

High Dimensional Robust M -Estimation: Arbitrary Corruption and Heavy Tails

Liu Liu Tianyang Li Constantine Caramanis
liuliu@utexas.edu lty@cs.utexas.edu constantine@utexas.edu

The University of Texas at Austin

Abstract

We consider the problem of sparsity-constrained M -estimation when both *explanatory and response* variables have heavy tails (bounded 4-th moments), or a fraction of arbitrary corruptions. We focus on the k -sparse, high-dimensional regime where the number of variables d and the sample size n are related through $n \sim k \log d$. We define a natural condition we call the Robust Descent Condition (RDC), and show that if a gradient estimator satisfies the RDC, then Robust Hard Thresholding (IHT using this gradient estimator), is guaranteed to obtain good statistical rates. The contribution of this paper is in showing that this RDC is a flexible enough concept to recover known results, and obtain new robustness results. Specifically, new results include: (a) For k -sparse high-dimensional linear- and logistic-regression with heavy tail (bounded 4-th moment) explanatory and response variables, a linear-time-computable median-of-means gradient estimator satisfies the RDC, and hence Robust Hard Thresholding is minimax optimal; (b) When instead of heavy tails we have $O(1/\sqrt{k} \log(nd))$ -fraction of arbitrary corruptions in explanatory and response variables, a near linear-time computable trimmed gradient estimator satisfies the RDC, and hence Robust Hard Thresholding is minimax optimal. We demonstrate the effectiveness of our approach in sparse linear, logistic regression, and sparse precision matrix estimation on synthetic and real-world US equities data.

1 Introduction

M -estimation is a standard technique for statistical estimation [vdV00]. The past decade has seen successful extensions of M -estimation to the high dimensional setting with sparsity (or other low-dimensional structure), e.g., using Lasso [Tib96, BvdG11, HTW15, Wai19]. Yet sparse modeling in high dimensions is NP-hard in the worst case [BDMS13, ZWJ14]. Thus theoretical sparse recovery guarantees for most computationally tractable approaches (e.g., ℓ_1 minimization [Don06, CRT04, Wai09], Iterative Hard Thresholding [BD09]) rely on strong assumptions on the probabilistic models of the data, such as sub-Gaussianity. Under such assumptions, these approaches achieve the minimax rate for sparse regression [RWY11].

Meanwhile, statistical estimation with heavy tailed outliers or even arbitrary corruptions has long been a focus in robust statistics [Box53, Tuk75, Hub11, HRRS11].¹ But heavy-tails and arbitrary corruptions in the data violate the assumptions required for convergence of the usual algorithms. A central question then, is what assumptions are sufficient to enable efficient and robust algorithms for high dimensional M -estimation under heavy tails or arbitrary corruption.

Huber's seminal work [Hub64] and more modern followup work [Loh17] has considered replacing the classical least squared risk minimization objective with a robust counterpart (e.g., Huber loss). Other approaches (e.g., [Li13]) considered regularization-based robustness approaches. However, when there are outliers in the explanatory variables (covariates), these approaches do not seem to succeed [CCM13]. Meanwhile, approaches combining recent advances in robust mean estimation and gradient descent have proved remarkably powerful in the low-dimensional setting [PSBR18, KKM18, DKK⁺18],

¹Following [Min18, FWZ16], by *heavy-tail* we mean satisfying only weak moment bounds, specifically, bounded 4-th order moments (compared to sub-exponential or sub-Gaussian).

but for high dimensions, have so far only managed to address the setting where the covariance of the explanatory variables is the identity, or sparse [BDLS17, LSLC18]. Meanwhile, flexible and statistically optimal approaches ([Gao17]) have relied on intractable estimators such as Tukey-depth.

For the heavy-tail setting, another line of research considers estimators such as Median-of-Means (MOM) [NY83, JVV86, AMS99, Min15] and Catoni’s mean estimator [Cat12, Min18] only use weak moment assumptions. [Min15, BJL15, HS16] generalized these ideas to M -estimation, yet it is not clear if these approaches apply to the high-dimensional setting with heavy tailed covariates.

Main Contributions. In this paper, we develop a sufficient condition that when satisfied, guarantees that an efficient algorithm (a variant of IHT) achieves the minimax optimal statistical rate. We show that our condition is flexible enough to apply to a number of important high-dimensional estimation problems under either heavy tails, or arbitrary corruption of the data. Specifically:

1. We consider two models. For our arbitrary corruption model, we assume that an adversary *replaces* an arbitrary ϵ -fraction of the authentic samples with arbitrary values (Definition 2.1). For the heavy-tailed model, we assume our data (response and covariates) satisfy only weak moment assumptions (Definition 2.2) without sub-Gaussian or sub-exponential concentration bounds.
2. We propose a notion that we call the Robust Descent Condition (RDC). Given any gradient estimator that satisfies the RDC, we define RHT – Robust Hard Thresholding (Algorithm 1) for sparsity constrained M -estimation, and prove that Algorithm 1 converges linearly to a minimax statistically optimal solution. Thus the RDC and Robust Hard Thresholding form the basis for a Deterministic Meta-Theorem (Theorem 3.1) that guarantees estimation error rates as soon as the RDC property of any gradient estimator can be certified.
3. We then obtain non-asymptotic bounds via certifying the RDC for different robust gradient estimators under various statistical models. (A) For corruptions in both response and explanatory variables, we show the trimmed gradient estimator satisfies the RDC. Thus our algorithm RHT has minimax-optimal statistical error, and tolerates $O(1/(\sqrt{k} \log(nd)))$ -fraction of outliers. This fraction is nearly independent of the d , which is important in the high dimension regime. (B) In the heavy tailed regime, we use the Median-of-Means (MOM) gradient estimator. Our RHT algorithm obtains the sharpest available error bound, in fact nearly matching the results in the sub-Gaussian case. With either of these gradient estimators, our algorithm is computationally efficient, nearly matching vanilla gradient descent. This is in particular much faster than algorithms relying on sparse PCA relaxations as subroutines ([BDLS17, LSLC18]).
4. We use Robust Hard Thresholding for neighborhood selection [MB06] for estimating Gaussian graphical models, and provide model selection guarantees under adversarial corruption of the data; our results share similar robustness guarantees with sparse regression.
5. We demonstrate the effectiveness of Robust Hard Thresholding on both arbitrarily corrupted/heavy tailed synthetic data and (unmodified) real data.

A concrete illustration of 3(B) above: Consider a sparse linear regression problem without noise (sparse linear equations), with scaling $n = O(k \log d)$. When the covariates are sub-Gaussian, Lasso succeeds in exact recovery with high probability (as expected). When the covariates have only 4-th moments, we do not expect Lasso to succeed, and indeed experiments indicate this. Moreover, to the best of our knowledge, no previous efficient algorithm with $O(k \log(d))$ samples can guarantee exact recovery in this observation model ([FWZ16] has a statistical rate depending on the norm of the parameter β^* , and thus exact recovery for $\sigma = 0$ is not guaranteed). Our contributions show that Robust Hard Thresholding using MOM achieves this (see also simulations in Figure 1(b)).

Related work

Sparse regression with arbitrary corruptions or heavy tails. Several works in robustness of high dimensional problems consider heavy tailed distributions or arbitrary corruptions only in the

response variables [Li13, BJK15, BJK17, Loh17, KP18, HS16, Min15, CLZL]. Yet these algorithms cannot be trivially extended to the setting with heavy tails or corruptions in explanatory variables. Another line [ACG13, VMX17, YLA18, SS18] focuses on alternating minimization approaches which extend Least Trimmed Squares [Rou84]. However, these methods only have local convergence guarantees, and cannot handle arbitrary corruptions.

[CCM13] was one of the first papers to provide guarantees for sparse regression with arbitrary outliers in both response and explanatory variables by trimming the design matrix. Similar trimming techniques are also used in [FWZ16] for heavy tails in response and explanatory variables. Those results are specific to sparse regression, however, and cannot be easily extended to general M -estimation problems. Moreover, even for linear regression, the statistical rates are not minimax optimal. [LM16] uses Median-of-Means tournaments to deal with heavy tails in the explanatory variables and obtains near optimal rates. However, Median-of-Means tournaments is not known to be computationally tractable. [LL17] deals with heavy tails and outliers in the explanatory variables, but they require higher moment bound (whose order is $O(\log(d))$) in the isotropic design case. [Gao17] optimizes Tukey depth [Tuk75, CGR18] for robust sparse regression under the Huber ϵ -contamination model, and their algorithm is minimax optimal and can handle a constant fraction of outliers. However, computing Tukey depth is intractable [JP78]. Recent results [BDLS17, LSLC18] leverage robust sparse mean estimation in robust sparse regression. Their algorithms are computationally tractable, and can tolerate $\epsilon = \text{const.}$, but they require very restrictive assumptions on the covariance matrix ($\Sigma = I_d$ or sparse), which precludes their use in applications such as graphical model estimation.

Robust M -estimation via robust gradient descent. Works in [CSX17, HI17] and later [YCRB18a] first leveraged the idea of using robust mean estimation in each step of gradient descent, using a subroutine such as geometric median. A similar approach using more sophisticated robust mean estimation methods was later proposed in [PSBR18, DKK⁺18, YCRB18b, SX18, Hol18] for robust gradient descent. These methods all focused on low dimensional robust M -estimation. Work in [LSLC18] extended the approach to the high-dimensional setting (though is limited to $\Sigma = I_d$ or sparse covariances). Even though the corrupted fraction ϵ can be independent of the ambient dimension d by using sophisticated robust mean estimation algorithms [DKK⁺16, LRV16, SCV17], or the sum-of-squares framework [KKM18], these algorithms (except [LSLC18]) are not applicable to the high dimensional setting ($n \ll d$), as they require at least $\Omega(d)$ samples.

Robust estimation of graphical models. A line of research using a robustified covariance matrix in Gaussian graphical models [LHY⁺12b, WG17, LT18] leverages GLasso [FHT08] or CLIME [CLL11] to estimate the sparse precision matrix. These robust methods are restricted to Gaussian graphical model estimation, and their techniques cannot be generalized to other M -estimation problems.

Notation. We denote the Hard Thresholding operator of sparsity k' by $P_{k'}$, and denote the Euclidean projection onto the ℓ_2 ball B by Π_B . We use $\mathbb{E}_{i \in_u \mathcal{S}}$ to denote the expectation operator obtained by the uniform distribution over all samples $\{i \in \mathcal{S}\}$.

2 Problem formulation

We now define the corruption and heavy tails model and sparsity constrained M -estimation.

Definition 2.1 (ϵ -corrupted samples). *Let $\{\mathbf{z}_i, i \in \mathcal{G}\}$ be i.i.d. observations with distribution P . We say that a collection of samples $\{\mathbf{z}_i, i \in \mathcal{S}\}$ is ϵ -corrupted if an adversary chooses an arbitrary ϵ -fraction of the samples in \mathcal{G} and modifies them with arbitrary values.*

This corruption model allows corruptions in both *explanatory and response* variables in regression problems where we observe $\mathbf{z}_i = (y_i, \mathbf{x}_i)$. Definition 2.1 also allows the adversary to select an ϵ -fraction of samples to *delete and corrupt*.

Definition 2.2 (heavy-tailed samples). *For a distribution P of $\mathbf{x} \in \mathbb{R}^d$ with mean $\mathbb{E}(\mathbf{x})$ and covariance Σ , we say that P has bounded $2k$ -th moment, if there is a universal constant C_{2k} such that, for a unit vector $\mathbf{v} \in \mathbb{R}^d$, we have $\mathbb{E}_P|\langle \mathbf{v}, \mathbf{x} - \mathbb{E}(\mathbf{x}) \rangle|^{2k} \leq C_{2k} \mathbb{E}_P(|\langle \mathbf{v}, \mathbf{x} - \mathbb{E}(\mathbf{x}) \rangle|^2)^k$.*

Definition 2.2 allows heavy tails in both *explanatory and response* variables for $\mathbf{z}_i = (y_i, \mathbf{x}_i)$. For example, in Corollary 4.3, we study linear regression with bounded 4-th moments for \mathbf{x} and bounded variance for y and noise.

Let $\ell : \mathbb{R}^d \times \mathcal{Z} \rightarrow \mathbb{R}$ be a convex and differentiable loss function. Our target is the unknown sparse population minimizer $\beta^* = \arg \min_{\beta \in \mathbb{R}^d, \|\beta\|_0 \leq k} \mathbb{E}_{\mathbf{z}_i \sim P} \ell_i(\beta; \mathbf{z}_i)$, and we write f as the population risk, $f(\beta) = \mathbb{E}_{\mathbf{z}_i \sim P} \ell_i(\beta; \mathbf{z}_i)$. Note that β^* 's definition allows model misspecification. The following Definition 2.3 provides general assumptions for the population risk.

Definition 2.3 (Strong convexity/smoothness). *For the population risk f , we assume $\mu_\alpha \|\beta_1 - \beta_2\|_2^2/2 \leq f(\beta_1) - f(\beta_2) - \langle \nabla f(\beta_2), \beta_1 - \beta_2 \rangle \leq \mu_L \|\beta_1 - \beta_2\|_2^2/2$, where μ_α is the strong-convexity parameter and μ_L is the smoothness parameter. The condition number is $\rho = \mu_L/\mu_\alpha \geq 1$.*

A well known result [NRWY12] considers ERM with convex relaxation from $\|\beta\|_0$ to $\|\beta\|_1$, by certifying the RSC condition for sub-Gaussian ensembles – this obtains uniform convergence of the empirical risk. From an optimization viewpoint, existing results reveal that gradient descent algorithms equipped with soft-thresholding [ANW12] or hard-thresholding [BD09, JTK14, SL17, YLZ18, LB18] have linear convergence rate, and achieve known minimax lower bounds in statistical estimation [RWY11, ZWJ14].

Given samples \mathcal{S} , running ERM on the entire input dataset: $\min_{\beta \in B, \|\beta\|_0 \leq k} \mathbb{E}_{i \in \mathcal{S}} \ell_i(\beta; \mathbf{z}_i)$, cannot guarantee uniform convergence of the empirical risk, and can be arbitrarily bad for ϵ -corrupted samples. The next two sections outline the main results of this paper, addressing this problem.

3 Robust sparse estimation via Robust Hard Thresholding

We introduce our meta-algorithm, Robust Hard Thresholding, that essentially uses a robust gradient estimator to run IHT. We require several definitions to specify the algorithm, and describe its results. We use $\widehat{\mathbf{G}}(\beta)$ as a placeholder for the estimate at β , obtained from whichever robust gradient estimator we are using. Let $\mathbf{G}(\beta) = \mathbb{E}_{\mathbf{z}_i \sim P} \nabla \ell_i(\beta; \mathbf{z}_i)$ denote the population gradient. We use $\widehat{\mathbf{G}}$ and \mathbf{G} when the context is clear.

Many previous works ([CSX17, HI17, PSBR18, DKK⁺18, YCRB18a, YCRB18b, SX18]) have provided algorithms for obtaining robust gradient estimators, then used as subroutines in robust gradient algorithms. However, those results require controlling $\|\widehat{\mathbf{G}} - \mathbf{G}\|_2$, and do not readily extend to high dimensions, as sufficiently controlling $\|\widehat{\mathbf{G}} - \mathbf{G}\|_2$ seems to require $n = \Omega(d)$. A recent work [LSLC18] on robust sparse linear regression uses a robust sparse mean estimator [BDLS17] to guarantee $\|\widehat{\mathbf{G}} - \mathbf{G}\|_2 = O(\delta_1 \|\beta - \beta^*\|_2 + \delta_2)$ with sample complexity $\Omega(k^2 \log(d))$. However, their algorithm requires the restrictive assumption $\Sigma = I_d$ or sparse, and thus cannot be extended to more general M -estimation problems.

To address this issue, we propose Robust Hard Thresholding (Algorithm 1), which uses hard thresholding after each robust gradient update². In line 7, we use a gradient estimator to obtain the robust gradient estimate $\widehat{\mathbf{G}}^t$. In line 8, we update the parameter by hard thresholding $\bar{\beta}^{t+1} = P_{k'}(\beta^t - \eta \widehat{\mathbf{G}}^t)$, where the hyper-parameter k' proportional to k is specified in Definition 2.3. A key observation in line 8 is that, in each step of IHT, the iterate β^t is sparse, and thus the *perturbation from outliers or heavy tails* only depends on IHT's sparsity k' instead of the ambient dimension d . Based on a careful analysis of the hard thresholding operator in each iteration, we show that rather than controlling $\|\widehat{\mathbf{G}} - \mathbf{G}\|_2$, it is enough to control a weaker quantity: this is what we call the Robust

²Our theory requires splitting samples across different iterations to maintain independence between iterations. We believe this is an artifact of the analysis, and do not use this in our experiments. [BWY17, PSBR18] use a similar approach for theoretical analysis.

