



Contents lists available at ScienceDirect

Econometrics and Statistics

journal homepage: www.elsevier.com/locate/ecosta

Robust empirical risk minimization via Newton's method

Eirini Ioannou^a, Muni Sreenivas Pydi^b, Po-Ling Loh^{c,*}^a University of Edinburgh, United Kingdom^b LAMSADE, University of Paris Dauphine-PSL, France^c University of Cambridge, United Kingdom

ARTICLE INFO

Article history:

Received 26 January 2023

Revised 14 July 2023

Accepted 16 July 2023

Available online xxx

ABSTRACT

A new variant of Newton's method for empirical risk minimization is studied, where at each iteration of the optimization algorithm, the gradient and Hessian of the objective function are replaced by robust estimators taken from existing literature on robust mean estimation for multivariate data. After proving a general theorem about the convergence of successive iterates to a small ball around the population-level minimizer, consequences of the theory in generalized linear models are studied when data are generated from Huber's epsilon-contamination model and/or heavy-tailed distributions. An algorithm for obtaining robust Newton directions based on the conjugate gradient method is also proposed, which may be more appropriate for high-dimensional settings, and conjectures about the convergence of the resulting algorithm are offered. Compared to robust gradient descent, the proposed algorithm enjoys the faster rates of convergence for successive iterates often achieved by second-order algorithms for convex problems, i.e., quadratic convergence in a neighborhood of the optimum, with a stepsize that may be chosen adaptively via backtracking linesearch.

© 2023 The Author(s). Published by Elsevier B.V. on behalf of EcoSta Econometrics and Statistics.

This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

1. Introduction

Statistical estimation via classical procedures often depends on strong model assumptions, which only hold in the absence of outliers and other deviations. However, many real-life data sets do not typically follow these model assumptions, necessitating the use of robust statistical methods (Huber and Ronchetti, 2011; Rousseeuw et al., 2011; Maronna et al., 2019), which remain reasonably accurate even under deviations from the model assumptions. In this paper, we focus on situations where data are sampled from a small ball around a parametric distribution, according to Huber's ϵ -contamination model. In other words, we have samples of the form $z_i \sim (1 - \epsilon)P_{\theta^*} + \epsilon Q$, where Q is an arbitrary distribution and the goal is to estimate the unknown parameter θ^* based on an observed data set $\{z_i\}_{i=1}^n$. We also analyze the behavior of the same algorithms in situations where data are generated from a heavy-tailed distribution. Although the parameter corresponds to the true data-generating distribution, "outliers" are observed in the data set due to random sampling, and the goal is to obtain an estimator with similar high-probability guarantees as in the case of standard parameter estimation techniques for lighter-tailed distributions.

* Corresponding author:

E-mail addresses: ei250@cantab.ac.uk (E. Ioannou), muni.pydi@lamsade.dauphine.fr (M.S. Pydi), pll28@cam.ac.uk (P.-L. Loh).

<https://doi.org/10.1016/j.ecosta.2023.07.003>

2452-3062/© 2023 The Author(s). Published by Elsevier B.V. on behalf of EcoSta Econometrics and Statistics. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

Classical robust statistics (Huber and Ronchetti, 2011) suggests the use of M -estimators, which involve optimizing an appropriate loss function over the space of parameters. More specifically, suppose we wish to estimate the parameter $\theta^* = \arg \min_{\theta \in \Theta} \mathcal{R}(\theta)$, where the risk $\mathcal{R}(\theta) = \mathbb{E}[\mathcal{L}(\theta, (x, y))]$ is the expectation of a loss function. In practice, one uses an empirical risk minimizer $\hat{\theta} \in \arg \min_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(\theta, (x_i, y_i))$. Standard theory of parametric statistics shows that the optimal choice of \mathcal{L} corresponds to the log-likelihood function when data are not contaminated. However, taking into account ϵ -contamination leads to the use of other losses such as the Huber loss, which can be shown to be optimal in a minimax sense when the uncontaminated data are normally distributed (Huber and Ronchetti, 2011). Similarly, while the least-squares loss corresponds to maximum likelihood for Gaussian errors, minimizing a different loss function may be advantageous in the case of heavy-tailed data.

In this paper, we adopt an alternative approach inspired by optimization methods (Boyd and Vandenberghe, 2004). Rather than seeking to design a robust loss, we introduce robustness into the estimation algorithm by implementing robust updates in an iterative second-order optimization procedure. Our work is directly inspired by the work of Prasad et al. (2020), who proposed and analyzed a first-order version of this method. Our algorithm, which we call “robust Newton’s method,” utilizes the *AgnosticMean* algorithm from Lai et al. (2016) in the Huber contamination setting to obtain robust gradient and Hessian estimates on each iterate of our algorithm. Moreover, given appropriate assumptions, we prove that the rate of convergence of this algorithm is faster than that of robust gradient descent, and successive iterates converge quadratically to a small ball around θ^* . Furthermore, a suitable stepsize may be determined adaptively using a robust variant of backtracking linesearch. Our analysis of the Newton iterates is fairly general, and can be used to derive convergence guarantees when alternative procedures are employed for gradient/Hessian estimation. We consequently propose a method based on the conjugate gradient method (Wright and Nocedal, 1999) for obtaining approximate Newton directions which may be useful in higher dimensions, and discuss some conjectures about the corresponding convergence rate on ϵ -contaminated data.

1.1. Related Work

Here, we discuss several other general approaches to robust empirical risk minimization (ERM) which have appeared in the literature. A variety of algorithms have been proposed based on median-of-means estimators, which give robust alternatives to mean estimators (a more detailed description is provided in Section 2.4 below). Median-of-mean tournaments (Lugosi and Mendelson, 2019c, 2019b, 2019a) provide a method for comparing pairs of candidate regression functions based on the number of blocks in which the empirical mean of the loss function is smaller for one function than the other. The final estimator is a function which “wins” the most pairwise matches among other candidate functions. Another use of median-of-means estimators derives an estimator by considering a “minimaximization” problem formed by increments of the objective function, where a median-of-means estimate is used in place of the expectation appearing in the population-level version of the problem (Lecué and Lerasle, 2019, 2020; Chinot et al., 2020). Finally, and more similar in spirit to the approach taken in our paper, we mention a method which involves modifying gradient descent by computing a gradient with respect to a median block on each iteration (Lecué et al., 2020). The median block is defined as the block with the smallest empirical mean (with respect to the objective function value) on the current iteration. Excess risk bounds are then derived for a class of binary classification problems, where a certain fraction of the data consists of arbitrarily generated outliers and the remaining points are drawn i.i.d. from the uncontaminated model.

The SEVER algorithm (Diakonikolas et al., 2019b) also operates via an appropriate modification of an iterative optimization procedure. It uses any “approximate learner” algorithm, which can find an approximate critical point of an empirical risk minimization problem, as a subroutine (e.g., gradient descent, stochastic gradient descent, or Newton’s method). On successive iterations, the SEVER algorithm filters out data points by applying the approximate learner to the currently remaining set of data points and then filtering out any points with outlying gradients computed at the parameter chosen by the approximate learner. Statistical error bounds are derived for the output of the SEVER algorithm on classification and regression problems, where data are drawn from a possibly heavy-tailed model and then corrupted by a small fraction of adversarial outliers.

Finally, our work is most closely related to the work of Prasad et al. (2020), which may be seen as a first-order version of our second-order algorithm. In that paper, the authors propose to perform parameter estimation by running a variant of gradient descent on the empirical risk objective, where successive gradients are computed by treating each gradient computation as an approximation of a population-level mean, and then applying a robust mean estimation procedure for multivariate data. As in our work, they use the mean estimation algorithm by Lai et al. (2016) for their multivariate estimation procedure in the case of Huber’s ϵ -contamination model. They also derive statistical error bounds for successive iterates, which hold with high probability. The main difference with our work is that we are able to derive faster rates of convergence due to the use of second-order algorithms, while enjoying the broad applicability of their approach.

1.2. Outline

The remainder of our paper is organized as follows: In Section 2, we discuss the setup of the problem we are aiming to solve. In Section 3, we introduce our novel robust Newton’s method and present two theorems concerning its convergence. In Section 4, we discuss applications of our general theory to generalized linear models. In Section 5, we provide some

illustrative numerical results and comparisons. In [Section 6](#), we present a version of robust Newton's method based on the conjugate gradient method and provide some conjectures. Finally, we conclude our paper with a discussion of open directions in [Section 7](#).

1.3. Notation

For a matrix $A \in \mathbb{R}^{p \times p}$, we use $\|A\|_2$ to denote the spectral norm, $\lambda_{\min}(A)$ to denote the minimum eigenvalue, and $\text{tr}(A)$ to denote the trace. We use $c, C, c_1, C_1, c_2, C_2, \dots$ to denote universal positive constants whose specific values may change from line to line. For functions $f(n)$ and $g(n)$, we write $f(n) = O(g(n))$ to mean that $f(n) \leq Cg(n)$ for some constant $C > 0$, and also write $f(n) \lesssim g(n)$ and $g(n) \gtrsim f(n)$. We write $f(n) \asymp g(n)$ when both inequalities hold simultaneously. We use $\tilde{O}(n)$ to hide logarithmic factors. We use the abbreviation ‘‘w.h.p.’’ for ‘‘with high probability,’’ meaning with probability tending to 1 as the sample size n tends to ∞ .

2. Background

We consider a parametric estimation problem, wherein the data $\{z_i\}_{i=1}^n \subseteq \mathcal{Z}$ sampled from a true distribution P are to be fit to a model with parameter $\theta \in \Theta$. A loss function $\mathcal{L} : \Theta \times \mathcal{Z} \rightarrow \mathbb{R}$ measures the goodness of fit of the model. The optimal parameter $\theta^* \in \Theta$ minimizes the population risk of the model, which is the expected loss incurred by the model over the true data distribution:

$$\theta^* = \underset{\theta \in \Theta}{\operatorname{argmin}} \mathcal{R}(\theta) := \mathbb{E}_{z \sim P}[\mathcal{L}(\theta, z)]. \quad (1)$$

Given n i.i.d. data points $\{z_i\}_{i=1}^n$ sampled from the true distribution P , the goal in empirical risk minimization is to estimate the parameter $\hat{\theta}_n$ that minimizes the empirical risk of the model, which is the average loss incurred by the model over the n data points:

$$\hat{\theta}_n = \underset{\theta \in \Theta}{\operatorname{argmin}} \hat{\mathcal{R}}_n(\theta) := \frac{1}{n} \sum_{i=1}^n \mathcal{L}(\theta, z_i). \quad (2)$$

2.1. Examples

Linear regression: In linear regression, data $z \in \mathcal{Z}$ are of the form $z = (x, y) \in \mathbb{R}^p \times \mathbb{R}$, where the covariate x and response y are related via

$$y = x^T \theta^* + w,$$

where $w \in \mathbb{R}$ is noise that is sampled independently from x and y . The loss function we use for this model is the squared loss function,

$$\mathcal{L}(\theta, (x, y)) = \frac{1}{2} (y - x^T \theta)^2.$$

Generalized linear models: In a generalized linear model (GLM), data $z = (x, y) \in \mathbb{R}^p \times \mathbb{R}$ are sampled from a true distribution P that satisfies the following relation on the conditional probability of y given x :

$$P(y|x) \propto \exp\left(\frac{yx^T \theta^* - \Phi(x^T \theta^*)}{c(\sigma)}\right), \quad (3)$$

where $c(\sigma)$ is the scale parameter and $\Phi : \mathbb{R} \rightarrow \mathbb{R}$ is a convex link function. The loss function we use for a GLM is the negative log-likelihood,

$$\mathcal{L}(\theta, (x, y)) = -yx^T \theta + \Phi(x^T \theta). \quad (4)$$

2.2. Optimization Algorithms

In practice, we seek efficient algorithms for solving the ERM problem (1). A popular algorithm is gradient descent ([Bertsekas, 2015](#)). Given an initial guess for the parameter $\theta_0 \in \Theta$ and a stepsize η , the gradient descent algorithm generates a sequence of iterates $\{\theta_t\}_{t=1}^\infty$, as follows:

$$\theta_{t+1} = \theta_t - \eta \nabla \mathcal{R}(\theta_t).$$

Another popular algorithm is Newton's method ([Wright and Nocedal, 1999](#); [Boyd and Vandenberghe, 2004](#)), whose iterates are given by the following update equation:

$$\theta_{t+1} = \theta_t - (\nabla^2 \mathcal{R}(\theta_t))^{-1} \nabla \mathcal{R}(\theta_t). \quad (5)$$

Whereas gradient descent uses only gradient information at the current iterate θ_t , Newton's method uses both gradient and Hessian information at the current iterate.

2.3. Huber's ϵ -Contamination Model

In Huber's ϵ -contamination model, samples are drawn from a mixture distribution of the form

$$P_\epsilon = (1 - \epsilon)P + \epsilon Q, \quad (6)$$

where P is the true data distribution and Q is an arbitrary noise distribution. The goal is to estimate a parameter $\theta^* \in \Theta$ corresponding to the uncontaminated component P , given n i.i.d. samples drawn from the corrupted distribution P_ϵ .

Huber's contamination model is a classical model studied in robust statistics (Huber and Ronchetti, 2011; Rousseeuw et al., 2011; Maronna et al., 2019), with many exciting theoretical breakthroughs in estimation and inference. More recently, as robust statistics received renewed attention in the theoretical computer science community, additional questions were raised, particularly concerning computational tractability for optimal robust estimators in high dimensions. The contemporaneous work of Lai et al. (2016) and Diakonikolas et al. (2019a) studied computationally tractable mean estimation in multivariate Gaussian settings, where the former paper studied contamination with respect to Huber's model and the latter paper studied a stronger form of "adversarial" contamination. The subroutine which we call Algorithm 1 comes from Lai et al. (2016)—we state it in the slightly adapted version studied in Prasad et al. (2020).

Algorithm 1 Huber Estimator

Require: Samples $S = \{s_i\}_{i=1}^n$, Corruption level ϵ , Dimension p , Failure probability δ

```

1: function HUBERESTIMATOR( $S = \{s_i\}_{i=1}^n, \epsilon, p, \delta$ )
2:   Set  $\tilde{S} = \text{HuberOutlierTruncation}(S, \epsilon, p, \delta)$ 
3:   if  $p = 1$  then
4:     return  $\text{mean}(\tilde{S})$ 
5:   else
6:     Compute  $\Sigma_{\tilde{S}}$ , the covariance matrix of  $\tilde{S}$ 
7:     Compute  $V$ , the span of the top  $p/2$  principal components of  $\Sigma_{\tilde{S}}$ , and  $W$ , its complement
8:     Set  $S_1 := P_V(\tilde{S})$ , where  $P_V$  is the projection operation onto  $V$ 
9:     Set  $\hat{\mu}_V := \text{uberEstimator}(S_1, \epsilon, p/2, \delta)$ 
10:    Set  $\hat{\mu}_W := \text{mean}(P_W \tilde{S})$ 
11:    Set  $\hat{\mu} \in \mathbb{R}^p$  such that  $P_V(\hat{\mu}) = \hat{\mu}_V$  and  $P_W(\hat{\mu}) = \hat{\mu}_W$ 
12:    return  $\hat{\mu}$ 
13:   end if
14: end function
15: function HUBEROUTLIERGRADIENTTRUNCATION( $S, \epsilon, p, \delta$ )
16:   if  $p = 1$  then
17:     Let  $[a, b]$  be the smallest interval containing an  $\left(1 - \epsilon - C\sqrt{\frac{\log(|S|/\delta)}{|S|}}\right)(1 - \epsilon)$  fraction of points
18:      $\tilde{S} \leftarrow S \cap [a, b]$ 
19:     return  $\tilde{S}$ 
20:   else
21:     Let  $[S]_i$  be the samples with the  $i^{\text{th}}$  coordinates only,  $[S]_i = \{\langle x, e_i \rangle | x \in S\}$ 
22:     for  $i = 1$  to  $p$  do
23:        $a[i] = \text{HuberEstimator}([S]_i, \epsilon, 1, \delta/p)$ 
24:     end for
25:     Let  $B(r, a)$  be the ball of smallest radius centered at  $a$  containing an  $\left(1 - \epsilon - C_p \left(\sqrt{\frac{p}{|S|} \log\left(\frac{|S|}{p\delta}\right)}\right)\right)(1 - \epsilon)$  frac-
tion of points in  $S$ 
26:      $\tilde{S} \leftarrow S \cap B(r, a)$ 
27:     return  $\tilde{S}$ 
28:   end if
29: end function

```

In terms of computational complexity, the initial screening step is coordinate-wise, hence $O(p)$. The dominant computation is to perform PCA (which has $O(p^3)$ complexity (Pan and Chen, 1999; Johnstone and Lu, 2009) $\log_2(p)$ times. Thus, the overall runtime of Algorithm 1 is $\tilde{O}(p^3)$.

