
High-dimensional Robust Mean Estimation via Gradient Descent

Yu Cheng^{*1} Ilias Diakonikolas^{*2} Rong Ge^{*3} Mahdi Soltanolkotabi^{*4}

Abstract

We study the problem of high-dimensional robust mean estimation in the presence of a constant fraction of adversarial outliers. A recent line of work has provided sophisticated polynomial-time algorithms for this problem with dimension-independent error guarantees for a range of natural distribution families. In this work, we show that a natural non-convex formulation of the problem can be solved directly by gradient descent. Our approach leverages a novel structural lemma, roughly showing that any approximate stationary point of our non-convex objective gives a near-optimal solution to the underlying robust estimation task. Our work establishes an intriguing connection between algorithmic high-dimensional robust statistics and non-convex optimization, which may have broader applications to other robust estimation tasks.

1. Introduction

Learning in the presence of outliers is an important goal in machine learning that has become a pressing challenge in a number of high-dimensional data analysis applications, including data poisoning attacks (Barreno et al., 2010; Biggio et al., 2012; Steinhardt et al., 2017) and exploratory analysis of real datasets with natural outliers, e.g., in biology (Rosenberg et al., 2002; Paschou et al., 2010; Li et al., 2008). In both these application domains, the outliers are not “random” but can be arbitrarily correlated, and could exhibit rather complex structures that is essentially impossible to accurately model. Hence, the goal in these settings is to design computationally efficient estimators that can tolerate

^{*}Equal contribution ¹University of Illinois at Chicago, Chicago, IL, USA ²University of Wisconsin-Madison, Madison, WI, USA ³Duke University, Durham, NC, USA ⁴University of Southern California, Los Angeles, CA, USA. Correspondence to: Yu Cheng <yucheng2@uic.edu>, Ilias Diakonikolas <ilias@cs.wisc.edu>, Rong Ge <rongge@cs.duke.edu>, Mahdi Soltanolkotabi <soltanol@usc.edu>.

a small constant fraction of arbitrary outliers.

Throughout this paper, we focus on the following data contamination model that generalizes several existing models, including Huber’s contamination model (Huber, 1964).

Definition 1.1 (Strong Contamination Model). *Given a parameter $0 < \epsilon < 1/2$ and a distribution family \mathcal{D} on \mathbb{R}^d , the adversary operates as follows: The algorithm specifies the number of samples N , and N samples are drawn from some unknown $D \in \mathcal{D}$. The adversary is allowed to inspect the samples, remove up to ϵN of them and replace them with arbitrary points. This modified set of N points is then given as input to the algorithm. We say that a set of samples is ϵ -corrupted if it is generated by the above process.*

The parameter ϵ in the above definition is the fraction of corrupted samples and quantifies the power of the adversary. Intuitively, among our samples, an unknown $(1 - \epsilon)$ fraction are generated from a distribution of interest and are called *inliers*, and the rest are called *outliers*.

The statistical foundations of outlier-robust estimation were laid out in early work by the robust statistics community, starting with the pioneering works of Tukey (1960) and Huber (1964). In contrast, until fairly recently, even the most basic algorithmic questions were poorly understood. Specifically, even for the basic task of high-dimensional mean estimation, all known robust estimators had runtime exponential in the dimension, rendering them ineffective in high-dimensional settings.

Recently, Diakonikolas et al. (2016); Lai et al. (2016) gave the first efficiently computable robust estimators for high-dimensional unsupervised learning tasks, including mean and covariance estimation. Specifically, Diakonikolas et al. (2016) obtained the first polynomial-time robust estimators with *dimension-independent* error guarantees, i.e., with error scaling only with the fraction of corrupted samples ϵ and not with the dimensionality of the data. Since the dissemination of these works, there has been a flurry of research activity on algorithmic aspects of high-dimensional robust statistics; see, e.g., Diakonikolas & Kane (2019) for a recent survey on the topic.

Despite this exciting progress, the design of efficient robust estimators in high dimensions remains challenging. The difficulty, of course, lies in the non-convexity of the under-

lying optimization problem. Prior work developed fairly sophisticated algorithmic tools, even for the task of robust mean estimation. These include convex relaxations (Diakonikolas et al., 2016) and quite subtle iterative spectral methods (Diakonikolas et al., 2016; Lai et al., 2016).

A natural and important goal is to understand to what extent such sophisticated methods are indeed necessary or whether much simpler robust learning algorithms exist. In this work, we take a direct optimization view of these problems and ask the following general question:

Is it possible to solve robust estimation tasks by standard first-order methods?

We believe that this question merits investigation in its own right. Moreover, its positive resolution may have significant implications in the practical adoption of robust estimation methods. Particularly so since prior algorithms are either (1) computationally prohibitive (relying on large convex relaxations), (2) involve carefully crafted parameters that require precise tuning for practical deployment, or (3) are challenging to extend to more sophisticated robust estimation tasks. A tantalizing possibility is the following: For a range of high-dimensional robust estimation tasks, there exists a (natural) non-convex formulation such that gradient descent efficiently converges to a near-optimal solution.

In this paper, we show that this premise is true for the task of high-dimensional robust mean estimation. In robust mean estimation, we are given a set of N ϵ -corrupted samples from an unknown distribution D in a known family \mathcal{D} , and we want to output a hypothesis vector $\hat{\mu}$ such that $\|\hat{\mu} - \mu^*\|_2$ is as small as possible, where μ^* is the mean of D . For simplicity, we will assume in this discussion that D is an unknown mean and identity covariance Gaussian on \mathbb{R}^d . We note that our results hold under more general distributional assumptions, as in Diakonikolas et al. (2016; 2017a).

The goal in robust mean estimation is to develop efficient algorithms whose ℓ_2 -error guarantee scales only with ϵ and not with the dimension d . In particular, for the identity covariance Gaussian case, Diakonikolas et al. (2016) gave polynomial-time algorithms for the problem that use $N = \tilde{\Omega}(d/\epsilon^2)$ samples and guarantee error $O(\epsilon\sqrt{\log(1/\epsilon)})$. This error guarantee matches known Statistical Query (SQ) lower bounds (Diakonikolas et al., 2017b).

1.1. Overview of Results and Contributions

In this paper, we consider a natural non-convex optimization formulation of high-dimensional robust mean estimation, and show that gradient descent¹ efficiently converges to a near-optimal solution. Specifically, we show that gradient

¹Throughout, we informally use the term ‘‘gradient descent’’ to refer to variations of gradient descent methods, which involve

descent converges in a polynomial number of iterations and matches the error guarantee of the best known polynomial-time algorithms for the problem. Our technical contribution lies in showing that *any* approximate stationary point of our non-convex objective suffices – in the sense that it gives a near-optimal solution for the underlying estimation problem.

To describe our non-convex formulation, we require some background. We use the following framework for robust mean estimation, introduced in Diakonikolas et al. (2016). The idea is to assign a non-negative weight to each data point and then find an appropriate combination of weights such that the weighted empirical mean is close to the true mean. The constraint on the chosen weights is that they represent at least a $(1 - \epsilon)$ -density fractional subset of the dataset. More formally, given datapoints $X_1, \dots, X_N \in \mathbb{R}^d$ with corresponding data matrix $X \in \mathbb{R}^{d \times N}$, the objective is to find a weight vector $w \in \mathbb{R}^N$ such that $\mu_w = Xw$ is close to μ^* . The constraint on w is that it belongs in the set

$$\Delta_{N,\epsilon} = \left\{ w \in \mathbb{R}^N : \|w\|_1 = 1 \text{ and } 0 \leq w_i \leq \frac{1}{(1-\epsilon)N} \forall i \right\},$$

which is the convex hull of all uniform distributions over subsets $S \subseteq [N]$ of size $|S| = (1 - \epsilon)N$.

