Robust Estimation via Robust Gradient Estimation

Adarsh Prasad[‡]   Arun Sai Suggala[‡]   Sivaraman Balakrishnan[†]   Pradeep Ravikumar[‡]

Machine Learning Department[‡]
Department of Statistics[†]
Carnegie Mellon University,
Pittsburgh, PA 15213.

Abstract

We provide a new computationally-efficient class of estimators for risk minimization. We show that these estimators are robust for general statistical models, under varied robustness settings, including in the classical Huber ϵ-contamination model, and in heavy-tailed settings. Our workhorse is a novel robust variant of gradient descent, and we provide conditions under which our gradient descent variant provides accurate estimators in a general convex risk minimization problem. We provide specific consequences of our theory for linear regression, logistic regression and for canonical parameter estimation in an exponential family. These results provide some of the first computationally tractable and provably robust estimators for these canonical statistical models. Finally, we study the empirical performance of our proposed methods on synthetic and real datasets, and find that our methods convincingly outperform a variety of baselines.

1   Introduction

In classical analyses of statistical estimators, statistical guarantees are derived under strong model assumptions, and in most cases these guarantees hold only in the absence of arbitrary outliers, and other deviations from the model assumptions. Strong model assumptions are rarely met in practice, and this has motivated the development of robust inferential procedures, and which has a rich history in statistics with seminal contributions due to Box [1], Tukey [2], Huber [3], Hampel [4] and several others. These have led to rich statistical concepts such as the influence function, the breakdown point, and the Huber ϵ-contamination model, to assess the robustness of estimators. Despite this progress however, the statistical methods with the strongest robustness guarantees are computationally intractable, for instance those based on non-convex M-estimators [3], ℓ_1 tournaments [5, 6, 7] and notions of depth [8, 9, 10].

In this paper, we present a general class of estimators that are computationally tractable, and have strong robustness guarantees. The estimators we propose are obtained by robustifying iterative updates of risk minimization, and are broadly applicable to a wide-range of parametric statistical models. In the risk minimization framework, the target parameter θ^* is defined as the solution to an optimization problem:

$$\theta^* = \operatorname*{argmin}_{\theta \in \Theta} \mathcal{R}(\theta) \equiv \operatorname*{argmin}_{\theta \in \Theta} \mathbb{E}_{z \sim P} \left[\overline{\mathcal{L}}(\theta; z) \right], \tag{1}$$

where $\overline{\mathcal{L}}$ is an appropriate loss-function, R is the population risk and Θ is the set of feasible parameters. The statistical inference problem within the risk minimization framework is then to compute an approximate minimizer to the above program when given access to samples $\mathcal{D}_n = \{z_1, \ldots, z_n\}$. A classical approach to do so is via *empirical risk minimization* (ERM), where we substitute the empirical expectation given the samples for the population expectation

in the specification of the risk objective. While most modern statistical estimators use the above empirical risk minimization framework, a standard assumption that is imposed on $\mathcal{D}_n$ is that the data has no outliers, and has no arbitrary deviations from model assumptions; i.e., it is typically assumed that each of the z_i's are independent and identically distributed according to the distribution P. Moreover, many analyses of risk minimization further assume that P follows a sub-gaussian distribution, or has otherwise well-controlled tails in order to appropriately control the deviation between the population risk and its empirical counterpart. Due in part to these caveats with ERM, the seminal work of M-estimation replaces the risk minimization objective with a robust counterpart, so that the minimizer of the empirical expectation of the robust counterpart is more robust than the ERM minimizer. As noted above, for strong statistical guarantees, these in turn require solving computationally intractable non-convex optimization programs.

In contrast to this classical work, we propose a class of estimators that have a shift in perspective: rather than specify a robust objective, we consider an algorithm, namely projected gradient descent, that directly optimizes the population risk objective in Eq. (1), and focus on making this algorithm robust. Thus, in contrast to specifying the robust parameter estimate as the solution to an optimization program as in M-estimation, which in turn could be computationally intractable, we specify the robust parameter estimate as the limit of a sequence of iterative updates that are individually robust as well as computationally tractable. We find that this shift in perspective leads to estimators that are both computationally tractable as well with strong robustness guarantees, that are as broadly applicable as ERM or M-estimators, and moreover with a unified statistical treatment for varied statistical models.

In addition to being applicable to a variety of statistical models, our general results are also applicable to a variety of *notions of robustness*. In this paper, we derive corollaries in particular for two canonical robustness settings:

(a) **Robustness to arbitrary outliers:** In this setting, we focus on Huber's ϵ-contamination model, where rather than observe samples directly from P in (1) we instead observe samples drawn from P_ϵ which for an *arbitrary* distribution Q is defined as:

$$P_\epsilon = (1 - \epsilon)P + \epsilon Q.$$

The distribution Q allows for arbitrary outliers, which may correspond to gross corruptions or more subtle deviations from the assumed model. This model can be equivalently viewed as model mis-specfication in the Total Variation (TV) metric.

(b) **Robustness to heavy-tails:** In this setting, we are interested in developing estimators under weak moment assumptions. We assume that the distribution P from which we obtain samples only has finite low-order moments (see Section 5.3 for a precise characterization). Such heavy tailed distributions arise frequently in the analysis of financial data and large-scale biological datasets (see for instance examples in [11, 12]). In contrast to classical analyses of empirical risk minimization [13], in this setting the empirical risk is not uniformly close to the population risk, and methods that directly minimize the empirical risk perform poorly (see Section 4).

While we provide corollaries demonstrating robustness with respect to the above deviations, we emphasize that our framework is more general. Below, we provide an outline of our results and contributions.

1. **Estimators.** Our first contribution is to introduce a new class of robust estimators for risk minimization (1). These estimators are based on robustly estimating gradients of

the population risk to then plug in to a projected gradient descent algorithm, and are computationally tractable by design. A crucial ingredient of our framework is the design of robust gradient estimators for the population risk in (1). Our main insight is that in this general risk minimization setting, the gradient of the population risk is simply a multivariate mean vector, and we can leverage prior work on mean estimation to design robust gradient estimators. Thus, for our two canonical robustness cases, we develop such robust gradient estimators building on prior work for robust mean estimation in the Huber model [14], and in the heavy-tailed model [15]. Another perspective of our framework is that it significantly generalizes the applicability of *mean estimation* methods to general parametric models.

2. **Empirical Investigations.** Our estimators are computationally practical, and accordingly, our second contribution is to conduct extensive numerical experiments on real and simulated data with our proposed class of estimators. We provide guidelines for tuning parameter selection, and compare the proposed estimators with several competitive baselines [3, 16, 17]. We find that our estimators consistently perform well across different settings, and across various metrics.

3. **Statistical Guarantees.** Finally, we provide rigorous robustness guarantees for the estimators we propose for a variety of classical statistical models: linear regression, logistic regression, and exponential family models. Our contributions in this direction are two-fold: building on prior work [18] we provide a general result on the stability of gradient descent for risk minimization, and show that under certain conditions, gradient descent can be quite tolerant to inaccurate gradient estimates. Subsequently, in concrete settings, we provide a careful analysis of the quality of gradient estimation afforded by our proposed gradient estimators, and combine these results to obtain guarantees on our final parameter estimates.

Broadly, as we discuss in the sequel, our work suggests that our class of estimators based on robust gradient estimation offer a variety of practical, conceptual, statistical and computational advantages for robust estimation. They provide the general applicability of classical M-estimators, together with computational practicality even for large-scale models, as well as strong robustness guarantees.

1.1 Related Work

There has been extensive work in the broad area of robust statistics (see for instance [4] and references therein); we focus this section on some lines of work that are most related to this paper. For the robustness setting of ϵ-contaminated models, several classical estimators have been developed that are optimally robust for a variety of inferential tasks, including hypothesis testing [19], mean estimation [8], general parametric estimation [6, 7, 20], and non-parametric estimation [5]. However, a major drawback with this classical line of work has been that most of the estimators with strong robustness guarantees are computationally intractable [2], while the remaining ones use heuristics and are consequently not optimal [21]. A complementary line of recent research [9, 10] has focused on providing minimax upper and lower bounds on the performance of estimators under ϵ-contamination model, without the constraint of computational tractability. Recently, there has been a flurry of research in theoretical computer science [14, 22, 23, 24, 25] designing provably robust estimators which are computationally tractable while achieving near-optimal contamination dependence, for

special classes of problems such as computing means and covariances. Some of the proposed algorithms are however not computationally practical as they rely on the ellipsoid algorithm or require solving semi-definite programs [22, 23, 24, 25] which can be slow for modern problem sizes.

While in the general ϵ-contamination setting, the contamination distribution could be arbitrary, there has been a lot of work in settings where the contamination distribution is restricted in various ways. For example, recent work in high-dimensional statistics (for instance [26, 27, 28, 29, 30]) have studied problems like principal component analysis and linear regression under the assumption that the corruptions are evenly spread throughout the dataset.

For the robustness setting of heavy tailed distributions, robust estimators aim to relax the sub-gaussian or sub-exponential distributional assumptions that are typically imposed on the target distribution, and allow it to be a heavy tailed distribution. Most approaches in this category substitute the *empirical mean* of the risk objective in risk minimization with robust mean estimators such as [31, 32] that exhibit sub-gaussian type concentration around the true mean for distributions satisfying mild moment assumptions. The median-of-means estimator [31] and Catoni's mean estimator [32] are two popular examples of such robust mean estimators. In particular, Hsu and Sabato [33] use the median-of-means estimator to solve the corresponding robust variant of ERM. Although this estimator has strong theoretical guarantees, and is computationally tractable, as noted by the authors in [33] it performs poorly in practice. In recent work Brownlees et al. [34] use the Catoni's mean estimator to solve the corresponding robust variant of ERM. The authors provide risk bounds similar to bounds one can achieve under sub-gaussian distributional assumptions. However, their estimator is not easily computable and the authors do not provide a practical algorithm to compute the estimator. Other recent work by Lerasle and Oliveira [31], Lugosi and Mendelson [35] use similar ideas to derive estimators that perform well theoretically, in heavy-tailed situations. However, these approaches involve optimization of complex objectives for which no computationally tractable algorithms exist. We emphasize that in contrast to our work, these works focus on robustly estimating the population risk which does not directly lead to a computable estimator. In contrast, we consider robustly estimating the *gradient* of the population risk, and embedding these estimates within the iterative algorithm of projected gradient descent, which leads naturally to a computionally practical estimator.

1.2 Outline

We conclude this section with a brief outline of the remainder of the paper. In Section 2, we provide some background on risk minimization and the running robustness settings of Huber contamination, and heavy-tailed noise models. In Section 3, we introduce our class of robust estimators, and provide concrete algorithms for the ϵ-contaminated and heavy-tailed settings. In Section 4 we study the empirical performance of our estimator on a variety of tasks and datasets. We complement our empirical results with theoretical guarantees in Sections 5, 6 and 7. We defer technical details to the Appendix. Finally, we conclude in Section 8 with a discussion of some open problems.

2 Background and Problem Setup

In this section we provide the necessary background on risk minimization, gradient descent, and introduce two notions of robustness that we consider in this work.

2.1 Risk Minimization and Parametric Estimation

In the setting of risk minimization, we assume that we have access to a differentiable loss function $\overline{\mathcal{L}} : \Theta \times \mathcal{Z} \mapsto \mathbb{R}$, where Θ is a convex subset of $\mathbb{R}^p$. Let $\mathcal{R}(\theta) = \mathbb{E}_{z \sim P}\left[\overline{\mathcal{L}}(\theta; z)\right]$ be the population loss, also known as the *risk*, and let θ^* be the minimizer of the population risk $\mathcal{R}(\theta)$, over the set Θ:

$$\theta^* = \operatorname*{argmin}_{\theta \in \Theta} \mathcal{R}(\theta). \tag{2}$$

The goal of risk minimization is to minimize the population risk $\mathcal{R}(\theta)$, given only n samples $\mathcal{D}_n = \{z_i\}_{i=1}^n$, in order to estimate the unknown parameter θ^*.

In this work we assume that the population risk is convex to ensure tractable minimization. Moreover, in order to ensure identifiability of the parameter θ^*, we impose two standard regularity conditions [36] on the population risk. These properties are defined in terms of the error of the first-order Taylor approximation of the population risk, i.e. defining, $\tau(\theta_1, \theta_2) := \mathcal{R}(\theta_1) - \mathcal{R}(\theta_2) - \langle \nabla \mathcal{R}(\theta_2), \theta_1 - \theta_2 \rangle$, we assume that

$$\frac{\tau_\ell}{2}\|\theta_1 - \theta_2\|_2^2 \leq \tau(\theta_1, \theta_2) \leq \frac{\tau_u}{2}\|\theta_1 - \theta_2\|_2^2, \tag{3}$$

where the parameters $\tau_\ell, \tau_u > 0$ denote the strong-convexity and smoothness parameters respectively.

2.2 Illustrative Examples of Risk Minimization

The framework of risk minimization is a central paradigm of statistical estimation and is widely applicable. In this section, we provide illustrative examples that fall under this framework.

2.2.1 Linear Regression

Here we observe paired samples $\{(x_1, y_1), \dots (x_n, y_n)\}$, where each $(x_i, y_i) \in \mathbb{R}^p \times \mathbb{R}$. We assume that the (x, y) pairs sampled from the true distribution P are linked via a linear model:

$$y = \langle x, \theta^* \rangle + w, \tag{4}$$

where w is drawn from a zero-mean distribution such as normal distribution with variance σ^2 ($\mathcal{N}(0, \sigma^2)$) or a more heavy-tailed distribution such as student-t or Pareto distribution. We suppose that under P the covariates $x \in \mathbb{R}^p$, have mean 0, and covariance Σ.

For this setting we use the squared loss as our loss function, which induces the following population risk:

$$\overline{\mathcal{L}}(\theta; (x, y)) = \frac{1}{2}\left(y - \langle x, \theta \rangle\right)^2, \quad \text{and} \quad \mathcal{R}(\theta) = \frac{1}{2}(\theta - \theta^*)^T \Sigma (\theta - \theta^*).$$

Note that the true parameter θ^* is the minimizer of the population risk $\mathcal{R}(\theta)$. The strong-convexity and smoothness assumptions from (3) in this setting require that $\tau_\ell \leq \lambda_{\min}(\Sigma) \leq \lambda_{\max}(\Sigma) \leq \tau_u$.

2.2.2 Generalized Linear Models

Here we observe paired samples $\{(x_1, y_1), \ldots (x_n, y_n)\}$, where each $(x_i, y_i) \in \mathbb{R}^p \times \mathcal{Y}$. We suppose that the (x, y) pairs sampled from the true distribution P are linked via a linear model such that when conditioned on the covariates x, the response variable has the distribution:

$$P(y|x) \propto \exp\left(\frac{y\langle x, \ \theta^*\rangle - \Phi(\langle x, \ \theta^*\rangle)}{c(\sigma)}\right) \tag{5}$$

Here $c(\sigma)$ is a fixed and known scale parameter and $\Phi : \mathbb{R} \mapsto \mathbb{R}$ is the link function. We focus on the random design setting where the covariates $x \in \mathbb{R}^p$, have mean 0, and covariance Σ. We use the negative conditional log-likelihood as our loss function, i.e.

$$\overline{\mathcal{L}}(\theta; (x, y)) = -y\langle x, \ \theta\rangle + \Phi(\langle x, \ \theta\rangle). \tag{6}$$

Once again, the true parameter θ^* is the minimizer of the resulting population risk $\mathcal{R}(\theta)$. It is easy to see that Linear Regression with Gaussian Noise lies in the family of generalized linear models. A popular instance of such GLMs is a logistic regression model.

Logistic Regression. In this case the (x, y) pairs are linked as:

$$y = \begin{cases} 1 & \text{with probability} \quad \frac{1}{1+\exp(-\langle x, \theta^*\rangle)}, \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

This corresponds to setting $\Phi(t) = \log(1 + \exp(t))$ and $c(\gamma) = 1$ in (5). The hessian of the population risk is given by

$$\nabla^2 \mathcal{R}(\theta) = \mathbb{E}\left[\frac{\exp\langle x, \theta\rangle}{(1 + \exp\langle x, \theta\rangle)^2} xx^T\right].$$

Note that as θ diverges, the minimum eigenvalue of the hessian approaches 0 and the loss is no longer strongly convex. To prevent this, in this case we take the parameter space Θ to be bounded.

2.2.3 Exponential Families and Canonical Parameters.

Finally we consider the case where the true distribution P is in exponential family with canonical parameters $\theta^* \in \mathbb{R}^p$, and a vector of sufficient statistics obtained from the map $\phi : \mathcal{Z} \mapsto \mathbb{R}^p$. Note that while the linear and logistic regression models are indeed in an exponential family, our interest in those cases was not in the canonical parameters.

In more details, we can write the true distribution P in this case as

$$P(z) = h(z) \exp\left(\langle\phi(z), \ \theta^*\rangle - A(\theta^*)\right),$$

where $h(z)$ is some base measure. The negative log-likelihood gives us the following loss function:

$$\overline{\mathcal{L}}(\theta; z) = -\langle\phi(z), \ \theta\rangle + A(\theta). \tag{8}$$

The strong-convexity and smoothness assumptions require that there are constants τ_ℓ, τ_u such that $\tau_\ell \leq \nabla^2 A(\theta) \leq \tau_u$, for $\theta \in \Theta$.

2.3 Empirical Risk Minimization

Given data $\mathcal{D}_n = \{z_i\}_{i=1}^n$, empirical risk minimization (ERM) substitutes the empirical expectation of the risk for the population risk in the risk minimization objective:

$$\widehat{\theta}_n = \operatorname*{argmin}_{\theta \in \Theta} \mathcal{R}_n(\theta) := \frac{1}{n} \sum_{i=1}^n \overline{\mathcal{L}}(\theta; z_i).$$

Most modern statistical estimators follow this ERM recipe above. When the loss is the log-likelihood of the statistical model, this reduces to the classical Maximum Likelihood Estimation (MLE) principle. The empirical risk minimizer is however a poor estimator of θ^* in the presence of outliers in the data: since ERM depends on the sample mean, outliers in the data can effect the sample mean and lead ERM to sub-optimal estimates. This observation has led to a large body of research that focuses on developing robust M-estimators, where we substitute in the empirical expectation of a robust counterpart of the loss function $\overline{\mathcal{L}}$; the resulting estimators have favorable statistical properties, but are often computationally intractable.

2.4 Projected Gradient Descent

A popular approach for solving the empirical risk minimization problem is projected gradient descent. Projected gradient descent generates a sequence of iterates $\{\theta^t\}_{t=0}^\infty$, by refining an initial parameter $\theta_0 \in \Theta$ via the update:

$$\theta^{t+1} = \mathcal{P}_\Theta \left(\theta^t - \eta \nabla \mathcal{R}_n(\theta^t) \right),$$

where $\eta > 0$ is the step size and $\mathcal{P}_\theta$ is the projection operator onto Θ. While this gradient descent method is simple, it is not however robust to various deviations, for general convex losses. Accordingly, we have a small shift in perspective: instead of performing gradient descent on the empirical risk, we perform gradient descent on the population risk. Our work then relies on the important observation that this gradient of the population risk ($\mathbb{E}_{z \sim P} \left[\nabla \overline{\mathcal{L}}(\theta; z) \right]$) is simply a mean vector: one that can be estimated robustly by leveraging recent advances in robust mean estimation [14, 15]. This leads to a general method for risk minimization based on embedding robust gradient estimation within a projected gradient descent algorithm (see Algorithm 1).

2.5 Robust Estimation

One of the goals of this work is to develop general statistical estimation methods that are robust under varied robustness settings. We derive corollaries in particular for two robustness models: Huber's ϵ-contamination model, and the heavy-tailed model. We now briefly review these two notions of robustness.

(a) **Huber's ϵ-contamination model:** Huber [19, 37] proposed the ϵ-contamination model where we observe samples that are obtained from a mixture of the form

$$P_\epsilon = (1 - \epsilon)P + \epsilon Q, \tag{9}$$

where P is the true distribution, ϵ is the expected fraction of outliers and Q is an *arbitrary* outlier distribution. Given i.i.d. observations drawn from P_ϵ, our objective is to estimate θ^*, the minimizer of the population risk $\mathcal{R}(\theta) = \mathbb{E}_{z \sim P} \left[\overline{\mathcal{L}}(\theta; z) \right]$, robust to the contamination from Q.

(b) **Heavy-tailed model:** In the heavy-tailed model it is assumed that the data follows a heavy-tailed distribution (i.e, P is heavy-tailed). While heavy-tailed distributions have various possible characterizations: in this paper we consider a characterization via gradients. For a fixed $\theta \in \Theta$ we let P_g^θ denote the multivariate distribution of the gradient of population loss, i.e. $\nabla \overline{\mathcal{L}}(\theta; z)$. We refer to a potentially heavy-tailed distribution as one for which our only assumption on P_g^θ is that it has finite second moments for any $\theta \in \Theta$. As we illustrate in Section 7, in various concrete examples this translates to relatively weak low-order moment assumptions on the data distribution P.

Given n i.i.d observations from P, our objective is to estimate the minimizer of the population risk. From a conceptual standpoint, the classical analysis of risk-minimization which relies on uniform concentration of the empirical risk around the true risk, fails in the heavy-tailed setting necessitating new estimators and analyses [33, 34, 35, 38].