2.4. Heavy-Tailed Model

In the heavy-tailed model, we assume that data are drawn i.i.d. from a distribution with some number of finite moments. Note that the heavy-tailed model does not involve a contaminating distribution Q . However, the i.i.d. data may still appear to have “outlier” points due to random sampling.

A popular approach for heavy-tailed mean estimation in the probably approximately correct (PAC) framework—obtaining high-probability deviation bounds which are as tight as possible under minimal distributional assumptions—is to use a median-of-means (MOM) estimator. Roughly speaking, data are randomly partitioned into blocks, the mean of each block is computed, and the median of all of the block means is returned as the estimator. In multiple dimensions, different notions of medians exist, leading to different flavors of MOM estimators. For a more detailed overview, see the survey (Lugosi and Mendelson, 2019a) and the references cited therein. The MOM algorithm is summarized in Algorithm 2. In particular, we will employ a version of the algorithm from Minsker (2015), which combines the mean estimates using the geometric median, i.e., the point which minimizes the sum of ℓ_2 -distances to the block means.

Algorithm 2 Heavy-Tailed Estimator

Require: Samples $S = \{s_i\}_{i=1}^n$, Failure probability δ

- 1: **function** HEAVYTAILED ESTIMATOR($S = \{s_i\}_{i=1}^n, \delta$)
 - 2: Set $b = 1 + \lfloor 3.5 \log 1/\delta \rfloor$, the number of buckets
 - 3: Partition S into b blocks B_1, \dots, B_b , each of size $\lfloor n/b \rfloor$
 - 4: **for** $i = 1, \dots, b$ **do**
 - 5: $\hat{\mu}_i = \frac{1}{|B_i|} \sum_{s \in B_i} s$
 - 6: **end for**
 - 7: Set $\hat{\mu} = \operatorname{argmin}_{\mu} \sum_{i=1}^b \|\mu - \hat{\mu}_i\|_2$ **return** $\hat{\mu}$
 - 8: **end function**
-

Since the runtime of the geometric median computation on n data points in p dimensions is $\tilde{O}(np)$ (Cohen et al., 2016), the runtime of Algorithm 2 is $\tilde{O}(n + bp) = \tilde{O}(n + p)$.

3. Robust Newton’s Method

We now present our variant of robust Newton’s method. At each iterate, we will use gradient and Hessian estimates ($g(\theta), H(\theta)$) in place of $(\nabla \mathcal{R}(\theta), \nabla^2 \mathcal{R}(\theta))$ in the update Eq. (5). We assume that these estimates satisfy the conditions described in the following definitions:

Definition 1 Prasad et al. (2020). A function $g(\theta)$ is a *robust gradient estimator* for a data set $S = \{z_i\}_{i=1}^n$ if for functions α_g and β_g , with probability at least $1 - \delta$, at any fixed $\theta \in \Theta$, the estimator satisfies the following inequality:

$$\|g(\theta) - \nabla \mathcal{R}(\theta)\|_2 \leq \alpha_g(n, \delta) \|\theta - \theta^*\|_2 + \beta_g(n, \delta). \quad (7)$$

Definition 2. A function $H(\theta)$ is a *robust Hessian estimator* for a data set $S = \{z_i\}_{i=1}^n$ if for functions α_h and β_h , with probability at least $1 - \delta$, at any fixed $\theta \in \Theta$, the estimator satisfies the following inequality:

$$\|H(\theta) - \nabla^2 \mathcal{R}(\theta)\|_2 \leq \alpha_h(n, \delta) \|\theta - \theta^*\|_2 + \beta_h(n, \delta). \quad (8)$$

Successive iterates then take the form

$$\theta_{t+1} = \theta_t - \alpha_t H(\theta_t)^{-1} g(\theta_t),$$

where α_t is chosen via a version of backtracking linesearch (Boyd and Vandenberghe, 2004). The exit condition of backtracking linesearch differs from its non-robust version in that function evaluations are replaced by robust estimates (cf. Lemmas 7 and 8 below) and an extra tolerance parameter ζ is included. The full algorithm is provided in Algorithm 3.

Remark 1. The statements of Definitions 1 and 2 are written quite generally; in Section 3.2 below, we provide algorithms for obtaining robust gradient and Hessian estimators under both of our contamination models which can help elucidate the form of the bounds. See also Propositions 3 and 4 in Appendix C and Appendix D, which provide explicit values of the parameters $(\alpha_g, \beta_g, \alpha_h, \beta_h)$ that are suitable for GLMs.

Algorithm 3 Robust Newton's Method

Require: Data samples $S = \{z_i\}_{i=1}^n$, Number of iterations T , Initial guess $\theta_0 \in \Theta$, Backtracking linesearch parameters $\kappa_1 \in (0, 0.5)$, $\kappa_2 \in (0, 1)$, and ζ

```

1: function ROBUSTNEWTON( $S, \theta_0, \kappa_1, \kappa_2, \zeta$ )
2:   for  $t = 0$  to  $T - 1$  do
3:     Compute losses  $\{\mathcal{L}(\theta_t, z_i)\}_{i=1}^n$  and gradients  $\{\nabla \mathcal{L}(\theta_t, z_i)\}_{i=1}^n$ 
4:     Compute gradient estimate  $g(\theta_t) = \text{RobustGradientEstimate}(S, \theta_t)$ 
5:     Compute Hessian estimate  $H(\theta_t) = \text{RobustHessianEstimate}(S, \theta_t)$ 
6:     Compute Newton step  $\Delta\theta_{nt} = -H(\theta_t)^{-1}g(\theta_t)$ 
7:     Compute stepsize  $\alpha = \text{BacktrackingLineSearch}(S, \theta_t, \Delta\theta_{nt}, g(\theta_t), \kappa_1, \kappa_2, \zeta)$ 
8:     Update  $\theta_{t+1} = \theta_t + \alpha \Delta\theta_{nt}$ 
9:   end for
10:  return  $\theta_T$ 
11: end function
12: function BACKTRACKINGLINESEARCH( $S, \theta, \Delta\theta_{nt}, g(\theta), \kappa_1, \kappa_2, \zeta$ )
13:  Set  $\alpha = 1$ 
14:  while  $\text{RobustEstimate}(\{\mathcal{L}(\theta + \alpha \Delta\theta_{nt}, z_i)\}_{i=1}^n) > \text{RobustEstimate}(\{\mathcal{L}(\theta, z_i)\}_{i=1}^n) + \kappa_1 \alpha g(\theta) \Delta\theta_{nt} + \zeta$  do
15:    Update  $\alpha = \kappa_2 \alpha$ 
16:  end while
17:  return  $\alpha$ 
18: end function

```

3.1. General Analysis for Robust Newton's Method

For the results of this section, we assume that $f(\theta) := \mathcal{R}(\theta)$ is twice-differentiable and satisfies the Lipschitz condition $\|\nabla^2 f(\theta_1) - \nabla^2 f(\theta_2)\|_2 \leq L \|\theta_1 - \theta_2\|_2$, for all θ_1, θ_2 . We also assume that f satisfies the strong convexity and smoothness conditions $mI \preceq \nabla^2 f(\theta) \preceq MI$, for all θ close enough to the initialization θ_0 . (We will verify these conditions for GLMs in [Propositions 1](#) and [2](#) below.) Finally, we will assume that at each iterate, the gradient and Hessian estimates $g(\theta)$ and $H(\theta)$ satisfy inequalities [\(7\)](#) and [\(8\)](#), respectively. As demonstrated in [Theorems 3](#) and [4](#) later, the last condition can typically be justified w.h.p. via a union bound. Observe that in this setting, the unique global minimum of f is the true parameter θ^* .

The first result shows that if $\|\nabla f(\theta_0)\|_2$ is sufficiently small, the backtracking linesearch procedure will always choose stepsize 1. (This is known as the “pure Newton” phase.) Furthermore, successive iterates converge at a geometric rate to a small ball around θ^* . Recall that the parameters (κ_1, κ_2) of backtracking linesearch are defined as in [Algorithm 3](#).

Theorem 1. Suppose $\|\nabla f(\theta_0)\|_2 < \eta$, where

$$\eta := \frac{m^2}{8L} \cdot \min\{3(1 - 2\kappa_1), 2\}. \quad (9)$$

Suppose the gradient and Hessian errors satisfy the bounds

$$\gamma_g := \frac{2\eta\alpha_g}{m} + \beta_g \leq \eta, \quad \text{and} \quad \gamma_h := \frac{2\eta\alpha_h}{m} + \beta_h \leq \frac{m}{2}. \quad (10)$$

Also suppose the robust estimates satisfy

$$\left| \text{RobustEstimate}(\{\mathcal{L}(\theta_t + \alpha \Delta\theta_t, z_i)\}_{i=1}^n) - f(\theta_t + \alpha \Delta\theta_t) \right| \leq \frac{\zeta}{4}, \quad (11)$$

for each evaluation of backtracking linesearch, where we set the linesearch parameter to be

$$\zeta \geq \frac{8\gamma_g\eta}{m} + \frac{16\gamma_h\eta^2}{m^2}. \quad (12)$$

Then backtracking linesearch chooses unit steps on all successive iterates, and we have $\|\nabla f(\theta_t)\|_2 < \eta$ and

$$\|\theta_t - \theta^*\|_2 \leq \frac{m}{L} \left(\frac{1}{2}\right)^{2t} + \frac{6c_2}{m} \quad (13)$$

for all $t \geq 1$, where

$$c_2 = \eta \left(\frac{4\gamma_g L}{m^2} + \frac{2\gamma_h}{m} \right) + \frac{2L\gamma_g^2}{m^2} + \gamma_g + \frac{2\gamma_g\gamma_h}{m},$$

and (γ_g, γ_h) are small enough so that

$$c_2 \leq \min \left\{ \frac{\eta}{2}, \frac{m^2}{24L} \right\}. \quad (14)$$

The proof of [Theorem 1](#) is found in [A.1](#).

Next, we show that after a finite number of steps, the iterates will indeed satisfy $\|\nabla f(\theta_t)\|_2 < \eta$, for an appropriate η . We can then apply [Theorem 1](#) to the first iterate satisfying this condition, relabeling it as θ_0 , to obtain estimation error bounds on the overall trajectory of robust Newton's method. The proof of the following result is provided in [A.2](#).

Theorem 2. *Suppose the parameters η and ζ are as defined in [Theorem 1](#). Define*

$$\gamma := \kappa_1 \kappa_2 \frac{m}{M} \left(\frac{1}{2} - \kappa_1 \right) \frac{\eta^2}{4\sqrt{2M}}.$$

Suppose (γ_g, γ_h) are chosen small enough such that $\zeta \leq \frac{\gamma}{2}$ and conditions [\(10\)](#) and [\(11\)](#) are satisfied. Also suppose

$$\frac{2\alpha_g}{m} \sqrt{2M(f(\theta_0) - f(\theta^*))} + \beta_g \leq \min \left\{ \frac{\eta}{2}, \sqrt{\frac{m}{2}} \cdot \frac{1}{2} \sqrt{\frac{\eta^2}{4\sqrt{2M}}} \right\}, \quad (15)$$

and

$$\frac{2\alpha_h}{m} \sqrt{2M(f(\theta_0) - f(\theta^*))} + \beta_h \leq M. \quad (16)$$

If $\|\nabla f(\theta_t)\|_2 \geq \eta$ for an iterate $t \geq 0$, then $f(\theta_t) \leq f(\theta_0)$ and $f(\theta_{t+1}) - f(\theta_t) < -\frac{\gamma}{2}$.

The preceding theorem directly implies that after a finite number of steps (known as the ‘‘damped Newton’’ phase), all successive iterates of the algorithm satisfy $\|\nabla f(\theta_t)\|_2 < \eta$. Indeed, [Theorem 2](#) guarantees that $f(\theta_{t+1}) - f(\theta_t) < -\frac{\gamma}{2}$ whenever $\|\nabla f(\theta_t)\|_2 \geq \eta$, where $\gamma > 0$. Since $\theta^* \in \arg \min_{\theta} f(\theta)$, we clearly could not have $\|\nabla f(\theta_t)\|_2 \geq \eta$ for all $0 \leq t \leq T$, where $T = \lceil \frac{f(\theta_0) - f(\theta^*)}{\gamma/2} \rceil$, or else

$$f(\theta_T) - f(\theta^*) = (f(\theta_0) - f(\theta^*)) + \sum_{t=0}^{T-1} (f(\theta_{t+1}) - f(\theta_t)) < (f(\theta_0) - f(\theta^*)) - \frac{T\gamma}{2} < 0,$$

contradicting the minimality of θ^* .

3.2. Robust Estimation of Gradients and Hessians

In this subsection, we explain how robust estimators for gradients and Hessians can be obtained under two models of contamination, namely the Huber ϵ -contamination model and the heavy-tailed model.

3.2.1. Robust Gradient Estimation

For the ϵ -contamination model, we obtain a robust gradient estimate by applying [Algorithm 1](#) to the gradients computed on each of the n sampled data points. Similarly, for the heavy-tailed model, we use [Algorithm 2](#) to obtain a robust gradient estimate. For completeness, we summarize this procedure in [Algorithm 4](#).

Algorithm 4 Robust Gradient Estimator

Require: Samples $S = \{z_i\}_{i=1}^n$, Parameter θ , Contamination type Type

Require: (If Type = Huber) Corruption Level ϵ , Dimension p , Failure probability δ

Require: (If Type = Heavy-tail) Failure probability δ

- 1: **function** ROBUSTGRADIENTESTIMATOR($S, \theta, \text{TYPE}, \epsilon, p, \delta$)
 - 2: Compute $\{\nabla \mathcal{L}(\theta, z_i)\}_{i=1}^n$, the gradient of the loss at each data point in S
 - 3: **if** Type = Huber **then**
 - 4: **return** HUBERESTIMATOR($\{\nabla \mathcal{L}(\theta, z_i)\}_{i=1}^n, \epsilon, p, \delta$)
 - 5: **end if**
 - 6: **if** Type = Heavy-tail **then**
 - 7: **return** HEAVYTAILEDESTIMATOR($\{\nabla \mathcal{L}(\theta, z_i)\}_{i=1}^n, \delta$)
 - 8: **end if**
 - 9: **end function**
-

The following lemmas, borrowed from [Prasad et al. \(2020\)](#), show that [Algorithm 4](#) returns a robust gradient estimator that satisfies [Definition 1](#).

Lemma 1 (Lemma 1 of Prasad et al. (2020)). Let $\{z_i\}_{i=1}^n$ be n i.i.d. samples drawn from a Huber ϵ -contaminated distribution (6). Let the true distribution of gradients $\nabla\mathcal{L}(\theta, z)$, with z drawn from P , have bounded fourth moments. Then Algorithm 4 with $S = \{z_i\}_{i=1}^n$, Type = Huber, and any $\theta \in \Theta$ returns a gradient estimate $g(\theta)$ that satisfies

$$\|g(\theta) - \mathbb{E}[\nabla\mathcal{L}(\theta, z)]\|_2 \leq C_1(\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon))\sqrt{\|\text{Cov}(\nabla\mathcal{L}(\theta, z))\|_2 \log p}, \quad (17)$$

with probability at least $1 - \delta$, where $C_1 > 0$ is a constant and γ is given by

$$\gamma(n, p, \delta, \epsilon) = \left(\frac{p \log(p) \log(n/(p\delta))}{n}\right)^{3/8} + \left(\frac{\epsilon p^2 \log(p) \log(p \log(p)/\delta)}{n}\right)^{1/4}. \quad (18)$$

Lemma 2 (Lemma 2 of Prasad et al. (2020)). Let $\{z_i\}_{i=1}^n$ be n i.i.d. samples drawn from a heavy-tailed distribution P such that the true distribution of gradients $\nabla\mathcal{L}(\theta, z)$ has bounded second moments. Then Algorithm 4 with $S = \{z_i\}_{i=1}^n$, Type = Heavy-tail, and any $\theta \in \Theta$ returns a gradient estimate $g(\theta)$ that satisfies

$$\|g(\theta) - \mathbb{E}[\nabla\mathcal{L}(\theta, z)]\|_2 \leq 11\sqrt{\frac{\text{tr}(\text{Cov}(\nabla\mathcal{L}(\theta, z))) \log(1.4/\delta)}{n}}, \quad (19)$$

with probability at least $1 - \delta$.