Diakonikolas et al. (2016) established a key structural lemma (Lemma 2.1), which formed the basis of their algorithms. Roughly speaking, the lemma states that any weight vector w is a good solution if the spectral norm of the weighted empirical covariance, $\Sigma_w = \sum_{i=1}^N w_i (X_i - \mu_w)(X_i - \mu_w)^\top$, is small. This lemma directly motivates the following non-convex optimization formulation:

$$\text{Min } \|\Sigma_w\|_2 \text{ subject to } w \in \Delta_{N,2\epsilon} \quad (1)$$

It follows from the aforementioned structural lemma that a near-optimal solution w to (1) gives an μ_w that is close to μ^* . The challenge is that the objective function is not convex, hence it is unclear how to efficiently optimize. Faced with this difficulty, prior works on the topic (Diakonikolas et al., 2016; 2017a) developed various sophisticated algorithms.

In this paper, we work directly with the natural formulation (1). Despite its non-convexity, we are able to leverage the structure of the problem to show that gradient descent efficiently converges to a good vector w . In more detail, we prove a novel result about the structure of approximate stationary points of this objective.

Theorem 1.2 (informal statement). *Any approximate stationary point w of (1) defines an μ_w that is close to μ^* .*

See Theorem 3.1 for a detailed formal statement. Technically speaking, our statement is more subtle for various reasons, including the fact that the objective function is not

updates based on a generalized notion of a gradient, e.g., sub-gradient for non-differentiable functions.

differentiable and the domain is constrained. As a result, we require a careful definition of stationarity in our setting.

Given Theorem 1.2, we proceed to show that projected sub-gradient descent converges to an approximate stationary point in a polynomial number of iterations. This step is also somewhat intricate as the function is non-convex, non-smooth and the optimization problem (1) involves constraints. In summary, we establish the following theorem:

Theorem 1.3. *After $\tilde{O}(N^2d^4)$ iterations, projected sub-gradient descent on (1) outputs a point w such that with high probability $\|\mu_w - \mu^*\|_2 = O(\epsilon\sqrt{\log(1/\epsilon)})$.*

The bound we establish on the convergence rate on the spectral norm objective (1) is polynomially bounded, but relatively slow. Our second main contribution involves considering the “softmax” version of the spectral norm, which has better smoothness properties. An analogous lemma about the structure of stationary points allows us to show a faster rate of convergence for this modified objective.

Theorem 1.4. *After $\tilde{O}(Nd^3/\epsilon)$ iterations, projected gradient descent on the softmax objective outputs a point w such that with high probability $\|\mu_w - \mu^*\|_2 = O(\epsilon\sqrt{\log(1/\epsilon)})$.*

As evident from the above result, the additional smoothness of the “softmax” objective allows us to establish a significantly improved bound on the number of iterations.

1.2. Related Work

The algorithmic question of designing efficient robust mean estimators in high-dimensions has been extensively studied in recent years. After the initial papers (Diakonikolas et al., 2016; Lai et al., 2016), a number of works (Diakonikolas et al., 2017a; Steinhardt et al., 2018; Cheng et al., 2018; Dong et al., 2019; Depersin & Lecue, 2019; Cheng et al., 2019) have obtained algorithms with improved asymptotic worst-case runtimes that work under weaker distributional assumptions on the good data. Moreover, efficient high-dimensional robust mean estimators have been used as primitives for robustly solving a range of machine learning tasks that can be expressed as stochastic optimization problems (Prasad et al., 2018; Diakonikolas et al., 2019a).

We compare our approach with the works of Cheng et al. (2018) and Dong et al. (2019) that give the asymptotically fastest known algorithms for robust mean estimation. At a high-level, Cheng et al. (2018), building on the convex programming relaxation of Diakonikolas et al. (2016), proposed a primal-dual approach for robust mean estimation that reduces the problem to a poly-logarithmic number of packing and covering SDPs. Each such SDP is known to be solvable in time $O(Nd)$, using mirror descent (Allen-Zhu et al. (2016); Peng et al. (2016)). Dong et al. (2019) build on the iterative spectral approach of Diakonikolas et al. (2016). That work uses the matrix multiplicative weights

update method with a specific regularization and dimension-reduction to improve the worst-case runtime.

In contrast to all of the above, we use a natural non-convex formulation of the robust mean estimation task, and show that a standard first-order method provably and efficiently converges to a near-optimal solution. Even though the convergence rates that we establish in this work do not yield the fastest known asymptotic runtimes for the problem, we believe that our approach is conceptually interesting for a number of reasons. First, our theorem regarding stationary points provides novel structural understanding about robust mean estimation and can be viewed as an explanation as to why this problem is polynomially solvable. Second, it is plausible that gradient descent applied in this context is more stable than previously known algorithms and may facilitate the adoption of robust estimation methods in practice. We hope that this work will serve as the starting point for solving other robust estimation tasks via first-order methods.

Finally, we note that there is an increasing literature on developing rigorous guarantees for non-convex optimization problems via gradient descent, e.g., see the recent survey (Jain & Kar, 2017) for a review of this literature. With a few exceptions (Loh & Wainwright, 2011; Hassani et al., 2017), this literature mostly focuses on showing that gradient descent converges to a global optimum starting from a spectral (Keshavan et al., 2010; Candes et al., 2015; Tu et al., 2015) or random initialization (Ge et al., 2015) in settings where there are no bad local optima. In contrast to most of this literature, in this paper we show that any *stationary* point has good approximation properties so that no specialized or random initialization is necessary. We believe that such a perspective may enable rigorous analysis of many other non-convex optimization problems.

1.3. Roadmap

In Section 2, we set up the necessary notation and provide some background on robust mean estimation. In the next two sections, we focus on the spectral norm objective. In Section 3, we prove our main structural result showing that any stationary point of the spectral norm objective yields a good solution. We also extend this result in Appendix B, showing that in fact, any approximate stationary point yields a sufficiently good solution. In Section 4, we show that gradient descent converges to an approximate stationary point and hence yields a good solution in a polynomial number of iterations. In Appendix C, we prove structural and algorithmic results for the softmax objective, showing that any approximate stationary point of the softmax objective yields a good solution, and we can find an approximate stationary point using projected gradient descent in a polynomial number of iterations. We conclude with future directions in Section 5.

2. Preliminaries and Background

Notation. For $N \in \mathbb{Z}_+$, we denote $[N] := \{1, \dots, N\}$. For a vector x , we use $\|x\|_1$, $\|x\|_2$, and $\|x\|_\infty$ to denote the ℓ_1 , ℓ_2 , and ℓ_∞ norm of x respectively. For a matrix A , we use $\|A\|_2$ to denote the spectral norm of A .

For two vectors $x, y \in \mathbb{R}^n$, we use $x^\top y = \sum_{i=1}^n x_i y_i$ to denote the inner product of x and y , and we use $x \odot y \in \mathbb{R}^n$ to denote entrywise product of x and y . For a vector $x \in \mathbb{R}^n$, let $\text{diag}(x) \in \mathbb{R}^{n \times n}$ denote a diagonal matrix with x on the diagonal. For a matrix $A \in \mathbb{R}^{n \times n}$, let $\text{diag}(A) \in \mathbb{R}^n$ denote a column vector with the diagonal entries of A .