3 Robust Gradient Descent via Gradient Estimation

Gradient descent and its variants are at the heart of modern optimization and are well-studied in the literature. Suppose we have access to the true distribution P_{θ^*}. Then to minimize the population risk $\mathcal{R}(\theta)$, we can use projected gradient descent, where starting at some initial θ^0 and for an appropriately chosen step-size η, we update our estimate according to:

$$\theta^{t+1} \leftarrow \mathcal{P}_\Theta(\theta^t - \eta \nabla \mathcal{R}(\theta^t)). \tag{10}$$

However, we only have access to n samples $\mathcal{D}_n = \{z_i\}_{i=1}^n$. The key technical challenges are then to estimate the gradient of $\mathcal{R}(\theta)$ from samples $\mathcal{D}_n$, and to ensure that an appropriate modification of gradient descent is stable to the resulting estimation error.

To address the first challenge we observe that the gradient of the population risk at any point θ is the mean of a multivariate distribution, *i.e.* $\nabla \mathcal{R}(\theta) = E_{z \sim P}\left[\nabla \overline{\mathcal{L}}(\theta; z)\right]$. Accordingly, the problem of gradient estimation can be reduced to a multivariate mean estimation problem, where our goal is to *robustly* estimate the true mean $\nabla \mathcal{R}(\theta)$ from n samples $\{\nabla \overline{\mathcal{L}}(\theta; z_i)\}_{i=1}^n$. For a given sample-size n and confidence parameter $\delta \in (0,1)$ we define a gradient estimator:

Definition 1. *A function $g(\theta; \mathcal{D}_n, \delta)$ is a gradient estimator, if for functions α and β, with probability at least $1 - \delta$, at any fixed $\theta \in \Theta$, the estimator satisfies the following inequality:*

$$\|g(\theta; \mathcal{D}_n, \delta) - \nabla \mathcal{R}(\theta)\|_2 \le \alpha(n, \delta)\|\theta - \theta^*\|_2 + \beta(n, \delta). \tag{11}$$

In subsequent sections, we will develop conditions under which we can obtain gradient estimators with strong control on the functions $\alpha(n, \delta)$ and $\beta(n, \delta)$ in the Huber and heavy-tailed models. Furthermore, by investigating the stability of gradient descent we will develop sufficient conditions on these functions such that gradient descent with an inaccurate gradient estimator still returns an accurate estimate.

To minimize $\mathcal{R}(\theta)$, we replace $\nabla \mathcal{R}(\theta)$ in equation (10) with the gradient estimator $g(\theta; \mathcal{D}_n, \delta)$ and perform projected gradient descent. In order to avoid complex statistical dependency issues that can arise in the analysis of gradient descent, for our theoretical results we consider a sample-splitting variant of the algorithm where each iteration is performed on a fresh batch of samples. We summarize the overall robust gradient descent algorithm via gradient estimation in Algorithm 1. In contrast to M-estimation where we use robust estimates of the overall loss function, here we use robust estimates of the gradient, a small shift in perspective, but which

---
Algorithm 1 Robust Gradient Descent
---
function RGD (GRADIENT ESTIMATOR $g(\cdot)$, DATA $\{z_1, \ldots, z_n\}$, STEP SIZE η, NUMBER
OF ITERATIONS T, CONFIDENCE δ)
    Split samples into T subsets $\{\mathcal{Z}_t\}_{t=1}^{T}$ of size $\widetilde{n}$.
    for $t = 0$ to $T - 1$ **do**
        $\theta^{t+1} = \operatorname{argmin}_{\theta \in \Theta} \| \theta - \left(\theta^t - \eta \, g(\theta^t; \mathcal{Z}_t, \widetilde{\delta}) \right) \|_2^2$.
    end for
end function
---

has strong statistical and computational consequences: we obtain a computationally practical algorithm, and moreover with strong robustness guarantees via careful statistical analyses of the stability of the resulting biased and inexact gradient descent iterates.

We further assume that the number of gradient iterations T is specified a-priori, and accordingly we define:

$$\widetilde{n} = \left\lfloor \frac{n}{T} \right\rfloor \quad \text{and} \quad \widetilde{\delta} = \frac{\delta}{T}. \tag{12}$$

We discuss methods for selecting T, and the impact of sample-splitting in later sections. As confirmed in our experiments (see Section 4), sample-splitting should be viewed as a device introduced for theoretical convenience which can likely be eliminated via more complex uniform arguments (see for instance the work [18]).

It can be seen that the key ingredient in the robust gradient descent Algorithm 1 is a robust estimator of the gradients. Next, we consider the two notions of robustness described in Section 2, and derive specific gradient estimators for each of the models using the framework described above. Although we derive corollaries of our general results for these two settings of Huber contamination and heavy-tailed models, we emphasize that our class of estimators are more general and are not restricted to these two notions of robustness.

3.1 Gradient Estimation in Huber's ϵ-contamination model

There has been a flurry of recent interest [14, 22, 23, 24, 25] in designing mean estimators which, under the Huber contamination model, can robustly estimate the mean of a random vector. While some of these results are focused on the case where the uncorrupted distribution is Gaussian, or isotropic, we are more interested in robust mean oracles for more general distributions. Lai et al. [14] proposed a robust mean estimator for general distributions, satisfying weak moment assumptions, and we leverage the existence of such an estimator to design a *Huber gradient estimator* $g(\theta; \mathcal{D}_n, \delta)$ which works in the Huber contamination model.

The estimator builds upon the fact that with a single dimension, it is relatively easy to estimate the gradient robustly. In higher dimensions, the crucial insight of Lai et al. [14] is that the effect of the contamination distribution Q on the mean of uncontaminated distribution P is effectively one-dimensional provided we can accurately estimate the direction along which the mean is shifted. In our context, if we can compute the gradient shift direction, i.e. the direction of the difference between the sample (corrupted) mean gradient and the true (population) gradient, then the true gradient can be estimated by using a robust 1D-mean algorithm along the gradient-shift direction and a non-robust sample-gradient in the orthogonal direction since the contamination has no effect on the gradient in this orthogonal direction.

In order to identify this gradient shift direction, we follow Lai et al. [14] and use a recursive Singular Value Decomposition (SVD) based algorithm. In each stage of the recursion, we first remove gross-outliers via a truncation algorithm (described in more detail in the Appendix, and termed HUBEROUTLIERGRADIENTTRUNCATION in Algorithm 2). We subsequently identify two subspaces using an SVD – a clean subspace where the contamination has a small effect on the mean and another subspace where the contamination has a potentially larger effect. We use a simple sample-mean estimator in the clean subspace and recurse our computation on the other subspace. Building on the work of Lai et al. [14], in Lemma 1 and Appendix K we provide a careful non-asymptotic analysis of this gradient estimator.

Algorithm 2 presents the overall *Huber gradient estimator* $g(\theta; \mathcal{D}_n, \delta)$.

Algorithm 2 Huber Gradient Estimator

function HUBERGRADIENTESTIMATOR(SAMPLE GRADIENTS $S = \{\nabla\overline{\mathcal{L}}(\theta; z_i)\}_{i=1}^n$, COR-RUPTION LEVEL ϵ, DIMENSION p, δ)

$\quad \widetilde{S} = $ HUBEROUTLIERGRADIENTTRUNCATION(S, ϵ, p, δ).

$\quad$ **if** p=1 **then**

$\quad\quad$ **return** mean$(\widetilde{S})$

$\quad$ **else**

$\quad\quad$ Let $\Sigma_{\widetilde{S}}$ be the covariance matrix of $\widetilde{S}$.

$\quad\quad$ Let V be the span of the top $p/2$ principal components of $\Sigma_{\widetilde{S}}$ and W be its complement.

$\quad\quad$ Set $S_1 := P_V(\widetilde{S})$ where P_V is the projection operation on to V.

$\quad\quad$ Let $\widehat{\mu}_V := $ HUBERGRADIENTESTIMATOR$(S_1, \epsilon, p/2, \delta)$.

$\quad\quad$ Let $\widehat{\mu}_W := $ mean$(P_W \widetilde{S})$.

$\quad\quad$ Let $\widehat{\mu} \in \mathbb{R}^p$ be such that $P_V(\widehat{\mu}) = \widehat{\mu}_V$, and $P_W(\widehat{\mu}) = \widehat{\mu}_W$.

$\quad$ **return** $\widehat{\mu}$.

$\quad$ **end if**

end function

3.2 Gradient Estimation in the Heavy-Tailed model

To design gradient estimators for the heavy-tailed model, we leverage recent work on designing robust *mean* estimators in this setting. These robust mean estimators build on the classical work of Alon et al. [39], Nemirovski and Yudin [40] and Jerrum et al. [41] on the so-called median-of-means estimator. For the problem of one-dimensional mean estimation, Lerasle and Oliveira [31], Catoni et al. [32] propose robust mean estimators that achieve exponential concentration around the true mean for any distribution with bounded second moment. In this work we require mean estimators for multivariate distributions. Several recent works ([15, 33, 38]) extend the median-of-means estimator of to general metric spaces. In this paper we use the geometric median-of-means estimator (GMOM), which was originally proposed and analyzed by Minsker [15], to design the gradient estimator $g(\theta; \mathcal{D}_n, \delta)$.

The basic idea behind the GMOM estimator is to first split the samples into non-overlapping subsamples and estimate the sample mean of each of the subsamples. Then the GMOM estimator is given by the median-of-means of the subsamples. Formally, let $\{x_i \ldots x_n\} \in \mathbb{R}$ be n i.i.d random variables sampled from a distribution P. Then the GMOM estimator for estimating the mean of P can be described as follows. Partition the n samples into b blocks $B_1, \ldots B_b$ each of size $\lfloor n/b \rfloor$. Let $\{\widehat{\mu}_1, \ldots, \widehat{\mu}_b\}$ be the sample means in each block, where $\widehat{\mu}_i = \frac{1}{|B_i|} \sum_{x_j \in B_i} x_j$. Then the GMOM estimator is given by median$\{\widehat{\mu}_1, \ldots \widehat{\mu}_b\}$. In high dimensions where different

notions of the median have been considered Minsker [15] uses geometric median:

$$\widehat{\mu} = \operatorname*{argmin}_{\mu} \sum_{i=1}^{b} \|\mu - \widehat{\mu}_i\|_2.$$

Algorithm 3 presents the gradient estimator $g(\theta; \mathcal{D}_n, \delta)$ obtained using GMOM as the mean estimator.

Algorithm 3 Heavy Tailed Gradient Estimator

function HEAVYTAILEDGRADIENTESTIMATOR(SAMPLE GRADIENTS $S = \{\nabla\overline{\mathcal{L}}(\theta; z_i)\}_{i=1}^{n}$, δ)

    Define number of buckets $b = 1 + \lfloor 3.5 \log 1/\delta \rfloor$.
    Partition S into b blocks $B_1, \dots B_b$ each of size $\lfloor n/b \rfloor$.
    for $i = 1 \dots n$ **do**
        $\widehat{\mu}_i = \frac{1}{|B_i|} \sum_{s \in B_i} s.$
    end for
    Let $\widehat{\mu} = \operatorname*{argmin}_{\mu} \sum_{i=1}^{b} \|\mu - \widehat{\mu}_i\|_2.$
    return $\widehat{\mu}$.
end function

3.3 Choice of Hyper-Parameters

In this section, we discuss how to tune the hyperparameters for our algorithms. In particular, note that the gradient estimators described in Algorithms 2, 3 depend on corruption level ϵ, and on confidence δ, which are not known in advance.

Since the standard hyper-parameter selection techniques such as cross validation, hold-out validation, pick hyper-parameters that minimize the empirical mean of the loss on validation data, they can't be used in the presence of outliers in the data. One criteria we could use in such cases is to choose hyper-parameters that minimize a robust estimate of the population risk on validation data. However, we can't use any of the existing robust mean estimators to estimate the population risk because they themselves depend on hyper-parameters such as corruption level ϵ.

Huber Contamination. We now consider the Huber contamination model and propose a heuristic based on Scheffe's tournament estimator [6, 42] for hyper-parameter selection. In particular, we consider the gradient descent procedure described in Algorithm 2 and explain our technique for choosing ϵ, δ using hold out cross validation. Note that our goal is to pick hyper-parameters that minimize the population risk $\mathcal{R}(\theta)$. Under the assumption of strong convexity of $\mathcal{R}(\theta)$, this is equivalent to picking hyper-parameters that minimize the parameter error $\|\theta - \theta^*\|_2$.

We begin with the problem of density estimation, where we are given n i.i.d samples $\{z_i\}_{i=1}^{n}$ from $(1 - \epsilon)P_{\theta^*} + \epsilon Q$, where P_{θ^*} belongs to the class of distributions $\{P_\theta : \theta \in \Theta\}$, and Q is an arbitrary distribution. Our goal is to estimate $\theta^* \in \Theta$ from the samples. Suppose $\left\{P_{\widehat{\theta}_1}, P_{\widehat{\theta}_2}, \dots, P_{\widehat{\theta}_m}\right\}$ are the candidate solutions returned by Algorithm 2 for different settings

of ϵ, δ. Consider the following pairwise test function:

$$\phi_{jk} = \mathbb{I}\left\{ \left| \frac{1}{n_{val}} \sum_{i=1}^{n_{val}} \mathbb{I}\left\{ p_{\widehat{\theta}_j}(z_i') > p_{\widehat{\theta}_k}(z_i') \right\} - P_{\widehat{\theta}_j}\left(p_{\widehat{\theta}_j}(z) > p_{\widehat{\theta}_k}(z) \right) \right| > \right.$$
$$\left. \left| \frac{1}{n_{val}} \sum_{i=1}^{n_{val}} \mathbb{I}\left\{ p_{\widehat{\theta}_j}(z_i') > p_{\widehat{\theta}_k}(z_i') \right\} - P_{\widehat{\theta}_k}\left(p_{\widehat{\theta}_j}(z) > p_{\widehat{\theta}_k}(z) \right) \right| \right\}, \tag{13}$$

where $p_{\widehat{\theta}_j}$ is the probability density of $P_{\widehat{\theta}_j}$, $\{z_i'\}_{i=1}^{n_{val}}$ is the validation data and $\mathbb{I}\{.\}$ is the indicator function. When $\phi_{jk} = 1$, then $\widehat{\theta}_k$ is favored over $\widehat{\theta}_j$ and when $\phi_{jk} = 0$, then $\widehat{\theta}_j$ is favored over $\widehat{\theta}_k$. Then, the final estimate $P_{\widehat{\theta}_{j*}}$ is given by

$$j^* = \operatorname*{argmin}_{j \in [m]} \sum_{\substack{k=1 \\ k \neq j}}^{m} \phi_{jk}$$

Following [42], it can be shown that the above procedure picks a j^* such that $P_{\widehat{\theta}_{j*}}$ is close to P_{θ^*} in TV metric. For distributions $\{P_\theta : \theta \in \Theta\}$ whose TV metric is roughly equivalent to the parameter error, the above procedure results in hyper-parameters which minimize the parameter error. This procedure can be extended to supervised learning problems such as regression and classification.

Heavy-Tailed Distribution. For the Heavy-Tailed setting we experimented with (a) empirical mean of the risk on validation data: $\frac{1}{n_{val}} \sum_{i=1}^{n_{val}} \overline{\mathcal{L}}(\theta; z_i')$ where $\{z_i'\}_{i=1}^{n_{val}}$ is the validation data, which does not require any tuning parameters, as well as (b) median of means based mean of the risk on validation data, for various confidence levels δ. However, both the techniques in the context of hold-out validation resulted in models with similar performance. So, in our experiments with heavy tailed distributions, we present results obtained using the empirical risk as in (a) on hold-out validation data.

4 Experiments

In this section we demonstrate our proposed methods for the Huber contamination and heavy-tailed models, on a variety of simulated and real data examples.

4.1 Huber Contamination

We first consider the Huber contamination model and demonstrate the practical utility of gradient-descent based robust estimator described in Algorithms 1 and 2.

4.1.1 Synthetic Experiments: Linear Regression

Recall the linear regression model described in (4) where we observe paired samples $\{(x_i, y_i)\}_{i=1}^{n}$. We assume that the (x, y) pairs sampled from the true distribution P are linked via a linear model:

$$y = \langle x, \theta^* \rangle + w.$$

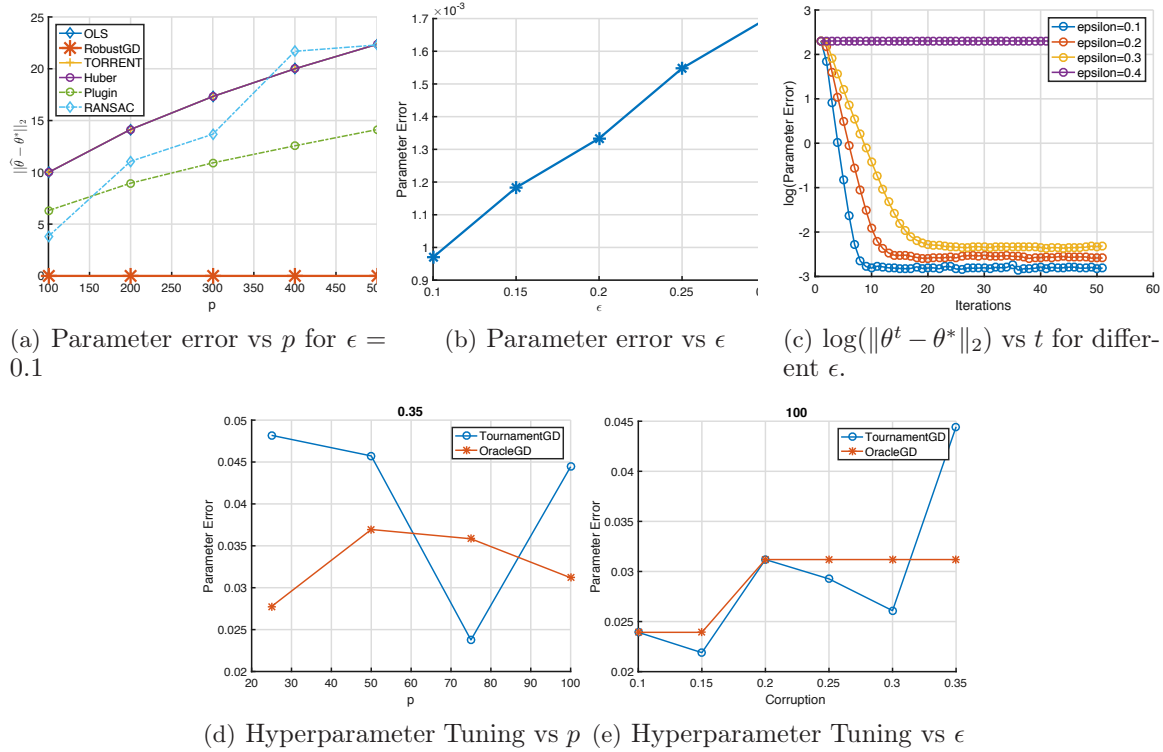We now describe the experiment setup, the data model and present the results.

(a) Parameter error vs p for $\epsilon = 0.1$

(b) Parameter error vs ϵ

(c) $\log(\|\theta^t - \theta^*\|_2)$ vs t for different ϵ.

(d) Hyperparameter Tuning vs p (e) Hyperparameter Tuning vs ϵ

Figure 1: Robust Linear Regression.

Setup. We fix the contamination level $\epsilon = 0.1$ and $\sigma^2 = 0.1$. Next, we generate $(1 - \epsilon)n$ *clean* covariates from a multivariate Gaussian $x \sim \mathcal{N}(0, \mathcal{I}_p)$, and we generate the corresponding clean responses using $y = \langle x, \theta^* \rangle + w$ where $\theta^* = [1, \ldots, 1]^T$ and $w \sim \mathcal{N}(0, \sigma^2)$. We simulate an outlier distribution by drawing the covariates from $\mathcal{N}(0, p^2 \mathcal{I}_p)$, and setting the responses to 0. The total number of samples is set to be $(10\frac{p}{\epsilon^2})$. We note that the sample size we choose increases with the dimension. This scaling is used to ensure that the statistical (minimax) error, in the absence of any contamination, is roughly 0.001. An optimally robust method should have error close to 0.1 (roughly equal to corruption level), which ours does (see Figure 1).

Metric. We measure the parameter error in ℓ_2-norm. We also study the convergence properties of our proposed method, for different contamination levels ϵ. We use code provided by Lai et al. [14] to implement our gradient estimator.

Baselines. As our baselines, we use OLS, TORRENT [16], the Huber-loss M-estimator, RANSAC and a plugin estimator (detailed further in Section sec:linregtheory, and which in a nutshell robustly estimates the sufficient statistics required for the OLS estimator). TORRENT is an iterative hard-thresholding based alternating minimization algorithm, where in one step, it calculates an active set of examples by keeping only $(1-\epsilon)n$ samples which have the smallest absolute values of residual $r = y - \langle x, \theta^t \rangle$, and in the other step it updates the current estimates by solving OLS on the active set. Bhatia et al. [16] have shown the superiority of TORRENT over a variety of other convex-penalty based outlier techniques, hence, we do not compare against those methods. The plugin estimator is implemented using Algorithm 2 to estimate both the mean vector $\frac{1}{n}\sum_{i=1}^{n} y_i x_i$ and the covariance matrix $\frac{1}{n}\sum_{i=1}^{n} x_i x_i^T$, which are

the required sufficient statistics for the OLS estimator.

Results. We summarize our main findings here.

