} 22 \mathrm{a}\right]$ after sufficient effort, a notable contribution of our work is a general and conceptually simpler framework for list-decodable estimation, which can easily be adapted to incorporate various structural constraints.

### 1.1 Related Work

Efficient estimators for high-dimensional robust statistics with sparsity constraints have been recently developed for various problems, such as mean estimation and PCA (see, e.g., [BDLS17, DKK ${ }^{+}$19, Li17, $\left.\mathrm{DKK}^{+} 22 \mathrm{a}\right]$ ). The problem of list-decodable mean estimation was first introduced in [CSV17], which achieved an error guarantee of $\tilde{O}(1 / \sqrt{\alpha})$ for distributions with bounded second moment; this guarantee turned out to be optimal for this distributional assumption (see [DKS18a]). Subsequent work improved the algorithmic guarantees for this problem (see, e.g., [CMY20, DKK20, DKK ${ }^{+}$21]).
To achieve better error guarantees, it is necessary to make further assumptions on the distribution, e.g., Gaussianity or bounded higher moments. In terms of the minimax optimal rate ${ }^{1}$, [DKS18a] showed that the optimal error for Gaussians is $\Theta(\sqrt{\log (1 / \alpha)})$. They also showed that any SQ algorithm that achieves the optimal error for Gaussians must take either super-polynomial time or samples, and presented an algorithm with matching guarantees. When the distribution $D$ has bounded $t$-th moment for some even $t>2$ ( i.e., $\mathbf{E}\left[\langle v, X-\mathbf{E}[X]\rangle^{t}\right]$ is bounded for all unit vectors $v$ ), [DKS18a] proved that the minimax optimal rate is $\Theta\left(\alpha^{-1 / t}\right)$. For distributions with certifiably bounded $t$-th moments, [KS17, RY20a] provided algorithms obtaining an error rate of $O\left((1 / \alpha)^{O(1 / t)}\right)$ with sample complexities $m=\operatorname{poly}\left(n^{t} / \alpha\right)$, and runtimes poly $\left(m^{t} n^{t}\right)$ and poly $(m n)^{\text {poly }(t, 1 / \alpha)}$, respectively. Recently, in the context of moment estimation and clustering

[^0]problems, [ST21] showed how to improve the dependence on $m$ in the runtime of algorithms that depend on certifiably bounded moments from poly $\left(m^{t} n^{t}\right)$ to poly $\left(m n^{t}\right)$. While it is possible that their result might be applied to [KS17, RY20a], our algorithmic technique naturally lends itself to achieve runtime poly $\left(m n^{t}\right)$ for the dense as well as sparse settings. We provide detailed comparisons with (see, e.g., [DKS18a, KS17, RY20a]) in Section 1.3. Finally, we mention that the list-decodable setting has also been studied in the context of linear regression (see, e.g., [KKK19, RY20a, DKP ${ }^{+}$21]) and subspace recovery (see, e.g., [BK21, RY20b]).

In an independent and concurrent work [ZS22], Zeng and Shen also study list-decodable sparse mean estimation when the underlying distribution is spherical Gaussian.

### 1.2 Our Results

We demonstrate an algorithm to perform list-decodable sparse mean estimation with $(k \log n)^{O(t)}$ samples, when the mean $\mu \in \mathbb{R}^{n}$ is known to be $k$-sparse. For this to be possible, we will require some assumptions on the underlying distribution of inliers $D$. Prior work in the dense setting ([KS17, RY20a]) assumed that the inlier distribution $D$ in Definition 1.1 satisfies $d$-certifiably bounded $t$-th moments in every direction (i.e., for some moment bound $M>0, M\|v\|_{2}^{t}-\mathbf{E}_{X \sim D}\left[\langle v, X-\mu\rangle^{t}\right]$ can be expressed as a sum of square polynomials of degree at most $d=O(t)$ in the entries of $v$ ), and $D$ has light tails. We highlight that our algorithmic technique can also be used in the dense case, under the same assumption, and provides qualitatively similar error guarantees with much simpler arguments and improved runtime. Below we apply our technique to the sparse setting.
Our results hold when (i) the $t$-th moment of $D$ is $d$-certifiably bounded for every $v$ that is $k$-sparse, with $d=O(t)$, and (ii) $D$ has light tails. For ease of exposition, we state the result assuming $D$ has subexponential tails (i.e., for some universal constant $c$, for all unit vectors $v$ and $p \in \mathbb{N}$, $\left.\mathbf{E}_{X \sim D}\left[|\langle v, X-\mu\rangle|^{p}\right]^{1 / p} \leq c p\right) .{ }^{2}$
Theorem 1.2 (List-Decodable Sparse Mean Estimation). Let t be an integer power of two. Let $D$ be a distribution over $\mathbb{R}^{n}$ with $k$-sparse mean $\mu$. Suppose that $D$ has $t$-th moments $d$-certifiably bounded in $k$-sparse directions by $M$ for some $d=O(t)$ (cf. Definition 2.3) and subexponential tails in the standard basis directions. There is an algorithm which, given $\alpha, M, t, k$, and a ( $1-\alpha$ )-corrupted set of $m=(t k \log n)^{O(t)} \max \left(1, M^{-2}\right) / \alpha$ samples from $D$, runs in time poly $\left(m n^{t}\right)$ and returns a vector $\hat{\mu} \in \mathbb{R}^{n}$ such that with probability $\Omega(\alpha)$ it is the case that $\|\hat{\mu}-\mu\|_{2}=O_{t}\left(M^{1 / t} / \alpha^{O(1) / t}\right)$.

Note that with high probability over the inliers and for any choice of outliers, with probability $\Omega(\alpha)$ over the internal randomness of the algorithm, the algorithm of Theorem 1.2 outputs an estimate $\hat{\mu}$ close to $\mu$. By running our algorithm $O(1 / \alpha)$ times, we can generate a list of size $O(1 / \alpha)$ such that with probability 0.9 the list contains the desired estimate $\hat{\mu}$.

Notably, for the important special case of Gaussian $\mathcal{N}(\mu, I)$ inliers, our algorithm achieves the information-theoretically optimal error rate. This is because $\mathcal{N}(\mu, I)$ has its $t$-th moment certifiably bounded by $t^{t / 2}$ in all directions. Specifically, for a large enough constant $C>0$, we obtain the following result: Given $\alpha$, $t$, and a $(1-\alpha)$-corrupted set of $m \geq(t k \log n)^{C t}$ samples from $\mathcal{N}(\mu, I)$ for a $k$-sparse vector $\mu$, our algorithm runs in time poly $\left(m n^{t}\right)$ and with probability $\Omega(\alpha)$ outputs a vector $\widehat{\mu}$ such that $\|\widehat{\mu}-\mu\|_{2} \leq O\left(\sqrt{t} / \alpha^{C / t}\right)$. Thus, by taking $t=C \log (1 / \alpha)$, we obtain the optimal error of $\Theta(\sqrt{\log (1 / \alpha)})$ in quasi-polynomial sample and time complexity.
We also note that a broad and natural class of distributions satisfying Definition 2.3 is the class of $\sigma$-Poincare distributions (see, e.g., $[\mathrm{KS} 17]$ ). A distribution is said to be $\sigma$-Poincare if for all differentiable functions $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$, we have that $\operatorname{Var}_{X \sim D}[f(X)] \leq \sigma^{2} \mathbf{E}_{X \sim D}\left[\|\nabla f(X)\|_{2}^{2}\right]$.
We complement our algorithm of Theorem 1.2 with a qualitatively matching lower bound in the Statistical Query (SQ) model [Kea98]. Instead of directly accessing samples, SQ algorithms are only allowed to perform adaptive queries of expectations of bounded functions of the underlying distribution, up to some desired tolerance (c.f. Definition C.1). The class of SQ algorithms is fairly broad: a wide range of known algorithmic techniques in machine learning are known to be implementable in the SQ model (see, e.g., [FGR $\left.{ }^{+} 13\right]$ ).
An SQ lower bound is an unconditional statement that for any SQ algorithm, either the number of queries $q$ must be large or the tolerance, $\tau$, of some query must be small. Since simulating a query of

[^1]tolerance $\tau$ by averaging i.i.d. samples may need up to $\Omega\left(1 / \tau^{2}\right)$ many of them, SQ lower bounds are naturally interpreted as a tradeoff between runtime $\Omega(q)$ and sample complexity $\Omega\left(1 / \tau^{2}\right)$. An adaptation of the result in [DKS18a] yields Theorem 1.3, which indicates that the $k^{O(t)}$ factor in the sample complexity of Theorem 1.2 might be necessary for efficient algorithms, even for Gaussian inliers.