Algorithm 1 Robust Hard Thresholding

- 1: **Input:** Data samples $\{y_i, \mathbf{x}_i\}_{i=1}^N$, gradient estimator $\widehat{\mathbf{G}}$.
 - 2: **Output:** The estimation $\widehat{\boldsymbol{\beta}}$.
 - 3: **Parameters:** Hard thresholding parameter $k' = 4\rho^2k$.
-
- 4: Split samples into T subsets each of size n . Initialize with $\boldsymbol{\beta}^0 = \mathbf{0}_d$.
 - 5: **for** $t = 0$ to $T - 1$, **do**
 - 6: At current $\boldsymbol{\beta}^t$, calculate all gradients for current n samples: $\mathbf{g}_i^t = \nabla \ell_i(\boldsymbol{\beta}^t)$, $i \in [n]$.
 - 7: For $\{\mathbf{g}_i^t\}_{i=1}^n$, we obtain $\widehat{\mathbf{G}}^t$
 - 8: Update the parameter: $\overline{\boldsymbol{\beta}}^{t+1} = P_{k'}(\boldsymbol{\beta}^t - \eta \widehat{\mathbf{G}}^t)$. Then project: $\boldsymbol{\beta}^{t+1} = \Pi_B(\overline{\boldsymbol{\beta}}^{t+1})$.
 - 9: **end for**
 - 10: Output the estimation $\widehat{\boldsymbol{\beta}} = \boldsymbol{\beta}^T$.
-

Descent Condition [Definition 3.1](#) and we define it next; it plays a key role in obtaining *sharp* rates of convergence for various types of statistical models.

Robust Descent Condition

The Robust Descent Condition [eq. \(1\)](#) provides an upper bound on the inner product of the robust gradient estimate and the distance to the population optimum. This is a natural notion to control the potential progress obtained by using a robust gradient update instead of the population gradient.

Definition 3.1 ((α, ψ) -Robust Descent Condition (RDC)). *For the population gradient \mathbf{G} at $\boldsymbol{\beta}$, a robust gradient estimator $\widehat{\mathbf{G}}(\boldsymbol{\beta})$ satisfies the robust descent condition if for any sparse $\boldsymbol{\beta}, \widetilde{\boldsymbol{\beta}} \in \mathbb{R}^d$,*

$$\left| \left\langle \widehat{\mathbf{G}}(\boldsymbol{\beta}) - \mathbf{G}(\boldsymbol{\beta}), \widetilde{\boldsymbol{\beta}} - \boldsymbol{\beta}^* \right\rangle \right| \leq \left(\alpha \|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|_2 + \psi \right) \left\| \widetilde{\boldsymbol{\beta}} - \boldsymbol{\beta}^* \right\|_2. \quad (1)$$

We begin with a Meta-Theorem for [Algorithm 1](#) that holds under the Robust Descent Condition [Definition 3.1](#) and assumptions on population risk [Definition 2.3](#). In [Theorem 3.1](#), we prove [Algorithm 1](#)'s global convergence and its statistical guarantees. The proofs are collected in [Appendix B](#).

Theorem 3.1 (Meta-Theorem). *Suppose we observe samples from a statistical model with population risk f satisfying [Definition 2.3](#). If a robust gradient estimator satisfies (α, ψ) -Robust Descent Condition ([Definition 3.1](#)) where $\alpha \leq \frac{1}{32}\mu_\alpha$, then [Algorithm 1](#) with $\eta = 1/\mu_L$ outputs $\widehat{\boldsymbol{\beta}}$ such that $\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2 = O(\psi/\mu_\alpha)$, by setting $T = O(\rho \log(\mu_\alpha \|\boldsymbol{\beta}^*\|_2/\psi))$.*

We note that [Theorem 3.1](#) is deterministic in nature. In the sequel, we omit the log term in the sample complexity due to sample splitting. We obtain high probability results via certifying that the RDC holds for certain robust gradient estimators under various statistical models. To obtain the minimax estimation error rate in [Theorem 3.1](#), the key step is providing a robust gradient estimator with sufficiently small ψ , in the definition of RDC.

[Section 4](#) uses the RDC and [Theorem 3.1](#) to obtain new results for sparse regression under heavy tails or arbitrary corruption. Before we move to this, we observe that we can use the RDC and [Theorem 3.1](#) to recover existing results in the literature. Some immediate examples are as follows:

Uncorrupted gradient satisfies the RDC. Suppose the samples follow from sparse linear regression with sub-Gaussian covariates and noise $\mathcal{N}(0, \sigma^2)$. The empirical average of gradient samples satisfies [eq. \(1\)](#) with $\psi = O(\sigma \sqrt{k \log(d)/n})$, by assuming ℓ_1 constraint on $\boldsymbol{\beta}$ and $\widetilde{\boldsymbol{\beta}}$ [[LW11](#)]. Plugging in this ψ to [Theorem 3.1](#) recovers the well-known minimax rates for sparse linear regression [[RWY11](#)].

RSGE implies RDC. When $\boldsymbol{\Sigma} = \mathbf{I}_d$ or is sparse, [[BDLS17](#)] and [[LSLC18](#)], respectively, provide robust sparse gradient estimators (RSGE) which upper bound $\|\widehat{\mathbf{G}}(\boldsymbol{\beta}) - \mathbf{G}(\boldsymbol{\beta})\|_2 \leq \alpha \|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|_2 + \psi$, for a constant fraction ϵ of corrupted samples. Noting that $|\langle \widehat{\mathbf{G}}(\boldsymbol{\beta}) - \mathbf{G}(\boldsymbol{\beta}), \widetilde{\boldsymbol{\beta}} - \boldsymbol{\beta}^* \rangle| \leq \|\widehat{\mathbf{G}}(\boldsymbol{\beta}) - \mathbf{G}(\boldsymbol{\beta})\|_2 \|\widetilde{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2$

$\mathbf{G}(\boldsymbol{\beta})\|_2\|\tilde{\boldsymbol{\beta}}-\boldsymbol{\beta}^*\|_2$, we observe that RSGE implies RDC. Hence any RSGE can be used in [Algorithm 1](#). The RSGE for $\boldsymbol{\Sigma} = I$ in [\[BDLS17\]](#) guarantees an RDC with $\psi = O(\sigma\epsilon)$ when $n = \Omega(k^2 \log d/\epsilon^2)$, and the RSGE for unknown sparse $\boldsymbol{\Sigma}$ from [\[LSLC18\]](#) guarantees $\psi = O(\sigma\sqrt{\epsilon})$ when $n = \Omega(k^2 \log d/\epsilon)$. Again plugging these values for ψ into our theorem, recovers the results in those papers.³

4 Main Results: Using the RDC and [Algorithm 1](#)

In the remainder of our paper, we use [Theorem 3.1](#) and the RDC to analyze two well-known and computationally efficient robust mean estimation subroutines that have been used in the low-dimensional setting: the trimmed mean estimator and the MOM estimator. We show that these two can obtain a sufficiently small ψ in the definition of the RDC. This leads to the minimax estimation error in the case of arbitrary corruptions or heavy tails.

4.1 Gradient estimation

The trimmed mean and MOM estimators have been successfully applied to robustify gradient descent [\[YCRB18a, PSBR18\]](#) in the low dimensional setting. They have not been used in the high dimensional regime, however, because until now we have not had the machinery to analyze their algorithmic convergence, statistical rates and minimax optimality in the high dimensional setting.

To show they satisfy the RDC with a sufficiently small ψ , we observe that by using Hölder’s inequality on the LHS of [eq. \(1\)](#), we have $|\langle \widehat{\mathbf{G}}(\boldsymbol{\beta}) - \mathbf{G}(\boldsymbol{\beta}), \tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}^* \rangle| \leq \|\widehat{\mathbf{G}}(\boldsymbol{\beta}) - \mathbf{G}(\boldsymbol{\beta})\|_\infty \|\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_1$. Using [Algorithm 1](#), the Hard Thresholding step enforces sparsity of $\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}^*$. Therefore, controlling ψ amounts to bounding the infinity norm of the robust gradient estimate.

In [Section 4.2](#), we show that by using coordinate-wise robust mean estimation, we can certify the RDC with sufficiently small ψ to guarantee minimax rates. Specifically, we show this for the trimmed gradient estimator for arbitrary corruption, and the MOM gradient estimator for heavy tailed distributions.

Definition 4.1. Given gradients samples $\{\nabla\ell_i(\boldsymbol{\beta}; \mathbf{z}_i) \in \mathbb{R}^d, i \in \mathcal{S}\}$, for each dimension $j \in [d]$,

(♠): **Trimmed gradient estimator** removes the largest and smallest α fraction of elements in $\{\nabla\ell_i(\boldsymbol{\beta}; \mathbf{z}_i)_j \in \mathbb{R}, i \in \mathcal{S}\}$, and calculates the mean of the remaining terms. We choose $\alpha = c_0\epsilon$ for constant $c_0 \geq 1$, and require $\alpha \leq 1/2 - c_1$ for a small constant $c_1 > 0$.

(♣): **MOM gradient estimator** partitions \mathcal{S} into $4.5\lceil\log(d)\rceil$ blocks and computes the sample mean of $\{\nabla\ell_i(\boldsymbol{\beta}; \mathbf{z}_i)_j \in \mathbb{R}\}$ within each block, and then take the median of these means.⁴

4.2 Statistical guarantees

In this section, we consider some typical models for general M -estimation.

Model 4.1 (Sparse linear regression). Samples $\mathbf{z}_i = (y_i, \mathbf{x}_i)$ are drawn from a linear model P : $y_i = \mathbf{x}_i^\top \boldsymbol{\beta}^* + \xi_i$, with $\boldsymbol{\beta}^* \in \mathbb{R}^d$ being k -sparse. We assume that \mathbf{x} ’s are i.i.d. with normalized covariance matrix $\boldsymbol{\Sigma}$, with $\boldsymbol{\Sigma}_{jj} \leq 1 \forall j$, and the stochastic noise ξ has mean 0 and variance σ^2 .

Model 4.2 (Sparse logistic regression). Samples $\mathbf{z}_i = (y_i, \mathbf{x}_i)$ are drawn from a binary classification model P , where the binary label $y_i \in \{-1, +1\}$ follows the conditional probability distribution $\Pr(y_i|\mathbf{x}_i) = 1/(1 + \exp(-y_i\mathbf{x}_i^\top \boldsymbol{\beta}^*))$, with $\boldsymbol{\beta}^* \in B \subset \mathbb{R}^d$ being k -sparse. We assume that \mathbf{x} ’s are i.i.d. with normalized covariance matrix $\boldsymbol{\Sigma}$, where $\boldsymbol{\Sigma}_{jj} \leq 1$ for all j .

To obtain the following corollaries, we first certify the RDC for a certain robust gradient estimator over random ensembles with corruption or heavy tails, and then use them in [Theorem 3.1](#). We collect the results for gradient estimation in [Appendix A](#), and the proofs for corollaries in [Appendix B](#).

³It remains an open question to obtain a RSGE for a constant fraction of outliers for robust sparse regression with arbitrary covariance $\boldsymbol{\Sigma}$.

⁴Without loss of generality, we assume the number of blocks divides n , and $4.5\lceil\log(d)\rceil$ is chosen in [\[HS16\]](#).

Arbitrary corruption case. Based on [Theorem 3.1](#), we first provide concrete results for arbitrary corruption case [Definition 2.1](#), where the covariates and response variables in the authentic distribution P are assumed to be sub-Gaussian.

Corollary 4.1. *Suppose we observe n ϵ -corrupted ([Definition 2.1](#)) sub-Gaussian samples from sparse linear regression model ([Model 4.1](#)). Under the condition $n = \Omega(\rho^4 k \log d)$, and $\epsilon = O\left(\frac{1}{\rho^2 \sqrt{k} \log(nd)}\right)$, with probability at least $1 - d^{-2}$, [Algorithm 1](#) with trimmed gradient estimator satisfies the RDC with $\psi = O(\rho \sigma \sqrt{k}(\epsilon \log(nd) + \sqrt{\log d/n}))$, and thus [Theorem 3.1](#) provides $\|\hat{\beta} - \beta^*\|_2 = O(\rho^2 \sigma(\epsilon \sqrt{k} \log(nd) + \sqrt{k \log d/n}))$.*

Time complexity. [Corollary 4.1](#) has a global linear convergence rate. In each iteration, we only use $O(nd \log n)$ operations complexity to calculate trimmed mean. We incur logarithmic overhead compared to normal gradient descent [[Bub15](#)].

Statistical accuracy and robustness. Compared with [[CCM13](#), [BDLS17](#)], our statistical error rate is minimax optimal [[RWY11](#), [ZWJ14](#)], and has no dependencies on $\|\beta^*\|_2$. Furthermore, the upper bound on ϵ is nearly independent of d , which guarantees [Algorithm 1](#)'s robustness in high dimensions.

Corollary 4.2. *Suppose we observe n ϵ -corrupted ([Definition 2.1](#)) sub-Gaussian samples from sparse logistic regression model ([Model 4.2](#)). With probability at least $1 - d^{-2}$, [Algorithm 1](#) with trimmed gradient estimator satisfies the RDC with $\psi = O(\rho \sqrt{k}(\epsilon \log(nd) + \sqrt{\log d/n}))$, and thus [Theorem 3.1](#) provides $\|\hat{\beta} - \beta^*\|_2 = O(\rho^2(\epsilon \sqrt{k} \log(nd) + \sqrt{k \log d/n}))$.*

Statistical accuracy and robustness. Under the sparse Gaussian linear discriminant analysis model (a typical example of [Model 4.2](#)), [Algorithm 1](#) achieves the statistical minimax rate [[LPR15](#), [LYCR17](#)].

Heavy-tailed distribution case. We next turn to the heavy tailed distribution case [Definition 2.2](#).

Corollary 4.3. *Suppose we observe n samples from sparse linear regression model ([Model 4.2](#)) with bounded 4-th moment covariates. Under the condition $n = \Omega(\rho^6 k \log d)$, with probability at least $1 - d^{-2}$, [Algorithm 1](#) with MOM gradient estimator satisfies the RDC with $\psi = O(\rho^{3/2} \sigma \sqrt{k \log d/n})$, and thus [Theorem 3.1](#) provides $\|\hat{\beta} - \beta^*\|_2 = O(\rho^{5/2} \sigma \sqrt{k \log d/n})$.*

Time complexity. Similar to [Corollary 4.1](#), [Corollary 4.3](#) has a global linear convergence. In each iteration, we only use $O(nd)$ operations complexity – the same as normal gradient descent [[Bub15](#)].

Statistical accuracy. [[LM16](#)] uses Median-of-Means tournaments to deal with sparse linear regression with bounded moment assumptions for the covariates, and they obtain near optimal rates. We obtain similar rates, however our algorithm is efficient, where as Median-of-Means tournaments is not known to be computationally tractable. [[FWZ16](#), [Zhu17](#)] deal with the same problem by truncating and shrinking the data to certify the RSC condition. Their results require boundedness of higher moments of the noise ξ , and the final error depends on $\|\beta^*\|_2$. Our estimation error bounds exactly recover optimal sub-Gaussian bounds for sparse regression [[NRWY12](#), [Wai19](#)], and moreover, we obtain exact recovery when ξ 's variance $\sigma^2 \rightarrow 0$.

Corollary 4.4. *Suppose we observe n samples from sparse logistic regression model ([Model 4.2](#)). With probability at least $1 - d^{-2}$, [Algorithm 1](#) with MOM gradient estimator satisfies the RDC with $\psi = O(\rho^{3/2} \sqrt{k \log d/n})$, and thus [Theorem 3.1](#) provides $\|\hat{\beta} - \beta^*\|_2 = O(\rho^{5/2} \sqrt{k \log d/n})$.*

4.3 Sparsity recovery and Gaussian graphical model estimation

We next demonstrate the sparsity recovery performance of [Algorithm 1](#) for graphical model learning [[MB06](#), [Wai09](#), [RWL10](#), [RWRY11](#), [BvdG11](#), [HTW15](#)]. Our sparsity recovery guarantees hold for both heavy tails and arbitrary corruption, though we only present results in the case of arbitrary corruption in this section.