- **Error vs dimension p:** All estimators except our proposed algorithm perform poorly with increasing dimension, as shown in Figure 1(a). Note that the TORRENT algorithm has strong guarantees when only the response y is corrupted but performs poorly in the Huber contamination model where both x and y may be contaminated. We find that the error for the robust plugin estimator increases with dimension. We investigate this theoretically in Section 6.1, where we find that the error of the plugin estimator grows with the norm of θ^*. In our experiments, we choose $\|\theta^*\|_2 = \sqrt{p}$, and thus Figure 1(a) corroborates Corollary 3 in Section 6.1.

- **Error vs ϵ:** In Figure 1(b) we find that the parameter error $\|\hat{\theta} - \theta^*\|_2$ increases linearly with the contamination rate ϵ and we study this further in Section 6.1.

- **Error vs iteration t:** Finally, Figure 1(c) shows that the convergence rate decreases with increasing contamination ϵ and after ϵ is high enough, the algorithm remains stuck at θ_0, corroborating Lemma 8 (in the Appendix).

- **Hyper-parameter Tuning:** In Figures 1(d) and 1(e), we find the final solution chosen by our tournament based heuristic for hyper-parameter selection (TournamentGD) has roughly the same performance as the algorithm which knows the true value of ϵ (OracleGD). In particular, our final error does not scale with p.

Next, we study the performance of our proposed method in the context of classification.

4.1.2 Synthetic Experiments: Logistic Regression

(a) 0-1 Error vs p at $\epsilon = 0.1$ (b) 0-1 error vs ϵ (c) $\log(\|\theta^t - \theta^*\|_2)$ vs t for different ϵ
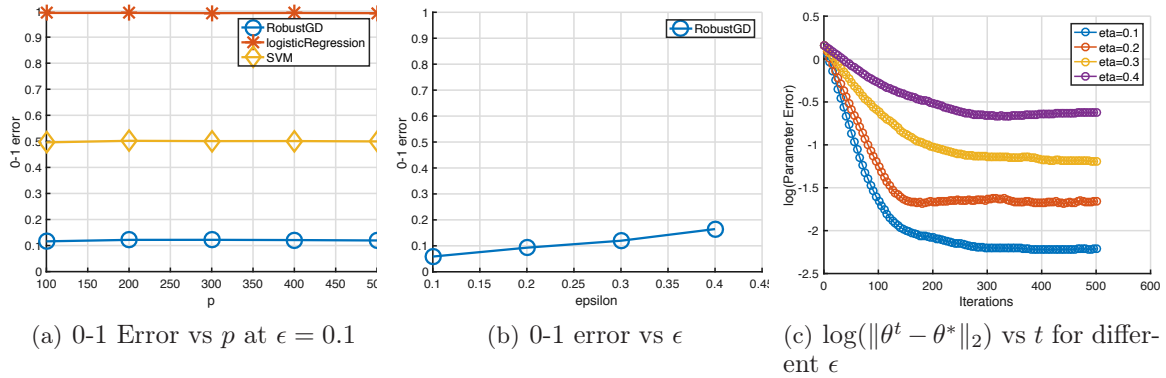
Figure 2: Robust Logistic Regression.

Setup. We simulate a linearly separable classification problem, where the *clean* covariates are sampled from $\mathcal{N}(0, \mathcal{I}_p)$, the corresponding clean responses are computed as $y = \text{sign}(\langle x, \theta^* \rangle)$ where $\theta^* = [1/\sqrt{p}, \ldots, 1/\sqrt{p}]^T$. We simulate the outlier distribution by adding asymmetric noise, *i.e.* we flip the labels of one class, and increase the variance of the corresponding covariates by multiplying them by p^2. The total number of samples are set to be $(10\frac{p}{\epsilon^2})$.

14

Table 1: Fitting to original image error.

	Best Possible	Proposed	TORRENT	OLS	SCRRR
Mean RMSE	0.05	0.09	0.175	0.21	0.13

Metric. We measure the 0-1 classification error on a held-out (clean) test set. We study how the 0-1 error changes with p and ϵ and the parameter estimation error of our proposed method for different contamination levels ϵ.

Baselines. We use the logistic regression MLE and the linear Support Vector Machine (SVM) as our baselines.

Results. We summarize our main findings below:

- **0/1 Error vs dimension p:** In Figure 2(a) we observe that both the SVM and logistic regression MLE perform poorly with increasing dimension. The logistic regression MLE completely flips the labels and has a 0-1 error close to 1, whereas the linear SVM outputs a random hyperplane classifier that flips the label for roughly half of the dataset.

- **0/1 Error vs ϵ and t:** Figures 2(b) and 2(c) show qualitatively similar results to the linear regression setting, i.e. that the error of our proposed estimator degrades gracefully (and grows linearly) with the contamination level ϵ and that the gradient descent iterates converge linearly.

4.1.3 Robust Face Reconstruction

Setup. In this experiment, we show the efficacy of our algorithm by attempting to reconstruct face images that have been corrupted with heavy occlusion, where the occluding pixels play the role the outliers. We use the data from the Cropped Yale Dataset [43] . The dataset contains 38 subjects, and each image has 192×168 pixels. Following the methodology of Wang et al. [44], we choose 8 face images per subject, taken under mild illumination conditions and computed an eigenface set with 20 eigenfaces. Then given a new corrupted face image of a subject, the goal is to get the best reconstruction/approximation of the true face. To remove scaling effects, we normalized all images to $[0, 1]$ range. One image per person was used to test reconstruction. Occlusions were simulated by randomly placing 10 blocks of size 30×30. We repeated this 10 times for each test image. In this setting, each image of a subject corresponds to a different task; *i.e.* X is a common fixed eigenface basis, y is an observed(occluded) image, and the goal is to reconstruct(de-noise) the given image using the given basis. Note that in this example, we use a linear regression model as the uncontaminated statistical model, which is almost certainly not an exact match for the unknown ground truth distribution. Despite this model misspecification, as our results show, that robust mean based gradient algorithms do well.

Metric. We use Root Mean Square Error (RMSE) between the original and reconstructed image to evaluate the performance of the algorithms. We also compute the best possible reconstruction of the original face image by using the 20 eigenfaces.

Methods. Wang et al. [44] implemented popular robust estimators such as RANSAC, Huber Loss *etc.* and showed their poor performance. Wang et al. [44] then proposed an alternate robust regression algorithm called Self Scaled Regularized Robust Regression(SCRRR). Hence, use TORRENT, SCRRR and OLS as baselines. We also compare against the best possible RMSE obtained by reconstructing the un-occluded image using the eigenfaces.

Results. Table 1 shows that the mean RMSE is best for our proposed gradient descent based method and that the recovered images are in most cases closer to the un-occluded original image. (Figure 4.1.3). Figure 3(c) shows a case when none of the methods succeed in
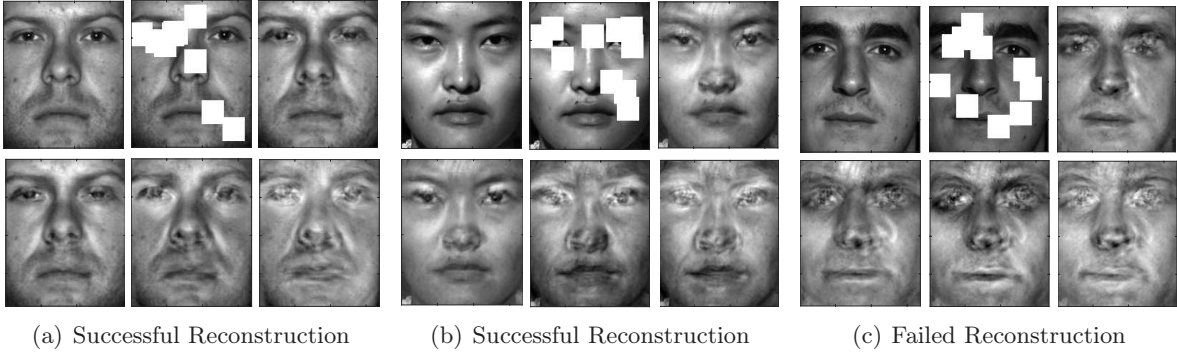

Figure 3: Robust Face recovery results: Top; in order from L to R: original image, occluded image, best possible recovery with given basis. Bottom; in order from L to R: Reconstructions using our proposed algorithm, TORRENT and ordinary least squares (OLS).

4.2 Heavy-tailed Estimation

We now consider the heavy-tailed model and present experimental results on synthetic datasets comparing the gradient descent based robust estimator described in Algorithms 1 and 3 (which we call RobustGD) with ERM and several other recent proposals. In these experiments we focus on the problem of linear regression which is described in Section 4.1 and work with heavy-tailed noise distributions.

4.2.1 Synthetic Experiments: Simple Linear Regression

Setup. The covariate $x \in \mathbb{R}^p$ is sampled from a zero-mean isotropic gaussian distribution. We set each entry of θ^* to $1/\sqrt{p}$. The noise w is sampled from a Pareto distribution, with mean zero, variance σ^2 and tail parameter β. The tail parameter β determines the moments of the Pareto random variable. More specifically, the moment of order k exists only if $k < \beta$, hence, smaller the β the more heavy-tailed the distribution. In this setup, we keep the dimension p fixed to 128 and vary n, σ and β. We always maintain the sample-size n to be at least $4p$.

Methods. We compare RobustGD with several baselines. Since we are always in the low-dimensional ($n \geq p$) setting, the solution to ERM has a closed form expression and is simply the OLS solution. We also study OLS-GD, which performs a gradient descent on ERM and is equivalent to using empirical mean as the gradient oracle in our framework. We also compare against the robust estimation techniques of Hsu and Sabato [33] and Duchi and Namkoong [45], which we refer to as RobustHS, RobustDN and two classical techniques namely the

LASSO [46] and ridge regression. In our experiments, all the iterative techniques are run until convergence.

Metrics. We use two metrics to compare the performance of various approaches: a) parameter error which is defined as $\|\theta - \theta^*\|_2$ and b) to compare the performance of two estimators $\widehat{\theta}_1, \widehat{\theta}_2$, we define the notion of relative efficiency:

$$\text{RELEFF}(\widehat{\theta}_1, \widehat{\theta}_2) = \frac{\|\widehat{\theta}_2 - \theta^*\|_2 - \|\widehat{\theta}_1 - \theta^*\|_2}{\|\widehat{\theta}_1 - \theta^*\|_2}.$$

Roughly, this corresponds to the percentage improvement in the parameter error obtained using $\widehat{\theta}_1$ over $\widehat{\theta}_2$. Whenever $\text{RELEFF}(\widehat{\theta}_1, \widehat{\theta}_2) > 0$, $\widehat{\theta}_1$ has a lower parameter error, and higher the value, the more the fractional improvement.

Results. To reduce the variance in the plots presented here, we averaged results over 20 repetitions. Figure 4 shows the benefits of using RobustGD over other baselines.

- **Error vs number of iterations:** In Figures 4(a), 4(b) we plot the excess risk of various approaches against the number of iterations (for OLS, LASSO, ridge regression and the method of Hsu and Sabato [33] we only plot the excess risk of the final iterate). We see that upon convergence RobustGD has a much lower parameter error. As expected, OLS-GD converges to OLS.

- **Error vs number of samples:** Next, in Figures 4(c), 4(d) we plot the parameter error as n/p increases. We see that RobustGD is always better than other baselines, even when the number of samples is 12 times the dimension p.

- **Relative Efficiency vs β, and σ:** In Figure 4(e), we plot the relative efficiency against β, the moment bound of Pareto distribution. This shows that the percentage improvement in the excess risk by RobustGD decreases as the moment bound β increases. This behavior is expected because as we increase the moment bound the tails of the noise distribution become lighter. This shows that there is more benefit in using RobustGD in the heavy tailed setting. We do a similar study to see the relative efficiency against the variance of the noise distribution. Figure 4(f) plots relative efficiency against standard deviation of the noise distribution.

5 Theoretical Preliminaries

In this section we develop some theoretical preliminaries. We first develop a general theory on convergence of projected gradient descent in Section 5.1. Next we analyze the gradient estimators defined in Algorithms 2 and 3 in Sections 5.2 and 5.3 respectively. Finally in Sections 6 and 7 we present consequences of our general theory for the canonical examples of risk minimization described in Section 2.2, under Huber contamination and heavy-tailed models.

For some of our examples, we will assume certain mild moment conditions. Concretely, for a random vector $x \in \mathbb{R}^p$, let $\mu = \mathbb{E}[x]$ and Σ be the covariance matrix. Then x has bounded $2k^{\text{th}}$ moments if there exists a constant C_{2k} such that for every unit vector v we have that

$$\mathbb{E}\left[\langle x - \mu,\ v\rangle^{2k}\right] \leq C_{2k}\left(\mathbb{E}\left[\langle x - \mu,\ v\rangle^2\right]\right)^k. \tag{14}$$

(a) $n = 512, p = 128, \sigma = 0.75. \beta = 3$

(b) $n = 1024, p = 128, \sigma = 0.75, \beta = 3$

(c) $\sigma = 0.75, \beta = 3$

(d) $\sigma = 1, \beta = 3$

(e) $n = 512, p = 128, \sigma = 0.75$

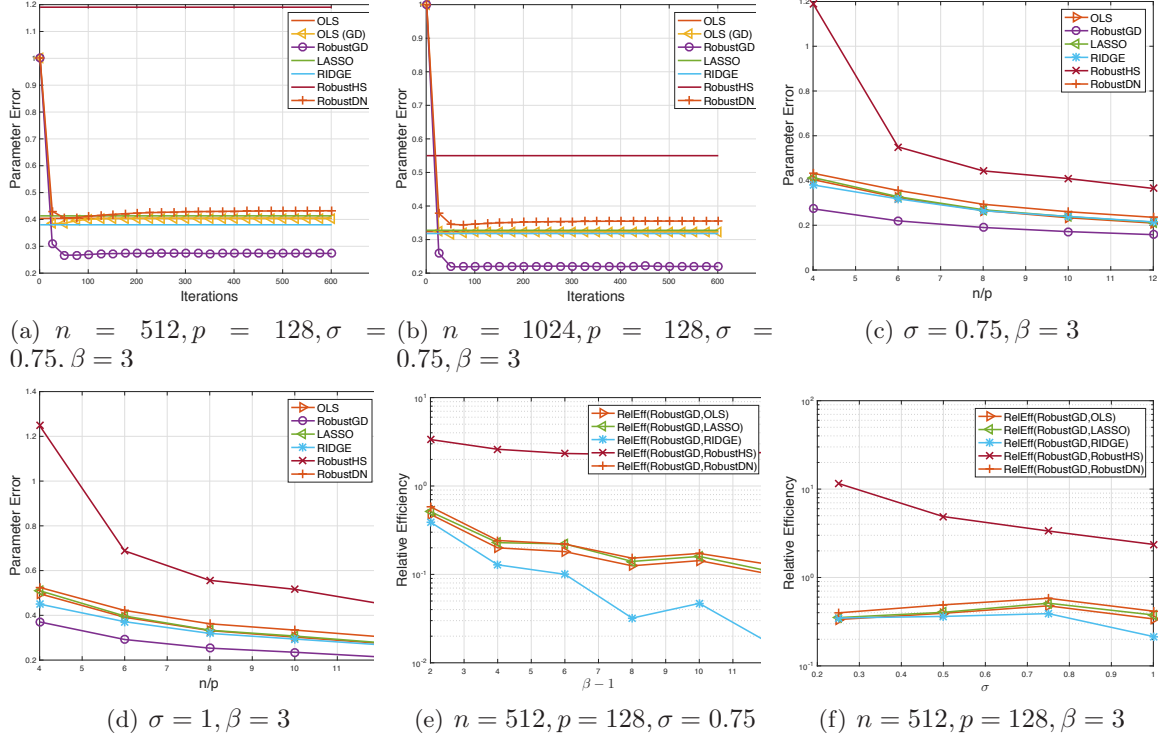(f) $n = 512, p = 128, \beta = 3$

Figure 4: Linear Regression: Performance comparison of RobustGD against baselines.

5.1 Stability of Gradient Descent

In this section we develop a general theory for the convergence of the projected gradient descent described in Algorithm 1. Note that our gradient estimators could be biased and are not guaranteed to be consistent estimators of the true gradient $\nabla \mathcal{R}(\theta)$. This is especially true in the Huber contamination model where it is impossible to obtain consistent estimators of the gradient of the risk because of the non-vanishing bias caused by the contaminated samples. Hence, we turn our attention to understanding the behavior of projected gradient descent with a biased, inexact, gradient estimator of the form in (11). Before we present our main result, we define the notion of stability of a gradient estimator, which plays a key role in the convergence of gradient descent.

Definition 2 (Stability). *A gradient estimator is* stable *for a given risk function* $\mathcal{R} : \Theta \mapsto \mathbb{R}$ *if for some* $\phi \in [0, \tau_\ell)$,

$$\alpha(\widetilde{n}, \widetilde{\delta}) < \tau_\ell - \phi.$$

We denote by κ the following contraction parameter:

$$\kappa := \sqrt{1 - \frac{2\eta\tau_\ell\tau_u}{\tau_\ell + \tau_u}} + \eta\alpha(\widetilde{n}, \widetilde{\delta}),$$

and note that $\kappa < 1$. With these definitions in place we state our main result on the stability of gradient descent:

Theorem 1. *Suppose that the gradient estimator satisfies the condition in* (11) *and is stable for the risk function* $\mathcal{R} : \Theta \mapsto \mathbb{R}$. *Then Algorithm 1 initialized at* θ^0 *with step-size* $\eta = 2/(\tau_\ell + \tau_u)$,

returns iterates $\{\widehat{\theta}^t\}_{t=1}^T$ such that with probability at least $1 - \delta$ for the contraction parameter κ above we have that,

$$\|\widehat{\theta}^t - \theta^*\|_2 \leq \kappa^t \|\theta^0 - \theta^*\|_2 + \frac{1}{1-\kappa}\beta(\widetilde{n}, \widetilde{\delta}). \tag{15}$$

We defer a proof of this result to the Appendix. For the bound (15), the first term is decreasing in T, while the second term is increasing in T. This suggests that for a given n and δ, we need to run just enough iterations for the first term to be bounded by the second. Hence, we can fix the number of iterations T^* as the smallest positive integer such that:

$$T \geq \log_{1/\kappa} \frac{(1-\kappa)\|\theta^0 - \theta^*\|_2}{\beta(\widetilde{n}, \widetilde{\delta})}.$$

Since we obtain linear convergence, i.e. $\kappa < 1$, typically a logarithmic number of iterations suffice to obtain an accurate estimate.

Theorem 1 provides a general result for risk minimization and parameter estimation, and requires bounds on $\alpha(\widetilde{n}, \widetilde{\delta}), \beta(\widetilde{n}, \widetilde{\delta})$ which capture the the error suffered by the gradient estimator for a given risk minimization problem. In any concrete instantiation for a given gradient estimator, risk pair, we first estimate these gradient estimator error bounds by studying the distribution of the gradient of the risk, and then apply Theorem 1. In the next two sections, we provide some general analyses of the gradient estimator in Algorithm 2 for the Huber contamination model, and the gradient estimator in Algorithm 3 for the heavy-tailed model, and which apply to any risk minimization problem. In Sections 6,7 we then instantiate these gradient estimator error results for various illustrative statistical models such as linear regression, logistic regression, and general exponential families. Plugging these into Theorem 1, we then get consequences of our robustness guarantees for various statistical model, robustness setting pairs.

5.2 General Analysis of Huber Contamination Gradient Estimator in Algorithm 2

We now analyze the gradient estimator described in Algorithm 2 for Huber contamination model and study the error suffered by it. As stated before, Algorithm 2 uses the robust mean estimator of Lai et al. [14]. Hence, while our proof strategy mimics that of Lai et al. [14], we present a different result which is obtained by a more careful non-asymptotic analysis of the algorithm.

We define:

$$\gamma(n, p, \delta, \epsilon) := \left(\frac{p \log p \log\left(n/(p\delta)\right)}{n} \right)^{3/8} + \left(\frac{\epsilon p^2 \log p \log\left(\frac{p \log(p)}{\delta}\right)}{n} \right)^{1/4}, \tag{16}$$

and with this definition in place we have the following result:

Lemma 1. *Let P be the true probability distribution of z and let P_θ be the true distribution of the gradients $\nabla\overline{\mathcal{L}}(\theta; z)$ on $\mathbb{R}^p$ with mean $\mu_\theta = \nabla\mathcal{R}(\theta)$, covariance Σ_θ, and bounded fourth moments. There exists a positive constant $C_1 > 0$, such that given n samples from the distribution in (9), the Huber Gradient Estimator described in Algorithm 2 when instantiated with the contamination level ϵ, with probability at least $1 - \delta$, returns an estimate $\widehat{\mu}$ of μ_θ such that,*

$$\|\widehat{\mu} - \mu_\theta\|_2 \leq C_1\left(\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon)\right)\|\Sigma_\theta\|_2^{\frac{1}{2}}\sqrt{\log p}.$$

We note in particular, if $n \to \infty$ (with other parameters held fixed) then $\gamma(n, p, \delta, \epsilon) \to 0$ and the error of our gradient estimator satisfies

$$\|\widehat{\mu} - \mu_\theta\|_2 \leq C\sqrt{\|\Sigma_\theta\|_2 \epsilon \log p},$$

and has only a weak dependence on the dimension p.