3.2.2. Robust Hessian Estimation: The Vectorizing Approach

The procedure for obtaining a robust Hessian, summarized in Algorithm 5, is similar to that of Algorithm 4, except that

Algorithm 5 Robust Hessian Estimator

Require: Samples $S = \{z_i\}_{i=1}^n$, Parameter θ , Contamination type Type

Require: (If Type = Huber) Corruption Level ϵ , Dimension p , Failure probability δ

Require: (If Type = Heavy-tail) Failure probability δ

- 1: **function** ROBUSTHESSIANESTIMATOR($S, \theta, \text{TYPE}, \epsilon, p, \delta$)
 - 2: Compute $\{\nabla^2\mathcal{L}(\theta, z_i)\}_{i=1}^n$, the Hessian of the loss at each data point in S
 - 3: **if** Type = Huber **then**
 - 4: **return** unflatten(HUBERESTIMATOR($\{\text{flatten}(\nabla^2\mathcal{L}(\theta, z_i))\}_{i=1}^n, \epsilon, p, \delta$))
 - 5: **if** Type = Heavy-tail **then**
 - 6: **return** unflatten(HEAVYTAILEDESTIMATOR($\{\text{flatten}(\nabla^2\mathcal{L}(\theta, z_i))\}_{i=1}^n, \delta$))
 - 7: **end if**
 - 8: **end function**
-

the appropriate multivariate estimation procedure is applied to a vectorized version of the Hessian matrix (where we use $\text{flatten}(A)$ to denote a vectorized version of the matrix A , and use $\text{unflatten}()$ to denote the inverse function).

The next two lemmas follow immediately from the arguments used to derive Lemmas 1 and 2:

Lemma 3. Let $\{z_i\}_{i=1}^n$ be n i.i.d. samples drawn from a Huber ϵ -contaminated distribution (6). Suppose $\text{Cov}(\text{flatten}(\nabla^2\mathcal{L}(\theta, z)))$ is finite and $\text{flatten}(\nabla^2\mathcal{L}(\theta, z))$ has bounded fourth moments. Then Algorithm 4 with $S = \{z_i\}_{i=1}^n$, Type = Huber, and any $\theta \in \Theta$ returns a Hessian estimate $H(\theta)$ that satisfies

$$\|H(\theta) - \mathbb{E}[\nabla^2\mathcal{L}(\theta, z)]\|_2 \leq C_2(\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon))\sqrt{\|\text{Cov}(\text{flatten}(\nabla^2\mathcal{L}(\theta, z)))\|_2 \log p}, \quad (20)$$

with probability at least $1 - \delta$, where $C_2 > 0$ is a constant and γ is given by equation (18).

Lemma 4. Let $\{z_i\}_{i=1}^n$ be n i.i.d. samples drawn from a heavy-tailed distribution P . Suppose $\text{Cov}(\text{flatten}(\nabla^2\mathcal{L}(\theta, z)))$ is finite. Then Algorithm 4 with $S = \{z_i\}_{i=1}^n$, Type = Heavy-tail, and any $\theta \in \Theta$ returns a Hessian estimate $H(\theta)$ that satisfies

$$\|H(\theta) - \mathbb{E}[\nabla^2\mathcal{L}(\theta, z)]\|_2 \leq C_3\sqrt{\frac{\text{tr}(\text{Cov}(\text{flatten}(\nabla^2\mathcal{L}(\theta, z)))) \log(1.4/\delta)}{n}}, \quad (21)$$

with probability at least $1 - \delta$, where $C_3 > 0$ is a constant.

4. Application to GLMs

In this section, we apply the robust Newton method to parametric estimation in GLMs. We consider the Huber ϵ -contamination model in Section 4.2, and we consider the heavy-tailed contamination model in Section 4.3.

Throughout this section, we will assume that the uncontaminated model is a GLM of the form (3). Consider the loss function in Eq. (4). We assume that the link function Φ of the GLM satisfies the following bounds:

$$\mathbb{E}[|\Phi'(x_i^T\theta) - \Phi'(x_i^T\theta^*)|^{2k}] \leq L_{\Phi, 2k}\|\theta - \theta^*\|_2^{2k} + B_{\Phi, 2k}, \quad \forall \theta \in \Theta, \quad (22)$$

$$\mathbb{E}[|\Phi^{(t)}(x_i^T \theta^*)|^k] \leq M_{\Phi, t, k}, \quad (23)$$

and

$$\|\Phi^{(t)}\|_{\infty} \leq \bar{M}_{\Phi, t}, \quad (24)$$

for pairs (k, t) to be specified in the sequel, where $\Phi^{(t)}$ is the t^{th} derivative of Φ .

We also make assumptions on the boundedness of moments of $x_i \in \mathbb{R}^p$. We say that x_i has bounded $2k^{\text{th}}$ moments if there is a constant $\tilde{C}_{2k} > 0$ such that for every unit vector $v \in \mathbb{R}^d$, we have $\mathbb{E}[(x_i^T v)^{2k}] \leq \tilde{C}_{2k} (\mathbb{E}[x_i^T v]^2)^k$.

Assumption 1. Suppose the distribution of the x_i 's has bounded eighth moments. Let Σ_x denote the finite covariance matrix of the x_i 's.

In order to apply [Theorems 1](#) and [2](#) to GLMs, we need the Hessian $\nabla^2 \mathcal{R}(\theta)$ to be Lipschitz smooth and satisfy $mI \leq \nabla^2 \mathcal{R}(\theta) \leq MI$ for all $\theta \in \Theta$ close enough to the initialization θ_0 . We now verify these assumptions. The following results are proved in [B.2](#) and [B.3](#).

Proposition 1. Let the link function Φ satisfy inequality (24) for $t = 3$, and suppose [Assumption 1](#) is satisfied. Then the Hessian $\nabla^2 \mathcal{R}(\theta)$ is L -Lipschitz and satisfies $\nabla^2 \mathcal{R}(\theta) \leq MI$ with

$$L := \sqrt{\tilde{C}_4 \bar{M}_{\Phi, 3}} \|\Sigma_x\|_2, \quad M := \bar{M}_{\Phi, 2} \sqrt{\tilde{C}_4} \|\Sigma_x\|_2. \quad (25)$$

Proposition 2. Suppose there exist constants $B, \tau > 0$ such that for any $\theta \in \mathbb{R}^p$ such that $\|\theta\|_2 \leq B$, we have

$$\tilde{C}_4 \|\Sigma_x\|_2^2 \cdot \mathbb{P}(|x_i^T \theta| > \tau) \leq \frac{1}{4} \lambda_{\min}^2(\Sigma_x). \quad (26)$$

Define $b_{\tau} := \inf_{|u| \leq \tau} \Phi''(u)$. Then $\frac{b_{\tau}}{2} \lambda_{\min}(\Sigma_x) I \leq \nabla^2 \mathcal{R}(\theta)$ for all $\theta \in \mathbb{R}^p$ such that $\|\theta\|_2 \leq B$.

Remark 2. Note that when the covariates are sub-Gaussian, we can certainly guarantee that the tail condition (26) is satisfied for sufficiently large τ , since $x_i^T \theta$ is sub-Gaussian with parameter scaling with B and the sub-Gaussian parameter σ_x^2 of x_i . Thus, we have

$$\mathbb{P}(|x_i^T \theta| > \tau) \leq c_1 \exp\left(-\frac{c_2 \tau^2}{B^2 \sigma_x^2}\right),$$

and it suffices to take $\tau = c_3 B \sigma_x \log^{1/2}\left(\frac{c_4 \|\Sigma_x\|_2^2}{\lambda_{\min}^2(\Sigma_x)}\right)$. Furthermore, in the proofs of [Theorems 1](#) and [2](#) (cf. inequalities (A.2) and (A.13), respectively), we show that $\|\theta_t - \theta^*\|_2$ remains bounded (where the bound depends on θ_0 and the problem parameters).

In the case of logistic regression, we have $\Phi''(u) = \frac{e^u}{(1+e^u)^2}$, and it is easy to see that $b_{\tau} > 0$ for any value of τ .

4.1. Preliminary Error Bounds

From [Lemmas 1](#) and [2](#), we see that the term $\text{Cov}(\nabla \mathcal{L}(\theta, z))$ plays a crucial role in proving that our gradient estimates are robust. Likewise, [Lemmas 3](#) and [4](#) show the importance of the term $\text{Cov}(\text{flatten}(\nabla^2 \mathcal{L}(\theta, z)))$ in proving that the Hessian estimates are robust. The following two lemmas provide upper bounds on these two terms for the specific case of GLMs:

Lemma 5 (Lemma 4 in [Prasad et al. \(2020\)](#)). Let $\{z_i\}_{i=1}^n$ be n i.i.d. samples drawn from a distribution that satisfies the GLM model (3). Let the link function Φ satisfy inequalities (22) and (23) for $k \in \{1, 2\}$ and $t \in \{2, 4\}$, and suppose [Assumption 1](#) is satisfied. Then the true distribution of gradients $\nabla \mathcal{L}(\theta, z)$ has bounded fourth moments. Moreover,

$$\begin{aligned} \|\text{Cov}(\nabla \mathcal{L}(\theta, z))\|_2 &\leq C_1 \|\Sigma_x\|_2 \left(\sqrt{L_{\Phi, 4}} + L_{\Phi, 2} \right) \|\theta - \theta^*\|_2^2 \\ &\quad + C_2 \|\Sigma_x\|_2 \left(B_{\Phi, 2} + \sqrt{B_{\Phi, 4}} + c(\sigma) \sqrt{M_{\Phi, 2, 2}} + \sqrt{c(\sigma)^3 M_{\Phi, 4, 1}} \right), \end{aligned} \quad (27)$$

where $C_1, C_2 > 0$ are constants.

Lemma 6. Let $\{z_i\}_{i=1}^n$ be n i.i.d. samples drawn from a distribution that satisfies the GLM model (3). Let the link function Φ satisfy inequality (24) for $t = 2$, and suppose [Assumption 1](#) is satisfied. Then the distribution of the flattened Hessian $\text{flatten}(\nabla^2 \mathcal{L}(\theta, z))$ has bounded fourth moments. Moreover, we have

$$\text{tr}(\text{Cov}(\text{flatten}(\nabla^2 \mathcal{L}(\theta, z)))) \leq \bar{M}_{\Phi, 2}^2 \tilde{C}_4 p^2 \|\Sigma_x\|_2^2. \quad (28)$$

The proof of [Lemma 6](#) is contained in [B.1](#).

Remark 3. Note that under additional assumptions (e.g., 4-wise independence of the components of the x_i 's), we can prove that

$$\|\text{Cov}(\text{flatten}(\nabla^2 \mathcal{L}(\theta, z)))\|_2 \leq C\tilde{C}_4 \|\Sigma_x\|_2^2,$$

for some constant $C > 0$, which avoids an extra dimension-dependent factor in comparison to inequality (28) (cf. Proposition 4.2 in Lai et al. (2016)) for the Huber contamination setting. Indeed, only the spectral norm of the covariance of the flattened Hessian appears in the deviation bound of Lemma 3 (Huber's ϵ -contamination model); the trace of the covariance appears in Lemma 4 (heavy-tailed model).

For applying Theorems 1 and 2, we also need the robust estimate of the losses to be close to the population risk, as in inequality (11). In the following two lemmas, we show that this assumption holds with high probability for the robust estimates obtained by applying Algorithms 1 and 2 on the losses. Further note that the following lemmas require boundedness of higher-order moments of $\mathcal{L}(\theta, z)$, which can be justified in our scenario if θ is bounded. As mentioned in Remark 2, we can indeed assume that the iterates $\{\theta_t\}$, to which Lemmas 7 and 8 are applied in the sequel, are bounded.

The following result is a consequence of Lemma 14 in Prasad et al. (2020):

Lemma 7. Let $\{z_i\}_{i=1}^n$ be n i.i.d. samples drawn from a Huber ϵ -contaminated distribution (6), where the true distribution satisfies the GLM model (3). Let $\mathcal{L}(\theta, z)$ have bounded fourth moments. Then with probability at least $1 - \delta$, the robust estimate returned by Algorithm 1 satisfies

$$|\text{HuberEstimate}(\{\mathcal{L}(\theta, z_i)\}_{i=1}^n) - \mathcal{R}(\theta)| \leq C_1 \left(\epsilon + \sqrt{\frac{\log(n/\delta)}{n}} \right)^{\frac{3}{4}} + C_2 \left(\epsilon + \sqrt{\frac{\log(n/\delta)}{n}} \right)^{\frac{1}{2}} \frac{\log(1/\delta)}{n},$$

where $C_1, C_2 > 0$ are constants.

The next result follows from similar arguments to those in Lemma 2:

Lemma 8. Let $\{z_i\}_{i=1}^n$ be n i.i.d. samples drawn from a heavy-tailed distribution P that satisfies the GLM model (3). Let $\mathcal{L}(\theta, z)$ have bounded second moments. Then with probability at least $1 - \delta$, the robust estimate returned by Algorithm 2 satisfies

$$|\text{HeavyTailedEstimate}(\{\mathcal{L}(\theta, z_i)\}_{i=1}^n) - \mathcal{R}(\theta)| \leq C \sqrt{\frac{\log\left(\frac{1.4}{\delta}\right)}{n}},$$

with probability at least $1 - \delta$, where $C > 0$ is a constant.

4.2. Huber Contamination

Throughout this subsection, we work under the following assumptions:

Assumption 2. Suppose the link function Φ satisfies inequalities (22) and (23) for $k \in \{1, 2, 4\}$ and $t \in \{2, 4\}$, and inequality (24) for $t \in \{2, 3\}$. Also suppose $\mathcal{R}(\theta)$ is m -strongly convex, i.e., $mI \leq \nabla^2 \mathcal{R}(\theta)$ uniformly over θ , and L and M are defined as in Eq. (25).

We then have the following result, proved in C.2:

Theorem 3. Let $\{z_i\}_{i=1}^n$ be n i.i.d. samples drawn from a Huber ϵ -contaminated distribution (6), where the true distribution satisfies the GLM model (3) and the conditions of Assumptions 1 and 2 are satisfied. Define $(\eta, c_2, \gamma_g, \gamma_h)$ as in Theorem 1 and γ as in Theorem 2. Suppose

$$\zeta = C \max \left\{ \frac{\gamma_g \eta}{m} + \frac{\gamma_h \eta^2}{m^2}, \left(\epsilon + \sqrt{\frac{\log(n/\delta)}{n}} \right)^{\frac{3}{4}} + \left(\epsilon + \sqrt{\frac{\log(n/\delta)}{n}} \right)^{\frac{1}{2}} \frac{\log(1/\delta)}{n} \right\} \leq \frac{\gamma}{2}. \quad (29)$$

Let $\delta > 0$. Define

$$\hat{\gamma} := \min \left\{ \frac{\hat{\alpha}_g}{c_1 \sqrt{\|\Sigma_x\|_2 \log p}}, \frac{\hat{\beta}_g}{c_2 \sqrt{\|\Sigma_x\|_2 \log p}}, \frac{\hat{\beta}_h}{c_3 \|\Sigma_x\|_2 p \sqrt{\log p}} \right\}, \quad (30)$$

where $(\hat{\alpha}_g, \hat{\beta}_h, \hat{\beta}_g)$ are as defined in Lemma 9. Suppose n and ϵ are such that

$$\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon) < \hat{\gamma}. \quad (31)$$

Then applying Algorithm 3 on $\{z_i\}_{i=1}^n$ with initialization $\theta_0 \in \Theta$ and number of iterations

$$T \geq \frac{\mathcal{R}(\theta_0) - \mathcal{R}(\theta^*)}{\gamma/2} + \log_2 \log_2 \left(\frac{6c_2 L}{m^2} \right)$$

returns an output such that

$$\|\theta_T - \theta^*\|_2 \leq \frac{12c_2}{m} = O\left(p\sqrt{\epsilon \log p}\right),$$

with probability at least $1 - T\delta \left(5 + \left\lceil \frac{\log\left(\frac{m}{M}\left(\frac{1}{2} - \kappa_1\right)\right)}{\log \kappa_2} \right\rceil\right)$.

Remark 4. Examining the condition (31), we see that (assuming $(m, M, L, \eta, \|\Sigma_x\|_2)$ are all constants) we have a required upper bound of $\hat{\gamma}^2 \asymp \frac{1}{\log p}$ on the contamination proportion ϵ . Furthermore, from the expression (18), we have (ignoring log factors) $n \gtrsim \max\{p, \epsilon p^2\}$. The condition (29) likewise gives a minimum sample size requirement on n in terms of δ .

