

SEVER: A Robust Meta-Algorithm for Stochastic Optimization

Ilias Diakonikolas ^{*}
 CS, USC
 diakonik@usc.edu

Gautam Kamath [†]
 EECS & CSAIL, MIT
 g@csail.mit.edu

Daniel M. Kane [‡]
 CSE & Math, UCSD
 dakane@cs.ucsd.edu

Jerry Li [§]
 EECS & CSAIL, MIT
 jerryzli@mit.edu

Jacob Steinhardt [¶]
 CS, Stanford University
 jsteinha@stanford.edu

Alistair Stewart ^{||}
 CS, USC
 stewart.al@gmail.com

March 8, 2018

Abstract

In high dimensions, most machine learning methods are brittle to even a small fraction of structured outliers. To address this, we introduce a new meta-algorithm that can take in a *base learner* such as least squares or stochastic gradient descent, and harden the learner to be resistant to outliers. Our method, SEVER, possesses strong theoretical guarantees yet is also highly scalable—beyond running the base learner itself, it only requires computing the top singular vector of a certain $n \times d$ matrix. We apply SEVER on a drug design dataset and a spam classification dataset, and find that in both cases it has substantially greater robustness than several baselines. On the spam dataset, with 1% corruptions, we achieved 7.4% test error, compared to 13.4% – 20.5% for the baselines, and 3% error on the uncorrupted dataset. Similarly, on the drug design dataset, with 10% corruptions, we achieved 1.42 mean-squared error test error, compared to 1.51-2.33 for the baselines, and 1.23 error on the uncorrupted dataset.

1 Introduction

Learning in the presence of outliers is a ubiquitous challenge in machine learning; nevertheless, most machine learning methods are very sensitive to outliers in high dimensions. The focus of this work is on designing algorithms that are outlier robust while remaining competitive in terms of accuracy and running time.

We highlight two motivating applications. The first is biological data (such as gene expression data), where mislabeling or measurement errors can create systematic outliers [RPW⁺02, LAT⁺08] that require painstaking manual effort to remove [PLJD10]. Detecting outliers in such settings is often important either because the outlier observations are of interest themselves or because they might contaminate the downstream statistical analysis. The second motivation is machine learning security, where outliers can be introduced through *data poisoning* attacks [BNJT10] in which an adversary inserts fake data into the training set (e.g., by creating a fake user account). Recent work has shown that for high-dimensional datasets, even a small fraction of outliers can substantially degrade the learned model [BNL12, NPXNR14, KL17, SKL17].

^{*}Supported by NSF Award CCF-1652862 (CAREER) and a Sloan Research Fellowship.

[†]Supported by NSF Award CCF-1617730, CCF-1650733, CCF-1741137, and ONR N00014-12-1-0999.

[‡]Supported by NSF Award CCF-1553288 (CAREER) and a Sloan Research Fellowship.

[§]Supported by NSF Award CCF-1453261 (CAREER), CCF-1565235, a Google Faculty Research Award, and an NSF Graduate Research Fellowship.

[¶]Supported by a Fannie & John Hertz Foundation Fellowship, a NSF Graduate Research Fellowship, and a Future of Life Institute grant.

^{||}Supported by a USC startup grant.

Authors are in alphabetical order.

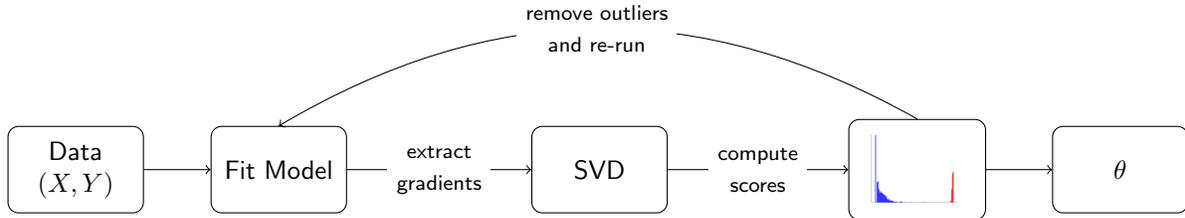


Figure 1: Illustration of the SEVER pipeline. We first use any machine learning algorithm to fit a model to the data. Then, we extract gradients for each data point at the learned parameters, and take the singular value decomposition of the gradients. We use this to compute an outlier score for each data point. If we detect outliers, we remove them and re-run the learning algorithm; otherwise, we output the learned parameters.

Crucially, in both the biological and security settings above, the outliers are not “random” but are instead highly correlated, and could have a complex internal structure that is difficult to model. This leads us to the following conceptual question underlying the present work: *Can we design training algorithms that are robust to the presence of an ε -fraction of arbitrary (and potentially adversarial) outliers?*

Estimation in the presence of outliers is a prototypical goal in robust statistics and has been systematically studied since the pioneering work of Tukey [Tuk60]. Nevertheless, until recently, all known efficient estimators suffered from large errors in the presence of a small fraction of outliers, even for the most basic high-dimensional tasks that are common in machine learning. Recent work in theoretical computer science has resulted in efficient robust estimators for classical problems such as linear classification [KLS09, ABL14], mean and covariance estimation [DKK⁺16, LRV16], clustering [CSV17], and regression [BJK15, BJKK17, BDLS17]. (See Section 1.1 for a detailed discussion of prior and related work.)

Despite this recent algorithmic progress, the promise of efficient high-dimensional robust estimation is yet to be realized; indeed, the aforementioned results generally suffer from one of two shortcomings—either they use sophisticated convex optimization algorithms that do not scale to large datasets, or they are tailored to specific problems of interest or specific distributional assumptions on the data, and hence do not have good accuracy on real data.

In this work, we address these shortcomings. We propose an algorithm, SEVER, that is:

- **Robust:** it can handle arbitrary outliers with only a small increase in error, even in high dimensions.
- **General:** it can be applied to most common learning problems including regression and classification, and handles non-convex models such as neural networks.
- **Practical:** the algorithm can be implemented with standard machine learning libraries.

At a high level, our algorithm (depicted in Figure 1 and described in detail in Section 2.2) is a simple “plug-in” outlier detector—first, run whatever learning procedure would be run normally (e.g., least squares in the case of linear regression). Then, consider the matrix of gradients at the optimal parameters, and compute the top singular vector of this matrix. Finally, remove any points whose projection onto this singular vector is too large (and re-train if necessary).

Despite its simplicity, our algorithm possesses strong theoretical guarantees: As long as the data is not too heavy-tailed, SEVER is provably robust to outliers—see Section 2 for detailed statements of the theory. At the same time, we show that our algorithm works very well in practice and outperforms a number of natural baseline outlier detectors. In line with our original motivating biological and security applications, we implement our method on two tasks—a linear regression task for predicting protein activity levels [OSB⁺18], and a spam classification task based on e-mails from the Enron corporation [MAP06]. Even with a small fraction of outliers, baseline methods perform extremely poorly on these datasets; for instance, on the Enron spam dataset with a 1% fraction of outliers, baseline errors range from 13.4% to 20.5%, while SEVER incurs only 7.3% error (in comparison, the error is 3% in the absence of outliers). Similarly, on the drug design

dataset, with 10% corruptions, we achieved 1.42 mean-squared error test error, compared to 1.51-2.33 for the baselines, and 1.23 error on the uncorrupted dataset.

1.1 Comparison to Prior Work

In the context of supervised learning, [KLS09] gave the first efficient algorithms for robustly learning origin-centered linear separators with respect to the 0-1 loss under the assumption that the distribution of the good data is isotropic and log-concave. Subsequently, [ABL14] obtained an improved robust algorithm that achieves nearly tight error bounds for this learning problem.

In the context of unsupervised learning, two concurrent works [DKK⁺16, LRV16] gave the first provably robust and efficiently computable estimators for several basic high-dimensional unsupervised tasks, including robust mean and covariance estimation. Since the dissemination of [DKK⁺16, LRV16], there has been a flurry of research activity on algorithmic robust estimation in a variety of high-dimensional settings, including learning graphical models [DKS16], understanding computation-robustness tradeoffs [DKS17c, DKK⁺18], establishing connections to PAC learning [DKS17a], tolerating more noise by outputting a list of hypotheses [CSV17, MV17, DKS17b], robust estimation of discrete structures [Ste17, QV17, SCV18], and robust estimation via sum-of-squares [KS17b, HL17, KS17a].