5.3 General Analysis of Heavy-tailed Model Gradient Estimator in Algorithm 3

In this section we analyze the gradient estimator for heavy-tailed setting, described in Algorithm 3. The following result shows that the gradient estimate has exponential concentration around the true gradient, under the mild assumption that the gradient distribution has bounded second moment. Its proof follows directly from the analysis of geometric median-of-means estimator of Minsker [15]. We use $\text{tr}(\Sigma_\theta)$ to denote the trace of the matrix Σ_θ.

Lemma 2. *Let P be the probability distribution of z and P_θ be the distribution of the gradients $\nabla\overline{\mathcal{L}}(\theta; z)$ on $\mathbb{R}^p$ with mean $\mu_\theta = \nabla\mathcal{R}(\theta)$, covariance Σ_θ. Then the heavy tailed gradient estimator described in Algorithm 3 returns an estimate $\widehat{\mu}$ that satisfies the following exponential concentration inequality, with probability at least $1 - \delta$:*

$$\|\widehat{\mu} - \mu_\theta\|_2 \leq 11\sqrt{\frac{tr(\Sigma_\theta)\log(1.4/\delta)}{n}}.$$

The results of the Lemmas 1 and 2 effectively ensure that under relatively mild moment assumptions we can robustly estimate multivariate mean vectors and in subsequent sections we show how to leverage these strong guarantees for robust parametric estimation.

6 Consequences for Estimation under ϵ-Contaminated Model

We now turn our attention to the examples introduced earlier, and present specific applications of Theorem 1, for parametric estimation under Huber contamination model. As shown in Lemma 1, we need the added assumption that the true gradient distribution has bounded fourth moments, which suggests the need for additional assumptions. We make our assumptions explicit and defer the technical details to the Appendix.

6.1 Linear Regression

We assume that the covariates $x \in \mathbb{R}^p$ have bounded 8^{th}-moments and the noise w has bounded 4^{th} moments.

Theorem 2 (Robust Linear Regression)**.** *Consider the statistical model in equation (4), and suppose that the number of samples n is large enough such that $\gamma(\widetilde{n}, p, \widetilde{\delta}) < \frac{C_1\tau_\ell}{\|\Sigma\|_2\sqrt{\log p}}$ and the contamination level is such that*

$$\epsilon < \left(\frac{C_2\tau_\ell}{\|\Sigma\|_2\sqrt{\log p}} - \gamma(\widetilde{n}, p, \widetilde{\delta})\right)^2,$$

for some constants C_1 and C_2. Then, there are universal constants C_3, C_4, such that if Algorithm 1 is initialized at θ^0 with stepsize $\eta = 2/(\tau_u + \tau_\ell)$ and Algorithm 2 as gradient estimator,

then it returns iterates $\{\widehat{\theta}^t\}_{t=1}^T$ such that for a contraction parameter $\kappa < 1$, with probability at least $1 - \delta$,

$$\|\widehat{\theta}^t - \theta^*\|_2 \leq \kappa^t \|\theta^0 - \theta^*\|_2 + \frac{C_3 \sigma \sqrt{\|\Sigma\|_2 \log p}}{1 - \kappa}\left(\epsilon^{\frac{1}{2}} + \gamma(\widetilde{n}, p, \widetilde{\delta})\right). \tag{17}$$

In the asymptotic setting when the number of samples $n \to \infty$ (and other parameters are held fixed), we see that for the Huber Gradient Estimator, the corresponding maximum allowed contamination level is

$$\epsilon < \frac{C_1 \tau_\ell^2}{\tau_u^2 \log p},$$

i.e. the better conditioned the covariance matrix Σ, the higher the contamination level we can tolerate.

Plugin Estimation. For linear regression, the true parameter can be written in closed form as $\theta^* = \mathbb{E}[xx^T]^{-1}\mathbb{E}[xy]$. A non-iterative way to estimate θ^* is to separately estimate $\mathbb{E}[xx^T]$ and $\mathbb{E}[xy]$ using robust covariance and mean oracles respectively. Under the assumption that $x \sim \mathcal{N}(0, \mathcal{I}_p)$, one can reduce the problem to robustly estimating $\mathbb{E}[xy]$. Under this setting, we now present a result using our robust mean estimator (from Lemma 1) to directly estimate $\mathbb{E}[xy]$. Recall, the definition of γ in (16). We have the following result:

Corollary 3. *Consider the model in* (4) *with the covariates drawn from* $\mathcal{N}(0, \mathcal{I}_p)$ *and* $w \in \mathcal{N}(0, 1)$, *then there are universal constants* C_1, C_2 *such that if* $\epsilon < C_1$, *the plugin estimator* $\widehat{\theta}$ *of* $\mathbb{E}[xy]$ *described above with probability at least* $1 - \delta$ *satisfies:*

$$\|\widehat{\theta} - \theta^*\|_2 \leq C_2 \sqrt{(1 + 2\|\theta^*\|_2^2) \log p}\left(\epsilon^{\frac{1}{2}} + \gamma(n, p, \delta, \epsilon)\right). \tag{18}$$

Comparing bounds (17) and (18), we see that even when it does not have to estimate the covariance matrix, the error of the plugin estimator depends on $\|\theta^*\|_2$, which would make the estimator vacuous if $\|\theta^*\|_2$ scales with the dimension p. On the other hand, the asymptotic rate of our robust gradient estimator is independent of $\|\theta^*\|_2$. This disadvantage of plugin estimation is inescapable, and is seen for instance in known minimax results for robust mean estimation [10] that show that the dependence on $\|\theta^*\|_2$ is unavoidable for any oracle which estimates the mean of xy in the ϵ-contaminated setting. Next, we apply our estimator to generalized linear models.

6.2 Generalized Linear Models

Here we assume that the covariates have bounded 8^{th} moments. Additionally, we assume smoothness of $\Phi'(\cdot)$ around θ^*. In particular, we assume that there exist universal constants $L_{\Phi, 2k}$, B_{2k} such that

$$\mathbb{E}_x\left[\left|\Phi'(\langle x, \theta\rangle) - \Phi'(\langle x, \theta^*\rangle)\right|^{2k}\right] \leq L_{\Phi, 2k}\|\theta^* - \theta\|_2^{2k} + B_{\Phi, 2k}, \quad \text{for } k = 1, 2, 4$$

We also assume that $\mathbb{E}_x[|\Phi^{(t)}(\langle x, \theta^*\rangle)|^k] \leq M_{\Phi, t, k}$ where $\Phi^{(t)}(\cdot)$ is the t^{th}-derivative of $\Phi(\cdot)$.

Theorem 4 (Robust Generalized Linear Models)**.** *Consider the statistical model in equation* (5), *and suppose that the number of samples n is large enough such that*

$$\gamma(\widetilde{n}, p, \widetilde{\delta}) < \frac{C_1 \tau_\ell}{\sqrt{\log p}\|\Sigma\|_2^{\frac{1}{2}}[L_{\Phi,4}^{\frac{1}{4}} + L_{\Phi,2}^{\frac{1}{2}}]},$$

and the contamination level is such that,

$$\epsilon < \left(\frac{C_2 \tau_\ell}{\sqrt{\log p}\|\Sigma\|_2^{\frac{1}{2}}[L_{\Phi,4}^{\frac{1}{4}} + L_{\Phi,2}^{\frac{1}{2}}]} - \gamma(\widetilde{n}, p, \widetilde{\delta})\right)^2,$$

for some constants C_1 and C_2. Then, there are universal constants C_3, C_4, such that if Algorithm 1 *is initialized at θ^0 with stepsize $\eta = 2/(\tau_u + \tau_\ell)$ and Algorithm* 2 *as gradient estimator, then it returns iterates $\{\widehat{\theta}^t\}_{t=1}^T$ such that with probability at least $1 - \delta$*

$$
\begin{aligned}
\|\widehat{\theta}^t - \theta^*\|_2 \leq &\kappa^t \|\theta^0 - \theta^*\|_2 \\
&+ \frac{C_3\sqrt{\log p}\|\Sigma\|_2^{\frac{1}{2}}[B_{\Phi,4}^{\frac{1}{4}} + B_{\Phi,2}^{\frac{1}{2}} + c(\sigma)^{\frac{1}{2}}M_{\Phi,2,2}^{\frac{1}{4}} + c(\sigma)^{\frac{3}{4}}M_{\Phi,4,1}^{\frac{1}{4}}]}{1 - \kappa}\left(\epsilon^{\frac{1}{2}} + \gamma(\widetilde{n}, p, \widetilde{\delta})\right),
\end{aligned}
\tag{19}
$$

for some contraction parameter $\kappa < 1$.

Note that for the case of linear regression with gaussian noise, it is relatively straightforward to see that $L_{\Phi,2k} = C_{2k}\|\Sigma\|_2^k$, $B_{\Phi,2k} = 0$, $M_{\Phi,t,k} = 1 \ \ \forall (t \geq 2, k \in \mathcal{N})$ and $M_{\Phi,t,k} = 0 \ \ \forall (t \geq 3, k \in \mathcal{N})$ under the assumption of bounded 8^{th} moments of the covariates; which essentially leads to an equivalence between Theorem 2 and Theorem 4 for this setting. In the following section, we instantiate the above Theorem for logistic regression and compare and contrast our results to other existing methods.

6.2.1 Logistic Regression

By observing that $\Phi^{(t)}(\cdot)$ is bounded for logistic regression for all $t \geq 1$, we can see that $L_{\Phi,2k} = 0$, and that there exists a universal constant $C > 0$ such that $B_{\Phi,2k} < C$ and $M_{\Phi,t,k} < C \ \ \forall (t \geq 1, k \in \mathcal{N})$.

Corollary 5 (Robust Logistic Regression)**.** *Consider the model in equation*(7), *then there are universal constants C_1, C_2, such that if $\epsilon < C_1$, then Algorithm* 1 *initialized at θ^0 with stepsize $\eta = 2/(\tau_u + \tau_\ell)$ and Algorithm* 2 *as gradient estimator, returns iterates $\{\widehat{\theta}^t\}_{t=1}^T$, such that with probability at least $1 - \delta$*

$$\|\widehat{\theta}^t - \theta^*\|_2 \leq \kappa^t \|\theta^0 - \theta^*\|_2 + \frac{C_2\sqrt{\|\Sigma\|_2 \log p}}{1 - \kappa}\left(\epsilon^{\frac{1}{2}} + \gamma(\widetilde{n}, p, \widetilde{\delta})\right), \tag{20}$$

for some contraction parameter $\kappa < 1$.

Under the restrictive assumption that $x \sim \mathcal{N}(0, \mathcal{I}_p)$, Du et al. [23] exploited Stein's trick to derive a plugin estimator for logistic regression. However, similar to the linear regression, the error of the plugin estimator scales with $\|\theta^*\|_2$, which is avoided in our robust gradient descent algorithm. We also note that our algorithm extends to general covariate distributions.

6.3 Exponential Family

Here we assume that the random vector $\phi(z), z \sim P$ has bounded 4^{th} moments.

Theorem 6 (Robust Exponential Family). *Consider the model in equation(8), then there are universal constants C_1, C_2, such that if $\epsilon < C_1$, then Algorithm 1 initialized at θ^0 with stepsize $\eta = 2/(\tau_u + \tau_\ell)$ and Algorithm 2 as gradient oracle, returns iterates $\{\widehat{\theta}^t\}_{t=1}^T$, such that with probability at least $1 - \delta$*

$$\|\widehat{\theta}^t - \theta^*\|_2 \leq \kappa^t \|\theta^0 - \theta^*\|_2 + \frac{C_2 \sqrt{\tau_u \log p}}{1 - \kappa} \left(\epsilon^{\frac{1}{2}} + \gamma(\widetilde{n}, p, \widetilde{\delta}) \right), \tag{21}$$

for some contraction parameter $\kappa < 1$.

Plugin Estimation. Since the true parameter θ^* is the minimizer of the negative log-likelihood, we know that $\mathbb{E}[\nabla \overline{\mathcal{L}}(\theta^*)] = 0$, which implies that $\nabla A(\theta^*) = \mathbb{E}_{\theta^*}[\phi(Z)]$. This shows that the true parameter θ^* can be obtained by inverting the ∇A operator, whenever possible. In the robust estimation framework, we can use a robust mean of the sufficient statistics to estimate $\mathbb{E}_{\theta^*}[\phi(Z)]$. We instantiate this estimator using the mean estimator of [14] to estimate $\mathbb{E}_{\theta^*}[\phi(Z)]$:

Corollary 7. *Consider the model in equation(8), then there are universal constants C_1, C_2 such that if $\epsilon < C_1$, then [14] returns an estimate $\widehat{\mu}$ of $\mathbb{E}[\phi(z)]$, such that with probability at least $1 - \delta$*

$$\|\mathcal{P}_\Theta \left[(\nabla A)^{-1} \widehat{\mu} \right] - \theta^*\|_2 \leq C_2 \frac{\sqrt{\tau_u \log p}}{\tau_\ell} \left(\epsilon^{\frac{1}{2}} + \gamma(n, p, \delta, \epsilon) \right), \tag{22}$$

where $\mathcal{P}_\Theta [\theta] = \operatorname{argmin}_{y \in \Theta} \|y - \theta\|_2^2$ is the projection operator onto the feasible set Θ.

6.4 Discussion and Limitations

In the asymptotic setting of $n \to \infty$, Algorithm 1 with Algorithm 2 as gradient estimator converges to a point $\widehat{\theta}$ such that $\|\widehat{\theta} - \theta^*\|_2 = O(\sqrt{\epsilon \log p})$. Hence, our error scales only logarithmically with the dimension p. This dependency on the dimension p is a facet of using the estimator from Lai et al. [14] for gradient estimation. Using better oracles will only improve our performance. Next, we would like to point to the difference in the maximum allowed contamination ϵ^* between the three models. For logistic regression and exponential family, $\epsilon^* < C_1$, while for linear regression, $\epsilon^* < \frac{C_1 \tau_\ell^2}{\tau_u^2 \log p}$. These differences are in large part due to differing variances of the gradients, which naturally depend on the underlying risk function. This scaling of the variance of gradients for linear regression also provides insights into the limitations of our robust gradient descent approach in Algorithm 1. In the Appendix, we provide an upper bound for the contamination level ϵ based on the initialization point θ^0, above which, Algorithm 1 would not work for any gradient estimator.

7 Consequences for Heavy-Tailed Estimation

In this section we present specific applications of Theorem 1 for parametric estimation, under heavy tailed setting. The proofs of the results can be found in the Appendix.

7.1 Linear Regression

We first consider the linear regression model described in Equation (4). We assume that the covariates $x \in \mathbb{R}^p$ have bounded 4^{th}-moments and the noise w has bounded 2^{nd} moments. This assumption is needed to bound the error in the gradient estimator (see Lemma 2).

Theorem 8 (Heavy Tailed Linear Regression). *Consider the statistical model in equation (4). There are universal constants $C_1, C_2 > 0$ such that if*

$$\widetilde{n} > \frac{C_1 \tau_u^2}{\tau_l^2} p \log(1/\widetilde{\delta})$$

and if Algorithm 1 is initialized at θ^0 with stepsize $\eta = 2/(\tau_u + \tau_\ell)$ and Algorithm 3 as gradient estimator, then it returns iterates $\{\widehat{\theta}^t\}_{t=1}^T$ such that with probability at least $1 - \delta$

$$\|\widehat{\theta}^t - \theta^*\|_2 \leq \kappa^t \|\theta^0 - \theta^*\|_2 + \frac{C_2 \sigma \sqrt{\|\Sigma\|_2}}{1 - \kappa} \left(\sqrt{\frac{p \log(1/\widetilde{\delta})}{\widetilde{n}}} \right), \tag{23}$$

for some contraction parameter $\kappa < 1$.

7.2 Generalized Linear Models

In this section we consider generalized linear models described in Equation (5), where the covariate x is allowed to have a heavy tailed distribution. Here we assume that the covariates have bounded 4^{th} moment. Additionally, we assume smoothness of $\Phi'(\cdot)$ around θ^*. Specifically, we assume that there exist universal constants $L_{\Phi,2k}$, B_{2k} such that

$$\mathbb{E}_x \left[\left| \Phi'(\langle x, \ \theta \rangle) - \Phi'(\langle x, \ \theta^* \rangle) \right|^{2k} \right] \leq L_{\Phi,2k} \|\theta^* - \theta\|_2^{2k} + B_{\Phi,2k}, \quad \text{for } k = 1, 2$$

We also assume that $\mathbb{E}_x[|\Phi^{(t)}(\langle x, \theta^* \rangle)|^k] \leq M_{\Phi,t,k}$ for $t \in \{1, 2, 4\}$, where $\Phi^{(t)}(\cdot)$ is the t^{th}-derivative of $\Phi(\cdot)$.

Theorem 9 (Heavy Tailed Generalized Linear Models). *Consider the statistical model in equation (5). There are universal constants $C_1, C_2 > 0$ such that if*

$$\widetilde{n} > \frac{C_1 \|\Sigma\|_2 \left(\sqrt{L_{\Phi,4}} + L_{\Phi,2} \right)}{\tau_l^2} p \log 1/\widetilde{\delta},$$

and if Algorithm 1 is initialized at θ^0 with stepsize $\eta = 2/(\tau_u + \tau_\ell)$ and Algorithm 3 as gradient estimator, it returns iterates $\{\widehat{\theta}^t\}_{t=1}^T$ such that with probability at least $1 - \delta$

$$\|\widehat{\theta}^t - \theta^*\|_2 \leq \kappa^t \|\theta^0 - \theta^*\|_2$$
$$+ \frac{C_2 \|\Sigma\|_2^{\frac{1}{2}} \left[B_{\Phi,4}^{\frac{1}{4}} + B_{\Phi,2}^{\frac{1}{2}} + c(\sigma)^{\frac{1}{2}} M_{\Phi,2,2}^{\frac{1}{4}} + c(\sigma)^{\frac{3}{4}} M_{\Phi,4,1}^{\frac{1}{4}} \right]}{1 - \kappa} \left(\sqrt{\frac{p \log(1/\widetilde{\delta})}{\widetilde{n}}} \right), \tag{24}$$

for some contraction parameter $\kappa < 1$.

We now instantiate the above Theorem for logistic regression model.

Corollary 10 (Heavy Tailed Logistic Regression). *Consider the model in equation*(7). *There are universal constants $C_1, C_2 > 0$ such that if*

$$\widetilde{n} > \frac{C_1^2 \|\Sigma\|_2}{\tau_l^2} p \log 1/\widetilde{\delta}.$$

and if Algorithm 1 initialized at θ^0 with stepsize $\eta = 2/(\tau_u + \tau_\ell)$ and Algorithm 3 as gradient estimator, it returns iterates $\{\widehat{\theta}^t\}_{t=1}^T$ such that with probability at least $1 - \delta$

$$\|\widehat{\theta}^t - \theta^*\|_2 \le \kappa^t \|\theta^0 - \theta^*\|_2 + \frac{C_2\sqrt{\|\Sigma\|_2}}{1 - \kappa}\left(\sqrt{\frac{p\log(1/\widetilde{\delta})}{\widetilde{n}}}\right), \tag{25}$$

for some contraction parameter $\kappa < 1$.

7.3 Exponential Family

We now instantiate Theorem 1 for parameter estimation in heavy-tailed exponential family distributions. Here we assume that the random vector $\phi(z), z \sim P$ has bounded 2nd moments, and we obtain the following result:

Theorem 11 (Heavy Tailed Exponential Family). *Consider the model in equation*(8). *If Algorithm 1 is initialized at θ^0 with stepsize $\eta = 2/(\tau_u + \tau_\ell)$ and Algorithm 3 as gradient estimator, it returns iterates $\{\widehat{\theta}^t\}_{t=1}^T$, such that with probability at least $1 - \delta$*

$$\|\widehat{\theta}^t - \theta^*\|_2 \le \kappa^t \|\theta^0 - \theta^*\|_2 + \frac{1}{1 - \kappa}C\sqrt{\frac{\|\nabla^2 A(\theta^*)\|_2 p \log 1/\widetilde{\delta}}{\widetilde{n}}}, \tag{26}$$

for some contraction parameter $\kappa < 1$ and universal constant C.

8 Discussion

In this paper we introduced a broad class of robust estimators, that leverage the inherent robustness of gradient descent, together with the observation that for risk minimization in most statistical models, the gradient of the risk takes the form of a simple multivariate mean, which can be robustly estimated using recent work on robust mean estimation. In contrast to classical M-estimators that use robust estimates of the risk, our class of estimators employ a shift in perspective, and use robust estimates of gradients of the risk instead, which can then be embedded into a simple projected gradient descent iterative algorithm. Our class of robust gradient descent estimators work well in practice and in many cases outperform other robust (and non-robust) estimators. We also show that these estimators have strong robustness guarantees under varied robustness settings, including Huber's ϵ-contamination model and for heavy-tailed distributions.

There are several avenues for future work, including a better understanding of robust mean estimation, any improvement in which would immediately translate to improved guarantees for our robust gradient descent estimators. Finally, it would also be of interest to understand the extent to which we could replace gradient descent with other optimization methods such as accelerated gradient descent or Newton's method. We note however, that although these methods may have faster rates of convergence in the classical risk minimization settings, in our setup their stability to using inexact gradients is far more crucial and warrants further investigation.

9   Acknowledgements

References

[1] G. E. P. Box. Non-normality and tests on variances. *Biometrika*, 40(3-4):318–335, 1953.

[2] John W Tukey. Mathematics and the picturing of data. In *Proceedings of the international congress of mathematicians*, volume 2, pages 523–531, 1975.

[3] P. J. Huber. *Robust Statistics*. John Wiley & Sons, 1981.