Remark 5. It is instructive to compare the result of Theorem 3 to Theorem 4 in Prasad et al. (2020), which gives a convergence statement of the form

$$\|\theta^t - \theta^*\|_2 \leq \kappa^t \|\theta^0 - \theta^*\|_2 + \frac{C\|\Sigma_x\|_2^{1/2}\sqrt{\log p}}{1 - \kappa} (\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon))$$

for iterates $\{\theta^t\}_{t \geq 0}$ of robust gradient descent. For sufficiently large t , the second term dominates, leaving an error term of $O(\sqrt{\epsilon \log p})$. Our theorem has a dominant factor of $O(p\sqrt{\epsilon \log p})$, which can be reduced to $O(\sqrt{\epsilon \log p})$ if we assume 4-wise independence of the coordinates of the covariate distribution (cf. Remark 3 above). In terms of the convergence rate of the optimization procedure, however, we just need $T \asymp \log \log \frac{1}{\epsilon}$, compared to $T \asymp \log \frac{1}{\epsilon}$ in the case of robust gradient descent.

Linear regression is of course a special case of GLMs, for which Theorem 3 readily applies. On the other hand, note that a much more direct way to obtain a robust estimator for linear regression would be to directly robustify the estimator $\hat{\theta}_{OLS} = \left(\frac{X^T X}{n}\right)^{-1} \left(\frac{X^T y}{n}\right)$, where we apply Algorithm 1 to obtain robust estimates of $\mathbb{E}[y_i x_i]$ and $\mathbb{E}[x_i x_i^T]$ (the latter matrix being vectorized before applying the agnostic mean algorithm). Indeed, in the non-robust case, applying Newton's method to the ordinary least squares objective converges in a single step. A careful analysis of this so-called "robust plug-in estimator" would also give an error of $O(\sqrt{\epsilon})$ in the robust case, but a direct analysis would provide an error bound which depends on $\|\theta^*\|_2$, since $\text{Cov}(x_i, y_i)$ would scale with $\|\theta^*\|_2$ (cf. Corollary 3 in Prasad et al. (2020)). On the other hand, the guarantee of Theorem 3 for the full robust Newton's method does not involve $\|\theta^*\|_2$.

Remark 6. A natural question is whether the estimation error upper bounds in Theorem 3 are tight: For i.i.d. samples from a GLM with Huber ϵ -contamination, is it possible to derive estimators with error smaller than $C\sqrt{\epsilon}$? For the case of linear regression, this problem has been studied quite carefully, and it has been established that when the uncontaminated data are Gaussian with an isotropic covariance, the rate should be $\Theta(\epsilon)$ (Chen et al., 2016; Diakonikolas et al., 2019c; Pensia et al., 2020; Depersin, 2020), with no dependence at all on p . In the case when the covariates only follow a bounded fourth moment assumption, the rate improves to $\Theta(\sqrt{\epsilon})$ (Cherapanamjeri et al., 2020a,b). We are not aware of existing lower bounds in the literature for more general GLMs, though it is reasonable to conjecture that the optimal rates for estimation in the Huber contamination model can also be made dimension-independent.

In terms of computational complexity, the overall complexity of the robust Newton method is the number of iterations T multiplied by the computational complexity of robust gradient/Hessian computations. As mentioned at the end of Section 2.3, the runtime of Algorithm 1 is $\tilde{O}(p^3)$; since we would be applying this to the vectorized Hessian matrices, the computational complexity of the robust Newton method would then be $\tilde{O}(Tp^6)$ (note that T depends on ϵ rather than n and p).

4.3. Heavy-Tailed Distributions

Throughout this subsection, we work under the following assumptions:

Assumption 3. Suppose the link function Φ satisfies inequalities (22) and (23) for $k \in \{1, 2\}$ and $t \in \{2, 4\}$, and inequality (24) for $t \in \{2, 3\}$. Also suppose $\mathcal{R}(\theta)$ is m -strongly convex, i.e., $ml \leq \nabla^2 \mathcal{R}(\theta)$ uniformly over θ , and L and M are defined as in Eq. (25).

We then have the following result, proved in D.2:

Theorem 4. Let $\{z_i\}_{i=1}^n$ be n i.i.d. samples drawn from a heavy-tailed distribution P that satisfies the GLM in (3), and suppose the conditions of Assumptions 1 and 3 are satisfied. Define $(\eta, c_2, \zeta, \gamma_g, \gamma_h)$ as in Theorem 1 and γ as in Theorem 2. Let $\delta > 0$. Suppose n satisfies

$$n > C \max \left\{ \frac{p}{\hat{\alpha}_g^2}, \frac{p}{\hat{\beta}_g^2}, \frac{p^2 \|\Sigma_x\|_2^2}{\hat{\beta}_h^2}, \frac{1}{\zeta^2} \right\} \log \left(\frac{1.4}{\delta} \right), \quad (32)$$

where $(\widehat{\alpha}_g, \widehat{\beta}_g, \widehat{\beta}_h)$ are as defined in [Lemma 9](#). Then applying [Algorithm 3](#) on $\{z_i\}_{i=1}^n$, with initialization $\theta_0 \in \Theta$ and number of iterations

$$T \geq \frac{\mathcal{R}(\theta_0) - \mathcal{R}(\theta^*)}{\gamma/2} + \log_2 \log_2 \left(\frac{6c_2 L}{m^2} \right),$$

returns an output such that

$$\|\theta_T - \theta^*\|_2 \leq \frac{12c_2}{m} = O\left(\sqrt{\frac{p^2}{n}}\right),$$

with probability at least $1 - T\delta \left(5 + \left\lceil \frac{\log\left(\frac{m}{M}\left(\frac{1}{2} - \kappa_1\right)\right)}{\log \kappa_2} \right\rceil\right)$.

Remark 7. Again, assuming 4-wise independence of the coordinates of the covariate distribution, we can reduce the dimension-dependence of the bounds (cf. [Remark 3](#)). We then take $\beta_h \asymp \sqrt{p} \|\Sigma_x\|_2^2$, to obtain an estimation error bound of the form $\|\widehat{\theta} - \theta^*\|_2 = O\left(\sqrt{\frac{p}{n}}\right)$.

We also briefly present a separate line of analysis that allows us to improve the estimation error rates from $O\left(\sqrt{\frac{p^2}{n}}\right)$ to $O\left(\sqrt{\frac{p}{n}}\right)$ under milder assumptions than the distributional assumptions mentioned in [Remark 7](#). [Theorem 1](#) of [Minsker \(2018\)](#) discusses an estimator \widehat{T} for the mean $\mathbb{E}[Y]$ of i.i.d. observations $Y_1, \dots, Y_n \in \mathbb{R}^{p \times p}$, with sub-Gaussian rates; i.e.,

$$\mathbb{P}\left(\|\widehat{T} - \mathbb{E}[Y]\|_2 \geq \|\mathbb{E}[Y^2]\|_2^{1/2} \sqrt{\frac{t}{n}}\right) \leq 2p \exp\left(-\frac{t}{2}\right),$$

for any $t > 0$. In our setting, we have

$$\mathbb{E}[Y^2] = \mathbb{E}\left[(\Phi''(x_i^T \theta))^2 x_i x_i^T x_i x_i^T\right] \leq \overline{M}_{\Phi, 2}^2 \mathbb{E}\left[\|x_i\|_2^2 \cdot x_i x_i^T\right].$$

In particular, for any unit vector $v \in \mathbb{R}^p$, we have

$$v^T \mathbb{E}[Y^2] v \leq \overline{M}_{\Phi, 2}^2 \mathbb{E}\left[\|x_i\|_2^2 \cdot (x_i^T v)^2\right] \leq \overline{M}_{\Phi, 2}^2 \sqrt{\mathbb{E}\left[\|x_i\|_2^4\right] \cdot \mathbb{E}\left[(x_i^T v)^4\right]}.$$

By assumption, we have

$$\mathbb{E}\left[(x_i^T v)^4\right] \leq \widetilde{C}_4 \left(\mathbb{E}\left[(x_i^T v)^2\right]\right)^2 \leq \widetilde{C}_4 \cdot \lambda_{\max}(\Sigma_x)^2 = O(1).$$

Furthermore,

$$\mathbb{E}\left[\|x_i\|_2^4\right] = \mathbb{E}\left[\left(\sum_{i=1}^p x_{ij}^2\right)^2\right],$$

and as argued in the proof of [Lemma 6](#), this is $O(p^2)$. Altogether, we conclude that $\|\mathbb{E}[Y^2]\|_2 = O(p)$, so using this high-probability bound in place of [Lemma 4](#) leads to an improvement in [Proposition 4](#) with $\beta_h = O\left(\sqrt{\frac{p}{n}}\right)$.

As mentioned at the end of [Section 2.4](#), the computational complexity of [Algorithm 2](#) is $\widetilde{O}(n + p)$, leading to an overall runtime of $\widetilde{O}(T(n + p^2))$ for the robust Newton method.

5. Simulations

We note that in our simulations, we have implemented the code from [Lai et al. \(2016\)](#) for agnostic mean estimation. In particular, the outlier truncation step is slightly different from the one analyzed in [Prasad et al. \(2020\)](#), and consequently also in our theorems above.

5.1. Huber's Contamination Model

We begin with simulations for linear and logistic regression in Huber's contamination model.

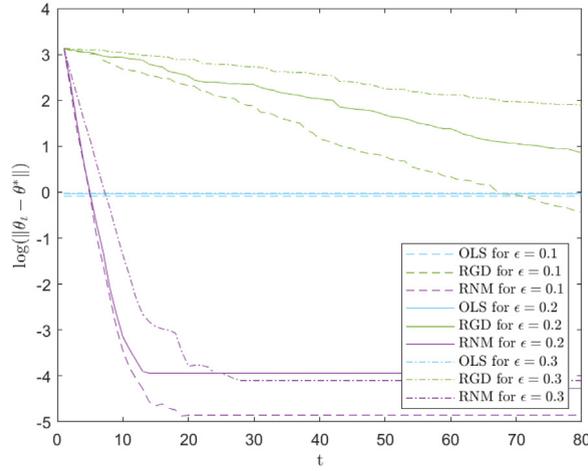


Fig. 1. Error $\log(\|\theta_t - \theta^*\|_2)$ with respect to each iteration of Robust Newton's Method (RNM), Robust Gradient Descent (RGD), and ordinary least squares (OLS), for linear regression with Huber contamination.

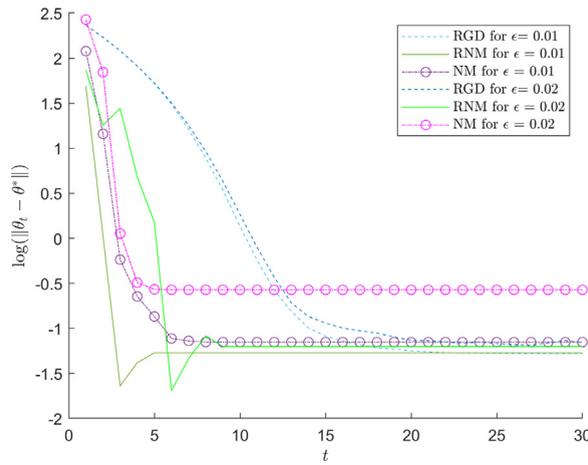


Fig. 2. Error $\log(\|\theta_t - \theta^*\|_2)$ with respect to each iteration for Robust Newton's Method (RNM) and Robust Gradient Descent (RGD), for logistic regression with Huber contamination. The behavior of the non-robust optimizer, found using Newton's method (NM), is also shown for reference.

5.1.1. Linear Regression

For our simulations, we set the dimension to be $p = 10$ and the number of data points to be $n = 1000$. We simulated the clean covariates as $x_i \sim N(0, I_p)$, with corresponding responses $y_i = x_i^T \theta^* + w_i$, where $w_i \sim N(0, 0.1)$ is i.i.d. noise and the true parameter is $\theta^* = \frac{1}{\sqrt{p}}(1, 1, \dots, 1)$. We simulated the outlier covariates as $x_i \sim N(0, p^2 I_p)$, with corresponding responses $y_i = 0$.

Figure 1 shows the results for Robust Newton's Method (RNM), Robust Gradient Descent (RGD), and ordinary least squares (OLS). We used the initialization $\theta_0 = (0.4, \dots, 0.4) + 10w$, with $w \sim \mathcal{N}(0, I_p)$, for both RNM and RGD. For RNM, we used the backtracking linesearch parameters $\kappa_1 = 0.01$, $\kappa_2 = 0.5$, and $\zeta = 10^{-8}$. For RGD, we used stepsize $\eta = 0.1$. We repeated the algorithm three times with contamination fractions $\epsilon = 0.1, 0.2$, and 0.3 . As seen in the figure, the statistical error indeed decreases quite quickly for RNM in comparison to RGD.

5.1.2. Logistic Regression

Next, we generated data from a logistic model with $p = 10$, $n = 1000$, and $\theta^* = (1/\sqrt{p}, \dots, 1/\sqrt{p})$, where we sampled the covariates as $x_i \sim \mathcal{N}(0, I_p)$ and sampled $y_i \in \{0, 1\}$ such that $p(y_i = 1|x_i) = \frac{1}{1 + e^{-x_i^T \theta^*}}$. We then randomly changed an ϵ fraction of the labels to be either 0 or 1, with equal probability. For various values of ϵ , we ran Robust Gradient Descent (RGD) and Robust Newton's Method (RNM), and plotted the parameter error in Figure 2. For RNM, we used the same backtracking linesearch parameters as in the case of linear regression with Huber contamination. For RGD, we used a stepsize of $\eta = 3$. As seen in the figure, the statistical error again decreases more quickly for RNM than for RGD.

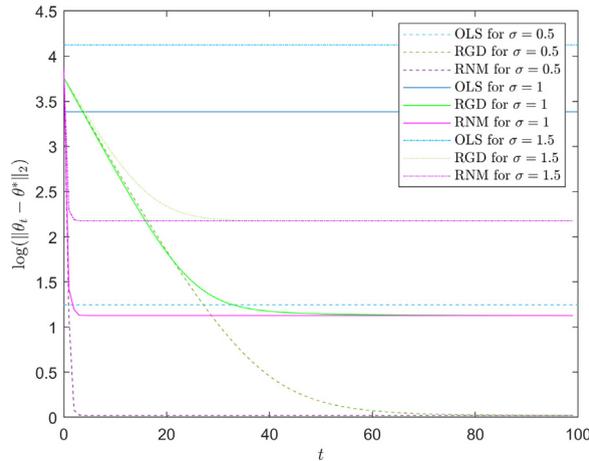


Fig. 3. Error $\log(\|\theta_t - \theta^*\|_2)$ with respect to each iteration for Robust Newton’s Method (RNM), Robust Gradient Descent (RGD), and ordinary least squares (OLS), for linear regression with heavy-tailed data.

5.2. Heavy-Tailed Data

For heavy-tailed data, we took $p = 10$ and $n = 1000$. We generated the covariates $x_i \sim N(0, I_p)$ and the corresponding responses $y_i = x_i^T \theta^* + w_i$, with w_i following a Pareto distribution with variance σ^2 and tail-index parameter β . We set the regression parameter $\theta^* = \frac{1}{\sqrt{\beta}}(1, 1, \dots, 1)$.

Figure 3 compares the results of Robust Newton’s Method (RNM), Robust Gradient Descent (RGD), and ordinary least squares (OLS). We used the initialization $\theta_0 = (10, 10, \dots, 10)$ for both RNM and RGD. For RNM, we used the backtracking linesearch parameters $\kappa_1 = 0.01$, $\kappa_2 = 0.5$, and $\zeta = 1000$. For RGD, we used stepsize $\eta = 0.1$. We repeated the algorithm three times, for $\sigma = 0.5, 1$, and 1.5 , all with $\beta = 1$. As seen in the figure, the statistical error again decreases more quickly for RNM than for RGD.

6. Robust Hessian Estimation: The Conjugate Gradient Approach

In this section, we discuss an alternative to Newton’s method (and present a robust variant thereof) which does not involve explicitly computing the Hessian. Inspired by Martens (2010), the idea is to estimate $\nabla^2 f(\theta)v$, for any vector v , using the approximation

$$h_v(\theta) = \frac{\nabla f(\theta + \delta v) - \nabla f(\theta)}{\delta}, \tag{33}$$

for some small $\delta > 0$. Note that in order to compute the Newton step $\Delta\theta_{nt}$, we need to solve the system $\nabla^2 f(\theta)\Delta\theta = -\nabla f(\theta)$, which we will do using the conjugate gradient algorithm, which provides an iterative method for solving a linear system of the form $Ax = b$ (Wright and Nocedal, 1999, Chapter 5). Our robust approach will involve using the robust gradient estimate $g(\theta)$ in place of $\nabla f(\theta)$.

The details of the algorithm are provided in Algorithm 6. Note that we have specified that the CGNewtonStep subroutine for finding the Newton direction on each iteration of CGRobustNewton runs for p steps, because in the noiseless case, the conjugate gradient method is known to terminate in at most p steps.