Let I denote the identity matrix. For a matrix $A \in \mathbb{R}^{n \times n}$, let $\text{tr}(A)$ denote the trace of A . For two matrices A and B of the same dimensions, let $A \bullet B = \langle A, B \rangle = \text{tr}(A^\top B)$ be the entry-wise inner product of A and B . We use $\exp(A)$ to denote the matrix exponential of A .

A symmetric matrix $A \in \mathbb{R}^{n \times n}$ is said to be positive semidefinite (PSD) if $x^\top A x \geq 0$ for all $x \in \mathbb{R}^n$. For two symmetric matrices A and B , we write $A \preceq B$ iff the matrix $B - A$ is positive semidefinite. Let $\Delta_{n \times n}$ be the set of all PSD matrices of trace 1.

Framework. We use N for the number of input samples, d for the dimension of the ground-truth distribution, and ϵ for the fraction of corrupted samples. Given N datapoints $X_1, \dots, X_N \in \mathbb{R}^d$, we use $X \in \mathbb{R}^{d \times N}$ to denote the sample matrix, where the i -th column of X is X_i .

Given $w \in \mathbb{R}^N$, let $\mu_w = Xw = \sum_{i=1}^N w_i X_i$ denote the weighted empirical mean and let $\Sigma_w = \sum_{i=1}^N w_i (X_i - \mu_w)(X_i - \mu_w)^\top$ denote the weighted empirical covariance. Let $\Delta_{N, \epsilon}$ denote the convex hull of all uniform distributions over subsets $S \subseteq [N]$ of size $|S| = (1 - \epsilon)N$:

$$\Delta_{N, \epsilon} = \left\{ w \in \mathbb{R}^N : \|w\|_1 = 1 \text{ and } 0 \leq w_i \leq \frac{1}{(1-\epsilon)N} \forall i \right\}.$$

Every weight vector $w \in \Delta_{N, \epsilon}$ corresponds to a fractional set of $(1 - \epsilon)N$ samples.

Background on Robust Mean Estimation. As mentioned in the introduction, our non-convex formulation is directly motivated by the following structural lemma:

Lemma 2.1 (Diakonikolas et al. (2016)). *Let S be an ϵ -corrupted set of $N = \tilde{\Omega}(d/\epsilon^2)$ samples from an unknown $\mathcal{N}(\mu^*, I)$ and $w \in \Delta_{N, 2\epsilon}$. If $\lambda_{\max}(\Sigma_w) \leq 1 + \delta$, for some $\delta \geq 0$, then with high probability, we have that $\|\mu^* - \mu_w\|_2 = O(\sqrt{\epsilon\delta} + \epsilon\sqrt{\log(1/\epsilon)})$.*

As in prior work, we will establish correctness for our algorithms under deterministic conditions on the inliers (good samples) that hold with high probability. Let G^* denote the original set of N good samples. Let $S = G \cup B$ denote

the input samples after the adversary replaced ϵ -fraction of the samples, where $G \subset G^*$ is the set of remaining good samples and B is the set of bad samples (outliers) added by the adversary. Note that $|G| = (1 - \epsilon)N$ and $|B| = \epsilon N$. Given $w \in \mathbb{R}^N$, let $w_G = \sum_{i \in G} w_i$ be the total weight on good samples, and w_B be the total weight on bad samples.

We require the following concentration bounds to hold for the original N good samples G^* (which happens with high probability when $N = \tilde{\Omega}(d/\epsilon^2)$). For all $\hat{w} \in \Delta_{N, 3\epsilon}$, we require the following condition to hold for $\delta = O(\epsilon \log(1/\epsilon))$:

$$\left\| \sum_{i \in G^*} \hat{w}_i (X_i - \mu^*)(X_i - \mu^*)^\top - I \right\|_2 \leq \delta. \quad (2)$$

Condition (2) on *original* samples G^* implies the following conditions on the *remaining* good samples G . For any weight vector $w \in \Delta_{N, 2\epsilon}$ on the ϵ -corrupted set of samples $S = G \cup B$:

$$\left\| \sum_{i \in G} w_i (X_i - \mu^*)(X_i - \mu^*)^\top - I \right\|_2 \leq \delta. \quad (3)$$

This is because we can define \hat{w} as follows: $\hat{w}_i = \frac{w_i}{w_G}$ for all $i \in G$ and $\hat{w}_i = 0$ for all $i \in B$. Since $w \in \Delta_{N, 2\epsilon}$, we have $\|\hat{w}\|_\infty \leq \frac{\|w\|_\infty}{w_G} = \frac{\|w\|_\infty}{1 - w_B} \leq \frac{\|w\|_\infty}{1 - |B| \cdot \|w\|_\infty} \leq \frac{1}{(1 - 3\epsilon)N}$. In other words, $\hat{w} \in \Delta_{N, 3\epsilon}$ and Condition (3) follows directly from Condition (2).

Remark 2.2 (Distributional Assumptions). For simplicity, in this paper we focus on the fundamental setting that the good data are drawn from an unknown mean and identity covariance Gaussian distribution. It should be noted that our structural and algorithmic results hold under more general distributional assumptions. Specifically, Theorem 4.1 immediately applies to identity covariance subgaussian distributions, with the same error guarantees, since it only relies on the concentration bounds (2) and (3) that only require subgaussian tails (see, e.g., (Diakonikolas et al., 2017a).) Moreover, one can modify the proof of our structural results (Theorems 3.1 and 3.2), mutatis-mutandis, to apply (1) for distributions with bounded covariance (i.e., $\Sigma \preceq I$) and match the optimal $O(\sqrt{\epsilon})$ approximation to the mean (Diakonikolas et al., 2017a); and, (2) more generally, under the (ϵ, δ) -stability condition of (Diakonikolas & Kane, 2019) to yield an $O(\delta)$ ℓ_2 -approximation to the mean.

Background and Definitions of Stationarity. Note that the spectral norm is not a differentiable function and therefore we need an alternative definition of stationarity. To address this issue, by the definition of spectral norm, we can define a function $F(w, u) = u^\top \Sigma_w u$ that takes two parameters as input: the weights $w \in \mathbb{R}^N$ and a unit vector $u \in \mathbb{R}^d$. Our non-convex objective $\min_w f(w) :=$

$\|\Sigma_w\|_2$ is then equivalent to solving the minimax problem $\min_w \max_u F(w, u)$. The function $\max_u F(w, u)$ is weakly-convex, and we use the following stationary point definition that is common in the weakly-convex optimization literature (Rockafellar, 1970; 1981; Drusvyatskiy, 2017; Davis & Drusvyatskiy, 2018; Jin et al., 2019).

Definition 2.3 (First-order stationary point). *Let $F(w, u)$ be a function that is differentiable with respect to w for all u . Let $f(w) = \max_u F(w, u)$. Consider the constrained optimization problem $\min_{w \in K} f(w)$, where K is a closed convex set. We say that $w \in K$ is a first-order stationary point if there exists some $u \in \arg \max_v F(w, v)$ such that*

$$(\nabla_w F(w, u))^\top (\tilde{w} - w) \geq 0 \text{ for all } \tilde{w} \in K .$$

We also need a notion of an *approximate* stationary point in the sense that the updates from one iteration to the next do not change much. In the unconstrained and differentiable case, such a point can be characterized by the gradient being small. However, the objective function we consider is both non-differentiable and has constraints, so that a proper definition of approximate stationarity is much more subtle. To overcome this, we appeal to tools from conic geometry and notions of stationarity for weakly convex functions (Rockafellar, 1970; 1981; Drusvyatskiy, 2017; Davis & Drusvyatskiy, 2018) to define an appropriate notion of approximate stationarity.