Theorem 1.3 (SQ Lower Bound, Informal). Consider the problem of list-decoding the mean of $\mathcal{N}(\mu, I)$, for a $k$-sparse vector $\mu \in \mathbb{R}^{n}$, up to error better than $O\left((t \alpha)^{-1 / t}\right)$. Any SQ algorithm that solves the problem does one of the following: (i) It returns a list of size $n^{k^{\Omega(1)}}$, (ii) it uses at least one query of tolerance $k^{-\Omega(t)} \exp \left(O(t \alpha)^{-2 / t}\right)$, or (iii) it makes at least $n^{k^{\Omega(1)}}$ queries.

Theorem 1.3 states that any SQ algorithm that runs in time less than $n^{k^{c}}$ for a constant $c>0$ and achieves the minimax-optimal error of $\sqrt{\log (1 / \alpha)}$ must use at least $k^{\Omega(\log (1 / \alpha) / \log \log (1 / \alpha))}$ samples. This follows by setting $t=\Omega(\log (1 / \alpha) / \log \log (1 / \alpha))$. A similar lower bound holds for the computational model of low-degree polynomial tests, as a consequence of the recently established relationship between the two models $\left[\mathrm{BBH}^{+} 21\right]$. See Appendix C for more details.

### 1.3 Overview of Techniques

We begin with a brief overview of the existing techniques for dense list-decodable mean estimation. We then highlight some of the obstacles in adapting these techniques to the sparse setting. Then, we provide an overview of our algorithmic approach and explain how it overcomes these obstacles.
Prior Work on List-Decodable Mean Estimation In the dense setting, prior algorithmic techniques for list-decoding the mean with error better than $\Omega\left(\alpha^{-1 / 2}\right)$ are quite complicated. [DKS18a] uses a multifilter-based technique for list-decoding spherical Gaussians, which relies critically on knowing the higher degree moments of the inliers (and thus does not generalize to less specific distribution families). Moreover, this method runs into technical difficulties related to being unable to determine the variance of higher degree polynomials on the inliers without knowing the mean ahead of time. The other approach in the literature (see, for example, [KS17, RY20a]) uses the Sum-of-Squares method (SoS) to find these clusters of points. The algorithm in [RY20a] involves solving a nested SoS program and then applying a complicated rounding procedure to get the final list. It should be noted that the runtime of [RY20a] is exponential in poly $(1 / \alpha)$, which can be quite large. Finally, [KS17] gave an SoS based list-decodable mean estimation algorithm for the dense case with error, sample complexity, and list size similar to the ones that are obtained in our work; but significantly worse runtime. The approach of [KS17] has some important differences. First, the clustering relaxation is conceptually harder, involving a more complex optimization problem for each filtering step and second, after the filtering ends, the error guarantee scales with the norm of the unknown mean; thus a complicated re-clustering step that combines ideas from [SCV18] is needed to reduce the error.
Here we present a significantly cleaner method to perform the outlier removal step; we avoid problems like not knowing the mean ahead of time by simply taking pairs of differences of our samples to make their mean zero. While it is likely that either of the prior techniques could be adapted to the sparse setting with sufficient effort, this would result in significantly more complicated algorithms. We briefly point out some difficulties below.

First, to ensure that the algorithms identify subsets of the samples that satisfy certifiably bounded moments in all $k$-sparse directions, requires additional variables and constraints to the algorithms. Additionally, one would need to replace the bounded moments in all directions condition by the corresponding condition for the sparse case, and ensure that all the proof steps can be modified to rely only on the latter - this would result in minor modifications of the original algorithms, such as thresholding of candidate solution vectors.
Second, at the end of this process, while the algorithm might qualitatively match the error guarantee that we achieve, the runtime would continue to be $(m n)^{O(t)}$ for [KS17] or $(1 / \alpha)^{\text {polylog }(1 / \alpha)} n^{O\left(\max \left\{1 / \alpha^{4}, t\right\}\right)}$ for [RY20a] - both of these are qualitatively worse than the runtime we achieve when $\alpha$ is sufficiently small. To obtain improved runtime using these prior techniques, one would require additional ideas, e.g., from [ST21], to be adapted to this setting, overall resulting in a far more complicated algorithm.
On the other hand, it seems difficult to adapt the multi-filtering technique from [DKS18a] to the setting we consider, without any introduction of an SoS component. We remind the reader that the [DKS18a]
algorithm depends critically on knowing the higher moments of the inliers exactly, and does not generalize to less specific distribution families. Even in the Gaussian setting, generalizing [DKS18a] might be difficult, since it would require the design of an efficiently verifiable notion of matching higher moments in $k$-sparse directions.

We believe that our novel list-decoding technique is significantly simpler. As a result of simplifying the optimization programs involved, our technique naturally improves the runtime from poly $\left(m^{t} n^{t}\right)$ in prior work to poly $\left(m n^{t}\right)$.

Novel List-Decodable Mean Estimation Algorithm All known list-decoding algorithms are based upon the following fundamental observation. Suppose that $S$ is a dataset that contains a subset $S_{\mathrm{g}}$ of samples with bounded moments. If we can find a subset $S^{\prime}$ of $S$ with bounded moments and large overlap with $S_{\mathrm{g}}$, then the means of $S^{\prime}$ and $S_{\mathrm{g}}$ cannot be too far apart (see Lemma 3.3). The goal of a list-decoding algorithm is to find such a subset $S^{\prime}$ (or, more precisely, a small number of hypotheses for such a subset). In order to generalize this to the sparse mean setting, it suffices for the subsets $S_{\mathrm{g}}$ and $S^{\prime}$ to have bounded moments only in all $k$-sparse directions. This will imply that $\left|\left\langle v, \mu_{\mathrm{g}}-\mu_{S^{\prime}}\right\rangle\right|$ is relatively small for all $k$-sparse unit vectors $v$, which will in turn imply that truncating $\mu_{S^{\prime}}$ to its $k$ largest entries will provide a suitable approximation to $\mu_{\mathrm{g}}$ (see Fact A.2).
The basic idea behind our novel list-decoding technique, which we call the difference of pairs filter, is the following: let $T$ be the set of differences of pairs of elements of $S$. The set $T$ will contain a relatively large subset, $T_{\mathrm{g}}$, (consisting of the pairwise differences of elements of $S_{\mathrm{g}}$ ) whose moments in $k$-sparse directions are bounded. Our goal will be to find a subset $T^{\prime} \subset T$ that has large overlap with $T_{\mathrm{g}}$ and also has bounded ( $k$-sparse) moments.
Naïvely, we can do this as follows. We start with $T^{\prime}=T$. Either this set has bounded $k$-sparse moments (in which case we are done) or there is some sparse direction $v$ in which the average value of $|\langle v, x\rangle|^{t}$ over $T^{\prime}$ is substantially larger than the average value over $T_{\mathrm{g}}$. By throwing away points $x$ from $T^{\prime}$ with probability proportional to $|\langle v, x\rangle|^{t}$, we eliminate mostly bad points. We repeat this until $T^{\prime}$ has bounded sparse moments. Unfortunately, while this approach can be shown to be correct, it does not suffice for our purposes because it is computationally infeasible in general to determine whether or not $T^{\prime}$ has bounded sparse moments. However, if we assume additionally that $S_{\mathrm{g}}$ has bounded $k$-sparse moments provable by a low-degree sum-of-squares proof (i.e., the moment bound inequality can be re-expressed as a sum of square polynomials being greater than 0 ), then so will $T_{\mathrm{g}}$. There is an efficient algorithm to determine whether or not $T^{\prime}$ has certifiably bounded moments in $k$-sparse directions as well. If $T^{\prime}$ does not have certifiably bounded moments, standard techniques for Sum-of-Squares programs imply that we can manufacture a non-negative polynomial $p$ so that the average value of $p$ on $T^{\prime}$ is substantially larger than the value of $p$ on any set with SoS-certifiable bounded moments in $k$-sparse directions. Thus, filtering out points $x$ in $T^{\prime}$ with probability proportional to $p(x)$ will likely remove mostly bad points. Using this idea, we can find such a set $T^{\prime}$ efficiently (see Theorem 3.1).