We use $\text{supp}(\mathbf{v}, k)$ to denote top k indexes of \mathbf{v} with the largest magnitude. Let \mathbf{v}_{\min} denote the smallest absolute value of nonzero element of \mathbf{v} . To control the false negative rate, [Corollary 4.5](#) shows that under the β_{\min} -condition, $\text{supp}(\widehat{\beta}, k)$ is exactly $\text{supp}(\beta^*)$. The proofs are given in [Appendix C](#). Sparsity recovery guarantee for sparse logistic regression is similar, and is omitted due to space constraints. Existing results on sparsity recovery for ℓ_1 regularized estimators [[Wai09](#), [LSRC15](#)] do not require the RSC condition, but instead require an irrepresentability condition, which is stronger. If $\epsilon \rightarrow 0$, [Corollary 4.5](#) has the same β_{\min} -condition as IHT for sparsity recovery [[YLZ18](#)].

Corollary 4.5. *Under the same condition as in [Corollary 4.1](#), and a β_{\min} -condition on β^* , $\beta_{\min}^* = \Omega(\rho^2 \sigma (\epsilon \sqrt{k} \log(nd) + \sqrt{k} \log d/n))$, [Algorithm 1](#) with trimmed gradient estimator guarantees that $\text{supp}(\widehat{\beta}, k) = \text{supp}(\beta^*)$, with probability at least $1 - d^{-2}$.*

We consider sparse precision matrix estimation for Gaussian graphical models. The sparsity pattern of its precision matrix $\Theta = \Sigma^{-1}$ matches the conditional independence relationships [[KFB09](#), [WJ08](#)].

Model 4.3 (Sparse precision matrix estimation). *Under the contamination model [Definition 2.1](#), authentic samples $\{\mathbf{x}_i\}_{i=1}^m$ are drawn from a multivariate Gaussian distribution $\mathcal{N}(0, \Sigma)$. We assume that each row of the precision matrix $\Theta = \Sigma^{-1}$ is $(k+1)$ -sparse – each node has at most k edges.*

For the uncorrupted samples drawn from the Gaussian graphical model, the neighborhood selection (NS) algorithm [[MB06](#)] solves a convex relaxation of the following sparsity constrained optimization to regress each variable against its neighbors

$$\widehat{\beta}_j = \arg \min_{\beta \in \mathbb{R}^{d-1}} \frac{1}{m} \sum_{i=1}^m (x_{ij} - \mathbf{x}_{i(j)}^\top \beta)^2, \quad \text{s.t. } \|\beta\|_0 \leq k, \quad \text{for each } j \in [d], \quad (2)$$

where x_{ij} denotes the j -th coordinate of $x_i \in \mathbb{R}^d$, and (j) denotes the index set $\{1, \dots, j-1, j+1, \dots, d\}$. Let $\boldsymbol{\theta}_{(j)} \in \mathbb{R}^{d-1}$ denote Θ 's j -th column with the diagonal entry removed. and $\Theta_{j,j} \in \mathbb{R}$ denote the j -th diagonal element of Θ . Then, the sparsity pattern of $\boldsymbol{\theta}_{(j)}$ can be estimated through $\widehat{\beta}_j$. Details on the connection between $\boldsymbol{\theta}_{(j)}$ and $\widehat{\beta}_j$ are given in [Appendix C](#).

However, given ϵ -corrupted samples from the Gaussian graphical model, this procedure will fail [[LHY+12b](#), [WG17](#)]. To address this issue, we propose Robust NS ([Algorithm 2](#) in [Appendix C](#)), which robustifies Neighborhood Selection [[MB06](#)] by using Robust Hard Thresholding (with least square loss) to robustify [eq. \(2\)](#). Similar to [Corollary 4.5](#), a $\boldsymbol{\theta}_{\min}$ -condition guarantees consistent edge selection.

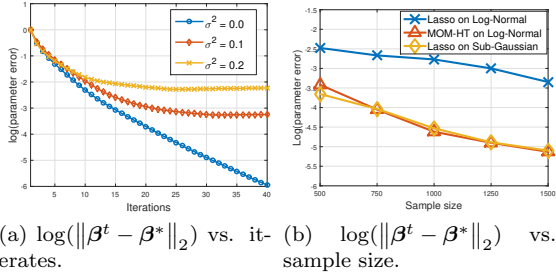
Corollary 4.6. *Under the same condition as in [Corollary 4.1](#), and a $\boldsymbol{\theta}_{\min}$ -condition for $\boldsymbol{\theta}_{(j)}$, $\boldsymbol{\theta}_{(j), \min} = \Omega(\Theta_{j,j}^{1/2} \rho^2 (\epsilon \sqrt{k} \log(nd) + \sqrt{k} \log d/n))$, Robust NS ([Algorithm 2](#)) achieves consistent edge selection, with probability at least $1 - d^{-1}$.*

Similar to [Corollary 4.1](#), the fraction ϵ is nearly independent of dimension d , which provides guarantees of Robust NS in high dimensions. Other Gaussian graphical model selection algorithms include GLasso [[FHT08](#)], CLIME [[CLL11](#)]. The experimental details comparing robustified versions of these algorithms are presented in [Appendix D.4](#).

5 Experiments


We provide the complete details for our experiment setup in [Appendix D](#).

Sparse regression with arbitrary corruption. We generate samples from a sparse regression model ([Model 4.1](#)) with a Toeplitz covariance Σ . Here, the stochastic noise $\xi \sim \mathcal{N}(0, \sigma^2)$, and we vary the noise level σ^2 in different simulations. We add outliers with $\epsilon = 0.1$, and track the parameter error $\|\beta^t - \beta^*\|_2$ in each iteration. Left plot of [Figure 1](#) shows [Algorithm 1](#)'s linear convergence, and



(a) $\log(\|\beta^t - \beta^*\|_2)$ vs. iterations. (b) $\log(\|\beta^t - \beta^*\|_2)$ vs. sample size.

Figure 1: In the left plot, the corruption level ϵ is fixed and we use Algorithm 1 with trimming for different noise level σ^2 . In the right plot, we consider log-normal samples, and we use Algorithm 1 with MOM for different sample size to compare with baselines (Lasso on heavy tailed data, and Lasso on sub-Gaussian data).



(a) Graph estimated by Algorithm 2. (b) Graph estimated by Vanilla NS approach.

Figure 2: Graph estimated from the S&P 500 stock data by Algorithm 2 and Vanilla NS approach. Variables are colored according to their sectors. In particular, the stocks from sector Information Technology are colored as purple.

the error curves flatten out at the final error level. Furthermore, Algorithm 1 can achieve machine precision when $\sigma^2 = 0$, which means exactly recovering of β^* .

Sparse regression with heavy tails. We consider a log-normal distribution (a typical example of heavy tails) in Model 4.1. More specifically, $x_i = \sqrt{\Sigma} \tilde{x}_i$, and $\xi_i = \sigma \tilde{\xi}_i$. Here, Σ is the same Toeplitz covariance, each entry of \tilde{x}_i and $\tilde{\xi}_i$ follows from $(Z - \mathbb{E}Z)/\sqrt{\text{Var}(Z)}$, where $Z \sim \log \mathcal{N}(0, 4)$. We fix k, d, σ , and vary sample size n . For log-normal samples, we run Algorithm 1 with MOM and vanilla Lasso. We then re-generate standard Gaussian samples using the same dimensions with Σ and run Vanilla Lasso. Each curve in the right plot of Figure 1 is the average of 50 trials. Algorithm 1 with MOM significantly improves vanilla Lasso on log-normal data, and has the same performance as Lasso on sub-Gaussian data

Real data experiments. We next apply Algorithm 2, to a US equities dataset [LHY⁺12a, ZLR⁺12], which is heavy-tailed and has many outliers [dP18]. The dataset contains 1,257 daily closing prices of 452 stocks (variables). It is well known that stocks from the same sector tend to be clustered together [Kin66]. Therefore, we use Robust NS (Algorithm 2) to construct an undirected graph among stocks. Graphs estimated by different algorithms are shown in Figure 2. We can see that stocks from the same sector are clustered together, and these clustering centers can be easily identified. We also compare Algorithm 2 to the baseline NS approach (as in the ideal setting). We can observe that stocks from Information Technology (colored by purple) are much better clustered by Algorithm 2.

References

- [ACG13] Andreas Alfons, Christophe Croux, and Sarah Gelper. Sparse least trimmed squares regression for analyzing high-dimensional large data sets. *The Annals of Applied Statistics*, pages 226–248, 2013.
- [AMS99] Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. *Journal of Computer and system sciences*, 58(1):137–147, 1999.
- [ANW12] Alekh Agarwal, Sahand Negahban, and Martin J. Wainwright. Fast global convergence of gradient methods for high-dimensional statistical recovery. *Ann. Statist.*, 40(5):2452–2482, 10 2012.
- [BD09] Thomas Blumensath and Mike E Davies. Iterative hard thresholding for compressed sensing. *Applied and computational harmonic analysis*, 27(3):265–274, 2009.
- [BDLS17] Sivaraman Balakrishnan, Simon S. Du, Jerry Li, and Aarti Singh. Computationally efficient robust sparse estimation in high dimensions. In *Proceedings of the 2017 Conference on Learning Theory*, 2017.
- [BDMS13] Afonso S. Bandeira, Edgar Dobriban, Dustin G. Mixon, and William F. Sawin. Certifying the restricted isometry property is hard. *IEEE Transactions on Information Theory*, 59(6):3448–3450, 2013.
- [BJK15] Kush Bhatia, Prateek Jain, and Purushottam Kar. Robust regression via hard thresholding. In *Advances in Neural Information Processing Systems*, pages 721–729, 2015.
- [BJK17] Kush Bhatia, Prateek Jain, and Purushottam Kar. Consistent robust regression. In *Advances in Neural Information Processing Systems*, pages 2107–2116, 2017.
- [BJL15] Christian T Brownlees, Emilien Joly, and Gabor Lugosi. Empirical risk minimization for heavy-tailed losses. *Annals of Statistics*, 43(6):2507–2536, 2015.
- [Box53] George EP Box. Non-normality and tests on variances. *Biometrika*, 40(3/4):318–335, 1953.
- [Bub15] Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- [BvdG11] Peter Bühlmann and Sara van de Geer. *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media, 2011.
- [BWY17] Sivaraman Balakrishnan, Martin J Wainwright, and Bin Yu. Statistical guarantees for the em algorithm: From population to sample-based analysis. *The Annals of Statistics*, 45(1):77–120, 2017.
- [Cat12] Olivier Catoni. Challenging the empirical mean and empirical variance: a deviation study. In *Annales de l’IHP Probabilités et statistiques*, volume 48, pages 1148–1185, 2012.
- [CCM13] Yudong Chen, Constantine Caramanis, and Shie Mannor. Robust sparse regression under adversarial corruption. In *International Conference on Machine Learning*, pages 774–782, 2013.
- [CGR18] Mengjie Chen, Chao Gao, and Zhao Ren. Robust covariance and scatter matrix estimation under hubers contamination model. *Ann. Statist.*, 46(5):1932–1960, 10 2018.
- [CLL11] Tony Cai, Weidong Liu, and Xi Luo. A constrained ℓ_1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association*, 106(494):594–607, 2011.
- [CLZL] Yuejie Chi, Yuanxin Li, Huishuai Zhang, and Yingbin Liang. Median-truncated gradient descent: A robust and scalable nonconvex approach for signal estimation.
- [CRT04] Emmanuel J. Candès, Justin K. Romberg, and Terence Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52:489–509, 2004.
- [CSX17] Yudong Chen, Lili Su, and Jiaming Xu. Distributed statistical machine learning in adversarial settings: Byzantine gradient descent. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 1(2):44, 2017.
- [DKK⁺16] Ilias Diakonikolas, Gautam Kamath, Daniel M Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robust estimators in high dimensions without the computational intractability. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 655–664. IEEE, 2016.

- [DKK⁺18] Ilias Diakonikolas, Gautam Kamath, Daniel M Kane, Jerry Li, Jacob Steinhardt, and Alistair Stewart. Sever: A robust meta-algorithm for stochastic optimization. *arXiv preprint arXiv:1803.02815*, 2018.
- [Don06] David L Donoho. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006.
- [dP18] M. L. de Prado. *Advances in Financial Machine Learning*. Wiley, 2018.
- [FHT08] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008.
- [FWZ16] Jianqing Fan, Weichen Wang, and Ziwei Zhu. A shrinkage principle for heavy-tailed data: High-dimensional robust low-rank matrix recovery. *arXiv preprint arXiv:1603.08315*, 2016.
- [Gao17] Chao Gao. Robust regression via multivariate regression depth. *arXiv preprint arXiv:1702.04656*, 2017.
- [HI17] Matthew J Holland and Kazushi Ikeda. Efficient learning with robust gradient descent. *arXiv preprint arXiv:1706.00182*, 2017.
- [Hol18] Matthew J Holland. Robust descent using smoothed multiplicative noise. *arXiv preprint arXiv:1810.06207*, 2018.
- [HRRS11] Frank R Hampel, Elvezio M Ronchetti, Peter J Rousseeuw, and Werner A Stahel. *Robust statistics: the approach based on influence functions*, volume 196. John Wiley & Sons, 2011.
- [HS16] Daniel Hsu and Sivan Sabato. Loss minimization and parameter estimation with heavy tails. *The Journal of Machine Learning Research*, 17(1):543–582, 2016.
- [HTW15] Trevor Hastie, Robert Tibshirani, and Martin Wainwright. *Statistical learning with sparsity: the lasso and generalizations*. CRC press, 2015.
- [Hub64] Peter J Huber. Robust estimation of a location parameter. *The annals of mathematical statistics*, pages 73–101, 1964.
- [Hub11] Peter J Huber. Robust statistics. In *International Encyclopedia of Statistical Science*, pages 1248–1251. Springer, 2011.
- [JP78] David Johnson and Franco Preparata. The densest hemisphere problem. *Theoretical Computer Science*, 6(1):93–107, 1978.
- [JTK14] Prateek Jain, Ambuj Tewari, and Purushottam Kar. On iterative hard thresholding methods for high-dimensional M-estimation. In *Advances in Neural Information Processing Systems*, pages 685–693, 2014.
- [JVV86] Mark R Jerrum, Leslie G Valiant, and Vijay V Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, 1986.
- [KFB09] Daphne Koller, Nir Friedman, and Francis Bach. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- [Kin66] Benjamin King. Market and industry factors in stock price behavior. *the Journal of Business*, 39(1):139–190, 1966.
- [KKM18] Adam Klivans, Pravesh K. Kothari, and Raghu Meka. Efficient Algorithms for Outlier-Robust Regression. *arXiv preprint arXiv:1803.03241*, 2018.
- [KP18] Sushrut Karmalkar and Eric Price. Compressed sensing with adversarial sparse noise via l1 regression. *arXiv preprint arXiv:1809.08055*, 2018.
- [LB18] Haoyang Liu and Rina Foygel Barber. Between hard and soft thresholding: optimal iterative thresholding algorithms. *arXiv preprint arXiv:1804.08841*, 2018.
- [LHY⁺12a] Han Liu, Fang Han, Ming Yuan, John Lafferty, and Larry Wasserman. The nonparanormal skeptic. *arXiv preprint arXiv:1206.6488*, 2012.
- [LHY⁺12b] Han Liu, Fang Han, Ming Yuan, John Lafferty, Larry Wasserman, et al. High-dimensional semiparametric gaussian copula graphical models. *The Annals of Statistics*, 40(4):2293–2326, 2012.