[4] Frank R Hampel, Elvezio M Ronchetti, Peter J Rousseeuw, and Werner A Stahel. *Robust statistics: the approach based on influence functions*, volume 114. John Wiley & Sons, 2011.

[5] L. Devroye and L. Györfi. *Nonparametric density estimation: the L1 view*. Wiley series in probability and mathematical statistics. Wiley, 1985.

[6] Yannis G. Yatracos. Rates of convergence of minimum distance estimators and kolmogorov's entropy. *Ann. Statist.*, 13(2):768–774, 06 1985.

[7] David L Donoho and Richard C Liu. The" automatic" robustness of minimum distance functionals. *The Annals of Statistics*, pages 552–586, 1988.

[8] Ivan Mizera. On depth and deep points: a calculus. *Annals of Statistics*, pages 1681–1736, 2002.

[9] Chao Gao. Robust regression via mutivariate regression depth. *arXiv preprint arXiv:1702.04656*, 2017.

[10] Mengjie Chen, Chao Gao, and Zhao Ren. Robust covariance matrix estimation via matrix depth. *arXiv preprint arXiv:1506.00691*, 2015.

[11] Jianqing Fan, Weichen Wang, and Ziwei Zhu. A shrinkage principle for heavy-tailed data: High-dimensional robust low-rank matrix recovery, 2016.

[12] Wen-Xin Zhou, Koushiki Bose, Jianqing Fan, and Han Liu. A new perspective on robust m-estimation: Finite sample theory and applications to dependence-adjusted multiple testing. *arXiv preprint arXiv:1711.05381*, 2017.

[13] S. van de Geer. *Empirical Processes in M-Estimation*. Cambridge University Press, 2000.

[14] Kevin A Lai, Anup B Rao, and Santosh Vempala. Agnostic estimation of mean and covariance. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 665–674. IEEE, 2016.

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

[16] Kush Bhatia, Prateek Jain, and Purushottam Kar. Robust regression via hard thresholding. In *Advances in Neural Information Processing Systems*, pages 721–729, 2015.

[17] Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6): 381–395, 1981.

[18] Sivaraman Balakrishnan, Martin J Wainwright, and Bin Yu. Statistical guarantees for the em algorithm: From population to sample-based analysis. *The Annals of Statistics*, 45(1):77–120, 2017.

[19] Peter J Huber. A robust version of the probability ratio test. *The Annals of Mathematical Statistics*, 36(6):1753–1758, 1965.

[20] Mengjie Chen, Chao Gao, Zhao Ren, et al. A general decision theory for huber's epsilon-contamination model. *Electronic Journal of Statistics*, 10(2):3752–3774, 2016.

[21] Cecil Hastings Jr, Frederick Mosteller, John W Tukey, and Charles P Winsor. Low moments for small samples: a comparative study of order statistics. *The Annals of Mathematical Statistics*, pages 413–426, 1947.

[22] Ilias Diakonikolas, Gautam Kamath, Daniel M Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robust estimators in high dimensions without the computational intractability. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 655–664. IEEE, 2016.

[23] Simon S Du, Sivaraman Balakrishnan, and Aarti Singh. Computationally efficient robust estimation of sparse functionals. *Conference on Learning Theory*, 2017.

[24] Jerry Li. Robust sparse estimation tasks in high dimensions. *Conference on Learning Theory*, 2017.

[25] Moses Charikar, Jacob Steinhardt, and Gregory Valiant. Learning from untrusted data. In *STOC*, 2017.

[26] Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis. *Journal of the ACM (JACM)*, 58(3):11, 2011.

[27] Po-Ling Loh and Martin J Wainwright. High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity. In *Advances in Neural Information Processing Systems*, pages 2726–2734, 2011.

[28] Xinyang Yi, Dohyung Park, Yudong Chen, and Constantine Caramanis. Fast algorithms for robust PCA via gradient descent. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 4152–4160, 2016.

[29] Po-Ling Loh. Statistical consistency and asymptotic normality for high-dimensional robust m-estimators. *Ann. Statist.*, 45(2):866–896, 04 2017.

[30] Yudong Chen, Constantine Caramanis, and Shie Mannor. Robust sparse regression under adversarial corruption. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, pages 774–782, 2013.

[31] Matthieu Lerasle and Roberto I Oliveira. Robust empirical mean estimators. *arXiv preprint arXiv:1112.3914*, 2011.

[32] Olivier Catoni et al. Challenging the empirical mean and empirical variance: a deviation study. In *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, volume 48, pages 1148–1185. Institut Henri Poincaré, 2012.

[33] Daniel Hsu and Sivan Sabato. Loss minimization and parameter estimation with heavy tails. *Journal of Machine Learning Research*, 17(18):1–40, 2016.

[34] Christian Brownlees, Emilien Joly, and Gábor Lugosi. Empirical risk minimization for heavy-tailed losses. *The Annals of Statistics*, 43(6):2507–2536, 2015.

[35] Gabor Lugosi and Shahar Mendelson. Risk minimization by median-of-means tournaments. *arXiv preprint arXiv:1608.00757*, 2016.

[36] Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.

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

[38] Gábor Lugosi and Shahar Mendelson. Sub-gaussian estimators of the mean of a random vector. *The Annals of Statistics*, 2017.

[39] Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*, STOC '96, pages 20–29, New York, NY, USA, 1996. ACM.

[40] A.S. Nemirovski and D.B. Yudin. *Problem Complexity and Method Efficiency in Optimization*. A Wiley-Interscience publication. Wiley, 1983.

[41] Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169 – 188, 1986.

[42] Luc Devroye and Gábor Lugosi. *Combinatorial methods in density estimation*. Springer Science & Business Media, 2012.

[43] Kuang-Chih Lee, Jeffrey Ho, and David J Kriegman. Acquiring linear subspaces for face recognition under variable lighting. *IEEE Transactions on pattern analysis and machine intelligence*, 27 (5):684–698, 2005.

[44] Yin Wang, Caglayan Dicle, Mario Sznaier, and Octavia Camps. Self scaled regularized robust regression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3261–3269, 2015.

[45] John Duchi and Hongseok Namkoong. Variance-based regularization with convex objectives. *arXiv preprint arXiv:1610.02581*, 2016.

[46] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.

[47] S Kakade, Shai Shalev-Shwartz, and Ambuj Tewari. Applications of strong convexity–strong smoothness duality to learning with matrices. *CoRR, abs/0910.0610*, 2009.

[48] Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.

A   Proof of Theorem 1

In this section, we present the proof of our main result on projected gradient descent with an inexact gradient estimator. To ease the notation we will often omit $\{D_n, \delta\}$ from $g(\theta; D_n, \delta)$.

Proof. At any iteration step $t \in \{1, 2, \dots, T\}$, by assumption we have that with probability at least $1 - \frac{\delta}{T}$,

$$\|g(\theta^t; D_n, \delta/T) - \nabla \mathcal{R}(\theta^t)\|_2 \le \alpha(n/T, \delta/T)\|\theta - \theta^*\|_2 + \beta(n/T, \delta/T). \tag{27}$$

Taking union bound, (27) holds over all iteration steps $t \in \{1 \dots T\}$, with probability at least $1 - \delta$. For the remainder of the analysis, we assume this event to be true.

Notation. Let $g(\theta^k) = \nabla \mathcal{R}(\theta^k) + e_k$ be the noisy gradient. Let $\alpha = \alpha(n/T, \delta/T)$ and $\beta = \beta(n/T, \delta/T)$ for brevity.

We have the following Lemma from Bubeck [36].

Lemma 3. *[Lemma 3.11 [36]] Let f be M-smooth and m-strongly convex, then for all $x, y \in \mathbb{R}^p$, we have:*

$$\langle \nabla f(x) - \nabla f(y), x - y \rangle \geq \frac{mM}{m+M} \|x - y\|_2^2 + \frac{1}{m+M} \|\nabla f(y) - \nabla f(x)\|_2^2.$$

By assumptions we have that: $\|\nabla \mathcal{R}(\theta^k) - g(\theta^k)\|_2 = \|e_k\|_2 \leq \alpha \|\theta^k - \theta^*\|_2 + \beta$. Our update rule is $\theta^{k+1} = \mathbb{P}_\Theta \left[\theta^k - \eta g(\theta^k) \right]$. Then we have that:

$$
\begin{aligned}
\|\theta^{k+1} - \theta^*\|_2^2 = \|\mathbb{P}_\Theta[\theta^k - \eta g(\theta^k)] - \theta^*\|_2^2 &= \|\mathbb{P}_\Theta[\theta^k - \eta g(\theta^k)] - \mathbb{P}_\Theta[\theta^* - \eta \nabla R(\theta^*)]\|_2^2 \\
&\leq \|\theta^k - \eta g(\theta^k) - (\theta^* - \eta \nabla R(\theta^*))\|_2^2 \qquad (28) \\
&= \|\theta^k - \theta^* - \eta(\nabla \mathcal{R}(\theta^k) - \nabla R(\theta^*)) - \eta e_k\|_2^2 \\
&\leq \|\theta^k - \theta^* - \eta(\nabla \mathcal{R}(\theta^k) - \nabla R(\theta^*))\|_2^2 + \eta^2 \|e_k\|_2^2 \\
&\quad + 2\|e_k\|_2 \|\theta^k - \theta^* - \eta(\nabla \mathcal{R}(\theta^k) - \nabla R(\theta^*))\|_2, \qquad (29)
\end{aligned}
$$

where Equation (28) follows from contraction property of projections. Now, we can write $\|\theta^k - \theta^* - \eta(\nabla \mathcal{R}(\theta^k) - \nabla R(\theta^*))\|_2$ as

$$
\begin{aligned}
\|\theta^k - \theta^* - \eta(\nabla \mathcal{R}(\theta^k) - \nabla R(\theta^*))\|_2^2 &= \|\theta^k - \theta^*\|_2^2 + \eta^2 \|\nabla \mathcal{R}(\theta^k) - \nabla R(\theta^*)\|_2^2 - 2\eta \langle \nabla \mathcal{R}(\theta^k) - \nabla R(\theta^*), \theta^k - \theta^* \rangle \\
&\leq \|\theta^k - \theta^*\|_2^2 + \eta^2 \|\nabla \mathcal{R}(\theta^k) - \nabla R(\theta^*)\|_2^2 - 2\eta \left(\frac{\tau_\ell \tau_u}{\tau_\ell + \tau_u} \|\theta^k - \theta^*\|_2^2 + \frac{1}{\tau_\ell + \tau_u} \|\nabla \mathcal{R}(\theta^k) - \nabla R(\theta^*)\|_2^2 \right)
\end{aligned}
$$
$$(30)$$

$$= \|\theta^k - \theta^*\|_2^2 (1 - 2\eta \tau_\ell \tau_u/(\tau_\ell + \tau_u)) + \eta \|\nabla \mathcal{R}(\theta^k) - \nabla R(\theta^*)\|_2^2 (\eta - 2/(\tau_u + \tau_\ell)) \qquad (31)$$

$$= \|\theta^k - \theta^*\|_2^2 (1 - 2\eta \tau_\ell \tau_u/(\tau_\ell + \tau_u)), \qquad (32)$$

where the second step follows from Lemma 3 and the last step follows from the step size $\eta = 2/(\tau_\ell + \tau_u)$.

Now, combining Equations (29) and (32), and using our assumption that $\|e_k\|_2 \leq \alpha \|\theta^k - \theta^*\|_2 + \beta$, we get:

$$\|\theta^{k+1} - \theta^*\|_2^2 \leq \left(\|\theta^k - \theta^*\|_2 \sqrt{(1 - 2\eta \tau_\ell \tau_u/(\tau_\ell + \tau_u))} + \eta \|e_k\|_2 \right)^2$$

$$\|\theta^{k+1} - \theta^*\|_2 \leq \left[\sqrt{1 - 2\eta \tau_\ell \tau_u/(\tau_\ell + \tau_u)} + \eta \alpha \right] \|\theta^k - \theta^*\|_2 + \eta \beta.$$

Let $\kappa = \sqrt{1 - 2\eta \tau_\ell \tau_u/(\tau_\ell + \tau_u)} + \eta \alpha$. By the assumption on stability we have $\alpha < \tau_\ell$.

$$\kappa = \sqrt{1 - 2\eta \tau_\ell \tau_u/(\tau_\ell + \tau_u)} + \eta \alpha \qquad (33)$$

$$< \sqrt{1 - 2\eta \tau_\ell \tau_u/(\tau_\ell + \tau_u)} + \eta \tau_\ell. \qquad (34)$$

Since $\eta = 2/(\tau_\ell + \tau_u)$, we get that

$$\kappa < \sqrt{1 - 4\tau_u^2 \tau_\ell^2/(\tau_\ell + \tau_u)^2} + 2\tau_\ell/(\tau_u + \tau_\ell) \qquad (35)$$

$$\kappa < \frac{\tau_u - \tau_\ell}{\tau_u + \tau_\ell} + 2\tau_\ell/(\tau_u + \tau_\ell) \qquad (36)$$

$$\kappa < 1 \qquad (37)$$

Therefore, we have that,

$$\|\theta^{k+1} - \theta^*\|_2 \le \kappa \|\theta^k - \theta^*\|_2 + \eta\beta.$$

for some $\kappa < 1$. Solving the induction,we get:

$$\|\theta^k - \theta^*\|_2 \le \kappa^k \|\theta^0 - \theta^*\|_2 + \frac{1}{1-\kappa}\eta\beta.$$

$\square$

B Proof of Theorem 2

The proof of Theorem 2 follows from Theorem 4, where we study Generalized Linear Models, which include linear regression as a special case. For the case of linear regression with gaussian noise, it is relatively straightforward to see that the smoothness parameters satisfy $L_{\Phi,2k} = C_{2k}\|\Sigma\|_2^k$, $B_{\Phi,2k} = 0$, $M_{\Phi,t,k} = 1$ $\forall(t \ge 2, k \in \mathcal{N})$ and $M_{\Phi,t,k} = 0$ $\forall(t \ge 3, k \in \mathcal{N})$ under the assumption of bounded 8^{th} moments of the covariates. Substituting these values in Theorem 4 gives us the required result.

C Proof of Theorem 4

To prove our result on Robust Generalized Linear Models, we first study the distribution of gradients of the corresponding risk function.

Lemma 4. *Consider the model in Equation* (5), *then there exist universal constants* $C_1, C_2 > 0$ *such that*

$$\|\mathrm{Cov}(\nabla\overline{\mathcal{L}}(\theta)\|_2 \le C_1\|\Delta\|_2^2\|\Sigma\|_2\left(\sqrt{C_4}\sqrt{L_{\Phi,4}} + L_{\Phi,2}\right)$$
$$+ C_2\|\Sigma\|_2\left(B_{\Phi,2} + \sqrt{B_{\Phi,4}} + c(\sigma)\sqrt{3M_{\Phi,2,2}} + \sqrt{c(\sigma)^3 M_{\Phi,4,1}}\right)$$

Bounded fourth moments $\mathbb{E}\left[\left[(\nabla\overline{\mathcal{L}}(\theta) - \mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)])^T v\right]^4\right] \le C_2(\mathrm{Var}[\nabla\overline{\mathcal{L}}(\theta)^T v])^2.$

Proof. The gradient $\nabla\overline{\mathcal{L}}(\theta)$ and it's expectation can be written as:

$$\nabla\overline{\mathcal{L}}(\theta) = -y.x + u(\langle x, \theta\rangle).x$$
$$\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)] = \mathbb{E}[x\left(u(x^T\theta) - u(x^T\theta^*)\right)],$$

where $u(t) = \Phi'(t)$.

$$\begin{aligned}
\|\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)]\|_2 &= \sup_{y\in\mathbb{S}^{p-1}} y^T\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)]\\
&\le \sup_{y\in\mathbb{S}^{p-1}} \mathbb{E}[(y^Tx)\left(u(x^T\theta) - u(x^T\theta^*)\right)]\\
&\le \sup_{y\in\mathbb{S}^{p-1}} \sqrt{\mathbb{E}[(y^Tx)^2]}\sqrt{\mathbb{E}[(u(x^T\theta) - u(x^T\theta^*))^2]}\\
&\le C_1\|\Sigma\|_2^{\frac{1}{2}}\sqrt{L_{\Phi,2}\|\Delta\|_2^2 + B_{\Phi,2}}
\end{aligned}$$

where the last line follows from our assumption of smoothness.

Now, to bound the maximum eigenvalue of the $\text{Cov}(\nabla\overline{\mathcal{L}}(\theta))$,

$$
\begin{aligned}
\|\text{Cov}(\nabla\overline{\mathcal{L}}(\theta))\|_2 &= \sup_{z\in\mathbb{S}^{p-1}} z^T\left(\mathbb{E}\left[\nabla\overline{\mathcal{L}}(\theta)\nabla\overline{\mathcal{L}}(\theta)^T\right] - \mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)]\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)]^T\right)z \\
&\leq \sup_{z\in\mathbb{S}^{p-1}} z^T\left(\mathbb{E}\left[\nabla\overline{\mathcal{L}}(\theta)\nabla\overline{\mathcal{L}}(\theta)^T\right]\right)z + \sup_{z\in\mathbb{S}^{p-1}} z^T\left(\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)]\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)]^T\right)z \\
&\leq \sup_{z\in\mathbb{S}^{p-1}} z^T\left(\mathbb{E}\left[xx^T\left(u(x^T\theta)-y\right)^2\right]\right)z + \|\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)]\|_2^2 \\
&\leq \sup_{z\in\mathbb{S}^{p-1}} \mathbb{E}\left[z^T\left(xx^T\left(u(x^T\theta)-y\right)^2\right)z\right] + \|\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)]\|_2^2 \\
&\leq \sup_{z\in\mathbb{S}^{p-1}} \sqrt{\mathbb{E}\left[(z^Tx)^4\right]}\sqrt{\mathbb{E}\left[\left(u(x^T\theta)-y\right)^4\right]} + \|\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)]\|_2^2
\end{aligned}
$$

To bound $\mathbb{E}\left[\left(u(x^T\theta)-y\right)^4\right]$, we make use of the C_r inequality.

C_r **inequality**. If X and Y are random variables such that $\mathbb{E}|X|^r < \infty$ and $\mathbb{E}|Y|^4 < \infty$ where $r\geq 1$ then:

$$
\mathbb{E}|X+Y|^r \leq 2^{r-1}\left(\mathbb{E}|X|^r + \mathbb{E}|Y|^r\right)
$$

Using the C_r inequality, we have that

$$
\begin{aligned}
\mathbb{E}\left[\left(u(x^T\theta)-y\right)^4\right] &\leq 8\left(\mathbb{E}\left[\left(u(x^T\theta)-u(x^T\theta^*)\right)^4\right] + \mathbb{E}\left[\left(u(x^T\theta^*)-y\right)^4\right]\right) \\
&\leq C\left(L_{\Phi,4}\|\Delta\|_2^4 + B_{\Phi,4} + c(\sigma)^3 M_{\Phi,4,1} + 3c(\sigma)^2 M_{\Phi,2,2}\right)
\end{aligned}
$$

where the last line follows from our assumption that $P_{\theta^*}(y|x)$ is in the exponential family, hence, the cumulants are higher order derivatives of the log-normalization function.

$$
\begin{aligned}
\|\text{Cov}(\nabla\overline{\mathcal{L}}(\theta))\|_2 &\leq \sqrt{C}\sqrt{C_4}\|\Sigma\|_2\left(\sqrt{L_{\Phi,4}}\|\Delta\|_2^2 + \sqrt{B_{\Phi,4}} + c(\sigma)\sqrt{3M_{\Phi,2,2}} + \sqrt{c(\sigma)^3 M_{\Phi,4,1}}\right) + \|\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)]\|_2^2 \\
&\leq \sqrt{C}\sqrt{C_4}\|\Sigma\|_2\left(\sqrt{L_{\Phi,4}}\|\Delta\|_2^2 + \sqrt{B_{\Phi,4}} + c(\sigma)\sqrt{3M_{\Phi,2,2}} + \sqrt{c(\sigma)^3 M_{\Phi,4,1}}\right) + C_1^2\|\Sigma\|_2\left(L_{\Phi,2}\|\Delta\|_2^2 + B_{\Phi,2}\right) \\
&\leq C\|\Delta\|_2^2\|\Sigma\|_2\left(\sqrt{C_4}\sqrt{L_{\Phi,4}} + L_{\Phi,2}\right) + C_6\|\Sigma\|_2\left(B_{\Phi,2} + \sqrt{B_{\Phi,4}} + c(\sigma)\sqrt{3M_{\Phi,2,2}} + \sqrt{c(\sigma)^3 M_{\Phi,4,1}}\right)
\end{aligned}
$$

Bounded Fourth Moment. To show that the fourth moment of the gradient distribution is bounded, we have

$$
\begin{aligned}
\mathbb{E}\left[\left[(\nabla\overline{\mathcal{L}}(\theta)-\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)])^T v\right]^4\right] &\leq \mathbb{E}\left[\left[|(\nabla\overline{\mathcal{L}}(\theta)-\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)])^T v|\right]^4\right] \\
&\leq 8\left[\underbrace{\mathbb{E}[|\nabla\overline{\mathcal{L}}(\theta)]^T v|^4]}_{A} + \underbrace{\mathbb{E}[|\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)]^T v|^4]}_{B}\right].
\end{aligned}
$$

Control of A.

$$\mathbb{E}[|\nabla\overline{\mathcal{L}}(\theta)]^T v|^4] = \mathbb{E}[(x^T v)^4 (u(x^T \theta) - y)^4]$$

$$\leq \sqrt{\mathbb{E}[(x^T v)^8]}\sqrt{\mathbb{E}[(u(x^T \theta) - y)^8]}$$

$$\leq \sqrt{C_8}\|\Sigma\|_2^2 \sqrt{\mathbb{E}[(u(x^T \theta) - u(x^T \theta^*))^8] + \mathbb{E}[(u(x^T \theta^*) - y)^8]}$$

$$\leq \sqrt{C_8}\|\Sigma\|_2^2 \sqrt{L_{\Phi,8}\|\Delta\|_2^8 + B_{\Phi,8} + \sum_{t,k=2}^{8} g_{t,k} M_{\Phi,t,k}}$$

$$\leq \sqrt{C}\|\Sigma\|_2^2 \sqrt{L_{\Phi,8}}\|\Delta\|_2^4 + \sqrt{B_{\Phi,8}} + \sqrt{\sum_{t,k=2}^{8} g_{t,k} M_{\Phi,t,k}}$$

where the last step follows from the fact that the 8th central moment can be written as a polynomial involving the lower cumulants, which in turn are the derivatives of the log-normalization function.