6.1. Convergence

We sketch some ideas here; a rigorous proof giving rates of convergence of the robust conjugate gradient method is beyond the scope of this work. Focusing on the pure Newton phase, note that our analysis of the iterates of robust Newton’s method essentially hinges on the Newton step $\Delta\theta_{nt}$ satisfying the equation

$$\nabla f(\theta_t) = -\nabla^2 f(\theta_t)\Delta\theta_{nt} + \chi_t, \tag{34}$$

where the next iterate is then defined by $\theta_{t+1} = \theta_t + \Delta\theta_{nt}$ and χ_t is a small, bounded error (cf. inequalities (A.10) and (A.11)). In particular, we can bound χ_t using the fact that $\Delta\theta_{nt} = -H(\theta_t)^{-1}g(\theta_t)$, and $\|g(\theta_t) - \nabla f(\theta_t)\|_2$ and $\|H(\theta_t) - \nabla^2 f(\theta_t)\|_2$ are small. In the case of the robust conjugate gradient method, we can again think of the conjugate gradient method as providing an approximate solution of the form

$$\nabla f(\theta_t) = -\nabla^2 f(\theta_t)\Delta\tilde{\theta}_{nt} + \tilde{\chi}_t, \tag{35}$$

Algorithm 6 Conjugate Gradient Robust Newton's Method

Require: Data samples $S = \{z_i\}_{i=1}^n$, Number of iterations T , Initial guess $\theta_0 \in \Theta$, Backtracking linesearch parameters $\kappa_1 \in (0, 0.5)$, $\kappa_2 \in (0, 1)$, and ζ , Tolerance δ

```

1:
2: function CGROBUSTNEWTON( $S, \xi, \theta_0, \kappa_1, \kappa_2, \zeta$ )
3:   for  $t = 0$  to  $T - 1$  do
4:     Compute losses  $\{\mathcal{L}(\theta_t, z_i)\}_{i=1}^n$  and gradients  $\{\nabla \mathcal{L}(\theta_t, z_i)\}_{i=1}^n$ 
5:     Compute Newton step  $\Delta\theta_{nt} = \text{CGNEWTONSTEP}(\theta_t)$ 
6:     Compute stepsize  $\alpha = \text{BACKTRACKINGLINESEARCH}(S, \theta_t, \Delta\theta_{nt}, g(\theta_t), \kappa_1, \kappa_2, \zeta)$ 
7:     Update  $\theta_{t+1} = \theta_t + \alpha \Delta\theta_{nt}$ 
8:   end for
9:   return  $\theta_T$ 
10: end function
11: function CGNEWTONSTEP( $\theta$ )
12:   Randomly initialize  $\Delta\theta^{(0)} \in \Theta$ 
13:   Compute gradient estimate  $g(\theta) = \text{ROBUSTGRADIENTESTIMATE}(S, \theta)$ 
14:   Compute Hessian-vector product estimate  $h_{\Delta\theta^{(0)}}(\theta) = \text{HVPRODUCT}(\theta, \Delta\theta^{(0)})$ 
15:   Set  $r_0 = h_{\Delta\theta^{(0)}}(\theta) + g(\theta)$ 
16:   Set  $p_0 = -r_0$ 
17:   for  $k = 1$  to  $p - 1$  do
18:     Compute Hessian-vector product estimate  $h_{p_k}(\theta) = \text{HVPRODUCT}(\theta, p_k)$ 
19:     Set  $\alpha_k = \frac{r_k^T r_k}{p_k^T h_{p_k}(\theta)}$ 
20:     Set  $\Delta\theta^{(k+1)} = \Delta\theta^{(k)} + \alpha_k p_k$ 
21:     Set  $r_{k+1} = r_k + \alpha_k h_{p_k}(\theta)$ 
22:     Set  $\beta_{k+1} = \frac{r_{k+1}^T r_{k+1}}{r_k^T r_k}$ 
23:     Set  $p_{k+1} = -r_{k+1} + \beta_{k+1} p_k$ 
24:   end for
25:   return  $\Delta\theta^{(d)}$ 
26: end function
27: function HVPRODUCT( $\theta, v$ )
28:   Compute gradient estimate  $g(\theta) = \text{ROBUSTGRADIENTESTIMATE}(S, \theta)$ 
29:   Compute gradient estimate  $g(\theta + \delta v) = \text{ROBUSTGRADIENTESTIMATE}(S, \theta + \delta v)$ 
30:   return  $\frac{g(\theta + \delta v) - g(\theta)}{\delta}$ 
31: end function
32: function BACKTRACKINGLINESEARCH( $(S, \theta, \Delta\theta_{nt}), g(\theta), \kappa_1, \kappa_2, \zeta$ )
33:   Set  $\alpha = 1$ 
34:   while  $\text{ROBUSTESTIMATE}(\{\mathcal{L}(\theta + \alpha \Delta\theta_{nt}, z_i)\}_{i=1}^n) > \text{ROBUSTESTIMATE}(\{\mathcal{L}(\theta, z_i)\}_{i=1}^n) + \kappa_1 \alpha g(\theta) \Delta\theta_{nt} + \zeta$  do
35:     Update  $\alpha = \kappa_2 \alpha$ 
36:   end while
37:   return  $\alpha$ 
38: end function

```

where successive iterates are then defined by $\tilde{\theta}_{t+1} = \tilde{\theta}_t + \Delta\tilde{\theta}_{nt}$. Thus, the main challenge is to understand the propagation of errors when the conjugate gradient method is applied to solve the system $Ax = b$, but the matrix-vector pair (A, b) is replaced by (\tilde{A}, \tilde{b}) on each iteration. To the best of our knowledge, this is actually an open question in optimization (Greenbaum, 1989; Greenbaum and Strakos, 1992). We note, however, that since our ultimate statistical estimation error bounds are all up to a small radius of, e.g., $O(\sqrt{\epsilon})$, we only need the output of the conjugate gradient method to be correct up to this error. In particular, as it is known that the exact conjugate gradient method terminates after p steps (Wright and Nocedal, 1999), it would for instance suffice to show that an inexact conjugate gradient method, where the error of (\tilde{A}, \tilde{b}) is also $O(\sqrt{\epsilon})$, only accumulates $O(\sqrt{\epsilon})$ error after p steps. Alternatively, one could try to derive a geometric rate of convergence (cf. equation (5.36) of Wright and Nocedal (1999)), with an additional additive error term, for inexact conjugate gradient steps. Clearly, each iterate of the conjugate gradient method has computational complexity $O(p^2)$, since it involves a small handful of matrix/vector multiplications. Thus, the overall complexity of p iterations would be $O(p^3)$, as well, leading to a computational complexity of $O(Tp^3)$ when combined with Newton's method.

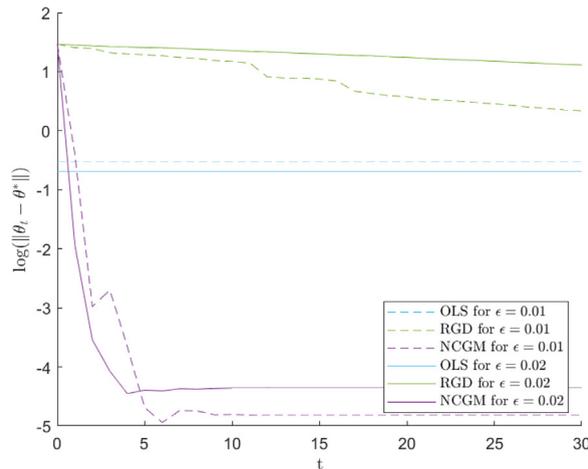


Fig. 4. Error $\log(\|\theta_t - \theta^*\|_2)$ with respect to each iteration for the Newton Conjugate Gradient Method (NCGM), Robust Gradient Descent (RGD), and ordinary least squares (OLS), for linear regression with Huber contamination.

We also need to quantify the error terms introduced to conjugate gradient steps due to inexactness. This depends on the increment δ used in the finite-difference approximation of the Hessian term (33). Note that by a Taylor expansion, we have

$$\nabla f(\theta + \delta v) = \nabla f(\theta) + \delta \nabla^2 f(\theta) v + C \delta^2,$$

for some constant C . Thus, we have the error bounds

$$\begin{aligned} \|h_v(\theta) - \nabla^2 f(\theta) v\|_2 &= \left\| \frac{g(\theta + \delta v) - g(\theta)}{\delta} - \nabla^2 f(\theta) v \right\|_2 \\ &= \left\| \frac{g(\theta + \delta v) - g(\theta)}{\delta} - \frac{\nabla f(\theta + \delta v) - \nabla f(\theta)}{\delta} - C \delta \right\|_2 \\ &\leq \frac{\|g(\theta + \delta v) - \nabla f(\theta + \delta v)\|_2}{\delta} + \frac{\|g(\theta) - \nabla f(\theta)\|_2}{\delta} + C \delta. \end{aligned}$$

If we had deviations bounds of the form (8), e.g., with $\alpha_g, \beta_g \asymp \sqrt{\epsilon}$, the optimal choice of δ would be $\delta \asymp \epsilon^{1/4}$.

In summary, we conjecture that the robust conjugate gradient method would allow us to incur an overall estimation error of $O(\epsilon^{1/4})$ in the case of Huber's ϵ -contamination model, again at a quadratic convergence rate for the successive Newton iterates. Although this is a slower rate than the one derived in Section 4 for GLMs, it may be applicable to a much wider range of settings. We also note that for the SEVER algorithm (Diakonikolas et al., 2019b), a rate of $O(\epsilon^{1/4})$ is derived for empirical risk minimization for a class of classification problems. If the above discussion could be made rigorous, it would also be extendable to the heavy-tailed setting in a straightforward manner.

6.2. Simulations

In Figure 4, we compare the Newton Conjugate Gradient Method (NCGM), Robust Gradient Descent (RGD), and ordinary least squares (OLS) on a linear model with Huber ϵ -contaminated data, with the same setup as in Section 5.1.1. We used the initial parameter $\theta_0 = (1, \dots, 1) + 2w$, with $w \sim N(0, I_p)$, for both NCGM and RGD. For NCGM, we used the backtracking linesearch parameters $\kappa_1 = 0.01$, $\kappa_2 = 0.5$, and $\zeta = 0.001$. We also used $\delta = 10^{-9}$ for the estimation of Hessian-vector products. For RGD, we used stepsize $\eta = 0.02$. We repeated the algorithm two times, with contamination fractions $\epsilon = 0.01$ and 0.02 .

In Figure 5, we compare the Newton Conjugate Gradient Method (NCGM), Robust Gradient Descent (RGD), and ordinary least squares (OLS) on a linear model with heavy-tailed data, again with the same setup as in Section 5.1.1. We used the initial parameter $\theta_0 = (1.5, \dots, 1.5) + 2w$, with $w \sim N(0, I_p)$, for both NCGM and RGD. For NCGM, we used the backtracking linesearch parameters $\kappa_1 = 0.01$, $\kappa_2 = 0.5$, and $\zeta = 0.00001$. We also used $\delta = 10^{-10}$ for the estimation of Hessian-vector products. For RGD, we used stepsize $\eta = 0.2$. We repeated the algorithm twice with $\sigma = 0.5$ and 0.25 , with $\beta = 0.7$.

In Figure 6, we compare the Newton Conjugate Gradient Method (NCGM) and Robust Gradient Descent (RGD) on a logistic model with Huber ϵ -contamination. To generate the contaminated logistic data, we used the same procedure outlined in Section 5.1.2. We also used the same hyperparameters for NCGM and RGD as in Section 5.1.2.

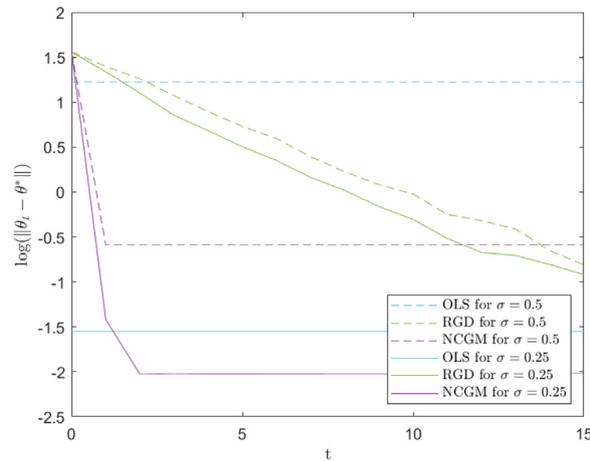


Fig. 5. Error $\log(\|\theta_n - \theta^*\|_2)$ with respect to each iteration for the Newton Conjugate Gradient Method (NCGD), Robust Gradient Descent (RGD), and ordinary least squares (OLS), for linear regression with heavy-tailed data.

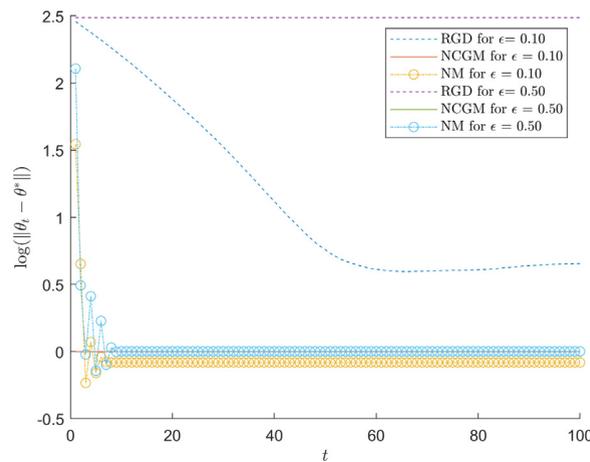


Fig. 6. Error $\log(\|\theta_n - \theta^*\|_2)$ with respect to each iteration for the Newton Conjugate Gradient Method (NCGM) and Robust Gradient Descent (RGD), for logistic regression with Huber contamination. The behavior of the non-robust optimizer, found using Newton's method (NM), is also shown for reference.

7. Discussion

We have presented a novel second-order method for robust parameter estimation, based on an adaptation of Newton's method where gradients and Hessians are computed in a robust manner on each iteration. In particular, we have shown that a variant of the backtracking linesearch algorithm will adaptively choose stepsizes in such a way that a finite number of iterates initially lie in a “damped” phase of the algorithm, after which the algorithm enters a “pure” phase where it only chooses stepsizes equal to 1 and converges quadratically to a small ball around the true parameter.

Under appropriate assumptions, our method shows clear computational advantages, both theoretically and empirically, in comparison to previously analyzed first-order methods. However, the general statements of [Theorems 3](#) and [4](#) leave much to be desired in terms of their dependence on p ; with an infinite computational budget, the rates from robust gradient descent depend only logarithmically on p in the Huber contamination model. It is thus natural to wonder whether improvements exist which either involve using less naive methods for computing robust Hessian matrices ([Minsker, 2018](#); [Minsker and Wang, 2022](#); [Cheng et al., 2019](#)) or a robust variant of a quasi-Newton method, which does not even attempt to estimate the Hessian matrices as closely ([Wright and Nocedal, 1999](#)). This opens up an interesting question of the “right” type of Hessian estimate which interpolates between robust gradient descent and the robust Newton method presented here, achieving the smallest number of iterations necessary for a desired level of statistical accuracy.

Another plausible extension of our analysis that could be studied under a similar theoretical framework would be to use robust gradient and Hessian estimators which employ the estimation procedures of [Diakonikolas and Kane \(2019\)](#) rather than those of [Lai et al. \(2016\)](#); we note that this would allow us to also handle the setting of adversarially contaminated data, rather than i.i.d. data from either an ϵ -contaminated or heavy-tailed model. We also note that the underlying assumption in our paper and all the aforementioned papers is that the clean data are drawn i.i.d. from a distribution. As pointed out

during the review process, extensions to the heteroscedastic case would be quite fascinating, but are beyond the scope of our current work. Even in the case of univariate mean estimation with no contamination, the analysis quickly becomes quite complicated (Pensia et al., 2022).

It would also be interesting and practically important to devise robust second-order algorithms appropriate for higher-dimensional data. For moderate to large p (even in settings where $p < n$), implementing the robust version of Newton's method can become more tedious, since it involves robustly computing $p \times p$ matrices and then inverting them on each iteration. In the truly high-dimensional case ($p > n$), even the canonical version of Newton's method must be modified, since the Hessian matrix becomes rank-deficient. This raises the question of whether it would be beneficial to analyze a robust inexact second-order algorithm, instead, where the Hessian matrix need not be approximated as closely. In the truly high-dimensional setting, combining this with regularization would be a natural direction for future work.