We are thus left with a set of differences of samples rather than a set of samples. At this point, we will need a rounding method that given a set $T^{\prime}$ of differences with bounded $k$-sparse moments guaranteed to have large overlap with $T_{\mathrm{g}}$, finds a small list of sets $S^{\prime}$ with bounded $k$-sparse moments so that at least one of them has a large overlap with $S_{\mathrm{g}}$. Note that $T^{\prime}$ might be the union of $T_{i}-T_{i}$, where each $T_{i}$ is an individual cluster drawn from a moment-bounded distribution. This demonstrates the necessity of a rounding step to identify the means of individual clusters. To achieve this, it is helpful to think of $T^{\prime}$ as consisting of a set of pairs of elements of $S$, or equivalently as a graph over $S$. Given $T^{\prime}$, our first task is to find some reasonably large subset of $S$ with bounded moments (ideally which has a large overlap with $S_{\mathrm{g}}$ ). It is not hard to see that it suffices to find any large clique in $T^{\prime}$ (see Lemma 3.3). Unfortunately, $T^{\prime}$ may well not have any large cliques, and even if it does, finding them may be computationally difficult. However, we are saved here by the observation that if $|\langle v, x-y\rangle|^{t}$ and $|\langle v, y-z\rangle|^{t}$ are both small, then so is $|\langle v, x-z\rangle|^{t}$. This means that if we replace $T^{\prime}$ by the graph $H$ of all pairs of vertices $(v, w)$, where $v$ and $w$ have many common neighbors in $T^{\prime}$, this new graph will also have relatively small $k$-sparse moments (see Lemma 3.5). As most elements of $S_{\mathrm{g}}$ are adjacent in $T^{\prime}$ to most other elements of $S_{\mathrm{g}}$, it is not hard to see that this new graph will have relatively large cliques; unfortunately, finding them may still be computationally difficult.

To find these large cliques efficiently, we need one final observation. If $v$ is a random vertex of a graph $G$, then there will (on average) be very few pairs of neighbors, $u$ and $w$, of $v$ so that $u$ and $w$ do not have a large number of common neighbors with each other in $G$ (see Lemma 3.6). This means
that if we pick a random sample $x$ in $S$, then most pairs of neighbors of $x$ in $T^{\prime}$ are neighbors in $H$. Using a densification procedure (see Lemma 3.7), it is not hard to find a large subset $S^{\prime}$ of these neighbors so that any two elements of $S^{\prime}$ have many neighbors in common in $G$. By Lemma 3.5, this implies that $S^{\prime}$ will in fact have bounded $k$-sparse moments. Furthermore, if we happened to pick $x$ from $S_{\mathrm{g}}$, it is not hard to see that it is likely that $S^{\prime}$ has large overlap with $S_{\mathrm{g}}$, and thus its mean provides us with a good estimate.
While the above describes our algorithmic approach, we also need to consider the sample complexity of our method. We know that the distribution $D$ has the property of bounded moments in $k$-sparse directions, and our algorithm requires that the property is also satisfied by the uniform distribution over the samples $S_{\mathrm{g}}$. In order for $S_{\mathrm{g}}$ to have it as well, it suffices that the $t^{t h}$ moment tensor of $S_{\mathrm{g}}-\mu$ be close to the corresponding moment tensor of $D-\mu$. It turns out that if these moment tensors are $\delta k^{-t}$-close coordinate-wise, which takes $O\left(t \log (n) k^{O(t)} / \delta^{2}\right)$ samples, this suffices to get the kind of certifiable concentration we require. This is captured by Lemma 2.4, a restatement from [DKK ${ }^{+} 22 \mathrm{a}$ ].

## 2 Preliminaries

Basic Notation. We use $\mathbb{Z}_{+}$to denote positive integers. For $n \in \mathbb{Z}_{+}$we denote $[n]:=\{1, \ldots, n\}$. We denote by $\mathbb{R}\left[x_{1}, \ldots, x_{n}\right]_{\leq d}$ the set of real-valued polynomials of degree at most $d$ in variables $x_{1}, \ldots, x_{n}$. We use poly $(\cdot)$ to indicate a quantity that is polynomial in its arguments. For an ordered set of variables $Q=\left\{x_{1}, \ldots, x_{n}\right\}$, we will denote $p(Q)$ to mean $p\left(x_{1}, \ldots, x_{n}\right)$. We use $I_{n}$ to denote the $n \times n$ identity matrix. For a vector $v$, we let $\|v\|_{2}$ denote its $\ell_{2}$-norm. We call a vector $k$-sparse if it has at most $k$ non-zero coordinates. We use $\langle v, u\rangle$ for the inner product of the vectors $u, v$. We use capital letters for random variables and $\mathbf{E}[\cdot]$ for expectation. When $S$ is a set, we will use the notation $X \sim S$ to mean that $X$ is distributed uniformly over $S$. For a graph $G$ we denote by $N_{G}(x)$ the neighborhood of $x$ in $G$. Throughout the paper, we will use the letter $n$ for the dimension, $m$ for the number of samples, $d$ for the degrees of the SoS proofs, and $t$ for the number of bounded moments.
SoS Preliminaries The following preliminaries are specific to the SoS part of this paper. We refer to [BS16, FKP19] for a more complete treatment of the SoS framework. Here, we review the basics.
Definition 2.1 (SoS Proof). Let $x_{1}, \ldots, x_{n}$ be indeterminates and let $\mathcal{A}$ be a set of polynomial equalities $\left\{p_{1}(x)=0, \ldots, p_{m}(x)=0\right\}$. An SoS proof of the inequality $r(x) \geq 0$ from axioms $\mathcal{A}$ is a set of polynomials $\left\{r_{i}(x)\right\}_{i \in[m]} \cup\left\{r_{0}(x)\right\}$ such that $r_{0}$ is a sum of square polynomial and $r_{i}$ 's are arbitrary, and $r(x)=r_{0}(x)+\sum_{i \in[m]} r_{i}(x) p_{i}(x)$. If the set of polynomials $\left\{r_{i}(x) \cdot p_{i}(x) \mid\right.$ $i \in[d]\} \cup\left\{r_{0}(x)\right\}$ have degree at most $d$, we say that this proof is of degree $d$ and denote it by $\left.\mathcal{A}\right|_{d} r(x) \geq 0$. When we need to emphasize what indeterminates are involved in a particular SoS proof, we denote it by $\mathcal{A} \left\lvert\, \frac{x}{d} r(x) \geq 0\right.$. When $\mathcal{A}$ is empty, we omit it, e.g., ${ }_{d} r(x) \geq 0$ or $\left\lvert\, \frac{x}{d} r(x) \geq 0\right.$.

We will also use the objects called pseudoexpectations.
Definition 2.2 (Pseudoexpectation). Let $x_{1}, \ldots, x_{n}$ be indeterminates. A degree-d pseudoexpectation $\tilde{\mathbf{E}}$ is a linear map $\tilde{\mathbf{E}}: \mathbb{R}\left[x_{1}, \ldots, x_{n}\right]_{\leq d} \rightarrow \mathbb{R}$ from degree-d polynomials to $\mathbb{R}$ such that $\tilde{\mathbf{E}}\left[p(x)^{2}\right] \geq 0$ for any $p$ of degree at most $d / 2$ and $\tilde{\mathbf{E}}[1]=1$. If $\mathcal{A}=\left\{p_{1}(x)=0, \ldots, p_{m}(x)=0\right\}$ is a set of polynomial inequalities, we say that a pseudoexpectation $\tilde{\mathbf{E}}$ satisfies $\mathcal{A}$ if for every $i \in[m]$, $\tilde{\mathbf{E}}\left[s(x) p_{i}(x)\right]=0$ for all polynomials $s(x)$ such that $s(x) p_{i}(x)$ has degree at most $d$.

It is well known (see, e.g., [BS16]) that pseudoexpectations are dual objects to SoS proofs in the following sense: given a set $\mathcal{P}$ of $r$ polynomial equalities in $n$ variables and a polynomial $q\left(x_{1}, \ldots, x_{n}\right)$, either there exists an SoS proof $\left.\mathcal{P}\right|_{\ell} ^{x} q(x) \geq 0$ or there exists a pseudoexpectation $\tilde{\mathbf{E}}$ of degree $\ell$ satisfying $\mathcal{P}$ but having $\tilde{\mathbf{E}}[q(x)]<0$. More importantly, there is an algorithm that runs in time $(r n)^{O(\ell)}$ and finds that pseudoexpectation when we are in the second case. Due to lack of space, we defer formal statements and a brief overview of SoS to Appendix A.1.