To discuss the notion of approximate stationarity that we use, we need to work with a smoothed variant of the objective known as the *Moreau envelope*.

Definition 2.4 (Moreau envelope). *For any function f and closed convex set \mathcal{K} , its associated Moreau envelope $f_\beta(w)$ is defined to be the function*

$$f_\beta(w) := \min_{\tilde{w} \in \mathcal{K}} f(\tilde{w}) + \beta \|w - \tilde{w}\|_2^2 .$$

The Moreau envelope can be thought of as a form of convolution between the original function f and a quadratic, so as to smoothen the landscape. In particular, when $f(w)$ takes the form of a maximization problem ($f(w) = \max_u F(w, u)$) with F a mapping that is β -smooth in the u parameter ($|\nabla_w F(w, \tilde{u}) - \nabla_w F(w, u)| \leq \beta \|\tilde{u} - u\|_2$), the Moreau envelope is also β -smooth (Drusvyatskiy, 2017). Therefore, the approximate stationarity of the Moreau envelope can be easily defined through its gradient allowing us to define the following notion of approximate stationarity.

Definition 2.5 (Approximate first-order stationary point). *For any function f and closed convex set \mathcal{K} consider its associated Moreau envelope $f_\beta(w)$ per Definition 2.4. we say that a point w is a ρ -approximately stationary point if $\|\nabla f_\beta(w)\|_2 \leq \rho$.*

As mentioned earlier, the spectral norm admits a minimax formulation of the form $f(w) = \max_u F(w, u)$. Further-

more, as detailed in Appendix B, the corresponding function $F(w, u)$ is β -smooth with $\beta = 2\|X\|_2^2$, so that this notion of approximate stationarity can be applied to the objective of interest in this paper.

3. Structural Result: Any Approximate Stationary Point Suffices

In this section, we establish our main structural result, which says that every *approximate* stationary point of (1) must give a μ_w that is close to μ^* . For simplicity of the exposition, in the main body of this paper, we state and prove a simpler theorem showing that every (exact) stationary point is a good solution.

Theorem 3.1 (Any stationary point is a good solution). *Let S denote an ϵ -corrupted set of N samples drawn from a d -dimensional Gaussian $\mathcal{N}(\mu^*, I)$ with unknown mean μ^* . Suppose that S satisfies Lemma 2.1 and Condition (3).*

Let $f(w)$ be the objective function defined in Equation (1). For any first-order stationary point $w \in \Delta_{N, 2\epsilon}$ of $f(w)$, we have $\|\mu_w - \mu^\|_2 = O(\epsilon\sqrt{\log(1/\epsilon)})$.*

We note that while Theorem 3.1 shows that any (exact) stationary point has small objective value, a stronger statement is required for our algorithmic results in the next section. Specifically, we require that any *approximate* stationary point — in the sense of Definition 2.5 — which gradient descent efficiently converges to, also has low objective value. This is accomplished in the next theorem which we prove in Appendix B. Specifically, by appealing to the gradient of the Moreau envelope from Definition 2.4, we extend the proof of Theorem 3.1 to show the following:

Theorem 3.2 (Any approximate stationary point suffices). *Consider the same setting as in Theorem 3.1. Consider the spectral norm objective $f(w) = \|\Sigma_w\|_2$ with $f_\beta(w)$ denoting the corresponding Moreau envelope function per Definition 2.4 with $\beta = 2\|X\|_2^2$. Then, for any $w \in \Delta_{N, 2\epsilon}$ satisfying*

$$\|\nabla f_\beta(w)\|_2 = O(\log(1/\epsilon)) ,$$

we have $\|\mu_w - \mu^\|_2 = O(\epsilon\sqrt{\log(1/\epsilon)})$.*

In the remainder of this section, we focus on proving Theorem 3.1 and briefly discuss how this proof can be generalized to prove Theorem 3.2. Our proof is carried out in two steps: (1) We establish a structural lemma which states that every stationary point w must satisfy a *bimodal subgradient* property; (2) We show any point satisfying such property must have a small objective value. Given these two steps, we can conclude any stationary points μ_w is close to μ^* , by Lemma 2.1.

For the first step, the bimodal subgradient property states that there exists a vector $\nu \in \partial f(w)$ (in the sub-gradient

of the function at that stationary point) whose entries divided in two groups of indices such that for any $i \in S^-$ and any $j \in S^+$ we have $\nu_i \leq \nu_j$. Intuitively, S^- contains all indices with positive w_i , so they can potentially be decreased; while S^+ contains all indices with $w_i < \frac{1}{(1-2\epsilon)N}$, so they can potentially be increased. If the bimodal sub-gradient property is violated, there must be indices $i \in S^-$, $j \in S^+$, where $\nu_i > \nu_j$. In this case, decreasing w_i and increasing w_j would decrease the objective and thus violate stationarity.

For the second step, recall that

$$\Sigma_w = (X \text{diag}(w)X^\top - Xww^\top X^\top)$$

and $F(w, u) = u^\top \Sigma_w u$. Let us first compute the sub-gradient $\nabla_w F(w, u)$ with respect to a vector u :

$$\nabla_w F(w, u) = X^\top u \odot X^\top u - 2(u^\top Xw)X^\top u. \quad (4)$$

Our key observation is that the sub-gradient at direction u is equivalent to the gradient of w for the one-dimensional problem with input $(X_i^\top u)_{i=1}^N$. This allows us to effectively reduce our problem to a one-dimensional robust mean estimation problem. This reduction allows us to show that when the objective function is large, then there must be some non-zero weights associated with the corrupted points that are far away from the mean (these points will be in S^-); while on the other hand, S^+ must contain at least ϵ -fraction of the good points. One can then select indices from these two sets to violate the bimodal sub-gradient property.

Fix a first-order stationary point $w \in \Delta_{N,2\epsilon}$. Definition 2.3 implies that there is a corresponding unit vector $u \in \mathbb{R}^d$ such that w is a stationary point of $F(w, u)$. We first state the bimodal sub-gradient property.

Lemma 3.3 (Bimodal sub-gradient property at stationarity). *Fix $w \in \Delta_{N,2\epsilon}$ and a unit vector u with $u^\top \Sigma_w u = \|\Sigma_w\|_2$. Let $S_- = \{i : w_i > 0\}$ and $S_+ = \{i : w_i < \frac{1}{(1-2\epsilon)N}\}$ denote the coordinates of w that can decrease and increase respectively. If w is a first-order stationary point of $F(w, u)$, then*

$$\nabla_w F(w, u)_i \leq \nabla_w F(w, u)_j,$$

for all $i \in S_-$ and $j \in S_+$.