Finally, we have proposed the robust conjugate gradient method as an alternative second-order algorithm which, though based on Newton's method, only requires computing robust gradients rather than needing to separately compute robust Hessians. This method could potentially enjoy the fast convergence benefits of Newton's method while bypassing some of the computational issues in higher dimensions. However, a rigorous analysis of the robust conjugate gradient method is beyond the current scope of this paper—in particular, it would involve carefully tracking the propagation of errors through iterates of the conjugate gradient method, which has remained a long-standing open problem. We note that any error bounds on successive conjugate gradient iterates could then easily be plugged into our proofs to obtain quadratic convergence to an appropriate ball around the true parameter.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The work of EI was supported by the Cantab Capital Institute for the Mathematics of Information via the Philippa Fawcett Internship programme (Faculty of Mathematics, University of Cambridge). The authors thank the associate editor and anonymous reviewers for their helpful feedback, which improved the quality of the paper.

Appendix A. Proofs of Optimization-Theoretic Results

We now provide the proofs of the convergence results stated in Section 3.1.

A1. Proof of Theorem 1

Our first step is to show that backtracking linesearch chooses unit steps whenever the gradient is small, i.e., $\|\nabla f(\theta_t)\|_2 < \eta$. In other words, we want to prove that

$$\tilde{f}(\theta + \Delta\theta_{nt}) \leq \tilde{f}(\theta) - \kappa_1 \tilde{\lambda}(\theta)^2 + \zeta,$$

where $\theta = \theta_t$ denotes the iterate, \tilde{f} denotes the robust estimate of f , and we have defined the noisy Newton decrement

$$\tilde{\lambda}(\theta) := \left(\mathbf{g}(\theta)^T H^{-1}(\theta) \mathbf{g}(\theta) \right)^{1/2}. \quad (\text{A.1})$$

Recall that $\Delta\theta_{nt} = -H(\theta_t)^{-1} \mathbf{g}(\theta_t)$. Note that since $\|\nabla f(\theta_t)\|_2 < \eta$, we have (Boyd and Vandenberghe, 2004, Equation (9.11))

$$\|\theta_t - \theta^*\|_2 \leq \frac{2}{m} \|\nabla f(\theta_t)\|_2 < \frac{2\eta}{m} := \gamma_0. \quad (\text{A.2})$$

In particular, this implies a bound of $\gamma_g := \alpha_g \gamma_0 + \beta_g$ on the error of the gradient, and a bound of $\gamma_h := \alpha_h \gamma_0 + \beta_h$ on the error of the Hessian, according to Definitions 1 and 2. We will show that

$$f(\theta + \Delta\theta_{nt}) \leq f(\theta) - \kappa_1 \tilde{\lambda}(\theta)^2 + \frac{\zeta}{2}, \quad (\text{A.3})$$

from which the desired result clearly follows by the accuracy bound (11) on the robust estimates and the triangle inequality.

Note that $\tilde{\lambda}(\theta)^2 = \Delta\theta_{nt}^T H(\theta) \Delta\theta_{nt}$, implying that

$$\tilde{\lambda}(\theta)^2 \geq (m - \gamma_h) \|\Delta\theta_{nt}\|_2^2 > \frac{m}{2} \|\Delta\theta_{nt}\|_2^2 \quad (\text{A.4})$$

(where we assume $\gamma_h \leq \frac{m}{2}$). Furthermore, by the Lipschitz condition, for $u \geq 0$, we have

$$\|\nabla^2 f(\theta + u\Delta\theta_{nt}) - \nabla^2 f(\theta)\|_2 \leq uL \|\Delta\theta_{nt}\|_2,$$

so

$$|\Delta\theta_{nt}^T(\nabla^2 f(\theta + u\Delta\theta_{nt}) - \nabla^2 f(\theta))\Delta\theta_{nt}| \leq uL\|\Delta\theta_{nt}\|_2^3. \quad (\text{A.5})$$

Defining $\tilde{f}(u) := f(\theta + u\Delta\theta_{nt})$, we have $\tilde{f}''(u) = \Delta\theta_{nt}^T \nabla^2 f(\theta + u\Delta\theta_{nt}) \Delta\theta_{nt}$, so we can rewrite inequality (A.5) as

$$|\tilde{f}''(u) - \tilde{f}''(0)| \leq uL\|\Delta\theta_{nt}\|_2^3,$$

implying that

$$\tilde{f}''(u) \leq \tilde{f}''(0) + uL\|\Delta\theta_{nt}\|_2^3 \leq \tilde{f}''(0) + uL\left(\frac{2}{m}\right)^{3/2} \tilde{\lambda}(\theta)^3,$$

using inequality (A.4). Integrating with respect to u gives

$$\tilde{f}'(u) \leq \tilde{f}'(0) + u\tilde{f}''(0) + \frac{u^2 L}{2} \left(\frac{2}{m}\right)^{3/2} \tilde{\lambda}(\theta)^3,$$

and a second integration gives

$$\tilde{f}(u) \leq \tilde{f}(0) + u\tilde{f}'(0) + \frac{u^2}{2}\tilde{f}''(0) + \frac{u^3 L}{6} \left(\frac{2}{m}\right)^{3/2} \tilde{\lambda}(\theta)^3. \quad (\text{A.6})$$

Now note that

$$\begin{aligned} \tilde{f}'(0) &= \nabla f(\theta)^T \Delta\theta_{nt} \\ &= -\nabla f(\theta)^T H^{-1}(\theta)g(\theta) \\ &= -\tilde{\lambda}(\theta)^2 + (g(\theta) - \nabla f(\theta))H^{-1}(\theta)g(\theta) \\ &\leq -\tilde{\lambda}(\theta)^2 + \gamma_g \frac{1}{m - \gamma_h} (\|\nabla f(\theta)\|_2 + \gamma_g) \\ &\leq -\tilde{\lambda}(\theta)^2 + \gamma_g \frac{1}{m - \gamma_h} (\eta + \gamma_g), \end{aligned} \quad (\text{A.7})$$

whereas

$$\begin{aligned} \tilde{f}''(0) &= \Delta\theta_{nt}^T \nabla^2 f(\theta) \Delta\theta_{nt} \\ &= \tilde{\lambda}(\theta)^2 + \Delta\theta_{nt}^T (\nabla^2 f(\theta) - H(\theta)) \Delta\theta_{nt} \\ &\leq \tilde{\lambda}(\theta)^2 + \gamma_h \|\Delta\theta_{nt}\|_2^2 \\ &\leq \tilde{\lambda}(\theta)^2 \left(1 + \frac{2\gamma_h}{m}\right), \end{aligned} \quad (\text{A.8})$$

using the bound (A.4) in the last inequality. Plugging inequalities (A.7) and (A.8) into inequality (A.6) (with $u = 1$) then gives

$$f(\theta + \Delta\theta_{nt}) \leq f(\theta) + \left(-\tilde{\lambda}(\theta)^2 + \frac{\gamma_g(\eta + \gamma_g)}{m - \gamma_h}\right) + \frac{\tilde{\lambda}(\theta)^2}{2} \left(1 + \frac{2\gamma_h}{m}\right) + \frac{L}{6} \left(\frac{2}{m}\right)^{3/2} \tilde{\lambda}(\theta)^3.$$

Finally, note that

$$\tilde{\lambda}(\theta) \leq \|H^{-1/2}(\theta)\|_2 \cdot \|g(\theta)\|_2 \leq \frac{\|\nabla f(\theta)\|_2 + \gamma_g}{\sqrt{m - \gamma_h}} \leq \frac{\eta + \gamma_g}{\sqrt{m - \gamma_h}}.$$

Since

$$\zeta \geq 2 \left(\frac{\gamma_g \cdot 2\eta}{m/2} + \frac{\gamma_h}{m} \left(\frac{2\eta}{\sqrt{m/2}} \right)^2 \right)$$

and using the assumptions $\gamma_h \leq \frac{m}{2}$ and $\gamma_g \leq \eta$, we then have

$$f(\theta + \Delta\theta_{nt}) \leq f(\theta) - \tilde{\lambda}(\theta)^2 \left(\frac{1}{2} - \frac{L}{6} \left(\frac{2}{m} \right)^{3/2} \tilde{\lambda}(\theta) \right) + \frac{\zeta}{2}.$$

In particular, if $\tilde{\lambda}(\theta) \leq \frac{3 - 6\kappa_1}{L(2/m)^{3/2}}$, which is guaranteed if η is chosen sufficiently small so that

$$\frac{\eta + \gamma_g}{\sqrt{m - \gamma_h}} \leq \frac{3 - 6\kappa_1}{L(2/m)^{3/2}},$$

then inequality (A.3) is indeed satisfied. We can guarantee this last inequality by taking $\eta \leq \frac{3m^2(1-2\kappa_1)}{8L}$, assuming $\gamma_h \leq \frac{m}{2}$ and $\gamma_g \leq \eta$.

To derive the geometric convergence rate (13), we will use induction. We first establish an inequality of the form

$$\|\nabla f(\theta + \Delta\theta_{nt})\|_2 \leq c_1 \|\nabla f(\theta)\|_2^2 + c_2, \tag{A.9}$$

assuming $\|\nabla f(\theta)\|_2 < \eta$. Note that

$$\begin{aligned} \|\nabla f(\theta + \Delta\theta_{nt})\|_2 &= \|\nabla f(\theta + \Delta\theta_{nt}) - g(\theta) - H(\theta)\Delta\theta_{nt}\|_2 \\ &\leq \|\nabla f(\theta + \Delta\theta_{nt}) - \nabla f(\theta) - \nabla^2 f(\theta)\Delta\theta_{nt}\|_2 + \gamma_g + \gamma_h \|\Delta\theta_{nt}\|_2 \\ &= \left\| \int_0^1 (\nabla^2 f(\theta + u\Delta\theta_{nt}) - \nabla^2 f(\theta))\Delta\theta_{nt} du \right\|_2 + \gamma_g + \gamma_h \|\Delta\theta_{nt}\|_2 \\ &\leq \frac{L}{2} \|\Delta\theta_{nt}\|_2^2 + \gamma_g + \gamma_h \|\Delta\theta_{nt}\|_2, \end{aligned} \tag{A.10}$$

using the Lipschitz condition in the second inequality. Next, we use the bound

$$\|\Delta\theta_{nt}\|_2 = \|H^{-1}(\theta)g(\theta)\|_2 \leq \frac{1}{m - \gamma_h} (\|\nabla f(\theta)\|_2 + \gamma_g) \leq \frac{2}{m} (\|\nabla f(\theta)\|_2 + \gamma_g),$$

assuming $\gamma_h \leq \frac{m}{2}$. Plugging back into inequality (A.10) gives

$$\begin{aligned} \|\nabla f(\theta + \Delta\theta_{nt})\|_2 &\leq \frac{L}{2} \left(\frac{2(\|\nabla f(\theta)\|_2 + \gamma_g)}{m} \right)^2 + \gamma_g + \gamma_h \left(\frac{2(\|\nabla f(\theta)\|_2 + \gamma_g)}{m} \right) \\ &= \frac{2L}{m^2} \|\nabla f(\theta)\|_2^2 + \|\nabla f(\theta)\|_2 \left(\frac{4\gamma_g L}{m^2} + \frac{2\gamma_h}{m} \right) + \left(\frac{2L\gamma_g^2}{m^2} + \gamma_g + \frac{2\gamma_g\gamma_h}{m} \right) \\ &\leq \frac{2L}{m^2} \|\nabla f(\theta)\|_2^2 + \eta \left(\frac{4\gamma_g L}{m^2} + \frac{2\gamma_h}{m} \right) + \frac{2L\gamma_g^2}{m^2} + \gamma_g + \frac{2\gamma_g\gamma_h}{m}, \end{aligned} \tag{A.11}$$

giving inequality (A.9) with $c_1 = \frac{2L}{m^2}$ and $c_2 = \eta \left(\frac{4\gamma_g L}{m^2} + \frac{2\gamma_h}{m} \right) + \frac{2L\gamma_g^2}{m^2} + \gamma_g + \frac{2\gamma_g\gamma_h}{m}$. In particular, c_2 can be made small if we choose γ_g and γ_h small enough, and we will assume that $c_2 \leq \frac{\eta}{2}$. We will also assume that $c_1 c_2 \leq \frac{1}{12}$.

We are now ready for our induction. Using the notation $y_t := c_1 \|\nabla f(\theta_t)\|_2$, we will prove that $y_t < c_1 \eta$ and $y_t \leq y_0^{2^t} + c_1 c_2$ for all $t \geq 1$. For the base case $t = 1$, note that

$$y_1 \leq y_0^2 + c_1 c_2 = y_0^{2^1} + c_1 c_2,$$

using inequality (A.9). Furthermore, since $y_0 < c_1 \eta < \frac{1}{2}$ and $c_2 \leq \frac{\eta}{2}$ by assumption, we have

$$y_1 \leq \frac{y_0}{2} + \frac{c_1 \eta}{2} < \frac{c_1 \eta}{2} + \frac{c_1 \eta}{2} = c_1 \eta.$$

For the inductive step, suppose $t \geq 1$, and we have $y_s < c_1 \eta$ and $y_s \leq y_0^{2^s} + 3c_1 c_2$ for all $s \leq t$. Then by inequality (A.9), we have

$$y_{t+1} \leq y_t^2 + c_1 c_2. \tag{A.12}$$

Furthermore, using the assumption $\eta \leq \frac{m^2}{4L}$, we have $y_t < \frac{1}{2}$, so if $c_2 \leq \frac{\eta}{2}$, this implies that

$$y_{t+1} \leq \frac{y_t}{2} + \frac{c_1 \eta}{2} < \frac{c_1 \eta}{2} + \frac{c_1 \eta}{2} \leq c_1 \eta.$$

By inequality (A.12) and the induction hypothesis, we now write

$$\begin{aligned} y_{t+1} &\leq y_t^2 + c_1 c_2 \\ &\leq \left(y_0^{2^t} + 3c_1 c_2 \right)^2 + c_1 c_2 \\ &= y_0^{2^{t+1}} + 6c_1 c_2 y_0^{2^t} + 9c_1^2 c_2^2 + c_1 c_2 \\ &\leq y_0^{2^{t+1}} + \frac{3}{2} c_1 c_2 + \frac{3}{4} c_1 c_2 + c_1 c_2 \\ &\leq y_0^{2^{t+1}} + 3c_1 c_2, \end{aligned}$$

using the assumption $c_1 c_2 \leq \frac{1}{12}$. This completes the induction.

Thus,

$$c_1 \|\nabla f(\theta_t)\|_2 \leq \left(\frac{1}{2}\right)^{2t} + 3c_1 c_2.$$

Applying inequality (A.2) then gives the convergence rate

$$\|\theta_{t+1} - \theta^*\|_2 \leq \frac{2}{m} \cdot \frac{m^2}{2L} \left(\left(\frac{1}{2}\right)^{2^{t+1}} + \frac{6L}{m^2} c_2 \right),$$

completing the proof.

A2. Proof of Theorem 2

First, we show that we have an upper bound $\gamma'_0 := \frac{2}{m} \sqrt{2M(f(\theta_0) - f(\theta^*))}$ on $\|\theta_t - \theta^*\|_2$. We can then translate this into upper bounds $\gamma'_g := \alpha_g \gamma'_0 + \beta_g$ and $\gamma'_h := \alpha_h \gamma'_0 + \beta_h$ on the gradient and Hessian deviations, respectively.

By the result of Theorem 1, we must have $\|\nabla f(\theta_s)\|_2 \geq \eta$ for all $0 \leq s \leq t$. Indeed, suppose $\|\nabla f(\theta_s)\|_2 < \eta$ for some $s < t$. Then by Theorem 1 (with the iterate θ_s relabeled as θ_0), all successive iterates $\theta_{s+1}, \theta_{s+2}, \dots$, including θ_t , would also need to have the norm of the gradient bounded by η , which contradicts the assumption that $\|\nabla f(\theta_t)\|_2 \geq \eta$. We now show by induction that:

1. $f(\theta_s) \leq f(\theta_0)$, and
2. $\|\theta_s - \theta^*\|_2 \leq \gamma'_0$,

for all $0 \leq s \leq t$. For the base case $s = 0$, note that claim (1) is obvious. We can establish claim (2) by noting that

$$\|\theta_s - \theta^*\|_2 \leq \frac{2}{m} \|\nabla f(\theta_s)\|_2 \leq \frac{2}{m} \sqrt{2M(f(\theta_s) - f(\theta^*))} \leq \frac{2}{m} \sqrt{2M(f(\theta_0) - f(\theta^*))} = \gamma'_0, \quad (\text{A.13})$$

using inequality (A.2), inequality (9.14) of Boyd and Vandenberghe (2004), and claim (1).