Control of B.

$$\mathbb{E}[|\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)]^T v|^4] \leq \|\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)]\|_2^4 \leq C_1 \|\Sigma\|_2^2 \left(L_{\Phi,2}^2\|\Delta\|_2^2 + B_{\Phi,2}^2\right)$$

By assumption $L_{\Phi,k}, B_{\Phi,k}, M_{\Phi,t,k}$ are all bounded for $k, t \leq 8$, which implies that there exist constants $c_1, c_2 > 0$ such that

$$\mathbb{E}\left[\left[(\nabla\overline{\mathcal{L}}(\theta) - \mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)]^T v\right]^4\right] \leq c_1 \|\Sigma\|_2^2 \|\Delta\|_2^4 + c_2 \tag{38}$$

Previously, we say that $\|\text{Cov}\nabla\overline{\mathcal{L}}(\theta)\|_2 \leq c_3\|\Sigma\|_2\|\Delta\|_2^2 + c_4$, for some universal constants $c_3, c_4 > 0$, hence the gradient $\nabla\overline{\mathcal{L}}(\theta)$ has bounded fourth moments. □

Having studied the distribution of the gradients, we use Lemma 1 to characterize the stability of Huber Gradient estimator. Using Lemma 1, we know that at any point θ, the Huber Gradient Estimator $g(\theta, \delta/T)$ satisfies that with probability $1 - \delta/T$,

$$\|g(\theta, \delta/T) - \nabla\mathcal{R}(\theta)\|_2 \leq C_2\left(\epsilon^{\frac{1}{2}} + \gamma(\widetilde{n}, p, \widetilde{\delta})\right)\|\text{Cov}(\nabla\overline{\mathcal{L}}(\theta))\|_2^{\frac{1}{2}}\sqrt{\log p}.$$

Substituting the upper bound on $\|\text{Cov}(\nabla\overline{\mathcal{L}}(\theta))\|_2$ from Lemma 4, we get that there are universal constants C_1, C_2 such that with probability at least $1 - \delta/T$

$$\|g(\theta) - \nabla\mathcal{R}(\theta)\|_2 \leq \underbrace{C_1\left(\epsilon^{\frac{1}{2}} + \gamma(\widetilde{n}, p, \widetilde{\delta})\right)\sqrt{\log p}\|\Sigma\|_2^{\frac{1}{2}}[L_{\Phi,4}^{\frac{1}{4}} + L_{\Phi,2}^{\frac{1}{2}}]\|\Delta\|_2}_{\alpha(\widetilde{n},\widetilde{\delta})} \tag{39}$$

$$+ \underbrace{C_2\left(\epsilon^{\frac{1}{2}} + \gamma(\widetilde{n}, p, \widetilde{\delta})\right)\sqrt{\log p}\|\Sigma\|_2^{\frac{1}{2}}[B_{\Phi,4}^{\frac{1}{4}} + B_{\Phi,2}^{\frac{1}{2}} + c(\sigma)^{\frac{1}{2}}M_{\Phi,2,2}^{\frac{1}{4}} + c(\sigma)^{\frac{3}{4}}M_{\Phi,4,1}^{\frac{1}{4}}]}_{\beta(\widetilde{n},\widetilde{\delta})}$$

$$\tag{40}$$

To ensure stability of gradient descent, we need that $\alpha(\widetilde{n}, \widetilde{\delta}) < \tau_\ell$. Using Equation (39), we get that gradient descent is stable as long as the number of samples n is large enough such that $\gamma(\widetilde{n}, p, \widetilde{\delta}) < \frac{C_1 \tau_\ell}{\sqrt{\log p}\|\Sigma\|_2^{\frac{1}{2}}[L_{\Phi,4}^{\frac{1}{4}} + L_{\Phi,2}^{\frac{1}{2}}]}$, and the contamination level is such that

$$\epsilon < \left(\frac{C_2 \tau_\ell}{\sqrt{\log p} \|\Sigma\|_2^{\frac{1}{2}} [L_{\Phi,4}^{\frac{1}{4}} + L_{\Phi,2}^{\frac{1}{2}}]} - \gamma(\widetilde{n}, p, \widetilde{\delta}) \right)^2 \quad \text{for some constants } C_1 \text{ and } C_2. \text{ Plugging the corre-}$$

sponding ϵ and $\beta(\widetilde{n}, \widetilde{\delta})$ into Theorem 1, we get back the result of Theorem 4.

D   Proof of Corollary 3

We begin by studying the distribution of the random variable $xy = xx^T\theta^* + x.w$.

Lemma 5. *Consider the model in Equation* (4)*, with* $x \sim \mathcal{N}(0, \mathcal{I}_p)$ *and* $w \sim \mathcal{N}(0, 1)$ *then there exist universal constants* C_1, C_2 *such that*

$$\mathbb{E}[xy] = \theta^*$$
$$\|\mathrm{Cov}(xy)\|_2 = 1 + 2\|\theta^*\|_2^2$$
$$\text{Bounded fourth moments} \quad \mathbb{E}\left[\left[(xy - \mathbb{E}[xy])^T v\right]^4 \right] \le C_2 (\mathrm{Var}[(xy)^T v])^2.$$

Proof. **Mean.**

$$xy = xx^T\theta^* + x.w \tag{41}$$
$$\mathbb{E}[xy] = \mathbb{E}[xx^T\theta^* + x.w] \tag{42}$$
$$\mathbb{E}[xy] = \theta^*. \tag{43}$$

Covariance.

$$\mathrm{Cov}(xy) = \mathbb{E}[(xx^T - I)\theta^* + x.w)((xx^T - I)\theta^* + x.w)^T)] \tag{44}$$
$$\mathrm{Cov}(xy) = \mathbb{E}[(xx^T - I)\theta^*\theta^{*T}(xx^T - I)] + I_p. \tag{45}$$

Now, $Z = (xx^T - I)\theta^*$ can be written as:

$$(xx^T - I)\theta^* = \begin{bmatrix} (x_1^2 - 1) & x_1 x_2 & \dots & x_1 x_p \\ x_1 x_2 & (x_2^2 - 1) & \dots & x_2 x_p \\ \vdots & \vdots & \vdots & \vdots \\ x_1 x_p & x_2 x_p & \dots & (x_p^2 - 1) \end{bmatrix} \begin{bmatrix} \theta_1^* \\ \theta_2^* \\ \vdots \\ \theta_p^* \end{bmatrix} = \begin{bmatrix} \theta_1^*(x_1^2 - 1) + x_1 x_2 \theta_2^* + \dots + x_1 x_p \theta_p^* \\ x_1 x_2 \theta_1^* + (x_2^2 - 1)\theta_2^* + \dots + x_2 x_p \theta_p^* \\ \vdots \\ x_1 x_p \theta_1^* + x_2 x_p \theta_2^* + \dots + (x_p^2 - 1)\theta_p^* \end{bmatrix}.$$

Then,

$$\mathbb{E}\left[ZZ^T\right] = \begin{bmatrix} 2\theta_1^{*2} + \theta_2^{*2} + \dots + \theta_p^{*2} & \theta_1^*\theta_2^* & \dots & \theta_1^*\theta_p^* \\ \theta_1^*\theta_2^* & \theta_1^{*2} + 2\theta_2^{*2} + \dots + \theta_p^{*2} & \dots & \theta_2^*\theta_p^* \\ \vdots & \vdots & \ddots & \vdots \\ \theta_p^*\theta_1^* & \theta_2^*\theta_p^* & \dots & \theta_1^{*2} + \theta_2^{*2} + \dots + 2\theta_p^{*2} \end{bmatrix}.$$

Hence the covariance matrix can be written as:

$$\mathrm{Cov}(xy) = I_p(1 + \|\theta^*\|_2^2) + \theta^*\theta^{*T}.$$

Therefore $\|\mathrm{Cov}(xy)\|_2 = 1 + 2\|\theta^*\|_2^2$.

Bounded Fourth Moment. We start from the LHS

$$\mathbb{E}\left[\left[(xy - \mathbb{E}[xy])^T v\right]^4\right] \leq \mathbb{E}\left[\left[|(xy - \mathbb{E}[xy])^T v|\right]^4\right] \tag{46}$$

$$= \mathbb{E}\left[\left|((xx^T - I)\theta^* + wx)^T v\right|\right]^4 \tag{47}$$

$$= \mathbb{E}\left[\left|(\theta^{*T}x)(x^T v) - \theta^{*T}v + wv^T x\right|\right]^4 \tag{48}$$

$$\leq 8\left[8\left[\underbrace{\mathbb{E}\left|(\theta^{*T}x)(x^T v)\right|^4}_{A} + \underbrace{\mathbb{E}\left|\theta^{*T}v\right|^4}_{B}\right] + \underbrace{\mathbb{E}\left|w(x^T v)\right|^4}_{C}\right] \tag{49}$$

The last line follows from two applications of the following inequality:

C_r **inequality.** If X and Y are random variables such that $\mathbb{E}|X|^r < \infty$ and $\mathbb{E}|Y|^4 < \infty$ where $r \geq 1$ then:

$$\mathbb{E}|X + Y|^r \leq 2^{r-1}\left(\mathbb{E}|X|^r + \mathbb{E}|Y|^r\right).$$

Now to control each term on:

- **Control of** A. Using Cauchy Schwartz, and normality of 1D projections of normal distribution

$$A \leq \sqrt{\mathbb{E}[|\theta^{*T}x|^8]}\sqrt{\mathbb{E}[|x^T v|^8]} \tag{50}$$

$$\precsim \|\theta^*\|_2^4. \tag{51}$$

- **Control of** B, $B \leq \|\theta^*\|_2^4$.

- **Control of** C, $C = O(1)$, using independence of w and normality of 1D projections of normal distribution.

Therefore the $\mathbb{E}\left[\left[(xy - \mathbb{E}[xy])^T v\right]^4\right] \precsim c + \|\theta^*\|_2^4$.

For the RHS:

$$\text{Var}((xy)^T v)^2 = (v^T \text{Cov}(xy)v)^2 \leq \|\text{Cov}(xy)\|_2^2.$$

We saw that the $\|\text{Cov}(xy)\|_2 \precsim c + \|\theta^*\|_2^2$, so both the LHS and RHS scale with $\|\theta^*\|_2^4$. Hence, xy has bounded fourth moments. $\square$

Now that we've established that xy has bounded fourth moments implies that we can use [14] as a mean estimation oracle. Using Theorem 1.3 [14], we know that the oracle of [14] outputs an estimate $\widehat{\theta}$ of $\mathbb{E}[xy]$ such that with probability at least $1 - 1/p^{C_1}$, we have:

$$\|\widehat{\theta} - \theta^*\|_2 \leq C_2\sqrt{\|\text{Cov}(xy)\|_2 \log p}\left(\epsilon^{\frac{1}{2}} + \gamma(n, p, \delta, \epsilon)\right)$$

Using Lemma 5 to subsitute $\|\text{Cov}(xy)\|_2 \leq 1 + 2\|\theta^*\|_2^2$, we recover the statement of Corollary 3.

E   Proof of Theorem 6

To prove our result on Robust Exponential Family, we first study the distribution of gradients of the corresponding risk function.

Lemma 6. *Consider the model in Equation* (8), *then there exists a universal constant C_1 such that*

$$\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)] = \nabla A(\theta) - \nabla A(\theta^*) \tag{52}$$

$$\|\mathrm{Cov}[\nabla\overline{\mathcal{L}}(\theta)]\|_2 = \|\nabla^2 A(\theta^*)\|_2 \tag{53}$$

Bounded fourth moments $\mathbb{E}\left[\left[(\nabla\overline{\mathcal{L}}(\theta) - \mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)])^T v\right]^4\right] \leq C_1(\mathrm{Var}[\nabla\overline{\mathcal{L}}(\theta)^T v])^2.$ (54)

Proof. By Fisher Consistency of the negative log-likelihood, we know that

$$\mathbb{E}_{\theta^*}[\nabla\overline{\mathcal{L}}(\theta^*)] = 0 \tag{55}$$

$$\implies \nabla A(\theta^*) - \mathbb{E}_{\theta^*}[\phi(z)] = 0 \tag{56}$$

$$\implies \nabla A(\theta^*) = \mathbb{E}_{\theta^*}[\phi(z)]. \tag{57}$$

For the mean,

$$\nabla\overline{\mathcal{L}}(\theta) = \nabla A(\theta) - \phi(z) \tag{58}$$

$$\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)] = \nabla A(\theta) - \mathbb{E}_{\theta^*}[\phi(z)] \tag{59}$$

$$\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)] = \nabla A(\theta) - \nabla A(\theta^*). \tag{60}$$

Now, for the covariance:

$$\begin{aligned}
\mathrm{Cov}_{\theta^*}[\nabla\overline{\mathcal{L}}(\theta)] &= \mathbb{E}_{\theta^*}\left[\left(\nabla\overline{\mathcal{L}}(\theta) - \mathbb{E}_{\theta^*}[\nabla\overline{\mathcal{L}}(\theta)]\right)\left(\nabla\overline{\mathcal{L}}(\theta) - \mathbb{E}_{\theta^*}[\nabla(\overline{\mathcal{L}}(\theta)]\right)^T\right] \\
&= \mathbb{E}_{\theta^*}\left[\left(\nabla A(\theta^*) - \phi(z)\right)\left(\nabla A(\theta^*) - \phi(z)\right)^T\right] \\
&= \mathrm{Cov}_{\theta^*}\left[\nabla\overline{\mathcal{L}}(\theta^*)\right] = \nabla^2 A(\theta^*).
\end{aligned}$$

Bounded moments follows from our assumption that the sufficient statistics have bounded 4th moments. □

Having studied the distribution of the gradients, we use Lemma 1 to characterize the stability of Huber Gradient estimator. Using Lemma 1, we know that at any point θ, the Huber Gradient Estimator $g(\theta, \delta/T)$ satisfies that with probability $1 - \delta/T$,

$$\|g(\theta, \delta/T) - \nabla\mathcal{R}(\theta)\|_2 \leq C_2\left(\epsilon^{\frac{1}{2}} + \gamma(\widetilde{n}, p, \widetilde{\delta})\right)\|\mathrm{Cov}(\nabla\overline{\mathcal{L}}(\theta))\|_2^{\frac{1}{2}}\sqrt{\log p}.$$

Substituting the upper bound on $\|\mathrm{Cov}(\nabla\overline{\mathcal{L}}(\theta))\|_2$ from Lemma 6, we get that there are universal constants C_1, C_2 such that

$$\|g(\theta) - \nabla\mathcal{R}(\theta)\|_2 \leq \underbrace{C_1\left(\epsilon^{\frac{1}{2}} + \gamma(\widetilde{n}, p, \widetilde{\delta})\right)\sqrt{\log p}\sqrt{\tau_u}}_{\beta(\widetilde{n}, \widetilde{\delta})}.$$

In this case we have that $\alpha(\widetilde{n}, \widetilde{\delta}) = 0 < \tau_\ell$ by assumption. Therefore we just have that $\epsilon < C_1$ for some universal constant C_1. Plugging the corresponding ϵ and $\beta(\widetilde{n}, \widetilde{\delta})$ into Theorem 1, we get back the result of Corollary 6.

F   Proof of Corollary 7

Using the contraction property of projections, we know that

$$\|\mathcal{P}_\Theta\left[(\nabla A)^{-1}\widehat{\mu}\right] - \theta^*\|_2 = \|\mathcal{P}_\Theta\left[(\nabla A)^{-1}\widehat{\mu}\right] - \mathcal{P}_\Theta\left[\theta^*\right]\|_2 \leq \|(\nabla A)^{-1}\widehat{\mu} - \theta^*\|_2.$$

By Fisher Consistency of the negative log-likelihood, we know that

$$\nabla A(\theta^*) = \mathbb{E}_{\theta^*}[\phi(z)].$$

The true parameter θ^* can be obtained by inverting the ∇A operator whenever possible.

$$\|(\nabla A)^{-1}\widehat{\mu} - \theta^*\|_2 = \|(\nabla A)^{-1}\widehat{\mu} - (\nabla A)^{-1}\mathbb{E}_{\theta^*}[\phi(z)]\|_2 \tag{61}$$

$$= \|\nabla A^*\widehat{\mu} - \nabla A^*\mathbb{E}_{\theta^*}[\phi(z)]\|_2. \tag{62}$$

where A^* is the convex conjugate of A. We can use the following result to control the Lipschitz smoothness A^*.

Theorem 12. *(Strong/Smooth Duality) Assume $f(\cdot)$ is closed and convex. Then $f(\cdot)$ is smooth with parameter M if and only if its convex conjugate $f(\cdot)$ is strongly convex with parameter $m = \frac{1}{M}$.*

A proof of the above theorem can be found in [47]. Hence, we have that:

$$\|\mathcal{P}_\Theta\left[(\nabla A)^{-1}\widehat{\mu}\right] - \theta^*\|_2 \leq \frac{1}{\tau_\ell}\|\widehat{\mu} - \mathbb{E}_{\theta^*}[\phi(z)]\|_2 \tag{63}$$

By assumption, we have that the fourth moments of the sufficient statistics are bounded. We also know that $\text{Cov}(\phi(z) = \nabla^2 A(\theta^*)$ which implies that we can use [14] as our oracle. Using Lemma 1, we get that, there exists universal constants C_1, C_2 such that with probability at least $1 - 1/p^{C_1}$,

$$\|\widehat{\mu} - \mathbb{E}_{\theta^*}[\phi(z)]\|_2 \leq C_2\sqrt{\tau_u \log p}\left(\epsilon^{\frac{1}{2}} + \gamma(n, p, \delta, \epsilon)\right).$$

Combining the above with Equation (63) recovers the result of Corollary 7.

G   Proof of Theorem 8

Before we present the proof of Theorem 8, we first study the distribution of gradients of the loss function. This will help us bound the error in the gradient estimator.

Lemma 7. *Consider the model in Equation (4). Suppose the covariates $x \in \mathbb{R}^p$ have bounded 4^{th}-moments and the noise w has bounded 2^{th} moments. Then there exist universal constants C_1, C_2 such that*

$$\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)] = \Sigma\Delta$$

$$\|\text{Cov}(\nabla\overline{\mathcal{L}}(\theta)\|_2 \leq \sigma^2\|\Sigma\|_2 + C_1\|\Delta\|_2^2\|\Sigma\|_2^2,$$

where $\Delta = \theta - \theta^*$ and $E[xx^T] = \Sigma$.

Proof. We start by deriving the results for $\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)]$.

$$\overline{\mathcal{L}}(\theta) = \frac{1}{2}(y - x^T\theta)^2 = \frac{1}{2}(x^T(\Delta) - w)^2 \tag{64}$$

$$\nabla\overline{\mathcal{L}}(\theta) = xx^T\Delta - x.w \tag{65}$$

$$\mathbb{E}[\nabla\overline{\mathcal{L}}(\theta)] = \Sigma\Delta. \tag{66}$$

Next, we bound the operator norm of the covariance of the gradients $\nabla\overline{\mathcal{L}}(\theta)$ at any point θ.