Proof. Suppose there is some $i \in S_-$ and $j \in S_+$ such that $\nabla_w F(w, u)_i > \nabla_w F(w, u)_j$, then intuitively we can make $f(w)$ smaller by decreasing w_i and increasing w_j . Formally, let $\tilde{w} = w + \min(w_i, \frac{1}{(1-2\epsilon)N} - w_j)(e_j - e_i)$ where e_i is the i -th basis vector. We have $\tilde{w} \in \Delta_{N,2\epsilon}$ and $(\nabla_w F(w, u))^\top (\tilde{w} - w) < 0$, which violates the assumption that w is a stationary point (Definition 2.3). \square

Given Lemma 3.3, we prove Theorem 3.1 by contradiction. We show that if μ_w is far from μ^* , then w violates the

property stated in Lemma 3.3 and therefore cannot be a stationary point. More specifically, we show that, if μ_w is far from μ^* , then there exists a bad sample with index $j \in S_-$ whose gradient is large (Lemma 3.4). Meanwhile, the concentration bounds in Condition (3) guarantee that there exists a good sample with index $i \in S_+$ whose gradient is small (Lemma 3.5).

Lemma 3.4 (Bad sample with large gradient). *Assume that Condition (3) and Lemma 2.1 hold. Fix $w \in \Delta_{N,2\epsilon}$ and a unit vector u with $u^\top \Sigma_w u = \|\Sigma_w\|_2$. Let $r = \|\mu_w - \mu^*\|_2$ and suppose $r \geq c_2 \epsilon \sqrt{\ln(1/\epsilon)}$. Then there exists some $i \in (B \cap S_-)$ such that*

$$\nabla_w F(w, u)_i - u^\top \mu^* (\mu^* - 2\mu_w)^\top u > 2c_3 \cdot \frac{r^2}{\epsilon^2}.$$

Here, c_2 and c_3 are universal positive constants.

Lemma 3.5 (Good sample with small gradient). *Consider the same setting as in Lemma 3.4. There is some $j \in (G \cap S_+)$ such that*

$$\nabla_w F(w, u)_j - u^\top \mu^* (\mu^* - 2\mu_w)^\top u \leq c_3 \cdot \frac{r^2}{\epsilon^2}.$$

We defer the proofs of Lemmas 3.4 and 3.5 to Sections 3.1 and 3.2, and we first use these two lemmas to prove Theorem 3.1.

Proof of Theorem 3.1. Suppose that $w \in \Delta_{N,2\epsilon}$ is a first-order stationary point of $f(w)$, and moreover, w is a bad solution where $\|\mu_w - \mu^*\|_2 \geq c_2 \epsilon \sqrt{\ln(1/\epsilon)}$. By Definition 2.3, there exists a unit vector $u \in \mathbb{R}^d$ such that w is a stationary point of $F(w, u)$.

Fix such a vector u . Since Condition (3) and Lemma 2.1 both hold, we can invoke Lemmas 3.4 and 3.5 on (w, u) to find two coordinates $i \in S_-$ and $j \in S_+$ that violate the bimodal subgradient condition in Lemma 3.3. Consequently, w cannot be a stationary point of $F(w, u)$. This leads to a contradiction, and therefore, all first-order stationary points of $f(w)$ are good solutions. \square

We now briefly comment on the modifications required to prove Theorem 3.2 (see Appendix B). Theorem 3.2 is proven by first showing (using conic geometry) that for such an approximate stationary point an approximate bimodal sub-gradient property holds. Specifically, we show that the bimodal sub-gradient property (Lemma 3.3) is *stable* in the sense that for an approximate stationary point an *approximate bimodal sub-gradient* property holds, i.e., $\nu_i \leq \nu_j + \delta$. Further, for any point obeying such an approximate bimodal property, the objective is small and has good approximation guarantees. The last two steps when combined show that any approximate stationary point has good approximation guarantees (similar to the proof of Theorem 3.1 for exact stationary points).

3.1. Finding a Bad Sample With Large Gradient

In this subsection, we prove Lemma 3.4.

Lemma 3.4 states that when μ_w is far from μ^* , there exists an index $i \in (B \cap S_-)$ such that the gradient $\nabla_w F(w, u)_i$ is relatively large.

Recall that $\nabla_w F(w, u)$ in Equation (4) is the same as the gradient of the variance (weighted by w) of the one-dimensional samples $(X_i^\top u)_{i=1}^N$. Roughly speaking, for this one-dimensional problem, a sample far from the (projected) true mean should have large gradient. Our objective is to find such a sample with positive weight.

More specifically, since w is a bad solution and u is in the top eigenspace of Σ_w , the weighted empirical variance of the projected samples is very large. Because the good samples cannot have this much variance, most of the variance comes from the bad samples. We show that among the bad samples that contribute a lot to the variance, one of them must be very far from the (projected) true mean.

In this section and Section 3.2, we use c_1, \dots, c_4 to denote universal constants that are independent of N, d , and ϵ . We give a detailed description of how to set these constants in Appendix A.

Proof of Lemma 3.4. We first show that the variance of one-dimensional samples $(X_i^\top u)_{i=1}^N$ is relatively large.

By Lemma 2.1, we know that if $\|\mu_w - \mu^*\|_2 \geq r$ and $r \geq c_2 \epsilon \sqrt{\ln(1/\epsilon)}$, then

$$\lambda_{\max}(\Sigma_w) \geq 1 + c_4 \cdot \frac{r^2}{\epsilon}$$

for some universal constant c_4 .

Because u is a unit vector that maximizes $u^\top \Sigma_w u$, we have

$$u^\top \Sigma_w u = \lambda_{\max}(\Sigma_w) \geq 1 + \frac{c_4 r^2}{\epsilon}.$$

Recall that $\Sigma_w = \sum_{i=1}^N w_i (X_i - \mu_w)(X_i - \mu_w)^\top$. If we replace μ_w with μ^* , we have

$$\sum_{i=1}^N w_i (X_i - \mu^*)(X_i - \mu^*)^\top \succeq \Sigma_w,$$

and therefore,

$$u^\top \left(\sum_{i=1}^N w_i (X_i - \mu^*)(X_i - \mu^*)^\top \right) u \geq 1 + \frac{c_4 r^2}{\epsilon}.$$

Next we show that most of this variance is due to bad samples. By Condition (3),

$$u^\top \left(\sum_{i \in G} w_i (X_i - \mu^*)(X_i - \mu^*)^\top \right) u \leq 1 + c_1 \epsilon \ln(1/\epsilon).$$

Consequently,

$$\begin{aligned} & u^\top \left(\sum_{i \in B} w_i (X_i - \mu^*)(X_i - \mu^*)^\top \right) u \\ & \geq \frac{c_4 r^2}{\epsilon} - c_1 \epsilon \ln(1/\epsilon) \geq 0.98 \cdot c_4 \cdot \frac{r^2}{\epsilon}. \end{aligned}$$

The last step is because $r \geq c_2 \cdot \epsilon \sqrt{\ln(1/\epsilon)}$ and we can choose c_4 to be sufficiently large.