Certifiably Bounded Moments in $k$-sparse Directions Our algorithm succeeds whenever the uncorrupted samples have certifiably bounded moments in $k$-sparse directions, defined as in [DKK $\left.{ }^{+} 22 \mathrm{a}\right]$ :

Definition 2.3 (( $M, t, d)$ Certifiably Bounded Moments in $k$-sparse Directions). Let $Q:=$ $\left\{v_{1}, \ldots, v_{n}, z_{1}, \ldots, z_{n}\right\}$ and $\mathcal{A}_{k \text {-sparse }}:=\left\{z_{i}^{2}=z_{i}\right\}_{i \in[n]} \cup\left\{v_{i} z_{i}=v_{i}\right\}_{i \in[n]} \cup\left\{\sum_{i=1}^{n} z_{i}=k\right\} \cup$
$\left\{\sum_{i=1}^{n} v_{i}^{2}=1\right\}$. For an $M>0$ and even $t \in \mathbb{N}$, a distribution $D$ with mean $\mu$ satisfies $(M, t, d)$ certifiably bounded moments in $k$-sparse directions if $\mathcal{A}_{k \text {-sparse }} \left\lvert\, \frac{Q}{d} \mathbf{E}_{X \sim D}\left[\langle v, X-\mu\rangle^{t}\right]^{2} \leq M^{2}\right.$.

The definition of $\mathcal{A}_{k \text {-sparse }}$ is based on the fact that a vector $v=\left(v_{1}, \ldots, v_{n}\right)$ is $k$-sparse if and only if there exists $z=\left(z_{1}, \ldots, z_{n}\right)$ such that $v, z$ satisfy $\mathcal{A}_{k \text {-sparse }}$.
We will use the following lemma proved in $\left[\mathrm{DKK}^{+} 22 \mathrm{a}\right]$ to bound the number of samples it takes to certify bounded moments in $k$-sparse directions. Although this is stated for subexponential distributions, it needs the distribution to have only bounded $t^{2} \log (d)$ moments (see Lemma A.10).
Lemma 2.4 ( $\left.\left[\mathrm{DKK}^{+} 22 \mathrm{a}\right]\right)$. Let $D$ be a distribution over $\mathbb{R}^{n}$ with mean $\mu$. Suppose that $D$ has $c$ -sub-exponential tails around $\mu$ for a constant $c$ and that $\mathcal{A}_{k \text {-sparse }} \left\lvert\, \frac{Q}{O(t)} \mathbf{E}_{X \sim D}\left[\langle v, X-\mu\rangle^{t}\right]^{2} \leq M^{2}\right.$. Let $S=\left\{X_{1}, \ldots, X_{m}\right\}$ be a set of $m$ i.i.d. samples from $D, D^{\prime}$ be the uniform distribution over $S$, and $\bar{\mu}:=\mathbf{E}_{X \sim D^{\prime}}[X]$. If $m>C(t k(\log n))^{5 t} \max \left(1, M^{-2}\right)$ for a sufficiently large constant $C$, then with probability at least 0.9 we have the following: (i) $\mathcal{A}_{k \text {-sparse }} \frac{Q}{O(t)} \mathbf{E}_{X \sim D^{\prime}}\left[\langle v, X-\bar{\mu}\rangle^{t}\right]^{2} \leq 8 M^{2}$ and (ii) $\langle v, \bar{\mu}-\mu\rangle \leq M^{1 / t} / \alpha^{6 / t}$ for every $k$-sparse unit vector $v$.

## 3 Main Result: Proof of Theorem 1.2

Recall our setting: we are given $\alpha \in(0,1 / 2)$ and a multiset $S:=\left\{x_{1}, \ldots, x_{m}\right\}$ such that an unknown subset of $\lfloor\alpha m\rfloor$ many of these points satisfy ( $M, t, d$ ) certifiably bounded central moments in $k$-sparse directions, and the remaining are arbitrary. The goal is to recover a candidate that is close to $\mu:=\mathbf{E}_{X \sim D}[X]$ with probability $\Omega(\alpha)$. In what follows, $t$ will always be $2^{\ell}$ for some $\ell \in \mathbb{Z}_{+}$.

### 3.1 The SoS-based Filter

Let $T$ be the set of pairwise differences of all samples in $S$ (similarly denote by $T_{\mathrm{g}}$ the subset of $T$ corresponding to inliers $S_{\mathrm{g}}$ ). We would like to either detect that the moments of $T$ are already bounded in all $k$-sparse directions or find a direction that violates this and filter out mostly outliers. Unfortunately, this kind of check is computationally infeasible, but it can be done efficiently if we check for moment bounds that are certified by SoS proofs. This is done in Algorithm 1, which takes as input the set $T$ along with the parameters $t, d, M$ and performs filtering until the resulting set $T^{\prime}$ has ( $M, t, d$ )-bounded moments.

```
Algorithm 1 SoS-based filter for list-decodable mean estimation
    function LDMEAN-SoS-Filter \(\left(T:=\left\{x_{1}, \ldots, x_{m}\right\}, t, d, M\right)\)
        Let \(Q=\left\{v_{1}, \ldots, v_{n}, z_{1}, \ldots, z_{n}\right\}\) and \(T^{\prime}=T\).
        while there is no SoS proof of \(\mathcal{A}_{k \text {-sparse }} \frac{\mid Q}{d} \sum_{x \in T^{\prime}}\langle v, x\rangle^{t} \leq 6 M|T|\) do
            Find a degree- \(d \tilde{\mathbf{E}}\) on \(Q\) satisfying \(\mathcal{A}_{k \text {-sparse }}\) and \(\tilde{\mathbf{E}}\left[\sum_{x \in T}\langle v, x\rangle^{t}\right]>6 M|T|\).
            Throw out \(x \in T^{\prime}\) with probability \(\tilde{\mathbf{E}}\left[\langle v, x\rangle^{t}\right] / \max _{x \in T^{\prime}} \tilde{\mathbf{E}}\left[\langle v, x\rangle^{t}\right]\).
        end while
        return \(T^{\prime}\)
    end function
```

Theorem 3.1 (Filter Identifies a Subset Satisfying Bounded Moments). Let $T$ be a multiset of points in $\mathbb{R}^{n}$ for which there exists a subset $T_{g} \subset T$ with $\left|T_{g}\right|=\alpha^{2}|T|$ for some $\alpha>0$. Furthermore assume that $T_{g}$ has zero mean and $(M, t, d)$-certifiably bounded moments in $k$-sparse directions for some $M>0, d \in \mathbb{Z}_{+}$, and even $t$. Then Algorithm 1, given $T, M, t, d$, returns a subset $T^{\prime} \subseteq T$ in time $\operatorname{poly}\left(m n^{d}\right)$ so that, with probability at least $2 / 3$, the following holds: (i) For any $k$-sparse unit vector $v$, we have that $\sum_{x \in T^{\prime}}\langle v, x\rangle^{t} \leq 6 M|T|$, and (ii) $\left|T^{\prime} \cap T_{g}\right| \geq\left|T_{g}\right| / 2$.