- [Li13] Xiaodong Li. Compressed sensing and matrix completion with constant proportion of corruptions. *Constructive Approximation*, 37(1):73–99, 2013.
- [LL17] Guillaume Lecué and Matthieu Lerasle. Robust machine learning by median-of-means: theory and practice. *arXiv preprint arXiv:1711.10306*, 2017.
- [LM16] Gabor Lugosi and Shahar Mendelson. Risk minimization by median-of-means tournaments. *arXiv preprint arXiv:1608.00757*, 2016.
- [Loh17] Po-Ling Loh. Statistical consistency and asymptotic normality for high-dimensional robust m -estimators. *The Annals of Statistics*, 45(2):866–896, 2017.
- [LPR15] Tianyang Li, Adarsh Prasad, and Pradeep K Ravikumar. Fast classification rates for high-dimensional gaussian generative models. In *Advances in Neural Information Processing Systems*, pages 1054–1062, 2015.
- [LRV16] Kevin A Lai, Anup B Rao, and Santosh Vempala. Agnostic estimation of mean and covariance. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 665–674. IEEE, 2016.
- [LSLC18] Liu Liu, Yanyao Shen, Tianyang Li, and Constantine Caramanis. High dimensional robust sparse regression. *arXiv preprint arXiv:1805.11643*, 2018.
- [LSRC15] Yen-Huan Li, Jonathan Scarlett, Pradeep Ravikumar, and Volkan Cevher. Sparsistency of 1-regularized M-estimators. In *AISTATS*, 2015.
- [LT18] Po-Ling Loh and Xin Lu Tan. High-dimensional robust precision matrix estimation: Cellwise corruption under ϵ -contamination. *Electronic Journal of Statistics*, 12(1):1429–1467, 2018.
- [LW11] Po-Ling Loh and Martin J Wainwright. High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity. In *Advances in Neural Information Processing Systems*, pages 2726–2734, 2011.
- [LYCR17] Tianyang Li, Xinyang Yi, Constantine Caramanis, and Pradeep Ravikumar. Minimax gaussian classification & clustering. In *Artificial Intelligence and Statistics*, pages 1–9, 2017.
- [MB06] Nicolai Meinshausen and Peter Bühlmann. High-dimensional graphs and variable selection with the lasso. *The annals of statistics*, 34(3):1436–1462, 2006.
- [Min15] Stanislav Minsker. Geometric median and robust estimation in banach spaces. *Bernoulli*, 21(4):2308–2335, 2015.
- [Min18] Stanislav Minsker. Sub-gaussian estimators of the mean of a random matrix with heavy-tailed entries. *The Annals of Statistics*, 46(6A):2871–2903, 2018.
- [NRWY12] Sahand N Negahban, Pradeep Ravikumar, Martin J Wainwright, and Bin Yu. A unified framework for high-dimensional analysis of m -estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, 2012.
- [NY83] Arkadii Semenovich Nemirovsky and David Borisovich Yudin. Problem complexity and method efficiency in optimization. 1983.
- [PSBR18] Adarsh Prasad, Arun Sai Suggala, Sivaraman Balakrishnan, and Pradeep Ravikumar. Robust estimation via robust gradient estimation. *arXiv preprint arXiv:1802.06485*, 2018.
- [Rou84] Peter J Rousseeuw. Least median of squares regression. *Journal of the American statistical association*, 79(388):871–880, 1984.
- [RWL10] Pradeep Ravikumar, Martin J Wainwright, and John D Lafferty. High-dimensional ising model selection using ℓ_1 -regularized logistic regression. *The Annals of Statistics*, 38(3):1287–1319, 2010.
- [RWRY11] Pradeep Ravikumar, Martin J Wainwright, Garvesh Raskutti, and Bin Yu. High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011.
- [RWY10] Garvesh Raskutti, Martin J Wainwright, and Bin Yu. Restricted eigenvalue properties for correlated gaussian designs. *Journal of Machine Learning Research*, 11(Aug):2241–2259, 2010.
- [RWY11] Garvesh Raskutti, Martin J Wainwright, and Bin Yu. Minimax rates of estimation for high-dimensional linear regression over ℓ_q -balls. *IEEE transactions on information theory*, 57(10):6976–6994, 2011.

- [SCV17] Jacob Steinhardt, Moses Charikar, and Gregory Valiant. Resilience: A criterion for learning in the presence of arbitrary outliers. *arXiv preprint arXiv:1703.04940*, 2017.
- [SL17] Jie Shen and Ping Li. A tight bound of hard thresholding. *The Journal of Machine Learning Research*, 18(1):7650–7691, 2017.
- [SS18] Yanyao Shen and Sujay Sanghavi. Iteratively learning from the best. *arXiv preprint arXiv:1810.11874*, 2018.
- [SX18] Lili Su and Jiaming Xu. Securing distributed machine learning in high dimensions. *arXiv preprint arXiv:1804.10140*, 2018.
- [Tib96] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- [Tuk75] John W Tukey. Mathematics and the picturing of data. In *Proceedings of the International Congress of Mathematicians, Vancouver, 1975*, volume 2, pages 523–531, 1975.
- [vdV00] Aad van der Vaart. *Asymptotic statistics*. Cambridge University Press, 2000.
- [VMX17] Daniel Vainsencher, Shie Mannor, and Huan Xu. Ignoring is a bliss: Learning with large noise through reweighting-minimization. In *Conference on Learning Theory*, pages 1849–1881, 2017.
- [Wai09] Martin J Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (lasso). *IEEE transactions on information theory*, 55(5):2183–2202, 2009.
- [Wai19] Martin Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*. Cambridge University Press, 2019.
- [WG17] Lingxiao Wang and Quanquan Gu. Robust gaussian graphical model estimation with arbitrary corruption. In *International Conference on Machine Learning*, pages 3617–3626, 2017.
- [WJ08] Martin Wainwright and Michael Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008.
- [YCRB18a] Dong Yin, Yudong Chen, Kannan Ramchandran, and Peter Bartlett. Byzantine-robust distributed learning: Towards optimal statistical rates. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5650–5659. PMLR, 10–15 Jul 2018.
- [YCRB18b] Dong Yin, Yudong Chen, Kannan Ramchandran, and Peter Bartlett. Defending against saddle point attack in byzantine-robust distributed learning. *arXiv preprint arXiv:1806.05358*, 2018.
- [YL15] Eunho Yang and Aurélie C Lozano. Robust gaussian graphical modeling with the trimmed graphical lasso. In *Advances in Neural Information Processing Systems*, pages 2602–2610, 2015.
- [YLA18] Eunho Yang, Aurélie C Lozano, and Aleksandr Aravkin. A general family of trimmed estimators for robust high-dimensional data analysis. *Electronic Journal of Statistics*, 12(2):3519–3553, 2018.
- [YLZ18] Xiao-Tong Yuan, Ping Li, and Tong Zhang. Gradient hard thresholding pursuit. *Journal of Machine Learning Research*, 18(166):1–43, 2018.
- [Zhu17] Ziwei Zhu. Taming the heavy-tailed features by shrinkage and clipping. *arXiv preprint arXiv:1710.09020*, 2017.
- [ZLR⁺12] Tuo Zhao, Han Liu, Kathryn Roeder, John Lafferty, and Larry Wasserman. The huge package for high-dimensional undirected graph estimation in r. *Journal of Machine Learning Research*, 13(Apr):1059–1062, 2012.
- [ZWJ14] Yuchen Zhang, Martin J Wainwright, and Michael I Jordan. Lower bounds on the performance of polynomial-time algorithms for sparse linear regression. In *Conference on Learning Theory*, pages 921–948, 2014.

Notations in Appendix. In our proofs, the exponent -10 in tail bounds is arbitrary, and can be changed to other larger constant without affecting the results. $\{c_j\}_{j=0}^3$ denote universal constants, and they may change line by line.

A Proofs for the gradient estimators

In Robust Hard Thresholding (Algorithm 1), we use trimmed gradient estimator or MOM gradient estimator. And in Theorem 3.1, the key quantity to control the statistical rates of convergence is the Robust Descent Condition (Definition 3.1).

By Holder inequality, we have

$$\left| \left\langle \widehat{\mathbf{G}}(\boldsymbol{\beta}) - \mathbf{G}(\boldsymbol{\beta}), \widetilde{\boldsymbol{\beta}} - \boldsymbol{\beta}^* \right\rangle \right| \leq \left\| \widehat{\mathbf{G}}(\boldsymbol{\beta}) - \mathbf{G}(\boldsymbol{\beta}) \right\|_{\infty} \left\| \widetilde{\boldsymbol{\beta}} - \boldsymbol{\beta}^* \right\|_1.$$

In this section, we provide one direct route for obtaining upper bound of Robust Descent Condition via bounding the infinity norm of the robust gradient estimate (Proposition A.1 and Proposition A.2).

Later, in Appendix B, we will leverage Proposition A.1 and Proposition A.2 in verifying the Robust Descent Condition for trimmed/MOM gradient estimator under sparse linear/logistic regression. Together with Theorem 3.1, this will complete Corollary 4.1 – Corollary 4.4.

Proposition A.1. *Suppose we observe n ϵ -corrupted sub-Gaussian samples (Definition 2.1). With probability at least $1 - d^{-3}$, the coordinate-wise trimmed gradient estimator can guarantee*

- $\|\widehat{\mathbf{G}} - \mathbf{G}\|_{\infty} = O\left(\sqrt{\|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|_2^2 + \sigma^2} \left(\epsilon \log(nd) + \sqrt{\log d/n}\right)\right)$ for sparse linear regression (Model 4.1).
- $\|\widehat{\mathbf{G}} - \mathbf{G}\|_{\infty} = O\left(\epsilon \log(nd) + \sqrt{\log d/n}\right)$ for sparse logistic regression (Model 4.2).

Proposition A.2. *Suppose we observe n samples from the heavy tailed model with bounded 4-th moment covariates. With probability at least $1 - d^{-3}$, the coordinate-wise Median of Means gradient estimator can guarantee*

- $\|\widehat{\mathbf{G}} - \mathbf{G}\|_{\infty} = O\left(\sqrt{\rho^2 \|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|_2^2 + \rho \sigma^2} \sqrt{\log d/n}\right)$ for sparse linear regression;
- $\|\widehat{\mathbf{G}} - \mathbf{G}\|_{\infty} = O\left(\sqrt{\rho \log d/n}\right)$ for sparse logistic regression.

A.1 Proofs for the MOM gradient estimator

We first prove Proposition A.2. Proposition A.1 of trimmed gradient estimator for ϵ -corrupted sub-Gaussian samples has the same dependency on $\|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|_2$. The proof of Proposition A.1 leverages standard concentration bound for sub-Gaussian samples, and then uses trimming to control the effect of outliers.

Proof of Proposition A.2. For ℓ_2 loss function, we have $\mathbf{g}(\boldsymbol{\beta}) = \mathbf{x}(\mathbf{x}^\top \boldsymbol{\beta} - y)$, where we omit the subscript i in the proof. We denote $\Delta := \boldsymbol{\beta} - \boldsymbol{\beta}^*$, and bound the operator norm of the covariance of

gradient samples

$$\begin{aligned}
\|\mathbb{E}(\mathbf{g} - \mathbf{G})(\mathbf{g} - \mathbf{G})^\top\|_{\text{op}} &\leq \|\mathbb{E}((\mathbf{x}\mathbf{x}^\top - \Sigma)\Delta\Delta^\top(\mathbf{x}\mathbf{x}^\top - \Sigma))\|_{\text{op}} + \|\mathbb{E}(\xi^2\mathbf{x}\mathbf{x}^\top)\|_{\text{op}} \\
&\leq \sup_{\mathbf{v}_1 \in \mathcal{S}^{d-1}} \mathbf{v}_1^\top \mathbb{E}((\mathbf{x}\mathbf{x}^\top - \Sigma)\Delta\Delta^\top(\mathbf{x}\mathbf{x}^\top - \Sigma))\mathbf{v}_1 + \sigma^2 \|\Sigma\|_{\text{op}} \\
&\leq \sup_{\mathbf{v}_1 \in \mathcal{S}^{d-1}} \langle \Delta\Delta^\top, \mathbb{E}(\mathbf{x}\mathbf{x}^\top - \Sigma)\mathbf{v}_1\mathbf{v}_1^\top(\mathbf{x}\mathbf{x}^\top - \Sigma) \rangle + \sigma^2 \|\Sigma\|_{\text{op}} \\
&\stackrel{(i)}{\leq} \|\Delta\|_2^2 \sup_{\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{S}^{d-1}} \mathbb{E}(\mathbf{v}_2^\top(\mathbf{x}\mathbf{x}^\top - \Sigma)\mathbf{v}_1)^2 + \sigma^2 \|\Sigma\|_{\text{op}} \\
&\leq 2\|\Delta\|_2^2 \sup_{\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{S}^{d-1}} (\mathbb{E}(\mathbf{v}_2^\top(\mathbf{x}\mathbf{x}^\top)\mathbf{v}_1)^2 + \|\Sigma\|_{\text{op}}^2) + \sigma^2 \|\Sigma\|_{\text{op}} \\
&\leq 2\|\Delta\|_2^2 \sup_{\mathbf{v}_1, \mathbf{v}_2 \in \mathcal{S}^{d-1}} (\sqrt{\mathbb{E}(\mathbf{v}_2^\top\mathbf{x})^4} \sqrt{\mathbb{E}(\mathbf{x}^\top\mathbf{v}_1)^4} + \|\Sigma\|_{\text{op}}^2) + \sigma^2 \|\Sigma\|_{\text{op}} \\
&\stackrel{(ii)}{\leq} 2(C_4 + 1) \|\Sigma\|_{\text{op}}^2 \|\Delta\|_2^2 + \sigma^2 \|\Sigma\|_{\text{op}},
\end{aligned}$$

where (i) follows from the Holder inequality, and (ii) follows from the 4-th moment bound assumption.

Hence, by using coordinate-wise Median of Means gradient estimator, we have

$$\sup_{\mathbf{v} \in \mathcal{S}^{d-1}} \mathbf{v}^\top (\widehat{\mathbf{G}} - \mathbf{G}) \stackrel{(i)}{=} O\left(\sqrt{\|\text{Cov}(\mathbf{g})\|_{\text{op}} \log d/n}\right) = O\left(\sqrt{\rho^2 \|\beta - \beta^*\|_2^2 + \rho\sigma^2 \sqrt{\log d/n}}\right)$$

with probability at least $1 - d^{-4}$, where (i) follows from Proposition 5 in [HS16]. Applying union bounds on all d indexes, we have $\|\widehat{\mathbf{G}} - \mathbf{G}\|_\infty = O\left(\sqrt{\rho^2 \|\beta - \beta^*\|_2^2 + \rho\sigma^2 \sqrt{\log d/n}}\right)$ with probability at least $1 - d^{-3}$.

For logistic loss, the gradient can be computed as: $\mathbf{g} = \frac{-y\mathbf{x}}{1 + \exp(y\mathbf{x}^\top\beta)}$, where we omit the subscript i in the proof.

Since $y \in \{-1, +1\}$, and $1 + \exp(y\mathbf{x}^\top\beta) \geq 1$, we directly have $\|\mathbb{E}(\mathbf{g} - \mathbf{G})(\mathbf{g} - \mathbf{G})^\top\|_{\text{op}} \leq \|\Sigma\|_{\text{op}}$. Similar to the case of ℓ_2 loss, we have $\|\widehat{\mathbf{G}} - \mathbf{G}\|_\infty = O\left(\sqrt{\rho \log d/n}\right)$, with probability at least $1 - d^{-3}$. \square

A.2 Proofs for the trimmed gradient estimator

We then turn to the trimmed gradient estimator for the case of arbitrary corruption. Before we proceed to the trimmed estimator, let us first visit the definition and tail bounds of sub-exponential random variable, as it will be used in sparse linear regression, where the gradient's distribution is indeed sub-exponential under the sub-Gaussian assumptions in Model 4.1.

We first present standard concentration inequalities ([Wai19]).

Definition A.1 (Sub-exponential random variables). *A random variable X with mean μ is sub-exponential if there are non-negative parameters ν such that*

$$\mathbb{E}[\exp(t(X - \mu))] \leq \exp\left(\frac{\nu^2 t^2}{2}\right), \quad \text{for all } |t| < \frac{1}{\nu}.$$

Lemma A.1 (Bernstein’s inequality). *Suppose that $X_i, i = 1, \dots, n$, are i.i.d. sub-exponential random variables with parameters ν . Then*

$$\Pr\left(\frac{1}{n}\sum_{i=1}^n X_i \geq \mu + t\right) \leq \begin{cases} \exp\left(-\frac{nt^2}{2\nu^2}\right) & \text{if } 0 \leq t \leq \nu, \text{ and} \\ \exp\left(-\frac{nt}{2\nu}\right) & \text{for } t > \nu. \end{cases}$$

We also have a two-sided tail bound

$$\Pr\left(\left|\frac{1}{n}\sum_{i=1}^n X_i - \mu\right| \geq t\right) \leq 2 \exp\left(-n \min\left(\frac{t^2}{2\nu^2}, \frac{t}{2\nu}\right)\right).$$

We define α -trimmed mean estimator for one dimensional samples, and denote it as $\text{trmean}_\alpha(\cdot)$.