Despite this recent theoretical progress, prior works focus on designing specialized algorithms for specific settings (such as mean estimation or linear classification) rather than on designing general algorithms. The only exception is [CSV17], which provides a robust meta-algorithm for stochastic convex optimization in a similar setting to ours. However, that algorithm (i) requires solving a large semidefinite program and (ii) incurs a significant loss in performance relative to standard training *even in the absence of outliers*. On the other hand, [DKK⁺17] provide a practical implementation of the robust mean and covariance estimation algorithms of [DKK⁺16], but it is not clear how to extend their approach to a broader class of learning tasks.

In the concrete settings of regression and classification considered in our experiments, there has been a substantial number of related works. A number of papers [NTN11, NT13, BJK15, BJKK17] have proposed efficient algorithms for a type of robust linear regression. However, these works consider a restrictive corruption model that only allows adversarial corruptions to the responses (but not the covariates). On the other hand, [BDLS17] studies (sparse) linear regression and, more broadly, generalized linear models (GLMs) under a robustness model very similar to the one considered here. The main issues with this algorithm are that (i) it requires running the ellipsoid method (hence does not scale) and (ii) it crucially assumes Gaussianity of the covariates, which is unlikely to hold in practice.

In a related direction, [SKL17] study outlier detectors for linear classification, and provide a method for either numerically certifying that a given outlier detector imparts robustness to the learning algorithm, or for providing an attack if the learner is not robust. They show that the outlier detector they consider is very brittle in high dimensions, motivating the need for the robust algorithms presented in the current work.

Concurrent Works. Contemporaneous work and independent work [PSBR18] has obtained a robust algorithm for stochastic convex optimization by combining gradient descent with robust mean estimation. For the case of linear regression, [DKS18] provide efficient robust algorithms with near-optimal error guarantees under various distributional assumptions and establish matching computational-robustness tradeoffs.

2 Framework and Algorithm

In this section, we describe our formal framework as well as the SEVER algorithm.

2.1 Formal Setting

We will consider stochastic optimization tasks, where there is some true distribution p^* over functions $f : \mathcal{H} \rightarrow \mathbb{R}$, and our goal is to find a parameter vector $w^* \in \mathcal{H}$ minimizing $\bar{f}(w) \stackrel{\text{def}}{=} \mathbb{E}_{f \sim p^*}[f(w)]$. Here we assume $\mathcal{H} \subseteq \mathbb{R}^d$ is a space of possible parameters. As an example, we consider linear regression, where

$f(w) = \frac{1}{2}(w \cdot x - y)^2$ for (x, y) drawn from the data distribution; or support vector machines, where $f(w) = \max\{0, 1 - y(w \cdot x)\}$.

To help us learn the parameter vector w^* , we have access to a *training set* of n functions $f_{1:n} \stackrel{\text{def}}{=} \{f_1, \dots, f_n\}$. (For linear regression, we would have $f_i(w) = \frac{1}{2}(w \cdot x_i - y_i)^2$, where (x_i, y_i) is an observed data point.) However, unlike the classical (uncorrupted) setting where we assume that $f_1, \dots, f_n \sim p^*$, we allow for an ε -fraction of the points to be arbitrary outliers:

Definition 2.1 (ε -contamination model). Given $\varepsilon > 0$ and a distribution p^* over functions $f : \mathcal{H} \rightarrow \mathbb{R}$, data is generated as follows: first, n clean samples f_1, \dots, f_n are drawn from p^* . Then, an *adversary* is allowed to inspect the samples and replace any εn of them with arbitrary samples. The resulting set of points is then given to the algorithm. We will call such a set of samples ε -*corrupted* (with respect to p^*).

In the ε -contamination model, the adversary is allowed to both add and remove points. Our theoretical results hold in this strong robustness model. Our experimental evaluation uses corrupted instances in which the adversary is only allowed to add corrupted points. Additive corruptions essentially correspond to Huber’s contamination model [Hub64] in robust statistics.

Finally, we will often assume access to a black-box learner, which we denote by \mathcal{L} , which takes in functions f_1, \dots, f_n and outputs a parameter vector $w \in \mathcal{H}$. We want to stipulate that \mathcal{L} approximately minimizes $\frac{1}{n} \sum_{i=1}^n f_i(w)$. For this purpose, we introduce the following definition:

Definition 2.2 (γ -approximate critical point). Given a function $f : \mathcal{H} \rightarrow \mathbb{R}$, a γ -approximate critical point of f , is a point $w \in \mathcal{H}$ so that for all unit vectors v where $w + \delta v \in \mathcal{H}$ for arbitrarily small positive δ , we have that $v \cdot \nabla f(w) \geq -\gamma$.

Essentially, the above definition means that the value of f cannot be decreased much by changing the input w locally, while staying within the domain. The condition enforces that moving in any direction v either causes us to leave \mathcal{H} or causes f to decrease at a rate at most γ . It should be noted that when $\mathcal{H} = \mathbb{R}^d$, our above notion of approximate critical point reduces to the standard notion of approximate stationary point (i.e., a point where the magnitude of the gradient is small).

We are now ready to define the notion of a γ -*approximate* learner:

Definition 2.3 (γ -approximate learner). A learning algorithm \mathcal{L} is called γ -*approximate* if, for any functions $f_1, \dots, f_n : \mathcal{H} \rightarrow \mathbb{R}$ each bounded below on a closed domain \mathcal{H} , the output $w = \mathcal{L}(f_{1:n})$ of \mathcal{L} is a γ -approximate critical point of $f(x) := \frac{1}{n} \sum_{i=1}^n f_i(x)$.

In other words, \mathcal{L} always finds an approximate critical point of the empirical learning objective. We note that most common learning algorithms (such as stochastic gradient descent) satisfy the γ -approximate learner property.

2.2 Algorithm and Theory

As outlined in Figure 1, our algorithm works by post-processing the gradients of a black-box learning algorithm. The basic intuition is as follows: we want to ensure that the outliers do not have a large effect on the learned parameters. Intuitively, for the outliers to have such an effect, their corresponding gradients should be (i) large in magnitude and (ii) systematically pointing in a specific direction. We can detect this via singular value decomposition—if both (i) and (ii) hold then the outliers should be responsible for a large singular value in the matrix of gradients, which allows us to detect and remove them.

This is shown more formally via the pseudocode in Algorithm 1.

Theoretical Guarantees. Our first theoretical result says that as long as the data is not too heavy-tailed, SEVER will find an approximate critical point of the true function \bar{f} , even in the presence of outliers.