Proof. Let $Q=\left\{z_{1}, \ldots, z_{n}, v_{1}, \ldots, v_{n}\right\}$. In Algorithm 1, Lines 3 and 4 use the separation oracle of Theorem A. 4 with $\mathcal{P}=\mathcal{A}_{k \text {-sparse }}$. Since $T_{\mathrm{g}}$ has zero mean and ( $M, t, d$ ) bounded central moments in $k$-sparse directions, we have that

$$
\begin{equation*}
\mathcal{A}_{k \text {-sparse }} \frac{Q}{d} M-\underset{X \sim T_{\mathrm{g}}}{\mathbf{E}}\left[\langle v, X\rangle^{t}\right] \geq 0 . \tag{1}
\end{equation*}
$$

Thus, if $\mathcal{A}_{k \text {-sparse }} \frac{Q}{d} \sum_{x \in T^{\prime}}\langle v, x\rangle^{t} \leq 6 M|T|$ the algorithm identifies that and stops (in which case we have the desired conclusion that for any $k$-sparse vector $v, \sum_{x \in T^{\prime}}\langle v, x\rangle^{t} \leq 6 M|T|$; otherwise it finds a degree- $d$ pseudo-expectation $\tilde{\mathbf{E}}$ on $Q$ satisfying $\mathcal{A}_{k \text {-sparse }}$ and $\tilde{\mathbf{E}}\left[-6 M|T|+\sum_{x \in T^{\prime}}\langle v, x\rangle^{t}\right]>$ 0 , in which case we can create a filter: Using Equation (1) and $\left|T_{\mathrm{g}}\right|=\alpha^{2}|T|$, we have that $M \geq$ $\mathbf{E}_{X \sim T_{\mathrm{g}}}\left[\tilde{\mathbf{E}}\left[\langle v, X\rangle^{t}\right]\right] \geq \frac{\alpha^{-2}}{|T|} \sum_{x \in T_{\mathrm{g}} \cap T^{\prime}} \tilde{\mathbf{E}}\left[\langle v, x\rangle^{t}\right]$. Since $\tilde{\mathbf{E}}\left[-6 M|T|+\sum_{x \in T^{\prime}}\langle v, x\rangle^{t}\right]>0$ and $\tilde{\mathbf{E}}$ is linear, we see that $\frac{\sum_{x \in T_{g} \cap T^{\prime}} \tilde{\mathbf{E}}\left[\langle v, x\rangle^{t}\right]}{\sum_{x \in T^{\prime}} \tilde{\mathbf{E}}\left[\langle v, x\rangle^{t}\right]} \leq \alpha^{2} / 6$. This means if we throw out sample $x$ with probability $\tilde{\mathbf{E}}\left[\langle v, x\rangle^{t}\right] / \max _{x \in T^{\prime}} \tilde{\mathbf{E}}\left[\langle v, x\rangle^{t}\right]$ (which is indeed in $[0,1]$ since $\langle v, x\rangle^{t}$ is SoS, so its pseudoexpectation is a non-negative value), on average, only an $\alpha^{2} / 6$ fraction of the points that are removed will be from $T_{\mathrm{g}}$. Since the $x$ with the largest value of $\tilde{\mathbf{E}}\left[\langle v, x\rangle^{t}\right]$ will be removed, the algorithm will terminate in polynomial time. We now analyze the size of $\left|T^{\prime} \cap T_{\mathrm{g}}\right|$. By the above analysis, at each step the expected number of samples thrown out from $T_{\mathrm{g}}$ is at most $\alpha^{2} / 6$ times the expected total number of samples removed. Thus, the potential function $\Delta=\left(\left|T_{\mathrm{g}} \cap T^{\prime}\right|-\right.$ $\left.\left(\alpha^{2} / 6\right)\left|T^{\prime}\right|\right) /|T|$ is a submartingale. Note that always $0 \leq \Delta / \alpha^{2} \leq 1$, and initially we had $\Delta / \alpha^{2} \geq$ $5 / 6$. By Doob's martimgale inequality (Proposition A. 9 applied with $t=1 / 2$ to the submartingale $\left.\left(\Delta / \alpha^{2}\right)\right)$, the probability that $\Delta / \alpha^{2}$ remains at least $\alpha^{2} / 2$ throughout the execution of the algorithm is at least $2 / 3$. Thus, we will have $\left|T_{\mathrm{g}} \cap T^{\prime}\right| \geq\left(\alpha^{2} / 2\right)|T|=\left|T_{\mathrm{g}}\right| / 2$ throughout the execution.

### 3.2 Identifying a Subset of Samples with Bounded Moments

Having identified a subset $T^{\prime} \subset T$ satisfying the conclusions of Theorem 3.1, we want to extract from $T^{\prime}$ a vector that is close to the original mean. Since the average of the set of differences is likely to be close to zero regardless of the true mean, we will need to use the information about the pairs that we get from $T^{\prime}$ to find subsets of the original samples that satisfy the appropriate concentration bounds. We will need the following definition.

Definition 3.2. Let $S \subset \mathbb{R}^{n}$. A graph $(V, E)$ on $S$ with $V=S$ is said to have $(M, t)$ - bounded moments in $k$-sparse directions if for all $k$-sparse unit vectors $v,\left(1 /|S|^{2}\right) \sum_{(x, y) \in E}\langle v, x-y\rangle^{t} \leq M$.

By the guarantee of our filter, if $T^{\prime}$ is the set returned, the graph $G$ with edges $(x, y)$ for which $x-y$ or $y-x$ belongs to $T^{\prime}$ will have bounded moments in the above sense. If $G$ contains a clique $C$ which intersects with an $\alpha$-fraction of the target samples $C_{\mathrm{g}}$, the means of $C$ and $C_{\mathrm{g}}$ are close.
Lemma 3.3. Let $S \subset \mathbb{R}^{n}$ and $G$ be a graph on $S$ satisfying Definition 3.2. Let $C \subset S$ be a clique in $G$. Let $C_{g} \subset C$ be a subset with $\left|C_{g}\right| \geq \alpha|S|$. If $\mu_{C}$ and $\mu_{g}$ denote the means of $C$ and $C_{g}$ respectively, then $\left\langle v, \mu_{C}-\mu_{g}\right\rangle^{t} \leq 2 M / \alpha^{2}$ for all $k$-sparse unit vectors $v$.

Proof. By using the fact that $t$ is even and the fact that $G$ satisfies Definition 3.2, we see that $M|S|^{2} \geq \sum_{(x, y) \in E}\langle v, x-y\rangle^{t} \geq \frac{1}{2} \sum_{x, y \in C}\langle v, x-y\rangle^{t} \geq \frac{1}{2} \sum_{x \in C_{g}, y \in C}\langle v, x-y\rangle^{t}$. Using Jensen's inequality we obtain, $M|S|^{2} \geq \sum_{x \in C_{g}, y \in C} \frac{\langle v, x-y\rangle^{t}}{2} \geq \frac{\left|C_{\mathrm{g}}\right||C|\left\langle v, \mu_{C}-\mu_{g}\right\rangle^{t}}{2} \geq \frac{|S|^{2} \alpha^{2}\left\langle v, \mu_{C}-\mu_{g}\right\rangle^{t}}{2}$.

Unfortunately, even the inliers might not form a clique in $G$. However, the guarantee that $\left|T^{\prime} \cap T_{g}\right| \geq$ $\left|T_{g}\right| / 2$ implies that the inliers share many neighbors in the graph $G$. Thus we look at the overlap graph defined below, in the hope that this graph will be more dense.
Definition 3.4 (Overlap Graph). Let $G=(V, E)$ be a graph and $\gamma>0$. The overlap graph $R_{\gamma}(G)$ is defined to be the graph with the vertex set $V$ where each $(x, y)$ is an edge in the graph iff $\left|N_{G}(x) \cap N_{G}(y)\right| \geq \gamma|V|$, where $N_{G}(x)$ denotes the neighborhood of the vertex $x$ in $G$.

The following result shows that if $G$ has bounded moments, then so does $R_{\gamma}(G)$. While this is not sufficient by itself, we will subsequently modify $R_{\gamma}(G)$ to ensure that we end up with a graph having bounded moments, as well as an identifiable clique.
Lemma 3.5 (If $G$ has bounded moments, then $R_{\gamma}(G)$ has bounded moments). Let $S$ be a set of points and $G$ be a graph with $(M, t)$-bounded moments in $k$-sparse directions. Then for $\gamma>0$, $R_{\gamma}(G)$ has $\left(2 \cdot 2^{t} M / \gamma, t\right)$-bounded moments in $k$-sparse directions.

Proof. For any $x, y$ in $R_{\gamma}(G)$, the triangle inequality implies $\langle v,(x-a)-(a-y)\rangle^{t} \leq 2^{t}[\langle v, x-$ $\left.a\rangle^{t}+\langle v, y-a\rangle^{t}\right]$. By taking a sum over all $a$ in $N_{G}(x) \cap N_{G}(y)$, we have

$$
\langle v, x-y\rangle^{t}=\sum_{a \in N_{G}(x) \cap N_{G}(y)} \frac{\langle v,(x-a)-(a-y)\rangle^{t}}{\left|N_{G}(x) \cap N_{G}(y)\right|} \leq\left(2^{t} / \gamma|S|\right) \sum_{\substack{a \in N_{G}(x) \cap N_{G}(y)}}\left[\langle v, x-a\rangle^{t}+\langle v, y-a\rangle^{t}\right] .
$$

Denote $\Gamma(\alpha, x):=\left\{y\right.$ : neighbor of $x$ in $R_{\gamma}(G)$ and neighbor of $a$ in $\left.G\right\}$. Summing over all the edges $(x, y)$ in $R_{\gamma}(G)$, and since $t$ is even and $G$ has $(M, t)$ bounded moments in $k$-sparse directions, we see: $\sum_{(x, y) \in R_{\gamma}(G)}\langle v, x-y\rangle^{t} \leq \frac{2^{t}}{\gamma|S|} \sum_{(x, y) \in R_{\gamma}(G)} \sum_{a \in N_{G}(x) \cap N_{G}(y)}\left[\langle v, x-a\rangle^{t}+\langle v, y-a\rangle^{t}\right]=$ $\frac{2 \cdot 2^{t}}{\gamma|S|} \sum_{(a, x) \in E} \sum_{y \in \Gamma(\alpha, x)}\langle v, x-a\rangle^{t} \leq \frac{2 \cdot 2^{t} \cdot|S|}{\gamma|S|} \sum_{(a, x) \in E}\langle v, a-x\rangle^{t} \leq \frac{2 \cdot 2^{t} M|S|^{2}}{\gamma}$.