Definition A.2 (α -trimmed mean estimator). *Given a set of ϵ -corrupted samples $\{\mathbf{z}_i \in \mathbb{R}, i \in \mathcal{S}\}$, the coordinate-wise trimmed mean estimator $\text{trmean}_\alpha(\cdot)$ removes the largest and smallest α fraction of elements in $\{\mathbf{z}_i \in \mathbb{R}, i \in \mathcal{S}\}$, and calculate the mean of the remaining terms. We choose $\alpha = c_0\epsilon$, for a constant $c_0 \geq 1$. We also require that $\alpha \leq 1/2 - c_1$, for some small constant $c_1 > 0$.*

Lemma A.2 shows the guarantees for this robust gradient estimator in each coordinate. We note that **Lemma A.2** is stronger than guarantees for trimmed mean estimator (Lemma 3) in [YCRB18a]. In our contamination model **Definition 2.1**, the adversary may *delete* ϵ -fraction of authentic samples, and then add arbitrary outliers. And **Lemma A.2** provides guarantees for trimmed mean estimator on sub-exponential random variables. The trimmed mean estimator is robust enough, that it allows the adversary to arbitrarily remove ϵ -fraction of data points. We use \mathcal{G}^j to denote the \mathbb{R}^1 samples at the j -th coordinate of \mathcal{G} . We can also define \mathcal{S}^j in the same way.

Lemma A.2. *Suppose we observe $n = \Omega(\log d)$ ϵ -corrupted samples from **Definition 2.1**. For each dimension $j \in \{1, 2, \dots, d\}$, we assume the samples in \mathcal{G}^j are i.i.d. ν -sub-exponential with mean $\boldsymbol{\mu}^j$. After the contamination, we have the j -th \mathbb{R}^1 samples as \mathcal{S}^j . Then, we can guarantee the trimmed mean estimator on j -th dimension that*

$$|\text{trmean}_\alpha\{x_i : i \in \mathcal{S}^j\} - \boldsymbol{\mu}^j| = O\left(\nu \left(\epsilon \log(nd) + \sqrt{\frac{\log d}{n}}\right)\right)$$

with probability at least $1 - d^{-4}$.

We leave the proof of **Lemma A.2** at the end of this section. Then, we present analysis of trimmed gradient estimator for sparse linear regression and sparse logistic regression by using **Lemma A.2**. For sparse linear regression model with sub-Gaussian covariates, the distribution of authentic gradients are sub-exponential instead of sub-Gaussian. More specifically, we first prove that when the current parameter iterate is $\boldsymbol{\beta}$, the sub-exponential parameter of all authentic gradient is $O((\|\Delta\|_2^2 + \sigma^2)^{1/2})$, where $\Delta := \boldsymbol{\beta} - \boldsymbol{\beta}^*$.

To gain some intuition for this, we can consider the sparse linear equation problem, where $\sigma^2 = 0$. When $\boldsymbol{\beta} = \boldsymbol{\beta}^*$ ($\|\Delta\|_2^2 = 0$), we exactly recover $\boldsymbol{\beta}^*$, and all stochastic gradients of authentic samples are actually zero vectors, as all observations are noiseless. It is clear that we will have sub-exponential parameter as 0.

Proof of Proposition A.1. For any β , the gradient for one sample can be written as

$$\mathbf{g} = \mathbf{x} (\mathbf{x}^\top \beta - y), \text{ and } \mathbf{G} = \mathbb{E}(\mathbf{g}) = \Sigma (\beta - \beta^*),$$

where we omit the subscript i in the proof. For any fixed standard basis vector $\mathbf{v} \in \mathbb{S}^{d-1}$, and define $\Delta = \beta - \beta^*$, we have

$$\mathbf{v}^\top \mathbf{g} = \mathbf{v}^\top \mathbf{x} \mathbf{x}^\top \Delta - \mathbf{v}^\top \mathbf{x} \xi, \text{ and } \mathbf{v}^\top \mathbf{G} = \mathbf{v}^\top \Sigma \Delta. \quad (4)$$

To characterize the tail bounds of $\mathbf{v}^\top \mathbf{g}$, we study the moment generating function:

$$\mathbb{E}[\exp(t(\mathbf{v}^\top \mathbf{g} - \mathbf{v}^\top \mathbf{G}))] = \mathbb{E}[\exp(t(\mathbf{v}^\top (\mathbf{x} \mathbf{x}^\top - \Sigma) \Delta - \mathbf{v}^\top \mathbf{x} \xi))].$$

We denote $\gamma \in \{-1, +1\}$ as a Rademacher random variable, which is independent of \mathbf{x} and ξ . Then we can use a standard symmetrization technique [Wai19],

$$\begin{aligned} \mathbb{E}_{\mathbf{x}, \xi}[\exp(t(\mathbf{v}^\top (\mathbf{x} \mathbf{x}^\top - \Sigma) \Delta - \mathbf{v}^\top \mathbf{x} \xi))] &\leq \mathbb{E}_{\mathbf{x}, \xi, \gamma}[\exp(2t\gamma(\mathbf{v}^\top \mathbf{x} \mathbf{x}^\top \Delta - \mathbf{v}^\top \mathbf{x} \xi))] \\ &\stackrel{(i)}{=} \sum_{k=0}^{\infty} \frac{1}{k!} (2t)^k \mathbb{E}[\gamma^k (\mathbf{v}^\top \mathbf{x} \mathbf{x}^\top \Delta - \mathbf{v}^\top \mathbf{x} \xi)^k] \\ &\stackrel{(ii)}{=} 1 + \sum_{l=1}^{\infty} \frac{1}{(2l)!} (2t)^{2l} \mathbb{E}[(\mathbf{v}^\top \mathbf{x})^{2l} (\mathbf{x}^\top \Delta - \xi)^{2l}], \end{aligned}$$

where (i) follows from the exponential function's power series expansion, and (ii) follows from the independence of γ , together with the fact that all odd moments of the γ terms have zero means.

By the Cauchy-Schwarz inequality, we have

$$\mathbb{E}[(\mathbf{v}^\top \mathbf{x})^{2l} (\mathbf{x}^\top \Delta - \xi)^{2l}] \leq \sqrt{\mathbb{E}[(\mathbf{v}^\top \mathbf{x})^{4l}] \mathbb{E}[(\mathbf{x}^\top \Delta - \xi)^{4l]}.$$

It is clear that ξ is a sub-Gaussian random variable with parameter σ . Since $\mathbf{x} \sim \mathcal{N}(0, \Sigma)$, we have $\mathbf{v}^\top \mathbf{x} \sim \mathcal{N}(0, \mathbf{v}^\top \Sigma \mathbf{v})$. For any fixed standard basis vector $\mathbf{v} \in \mathbb{S}^{d-1}$, we can conclude that $\mathbf{v}^\top \mathbf{x}$ is sub-Gaussian with parameter at most 1 based on Model 4.1. By basic properties of sub-Gaussian random variables [Wai19], we have

$$\begin{aligned} \sqrt{\mathbb{E}[(\mathbf{v}^\top \mathbf{x})^{4l}] &\leq \sqrt{\frac{(4l)!}{2^{2l} (2l)!}} (\sqrt{8e})^{2l} \\ \sqrt{\mathbb{E}[(\mathbf{x}^\top \Delta - \xi)^{4l}] &\stackrel{(i)}{\leq} \sqrt{\frac{(4l)!}{2^{2l} (2l)!}} \left(8e^2 (\|\Delta\|_2^2 + \sigma^2)\right)^l, \end{aligned}$$

where (i) follows from the fact that $\mathbf{x}^\top \Delta - \xi$ is the weighted summation of two independent sub-

Gaussian random variables. Hence, we have

$$\begin{aligned}
\mathbb{E}[\exp(t(\mathbf{v}^\top \mathbf{g} - \mathbf{v}^\top \mathbf{G}))] &\leq 1 + \sum_{l=1}^{\infty} \frac{1}{(2l)!} (2t)^{2l} \frac{(4l)!}{2^{2l} (2l)!} (\sqrt{8e})^{4l} (\|\Delta\|_2^2 + \sigma^2)^l \\
&\stackrel{(i)}{\leq} 1 + \sum_{l=1}^{\infty} (4t)^{2l} (\sqrt{8e})^{4l} (\|\Delta\|_2^2 + \sigma^2)^l \\
&= 1 + \sum_{l=1}^{\infty} (4t)^{2l} (8e^2)^{2l} \left(\sqrt{\|\Delta\|_2^2 + \sigma^2}\right)^{2l}, \tag{5}
\end{aligned}$$

where (i) follows from $(4l)! \leq 2^{4l} ((2l)!)^2$ (proof by mathematical induction). When we have $f(t) = 32te^2 \sqrt{\|\Delta\|_2^2 + \sigma^2} < 1$, eq. (5) converges to $\frac{1}{1-f^2(t)}$. Hence,

$$\mathbb{E}[\exp(t(\mathbf{v}^\top \mathbf{g} - \mathbf{v}^\top \mathbf{G}))] \leq \frac{1}{1-f^2(t)} \leq \exp(f^2(t)).$$

That being said, $\mathbf{v}^\top \mathbf{g}$ is a sub-exponential random variable. By choosing \mathbf{v} as each coordinate in \mathbb{R}^d , each coordinate of gradient has sub-exponential parameter as $32\sqrt{2}e^2 \sqrt{\|\Delta\|_2^2 + \sigma^2}$.

Then, applying Lemma A.2 on this collection of corrupted sub-exponential random variables, we have

$$|\text{trmean}_\alpha\{x_i : i \in \mathcal{S}^j\} - \mu^j| = O\left(\sqrt{\|\Delta\|_2^2 + \sigma^2} \left(\epsilon \log(nd) + \sqrt{\frac{\log d}{n}}\right)\right), \tag{6}$$

with probability at least $1 - d^{-4}$.

Applying union bounds on eq. (6) for all d indexes, we have

$$\|\widehat{\mathbf{G}} - \mathbf{G}\|_\infty = O\left(\sqrt{\|\Delta\|_2^2 + \sigma^2} \left(\epsilon \log(nd) + \sqrt{\frac{\log d}{n}}\right)\right),$$

with probability at least $1 - d^{-3}$.

In this subsection, we use Lemma A.2 to bound $\|\widehat{\mathbf{G}} - \mathbf{G}\|_\infty$ for sparse logistic regression. The technique for sparse logistic regression is similar to linear regression. Since we can directly show the sub-Gaussian distribution of gradient in this case, applying Lemma A.2 leads to the bound for $\|\widehat{\mathbf{G}} - \mathbf{G}\|_\infty$.

Under the statistical model of sparse logistic regression, the gradient can be computed as:

$$\mathbf{g} = \frac{-y\mathbf{x}}{1 + \exp(y\mathbf{x}^\top \boldsymbol{\beta})},$$

where we omit the subscript i in the proof.

Since $y \in \{-1, +1\}$, and $1 + \exp(y\mathbf{x}^\top \boldsymbol{\beta}) \geq 1$, then for any fixed standard basis vector $\mathbf{v} \in \mathbb{S}^{d-1}$, $\mathbf{v}^\top \mathbf{g}$ is sub-Gaussian with parameter at most 1 based on Model 4.2. Notice that ν -sub-Gaussian random variables are still ν -sub-exponential. Applying Lemma A.2 again, we have

$$|\text{trmean}_\alpha\{x_i : i \in \mathcal{S}^j\} - \mu| = O\left(\epsilon \log(nd) + \sqrt{\frac{\log d}{n}}\right) \tag{7}$$

with probability at least $1 - d^{-4}$.

Applying union bounds on eq. (7) for all d indexes, we have

$$\left\| \widehat{\mathbf{G}} - \mathbf{G} \right\|_{\infty} = O \left(\epsilon \log(nd) + \sqrt{\frac{\log d}{n}} \right)$$

with probability at least $1 - d^{-3}$. □

A.3 Trimmed mean estimator for strong contamination model

Now, it only remains to prove [Lemma A.2](#). The proof technique is as follow: even though an adversary may delete samples from \mathcal{G}^j , we can still show the concentration inequalities for remaining authentic \mathbb{R}^1 samples (denoting as $\widetilde{\mathcal{G}}^j$ in the proof). Then, we show that by using trimmed mean estimator, either the abnormal outliers will be removed, or their effect is controlled.

Proof of Lemma A.2. Without loss of generality, we assume $\boldsymbol{\mu} = 0$ throughout the proof.

For each dimension $j \in \{1, 2, \dots, d\}$, we can split the j -th one-dimensional samples as $\mathcal{S}^j = \widetilde{\mathcal{G}}^j \cup \mathcal{B}^j$. To study the performance of $\text{trmean}_{\alpha}\{x_i : i \in \mathcal{S}^j\}$, we first show a concentration inequality of the sub-exponential variables in $\widetilde{\mathcal{G}}^j$, without worrying about removing points from \mathcal{G}^j . This part of our proof is similar to Lemma 4.5 in [\[DKK+16\]](#).

Concentration inequality for $\widetilde{\mathcal{G}}^j$ We consider the set $\{x_i : i \in \widetilde{\mathcal{G}}^j\}$ in \mathbb{R}^1 . Since $\widetilde{\mathcal{G}}^j$ is a subset of \mathcal{G}^j , by triangle inequality we have,

$$\left| \mathbb{E}_{i \in \widetilde{\mathcal{G}}^j} x_i \right| = \left| \frac{\sum_{i \in \widetilde{\mathcal{G}}^j} x_i}{(1 - \epsilon)n} \right| \leq \underbrace{\left| \frac{\sum_{i \in \mathcal{G}^j} x_i}{(1 - \epsilon)n} \right|}_{A_1} + \underbrace{\left| \frac{\sum_{i \in \mathcal{G}^j \setminus \widetilde{\mathcal{G}}^j} x_i}{(1 - \epsilon)n} \right|}_{A_2}.$$

The first term A_1 is simply the average of i.i.d. sub-exponential random variables. By [Lemma A.1](#), we have

$$\Pr \left(\left| \frac{\sum_{i \in \widetilde{\mathcal{G}}^j} x_i}{(1 - \epsilon)n} \right| \geq c_0 \nu \sqrt{\frac{\log d}{n}} \right) \leq 2 \exp \left(-c_1 n \min \left(\sqrt{\frac{\log d}{n}}, \frac{\log d}{n} \right) \right) \quad (8)$$

$$\leq c_2 d^{-10}. \quad (9)$$

For the second term A_2 , We now wish to show that with probability $1 - \tau$, there does not exist a subset $\widetilde{\mathcal{G}}^j$ so that the A_2 is more than δ_0 . This event is equivalent to

$$\left| \frac{\sum_{i \in \mathcal{G}^j \setminus \widetilde{\mathcal{G}}^j} x_i}{(1 - \epsilon)n} \right| = \left| \frac{\sum_{i \in \mathcal{G}^j \setminus \widetilde{\mathcal{G}}^j} x_i}{\epsilon n} \right| \frac{\epsilon}{1 - \epsilon} \geq \delta_0$$

Let $\delta_1 = \frac{1 - \epsilon}{\epsilon} \delta_0$. For one subset $\mathcal{G}^j \setminus \widetilde{\mathcal{G}}^j$, by [Lemma A.1](#), we have

$$\Pr \left(\left| \frac{\sum_{i \in \mathcal{G}^j \setminus \widetilde{\mathcal{G}}^j} x_i}{\epsilon n} \right| \geq \delta_1 \right) \leq 2 \exp \left(-\epsilon n \min \left(\frac{\delta_1^2}{2\nu^2}, \frac{\delta_1}{2\nu} \right) \right).$$

Then, we take union bounds over all possible $\mathcal{G}^j \setminus \tilde{\mathcal{G}}^j$, which have $\binom{n}{\epsilon n}$ events. Hence, the tail probability of A_2 can be bounded as

$$\begin{aligned} \tau &\leq 2 \binom{n}{\epsilon n} \exp\left(-\epsilon n \min\left(\frac{\delta_1^2}{2\nu^2}, \frac{\delta_1}{2\nu}\right)\right) \\ &\stackrel{(i)}{\leq} c_0 \exp\left(nH(\epsilon) - \epsilon n \min\left(\frac{\delta_1^2}{2\nu^2}, \frac{\delta_1}{2\nu}\right)\right) \end{aligned} \quad (10)$$

where (i) follows from the fact that $\log \binom{n}{\epsilon n} = O(nH(\epsilon))$ for n large enough, and $H(\cdot)$ is the binary entropy function. Choosing $\delta_1 = c_1\nu \log(nd)$, and hence $\delta_0 = c_1\nu\epsilon \log(nd)$, we have $\tau \leq c_0 \exp(-c_2 n\epsilon \log(nd)) \leq c_3 d^{-10}$.