Turning to the inductive step, suppose claims (1) and (2) hold for all $s \leq s'$, where $0 \leq s' < t$. We wish to establish the claims for $s = s' + 1$. Note that if we prove claim (1), then claim (2) follows by the same chain of inequalities (A.13). Thus, it remains to establish claim (1).

Assuming $\gamma'_g \leq \frac{\eta}{2}$, we have $\|g(\theta_s)\|_2 \geq \frac{\eta}{2}$ for all $0 \leq s \leq s'$ by claim (2), the fact that $\|\nabla f(\theta_s)\|_2 \geq \eta$, and the triangle inequality. Using the same notation for the Newton decrement (A.1), we note that

$$\tilde{\lambda}(\theta_s)^2 \geq \frac{1}{\sqrt{M + \gamma'_h}} \|g(\theta_s)\|_2^2 \geq \frac{\eta^2}{4\sqrt{2M}}, \quad (\text{A.14})$$

where we assume $\gamma'_h \leq M$.

First, we will prove that the exit condition of the *BacktrackingLineSearch* function will be satisfied, i.e., we want to prove that

$$\tilde{f}(\theta_{s'} + \alpha \Delta \theta_{nt}) \leq \tilde{f}(\theta) - \kappa_1 \alpha \tilde{\lambda}(\theta_{s'})^2 + \zeta \quad (\text{A.15})$$

holds for small enough α , where \tilde{f} is the robust estimate of f . For convenience, we use the notation $\theta = \theta_{s'}$ in what follows. In fact, we will show that

$$f(\theta + \alpha \Delta \theta_{nt}) \leq f(\theta) - \kappa_1 \alpha \tilde{\lambda}(\theta)^2 \quad (\text{A.16})$$

for small enough α , which clearly then implies inequality (A.15) by the triangle inequality and the condition (11).

Consider the following:

$$\begin{aligned} f(\theta + \alpha \Delta \theta_{nt}) &\leq f(\theta) + \alpha \nabla f(\theta)^T \Delta \theta_{nt} + \frac{M}{2} \|\Delta \theta_{nt}\|_2^2 \alpha^2 \\ &\leq f(\theta) + \alpha g(\theta)^T \Delta \theta_{nt} + \alpha \gamma'_g \|\Delta \theta_{nt}\|_2 + \frac{M}{2} \tilde{\lambda}(\theta)^2 \frac{2}{m} \alpha^2 \\ &= f(\theta) - \alpha \tilde{\lambda}(\theta)^2 + \alpha \gamma'_g \|\Delta \theta_{nt}\|_2 + \frac{M}{m} \tilde{\lambda}(\theta)^2 \alpha^2 \\ &\leq f(\theta) - \alpha \tilde{\lambda}(\theta)^2 + \alpha \gamma'_g \tilde{\lambda}(\theta) \sqrt{\frac{2}{m}} + \frac{M}{m} \tilde{\lambda}(\theta)^2 \alpha^2, \end{aligned}$$

where we use the relation $-\tilde{\lambda}(\theta)^2 = g(\theta)^T \Delta \theta_{nt}$ and inequality (A.4). Assuming $\gamma'_g \sqrt{\frac{2}{m}} \leq \frac{1}{2} \sqrt{\frac{\eta^2}{4\sqrt{2M}}}$ and using inequality (A.14), the last expression is upper-bounded as

$$f(\theta + \alpha \Delta \theta_{nt}) \leq f(\theta) - \frac{\alpha}{2} \tilde{\lambda}(\theta)^2 + \frac{M}{m} \tilde{\lambda}(\theta)^2 \alpha^2$$

$$= f(\theta) - \tilde{\lambda}(\theta)^2 \alpha \left(\frac{1}{2} - \frac{M}{m} \alpha \right). \quad (\text{A.17})$$

Hence, the condition (A.16) is indeed satisfied for sufficiently small α , i.e., $\alpha \leq \frac{m}{M}(\frac{1}{2} - \kappa_1)$, and in particular, the linesearch procedure must return a stepsize satisfying $\alpha \geq \kappa_2 \frac{m}{M}(\frac{1}{2} - \kappa_1)$. Plugging such a stepsize into inequality (A.15), we have

$$f(\theta_{s'+1}) \leq f(\theta_{s'}) - \kappa_1 \cdot \kappa_2 \frac{m}{M} \left(\frac{1}{2} - \kappa_1 \right) \cdot \frac{\eta^2}{4\sqrt{2}M} + \zeta \leq f(\theta_0) - \gamma + \zeta \leq f(\theta_0) - \frac{\gamma}{2}, \quad (\text{A.18})$$

using inequality (A.14) and the induction hypothesis. This implies that claim (1) is true, completing the induction.

Finally, note that the inequality $f(\theta_{t+1}) - f(\theta_t) < -\frac{\gamma}{2}$ follows by the same argument in inequality (A.18) with $\theta = \theta_t$, completing the proof.

Appendix B. Proofs of Auxiliary Results for GLMs

In this appendix, we prove some auxiliary results appearing in Section 4.

B1. Proof of Lemma 6

From the definition of the loss function (4), we have $\mathbb{E}[\nabla^2 \mathcal{L}(\theta, z)] = \mathbb{E}[\Phi''(x_i^T \theta) x_i x_i^T]$. By our assumptions on the boundedness of Φ'' and bounded eighth moments of x_i , we see that the distribution of the flattened Hessian $\text{flatten}(\nabla^2 \mathcal{L}(\theta, z))$ has bounded fourth moments. We then write

$$\begin{aligned} \text{tr}(\text{Cov}(\text{flatten}(\nabla^2 \mathcal{L}(\theta, z)))) &= \text{tr}(\text{Cov}(\text{flatten}(\Phi''(x_i^T \theta) x_i x_i^T))) \\ &\leq \text{tr}(\mathbb{E}[\text{flatten}(\Phi''(x_i^T \theta) x_i x_i^T) \text{flatten}(\Phi''(x_i^T \theta) x_i x_i^T)^T]) \\ &\leq \overline{M}_{\Phi,2}^2 \text{tr}(\mathbb{E}[\text{flatten}(x_i x_i^T) \text{flatten}(x_i x_i^T)^T]) \\ &= \overline{M}_{\Phi,2}^2 \sum_{j,k=1}^p \mathbb{E}[x_{ij}^2 x_{ik}^2] \\ &\leq \overline{M}_{\Phi,2}^2 \mathbb{E} \left[\left(\sum_{j=1}^p x_{ij} \right)^4 \right] \\ &\leq \overline{M}_{\Phi,2}^2 \mathbb{E} \left[(x_i^T \mathbf{1})^4 \right], \end{aligned}$$

where $\mathbf{1}$ denotes the all-ones vector. Finally, note that

$$\begin{aligned} \mathbb{E} \left[(x_i^T \mathbf{1})^4 \right] &\leq \tilde{C}_4 \mathbb{E} \left[(x_i^T \mathbf{1})^2 \right]^2 \\ &\leq \tilde{C}_4 p^2 \|\Sigma_x\|_2^2, \end{aligned}$$

implying the desired result.

B2. Proof of Proposition 1

For the Lipschitz condition, note that for any $\theta_1, \theta_2 \in \mathbb{R}^p$, we have

$$\begin{aligned} \|\nabla^2 \mathcal{R}(\theta_1) - \nabla^2 \mathcal{R}(\theta_2)\|_2 &= \|\nabla^2 \mathcal{L}(\theta_1, z) - \nabla^2 \mathcal{L}(\theta_2, z)\|_2 \\ &= \|\mathbb{E}[x_i x_i^T (\Phi''(x_i^T \theta_1) - \Phi''(x_i^T \theta_2))]\|_2 \\ &= \sup_{u \in \mathbb{S}^{p-1}} u^T \mathbb{E}[x_i x_i^T (\Phi''(x_i^T \theta_1) - \Phi''(x_i^T \theta_2))] u \\ &= \sup_{u \in \mathbb{S}^{p-1}} \mathbb{E}[(u^T x_i)^2 (\Phi''(x_i^T \theta_1) - \Phi''(x_i^T \theta_2))] \\ &\leq \sup_{u \in \mathbb{S}^{p-1}} \mathbb{E}[(u^T x_i)^4]^{\frac{1}{2}} \mathbb{E}[(\Phi''(x_i^T \theta_1) - \Phi''(x_i^T \theta_2))^2]^{\frac{1}{2}} \\ &\leq \sup_{u \in \mathbb{S}^{p-1}} \sqrt{\tilde{C}_4} \mathbb{E}[(u^T x_i)^2] \sqrt{\overline{M}_{\Phi,3}} \|\theta_1 - \theta_2\|_2 \\ &\leq \sqrt{\tilde{C}_4 \overline{M}_{\Phi,3}} \|\Sigma_x\|_2 \|\theta_1 - \theta_2\|_2, \end{aligned}$$

where we use the mean value theorem to upper-bound the expectation in the second-to-last inequality.

For any $\theta \in \mathbb{R}^p$, we have

$$\begin{aligned}
\|\nabla^2 \mathcal{R}(\theta)\|_2 &= \|\nabla^2 \mathcal{L}(\theta, z)\|_2 \\
&= \|\mathbb{E}[x_i x_i^T (\Phi''(x_i^T \theta))]\|_2 \\
&= \sup_{u \in \mathbb{S}^{p-1}} u^T \mathbb{E}[x_i x_i^T (\Phi''(x_i^T \theta))] u \\
&= \sup_{u \in \mathbb{S}^{p-1}} \mathbb{E}[(u^T x_i)^2 (\Phi''(x_i^T \theta))] \\
&\leq \sup_{u \in \mathbb{S}^{p-1}} \mathbb{E}[(u^T x_i)^4]^{\frac{1}{2}} \mathbb{E}[(\Phi''(x_i^T \theta))^2]^{\frac{1}{2}} \\
&\leq \sup_{u \in \mathbb{S}^{p-1}} \sqrt{\tilde{C}_4} \mathbb{E}[(u^T x_i)^2] \bar{M}_{\Phi, 2} \\
&\leq \bar{M}_{\Phi, 2} \sqrt{\tilde{C}_4} \|\Sigma_x\|_2.
\end{aligned}$$

B3. Proof of Proposition 2

Suppose $v \in \mathbb{R}^p$ is a unit vector. We write

$$\begin{aligned}
v^T \nabla^2 \mathcal{R}(\theta) v &= \mathbb{E}[(v^T x_i)^2 \cdot \Phi''(x_i^T \theta)] \\
&\geq \mathbb{E}[(v^T x_i)^2 \cdot b_\tau \mathbf{1}\{|x_i^T \theta| \leq \tau\}] \\
&= b_\tau (\mathbb{E}[(v^T x_i)^2] - \mathbb{E}[(v^T x_i)^2 \cdot \mathbf{1}\{|x_i^T \theta| > \tau\}]) \\
&\geq b_\tau \left(\lambda_{\min}(\Sigma_x) - \sqrt{\mathbb{E}[(v^T x_i)^4] \cdot \mathbb{P}(|x_i^T \theta| > \tau)} \right) \\
&\geq b_\tau \left(\lambda_{\min}(\Sigma_x) - \sqrt{\tilde{C}_4 \|\Sigma_x\|_2^2 \cdot \mathbb{P}(|x_i^T \theta| > \tau)} \right) \\
&\geq \frac{b_\tau}{2} \lambda_{\min}(\Sigma_x),
\end{aligned}$$

where we have used the fact that Φ'' is always nonnegative in the first inequality, applied Cauchy-Schwarz in the second inequality, and used the assumption (26) in the last inequality.

Appendix C. Proofs about Huber Contamination

In this appendix, we provide proofs of the results stated in Section 4.2.

C1. Bounds on Error Terms

Our first lemma shows how small the parameters $(\alpha_g, \beta_g, \alpha_h, \beta_h)$ in the robust gradient and Hessian estimates need to be in order to satisfy the assumptions of Theorems 1 and 2.

Lemma 9. Define

$$\begin{aligned}
\hat{\alpha}_g &:= \min \left\{ \frac{m}{64}, \frac{\eta m}{8\sqrt{2M}(f(\theta_0) - f(\theta^*))}, \sqrt{\frac{\eta^2 m}{8\sqrt{2M}}} \cdot \frac{m}{8\sqrt{2M}(f(\theta_0) - f(\theta^*))} \right\}, \\
\hat{\beta}_g &:= \min \left\{ \frac{\eta}{32}, \frac{\eta m}{8\sqrt{2M}(f(\theta_0) - f(\theta^*))}, \frac{1}{4} \sqrt{\frac{\eta^2 m}{8\sqrt{2M}}} \right\}, \\
\hat{\alpha}_h &:= \min \left\{ \frac{m^2}{256\eta}, \frac{mM}{4\sqrt{2M}(f(\theta_0) - f(\theta^*))} \right\}, \\
\hat{\beta}_h &:= \min \left\{ \frac{m}{128}, \frac{M}{2} \right\}.
\end{aligned}$$

Suppose $\alpha_g \leq \hat{\alpha}_g$, $\beta_g \leq \hat{\beta}_g$, $\alpha_h \leq \hat{\alpha}_h$, and $\beta_h \leq \hat{\beta}_h$. Then the bounds (10) and (14) of Theorem 1, as well as the bounds (15) and (16) of Theorem 2, are satisfied.

Proof. Under the assumptions, we have

$$\begin{aligned}\gamma_g &= \frac{2\eta\alpha_g}{m} + \beta_g \leq \frac{\eta}{32} + \frac{\eta}{32} = \frac{\eta}{16}, \\ \gamma_h &= \frac{2\eta\alpha_h}{m} + \beta_h \leq \frac{m}{128} + \frac{m}{128} = \frac{m}{64}, \\ \frac{2\alpha_g}{m} \sqrt{2M(f(\theta_0) - f(\theta^*))} + \beta_g &\leq \min \left\{ \frac{\eta}{2}, \sqrt{\frac{m}{2}} \cdot \frac{1}{2} \sqrt{\frac{\eta^2}{4\sqrt{2M}}} \right\}, \\ \frac{2\alpha_h}{m} \sqrt{2M(f(\theta_0) - f(\theta^*))} + \beta_h &\leq M.\end{aligned}$$

Hence, inequalities (10), (15), and (16) are satisfied. Using the fact that $\eta := \frac{m^2}{8L} \cdot \min\{3(1 - 2\kappa_1), 2\} \leq \frac{m^2}{4L}$, we have $\frac{L}{m^2} \leq \frac{1}{4\eta}$. Then

$$\begin{aligned}c_2 &= \eta \left(\frac{4\gamma_g L}{m^2} + \frac{2\gamma_h}{m} \right) + \frac{2L\gamma_g^2}{m^2} + \gamma_g + \frac{2\gamma_g\gamma_h}{m} \\ &\leq \eta \left(\frac{4\eta}{16} \cdot \frac{1}{4\eta} + \frac{2}{m} \cdot \frac{m}{64} \right) + \frac{2\eta^2}{256} \cdot \frac{1}{4\eta} + \frac{\eta}{16} + \frac{2}{m} \cdot \frac{\eta}{16} \cdot \frac{m}{64} \\ &< \frac{\eta}{6} \\ &\leq \frac{m^2}{24L}.\end{aligned}$$

Hence, inequality (14) is satisfied. \square

In Propositions 3 and 4, below, we derive expressions for $(\alpha_g, \beta_g, \alpha_h, \beta_h)$ for the Huber contamination and heavy-tailed models, which will then allow us to translate the conditions of Lemma 9 into assumptions involving the contamination level and/or minimum sample size required for our theoretical results to hold.

We begin with a result concerning the parameters $(\alpha_g, \beta_g, \alpha_h, \beta_h)$ controlling the robust gradient and Hessian errors.

Proposition 3. Under the assumptions above, the gradient and Hessian estimates with Type = Huber returned by Algorithms 4 and 5, respectively, satisfy the conditions of Definitions 1 and 2 with the following parameters:

$$\begin{aligned}\alpha_g &= c_1(\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon))\sqrt{\|\Sigma_x\|_2 \log p}, \\ \beta_g &= c_2(\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon))\sqrt{\|\Sigma_x\|_2 \log p}, \\ \alpha_h &= 0, \\ \beta_h &= c_3(\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon))\|\Sigma_x\|_2 p \sqrt{\log p},\end{aligned}$$

with probability at least $1 - \delta$.