Now that we know most of the variance is due to the bad samples, observe that the total weight w_B on the bad samples is at most $\epsilon N \cdot \frac{1}{(1-2\epsilon)^N} \leq 2\epsilon$. Therefore, there must be some $i \in B$ with $w_i > 0$ such that

$$\begin{aligned} u^\top ((X_i - \mu^*)(X_i - \mu^*)^\top) u & \geq \frac{0.98 \cdot c_4 \cdot r^2 \cdot \epsilon^{-1}}{w_B} \\ & \geq 0.49 \cdot c_4 \cdot \frac{r^2}{\epsilon^2}. \end{aligned}$$

In other words,

$$|u^\top (X_i - \mu^*)| \geq 0.7 \cdot \sqrt{c_4} \cdot \frac{r}{\epsilon}.$$

By definition, $i \in B \cap S_-$. It remains to show that $\nabla_w F(w, u)_i$ is large.

$$\begin{aligned} & \nabla_w F(w, u)_i - u^\top \mu^* (\mu^* - 2\mu_w)^\top u \\ & = u^\top ((X_i - \mu^*)(X_i - \mu^*)^\top) u \\ & \quad - 2u^\top ((X_i - \mu^*)(\mu_w - \mu^*)^\top) u \\ & \geq (u^\top (X_i - \mu^*))^2 - 2|u^\top (X_i - \mu^*)| \cdot \|\mu_w - \mu^*\|_2 \\ & \geq \frac{0.49 \cdot c_4 \cdot r^2}{\epsilon^2} - 2 \cdot \frac{0.7 \cdot \sqrt{c_4} \cdot r}{\epsilon} \cdot r > 2c_3 \cdot \frac{r^2}{\epsilon^2}. \end{aligned}$$

The first inequality is by Cauchy-Schwarz. The last step uses the fact that ϵ is sufficiently small. \square

3.2. Finding a Good Sample With Small Gradient

In this subsection, we prove Lemma 3.5.

Lemma 3.5 states that there exists an index $j \in (G \cap S_+)$ such that the gradient $\nabla_w F(w, u)_j$ is relatively small. Similar to the previous section, a sample close to the (projected) true mean should have small gradient. Our goal is to find such a sample for which we can increase its weight.

Recall that S^+ contains all samples whose weight can be increased. We first prove that there are at least ϵN good samples in S^+ . Among these ϵN good samples, the concentration bounds imply that some X_j must be very close to the (projected) true mean.

Proof of Lemma 3.5. Recall that S^+ contains every coordinate i where $w_i < \frac{1}{(1-2\epsilon)^N}$. Since at most $(1-2\epsilon)N$

samples can have the maximum weight $\frac{1}{(1-2\epsilon)N}$, we know that $|S^+| \geq 2\epsilon N$. Combining this with $|G| = (1-\epsilon)N$, we know that $|G \cap S^+| \geq \epsilon N$.

Fix a subset $G^+ \subseteq (G \cap S^+)$ of size $|G^+| = \epsilon N$. We first show that on average, samples in G^+ do not contribute much to the variance.

Let w' be the uniform weight vector on G , i.e., $w'_i = \frac{1}{(1-\epsilon)N}$ for all $i \in G$ and $w'_i = 0$ otherwise. Since $w' \in \Delta_{N,2\epsilon}$, by Condition (3),

$$\left\| \sum_{i \in G} \frac{1}{|G|} (X_i - \mu^*)(X_i - \mu^*)^\top - I \right\|_2 \leq c_1 \cdot \epsilon \ln(1/\epsilon).$$

Let w'' be the uniform weight vector on $S \setminus G^+ = (G \setminus G^+) \cup B$, i.e., $w''_i = \frac{1}{(1-\epsilon)N}$ for all $i \in ((G \setminus G^+) \cup B)$ and $w''_i = 0$ otherwise. Since $w'' \in \Delta_{N,2\epsilon}$, again by Condition (3), we have

$$\left\| \sum_{i \in G \setminus G^+} \frac{1}{|G|} (X_i - \mu^*)(X_i - \mu^*)^\top - I \right\|_2 \leq c_1 \epsilon \ln(1/\epsilon).$$

Combining the previous two concentration bounds,

$$\begin{aligned} & \left\| \sum_{i \in G^+} \frac{1}{|G|} (X_i - \mu^*)(X_i - \mu^*)^\top \right\|_2 \\ & \leq \left\| \sum_{i \in G} \frac{1}{|G|} (X_i - \mu^*)(X_i - \mu^*)^\top - I \right\|_2 \\ & \quad + \left\| \sum_{i \in G \setminus G^+} \frac{1}{|G|} (X_i - \mu^*)(X_i - \mu^*)^\top - I \right\|_2 \\ & \leq 2c_1 \cdot \epsilon \ln(1/\epsilon). \end{aligned}$$

Consequently, because u is a unit vector,

$$u^\top \left(\sum_{i \in G^+} \frac{1}{|G|} (X_i - \mu^*)(X_i - \mu^*)^\top \right) u \leq 2c_1 \epsilon \ln(1/\epsilon).$$

At this point, we know samples in G^+ do not contribute much to the variance. We now proceed to show that one of these samples satisfies the lemma.

Let $j = \arg \min_{i \in G^+} |u^\top (X_i - \mu^*)|$. We have

$$\begin{aligned} u^\top ((X_j - \mu^*)(X_j - \mu^*)^\top) u & \leq \frac{|G|}{|G^+|} \cdot 2c_1 \cdot \epsilon \ln(1/\epsilon) \\ & \leq 2c_1 \ln(1/\epsilon). \end{aligned}$$

Finally, because $|u^\top (X_j - \mu^*)| \leq \sqrt{2c_1 \ln(1/\epsilon)}$, we can

show that $\nabla_w F(w, u)_j$ is small:

$$\begin{aligned} & \nabla f(w)_j - \mu^{*\top} Y(\mu^* - 2\mu_w) \\ & = u^\top ((X_j - \mu^*)(X_j - \mu^*)^\top) u \\ & \quad + 2u^\top ((X_j - \mu^*)(\mu_w - \mu^*)^\top) u \\ & \leq 2c_1 \ln(1/\epsilon) + 2\sqrt{2c_1 \ln(1/\epsilon)} \cdot r \\ & \leq \frac{c_3}{2} \cdot \frac{r^2}{\epsilon^2} + \frac{c_3}{2} \cdot \frac{r}{\epsilon} \cdot r \leq c_3 \cdot \frac{r^2}{\epsilon^2}. \end{aligned}$$

The last step uses that c_3 is sufficiently large, as well as the fact that $\ln(1/\epsilon) \leq \frac{r^2}{\epsilon^2}$ because $r \geq c_2 \epsilon \sqrt{\ln(1/\epsilon)}$. \square

4. Algorithmic Result: Finding a Stationary Point via Gradient Descent

In this section, we show that a simple Projected Gradient Descent (PGD) algorithm (Algorithm 1) can efficiently find an approximate stationary point w of our spectral norm objective, and that w is a good solution to our robust mean estimation task.

Algorithm 1 Robust Mean Estimation via PGD

Input: ϵ -corrupted set of N samples $\{X_i\}_{i=1}^N$ on \mathbb{R}^d satisfying Condition (3), and $\epsilon < \epsilon_0$.

Output: $w \in \mathbb{R}^d$ with $\|\mu_w - \mu^*\|_2 \leq O(\epsilon \sqrt{\log(1/\epsilon)})$.

Let $F(w, u) = u^\top \Sigma_w u$.

Let w_0 be an arbitrary weight vector in $\Delta_{N,2\epsilon}$.

Let $T = \tilde{O}(N^2 d^4)$.

for $\tau = 0$ **to** $T - 1$ **do**

Find a unit vector $u_\tau \in \mathbb{R}^d$ such that $F(w_\tau, u_\tau) \geq (1 - \epsilon) \max_u F(w_\tau, u)$.

$w_{\tau+1} = \mathcal{P}_{\Delta_{N,2\epsilon}}(w_\tau - \eta \nabla_w F(w_\tau, u_\tau))$, where $\mathcal{P}_{\mathcal{K}}(\cdot)$ is the ℓ_2 projection operator onto \mathcal{K} .

end for

return w_{τ^*} where $\tau^* = \arg \min_{0 \leq \tau < T} \|\Sigma_{w_\tau}\|_2$.