Theorem 2.1. *Suppose that functions $f_1, \dots, f_n, \bar{f} : \mathcal{H} \rightarrow \mathbb{R}$ are bounded below on a closed domain \mathcal{H} , and suppose that they satisfy the following deterministic regularity conditions: There exists a set $I_{\text{good}} \subseteq [n]$ with $|I_{\text{good}}| \geq (1 - \varepsilon)n$ and $\sigma > 0$ such that*

Algorithm 1 SEVER($f_{1:n}, \mathcal{L}, \sigma$)

- 1: **Input:** Sample functions $f_1, \dots, f_n : \mathcal{H} \rightarrow \mathbb{R}$, bounded below on a closed domain \mathcal{H} , γ -approximate learner \mathcal{L} , and parameter $\sigma \in \mathbb{R}_+$.
 - 2: Initialize $S \leftarrow \{1, \dots, n\}$.
 - 3: **repeat**
 - 4: $w \leftarrow \mathcal{L}(\{f_i\}_{i \in S})$. \triangleright Run approximate learner on points in S .
 - 5: Let $\widehat{\nabla} = \frac{1}{|S|} \sum_{i \in S} \nabla f_i(w)$.
 - 6: Let $G = [\nabla f_i(w) - \widehat{\nabla}]_{i \in S}$ be the $|S| \times d$ matrix of centered gradients.
 - 7: Let v be the top right singular vector of G .
 - 8: Compute the vector τ of outlier scores defined via $\tau_i = \left((\nabla f_i(w) - \widehat{\nabla}) \cdot v \right)^2$.
 - 9: $S' \leftarrow S$
 - 10: $S \leftarrow \text{FILTER}(S', \tau, \sigma)$ \triangleright Remove some i 's with the largest scores τ_i from S ; see Algorithm 2.
 - 11: **until** $S = S'$.
 - 12: Return w .
-

Algorithm 2 FILTER(S, τ, σ)

- 1: **Input:** Set $S \subseteq [n]$, vector τ of outlier scores, and parameter $\sigma \in \mathbb{R}_+$.
 - 2: If $\sum_i \tau_i \leq c \cdot \sigma$, for some constant $c > 1$, return S \triangleright We only filter out points if the variance is larger than an appropriately chosen threshold.
 - 3: Draw T from the uniform distribution on $[0, \max_i \tau_i]$.
 - 4: Return $\{i \in S : \tau_i < T\}$.
-

$$(i) \text{Cov}_{I_{\text{good}}}[\nabla f_i(w)] \preceq \sigma^2 I, w \in \mathcal{H},$$

$$(ii) \|\nabla \hat{f}(w) - \nabla \bar{f}(w)\|_2 \leq \sigma \sqrt{\varepsilon}, w \in \mathcal{H}, \text{ where } \hat{f} \stackrel{\text{def}}{=} (1/|I_{\text{good}}|) \sum_{i \in I_{\text{good}}} f_i.$$

Then our algorithm SEVER applied to f_1, \dots, f_n, σ returns a point $w \in \mathcal{H}$ that, with probability at least 9/10, is a $(\gamma + O(\sigma\sqrt{\varepsilon}))$ -approximate critical point of \bar{f} .

The key take-away from Theorem 2.1 is that the error guarantee has no dependence on the underlying dimension d . In contrast, most natural algorithms incur an error that grows with d , and hence have poor robustness in high dimensions.

We show that under some niceness assumptions on p^* , the deterministic regularity conditions are satisfied with high probability with polynomially many samples:

Proposition 2.2 (Informal). *Let $\mathcal{H} \subset \mathbb{R}^d$ be a closed bounded set with diameter at most r . Let p^* be a distribution over functions $f : \mathcal{H} \rightarrow \mathbb{R}$ and $\bar{f} = \mathbb{E}_{f \sim p^*}[f]$. Suppose that for each $w \in \mathcal{H}$ and unit vector v we have $\mathbb{E}_{f \sim p^*}[(v \cdot (\nabla f(w) - \bar{f}(w)))^2] \leq \sigma^2$. Under appropriate Lipschitz and smoothness assumptions, for $n = \Omega(d \log(r/(\sigma\varepsilon))/(\sigma^2\varepsilon))$, an ε -corrupted set of functions drawn i.i.d. from p^* , f_1, \dots, f_n with high probability satisfy conditions (i) and (ii).*

The reader is referred to Proposition B.5 in the appendix for a detailed formal statement.

While Theorem 2.1 is very general and holds even for non-convex loss functions, we might in general hope for more than an approximate critical point. In particular, for convex problems, we can guarantee that we find an approximate global minimum. This follows as a corollary of Theorem 2.1:

Corollary 2.3. *Suppose that $f_1, \dots, f_n : \mathcal{H} \rightarrow \mathbb{R}$ satisfy the regularity conditions (i) and (ii), and that \mathcal{H} is convex with ℓ_2 -radius r . Then, with probability at least 9/10, the output of SEVER satisfies the following:*

$$(i) \text{ If } \bar{f} \text{ is convex, the algorithm finds a } w \in \mathcal{H} \text{ such that } \bar{f}(w) - \bar{f}(w^*) = O((\sigma\sqrt{\varepsilon} + \gamma)r).$$

$$(ii) \text{ If } \bar{f} \text{ is } \xi\text{-strongly convex, the algorithm finds a } w \in \mathcal{H} \text{ such that } \bar{f}(w) - \bar{f}(w^*) = O((\varepsilon\sigma^2 + \gamma^2)/\xi).$$

Practical Considerations. For our theory to hold, we need to use the randomized filtering algorithm shown in Algorithm 2 (which is essentially the robust mean estimation algorithm of [DKK⁺17]), and filter until the stopping condition in line 11 of Algorithm 1 is satisfied. However, in practice we found that the following simpler algorithm worked well: in each iteration simply remove the top p fraction of outliers according to the scores τ_i , and instead of using a specific stopping condition, simply repeat the filter for r iterations in total. This is the version of SEVER that we use in our experiments in Section 3.

Concrete Applications We also provide several concrete applications of our general theorem, particularly involved with optimization problems related to learning generalized linear models. In this setting, we are given a set of pairs (X, Y) where $X \in \mathbb{R}^d$ and Y is in some (usually discrete) set. One then tries to find some vector w that minimizes some appropriate loss function $L(w, (X, Y)) = \sigma_Y(w \cdot X)$. For example, the standard hinge-loss has $Y \in \{\pm 1\}$ and $L = \max(0, 1 - Y(w \cdot X))$. Similarly, the logistic loss function is $-\log(1 + \exp(-Y(w \cdot X)))$. In both cases, we show that an approximate minimizer to the empirical loss function can be found with a near-optimal number of samples, even under ε -corruptions (for exact theorem statements see Theorems E.8 and E.11).

Theorem 2.4 (Informal Statement). *Let $D_{X,Y}$ be a distribution over $\mathbb{R}^d \times \{\pm 1\}$ so that $\mathbb{E}[XX^T] \preceq I$ and so that not too many X values lie near any hyperplane. Let $(X_1, Y_1), \dots, (X_n, Y_n)$ be $n = \tilde{O}(d/\varepsilon)$ ε -corrupted samples from $D_{X,Y}$. Let L be either the hinge loss or logistic loss function. Then there exists a polynomial time algorithm that with probability 9/10 returns a vector w that minimizes $\mathbb{E}_{(X,Y) \sim D_{X,Y}}(L(w, (X, Y)))$ up to an additive $\tilde{O}(\varepsilon^{1/4})$ error.*

Another application allows us to use the least-squares loss function ($L(w, (X, Y)) = (Y - w \cdot X)^2$) to perform linear regression under somewhat more restrictive assumptions (see Theorem E.2 for the full statement):

Theorem 2.5 (Informal Statement). *Let $D_{X,Y}$ be a distribution over $\mathbb{R}^d \times \mathbb{R}$ where $Y = w^* \cdot X + e$ for some independent e with mean 0 and variance 1. Assume furthermore, that $\mathbb{E}[XX^T] \preceq I$ and that X has bounded fourth moments. Then there exists an algorithm that given $O(d^5/\varepsilon^2)$ ε -corrupted samples from D , computes a value $w \in \mathbb{R}^d$ so that with high probability $\|w - w^*\|_2 = O(\sqrt{\varepsilon})$.*

2.3 Overview of Sever and its Analysis

For simplicity of the exposition, we restrict ourselves to the important special case where the functions involved are convex. We have a probability distribution p^* over convex functions on some convex domain $\mathcal{H} \subseteq \mathbb{R}^d$ and we wish to minimize the function $\bar{f} = \mathbb{E}_{f \sim p^*}[f]$. This problem is well-understood in the absence of corruptions: Under mild assumptions, if we take sufficiently many samples from p^* , their average \hat{f} approximates \bar{f} pointwise with high probability. Hence, we can use standard methods from convex optimization to find an approximate minimizer for \hat{f} , which will in turn serve as an approximate minimizer for \bar{f} .

In the robust setting, stochastic optimization becomes quite challenging: Even for the most basic special cases of this problem (e.g., mean estimation, linear regression) a *single* adversarially corrupted sample can substantially change the location of the minimum for \hat{f} . Moreover, naive outlier removal methods can only tolerate a negligible fraction ε of corruptions (corresponding to $\varepsilon = O(d^{-1/2})$).