Proof. By Lemma 5, the true distribution of the gradients $\nabla\mathcal{L}(\theta, z)$ has bounded fourth moments. Moreover,

$$\begin{aligned}\|\text{Cov}(\nabla\mathcal{L}(\theta, z))\|_2 &\leq C_1 \|\Sigma_x\|_2 \left(\sqrt{L_{\Phi,4}} + L_{\Phi,2} \right) \|\theta - \theta^*\|_2^2 \\ &\quad + C_2 \|\Sigma_x\|_2 \left(B_{\Phi,2} + \sqrt{B_{\Phi,4}} + c(\sigma) \sqrt{M_{\Phi,2,2}} + \sqrt{c(\sigma)^3 M_{\Phi,4,1}} \right).\end{aligned}$$

Plugging the above bound into inequality (17) of Lemma 1, we obtain

$$\begin{aligned}\|g(\theta) - \mathbb{E}[\nabla\mathcal{L}(\theta, z)]\|_2 &\leq C'_1(\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon))\sqrt{\|\text{Cov}(\nabla\mathcal{L}(\theta, z))\|_2 \log p} \\ &\leq c_1(\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon))\sqrt{\|\Sigma_x\|_2 \log p} \cdot \|\theta - \theta^*\|_2 \\ &\quad + c_2(\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon))\sqrt{\|\Sigma_x\|_2 \log p}.\end{aligned}$$

Hence, the gradient estimate returned by Algorithm 4 satisfies Definition 1 with $\alpha_g = c_1(\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon))\sqrt{\|\Sigma_x\|_2 \log p}$ and $\beta_g = c_2(\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon))\sqrt{\|\Sigma_x\|_2 \log p}$.

By Lemma 6, the true distribution of the flattened Hessian $\text{flatten}(\nabla^2\mathcal{L}(\theta, z))$ has bounded fourth moments. Moreover, combining Lemma 6 with Lemma 3, we obtain

$$\begin{aligned}\|H(\theta) - \mathbb{E}[\nabla^2\mathcal{L}(\theta, z)]\|_2 &\leq C'_2(\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon))\sqrt{\|\text{Cov}(\text{flatten}(\nabla^2\mathcal{L}(\theta, z)))\|_2 \log p} \\ &\leq c_3(\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon))\|\Sigma_x\|_2 p \sqrt{\log p}.\end{aligned}$$

Hence, the Hessian estimate returned by Algorithm 5 satisfies Definition 2 with $\alpha_h = 0$ and $\beta_h = c_3(\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon))\|\Sigma_x\|_2 p \sqrt{\log p}$. \square

C2. Proof of Theorem 3

We will apply Theorems 1 and 2 to show that Algorithm 5 returns $\hat{\theta}_T$ such that $\|\theta_T - \theta^*\|_2 \leq \frac{12c_2}{m} = O\left(\epsilon \log(p) + \sqrt{\epsilon \log(p)}\right)$.

Under the assumption that $\sqrt{\epsilon} + \gamma(n, p, \delta, \epsilon) < \hat{\gamma}$, and using Proposition 3, the assumptions of Lemma 9 are satisfied. Note that the assumptions of Lemma 7 are likewise satisfied by the condition (29). Applying Theorem 2, the risk $\mathcal{R}(\theta_t)$ is reduced by at least $\frac{\gamma}{2}$ in each step of the damped Newton phase of the algorithm. Hence, the number of such iterations cannot exceed

$$T_{damp} := \frac{\mathcal{R}(\theta_0) - \mathcal{R}(\theta^*)}{\gamma/2}. \tag{C.1}$$

Define

$$T_{pure} := \log_2 \log_2 \left(\frac{6c_2 L}{m^2} \right). \tag{C.2}$$

Applying Theorem 1, we observe that after T_{pure} iterations in the pure Newton phase, we have $\frac{m}{L} \left(\frac{1}{2}\right)^{2^{T_{pure}}} < \frac{6c_2}{m}$. Therefore, from inequality (13), we have $\|\hat{\theta} - \theta^*\|_2 \leq \frac{12c_2}{m}$. Combining inequalities (C.1) and (C.2), we obtain the bound on the total number of iterations T .

From the preceding analysis on the robust gradient and Hessian estimators, observe that η , L , and m are independent of ϵ and p , while γ_g is $O(\sqrt{\epsilon \log p})$ and γ_h is $O(p\sqrt{\epsilon \log p})$. Hence, from inequality (13), we have $\|\theta_T - \theta^*\|_2 = O\left(p\sqrt{\epsilon \log p}\right)$.

We now compute the error probability of the algorithm via a union bound. For each of the T gradient and Hessian calculations, we have a possible error of δ . Furthermore, each call of backtracking linesearch incurs a possible error from the robust estimates, by Lemma 7; once at θ_t and once for each value of α used in the linesearch. This is a total of $2T_{pure}$ evaluations for the pure Newton steps, and a maximum of $T_{damp} \left(1 + \left\lceil \frac{\log\left(\frac{m}{M}\left(\frac{1}{2} - \kappa_1\right)\right)}{\log \kappa_2} \right\rceil\right)$ evaluations for the damped Newton steps. Thus, the overall probability of error is at most

$$2T\delta + 2T_{pure}\delta + T_{damp} \left(1 + \left\lceil \frac{\log\left(\frac{m}{M}\left(\frac{1}{2} - \kappa_1\right)\right)}{\log \kappa_2} \right\rceil\right) \delta \leq T\delta \left(5 + \left\lceil \frac{\log\left(\frac{m}{M}\left(\frac{1}{2} - \kappa_1\right)\right)}{\log \kappa_2} \right\rceil\right).$$

Appendix D. Proofs about Heavy-Tailed Contamination

In this appendix, we provide proofs of the results stated in Section 4.3.

D1. Bounds on Error Terms

The first result concerns the parameters $(\alpha_g, \beta_g, \alpha_h, \beta_h)$, which control the robust gradient and Hessian errors.

Proposition 4. Under the assumptions above, the gradient and Hessian estimates with Type = Heavy-tail returned by Algorithms 4 and 5, respectively, satisfy the conditions of Definitions 1 and 2 with the following parameters:

$$\begin{aligned} \alpha_g &= c_1 \sqrt{\frac{p \log(1.4/\delta)}{n}}, \\ \beta_g &= c_2 \sqrt{\frac{p \log(1.4/\delta)}{n}}, \\ \alpha_h &= 0, \\ \beta_h &= c_3 \|\Sigma_x\|_2 p \sqrt{\frac{\log(1.4/\delta)}{n}}, \end{aligned}$$

with probability at least $1 - \delta$.

Proof. By Lemma 5, the distribution of the gradients $\nabla \mathcal{L}(\theta, z)$ has bounded fourth moments. Moreover,

$$\begin{aligned} \|\text{Cov}(\nabla \mathcal{L}(\theta, z))\|_2 &\leq C_1 \|\Sigma_x\|_2 \left(\sqrt{L_{\Phi,4}} + L_{\Phi,2} \right) \|\theta - \theta^*\|_2^2 \\ &\quad + C_2 \|\Sigma_x\|_2 \left(B_{\Phi,2} + \sqrt{B_{\Phi,4}} + c(\sigma) \sqrt{M_{\Phi,2,2}} + \sqrt{c(\sigma)^3 M_{\Phi,4,1}} \right). \end{aligned}$$

Plugging this bound into inequality (19) of Lemma 2, we obtain

$$\begin{aligned} \|g(\theta) - \mathbb{E}[\nabla\mathcal{L}(\theta, z)]\|_2 &\leq 11\sqrt{\frac{\text{tr}(\text{Cov}(\nabla\mathcal{L}(\theta, z))) \log(1.4/\delta)}{n}} \\ &\leq 11\sqrt{\frac{p \text{Cov}(\nabla\mathcal{L}(\theta, z)) \log(1.4/\delta)}{n}} \\ &\leq c_1\sqrt{\frac{p \log(1.4/\delta)}{n}} \|\theta - \theta^*\|_2 + c_2\sqrt{\frac{p \log(1.4/\delta)}{n}}. \end{aligned}$$

Hence, the gradient estimate returned by Algorithm 4 satisfies Definition 1 with $\alpha_g = c_1\sqrt{\frac{p \log(1.4/\delta)}{n}}$ and $\beta_g = c_2\sqrt{\frac{p \log(1.4/\delta)}{n}}$.

By Lemma 6, the distribution of the flattened Hessian $\text{flatten}(\nabla^2\mathcal{L}(\theta, z))$ has bounded fourth moments. Moreover, combining Lemma 6 with Lemma 4, we obtain

$$\begin{aligned} \|H(\theta) - \mathbb{E}[\nabla^2\mathcal{L}(\theta, z)]\|_2 &\leq C_3\sqrt{\frac{\text{tr}(\text{Cov}(\text{flatten}(\nabla^2\mathcal{L}(\theta, z)))) \log(1.4/\delta)}{n}} \\ &\leq c_3\|\Sigma_x\|_2 p\sqrt{\frac{\log(1.4/\delta)}{n}}. \end{aligned}$$

Hence, the Hessian estimate returned by Algorithm 5 satisfies Definition 2 with $\alpha_h = 0$ and $\beta_h = c_3\|\Sigma_x\|_2\sqrt{\frac{\log(1.4/\delta)}{n}}$. \square

D2. Proof of Theorem 4

We will follow a similar outline as in the proof of Theorem 3.

Using the assumption on n and Proposition 4, it is straightforward to verify that the conditions of Lemma 9 are satisfied. Furthermore, the conditions of Lemma 8 are satisfied by inequality (32), as well. Applying Theorem 2, the risk $\mathcal{R}(\theta_t)$ is reduced by at least $\frac{\gamma}{2}$ in each step of the damped Newton phase of the algorithm. Hence, the number of such iterations cannot exceed T_{damp} , defined as in Eq. (C.1). Applying Theorem 1, we observe that after T_{pure} iterations (defined as in Eq. (C.2)) in the pure Newton phase, we have $\frac{m}{L}\left(\frac{1}{2}\right)^{2t} < \frac{6c_2}{m}$. Therefore, from inequality (13), we have $\|\hat{\theta} - \theta^*\|_2 \leq \frac{12c_2}{m}$. Combining inequalities (C.1) and (C.2), we obtain the bound on the total number of iterations T .

From the preceding analysis on the robust gradient and Hessian estimators (cf. Proposition 4), observe that $\gamma_g = O\left(\sqrt{\frac{p}{n}}\right)$ and $\gamma_h = O\left(\sqrt{\frac{p^2}{n}}\right)$. Hence, c_2 is $O\left(\sqrt{\frac{p^2}{n}}\right)$. From inequality (13), we then have $\|\theta_T - \theta^*\|_2 = O\left(\sqrt{\frac{p^2}{n}}\right)$.

Computing the error probability of the algorithm via a union bound is the same as in Theorem 3 with the use of appropriate gradient, Hessian, and robust estimates.

References

- Bertsekas, D., 2015. *Convex Optimization Algorithms*. Athena Scientific.
- Boyd, S.P., Vandenberghe, L., 2004. *Convex Optimization*. Cambridge University Press.
- Chen, M., Gao, C., Ren, Z., 2016. A general decision theory for Hubers ϵ -contamination model. *Electronic Journal of Statistics* 10 (2), 3752–3774.
- Cheng, Y., Diakonikolas, I., Ge, R., Woodruff, D.P., 2019. Faster algorithms for high-dimensional robust covariance estimation. In: *Conference on learning theory*. PMLR, pp. 727–757.
- Cherapanamjeri, Y., Aras, E., Tripuraneni, N., Jordan, M.I., Flammarion, N., Bartlett, P.L., 2020a. Optimal robust linear regression in nearly linear time. *arXiv preprint arXiv:2007.08137*.
- Cherapanamjeri, Y., Hopkins, S.B., Kathuria, T., Raghavendra, P., Tripuraneni, N., 2020b. Algorithms for heavy-tailed statistics: Regression, covariance estimation, and beyond. In: *Proceedings of the 52nd annual ACM SIGACT symposium on theory of computing*, pp. 601–609.
- Chinot, G., Lecué, G., Lerasle, M., 2020. Robust statistical learning with Lipschitz and convex loss functions. *Probability Theory and Related Fields* 176 (3), 897–940.
- Cohen, M.B., Lee, Y.T., Miller, G., Pachocki, J., Sidford, A., 2016. Geometric median in nearly linear time. In: *Proceedings of the forty-eighth annual ACM symposium on theory of computing*, pp. 9–21.
- Depersin, J., 2020. A spectral algorithm for robust regression with subgaussian rates. *arXiv preprint arXiv:2007.06072*.
- Diakonikolas, I., Kamath, G., Kane, D., Li, J., Moitra, A., Stewart, A., 2019a. Robust estimators in high-dimensions without the computational intractability. *SIAM Journal on Computing* 48 (2), 742–864.
- Diakonikolas, I., Kamath, G., Kane, D., Li, J., Steinhardt, J., Stewart, A., 2019b. SEVER: A robust meta-algorithm for stochastic optimization. In: *International conference on machine learning*. PMLR, pp. 1596–1606.
- Diakonikolas, I., Kane, D.M., 2019. Recent advances in algorithmic high-dimensional robust statistics. *arXiv preprint arXiv:1911.05911*.
- Diakonikolas, I., Kong, W., Stewart, A., 2019c. Efficient algorithms and lower bounds for robust linear regression. In: *Proceedings of the thirtieth annual ACM-SIAM symposium on discrete algorithms*. SIAM, pp. 2745–2754.
- Greenbaum, A., 1989. Behavior of slightly perturbed Lanczos and conjugate-gradient recurrences. *Linear Algebra and its Applications* 113, 7–63.
- Greenbaum, A., Strakos, Z., 1992. Predicting the behavior of finite precision Lanczos and conjugate gradient computations. *SIAM Journal on Matrix Analysis and Applications* 13 (1), 121–137.
- Huber, P.J., Ronchetti, E.M., 2011. *Robust Statistics*. Wiley Series in Probability and Statistics. Wiley.
- Johnstone, I.M., Lu, A.Y., 2009. Sparse principal components analysis. *arXiv preprint arXiv:0901.4392*.

- Lai, K.A., Rao, A.B., Vempala, S., 2016. Agnostic estimation of mean and covariance. In: 2016 IEEE 57th annual symposium on foundations of computer science (FOCS). IEEE, pp. 665–674.
- Lecué, G., Lerasle, M., 2019. Learning from MOMs principles: Le Cams approach. *Stochastic Processes and their Applications* 129 (11), 4385–4410.
- Lecué, G., Lerasle, M., 2020. Robust machine learning by median-of-means: Theory and practice. *The Annals of Statistics* 48 (2), 906–931.
- Lecué, G., Lerasle, M., Mathieu, T., 2020. Robust classification via MOM minimization. *Machine Learning* 109 (8), 1635–1665.
- Lugosi, G., Mendelson, S., 2019a. Mean estimation and regression under heavy-tailed distributions: A survey. *Foundations of Computational Mathematics* 19 (5), 1145–1190.
- Lugosi, G., Mendelson, S., 2019b. Regularization, sparse recovery, and median-of-means tournaments. *Bernoulli* 25 (3), 2075–2106.
- Lugosi, G., Mendelson, S., 2019c. Risk minimization by median-of-means tournaments. *Journal of the European Mathematical Society* 22 (3), 925–965.
- Maronna, R.A., Martin, R.D., Yohai, V.J., Salibián-Barrera, M., 2019. *Robust Statistics: Theory and Methods* (with R). John Wiley & Sons.
- Martens, J., 2010. Deep learning via Hessian-free optimization. In: *ICML*, vol. 27, pp. 735–742.
- Minsker, S., 2015. Geometric median and robust estimation in Banach spaces. *Bernoulli* 21 (4), 2308–2335.
- Minsker, S., 2018. Sub-Gaussian estimators of the mean of a random matrix with heavy-tailed entries. *The Annals of Statistics* 46 (6A), 2871–2903.
- Minsker, S., Wang, L., 2022. Robust estimation of covariance matrices: Adversarial contamination and beyond. *arXiv preprint arXiv:2203.02880*.
- Pan, V.Y., Chen, Z.Q., 1999. The complexity of the matrix eigenproblem. In: *Proceedings of the thirty-first annual ACM symposium on theory of computing*, pp. 507–516.
- Pensia, A., Jog, V., Loh, P., 2020. Robust regression with covariate filtering: Heavy tails and adversarial contamination. *arXiv preprint arXiv:2009.12976*.
- Pensia, A., Jog, V., Loh, P., 2022. Estimating location parameters in sample-heterogeneous distributions. *Information and Inference: A Journal of the IMA* 11 (3), 959–1036.
- Prasad, A., Suggala, A.S., Balakrishnan, S., Ravikumar, P., 2020. Robust estimation via robust gradient estimation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 82 (3), 601–627.
- Rousseeuw, P.J., Hampel, F.R., Ronchetti, E.M., Stahel, W.A., 2011. *Robust Statistics: The Approach Based on Influence Functions*. John Wiley & Sons.
- Wright, S., Nocedal, J., 1999. Numerical Optimization. *Springer Science* 35 (67–68), 7.