Combining the analysis on A_1 and A_2 (eq. (9) and eq. (10)), we have

$$\Pr\left(\left|\mathbb{E}_{i \in_u \tilde{\mathcal{G}}^j} x_i\right| \geq \nu \left(c_0 \sqrt{\frac{\log d}{n}} + c_1 \epsilon \log(nd)\right)\right) \leq c_2 d^{-10}. \quad (11)$$

This completes the concentration bounds on $\left|\mathbb{E}_{i \in_u \tilde{\mathcal{G}}^j} x_i\right|$ for all possible samples in $\tilde{\mathcal{G}}^j$ without worrying about sample removing.

Trimmed mean estimator for \mathcal{S}^j Then, we can consider the contribution of each part in $\mathcal{S}^j = \tilde{\mathcal{G}}^j \cup \mathcal{B}^j$. We denote the remaining set after trimming as \mathcal{R}^j , and the trimmed set as \mathcal{T}^j . Recall that we assume $\boldsymbol{\mu} = 0$, we only need to bound $|\text{trmean}_\alpha\{x_i : i \in \mathcal{S}^j\}|$, which is the empirical average of all samples in the remaining set $\{x_i : i \in \mathcal{R}^j\}$.

As \mathcal{R}^j can be easily separated by the union of two distinct set $\mathcal{B}^j \cap \mathcal{R}^j$ and $\tilde{\mathcal{G}}^j \cap \mathcal{R}^j$, we have the following inequalities,

$$\begin{aligned} |\text{trmean}_\alpha\{x_i : i \in \mathcal{S}^j\}| &= \left| \frac{1}{(1-2\alpha)n} \sum_{i \in \mathcal{R}^j} x_i \right| \\ &\leq \frac{1}{(1-2\alpha)n} \left| \sum_{i \in \tilde{\mathcal{G}}^j} x_i - \sum_{i \in \tilde{\mathcal{G}}^j \cap \mathcal{T}^j} x_i + \sum_{i \in \mathcal{B}^j \cap \mathcal{R}^j} x_i \right| \\ &\leq \frac{1}{(1-2\alpha)n} \left(\underbrace{\left| \sum_{i \in \tilde{\mathcal{G}}^j} x_i \right|}_{B_1} + \underbrace{\left| \sum_{i \in \tilde{\mathcal{G}}^j \cap \mathcal{T}^j} x_i \right|}_{B_2} + \underbrace{\left| \sum_{i \in \mathcal{B}^j \cap \mathcal{R}^j} x_i \right|}_{B_3} \right) \end{aligned}$$

For any $i \in \tilde{\mathcal{G}}^j$, by Lemma A.1, we have

$$\Pr(|x_i| \geq c_0\nu \log(nd)) \leq 2 \exp(-c_1 \min(\log(nd), \log^2(nd))).$$

Applying a union bound for all samples, we can control the maximum magnitude for any $i \in \tilde{\mathcal{G}}^j$,

$$\begin{aligned} \Pr\left(\max_{i \in \tilde{\mathcal{G}}^j} |x_i| \geq c_0\nu \log(nd)\right) &\leq 2(1-\epsilon)n \exp(-c_1 \min(\log(nd), \log^2(nd))), \\ &\leq c_1 d^{-10}. \end{aligned}$$

We can bound B_1 by applying eq. (11). For the trimmed good samples $\{i \in \tilde{\mathcal{G}}^j \cap \mathcal{T}^j\}$, we have $B_2 \leq 2\alpha n \max_{i \in \tilde{\mathcal{G}}^j} |x_i|$. Since we choose $\alpha \geq \epsilon$, we have $B_3 \leq \epsilon n \max_{i \in \tilde{\mathcal{G}}^j} |x_i|$.

Putting together the pieces, and choosing $\alpha = c\epsilon$ for some universal constant $c \geq 1$, we have

$$|\text{trmean}_\alpha\{x_i : i \in \mathcal{S}^j\} - \boldsymbol{\mu}^j| = O\left(\nu \left(\epsilon \log(nd) + \sqrt{\frac{\log d}{n}}\right)\right),$$

with probability at least $1 - d^{-4}$. This completes the proof for Lemma A.2. \square

B Statistical estimation via Robust Hard Thresholding

Here, we provide the Meta-Theorem Theorem 3.1 for statistical estimation performance of Algorithm 1 under statistical models.

We first introduce a supporting Lemma on the property of hard thresholding operator.

Lemma B.1 (Lemma 1 in [LB18]). *We set k' in hard thresholding operator as $k' = kc_\rho^2$, where $c_\rho \geq 1$, then we have*

$$\sup \left\{ \frac{\langle \boldsymbol{\beta}^* - P_{k'}(\mathbf{z}), \mathbf{z} - P_{k'}(\mathbf{z}) \rangle}{\|\boldsymbol{\beta}^* - P_{k'}(\mathbf{z})\|_2^2} : \boldsymbol{\beta}^*, \mathbf{z} \in \mathbb{R}^d, \|\boldsymbol{\beta}^*\|_0 \leq k, \boldsymbol{\beta}^* \neq P_{k'}(\mathbf{z}) \right\} = \frac{1}{2} \sqrt{\frac{k}{k'}} = \frac{1}{2c_\rho}.$$

Note that c_ρ will be specified as 2ρ later in the proof, as we choose $k' = 4\rho^2 k$ as in Definition 2.3.

Proof of Theorem 3.1. We first study the objective function gap $f(\bar{\boldsymbol{\beta}}^t) - f(\boldsymbol{\beta}^*)$. Since the population risk f satisfies μ_α -strong convexity and μ_L -smoothness (Definition 2.3), we have

$$f(\boldsymbol{\beta}^*) \stackrel{(i)}{\geq} f(\boldsymbol{\beta}^{t-1}) + \langle \mathbf{G}(\boldsymbol{\beta}^{t-1}), \boldsymbol{\beta}^* - \boldsymbol{\beta}^{t-1} \rangle + \frac{\mu_\alpha}{2} \|\boldsymbol{\beta}^* - \boldsymbol{\beta}^{t-1}\|_2^2, \quad (12)$$

$$f(\bar{\boldsymbol{\beta}}^t) \leq f(\boldsymbol{\beta}^{t-1}) + \langle \mathbf{G}(\boldsymbol{\beta}^{t-1}), \bar{\boldsymbol{\beta}}^t - \boldsymbol{\beta}^{t-1} \rangle + \frac{\mu_L}{2} \|\bar{\boldsymbol{\beta}}^t - \boldsymbol{\beta}^{t-1}\|_2^2, \quad (13)$$

where (i) follows from the fact that $\boldsymbol{\beta}^*, \boldsymbol{\beta}^{t-1} \in B$, and μ_α -strong convexity holds.

Combining these two inequalities, we obtain

$$f(\bar{\boldsymbol{\beta}}^t) - f(\boldsymbol{\beta}^*) \leq \langle \mathbf{G}(\boldsymbol{\beta}^{t-1}), \bar{\boldsymbol{\beta}}^t - \boldsymbol{\beta}^* \rangle - \frac{\mu_\alpha}{2} \|\boldsymbol{\beta}^* - \boldsymbol{\beta}^{t-1}\|_2^2 + \frac{\mu_L}{2} \|\bar{\boldsymbol{\beta}}^t - \boldsymbol{\beta}^{t-1}\|_2^2.$$

Expanding the last term, we also have

$$\begin{aligned} \frac{\mu_L}{2} \|\bar{\boldsymbol{\beta}}^t - \boldsymbol{\beta}^*\|_2^2 &= \frac{\mu_L}{2} \|\boldsymbol{\beta}^{t-1} - \boldsymbol{\beta}^*\|_2^2 - \frac{\mu_L}{2} \|\bar{\boldsymbol{\beta}}^t - \boldsymbol{\beta}^{t-1}\|_2^2 + \mu_L \langle \boldsymbol{\beta}^{t-1} - \bar{\boldsymbol{\beta}}^t, \boldsymbol{\beta}^* - \bar{\boldsymbol{\beta}}^t \rangle \\ &= \frac{\mu_L}{2} \|\boldsymbol{\beta}^{t-1} - \boldsymbol{\beta}^*\|_2^2 - \frac{\mu_L}{2} \|\bar{\boldsymbol{\beta}}^t - \boldsymbol{\beta}^{t-1}\|_2^2 \\ &\quad + \underbrace{\mu_L \langle (\boldsymbol{\beta}^{t-1} - \eta \widehat{\mathbf{G}}(\boldsymbol{\beta}^{t-1})) - \bar{\boldsymbol{\beta}}^t, \boldsymbol{\beta}^* - \bar{\boldsymbol{\beta}}^t \rangle}_{T_1} - \underbrace{\langle \widehat{\mathbf{G}}(\boldsymbol{\beta}^{t-1}), \bar{\boldsymbol{\beta}}^t - \boldsymbol{\beta}^* \rangle}_{T_2} \end{aligned}$$

For the term T_1 , recall that $\bar{\boldsymbol{\beta}}^t$ is obtained from hard thresholding, and $\bar{\boldsymbol{\beta}}^t = P_{k'}(\boldsymbol{\beta}^{t-1} - \eta \widehat{\mathbf{G}}(\boldsymbol{\beta}^{t-1}))$,

we apply Lemma B.1 with $\mathbf{z} = \boldsymbol{\beta}^{t-1} - \eta \widehat{\mathbf{G}}(\boldsymbol{\beta}^{t-1})$:

$$\left\langle \left(\boldsymbol{\beta}^{t-1} - \eta \widehat{\mathbf{G}}(\boldsymbol{\beta}^{t-1}) \right) - \bar{\boldsymbol{\beta}}^t, \boldsymbol{\beta}^* - \bar{\boldsymbol{\beta}}^t \right\rangle \leq \frac{1}{2c_\rho} \|\bar{\boldsymbol{\beta}}^t - \boldsymbol{\beta}^*\|_2^2.$$

The term T_2 can be bounded by using eq. (1) in Definition 3.1. We have

$$\begin{aligned} \langle \widehat{\mathbf{G}}(\boldsymbol{\beta}^{t-1}), \bar{\boldsymbol{\beta}}^t - \boldsymbol{\beta}^* \rangle &\geq \langle \mathbf{G}(\boldsymbol{\beta}^{t-1}), \bar{\boldsymbol{\beta}}^t - \boldsymbol{\beta}^* \rangle - \left| \langle \mathbf{G}(\boldsymbol{\beta}^{t-1}) - \widehat{\mathbf{G}}(\boldsymbol{\beta}^{t-1}), \bar{\boldsymbol{\beta}}^t - \boldsymbol{\beta}^* \rangle \right| \\ &\geq \langle \mathbf{G}(\boldsymbol{\beta}^{t-1}), \bar{\boldsymbol{\beta}}^t - \boldsymbol{\beta}^* \rangle - (\alpha \|\boldsymbol{\beta}^{t-1} - \boldsymbol{\beta}^*\|_2 + \psi) \|\bar{\boldsymbol{\beta}}^t - \boldsymbol{\beta}^*\|_2, \end{aligned}$$

with probability at least $1 - d^{-3}$.

We denote $\bar{\Delta}^t = \bar{\boldsymbol{\beta}}^t - \boldsymbol{\beta}^*$ and $\Delta^t = \boldsymbol{\beta}^t - \boldsymbol{\beta}^*$. Since, $\eta\mu_\alpha \geq \frac{1}{\mu_L} \cdot \mu_\alpha = \frac{1}{\rho}$, putting together the pieces, we have

$$f(\bar{\boldsymbol{\beta}}^t) - f(\boldsymbol{\beta}^*) \leq \frac{\mu_L}{2} \left[\left(1 - \frac{1}{\rho}\right) \|\Delta^{t-1}\|_2^2 - \left(1 - \frac{1}{c_\rho}\right) \|\bar{\Delta}^t\|_2^2 \right] + (\alpha \|\Delta^{t-1}\|_2 + \psi) \|\bar{\Delta}^t\|_2, \quad (14)$$

with probability at least $1 - d^{-3}$. Applying convexity, $f(\bar{\boldsymbol{\beta}}^t) - f(\boldsymbol{\beta}^*) \geq 0$, as $\boldsymbol{\beta}^*$ is the population minimizer. Hence, we have

$$0 \leq \frac{\mu_L}{2} \left[\left(1 - \frac{1}{\rho}\right) \|\Delta^{t-1}\|_2^2 - \left(1 - \frac{1}{c_\rho}\right) \|\bar{\Delta}^t\|_2^2 \right] + (\alpha \|\Delta^{t-1}\|_2 + \psi) \|\bar{\Delta}^t\|_2, \quad (15)$$

with probability at least $1 - d^{-3}$.

Notice that eq. (15) is a quadratic inequality for $\|\bar{\Delta}^t\|_2$, and we can use the root of eq. (15) to upper bound $\|\bar{\Delta}^t\|_2$:

$$\begin{aligned} \|\bar{\Delta}^t\|_2 &\leq \frac{(\alpha \|\Delta^{t-1}\|_2 + \psi) + \sqrt{(\alpha \|\Delta^{t-1}\|_2 + \psi)^2 + \left(\mu_L \left(1 - \frac{1}{c_\rho}\right)\right) \cdot \mu_L \left(1 - \frac{1}{\rho}\right) \|\Delta^{t-1}\|_2^2}}{\mu_L \left(1 - \frac{1}{c_\rho}\right)} \\ &\stackrel{(i)}{\leq} \frac{2(\alpha \|\Delta^{t-1}\|_2 + \psi) + \|\Delta^{t-1}\|_2 \sqrt{\left(\mu_L \left(1 - \frac{1}{c_\rho}\right)\right) \cdot \mu_L \left(1 - \frac{1}{\rho}\right)}}{\mu_L \left(1 - \frac{1}{c_\rho}\right)} \\ &= \|\Delta^{t-1}\|_2 \sqrt{\frac{1 - \frac{1}{\rho}}{1 - \frac{1}{c_\rho}}} + \frac{2\alpha \|\Delta^{t-1}\|_2}{\mu_L \left(1 - \frac{1}{c_\rho}\right)} + \frac{2\psi}{\mu_L \left(1 - \frac{1}{c_\rho}\right)} \end{aligned}$$

where (i) follows from the basic inequality $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for non-negative a, b .

We choose $c_\rho = 2\rho$, and this leads to $\sqrt{\left(1 - \frac{1}{\rho}\right) / \left(1 - \frac{1}{c_\rho}\right)} \leq 1 - \frac{1}{4\rho}$. Under the condition $\alpha \leq \frac{1}{32}\mu_\alpha$, we have $\frac{2\alpha}{\mu_L \left(1 - \frac{1}{c_\rho}\right)} \leq \frac{1}{8\rho}$. Then,

$$\|\bar{\Delta}^t\|_2 \leq \left(1 - \frac{1}{8\rho}\right) \|\Delta^{t-1}\|_2 + \frac{4\psi}{\mu_L} \quad (16)$$

Since $\beta^{t+1} = \Pi_B(\bar{\beta}^{t+1})$ is projection onto a convex set, by the property of Euclidean projection [Bub15], we have

$$\|\Delta^t\|_2 = \|\beta^t - \beta^*\|_2 \leq \|\bar{\beta}^t - \beta^*\|_2 = \|\bar{\Delta}^t\|_2. \quad (17)$$

Together with eq. (17), eq. (16) establishes global linear convergence of Δ^t .