Covariance.

$$\text{Cov}(\nabla\overline{\mathcal{L}}(\theta)) = \mathbb{E}[(xx^T - \Sigma)\Delta - x.w)((xx^T - \Sigma)\Delta - x.w)^T)] \tag{67}$$

$$\text{Cov}(\nabla\overline{\mathcal{L}}(\theta)) = \mathbb{E}[(xx^T - \Sigma)\Delta\Delta^T(xx^T - \Sigma)] + \sigma^2\Sigma. \tag{68}$$

Now, we want to bound $\|\text{Cov}(\nabla\overline{\mathcal{L}}(\theta))\|_2 = \lambda_{\max}(\text{Cov}(\nabla\overline{\mathcal{L}}(\theta)))$.

$$\lambda_{\max}(\text{Cov}(\nabla\overline{\mathcal{L}}(\theta))) \le \sigma^2\lambda_{\max}(\Sigma) + \lambda_{\max}\left(\mathbb{E}[(xx^T - \Sigma)\Delta\Delta^T(xx^T - \Sigma)]\right) \tag{69}$$

$$\le \sigma^2\lambda_{\max}(\Sigma) + \sup_{y\in\mathbb{S}^{p-1}} y^T\left(\mathbb{E}[(xx^T - \Sigma)\Delta\Delta^T(xx^T - \Sigma)]\right)y \tag{70}$$

$$\le \sigma^2\lambda_{\max}(\Sigma) + \sup_{y\in\mathbb{S}^{p-1}} y^T\left(\mathbb{E}[(xx^T - \Sigma)\Delta\Delta^T(xx^T - \Sigma)]\right)y \tag{71}$$

$$\le \sigma^2\lambda_{\max}(\Sigma) + \|\Delta\|_2^2 \sup_{y,z\in\mathbb{S}^{p-1}} \mathbb{E}\left[(y^T(xx^T - \Sigma)z)^2\right] \tag{72}$$

$$\le \sigma^2\lambda_{\max}(\Sigma) + \|\Delta\|_2^2 \sup_{y,z\in\mathbb{S}^{p-1}} \left(\mathbb{E}\left[2(y^Tx)^2(x^Tz)^2 + 2(y^T\Sigma z)^2\right]\right) \tag{73}$$

$$\le \sigma^2\lambda_{\max}(\Sigma) + 2\|\Delta\|_2^2 \sup_{y,z\in\mathbb{S}^{p-1}} \left(\mathbb{E}\left[(y^Tx)^2(x^Tz)^2\right] + \|\Sigma\|_2^2\right) \tag{74}$$

$$\le \sigma^2\lambda_{\max}(\Sigma) + 2\|\Delta\|_2^2 \sup_{y,z\in\mathbb{S}^{p-1}} \left(\mathbb{E}\left[(y^Tx)^2(x^Tz)^2\right] + \|\Sigma\|_2^2\right) \tag{75}$$

$$\le \sigma^2\lambda_{\max}(\Sigma) + 2\|\Delta\|_2^2 \sup_{y,z\in\mathbb{S}^{p-1}} \left(\sqrt{\mathbb{E}[(y^Tx)^4]}\sqrt{\mathbb{E}[(z^Tx)^4]} + \|\Sigma\|_2^2\right) \tag{76}$$

$$\le \sigma^2\|\Sigma\|_2 + 2\|\Delta\|_2^2(\|\Sigma\|_2^2 + C_4\|\Sigma\|_2^2), \tag{77}$$

where the second last step follows from Cauchy-Schwartz and the last step follows from our assumption of bounded 4^{th} moments (see Equation (14)). $\qquad\square$

We now proceed to the proof of Theorem 8. From Lemma 2, we know that at any point θ, the gradient estimator described in Algorithm 3, $g(\theta; D_{\widetilde{n}}, \widetilde{\delta})$, satisfies the following with probability at least $1 - \delta$,

$$\|g(\theta; D_{\widetilde{n}}, \widetilde{\delta}) - \nabla\mathcal{R}(\theta)\|_2 \le C\sqrt{\frac{\text{tr}(\text{Cov}(\nabla\overline{\mathcal{L}}(\theta)))\log 1/\widetilde{\delta}}{\widetilde{n}}}.$$

We substitute the upper bound for $\|\text{Cov}(\nabla\overline{\mathcal{L}}(\theta))\|_2$ from Lemma 7 in the above equation

$$
\begin{aligned}
\|g(\theta; D_{\widetilde{n}}, \widetilde{\delta}) - \nabla\mathcal{R}(\theta)\|_2 &\le C\sqrt{\frac{\text{tr}(\text{Cov}(\nabla\overline{\mathcal{L}}(\theta)))\log 1/\widetilde{\delta}}{\widetilde{n}}} \\
&\le C\sqrt{\frac{p\left(\sigma^2\|\Sigma\|_2 + 2\|\Delta\|_2^2(\|\Sigma\|_2^2 + C_4\|\Sigma\|_2^2)\right)\log 1/\widetilde{\delta}}{\widetilde{n}}} \\
&\le \underbrace{C_1\sqrt{\frac{\|\Sigma\|_2^2 p\log 1/\widetilde{\delta}}{\widetilde{n}}}}_{\alpha(\widetilde{n},\widetilde{\delta})}\|\theta - \theta^*\|_2 \\
&\quad + \underbrace{C_2\sigma\sqrt{\frac{\|\Sigma\|_2 p\log 1/\widetilde{\delta}}{\widetilde{n}}}}_{\beta(\widetilde{n},\widetilde{\delta})}.
\end{aligned}
$$

37

To complete the proof of this theorem, we use the results from Theorem 1. Note that the gradient estimator satisfies the stability condition if $\alpha(\widetilde{n}, \widetilde{\delta}) < \tau_l$. This holds when

$$\widetilde{n} > \frac{C_1^2 \tau_u^2}{\tau_l^2} p \log 1/\widetilde{\delta}.$$

Now suppose $\widetilde{n}$ satisfies the above condition, then plugging $\beta(\widetilde{n}, \widetilde{\delta})$ into Theorem 1 gives us the required result.

H   Proof of Theorem 9

To prove the Theorem we use the result from Lemma 4, where we derived the following expression for covariance of $\nabla \overline{\mathcal{L}}(\theta)$

$$\|\text{Cov}(\nabla\overline{\mathcal{L}}(\theta)\|_2 \leq C_1 \|\Delta\|_2^2 \|\Sigma\|_2 \left(\sqrt{C_4}\sqrt{L_{\Phi,4}} + L_{\Phi,2} \right)$$
$$+ C_2 \|\Sigma\|_2 \left(B_{\Phi,2} + \sqrt{B_{\Phi,4}} + c(\sigma)\sqrt{3M_{\Phi,2,2}} + \sqrt{c(\sigma)^3 M_{\Phi,4,1}} \right)$$

From Lemma 2, we know that at any point θ, the gradient estimator described in Algorithm 3, $g(\theta; D_{\widetilde{n}}, \widetilde{\delta})$, satisfies the following with probability at least $1 - \delta$,

$$\|g(\theta; D_{\widetilde{n}}, \widetilde{\delta}) - \nabla\mathcal{R}(\theta)\|_2 \leq C\sqrt{\frac{\text{tr}(\text{Cov}(\nabla\overline{\mathcal{L}}(\theta))) \log 1/\widetilde{\delta}}{\widetilde{n}}}.$$

Substituting the upper bound for $\|\text{Cov}(\nabla\overline{\mathcal{L}}(\theta))\|_2$ in the above equation, we get

$$\|g(\theta; D_{\widetilde{n}}, \widetilde{\delta}) - \nabla\mathcal{R}(\theta)\|_2 \leq C\sqrt{\frac{\text{tr}(\text{Cov}(\nabla\overline{\mathcal{L}}(\theta))) \log 1/\widetilde{\delta}}{\widetilde{n}}}$$
$$\leq \underbrace{C_1 \sqrt{\frac{\|\Sigma\|_2 \left(\sqrt{C_4}\sqrt{L_{\Phi,4}} + L_{\Phi,2} \right) p \log 1/\widetilde{\delta}}{\widetilde{n}}}}_{\alpha(\widetilde{n},\widetilde{\delta})} \|\theta - \theta^*\|_2$$
$$+ \underbrace{C_2 \sqrt{\frac{\|\Sigma\|_2 \left(B_{\Phi,2} + \sqrt{B_{\Phi,4}} + c(\sigma)\sqrt{3M_{\Phi,2,2}} + \sqrt{c(\sigma)^3 M_{\Phi,4,1}} \right) p \log 1/\widetilde{\delta}}{\widetilde{n}}}}_{\beta(\widetilde{n},\widetilde{\delta})}.$$

We now use the results from Theorem 1. The gradient estimator satisfies the stability condition if $\alpha(\widetilde{n}, \widetilde{\delta}) < \tau_l$. This holds when

$$\widetilde{n} > \frac{C_1^2 \|\Sigma\|_2 \left(\sqrt{C_4}\sqrt{L_{\Phi,4}} + L_{\Phi,2} \right)}{\tau_l^2} p \log 1/\widetilde{\delta}.$$

Now suppose $\widetilde{n}$ satisfies the above condition, then plugging $\beta(\widetilde{n}, \widetilde{\delta})$ into Theorem 1 gives us the required result.

38

I   Proof of Theorem 11

The proof proceeds along similar lines as the proof of Theorem 9. To prove the Theorem we utilize the result of Lemma 6, where we showed that $\|\text{Cov}[\nabla\overline{\mathcal{L}}(\theta)]\|_2 = \|\nabla^2 A(\theta^*)\|_2$. Combining this result with Lemma 2 we get that with probability at least $1 - \delta$

$$
\begin{aligned}
\|g(\theta; D_{\widetilde{n}}, \widetilde{\delta}) - \nabla\mathcal{R}(\theta)\|_2 &\leq C\sqrt{\frac{\text{tr}(\text{Cov}(\nabla\overline{\mathcal{L}}(\theta)))\log 1/\widetilde{\delta}}{\widetilde{n}}} \\
&\leq \underbrace{C\sqrt{\frac{\|\nabla^2 A(\theta^*)\|_2 p \log 1/\widetilde{\delta}}{\widetilde{n}}}}_{\beta(\widetilde{n},\widetilde{\delta})}.
\end{aligned}
$$

Since $\alpha(\widetilde{n}, \widetilde{\delta}) = 0$, the stability condition is always satisfied, as long as $\tau_l > 0$. Substituting $\beta(\widetilde{n}, \widetilde{\delta})$ into Theorem 1 gives us the required result.

J   Upper bound on Contamination Level

We provide a complementary result, which gives an upper bound for the contamination level ϵ based on the initialization point θ^0, above which, Algorithm 1 would not work. The key idea is that the error incurred by any mean estimation oracle is lower bounded by the variance of the distribution, and that if the zero vector lies within that error ball, then any mean oracle can be forced to output $\mathbf{0}$ as the mean. For Algorithm 1, this implies that, in estimating the mean of the gradient, if the error is high, then one can force the mean to be $\mathbf{0}$ which forces the algorithm to converge. For the remainder of the section we consider the case of linear regression with $x \sim \mathcal{N}(0, \mathcal{I}_p)$ in the asymptotic regime of $n \to \infty$.

Lemma 8. *Consider the model in equation(4) with $x \sim \mathcal{N}(0, \mathcal{I}_p)$ and $w \sim \mathcal{N}(0, 1)$, then there exists a universal constant C_1 such that if $\epsilon > C_1 \frac{\|\theta^0 - \theta^*\|_2}{\sqrt{1 + 2\|\theta^0 - \theta^*\|_2^2}}$, then for every gradient oracle, there exists a contamination distribution Q such that, Algorithm 1 will converge to θ^0 even when the number of samples $n \to \infty$.*

Proof. Using Lemma 5, we know that for any point θ,

$$
\begin{aligned}
\nabla\overline{\mathcal{L}}(\theta) &= xx^T\Delta - x.w \\
\mathbb{E}_{\theta^*}[\nabla\overline{\mathcal{L}}(\theta)] &= (\theta - \theta^*) = \Delta \\
\|\text{Cov}(\nabla\overline{\mathcal{L}}(\theta)\|_2 &= 1 + 2\|\Delta\|_2^2,
\end{aligned}
$$

where $\Delta = \theta - \theta^*$.

Let $P_{\nabla\overline{\mathcal{L}}(\theta)}$ represent the distribution $\nabla\overline{\mathcal{L}}\theta$. Similarly, let $P_{\epsilon, \nabla\overline{\mathcal{L}}(\theta), Q}$ represent the corresponding ϵ-contaminated distribution. Then, using Theorem 2.1 [10], we know that the minimax rate for estimating the mean of the distribution of gradients is given by:

$$
\inf_{\widehat{\mu}} \sup_{\theta \in \mathbb{R}^p, Q} P_{\epsilon, \nabla\overline{\mathcal{L}}(\theta), Q} \left\{ \|\widehat{\mu} - \mathbb{E}_{\theta^*}[\nabla\overline{\mathcal{L}}(\theta)]\|_2^2 \geq C\epsilon^2(1 + 2\|\Delta\|_2^2) \right\} \geq c.
$$

The above statement says that at any point θ, any mean oracle Ψ will always incur an error of $\Omega(\sqrt{C\epsilon^2(1 + 2\|\Delta\|_2^2)})$ in estimating the gradient $\mathbb{E}_{\theta^*}[\nabla\overline{\mathcal{L}}(\theta)]$.

$$
\|\Psi(\theta) - \mathbb{E}_{\theta^*}[\nabla\overline{\mathcal{L}}(\theta)]\|_2 \geq C\epsilon\sqrt{(1 + 2\|\Delta\|_2^2)} \quad \forall \; \Psi
$$

39

For any oracle Ψ, there exists some adversarial contamination Q, such that whenever $\|\mathbb{E}_{\theta^*}[\nabla\overline{\mathcal{L}}(\theta)]\|_2 < C\epsilon\sqrt{(1 + 2\|\Delta\|_2^2)}$, then $\|\Psi(\theta)\|_2 = 0$.

Suppose that the contamination level ϵ is such that,

$$\epsilon > \frac{1}{C}\frac{\|\mathbb{E}_{\theta^*}[\nabla\overline{\mathcal{L}}(\theta^0)]\|_2}{\sqrt{(1 + 2\|\theta^0 - \theta^*\|_2^2)}},$$

then for every oracle there exists a corresponding Q such that Algorithm 1 will remain stuck at θ^0.

Plugging $\mathbb{E}_{\theta^*}[\nabla\overline{\mathcal{L}}(\theta^0)] = \theta^0 - \theta^*$, we recover the statement of the lemma. $\qquad\square$

Chen et al. [10] provide a general minimax lower bound of $\Omega(\epsilon)$ for ϵ-contamination models in this setting. In contrast, using Algorithm 1 with [14] as oracle, we can only $O(\sqrt{\epsilon\log p})$ close to the true parameter even when the contamination is small, which implies that our procedure is not minimax optimal. Our approach is nonetheless the only practical algorithm for robust estimation of general statistical models.

K   Proof of Lemma 1

In this section we present a refined, non-asymptotic analysis of the robust mean estimator of [14], described in Algorithm 2. We begin by introducing some preliminaries. We subsequently analyze the algorithm in 1-dimension and finally turn our attention to the general algorithm.

K.1   Preliminaries

Unless otherwise stated, we assume throughout that the random variable X has bounded fourth moments, i.e. for every unit vector v,

$$\mathbb{E}\left[\langle X - \mu,\ v\rangle^4\right] \leq C_4\left[\mathbb{E}\left[\langle X - \mu,\ v\rangle^2\right]\right]^2.$$

We summarize some useful results from [14], which bound the deviation of the conditional mean/covariance from the true mean/covariance.

Lemma 9. *[Lemma 3.11 [14]] Let X be a univariate random variable with bounded fourth moments, and let A be any with event with probability $\mathbb{P}(A) = 1 - \gamma \geq \frac{1}{2}$. Then,*

$$|\mathbb{E}(X|A) - \mathbb{E}(X)| \leq \sigma\sqrt[4]{8C_4\gamma^3}.$$

Lemma 10. *[Lemma 3.12 [14]] Let X be a univariate random variable with $\mathbb{E}[X] = \mu$, $\mathbb{E}\left((X - \mu)^2\right) = \sigma^2$ and let $\mathbb{E}((X - \mu)^4) \leq C_4\sigma^4$. Let A be any with event with probability $\mathbb{P}(A) = 1 - \gamma \geq \frac{1}{2}$. Then,*

$$(1 - \sqrt{C_4\gamma})\sigma^2 \leq \mathbb{E}((X - \mu)^2|A) \leq (1 + 2\gamma)\sigma^2.$$

Corollary 13. *[Corollary 3.13 [14]] Let A be any event with probability $\mathbb{P}(A) = 1 - \gamma \geq \frac{1}{2}$, and let X be a random variable with bounded fourth moments. We denote $\Sigma|_A = \mathbb{E}(XX^T|A) - (\mathbb{E}(X|A))(\mathbb{E}(X|A))^T$ to be the conditional covariance matrix. We have that,*

$$(1 - \sqrt{C_4\gamma} - \sqrt{8C_4\gamma^3})\Sigma \preceq \Sigma|_A \preceq (1 + 2\gamma)\Sigma.$$

For random variables with bounded fourth moments we can use Chebyshev's inequality to obtain tail bounds.

Lemma 11. *[Lemma 3.14 [14]] Let X have bounded fourth moments, then for every unit vector v we have that,*

$$\mathbb{P}(|\langle X,\ v\rangle - \mathbb{E}[\langle X,\ v\rangle]| \geq t\sqrt{[\mathbb{E}\left[\langle X - \mu,\ v\rangle^2]\right]}) \leq \frac{C_4}{t^4}.$$

Our proofs also use the matrix Bernstein inequality for rectangular matrices. As a preliminary, we consider a finite sequence $\{Z_k\}$ of independent, random matrices of size $d_1 \times d_2$. We assume that each random matrix satisfies $\mathbb{E}(Z_k) = 0$, and $\|Z_k\|_{\text{op}} \leq R$ almost surely. We define:

$$\sigma^2 := \max\left\{\|\sum_k \mathbb{E}(Z_k Z_k^T)\|_{\text{op}}, \|\sum_k \mathbb{E}(Z_k Z_k^T)\|_{\text{op}}\right\}.$$

With these preliminaries in place we use the following result from [48].

Lemma 12. *For all $t \geq 0$,*

$$\mathbb{P}\left(\left\|\sum_k Z_k\right\|_{op} \geq t\right) \leq (d_1 + d_2)\exp\left(\frac{-t^2/2}{\sigma^2 + Rt/3}\right).$$

Equivalently, with probability at least $1 - \delta$,

$$\left\|\sum_k Z_k\right\|_{op} \leq \sqrt{2\sigma^2 \log\left(\frac{d_1 + d_2}{\delta}\right)} + \frac{2R}{3}\log\left(\frac{d_1 + d_2}{\delta}\right).$$

We let $\mathcal{I}$ denote the set of all intervals in $\mathbb{R}$. The following is a standard uniform convergence result.

Lemma 13. *Suppose $X_1, \ldots, X_n \sim \mathbb{P}$, then with probability at least $1 - \delta$,*

$$\sup_{I \in \mathcal{I}}\left|\mathbb{P}(I) - \frac{1}{n}\sum_{i=1}^{n}\mathbb{I}(X_i \in I)\right| \leq 2\sqrt{\frac{4\log(en) + 2\log(2/\delta)}{n}}.$$

Algorithm 4 Huber Outlier Gradients Truncation

function HUBEROUTLIERGRADIENTTRUNCATION(SAMPLE GRADIENTS S, CORRUPTION LEVEL ϵ, DIMENSION p,δ)

    if p=1 **then**

        Let $[a,b]$ be smallest interval containing $\left(1 - \epsilon - C_5\left(\sqrt{\frac{1}{|S|}\log\left(\frac{|S|}{\delta}\right)}\right)\right)(1-\epsilon)$ fraction of points.

        $\widetilde{S} \leftarrow S \cap [a,b]$.

        return $\widetilde{S}$

    else

        Let $[S]_i$ be the samples with the i^{th} co-ordinates only, $[S]_i = \{\langle x, e_i\rangle\,|x \in S\}$

        for $i = 1$ to p **do**

            $a[i] = $ HUBERGRADIENTESTIMATOR$([S]_i, \epsilon, 1, \delta/p)$.

        end for

        Let $B(r,a)$ be the ball of smallest radius centered at a containing $(1 - \epsilon - C_p\left(\sqrt{\frac{p}{|S|}\log\left(\frac{|S|}{p\delta}\right)}\right)(1-\epsilon)$ fraction of points in S.

        $\widetilde{S} \leftarrow S \cap B(r,a)$.

        return $\widetilde{S}$

    end if

end function

We now turn our attention to an analysis of Algorithm 2 for the 1-dimensional case.

K.2   The case when $p = 1$

Firstly, we analyze Algorithm 2 when $p = 1$.

Lemma 14. *Suppose that, P_θ^* is a distribution on $\mathbb{R}^1$ with mean μ, variance σ^2, and bounded fourth moments. There exist positive universal constants $C_1, C_2, C_8 > 0$, such that given n samples from the distribution in (9), the algorithm with probability at least $1 - \delta$, returns an estimate $\widehat{\mu}$ such that,*

$$\|\widehat{\mu} - \mu\|_2 \leq C_1 C_4^{\frac{1}{4}}\sigma\left(\epsilon + \sqrt{\frac{\log 3/\delta}{2n}} + t\right)^{\frac{3}{4}} + C_2\sigma\left(\epsilon + \sqrt{\frac{\log 3/\delta}{2n}} + t\right)^{\frac{1}{2}}\sqrt{\frac{\log(3/\delta)}{n}}$$

where $t = C_8\sqrt{\frac{1}{n}\log\left(\frac{n}{\delta}\right)}$. which can be further simplified to,

$$\|\widehat{\mu} - \mu\|_2 \leq C_1 C_4^{\frac{1}{4}}\sigma\left(\epsilon + C_8\sqrt{\frac{1}{n}\log\left(\frac{n}{\delta}\right)}\right)^{\frac{3}{4}} + C_2\sigma\left(\epsilon + C_8\sqrt{\frac{1}{n}\log\left(\frac{n}{\delta}\right)}\right)^{\frac{1}{2}}\sqrt{\frac{\log(1/\delta)}{n}}$$

Proof. By an application of Hoeffding's inequality we obtain that with probability at least $1 - \delta/3$, the fraction of corrupted samples (i.e. samples from the distribution Q) is less than $\epsilon + \sqrt{\frac{\log(3/\delta)}{2n}}$. We condition on this event through the remainder of this proof. We let η denote the fraction of corrupted samples. Further, we let S_P be the samples from the true distribution. Let n_P be the cardinality of this set, i.e. $n_P := |S_P|$.