A first idea to get around this obstacle is the following: We consider the standard (projected) gradient descent method used to find the minimum of \hat{f} . This algorithm would proceed by repeatedly computing the gradient of \hat{f} at appropriate points and using it to update the current location. The issue is that adversarial corruptions can completely compromise this algorithm’s behavior, since they can substantially change the gradient of \hat{f} at the chosen points. The key observation is that approximating the gradient of \bar{f} at a given point, given access to an ε -corrupted set of samples, can be viewed as a robust mean estimation problem. We can thus use the robust mean estimation algorithm of [DKK⁺17], which succeeds under fairly mild assumptions about the good samples. Assuming that the covariance matrix of $\nabla f(w)$, $f \sim p^*$, is bounded, we can thus “simulate” gradient descent and compute an approximate minimum for \bar{f} .

In summary, the first algorithmic idea is to use a robust mean estimation routine as a black-box in order to robustly estimate the gradient at *each* iteration of (projected) gradient descent. This yields a simple robust method for stochastic optimization with polynomial sample complexity and running time in a very general setting (See Appendix D for details.)

We are now ready to describe SEVER (Algorithm 1) and the main insight behind it. Roughly speaking, SEVER only calls our robust mean estimation routine (which is essentially the filtering method of [DKK⁺17] for outlier removal) each time the algorithm reaches an approximate critical point of \hat{f} . There are two main motivations for this approach: First, we empirically observed that if we iteratively filter samples, keeping the subset with the samples removed, then few iterations of the filter remove points. Second, an iteration of the filter subroutine (Algorithm 2) is more expensive than an iteration of gradient descent. Therefore, it is advantageous to run many steps of gradient descent on the current set of corrupted samples between consecutive filtering steps. This idea is further improved by using stochastic gradient descent, rather than computing the average at each step.

An important feature of our analysis is that SEVER does not use a robust mean estimation routine as a black box. In contrast, we take advantage of the performance guarantees of our filtering algorithm. The main idea for the analysis is as follows: Suppose that we have reached an approximate critical point w of \hat{f} and at this step we apply our filtering algorithm. By the performance guarantees of the latter algorithm we are in one of two cases: either the filtering algorithm removes a set of corrupted functions or it certifies that the gradient of \hat{f} is “close” to the gradient of \bar{f} at w . In the first case, we make progress as we produce a “cleaner” set of functions. In the second case, our certification implies that the point w is also an approximate critical point of \bar{f} and we are done.

3 Experiments

In this section we apply SEVER to regression and classification problems. As our base learners, we used ridge regression and an SVM, respectively. We implemented the latter as a quadratic program, using Gurobi [Gur16] as a backend solver and YALMIP [Löf04] as the modeling language.

In both cases, we ran the base learner and then extracted gradients for each data point at the learned parameters. We then centered the gradients and ran MATLAB’s `svds` method to compute the top singular vector v , and removed the top p fraction of points i with the largest *outlier score* τ_i , computed as the squared magnitude of the projection onto v (see Algorithm 1). We repeated this for r iterations in total. For classification, we centered the gradients separately (and removed points separately) for each class, which improved performance.

We compared our method to five baseline methods. These all have the same high-level form as SEVER (run the base learner then filter top p fraction of points with the largest score), but use a different definition of the score τ_i for deciding which points to filter:

- **noDefense**: no points are removed.
- **l2**: remove points where the covariate x has large ℓ_2 distance from the mean.
- **loss**: remove points with large loss (measured at the parameters output by the base learner).
- **gradient**: remove points with large gradient (in ℓ_2 -norm).
- **gradientCentered**: remove points whose gradients are far from the mean gradient in ℓ_2 -norm.

Note that **gradientCentered** is similar to our method, except that it removes large gradients in terms of ℓ_2 -norm, rather than in terms of projection onto the top singular vector. As before, for classification we compute these metrics separately for each class.

Both ridge regression and SVM have a single hyperparameter (the regularization coefficient). We optimized this based on the uncorrupted data and then kept it fixed throughout our experiments. In addition, since the data do not already have outliers, we added varying amounts of outliers (ranging from 0.5% to 10% of the clean data); this process is described in more detail below.

3.1 Ridge Regression

For ridge regression, we tested our method on a synthetic Gaussian dataset as well as a drug discovery dataset. The synthetic dataset consists of observations (x_i, y_i) where $x_i \in \mathbb{R}^{500}$ has independent standard Gaussian entries, and $y_i = \langle x_i, w^* \rangle + 0.1z_i$, where z_i is also Gaussian. We generated 5000 training points and 100 test points. The drug discovery dataset was obtained from the ChEMBL database and was originally curated by [OSB⁺18]; it consists of 4084 data points in 410 dimensions; we split this into a training set of 3084 points and a test set of 1000 points.

Centering We found that centering the data points decreased error noticeably on the drug discovery dataset, while scaling each coordinate to have variance 1 decreased error by a small amount on the synthetic data. To center in the presence of outliers, we used the robust mean estimation algorithm from [DKK⁺17].

Adding outliers. We devised a method of generating outliers that fools all of the baselines while still inducing high test error. At a high level, the outliers cause ridge regression to output $w = 0$ (so the model always predicts $y = 0$).

If (X, y) are the true data points and responses, this can be achieved by setting each outlier point $(X_{\text{bad}}, y_{\text{bad}})$ as

$$X_{\text{bad}} = \frac{1}{\alpha \cdot n_{\text{bad}}} y^\top X \quad \text{and} \quad y_{\text{bad}} = -\beta,$$

where n_{bad} is the number of outliers we add, and α and β are hyperparameters.

If $\alpha = \beta$, one can check that $w = 0$ is the unique minimizer for ridge regression on the perturbed dataset. By tuning α and β , we can then obtain attacks that fool all the baselines while damaging the model (we tune α and β separately to give an additional degree of freedom to the attack). To increase the error, we also found it useful to perturb each individual X_{bad} by a small amount of Gaussian noise.

In our experiments we found that this method generated successful attacks as long as the fraction of outliers was at least roughly 2% for synthetic data, and roughly 5% for the drug discovery data.

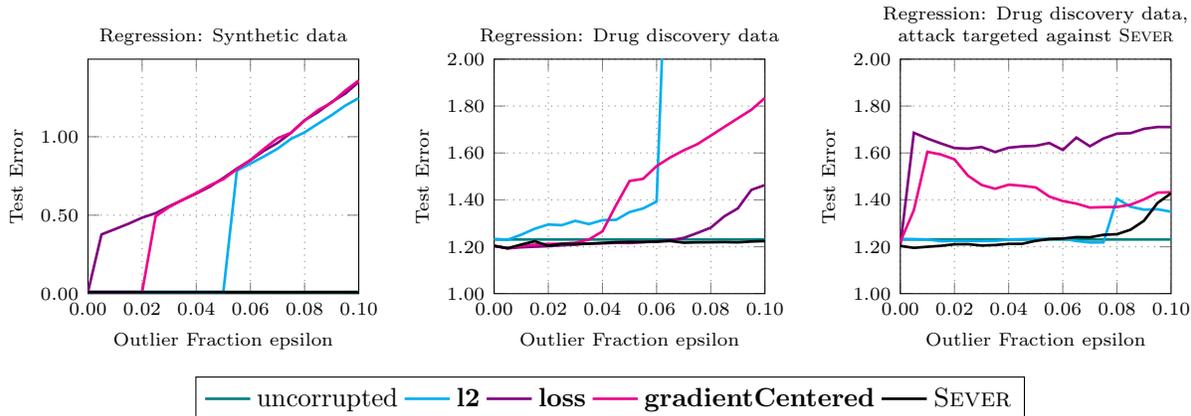


Figure 2: ϵ vs test error for baselines and SEVER on synthetic data and the drug discovery dataset. The left and middle figures show that SEVER continues to maintain statistical accuracy against our attacks which are able to defeat previous baselines. The right figure shows an attack with parameters chosen to increase the test error SEVER on the drug discovery dataset as much as possible. Despite this, SEVER still has relatively small test error.