We note that finding the unit vector u_τ required in the for loop of Algorithm 1 can be done in time $O(Nd \log(d)/\epsilon)$. Given the PSD matrix $A = \Sigma(w_\tau)$, we want to find a unit vector $u \in \mathbb{R}^d$ such that $u^\top A u \geq (1 - \epsilon) \max_v (v^\top A v)$. This is the (approximate) largest eigenvector problem which can be solved via power method in $O(\log(d)/\epsilon)$ iterations. Since the matrix-vector multiplication $Av = \Sigma_{w_\tau} v = (X \text{diag}(w_\tau) X^\top - X w_\tau w_\tau^\top X^\top) v$ can be computed in time $O(Nd)$, the running time for finding such a vector u_τ is $O(Nd \log(d)/\epsilon)$.

The main result of this section is the following theorem:

Theorem 4.1 (Gradient descent finds a good solution). *Let S be an ϵ -corrupted set of $N = \tilde{\Omega}(d/\epsilon^2)$ samples from a d -dimensional Gaussian $\mathcal{N}(\mu^*, I)$ with unknown mean μ^* . Suppose S satisfies Condition (3) and Lemma 2.1. Then,*

after $\tilde{O}(N^2 d^4)$ iterations, Algorithm 1 outputs a weight vector $w \in \mathbb{R}^N$ such that $\|\mu_w - \mu^*\|_2 = O(\epsilon \sqrt{\log(1/\epsilon)})$.

We first give a high-level overview of the proof. Our proof of Theorem 4.1 can be divided into two steps:

1. The first step is an immediate consequence of Theorem 3.2, which allows us to conclude that any approximate stationary point (in the sense of Definition 2.5) has good approximation guarantees.
2. To finalize the proof, in the second step we show that simple iterative procedures such as (sub)gradient descent can converge in a polynomial number of iterations to such an approximate stationary point. We prove such a result by utilizing a simple and well-known observation: a minimax optimization problem which is smooth in the minimization parameter is weakly convex (after maximization) in the minimization parameter. This connection allows us to leverage recent literature (Drusvyatskiy, 2017; Davis & Drusvyatskiy, 2018) that provides convergence guarantees for weakly convex optimization problems to prove our algorithm finds an approximate stationary point in a polynomial number of iterations.

To elaborate further, in the second step of our proof, we utilize and slightly generalize² the analysis of (Davis & Drusvyatskiy, 2018) and prove that projected sub-gradient descent can find an approximate stationary point.

Lemma 4.2. *Let \mathcal{K} be a closed convex set. Let $F(w, u)$ be a function which is L -Lipschitz and β -smooth with respect to w . Consider the following optimization problem $\min_{w \in \mathcal{K}} \max_{\|u\|_2=1} F(w, u)$.*

Starting from any initial point $w_0 \in \mathcal{K}$, we run iterative updates of the form:

$$\begin{aligned} \text{find } u_\tau \text{ with } F(w_\tau, u_\tau) &\geq (1 - \epsilon') \max_u F(w_\tau, u) \\ w_{\tau+1} &= \mathcal{P}_{\mathcal{K}}(w_\tau - \eta \nabla_w F(w_\tau, u_\tau)) \end{aligned}$$

for T iterations with step size $\eta = \frac{\gamma}{\sqrt{T}}$. Then, we have

$$\begin{aligned} \min_{0 \leq \tau < T} \|\nabla f_\beta(w_\tau)\|_2^2 \\ \leq \frac{2}{\sqrt{T}} \left(\frac{f_\beta(w_0) - \min_w f(w)}{\gamma} + \gamma \beta L^2 \right) + 4\beta \epsilon' \end{aligned}$$

where $f_\beta(w)$ is the Moreau envelope as in Definition 2.4.

As shown in Appendix B, $F(w, u)$ associated with $f(w)$ obeys the required Lipschitz and smoothness property, with $L = \tilde{O}(\sqrt{Nd})$ and $\beta = \tilde{O}(Nd)$. In addition, we have $0 \leq$

²The generalization is to deal with constraints and handle the fact that the inner maximization is not solved precisely.

$f(w) \leq \tilde{O}(d)$ for all $w \in \Delta_{N, 2\epsilon}$. Thus, we can apply the result above with the constraint $\mathcal{K} = \Delta_{N, 2\epsilon}$. Theorem 4.1 follows by combining Theorem 3.2 and Lemma 4.2. We defer the proofs to Appendix B.

5. Discussion

The main conceptual contribution of this work is to establish an intriguing connection between algorithmic high-dimensional robust statistics and non-convex optimization. Specifically, we showed that high-dimensional robust mean estimation can be efficiently solved by directly applying a first-order method to a natural non-convex formulation of the problem.

The main technical contribution of this paper is in showing that any approximate stationary point of our non-convex objective suffices to solve the underlying learning problem. Our novel structural result may be viewed as an explanation as to why robust mean estimation can be solved efficiently in high dimensions, despite its non-convexity. Specifically, we establish that the optimization landscape of our non-convex objective is well-behaved, in a precise sense.

There are a number of directions along which our results could be improved. At the technical level, it would be interesting to obtain faster convergence rates for gradient descent (or other first-order methods), with linear convergence as the ultimate goal. We note that our upper bound is fairly loose and we did not make an explicit effort to optimize the polynomial dependence.

A natural direction is to extend our approach to more general robust estimation tasks, including covariance estimation (Diakonikolas et al., 2016; Cheng et al., 2019), sparse PCA (Balakrishnan et al., 2017; Diakonikolas et al., 2019b), and robust regression (Klivans et al., 2018; Diakonikolas et al., 2019c). Such generalizations will appear in a followup work.

6. Acknowledgments

We thank Jelena Diakonikolas for sharing her expertise in optimization. Part of the work was done while Ilias Diakonikolas and Mahdi Soltanolkotabi were visiting the Simons Institute for the Theory of Computing during the Summer 2019 program on the Foundations of Deep Learning. Part of the work was done while Yu Cheng and Rong Ge were visiting the Institute for Advanced Study for the Special Year on Optimization, Statistics, and Theoretical Machine Learning. Ilias Diakonikolas is supported by NSF Award CCF-1652862 (CAREER), a Sloan Research Fellowship, and a DARPA Learning with Less Labels (LwLL) grant. Rong Ge is supported by NSF CCF-1704656, NSF CCF-1845171 (CAREER), NSF CCF-1934964, a Sloan Fellowship, and a

Google Faculty Research Award. Mahdi Soltanolkotabi is supported by the Packard Fellowship in Science and Engineering, a Sloan Research Fellowship in Mathematics, NSF CCF-CIF grants #1846369 and #1813877, AFOSR-YIP under award #FA9550-18-1-0078, DARPA Learning with Less Labels (LwLL) and Fast Network Interface Cards (FastNICs) programs, and a Google Faculty Research Award.