Let $I_{1-\eta}$ be the interval around μ containing $1 - \eta$ mass of P_θ^*. Then, using Lemma 11, we have that:

$$\text{length}\,(I_{1-\eta}) \leq \frac{C_4^{\frac{1}{4}}\sigma}{\eta^{\frac{1}{4}}}.$$

Using Lemma 13 we obtain that with probability at least $1 - \delta/3$ the number of samples from the distribution P that fall in the interval $I_{1-\eta}$ is at least $1 - \eta - t$ where t is upper bounded as:

$$t \leq 2\sqrt{\frac{4\log(en) + 2\log(6/\delta)}{n}}.$$

Now we let $\widetilde{S}$ be the set of points in the smallest interval containing $(1 - \eta - t)(1 - \eta)$ fraction of all the points.

- Using VC theory, we know that for every interval $I \subset \mathbb{R}$, there exists some universal constant C_3 such that

$$\mathbb{P}\left(|(P(x \in I | x \sim D) - P(x \in I | x \in_u S_D))| > t/2\right) \leq n_D^2 \exp(-n_D t^2 / 8) \qquad (78)$$

  This can be re-written as, that with probability at least $(1 - \delta/3)$, there exists a universal constant C_0 such that,

$$\sup_I |(P(x \in I | x \sim D) - P(x \in I | x \in_u S_D))| \leq C_0 \sqrt{\frac{1}{n_D}\log\left(\frac{n_D}{\delta}\right)} \leq \underbrace{C_5 \sqrt{\frac{1}{n}\log\left(\frac{n}{\delta}\right)}}_{t}$$

- Using Equation (78), we know that $(1 - \eta - t)$ fraction of S_D lie in $\mathcal{I}_{1-\eta}$.

  Let $\widetilde{S}$ be the set of points in the smallest interval containing $(1 - \eta - t)(1 - \eta)$ fraction of the points.

- We know that the length of minimum interval containing $(1 - \eta - t)(1 - \eta)$ fraction of the points of S is less than length of smallest interval containing $(1 - \eta - t)$ fraction of points of S_D, which in turn is less than length of $I_{1-\eta}$.

- Now, $I_{1-\eta}$ and minimum interval containing $(1 - \eta - t)$ fraction of points of S_D need to overlap. This is because, n is large enough such that $t < \frac{1}{2} - \eta$ hence, the extreme points for such an interval can be almost $2\text{length}\,(I_{1-\eta})$ away.

- Hence, the distance of all chosen noise-points from μ will be within the length $(I_{1-\eta})$.

- Moreover, the interval of minimum length with $(1-\eta-t)(1-\eta)$ fraction of S will contain at least $1 - 3\eta - t$ fraction of S_D.

- Hence, we can bound the error of $mean(\widetilde{S})$ by controlling the sources of error.

  - All chosen noise points are within length $(I_{1-\eta})$, and there are almost η of them, hence the maximum error can be $\eta \text{length}\,(I_{1-\eta})$.

43

- Next, the mean of chosen good points will converge to the mean of the conditional distribution. *i.e.* points sampled from D but conditioned to lie in the minimum length interval. The variance of these random variables is upper bounded using Lemma 10.
- To control the distance between the mean($E(X)$ and the conditional mean($E(X|A)$), where A is the event that a sample x is in the chosen interval. We know that $P(A) \geq 1 - 3\eta - t$, hence, using Lemma 3.11[14], we get that there exists a constant C_{13} such that,

$$|E[X] - E[X|A]| \leq C_{13}C_4^{\frac{1}{4}}\sigma(\eta + t)^{\frac{3}{4}}$$

- Hence, with probability at least $1 - \delta/3$, the mean of $\widetilde{S}$ will be within

$$\eta \times \text{length}\,(I_{1-\eta}) + C_{13}C_4^{\frac{1}{4}}\sigma(\eta + t)^{\frac{3}{4}} + C_6\sigma(1 + 2\eta)^{\frac{1}{2}}\sqrt{\frac{\log(3/\delta)}{n}}$$

- Taking union-bound over all conditioning statements, and upper bounding, η with $\epsilon + \sqrt{\frac{\log(3/\delta)}{2n}}$, we recover the statement of the lemma.

$\square$

K.3 The case when $p > 1$

To prove the case for $p > 1$, we use a series of lemmas. Lemma 15 proves that the outlier filtering constrains the points in a ball around the true mean. Lemma 17 controls the error in the mean and covariance the true distribution after outlier filtering ($\widetilde{D}$). Lemma 18 controls the error for the mean of $\widetilde{S}$ when projected onto the bottom span of the covariance matrix $\Sigma_{\widetilde{S}}$.

Lemma 15. *Suppose that, P_θ^* is a distribution on $\mathbb{R}^p$ with mean μ, covariance Σ, and bounded fourth moments. There exist positive universal constants $C_1, C_2, C_8 > 0$, such that given n samples from the distribution in Equation* (9), *we can find a vector $a \in \mathbb{R}^p$ such that with probability at least $1 - \delta$,*

$$\|a - \mu\|_2 \leq C_1 C_4^{\frac{1}{4}}\sqrt{tr\,(\Sigma)}\left(\epsilon + C_8\sqrt{\frac{1}{n}\log\left(\frac{np}{\delta}\right)}\right)^{\frac{3}{4}}$$

$$+ C_2\left(\epsilon + C_8\sqrt{\frac{1}{n}\log\left(\frac{np}{\delta}\right)}\right)^{\frac{1}{2}}\sqrt{tr\,(\Sigma)}\sqrt{\frac{\log(p/\delta)}{n}}$$

Proof. Pick n orthogonal directions $v_1, v_2, \ldots, v_n$, and use method for one-dimensions, and using union bound, we can recover the result. $\square$

Next, we prove the case when $p > 1$. Firstly, we prove that after the outlier step,

Lemma 16. *After the outlier removal step, there exists universal constants $C_{11} > 0$ such that with probability at least $1 - \delta$, every remaining point x satisfies,*

$$\|x - \mu\|_2 \leq r_1^* + 2r_2^*$$

where $r_1^* = C_{10}\dfrac{C_4^{\frac{1}{4}}\sqrt{p\|\Sigma\|_2}}{\eta^{\frac{1}{4}}}$ and $r_2^* = C_1 C_4^{\frac{1}{4}}\sqrt{p\|\Sigma\|_2}(\eta + t)^{\frac{3}{4}} + C_2\sqrt{p\|\Sigma\|_2}(\eta + t)^{\frac{1}{2}}\sqrt{\dfrac{\log(1/\delta)}{n}}$ and $t = C_8\sqrt{\dfrac{1}{n}\log\left(\dfrac{np}{\delta}\right)}$. Here $\eta \leq \epsilon + \sqrt{\dfrac{\log(1/\delta)}{2n}}$ is the fraction of samples corrupted.

Proof. • Let $\widetilde{S}$ be the set of points chosen after the outlier filtering. Let $\widetilde{S}_D$ be set of good points chosen after the outlier filtering. Let $\widetilde{S}_N$ be the set of bad points chosen after the outlier filtering.

• Using VC theory we know that for every closed ball $\mathcal{B}(\mu, r) = \{x | \|x - \mu\|_2 \leq r\}$, there exists a constant C_9 such that with probability at least $1 - \delta$

$$\sup_B |P(x \in B | x \sim D) - P(x \in B | x \in_u S_D)| \leq \underbrace{C_9 \sqrt{\frac{p}{n} \log\left(\frac{n}{p\delta}\right)}}_{t_2}$$

• Let $B^* = B(\mu, r^*)$ for $r_1^* = C_{10} \frac{C_4^{\frac{1}{4}}}{(\eta)^{\frac{1}{4}}} \sqrt{p \|\Sigma\|_2}$. Then, we claim that

$$P(x \in B^* | x \sim D) \geq 1 - \eta$$

– To see this, suppose we have some $x \in D$. Let $z = x - \mu$. Let $z_i = z^T v_i$ for some orthogonal directions $v_1, v_2, \ldots, v_p$. Let $Z^2 = \sum z_i^2 = \|z\|_2^2$.

–

$$P\left(Z^2 \geq \frac{C_4^{\frac{1}{2}} p \|\Sigma\|_2}{(\eta)^{\frac{1}{2}}}\right) = P\left(Z^4 \geq \frac{C_4 p^2 \|\Sigma\|_2^2}{(\eta)}\right) \leq \frac{(\eta) E(Z^4)}{C_4 p^2 \|\Sigma\|_2^2}$$

– Now, $E(Z^4) \leq p^2 \max_i E(z_i^4) \leq C_4 p^2 \|\Sigma\|_2^2$. Plugging this in the above, we have that $P(x \in B^* | x \sim D) \geq 1 - \eta$.

• Hence, we have that $P(x \in B^* | x \in_u S_D) \geq 1 - \eta - t_2$.

• Using Lemma 15, we have that at least $(1 - \eta - t_2)$ fraction of good points are $r_1^* + r_2^*$ away from a. Hence, we have that the minimum radius of the ball containing all the $(1 - \eta - t_2)(1 - \eta)$ has a radius of atmost $r_1^* + r_2^*$, which when combined with the triangle inequality recovers the statement of lemma. □

As before, let $\widetilde{S}$ be the set of points after outlier filtering. Let $\mu_{\widetilde{S}} = mean(\widetilde{S})$, $\mu_{\widetilde{S}_D} = mean(\widetilde{S}_D)$, $\mu_{\widetilde{S}_N} = mean(\widetilde{S}_N)$.

Lemma 17. *Let $\widetilde{S}_D$ be the set of clean points remaining after the outlier filtering. Then, with probability at least $1 - \delta$, we have that*

$$\|\mu_{\widetilde{S}_D} - \mu\|_2 \leq C_1 C_4^{\frac{1}{4}} (\eta + t_2)^{\frac{3}{4}} \sqrt{\|\Sigma\|_2} \left(1 + \frac{\log(p/\delta)}{n}\right) + \sqrt{\|\Sigma\|_2} \sqrt{(1 + 2(\eta + t_2))} \sqrt{\frac{1}{n} \log(p/\delta)}$$
$$+ C_{15} \frac{(r_1^* + 2r_2^*)}{n} \log(p/\delta)$$

and

$$\|\Sigma_{\widetilde{S}_D}\|_2 \leq \beta(n, \delta) \|\Sigma\|_2,$$

where

$$\beta(n, \delta) = \left(1 + 2C(\eta + t_2) + \left(1 + \frac{p\sqrt{C_4}}{\sqrt{\eta}} + \sqrt{C_4} p(\eta + t)^{\frac{3}{2}} + (\eta + t_2)^{\frac{3}{2}}\right) \left(\sqrt{\frac{\log(p/\delta)}{n}} + \frac{\log(p/\delta)}{n}\right)\right)$$

45

Proof. We first prove the bounds on the mean shift.

$$\|\mu_{\widetilde{S}_D} - \mu\|_2 \leq \underbrace{\|\mu_{\widetilde{S}_D} - \mu_{\widetilde{D}}\|_2}_{A} + \underbrace{\|\mu_{\widetilde{D}} - \mu\|_2}_{B}$$

- **Control of B.** We use Lemma 9 on $X = x^T \frac{\mu_{\widetilde{D}} - \mu}{\|\mu_{\widetilde{D}} - \mu\|_2}$ for $x \sim D$, and A be the event that x is not removed by the outlier filtering.

$$\|\mu_{\widetilde{D}} - \mu\|_2 \leq C_1 C_4^{\frac{1}{4}} (\eta + t_2)^{\frac{3}{4}} \sqrt{\|\Sigma\|_2} \tag{79}$$

- **Control of A**. Using Lemma 10, we have that $\|\Sigma_{\widetilde{D}}\|_2 \leq (1 + 2(\eta + t_2))\|\Sigma\|_2$. Now, we use Bernstein's inequality . Lemma 12 with $R = C(r_1^* + 2r_2^* + B)$, we get that, with probability at least $1 - \delta$,

$$\|\mu_{\widetilde{S}_D} - \mu_{\widetilde{D}}\|_2 \leq C_{14}\sqrt{\|\Sigma\|_2}\sqrt{(1 + 2(\eta + t_2)}\sqrt{\frac{1}{n}\log(p/\delta)} + C_{15}\frac{(r_1^* + 2r_2^* + B)}{n}\log(p/\delta) \tag{80}$$

Next, we prove the bound for covariance matrix.

$$\|\Sigma_{\widetilde{S}_D}\|_2 \leq \|\Sigma_{\widetilde{S}_D} - \Sigma_{\widetilde{D}}\|_2 + \underbrace{\|\Sigma_{\widetilde{D}} - \Sigma\|_2}_{\leq 2C(\eta + t_2)\|\Sigma\|_2 (\text{By Corollary } 13))} + \|\Sigma\|_2 \tag{81}$$

To control $\|\Sigma_{\widetilde{S}_D} - \Sigma_{\widetilde{D}}\|_2$, we use Bernstein's inequality, with $Z_k = \frac{(x_k - \mu_{\widetilde{D}})(x_k - \mu_{\widetilde{D}})^T - \Sigma_{\widetilde{D}}}{n}$. From, Lemma 16, we know that the points are constrained in a ball. Plugging this into Lemma 12,

$$\|\Sigma_{\widetilde{S}_D} - \Sigma_{\widetilde{D}}\|_2 \leq C(\|\Sigma\|_2 + R^2)\left(\sqrt{\frac{\log(p/\delta)}{n}} + \frac{\log(p/\delta)}{n}\right)$$

where $R^2 = C\left(r_1^{*2} + r_2^{*2} + B^2\right)$. $\qquad\square$

Plugging in the values, we get that,

$$\|\Sigma_{\widetilde{S}_D} - \Sigma_{\widetilde{D}}\|_2 \leq C\|\Sigma\|_2\left(1 + \frac{p\sqrt{C_4}}{\sqrt{\eta}} + \sqrt{C_4}p(\eta + t)^{\frac{3}{2}} + (\eta + t_2)^{\frac{3}{2}}\right)\left(\sqrt{\frac{\log(p/\delta)}{n}} + \frac{\log(p/\delta)}{n}\right)$$

Finally, we have that,

$$\|\Sigma_{\widetilde{S}_D}\|_2 \leq \|\Sigma\|_2 \underbrace{\left(1 + 2C(\eta + t_2) + \left(1 + \frac{p\sqrt{C_4}}{\sqrt{\eta}} + \sqrt{C_4}p(\eta + t)^{\frac{3}{2}} + (\eta + t_2)^{\frac{3}{2}}\right)\left(\sqrt{\frac{\log(p/\delta)}{n}} + \frac{\log(p/\delta)}{n}\right)\right)}_{\beta(n,\delta)}$$

Lemma 18. *Let W be the bottom $p/2$ principal components of the covariance matrix after filtering $\Sigma_{\widetilde{S}}$. Then there exists a universal constant $C > 0$ such that with probability at least $1 - \delta$, we have that*

$$\|\eta P_W \delta_\mu\|_2^2 \leq C\eta\left(\beta(n,\delta) + \gamma(n,\delta)C_4^{\frac{1}{2}})\|\Sigma\|_2\right),$$

46

where $\delta_\mu = \mu_{\widetilde{S}_N} - \mu_{\widetilde{S}_D}$, P_W is the projection matrix on the bottom $p/2$-span of $\Sigma_{\widetilde{S}}$, $\beta(n, \delta)$ is as defined in Lemma 17 and $\gamma(n, \delta) = \left(\eta^{\frac{1}{2}} + (\eta + t)^{5/2} + \eta(\eta + t)\frac{\log(1/\delta)}{n} \right)$

Proof. We have

$$\Sigma_{\widetilde{S}} = \underbrace{(1 - \eta)\Sigma_{\widetilde{S}_D}}_{E} + \underbrace{\eta\Sigma_{\widetilde{S}_N} + (\eta - \eta^2)\delta_\mu\delta_{\mu^T}}_{F} \tag{82}$$

$$\tag{83}$$

By Weyl's inequality we have that,

$$\lambda_{p/2}(\Sigma_{\widetilde{S}}) \leq \lambda_1(E) + \lambda_{p/2}(F)$$

- **Control of $\lambda_{p/2}(F)$.**

$$\begin{aligned}
\lambda_{p/2}(F) &\leq \frac{\operatorname{tr}(F)}{p/2} \\
&\leq C_{15}\eta\frac{((r_1^*)^2 + (r_2^*)^2) + B^2}{p/2} \\
&\leq C_{16}C_4^{\frac{1}{2}}\|\Sigma\|_2 \underbrace{\left(\eta^{\frac{1}{2}} + (\eta + t)^{5/2} + \eta(\eta + t)\frac{\log(1/\delta)}{n} \right)}_{\gamma(n,\delta)}
\end{aligned}$$

where $t = C_8\sqrt{\frac{1}{n}\log\left(\frac{np}{\delta}\right)}$.

- **Control of $\lambda_1(E)$.**

$$\lambda_1(E) \leq (1 - \eta)\beta\|\Sigma\|_2$$

Hence, we have that:

$$\lambda_{p/2}(\Sigma_{\widetilde{S}}) \leq (1 - \eta)\beta\|\Sigma\|_2 + C_{16}\gamma\sqrt{C_4}\|\Sigma\|_2$$

Using that W is the space spanned by the bottom $p/2$ eigenvectors of $\Sigma_{\widetilde{S}}$ and P_W is corresponding projection operator, we have that:

$$P_W^T\Sigma_{\widetilde{S}}P_W \preceq \left[(1 - \eta)\beta + C_{16}\gamma\sqrt{C_4} \right]\|\Sigma\|_2 I_p$$

Following some algebraic manipulation in [14], we get that,

$$\|\eta P_W\delta_\mu\|_2^2 \leq \eta\left((\beta(n, \delta) + \gamma C_4^{\frac{1}{2}})\|\Sigma\|_2 \right)$$

$\square$

Having established all required results, we are now ready to prove Lemma 1. We first present a result for general mean estimation. The proof of Lemma 1 then follows directly from this result.

Lemma 19. *Suppose that, P_θ^* is a distribution on $\mathbb{R}^p$ with mean μ, covariance Σ and bounded fourth moments. There exist positive universal constant $C > 0$, such that given n samples from the distribution in Equation* (9), *the algorithm with probability at least $1 - \delta$, returns an estimate $\widehat{\mu}$ such that,*

$$\|\widehat{\mu} - \mu\|_2 \leq C\|\Sigma\|_2^{\frac{1}{2}}(1 + \sqrt{\log p})\left(\sqrt{\eta} + C_4^{\frac{1}{4}}(\eta + t_2)^{\frac{3}{4}} + \left(\sqrt{\eta}pC_4^{\frac{1}{2}}\sqrt{\frac{\log p \log(p\log(p/\delta))}{n}}\right)^{\frac{1}{2}}\right)$$

where $\eta = \epsilon + \sqrt{\frac{\log(p)\log(\log p/\delta)}{2n}}$ and $t_2 = \sqrt{\frac{p\log(p)\log(n/(p\delta))}{n}}$.

Proof. We divide n samples into $\lfloor\log(p)\rfloor$ different sets. We choose the first set and keep that as our active set of samples. We run our outlier filtering on this set, and let the remaining samples after the outlier filtering be $\widetilde{S}_D$. By orthogonality of subspaces spanned by eigenvectors, coupled with triangle inequality and contraction of projection operators, we have that

$$\|\widehat{\mu} - \mu\|_2^2 \leq 2\|P_W(\widehat{\mu} - \mu_{\widetilde{S}_D})\|_2^2 + 2\|P_W(\mu_{\widetilde{S}_D} - \mu)\|_2^2 + \|\widehat{\mu}_V - P_V\mu\|_2^2$$

$$\|\widehat{\mu} - \mu\|_2^2 \leq 2\|P_W(\widehat{\mu} - \mu_{\widetilde{S}_D})\|_2^2 + 2\|(\mu_{\widetilde{S}_D} - \mu)\|_2^2 + \|\widehat{\mu}_V - P_V\mu\|_2^2$$

where V is the span of the top $p/2$ principal components of $\Sigma_{\widetilde{S}}$ and where $\widehat{\mu}_V$ is the mean vector of returned by the running the algorithm on the reduced dimensions $dim(V) = p/2$. From Lemma 18, both $\beta(n,\delta)$ and $\gamma(n,\delta)$ are monotonically increasing in the dimension; moreover the upper bound in Lemma 17 is also monotonically increasing in the dimension p, hence, the error at each step of the algorithm can be upper bounded by error incurred when running on dimension p, with $n/\log(p)$ samples, and probability of $\delta/\log p$. Hence, the overall error for the recursive algorithm can be upper bounded as,

$$\|\widehat{\mu} - \mu\|_2^2 \leq \left(2\|P_W(\widehat{\mu} - \mu_{\widetilde{S}_D})\|_2^2 + 2\|\mu_{\widetilde{S}_D} - \mu\|_2^2\right)(1 + \log p)$$

Combining Lemma 17 and Lemma 18 which are instantiated for $n/\log p$ samples and probability $\delta/\log p$, we get,

$$\|\widehat{\mu} - \mu\|_2 \leq C\|\Sigma\|_2^{\frac{1}{2}}\sqrt{\log p}\left(\sqrt{\eta} + C_4^{\frac{1}{4}}(\eta + t_2)^{\frac{3}{4}} + \left(\sqrt{\eta}pC_4^{\frac{1}{2}}\sqrt{\frac{\log p \log(p\log(p/\delta))}{n}}\right)^{\frac{1}{2}}\right)$$

$\square$