While $R_{\gamma}(G)$ may not have any cliques either, it is guaranteed to have fairly dense subgraphs. We will subsequently prune out points so that the resulting final graph will have a clique.
Lemma 3.6 $\left(R_{\gamma}(G)\right.$ has dense subgraphs). Let $G=(V, E)$ be a graph and $\gamma>0$. If $x$ is a randomly selected vertex of $G$, then the expected number of pairs $y, z \in N_{G}(x)$ so that $y$ and $z$ are not neighbors in $R_{\gamma}(G)$ is at most $\gamma|V|^{2}$.

Proof. The expectation in question is $1 /|V|$ times the number of triples $x, y, z \in V$ so that $y$ and $z$ are not neighbors in $R_{\gamma}(G)$, but are both neighbors of $x$ in $G$. By the definition of $R_{\gamma}(G)$, if $y$ and $z$ are not neighbors in $R_{\gamma}(G)$, they have at most $\gamma|V|$ common neighbors in $G$. Thus, the number of such triples is at most $\gamma|V|^{3}$, so the expectation in question is at most $\gamma|V|^{2}$.

As outlined above, the inliers in $R_{\gamma}(G)$ form a dense subgraph. The next procedure (PRUNING in Algorithm 2) prunes out points from a dense subgraph (inliers in $R_{\gamma}(G)$ for us) to find a clique.

```
Algorithm 2 Algorithms for clique creation and rounding
    function Pruning \((G=(V, E), W \subset V)\)
        Let \(W^{\prime}=W\)
        while \(\exists x \in W^{\prime}\) that is not connected to at least \(2|W| / 3\) vertices in \(W^{\prime}:\) Remove \(x\) from \(W^{\prime}\)
        return \(W^{\prime}\)
    end function
    function \(\operatorname{Rounding}(S, G=(V, E))\)
        Let \(\delta=\alpha^{3} / 4608\).
        Choose \(x \in S\) uniformly at random, let \(W=N_{G}(x)\), and let \(G^{\prime}=R_{\delta}(G)\)
        if the number of pairs of points in \(W\) that are not connected in \(G^{\prime}\) is more than \((8 \delta / \alpha)|V|^{2}\)
    or if \(|W| \leq(\alpha / 4)|V|\) return FAIL
        else Run Pruning on \(G^{\prime}\) and \(W\) to obtain \(W^{\prime}\)
        return \(\mathbf{E}_{X \sim W^{\prime}}[X]\).
    end function
```

Lemma 3.7 (Dense subgraphs can be pruned to obtain a clique). Let $G=(V, E)$ be a graph and let $W \subset V$ be a set of vertices with $|W|=\beta|V|$ and all but $\gamma|V|^{2}$ pairs of vertices in $W$ are connected in $G$, for $\beta, \gamma>0$ with $\gamma \leq \beta^{2} / 36$. There exists an algorithm (PRUNING in Algorithm 2) that given $G, W, \beta, \gamma$ runs in polynomial time and returns $a W^{\prime} \subset W$ so that $\left|W^{\prime}\right| \geq|W|-(6 \gamma / \beta)|V|$ and so that $\left|W^{\prime}\right|$ is a clique in $R_{\beta / 3}(G)$.

Proof. In Line 3 of the Algorithm, the point $x$ which is removed satisfies $\left|N_{G}(x) \cap W^{\prime}\right|<2 / 3|W|$. If we also have that $\left|W^{\prime}\right| \geq 5|W| / 6$ (something that we will verify later), the removal of $x$ decreases the number of pairs of unconnected elements in $W^{\prime}$ by $\left|W^{\prime}\right|-\left|N_{G}(x) \cap W^{\prime}\right| \geq\left|W^{\prime}\right|-(2 / 3)|W| \geq$ $|W| / 6=(\beta / 6)|V|$. This can happen at most $(6 \gamma / \beta)|V|$ times before we run out of unconnected pairs of elements in $W^{\prime}$, thus $\left|W^{\prime}\right| \geq|W|-(6 \gamma / \beta)|V|$ upon termination. Also, since it holds $(6 \gamma / \beta)|V| \leq(\beta / 6)|V|=|W| / 6$, we indeed have $\left|W^{\prime}\right| \geq 5|W| / 6$ as claimed at the start. Now note that each element of $W^{\prime}$ is connected to at least $2|W| / 3$ other elements of $W^{\prime}$ in $G$. Thus any pair of elements of $W^{\prime}$ have at least $|W| / 3$ common neighbors, and thus are adjacent in $R_{\beta / 3}(G)$.
We are finally ready to prove our main algorithmic result on rounding. The basic idea is that most of the inliers in $G$ (which are at least $\alpha$-fraction of vertices) are connected to many other inliers, and
thus if we start with an inlier, its neighborhood will also contain many inliers and will be dense in the overlap graph $G^{\prime}$ (Lemma 3.6). Thus, we can apply the pruning of Lemma 3.7 to obtain a large clique in the overlap graph of $G^{\prime}$, which also has bounded moments by two applications of Lemma 3.5.
Theorem 3.8 (Rounding). Let $S \subset \mathbb{R}^{n}$ and let $G=(V, E)$ be a graph with $V=S$ and $(M, t)$ bounded moments in $k$-sparse directions. Suppose there is a subset $S_{g} \subset S$ with $\left|S_{g}\right| \geq \alpha|S|$ and at least half of the pairs of points in $S_{g}$ are connected by an edge in $G$. Suppose that the $S_{g}$ has mean $\mu_{g}$ and $t$-th moment bounded by $M$ in $k$ sparse directions. Then, there exists a randomized algorithm that given $G, S$ and $\alpha$ runs in polynomial time and returns $a \widehat{\mu} \in \mathbb{R}^{n}$ such that with probability $\Omega(\alpha)$, for all $k$-sparse unit vectors $v,\left\langle v, \widehat{\mu}-\mu_{g}\right\rangle^{t}=O\left(10^{t} M \alpha^{-6}\right)$

Proof. The algorithm we consider is Rounding (Algorithm 2). Let $x$ and $G^{\prime}$ be as in Line 8 . We will claim that algorithm Rounding succeeds as long as the following hold: (i) $x \in S_{\mathrm{g}}$, (ii) $x$ has at least $S_{\mathrm{g}} / 4$ neighbors in $S_{\mathrm{g}}$ in $G^{\prime}$, and (iii) the number of pairs of neighbors of $x$ that are not neighbors in $G^{\prime}$ is at most $(8 \delta / \alpha)|V|^{2}$. First, we show that these conditions hold with probability $\Omega(\alpha)$. The first condition holds with probability at least $\alpha$ over the choice of $x$. Conditioned on $x \in S_{\mathrm{g}}$, the expected number of non-neighbors that $x$ has in $S_{\mathrm{g}}$ is at most $\left|S_{\mathrm{g}}\right| / 2$. Thus, the probability that it has more than $3\left|S_{\mathrm{g}}\right| / 4$ non-neighbors is at most $2 / 3$ by Markov's inequality. Thus, the first two conditions both hold with probability at least $\alpha / 3$. Finally, the expected number of pairs of neighbors of $x$ that are non-neighbors in $G^{\prime}$ is at most $\delta|V|^{2}$ by Lemma 3.6. Thus, by Markov's inequality, there will be more than $(8 \delta / \alpha)|V|^{2}$ such non-connected neighbors with probability at most $\alpha / 8$. Combining with the above, all three conditions hold with probability at least $\alpha / 24$.
Given these assumptions, we note that $|W| \geq\left|S_{\mathrm{g}}\right| / 4 \geq(\alpha / 4)|V|$, and at most $(8 \delta / \alpha)|V|^{2}$ of pairs in $W$ are not connected in $G^{\prime}$. This implies that we pass the condition in Line 9. We will now verify the conditions in Lemma 3.7. Since $\beta:=|W| /|V| \geq \alpha / 4$ and $\gamma$, the number of pairs of vertices in $W$ that are not connected in $G^{\prime}$, is at most $8 \delta / \alpha=\alpha^{2} / 576$, we have $\gamma \leq \beta^{2} / 36$, satisfying the assumptions of Lemma 3.7. Thus the returned $W^{\prime}$ is a clique in $R_{\beta / 3}\left(G^{\prime}\right)$ and satisfies

$$
|W|-\left|W^{\prime}\right| \leq(6 \gamma / \beta)|V| \leq(48 \delta / \alpha) /(\alpha / 4)|V| \leq(\alpha / 24)|V|
$$