We apply a union bound on T iterates. Since $1 - Td^{-3} \geq 1 - d^{-2}$ for sufficiently large d , we have

$$\|\Delta^T\|_2 \leq \left(1 - \frac{1}{8\rho}\right)^T \|\beta^*\|_2 + \frac{32\psi}{\mu_\alpha}$$

with probability at least $1 - d^{-2}$. Hence, we can achieve the final error

$$\|\hat{\beta} - \beta^*\|_2 = O(\psi/\mu_\alpha),$$

by setting $T = O\left(\rho \log\left(\frac{\mu_\alpha \|\beta^*\|_2}{\psi}\right)\right)$. □

B.1 Sparse linear regression

Proof of Corollary 4.1 and Corollary 4.3. With Proposition A.2, Proposition A.1 in hand, we prove the arbitrary corruption case Corollary 4.1, and the proof of heavy tailed distribution Corollary 4.3 is similar. We evaluate the RDC Definition 3.1 in Algorithm 1 for trimmed gradient estimator. With probability at least $1 - d^{-3}$, we have

$$\begin{aligned} \left| \langle \hat{\mathbf{G}}(\beta) - \mathbf{G}(\beta), \beta^{t-1} - \beta^* \rangle \right| &\stackrel{(i)}{\leq} \left\| \hat{\mathbf{G}}(\beta) - \mathbf{G}(\beta) \right\|_\infty \|\beta^{t-1} - \beta^*\|_1 \\ &\stackrel{(ii)}{\leq} \sqrt{k' + k} \left\| \hat{\mathbf{G}}(\beta) - \mathbf{G}(\beta) \right\|_\infty \|\beta^{t-1} - \beta^*\|_2 \\ &\stackrel{(iii)}{\leq} \sqrt{k' + k} \left(\sqrt{\|\beta - \beta^*\|_2^2 + \sigma^2} \left(\epsilon \log(nd) + \sqrt{\log d/n} \right) \right) \|\beta^{t-1} - \beta^*\|_2 \\ &\leq \sqrt{k' + k} \left(\epsilon \log(nd) + \sqrt{\log d/n} \right) (\|\beta - \beta^*\|_2 + \sigma) \|\beta^{t-1} - \beta^*\|_2 \\ &\leq (\alpha \|\beta - \beta^*\|_2 + \psi) \|\beta^{t-1} - \beta^*\|_2, \end{aligned}$$

where (i) follows from Holder inequality, (ii) follows from the sparsity of $\beta^{t-1} - \beta^*$ in Algorithm 1, (iii) follows from plugging in Proposition A.1, which yields $\alpha = \sqrt{k' + k} \left(\epsilon \log(nd) + \sqrt{\log d/n} \right)$, $\psi = \sigma \sqrt{k' + k} \left(\epsilon \log(nd) + \sqrt{\log d/n} \right)$.

We apply a union bound on T iterates, and $1 - Td^{-3} \geq 1 - d^{-2}$ for sufficiently large d . The condition $\alpha \leq \frac{1}{32}\mu_\alpha$ in Theorem 3.1 can be achieved if

$$n = \Omega\left(\frac{\rho^2 k \log d}{\mu_\alpha^2}\right), \text{ and } \epsilon = O\left(\frac{\mu_\alpha}{\rho \sqrt{k} \log(nd)}\right).$$

Since $\rho = \mu_L/\mu_\alpha$, and $\mu_L \geq 1$, these conditions can be expressed as

$$n = \Omega(\rho^4 k \log d), \text{ and } \epsilon = O\left(\frac{1}{\rho^2 \sqrt{k} \log(nd)}\right).$$

The final error can be expressed as

$$\left\| \widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^* \right\|_2 = O \left(\rho^2 \sigma \left(\underbrace{\epsilon \sqrt{k} \log(nd)}_{\text{robustness error}} + \underbrace{\sqrt{\frac{k \log d}{n}}}_{\text{statistical error}} \right) \right).$$

□

B.2 Sparse logistic regression

Proof of Corollary 4.2 and Corollary 4.4. We prove Corollary 4.2, and the proof of Corollary 4.4 is similar. With probability at least $1 - d^{-3}$, we have

$$\begin{aligned} \left| \left\langle \widehat{\mathbf{G}}(\boldsymbol{\beta}) - \mathbf{G}(\boldsymbol{\beta}), \boldsymbol{\beta}^{t-1} - \boldsymbol{\beta}^* \right\rangle \right| &\stackrel{(i)}{\leq} \sqrt{k' + k} \left(\epsilon \log(nd) + \sqrt{\log d/n} \right) \|\boldsymbol{\beta}^{t-1} - \boldsymbol{\beta}^*\|_2 \\ &\leq (\alpha \|\boldsymbol{\beta} - \boldsymbol{\beta}^*\|_2 + \psi) \|\boldsymbol{\beta}^{t-1} - \boldsymbol{\beta}^*\|_2, \end{aligned}$$

where (i) follows from the proof of Corollary 4.1 by using $\|\widehat{\mathbf{G}} - \mathbf{G}\|_\infty = O(\sqrt{\log d/n})$ in Proposition A.1, and $\alpha = 0$, $\psi = \sqrt{k' + k} (\epsilon \log(nd) + \sqrt{\log d/n})$.

Similar to the proof in sparse linear regression, this final error can be expressed as

$$\left\| \widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^* \right\|_2 = O \left(\rho^2 \left(\underbrace{\epsilon \sqrt{k} \log(nd)}_{\text{robustness error}} + \underbrace{\sqrt{\frac{k \log d}{n}}}_{\text{statistical error}} \right) \right).$$

□

C Sparsity recovery and sparse precision matrix estimation

C.1 Sparsity recovery guarantee

The same as the main text, we use $\text{supp}(\mathbf{v}, k)$ to denote top k indexes of \mathbf{v} with the largest magnitude. Let \mathbf{v}_{\min} denote the smallest absolute value of nonzero elements of \mathbf{v} .

Proof of Corollary 4.5. The sparsity recovery guarantee is similar to [YLZ18]. Since $\widehat{\boldsymbol{\beta}}$ is k' sparse ($k' \geq k$) by the definition of hard thresholding operator, we use $\widehat{\boldsymbol{\beta}}_k$ to denote $P_k(\widehat{\boldsymbol{\beta}})$. We use the technique proof by contradiction. If $\text{supp}(\widehat{\boldsymbol{\beta}}, k) \neq \text{supp}(\boldsymbol{\beta}^*)$, we at least have ℓ_2 error as $\boldsymbol{\beta}_{\min}^*$. Hence, $\boldsymbol{\beta}_{\min}^* \leq \left\| \widehat{\boldsymbol{\beta}}_k - \boldsymbol{\beta}^* \right\|_2 \stackrel{(i)}{\leq} 2 \left\| \widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^* \right\|_2 \stackrel{(ii)}{=} O \left(\rho^2 \sigma \left(\epsilon \sqrt{k} \log(nd) + \sqrt{\frac{k \log d}{n}} \right) \right)$, where (i) follows from the triangle inequality and definition of hard thresholding $\|\widehat{\boldsymbol{\beta}}_k - \boldsymbol{\beta}^*\|_2 \leq \|\widehat{\boldsymbol{\beta}}_k - \widehat{\boldsymbol{\beta}}\|_2 + \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2 \leq 2\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2$, and (ii) follows from the statistical guarantee in Corollary 4.1.

This contradicts with the $\boldsymbol{\beta}_{\min}$ -condition in Corollary 4.5, and hence we have the result in Corollary 4.5. □

Algorithm 2 Neighborhood Selection via Robust Hard Thresholding (Robust NS)

- 1: **Input:** Data samples $\{\mathbf{x}_i\}_{i=1}^m$.
 - 2: **Output:** The sparsity pattern estimation of Θ .
 - 3: **Parameters:** Hard thresholding parameter k' .
-
- 4: **for** each variable j , **do**
 - 5: Use X_j as response variable, and $\mathbf{X}_{(j)}$ as covariates.
 - 6: Run [Algorithm 1](#) with input $\{x_{ij}, \mathbf{x}_{i(j)}\}_{i=1}^m$. We set the parameter as k' , and the loss function as least square loss, and use trimmed gradient estimator.
 - 7: The output of [Algorithm 1](#) is denoted as $\hat{\beta}_j \in \mathbb{R}^{d-1}$.
 - 8: **end for**
 - 9: Aggregate the neighborhood support set of $\{\hat{\beta}_j\}_{j=1}^d$ via intersection or union.
-

C.2 Model selection for Gaussian graphical models

We then start to consider the sparsity recovery results for sparse precision matrix estimation – this is the part of [Corollary 4.6](#). We first use following notations for a Gaussian graphical model.

We use \mathbf{x}_i to denote the i -th samples of Gaussian graphical model, and X_j to denote the j -th random variable. Let (j) be the index set $\{1, \dots, j-1, j+1, \dots, d\}$. We use $\Sigma_{(j)} = \Sigma_{(j),(j)} \in \mathbb{R}^{(d-1) \times (d-1)}$ to denote the sub-matrix of covariance matrix Σ with both j -th row and j -th column removed, and use $\sigma_{(j)} \in \mathbb{R}^{d-1}$ to denote Σ 's j -th column with the diagonal entry removed. Also, we use $\theta_{(j)} \in \mathbb{R}^{d-1}$ to denote Θ 's j -th column with the diagonal entry removed. and $\Theta_{j,j} \in \mathbb{R}$ to denote the j -th diagonal element of Θ .

By basic probability computation, for each $j = 1, \dots, d$, the variable X_j conditioning $\mathbf{X}_{(j)}$ follows from a Gaussian distribution $\mathcal{N}(\mathbf{X}_{(j)}^\top \Sigma_{(j)}^{-1} \sigma_{(j)}, 1 - \sigma_{(j)}^\top \Sigma_{(j)}^{-1} \sigma_{(j)})$. Then we have the linear regression formulation $X_j = \mathbf{X}_{(j)}^\top \beta_j + \xi_j$, where $\beta_j = \Sigma_{(j)}^{-1} \sigma_{(j)}$ and $\xi_j \sim \mathcal{N}(0, 1 - \sigma_{(j)}^\top \Sigma_{(j)}^{-1} \sigma_{(j)})$. Notice the definition of precision matrix Θ , we have $\beta_j = -\theta_{(j)}/\Theta_{j,j}$, and $\Theta_{j,j} = 1/\text{Var}(\xi_j)$. Thus for the j -th variable, $\theta_{(j)}$ and β_j have the same sparsity pattern. Hence, the sparsity pattern of $\theta_{(j)}$ can be estimated through $\hat{\beta}_j$ via solving the optimization [eq. \(2\)](#) (Neighborhood Selection in [\[MB06\]](#)).

In [Algorithm 2](#), we robustify Neighborhood Selection by using Robust Hard Thresholding (with ℓ_2 loss and trimmed gradient estimator) to robustify [eq. \(2\)](#). In line 6, we use Robust Hard Thresholding to regress each variable against its neighbors. In line 9, the sparsity pattern of Θ can be estimated by aggregating the neighborhood support set of $\{\hat{\beta}_j\}_{j=1}^d$ via intersection or union. Similar to [Corollary 4.5](#), a θ_{\min} -condition guarantees consistent edge selection.

Proof of [Corollary 4.6](#). [Algorithm 2](#) iteratively uses [Algorithm 1](#) as a Neighborhood Selection approach for each variable. Hence, we can apply [Corollary 4.5](#) for each variable, and the sparsity patterns are the same according to $\theta_{(j)} = -\beta_j/\text{Var}(\xi_j)$. The stochastic noise term σ in sparse linear regression can be expressed as $1/\sqrt{\Theta_{j,j}}$. Hence, under the same condition as [Corollary 4.1](#), for each $j \in [d]$, we require a θ_{\min} -condition for $\theta_{(j)}$, $\theta_{(j),\min} = \Omega\left(\Theta_{j,j}^{1/2} \rho^2 \left(\epsilon \sqrt{k} \log(nd) + \sqrt{\frac{k \log d}{n}}\right)\right)$.

Using a union bound, we conclude that [Algorithm 2](#) is consistent in edge selection, with probability at least $1 - d^{-1}$. \square

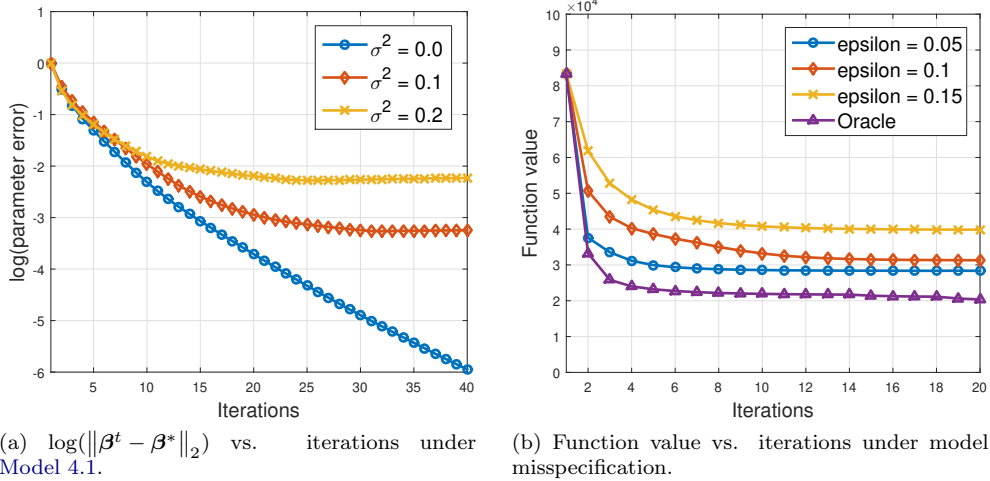


Figure 3: In the left plot, we use Algorithm 1 under Model 4.1 with different noise level σ^2 . In the right plot, we use Algorithm 1 under model misspecification with ϵ . The function value is defined as $F(\beta) = \sum_{i \in \bar{q}} (y_i - \mathbf{x}_i^\top \beta)^2$.

D Full experiments details

We study empirical performance of Robust Hard Thresholding (Algorithm 1 and Algorithm 2). And we present the complete details of experimental setup in Section 5.

D.1 Synthetic data – sparse linear models

We first consider the performance of Algorithm 1 under (generalized) linear models with ϵ -corrupted samples.

Sparse linear regression. In the first experiment, we consider an exact sparse linear regression model (Model 4.1). In this model, the stochastic noise $\xi \sim \mathcal{N}(0, \sigma^2)$, and we vary the noise level σ^2 in different simulations. We first generate authentic explanatory variables with parameters $k = 5, d = 1000, n = 300$, from a Gaussian distribution $\mathcal{N}(\mathbf{0}_d, \Sigma)$, where the covariance matrix Σ is a Toeplitz matrix with an exponential decay $\Sigma_{ij} = \exp^{-|i-j|}$. This design matrix is known to enjoy the RSC-condition [RWY10], which meets the requirement of Corollary 4.1. The entries of the k -sparse true parameter β^* are set to either $+1$ or -1 . Fixing the contamination level at $\epsilon = 0.1$, we set the covariates of the outliers as A , where A is a random ± 1 matrix of dimension $\frac{\epsilon}{1-\epsilon} \times d$, and the responses of outliers to $-A\beta^*$.

To show the performance of Algorithm 1 under different noise levels determined by σ^2 , we track the parameter error $\|\beta^t - \beta^*\|_2$ in each iteration. In the left plot of Figure 3, Algorithm 1 shows linear convergence, and the error curves flatten out at the level of the final error, which is consistent with our theory. Furthermore, Algorithm 1 can achieve machine precision when $\sigma^2 = 0$, which means exact recovery of β^* .

Misspecified model. For the second experiment, we use a sparse linear regression with model misspecification – the underlying authentic samples do not follow a linear model. We use the same Toeplitz covariates and true parameter β^* , but the corresponding y_i 's are calculated as $y_i = \sum_{j=1}^d \mathbf{x}_{ij}^3 \beta_j^*$. Although this is a non-linear function, sparse linear regression on these authentic

samples can still recover the support, as the cubic function is monotone and β^* is sparse. We generate outliers using the same distribution as the first experiment, but with a different fraction of corruptions ϵ .

For simplicity, we track the function evaluated on all authentic samples $F(\beta) = \sum_{i \in \mathcal{G}} (y_i - \mathbf{x}_i^\top \beta)^2$. In the right plot of [Figure 3](#), we show the performance of [Algorithm 1](#) under different ϵ , and the oracle curve means using IHT *only on* authentic samples. The right plot has similar convergence under different values of corrupted fraction ϵ , and shows the robustness of [Algorithm 1](#) without assuming an underlying linear model.