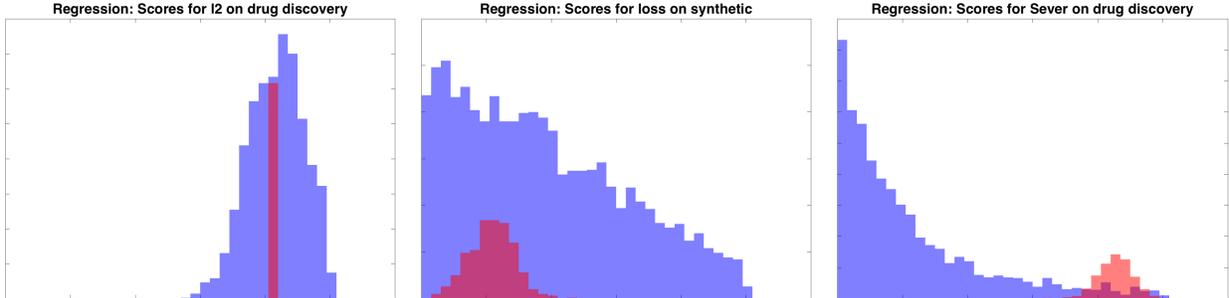


Figure 3: A representative set of histograms of scores for baselines and SEVER on synthetic data and a drug discovery dataset. From left to right: scores for the **l2** defense on the drug discovery dataset, scores for **loss** on synthetic data, and scores for SEVER on the drug discovery dataset, all with the addition of 10% outliers. The scores for the true dataset are in blue, and the scores for the outliers are in red. For the baselines, the scores for the outliers are inside the bulk of the distribution and thus hard to detect, whereas the scores for the outliers assigned by SEVER are clearly within the tail of the distribution and easily detectable.

Results. In Figure 2 we compare the test error of our defense against the baselines as we increase the fraction ε of added outliers. To avoid cluttering the figure, we only show the performance of **l2**, **loss**, **gradientCentered**, and SEVER; the performance of the remaining baselines is qualitatively similar to the baselines in Figure 2.

For both the baselines and our algorithms, we iterate the defense $r = 4$ times, each time removing the $p = \varepsilon/2$ fraction of points with largest score. For consistency of results, for each defense and each value of ε we ran the defense 3 times on fresh attack points and display the median of the 3 test errors.

When the attack parameters α and β are tuned to defeat the baselines (Figure 2 left and center), our defense substantially outperforms the baselines as soon as we cross $\varepsilon \approx 1.5\%$ for synthetic data, and $\varepsilon \approx 5.5\%$ for the drug discovery data. In fact, most of the baselines do worse than not removing any outliers at all (this is because they end up mostly removing good data points, which causes the outliers to have a larger effect). Even when α and β are instead tuned to defeat SEVER, its resulting error remains small (Figure 2 right).

To understand why the baselines fail to detect the outliers, in Figure 3 we show a representative sample of the histograms of scores of the uncorrupted points overlaid with the scores of the outliers, for both synthetic data and the drug discovery dataset with $\varepsilon = 0.1$, after one run of the base learner. The scores of the outliers lie well within the distribution of scores of the uncorrupted points. Thus, it would be impossible for the baselines to remove them without also removing a large fraction of uncorrupted points.

Interestingly, for small ε all of the methods improve upon the uncorrupted test error for the drug discovery data; this appears to be due to the presence of a small number of natural outliers in the data that all of the methods successfully remove.

3.2 Support Vector Machines

We next describe our experimental results for SVMs; we tested our method on a synthetic Gaussian dataset as well as a spam classification task. Similarly to before, the synthetic data consists of observations (x_i, y_i) , where $x_i \in \mathbb{R}^{500}$ has independent standard Gaussian entries, and $y_i = \text{sign}(\langle x_i, w^* \rangle + 0.1z_i)$, where z_i is also Gaussian and w^* is the true parameters (drawn at random from the unit sphere). The spam dataset comes from the Enron corpus [MAP06], and consists of 4137 training points and 1035 test points in 5116 dimensions.

To generate attacks, we used the data poisoning algorithm presented in [SKL17]; the authors provided us with an improved version of their algorithm that can circumvent the **l2** and **loss** baselines and partially circumvents the gradient baselines as well.

In contrast to ridge regression, we did not perform centering and rescaling for these datasets as it did not seem to have a large effect on results.

In all experiments for this section, each method removed the top $p = \frac{n_- + n_+}{\min\{n_+, n_-\}} \cdot \frac{\epsilon}{r}$ of highest-scoring points for each of $r = 2$ iterations, where n_+ and n_- are the number of positive and negative training points respectively. This expression for p is chosen in order to account for class imbalance, which is extreme in the case of the Enron dataset – if the attacker plants all the outliers in the smaller class, then a smaller value of p would remove too few points, even with a perfect detection method.

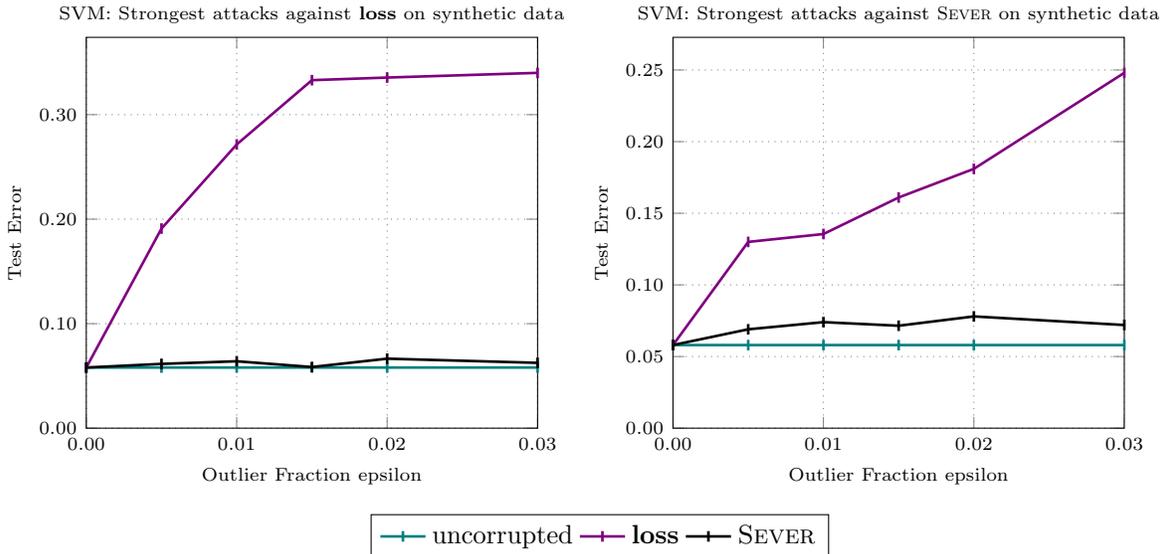


Figure 4: ϵ versus test error for **loss** baseline and SEVER on synthetic data. The left figure demonstrates that SEVER is accurate when outliers manage to defeat **loss**. The right figure shows the result of attacks which increased the test error the most against SEVER. Even in this case, SEVER performs much better than the baselines.

Synthetic results. We considered fractions of outliers ranging from $\epsilon = 0.005$ to $\epsilon = 0.03$. By performing a sweep across hyperparameters of the attack, we generated 56 distinct sets of attacks for each value of ϵ . In Figure 4, we show results for the attack where the **loss** baselines does the worst, as well as for the attack where our method does the worst. When attacks are most effective against **loss**, SEVER substantially outperforms it, nearly matching the test accuracy of 5.8% on the uncorrupted data, while **loss** performs worse than 30% error at just a 1.5% fraction of injected outliers. Even when attacks are most effective against SEVER, it still outperforms **loss**, achieving a test error of at most 9.05%. We note that other baselines behaved qualitatively similarly to **loss**, and the results are displayed in Section F.