References

- Allen-Zhu, Z., Lee, Y., and Orecchia, L. Using optimization to obtain a width-independent, parallel, simpler, and faster positive SDP solver. In *Proc. 27th Annual Symposium on Discrete Algorithms (SODA)*, pp. 1824–1831, 2016.
- Balakrishnan, S., Du, S. S., Li, J., and Singh, A. Computationally efficient robust sparse estimation in high dimensions. In *Proc. 30th Annual Conference on Learning Theory*, pp. 169–212, 2017.
- Barreno, M., Nelson, B., Joseph, A. D., and Tygar, J. D. The security of machine learning. *Machine Learning*, 81(2):121–148, 2010.
- Beck, A. *First-Order Methods in Optimization*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2017. doi: 10.1137/1.9781611974997. URL <https://epubs.siam.org/doi/abs/10.1137/1.9781611974997>.
- Biggio, B., Nelson, B., and Laskov, P. Poisoning attacks against support vector machines. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012*, 2012.
- Candes, E. J., Li, X., and Soltanolkotabi, M. Phase retrieval via Wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015.
- Cheng, Y., Diakonikolas, I., and Ge, R. High-dimensional robust mean estimation in nearly-linear time. *CoRR*, abs/1811.09380, 2018. URL <http://arxiv.org/abs/1811.09380>. Conference version in SODA 2019, p. 2755–2771.
- Cheng, Y., Diakonikolas, I., Ge, R., and Woodruff, D. P. Faster algorithms for high-dimensional robust covariance estimation. In *Conference on Learning Theory, COLT 2019*, pp. 727–757, 2019.
- Davis, D. and Drusvyatskiy, D. Stochastic subgradient method converges at the rate $o(k^{-1/4})$ on weakly convex functions. *arXiv preprint arXiv:1802.02988*, 2018.
- Depersin, J. and Lecue, G. Robust subgaussian estimation of a mean vector in nearly linear time. *CoRR*, abs/1906.03058, 2019.
- Diakonikolas, I. and Kane, D. M. Recent advances in algorithmic high-dimensional robust statistics. *CoRR*, abs/1911.05911, 2019. URL <http://arxiv.org/abs/1911.05911>.
- Diakonikolas, I., Kamath, G., Kane, D. M., Li, J., Moitra, A., and Stewart, A. Robust estimators in high dimensions without the computational intractability. In *Proc. 57th IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 655–664, 2016.
- Diakonikolas, I., Kamath, G., Kane, D. M., Li, J., Moitra, A., and Stewart, A. Being robust (in high dimensions) can be practical. In *Proc. 34th International Conference on Machine Learning (ICML)*, pp. 999–1008, 2017a.
- Diakonikolas, I., Kane, D. M., and Stewart, A. 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)*, pp. 73–84, 2017b.
- Diakonikolas, I., Kamath, G., Kane, D. M., Li, J., Steinhardt, J., and Stewart, A. SEVER: A robust meta-algorithm for stochastic optimization. In *Proc. 36th International Conference on Machine Learning (ICML)*, pp. 1596–1606, 2019a.
- Diakonikolas, I., Karmalkar, S., Kane, D., Price, E., and Stewart, A. Outlier-robust high-dimensional sparse estimation via iterative filtering. In *Advances in Neural Information Processing Systems 33, NeurIPS 2019*, pp. 10688–10699, 2019b.
- Diakonikolas, I., Kong, W., and Stewart, A. Efficient algorithms and lower bounds for robust linear regression. In *Proc. 30th Annual Symposium on Discrete Algorithms (SODA)*, pp. 2745–2754, 2019c.
- Dong, Y., Hopkins, S. B., and Li, J. Quantum entropy scoring for fast robust mean estimation and improved outlier detection. *CoRR*, abs/1906.11366, 2019. URL <http://arxiv.org/abs/1906.11366>. Conference version in NeurIPS 2019.
- Drusvyatskiy, D. The proximal point method revisited. *arXiv preprint arXiv:1712.06038*, 2017.
- Ge, R., Huang, F., Jin, C., and Yuan, Y. Escaping from saddle points—online stochastic gradient for tensor decomposition. In *Conference on Learning Theory*, pp. 797–842, 2015.

- Hassani, H., Soltanolkotabi, M., and Karbasi, A. Gradient methods for submodular maximization. In *Advances in Neural Information Processing Systems*, pp. 5841–5851, 2017.
- Huber, P. J. Robust estimation of a location parameter. *Ann. Math. Statist.*, 35(1):73–101, 03 1964.
- Jain, P. and Kar, P. Non-convex optimization for machine learning. *Foundations and Trends® in Machine Learning*, 10(3-4):142–336, 2017.
- Jin, C., Netrapalli, P., and Jordan, M. I. What is local optimality in nonconvex-nonconcave minimax optimization? *arXiv preprint arXiv:1902.00618*, 2019.
- Keshavan, R. H., Montanari, A., and Oh, S. Matrix completion from a few entries. *IEEE transactions on information theory*, 56(6):2980–2998, 2010.
- Klivans, A., Kothari, P., and Meka, R. Efficient algorithms for outlier-robust regression. In *Proc. 31st Annual Conference on Learning Theory (COLT)*, pp. 1420–1430, 2018.
- Lai, K. A., Rao, A. B., and Vempala, S. Agnostic estimation of mean and covariance. In *Proc. 57th IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 665–674, 2016.
- Li, J., Absher, D., Tang, H., Southwick, A., Casto, A., Ramachandran, S., Cann, H., Barsh, G., Feldman, M., Cavalli-Sforza, L., and Myers, R. Worldwide human relationships inferred from genome-wide patterns of variation. *Science*, 319:1100–1104, 2008.
- Loh, P. and Wainwright, M. J. High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity. In *Advances in Neural Information Processing Systems*, pp. 2726–2734, 2011.
- Paschou, P., Lewis, J., Javed, A., and Drineas, P. Ancestry informative markers for fine-scale individual assignment to worldwide populations. *Journal of Medical Genetics*, 47:835–847, 2010.
- Peng, R., Tangwongsan, K., and Zhang, P. Faster and simpler width-independent parallel algorithms for positive semidefinite programming. *arXiv preprint arXiv:1201.5135v3*, 2016.
- Prasad, A., Suggala, A. S., Balakrishnan, S., and Ravikumar, P. Robust estimation via robust gradient estimation. *arXiv preprint arXiv:1802.06485*, 2018.
- Rockafellar, R. T. *Convex analysis*. Number 28. Princeton university press, 1970.
- Rockafellar, R. T. Favorable classes of lipschitz continuous functions in subgradient optimization. 1981.
- Rockafellar, R. T. *Convex Analysis*. Princeton Landmarks in Mathematics and Physics. Princeton University Press, 2015. ISBN 9781400873173.
- Rosenberg, N., Pritchard, J., Weber, J., Cann, H., Kidd, K., Zhivotovsky, L., and Feldman, M. Genetic structure of human populations. *Science*, 298:2381–2385, 2002.
- Steinhardt, J., Koh, P. W., and Liang, P. S. Certified defenses for data poisoning attacks. In *Advances in Neural Information Processing Systems 30*, pp. 3520–3532, 2017.
- Steinhardt, J., Charikar, M., and Valiant, G. Resilience: A criterion for learning in the presence of arbitrary outliers. In *Proc. 9th Innovations in Theoretical Computer Science Conference (ITCS)*, pp. 45:1–45:21, 2018.
- Tu, S., Boczar, R., Simchowitz, M., Soltanolkotabi, M., and Recht, B. Low-rank solutions of linear matrix equations via Procrustes flow. *arXiv preprint arXiv:1507.03566*, 2015.
- Tukey, J. W. A survey of sampling from contaminated distributions. *Contributions to probability and statistics*, 2:448–485, 1960.
- Wilcox, R. M. Exponential operators and parameter differentiation in quantum physics. *Journal of Mathematical Physics*, 8(4):962–982, 1967.