D.2 Robust M -estimators via Robust Hard Thresholding

Classical robust M -estimators [[Loh17](#)] (such as empirical risk minimization using Huber loss) are widely used in robust statistics in the case where the error distribution is heavy tailed or when there are arbitrary outliers *only in the response variables*. In the high dimensional setting, given ϵ -corrupted samples [Definition 2.1](#), we can use

$$\min_{\beta \in \Omega} \mathbb{E}_{i \in_u \mathcal{S}} \ell_i(\beta; \mathbf{z}_i), \quad \text{s.t. } \|\beta\|_0 \leq k,$$

where $\ell_i(\beta; \mathbf{z}_i)$ can be chosen as Huber loss with parameter δ :

$$H_\delta(\beta; \mathbf{z}_i) = \begin{cases} \frac{1}{2}(y_i - \mathbf{x}_i^\top \beta)^2 & \text{for } |y_i - \mathbf{x}_i^\top \beta| \leq \delta, \\ \delta |y_i - \mathbf{x}_i^\top \beta| - \frac{1}{2}\delta^2 & \text{otherwise.} \end{cases}$$

[[Loh17](#)] studied robust M -estimators in high dimensions, and proposed a composite optimization using $\|\beta\|_1$ instead of $\|\beta\|_0$. They established local convergence guarantee for this composite optimization procedure, using a local RSC condition in a neighborhood around β^* . Yet their results do not trivially extend to settings with arbitrarily corrupted covariates.

In our experiments, we use Huber loss in Robust Hard Thresholding to deal with heavy-tailed error distribution. In addition to heavy-tailed noise, ϵ -fraction of $\{y_i, \mathbf{x}_i\}_{i=1}^n$ are still arbitrarily corrupted.

For the experiments, we use the same Toeplitz covariates and true parameter β^* as in previous experiments on sparse linear models with fixed dimension parameters $k = 5, d = 1000, n = 300$. The error distribution is a Cauchy distribution, which is a special case model misspecification, as it doesn't meet the sub-Gaussian requirement in [Model 4.1](#). For different contamination levels, we set the covariates of the outliers as A , where A is a random ± 1 matrix of dimension $\frac{\epsilon}{1-\epsilon} \times d$, and the responses of outliers to $-A\beta^*$.

Empirically, we observe linear convergence, and this is shown in [Figure 4](#). This linear convergence results validates the local RSC condition proposed in [[Loh17](#)], and we can still achieve this even with ϵ -fraction of corrupted covariates.

D.3 Sparse logistic regression

For binary classification problem, we generate samples from a sparse LDA problem, where the distributions of the explanatory variables conditioned on the response variables follow multivariate Gaussian distributions with the same covariance matrix but different means.




Figure 4: $\log \|\beta^t - \beta^*\|_2$ vs. iterations for Robust Hard Thresholding using Huber loss in the sparse linear model where the error distribution is a heavy-tailed Cauchy distribution.

We generate authentic samples \mathbf{x}_i from a Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}_+, \mathbf{I}_d)$ if $y_i = +1$, and another distribution $\mathcal{N}(\boldsymbol{\mu}_-, \mathbf{I}_d)$ if $y_i = -1$. The parameters are fixed $k = 5, d = 1000, n = 300$. We set $\boldsymbol{\mu}_+ = \mathbf{1}_d + \mathbf{v}$, where \mathbf{v} is k -sparse and its entries are set to be either $+1/\sqrt{k}$ or $-1/\sqrt{k}$. And we set $\boldsymbol{\mu}_- = \mathbf{1}_d - \mathbf{v}$. The Bayes classifier is $\beta^* = 2\mathbf{v}$. This is a special case of Model 4.2, and it is known that sparse logistic regression attains fast classification error rates [LPR15]. We then set the covariates of the outliers as A , where A is a matrix of dimension $\frac{\epsilon}{1-\epsilon} \times d$, where the entries are random ± 3 . The responses of outliers follow the distribution $\Pr(y_i | \mathbf{x}_i) = 1/(1 + \exp(y_i \mathbf{x}_i^\top \beta^*))$, which is exactly the opposite of Model 4.2.

We run Algorithm 1 with logistic loss under different levels of outlier fraction ϵ . In the left plot of Figure 5, we observe similar linear convergence as sparse linear regression. This is consistent with Corollary 4.2 for sparse logistic regression, and it is clear that we cannot exactly recover β^* unless the number of samples n is infinite.

We then compare Algorithm 1 with the Trimmed Lasso estimator for sparse logistic regression [YLA18]. Although they also use a trimming technique, their algorithm is totally different from Algorithm 1, as we use coordinate-wise trimmed mean estimator for gradients in hard thresholding, but they trim samples in each iteration according to the each sample’s loss. Under the same sparse LDA model, we set $k = \sqrt{d}, n = 15k$. In simulation, we increase d , and plot classification error (averaged over 50 trials on authentic test set) for different $\epsilon = 0.1, 0.2$. The right plot of Figure 5 shows that Robust Hard Thresholding is better than Trimmed Lasso.

D.4 Synthetic data – Gaussian graphical model

We generate Gaussian graphical model samples by huge [ZLR+12]. We choose the “cluster” sparsity pattern, where the clustering parameters are default values in the package where the number of clusters in the graph is $d/20$, the probability that a pair of nodes within a cluster are connected is 0.3, and there are no edges between nodes within different clusters. The off-diagonal elements of the precision matrix is denoted as v , which is an experiment parameter for SNR.

We then add an additional $\frac{\epsilon}{1-\epsilon}$ fraction of samples sampled from another distribution. Following the experimental design in [YL15, WG17], each outlier is generated by a mixture of d -dimensional

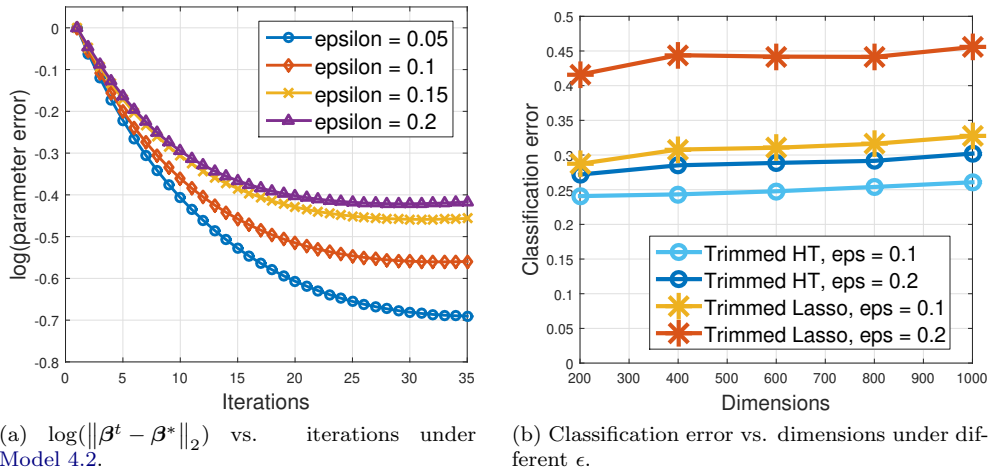


Figure 5: The left plot demonstrates linear convergence of [Algorithm 1](#) for sparse logistic regression. The right plot compares our Robust Hard Thresholding (denoted as Trimmed HT) to Trimmed Lasso [\[YLA18\]](#) in binary classification problem.

Gaussian distributions $\frac{1}{2}\mathcal{N}(\boldsymbol{\mu}^o, \boldsymbol{\Sigma}^o) + \frac{1}{2}\mathcal{N}(-\boldsymbol{\mu}^o, \boldsymbol{\Sigma}^o)$, where $\boldsymbol{\mu}^o = (1.5, 1.5, \dots, 1.5)^\top$, and $\boldsymbol{\Sigma}^o = \mathbf{I}_d$. We compare [Algorithm 2](#) with other existing methods: Trimmed GLasso [\[YLA18\]](#), RCLIME [\[WG17\]](#), Skeptic [\[LHY⁺12b\]](#), and Spearman [\[LT18\]](#). The latter two are based on robustifying the covariance matrix, and then using standard graphical model selection algorithms such as GLasso or CLIME. To directly compare these methods, we use CLIME for both of them.

To evaluate model selection performance, we use receiver operating characteristic (ROC) curves to compare our method to others over the full regularization paths. We generate regularization paths for other robust algorithms by tuning the λ in CLIME and GLasso. For [Algorithm 2](#), we explicitly tune different sparsity level k' to generate the regularization path.

We set $\epsilon = 0.1$, and vary (n, d) , and the SNR parameter v for off-diagonal elements. We use different $(n, d) = (100, 100), (200, 200)$. For different off-diagonal values, we set $v = 0.3$ (Low SNR), and $v = 0.6$ (High SNR). We show ROC curves to demonstrate model selection performance in [Figure 6](#). For the entire regularization path, our algorithm (denoted as Robust NS) has a better ROC compared to other algorithms.

In particular, Robust NS outperforms other methods with higher true positive rate when the false positive rate is small. This is the case where we use smaller hard thresholding sparsity in [Algorithm 2](#), and larger regularization parameter for $\|\Theta\|_1$ other methods based on GLasso and CLIME. This validates our theory in [Corollary 4.6](#), which guarantees sparsity recovery when hard thresholding hyper-parameter k' is suitably chosen to match $\boldsymbol{\beta}^*$'s sparsity k .

D.5 Real data experiments

Here, we present details of the experiment using US equities data [\[ZLR⁺12\]](#). We preprocess it by taking log-transformation and calculate the corresponding daily returns. Obvious outliers are removed by winsorizing each variable so that all samples are within five times the winsorized standard deviation from the winsorized mean. After preprocessing, we present example histograms and QQ plots from




Figure 6: ROC curves of different methods on cluster graphs with arbitrary corruptions. The curve Robust NS denotes Algorithm 2, and Oracle NS denotes the neighborhood selection Lasso only on authentic data.

the Information Technology sector. In Figure 7, we list the histograms of two typical companies in this sector. As we can see from Figure 8, even after preprocessing on these stock prices, they are still highly non-normal and heavy tailed. We do not add any manual outliers as financial data is already heavy tailed and have many outliers [dP18]. We also compare Algorithm 2 with the baseline NS approach (without consideration for corruptions or outliers).

We limit the number of edges to 2,000 for both methods. The cluster colored by purple denotes the Information Technology sector. In Figure 9, we can easily separate different clusters by using Robust NS. However, the Vanilla NS approach *cannot* distinguish the sector Information Technology (purple). Furthermore, we can observe that stocks from Information Technology (colored by purple) are much better clustered by Algorithm 2.

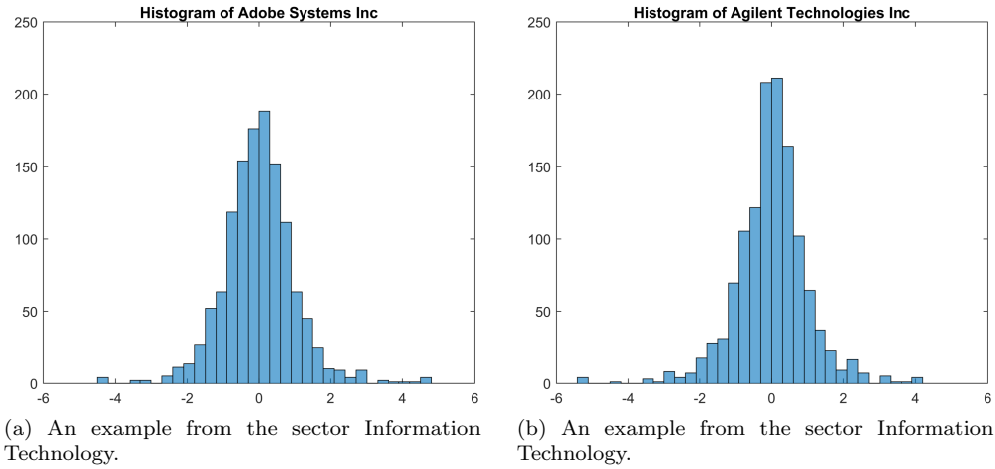


Figure 7: We winsorize the daily return of log price so that all samples are within five winsorized standard deviation from the winsorized mean. After preprocessing, we show histograms of the daily returns.

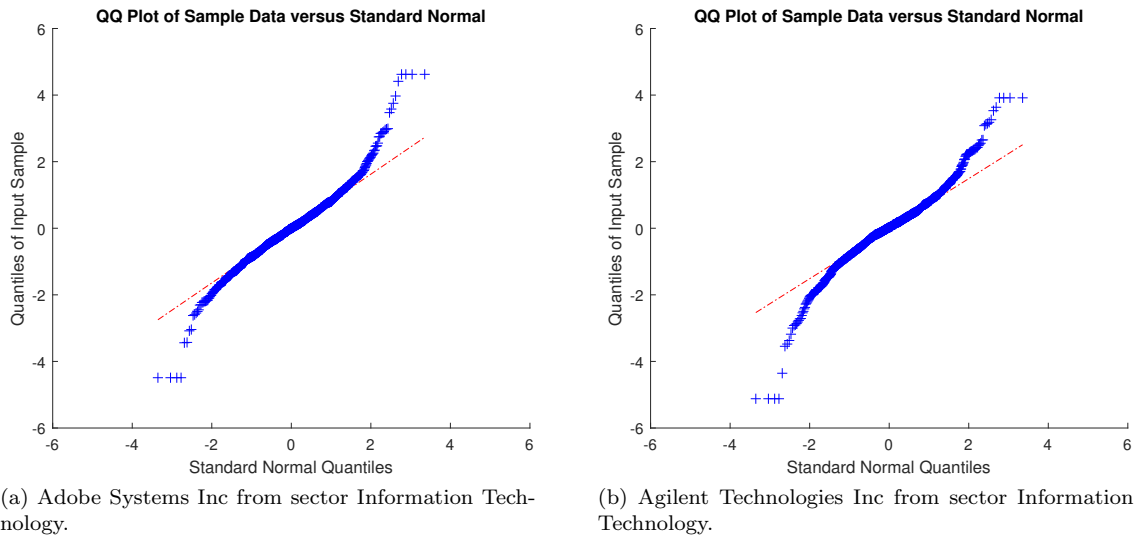


Figure 8: After the same preprocessing, we present the QQ plot of the daily returns versus standard normal.

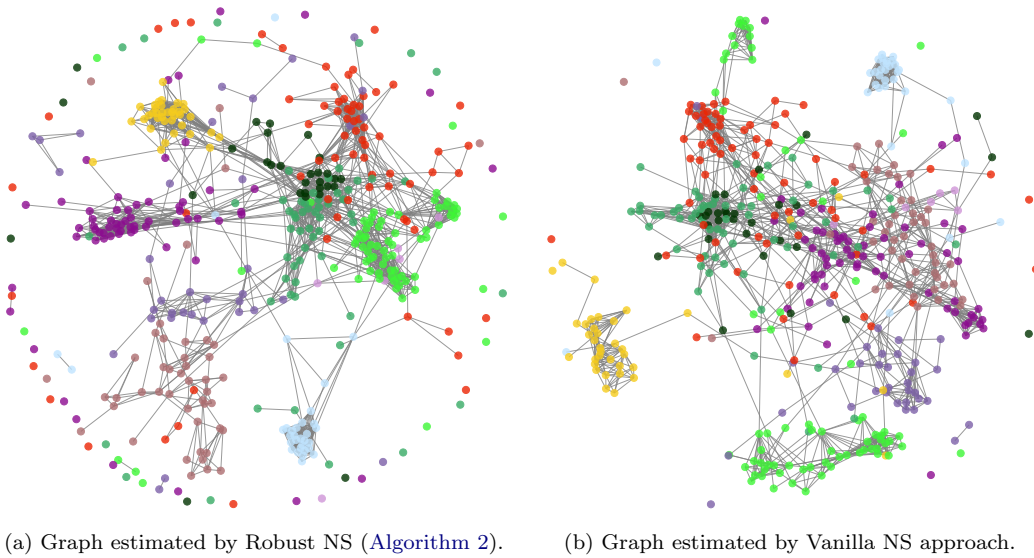


Figure 9: Graph estimated from the S&P 500 stock data by Algorithm 2 and Vanilla NS approach. Variables are colored according to their sector. In particular, the stocks from sector Information Technology are colored as purple.