Spam results. For results on Enron, we used the same values of ϵ , and considered 96 distinct hyperparameters for the attack. There was not a single attack that simultaneously defeated all of the baselines, so in Figure 5 we show two attacks that do well against different sets of baselines, as well as the attack that performs best against our method.

At $\epsilon = 0.01$, the worst performance of our method against all attacks was 7.34%, in contrast to 13.43% – 20.48% for the baselines (note that the accuracy is 3% in the absence of outliers). However, at $\epsilon = 0.03$, while we still outperform the baselines, our error is relatively large—13.53%.

To investigate this further, we looked at all 48 attacks and found that while on 42 out of 48 attacks our error never exceeded 7%, on 6 of the attacks (including the attack in Figure 5) the error was substantially higher. Figure 6 shows what is happening. The leftmost figure displays the scores assigned by SEVER

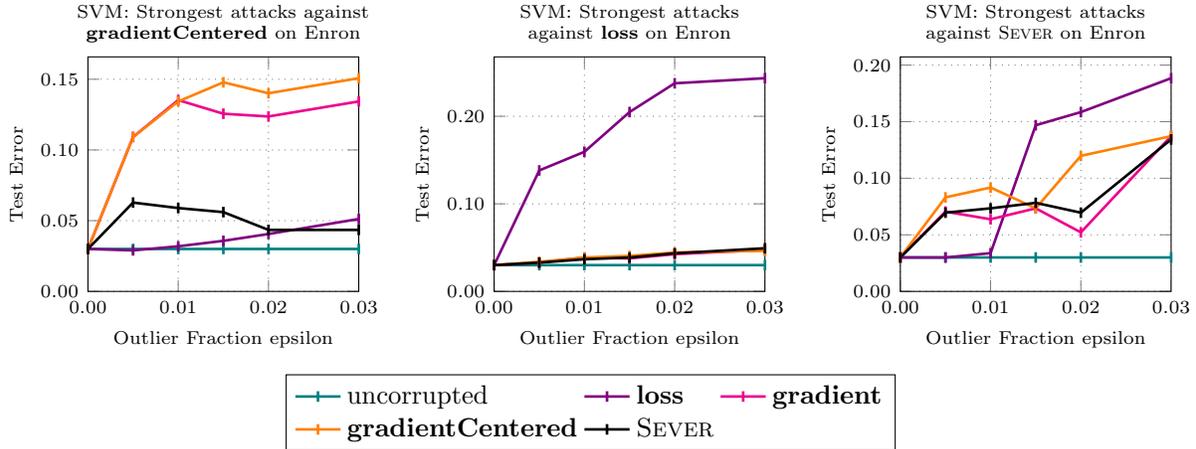


Figure 5: ϵ versus test error for baselines and SEVER on the Enron spam corpus. The left and middle figures are the attacks which perform best against two baselines, while the right figure performs best against SEVER. Though other baselines may perform well in certain cases, only SEVER is consistently accurate. The exception is for certain attacks at $\epsilon = 0.03$, which, as shown in Figure 6, require three rounds of outlier removal for any method to obtain reasonable test error – in these plots, our defenses perform only two rounds.

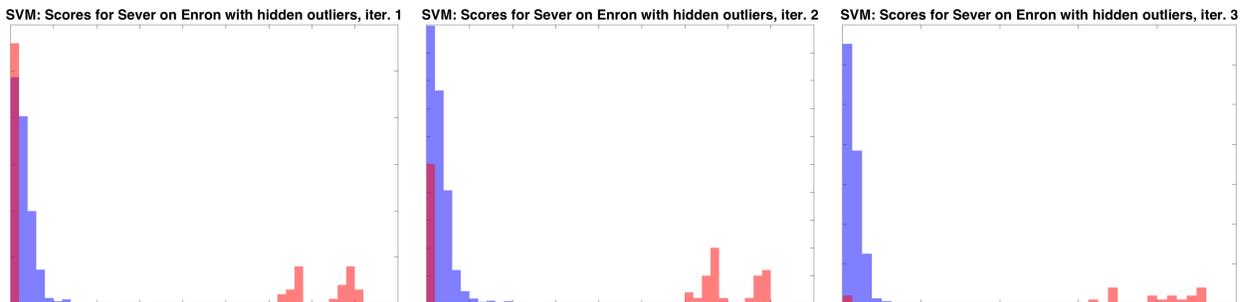


Figure 6: An illustration of why multiple rounds of filtering are necessary. Histograms of scores assigned by SEVER in three subsequent iterations of outlier removal. Inliers are blue, and outliers are red (scaled up by a factor of 10). In early iterations, a significant fraction of outliers may be “hidden” (i.e. have 0 loss) by being correctly classified in one iteration. However, once previous outliers are removed, these points may become incorrectly classified, thus significantly degrading the quality of our solution but simultaneously becoming evident to SEVER.

after the first iteration, where red bars indicate outliers. While some outliers are assigned extremely large scores and thus detected, several outliers are correctly classified and thus have 0 gradient. However, once we remove the first set of outliers, some outliers which were previously correctly classified now have large score, as displayed in the middle figure. Another iteration of this process produces the rightmost figure, where almost all the remaining outliers have large score and will thus be removed by SEVER. This demonstrates that some outliers may be hidden until other outliers are removed, necessitating multiple iterations.

Motivated by this, we re-ran our method against the 6 attacks using $r = 3$ iterations instead of 2 (and decreasing p as per the expression above). After this change, all 6 of the attacks had error at most 7.4%.

4 Discussion

In this paper we have presented an algorithm, SEVER, that has both strong theoretical robustness properties in the presence of outliers, and performs well on real datasets. SEVER is based on the idea that learning can often be cast as the problem of finding an approximate stationary point of the loss, which can in turn be cast as a robust mean estimation problem, allowing us to leverage existing techniques for efficient robust mean estimation.

There are a number of directions along which SEVER could be improved: first, it could be extended to handle more general assumptions on the data; second, it could be strengthened to achieve better error bounds in terms of the fraction of outliers; finally, one could imagine *automatically learning* a feature representation in which SEVER performs well. We discuss each of these ideas in detail below.

More general assumptions. The main underlying assumption on which SEVER rests is that the top singular value of the gradients of the data is small. While this appeared to hold true on the datasets we considered, a common occurrence in practice is for there to be *a few* large singular values, together with *many* small singular values. It would therefore be desirable to design a version of SEVER that can take advantage of such phenomena. In addition, it would be worthwhile to do a more detailed empirical analysis across a wide variety of datasets investigating properties that can enable robust estimation (the notion of *resilience* in [SCV18] could provide a template for finding such properties).

Stronger robustness to outliers. In theory, SEVER has a $O(\sqrt{\varepsilon})$ dependence in error on the fraction ε of outliers (see Theorem 2.1). While without stronger assumptions this is likely not possible to improve, in practice we would prefer to have a dependence closer to $O(\varepsilon)$. Therefore, it would also be useful to improve SEVER to have such an $O(\varepsilon)$ -dependence under stronger but realistic assumptions. Unfortunately, all existing algorithms for robust mean estimation that achieve error better than $O(\sqrt{\varepsilon})$ either rely on strong distributional assumptions such as Gaussianity [DKK⁺16, LRV16], or else require expensive computation involving e.g. sum-of-squares optimization [HL17, KS17a, KS17b]. Improving the robustness of SEVER thus requires improvements on the robust mean estimation algorithm that SEVER uses as a primitive.

Learning a favorable representation. Finally, we note that SEVER performs best when the features have small covariance and strong predictive power. One situation in particular where this holds is when there are many approximately independent features that are predictive of the true signal.

It would be interesting to try to learn a representation with such a property. This could be done, for instance, by training a neural network with some cost function that encourages independent features (some ideas along these general lines are discussed in [Ben17]). An issue is how to learn such a representation robustly; one idea is learn a representation on a dataset that is known to be free of outliers, and hope that the representation is useful on other datasets in the same application domain.