This means that $\left|W^{\prime} \cap S_{\mathrm{g}}\right| \geq\left|S_{\mathrm{g}}\right| / 4-(\alpha / 24)|V| \geq\left|S_{\mathrm{g}}\right| / 6$. On the other hand, we know that $G$ has ( $M, t$ ) bounded moments in $k$-sparse directions. Lemma 3.5 implies that $G^{\prime}$ has moments bounded by $O\left(2^{t} M / \alpha^{3}\right)$. Applying the lemma once more implies that $R_{\beta / 3}\left(G^{\prime}\right)$ has moments bounded by $O\left(4^{t} M /\left(\alpha^{3} \beta\right)\right)=O\left(4^{t} M / \alpha^{4}\right)$. Since $W^{\prime}$ is a clique in $R_{\beta / 3}\left(G^{\prime}\right)$, we have by Lemma 3.3 that if $\tilde{\mu}$ is the sample mean of $S_{\mathrm{g}} \cap W^{\prime}$, then

$$
\begin{equation*}
\langle v, \widehat{\mu}-\tilde{\mu}\rangle^{t} \leq O\left(4^{t} M \alpha^{-6}\right) \quad \text { for all } k \text {-sparse unit vectors } v . \tag{2}
\end{equation*}
$$

Since $\left|S_{\mathrm{g}} \cap W^{\prime}\right| \geq\left|S_{\mathrm{g}}\right| / 6$ and the $S_{\mathrm{g}}$ has bounded $t$-th moment along $k$-sparse directions, we have that $\left\langle v, \tilde{\mu}-\mu_{\text {good }}\right\rangle^{t} \leq O(M)$ (see Lemma B. 1 for a proof of this fact). Combining this with Equation (2) using triangle inequality completes the proof.

### 3.3 Proof Sketch of Theorem 1.2

We sketch the proof here, deferring the full proof to Appendix B.1. Let $S$ be the $(1-\alpha)$-corrupted set of samples, and $S_{\mathrm{g}}$ be the inliers. Let $T=\{x-y \mid x, y, \in S\}$ and $T_{\mathrm{g}}$ be its part due to inliers. To every subset $T^{\prime}$ of $T$, we can associate a graph $G_{T^{\prime}}$ having vertices $S$ and edges between the pairs included in $T^{\prime}$. Because of Lemma 2.4, $T_{\mathrm{g}}$ has certifiably bounded moments in $k$-sparse directions. By Theorem 3.1, the filtering step will return a subset $T^{\prime} \subset T$ that has sizable overlap with $T_{\mathrm{g}}$ and its graph $G_{T^{\prime}}$ has bounded moments in $k$-sparse directions. Finally, By Theorem 3.1, the rounding algorithm will return a $\hat{\mu}$ that is close to $\mu$ in all $k$-sparse directions. This $\hat{\mu}$ can be truncated to yield a vector close to $\mu$ in the standard $\ell_{2}$-norm (Fact A.2).

## 4 Acknowledgments

ID was supported by NSF Medium Award CCF-2107079, NSF Award CCF-1652862 (CAREER), a Sloan Research Fellowship, and a DARPA Learning with Less Labels (LwLL) grant. DK was supported by NSF Medium Award CCF-2107547, NSF Award CCF-1553288 (CAREER), a Sloan Research Fellowship, and a grant from CasperLabs. SK was supported by NSF under Grant \#2127309 to the Computing Research Association for the CIFellows 2021 Project. AP was supported by NSF grants NSF Award CCF-1652862 (CAREER), DMS-1749857, and CCF-1841190. TP was supported by NSF Award CCF-1652862 (CAREER) and NSF Medium Award CCF-2107079.