Beyond these specific questions, we view the general investigation of robust methods (both empirically and theoretically) as an important step as machine learning moves forwards. Indeed, as machine learning is applied in increasingly many situations and in increasingly automated ways, it is important to attend to robustness considerations so that machine learning systems behave reliably and avoid costly errors. While the bulk of recent work has highlighted the vulnerabilities of machine learning (e.g. [SZS⁺14, LWSV16, SKL17, EEF⁺17, CLL⁺17]), we are optimistic that practical algorithms backed by principled theory can finally patch these vulnerabilities and lead to truly reliable systems.

References

[ABL14] P. Awasthi, M. F. Balcan, and P. M. Long. The power of localization for efficiently learning linear separators with noise. In *Symposium on Theory of Computing (STOC)*, pages 449–458, 2014.

- [BDLS17] S. Balakrishnan, S. S. Du, J. Li, and A. Singh. Computationally efficient robust sparse estimation in high dimensions. In *Proceedings of the 30th Conference on Learning Theory, COLT 2017*, pages 169–212, 2017.
- [Ben17] Y. Bengio. The consciousness prior. *arXiv preprint arXiv:1709.08568*, 2017.
- [BJK15] K. Bhatia, P. Jain, and P. Kar. Robust regression via hard thresholding. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015*, pages 721–729, 2015.
- [BJKK17] K. Bhatia, P. Jain, P. Kamalaruban, and P. Kar. Consistent robust regression. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, pages 2107–2116, 2017.
- [BNJT10] M. Barreno, B. Nelson, A. D. Joseph, and J. D. Tygar. The security of machine learning. *Machine Learning*, 81(2):121–148, 2010.
- [BNL12] B. Biggio, B. Nelson, and P. Laskov. Poisoning attacks against support vector machines. In *International Conference on Machine Learning (ICML)*, pages 1467–1474, 2012.
- [CLL⁺17] X. Chen, C. Liu, B. Li, K. Lu, and D. Song. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*, 2017.
- [CSV17] M. Charikar, J. Steinhardt, and G. Valiant. Learning from untrusted data. In *Proceedings of STOC 2017*, pages 47–60, 2017.
- [DKK⁺16] I. Diakonikolas, G. Kamath, D. M. Kane, J. Li, A. Moitra, and A. Stewart. Robust estimators in high dimensions without the computational intractability. In *Proceedings of FOCS’16*, pages 655–664, 2016.
- [DKK⁺17] I. Diakonikolas, G. Kamath, D. M. Kane, J. Li, A. Moitra, and A. Stewart. Being robust (in high dimensions) can be practical. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017*, pages 999–1008, 2017. Full version available at <https://arxiv.org/abs/1703.00893>.
- [DKK⁺18] I. Diakonikolas, G. Kamath, D. M. Kane, J. Li, A. Moitra, and A. Stewart. Robustly learning a gaussian: Getting optimal error, efficiently. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*, pages 2683–2702, 2018. Full version available at <https://arxiv.org/abs/1704.03866>.
- [DKS16] I. Diakonikolas, D. M. Kane, and A. Stewart. Robust learning of fixed-structure bayesian networks. *CoRR*, abs/1606.07384, 2016.
- [DKS17a] I. Diakonikolas, D. M. Kane, and A. Stewart. Learning geometric concepts with nasty noise. *CoRR*, abs/1707.01242, 2017.
- [DKS17b] I. Diakonikolas, D. M. Kane, and A. Stewart. List-decodable robust mean estimation and learning mixtures of spherical gaussians. *CoRR*, abs/1711.07211, 2017.
- [DKS17c] I. Diakonikolas, D. M. Kane, and A. Stewart. Statistical query lower bounds for robust estimation of high-dimensional gaussians and gaussian mixtures. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*, pages 73–84, 2017. Full version available at <http://arxiv.org/abs/1611.03473>.
- [DKS18] I. Diakonikolas, W. Kong, and A. Stewart. Efficient algorithms and lower bounds for robust linear regression. Manuscript, 2018.
- [EEF⁺17] I. Evtimov, K. Eykholt, E. Fernandes, T. Kohno, B. Li, A. Prakash, A. Rahmati, and D. Song. Robust physical-world attacks on machine learning models. *arXiv*, 2017.

- [Gur16] Gurobi Optimization, Inc. Gurobi optimizer reference manual, 2016.
- [HL17] S. B. Hopkins and J. Li. Mixture models, robustness, and sum of squares proofs. *CoRR*, abs/1711.07454, 2017.
- [Hub64] P. J. Huber. Robust estimation of a location parameter. *Ann. Math. Statist.*, 35(1):73–101, 03 1964.
- [KL17] P. W. Koh and P. Liang. Understanding black-box predictions via influence functions. In *International Conference on Machine Learning (ICML)*, 2017.
- [KLS09] A. R. Klivans, P. M. Long, and R. A. Servedio. Learning halfspaces with malicious noise. *Journal of Machine Learning Research (JMLR)*, 10:2715–2740, 2009.
- [KS17a] P. K. Kothari and J. Steinhardt. Better agnostic clustering via relaxed tensor norms. *CoRR*, abs/1711.07465, 2017.
- [KS17b] P. K. Kothari and D. Steurer. Outlier-robust moment-estimation via sum-of-squares. *CoRR*, abs/1711.11581, 2017.
- [LAT⁺08] J.Z. Li, D.M. Absher, H. Tang, A.M. Southwick, A.M. Casto, S. Ramachandran, H.M. Cann, G.S. Barsh, M. Feldman, L.L. Cavalli-Sforza, and R.M. Myers. Worldwide human relationships inferred from genome-wide patterns of variation. *Science*, 319:1100–1104, 2008.
- [Löf04] J. Löfberg. YALMIP: A toolbox for modeling and optimization in MATLAB. In *CACSD*, 2004.
- [LRV16] K. A. Lai, A. B. Rao, and S. Vempala. Agnostic estimation of mean and covariance. In *Proceedings of FOCS’16*, 2016.
- [LWSV16] B. Li, Y. Wang, A. Singh, and Y. Vorobeychik. Data poisoning attacks on factorization-based collaborative filtering. In *Advances in Neural Information Processing Systems (NIPS)*, 2016.
- [MAP06] V. Metsis, I. Androutsopoulos, and G. Paliouras. Spam filtering with naive Bayes – which naive Bayes? In *CEAS*, volume 17, pages 28–69, 2006.
- [MV17] M. Meister and G. Valiant. A data prism: Semi-verified learning in the small-alpha regime. *arXiv*, 2017.
- [NPXNR14] A. Newell, R. Potharaju, L. Xiang, and C. Nita-Rotaru. On the practicality of integrity attacks on document-level sentiment analysis. In *Workshop on Artificial Intelligence and Security (AISec)*, pages 83–93, 2014.
- [NT13] N. H. Nguyen and T. D. Tran. Exact recoverability from dense corrupted observations via ℓ_1 -minimization. *IEEE Transactions on Information Theory*, 59(4):2017–2035, 2013.
- [NTN11] N. M. Nasrabadi, T. D. Tran, and N. Nguyen. Robust lasso with missing and grossly corrupted observations. In *Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [OSB⁺18] I. Olier, N. Sadawi, G. R. Bickerton, J. Vanschoren, C. Grosan, L. Soldatova, and Ross D. King. Meta-qsar: a large-scale application of meta-learning to drug design and discovery. *Machine Learning*, 107(1):285–311, Jan 2018.
- [PLJD10] P. Paschou, J. Lewis, A. Javed, and P. Drineas. Ancestry informative markers for fine-scale individual assignment to worldwide populations. *Journal of Medical Genetics*, 47:835–847, 2010.
- [PSBR18] A. Prasad, A. S. Suggala, S. Balakrishnan, and P. Ravikumar. Robust estimation via robust gradient estimation. *CoRR*, abs/1802.06485, 2018.
- [QV17] Mingda Qiao and Gregory Valiant. Learning discrete distributions from untrusted batches. *arXiv preprint arXiv:1711.08113*, 2017.