## References

[ABH ${ }^{+} 72$ ] D. F. Andrews, P. J. Bickel, F. R. Hampel, P. J. Huber, W. H. Rogers, and J. W. Tukey. Robust Estimates of Location: Survey and Advances. Princeton University Press, Princeton, NJ, USA, 1972.
[AK05] S. Arora and R. Kannan. Learning mixtures of separated nonspherical gaussians. The Annals of Applied Probability, 15(1A):69-92, 2005.
$\left[\mathrm{BBH}^{+} 21\right]$ M. Brennan, G. Bresler, S. Hopkins, J. Li, and T. Schramm. Statistical query algorithms and low degree tests are almost equivalent. In Conference on Learning Theory, 2021.
[BBV08] M. F. Balcan, A. Blum, and S. Vempala. A discriminative framework for clustering via similarity functions. In STOC, pages 671-680, 2008.
[BDLS17] S. Balakrishnan, S. S. Du, J. Li, and A. Singh. Computationally efficient robust sparse estimation in high dimensions. In Proc. 30th Annual Conference on Learning Theory, 2017.
[BK21] A. Bakshi and P. Kothari. List-decodable subspace recovery: Dimension independent error in polynomial time. In Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 1279-1297. SIAM, 2021.
[Bom98] I. M. Bomze. On standard quadratic optimization problems. Journal of Global Optimization, 13(4), 1998.
[BS16] B. Barak and D. Steurer. Proofs, beliefs, and algorithms through the lens of sum-ofsquares. 1, 2016.
[CGR16] M. Chen, C. Gao, and Z. Ren. A general decision theory for Huber's \$lepsilon\$contamination model. Electronic Journal of Statistics, 10(2):3752-3774, 2016.
[CMY20] Y. Cherapanamjeri, S. Mohanty, and M. Yau. List decodable mean estimation in nearly linear time. In 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS), pages 141-148. IEEE, 2020.
[CSV17] M. Charikar, J. Steinhardt, and G. Valiant. Learning from untrusted data. In Proc. 49th Annual ACM Symposium on Theory of Computing, pages 47-60, 2017.
[Das99] S. Dasgupta. Learning mixtures of Gaussians. In Proceedings of the 40th Annual Symposium on Foundations of Computer Science, pages 634-644, 1999.
[DG92] D. L. Donoho and M. Gasko. Breakdown properties of location estimates based on halfspace depth and projected outlyingness. Ann. Statist., 20(4):1803-1827, 121992.
[DK19] I. Diakonikolas and D. M. Kane. Recent advances in algorithmic high-dimensional robust statistics. arXiv preprint arXiv:1911.05911, 2019.
[DKK ${ }^{+}$16] I. Diakonikolas, G. Kamath, D. M. Kane, J. Li, A. Moitra, and A. Stewart. Robust estimators in high dimensions without the computational intractability. In Proc. 57th IEEE Symposium on Foundations of Computer Science (FOCS), pages 655-664, 2016.
[DKK ${ }^{+}$19] I. Diakonikolas, S. Karmalkar, D. M. Kane, E. Price, and A. Stewart. Outlier-robust high-dimensional sparse estimation via iterative filtering. In Advances in Neural Information Processing Systems 33, NeurIPS 2019, 2019.
[DKK20] I. Diakonikolas, D. M. Kane, and D. Kongsgaard. List-decodable mean estimation via iterative multi-filtering. Advances in Neural Information Processing Systems, 33:9312-9323, 2020.
[DKK $\left.{ }^{+} 21\right]$ I. Diakonikolas, D. M. Kane, D. Kongsgaard, J. Li, and K. Tian. List-decodable mean estimation in nearly-pca time. Advances in Neural Information Processing Systems, 34, 2021.
[DKK ${ }^{+} 22 \mathrm{a}$ ] I. Diakonikolas, D. M. Kane, S. Karmalkar, A. Pensia, and T. Pittas. Robust sparse mean estimation via sum of squares. 35th Annual Conference on Learning Theory (COLT), 2022.
$\left[\mathrm{DKK}^{+} 22 \mathrm{~b}\right]$ I. Diakonikolas, D. M. Kane, D. Kongsgaard, J. Li, and K. Tian. Clustering mixture models in almost-linear time via list-decodable mean estimation. In Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, pages 1262-1275, 2022.
[DKP ${ }^{+}$21] I. Diakonikolas, D. M. Kane, A. Pensia, T. Pittas, and A. Stewart. Statistical query lower bounds for list-decodable linear regression. Advances in Neural Information Processing Systems, 34, 2021.
[DKS17] I. Diakonikolas, D. M. Kane, and A. Stewart. Statistical query lower bounds for robust estimation of high-dimensional Gaussians and Gaussian mixtures. In Proc. 58th IEEE Symposium on Foundations of Computer Science (FOCS), pages 73-84, 2017.
[DKS18a] I. Diakonikolas, D. M. Kane, and A. Stewart. List-decodable robust mean estimation and learning mixtures of spherical gaussians. In Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, pages 1047-1060, 2018. Full version available at https://arxiv.org/abs/1711.07211.
[DKS18b] I. Diakonikolas, D. M. Kane, and A. Stewart. List-decodable robust mean estimation and learning mixtures of spherical gaussians. In Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, pages 1047-1060, 2018.
[DKS18c] I. Diakonikolas, D. M. Kane, and A. Stewart. List-decodable robust mean estimation and learning mixtures of spherical Gaussians. In Proc. 50th Annual ACM Symposium on Theory of Computing (STOC), pages 1047-1060, 2018.
[DL88] D. L. Donoho and R. C. Liu. The "Automatic" Robustness of Minimum Distance Functionals. The Annals of Statistics, 16(2):552-586, 1988.
[DS07] S. Dasgupta and L. J. Schulman. A probabilistic analysis of em for mixtures of separated, spherical gaussians. Journal of Machine Learning Research, 8:203-226, 2007.
[EK12] Y. C. Eldar and G. Kutyniok. Compressed sensing: theory and applications. Cambridge university press, 2012.
$\left[\mathrm{FGR}^{+} 13\right]$ V. Feldman, E. Grigorescu, L. Reyzin, S. Vempala, and Y. Xiao. Statistical algorithms and a lower bound for detecting planted cliques. In Proceedings of STOC'13, pages 655-664, 2013. Full version in Journal of the ACM, 2017.
[FKP19] N. Fleming, P. Kothari, and T. Pitassi. Semialgebraic proofs and efficient algorithm design. Foundations and Trends ${ }^{\circledR}$ in Theoretical Computer Science, 14(1-2):1-221, 2019.
[GLS81] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. Combinatorica, 1(2):169-197, 1981.
[Hop18] S. B. Hopkins. Clustering and sum of squares proofs: Six blog posts on unsupervised learning. 2018.
[HR09] P. J. Huber and E. M. Ronchetti. Robust statistics. Wiley New York, 2009.
[HRRS86] F. R. Hampel, E. M. Ronchetti, P. J. Rousseeuw, and W. A. Stahel. Robust statistics. The approach based on influence functions. Wiley New York, 1986.
[HTW15] T. Hastie, R. Tibshirani, and M. Wainwright. Statistical Learning with Sparsity: The Lasso and Generalizations. Chapman \& Hall/CRC, 2015.
[Kea98] M. Kearns. Efficient noise-tolerant learning from statistical queries. Journal of the ACM, 45(6):983-1006, 1998.
[KK10] A. Kumar and R. Kannan. Clustering with spectral norm and the k-means algorithm. In 2010 IEEE 51st Annual Symposium on Foundations of Computer Science, pages 299-308. IEEE, 2010.
[KKK19] S. Karmalkar, A. Klivans, and P. Kothari. List-decodable linear regression. In Advances in Neural Information Processing Systems, pages 7423-7432, 2019.
[KKM18] A. Klivans, P. Kothari, and R. Meka. Efficient algorithms for outlier-robust regression. In Proc. 31st Annual Conference on Learning Theory (COLT), pages 1420-1430, 2018.
[KS17] P. K. Kothari and J. Steinhardt. Better agnostic clustering via relaxed tensor norms. CoRR, abs/1711.07465, 2017.
[Las01] J. B. Lasserre. New positive semidefinite relaxations for nonconvex quadratic programs. In Advances in Convex Analysis and Global Optimization, pages 319-331. Springer, 2001.
[Li17] J. Li. Robust sparse estimation tasks in high dimensions. CoRR, abs/1702.05860, 2017.
[LRV16] K. A. Lai, A. B. Rao, and S. Vempala. Agnostic estimation of mean and covariance. In Proc. 57th IEEE Symposium on Foundations of Computer Science (FOCS), pages 665-674, 2016.
[MV18] M. Meister and G. Valiant. A data prism: Semi-verified learning in the small-alpha regime. In Conference On Learning Theory, COLT 2018, pages 1530-1546, 2018.
[Nes00] Y. Nesterov. Squared functional systems and optimization problems. In High performance optimization, pages 405-440. Springer, 2000.
[RV17] O. Regev and A. Vijayaraghavan. On learning mixtures of well-separated gaussians. In 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), pages 85-96. IEEE, 2017.
[RY20a] P. Raghavendra and M. Yau. List decodable learning via sum of squares. In Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 161180. SIAM, 2020.
[RY20b] P. Raghavendra and M. Yau. List decodable subspace recovery. In Conference on Learning Theory, COLT 2020, volume 125 of Proceedings of Machine Learning Research, pages 3206-3226. PMLR, 2020.
[SCV18] J. Steinhardt, M. Charikar, and G. Valiant. Resilience: A criterion for learning in the presence of arbitrary outliers. In Proc. 9th Innovations in Theoretical Computer Science Conference (ITCS), pages 45:1-45:21, 2018.
[Sho87] N.Z. Shor. Quadratic optimization problems. Soviet Journal of Computer and Systems Sciences, 1987.
[SKL17] J. Steinhardt, P. Wei Koh, and P. S. Liang. Certified defenses for data poisoning attacks. In Advances in Neural Information Processing Systems 30, pages 3520-3532, 2017.
[ST21] D. Steurer and S. Tiegel. Sos degree reduction with applications to clustering and robust moment estimation. In Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 374-393. SIAM, 2021.
[SVC16] J. Steinhardt, G. Valiant, and M. Charikar. Avoiding imposters and delinquents: Adversarial crowdsourcing and peer prediction. Advances in Neural Information Processing Systems, 29, 2016.
[Tuk60] J. W. Tukey. A survey of sampling from contaminated distributions. Contributions to probability and statistics, 2:448-485, 1960.
[van16] S. van de Geer. Estimation and Testing Under Sparsity. Number 2159 in École d'Été de Probabilités de Saint-Flour. Springer, 1st ed. 2016 edition, 2016.
[VW04] S. Vempala and G. Wang. A spectral algorithm for learning mixture models. Journal of Computer and System Sciences, 68(4):841-860, 2004.
[Yat85] Y. G. Yatracos. Rates of convergence of minimum distance estimators and Kolmogorov's entropy. Annals of Statistics, 13:768-774, 1985.
[ZS22] Shiwei Zeng and Jie Shen. List-decodable sparse mean estimation, 2022.

## Checklist

1. For all authors...
(a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
(b) Did you describe the limitations of your work? [Yes]
(c) Did you discuss any potential negative societal impacts of your work? [N/A]
(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
(a) Did you state the full set of assumptions of all theoretical results? [Yes]
(b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
(a) If your work uses existing assets, did you cite the creators? [N/A]
(b) Did you mention the license of the assets? [N/A]
(c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

[^0]:    ${ }^{1}$ Informally speaking, we say that the minimax optimal rate is $\gamma$ if (i) no algorithm (regardless of sample size and runtime) has error $o(\gamma)$ with a list of size independent of dimension $n$, and (ii) there is an algorithm with error $O(\gamma)$ with a list of size independent of $n$; in our case, poly $(n / \alpha)$ samples and $O(1 / \alpha)$ list size suffice.

[^1]:    ${ }^{2}$ It suffices for $D$ to have poly $(t \log (n))$ bounded moments in axis-aligned directions; see Section 2.