- [RPW⁺02] N. Rosenberg, J. Pritchard, J. Weber, H. Cann, K. Kidd, L.A. Zhivotovsky, and M.W. Feldman. Genetic structure of human populations. *Science*, 298:2381–2385, 2002.
- [SCV18] J. Steinhardt, M. Charikar, and G. Valiant. Resilience: A criterion for learning in the presence of arbitrary outliers. In *Innovations in Theoretical Computer Science (ITCS)*, 2018.
- [SKL17] J. Steinhardt, P. W. Koh, and P. Liang. Certified defenses for data poisoning attacks. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [Ste17] J. Steinhardt. Does robustness imply tractability? A lower bound for planted clique in the semi-random model. *arXiv*, 2017.
- [SZS⁺14] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations (ICLR)*, 2014.
- [Tuk60] J. W. Tukey. A survey of sampling from contaminated distributions. *Contributions to probability and statistics*, 2:448–485, 1960.

Appendix

The appendix is organized as follows. In Section A, we describe additional preliminaries required for our technical arguments. In Section B, we analyze our main algorithm, SEVER. In Section C, we specialize our analysis to the important case of Generalized Linear Models (GLMs). In Section D, we describe a variant of our algorithm which performs robust filtering on each iteration of projected gradient descent, and works under more general assumptions. In Section E, we describe concrete applications of SEVER – in particular, how it can be used to robustly optimize in the settings of linear regression, logistic regression, and support vector machines. Finally, in Section F, we provide additional plots from our experimental evaluations.

A Preliminaries

In this section, we formally introduce our setting for robust stochastic optimization.

Notation. For $n \in \mathbb{Z}_+$, we will denote $[n] \stackrel{\text{def}}{=} \{1, \dots, n\}$. For a vector v , we will let $\|v\|_2$ denote its Euclidean norm. For any $r \geq 0$ and any $x \in \mathbb{R}^d$, let $B(x, r)$ be the ℓ_2 ball of radius r around x . If M is a matrix, we will let $\|M\|_2$ denote its spectral norm and $\|M\|_F$ denote its Frobenius norm. We will write $X \sim_u S$ to denote that X is drawn from the empirical distribution defined by S . We will sometimes use the notation \mathbb{E}_S , instead of $\mathbb{E}_{X \sim S}$, for the corresponding expectation. We will also use the same convention for the covariance, i.e. we let Cov_S denote the covariance over the empirical distribution. [Anything else?]

Setting. We consider a stochastic optimization setting with outliers. Let $\mathcal{H} \subseteq \mathbb{R}^d$ be a space of parameters. We observe n functions $f_1, \dots, f_n : \mathcal{H} \rightarrow \mathbb{R}$ and we are interested in (approximately) minimizing some target function $\bar{f} : \mathcal{H} \rightarrow \mathbb{R}$, related to the f_i 's. We will assume for simplicity that the f_i 's are differentiable with gradient ∇f_i . (Our results can be easily extended for the case that only a sub-gradient is available.)

In most concrete applications we will consider, there is some true underlying distribution p^* over functions $f : \mathcal{H} \rightarrow \mathbb{R}$, and our goal is to find a parameter vector $w^* \in \mathcal{H}$ minimizing $\bar{f}(w) \stackrel{\text{def}}{=} \mathbb{E}_{f \sim p^*}[f(w)]$. Unlike the classical realizable setting, where we assume that $f_1, \dots, f_n \sim p^*$, we allow for an ε -fraction of the points to be arbitrary outliers. This is captured in the following definition (Definition 2.1) that we restate for convenience:

Definition A.1 (ε -corruption model). Given $\varepsilon > 0$ and a distribution p^* over functions $f : \mathcal{H} \rightarrow \mathbb{R}$, data is generated as follows: first, n clean samples f_1, \dots, f_n are drawn from p^* . Then, an *adversary* is allowed to inspect the samples and replace any εn of them with arbitrary samples. The resulting set of points is then given to the algorithm.

In addition, some of our bounds will make use of the following quantities:

- The ℓ_2 -radius r of the domain \mathcal{H} : $r = \max_{w \in \mathcal{H}} \|w\|_2$.
- The strong convexity parameter ξ of \bar{f} , if it exists. This is the maximal ξ such that $\bar{f}(w) \geq \bar{f}(w_0) + \langle w - w_0, \nabla \bar{f}(w_0) \rangle + \frac{\xi}{2} \|w - w_0\|_2^2$ for all $w, w_0 \in \mathcal{H}$.
- The strong smoothness parameter β of \bar{f} , if it exists. This is the minimal β such that $\bar{f}(w) \leq \bar{f}(w_0) + \langle w - w_0, \nabla \bar{f}(w_0) \rangle + \frac{\beta}{2} \|w - w_0\|_2^2$ for all $w, w_0 \in \mathcal{H}$.
- The Lipschitz constant L of \bar{f} , if it exists. This is the minimal L such that $\bar{f}(w) - \bar{f}(w_0) \leq L \|w - w_0\|_2$ for all $w, w_0 \in \mathcal{H}$.

B General Analysis of Sever

This section is dedicated to the analysis of Algorithm 1, where we do not make convexity assumptions about the underlying functions f_1, \dots, f_n . In this case, we can show that our algorithm finds an approximate critical point of \bar{f} . When we specialize to convex functions, this immediately implies that we find an approximate minimal point of \bar{f} .

Our proof proceeds in two parts. First, we define a set of deterministic conditions under which our algorithm finds an approximate minimal point of \bar{f} . We then show that, under mild assumptions on our functions, this set of deterministic conditions holds with high probability after polynomially many samples.

For completeness, we recall the definitions of a γ -approximate critical point and a γ -approximate learner:

Definition 2.2 (γ -approximate critical point). Given a function $f : \mathcal{H} \rightarrow \mathbb{R}$, a γ -approximate critical point of f , is a point $w \in \mathcal{H}$ so that for all unit vectors v where $w + \delta v \in \mathcal{H}$ for arbitrarily small positive δ , we have that $v \cdot \nabla f(w) \geq -\gamma$.

Definition 2.3 (γ -approximate learner). A learning algorithm \mathcal{L} is called γ -approximate if, for any functions $f_1, \dots, f_n : \mathcal{H} \rightarrow \mathbb{R}$ each bounded below on a closed domain \mathcal{H} , the output $w = \mathcal{L}(f_{1:n})$ of \mathcal{L} is a γ -approximate critical point of $f(x) := \frac{1}{n} \sum_{i=1}^n f_i(x)$.

Deterministic Regularity Conditions We first explicitly demonstrate a set of deterministic conditions on the (uncorrupted) data points. Our deterministic regularity conditions are as follows:

Assumption B.1. Fix $0 < \varepsilon < 1/2$. There exists an unknown set $I_{\text{good}} \subseteq [n]$ with $|I_{\text{good}}| \geq (1 - \varepsilon)n$ of “good” functions $\{f_i\}_{i \in I_{\text{good}}}$ and parameters $\sigma_0, \sigma_1 \in \mathbb{R}_+$ such that:

$$\left\| \mathbb{E}_{I_{\text{good}}} [(\nabla f_i(w) - \nabla \bar{f}(w))(\nabla f_i(w) - \nabla \bar{f}(w))^T] \right\|_2 \leq (\sigma_0 + \sigma_1 \|w^* - w\|_2)^2, \text{ for all } w \in \mathcal{H}, \quad (1)$$

and

$$\|\nabla \hat{f}(w) - \nabla \bar{f}(w)\|_2 \leq (\sigma_0 + \sigma_1 \|w^* - w\|_2) \sqrt{\varepsilon}, \text{ for all } w \in \mathcal{H}, \text{ where } \hat{f} \stackrel{\text{def}}{=} \frac{1}{|I_{\text{good}}|} \sum_{i \in I_{\text{good}}} f_i. \quad (2)$$

In Section B.1, we prove the following theorem, which shows that under Assumption B.1 our algorithm succeeds:

Theorem B.2. Suppose that the functions $f_1, \dots, f_n, \bar{f} : \mathcal{H} \rightarrow \mathbb{R}$ are bounded below, and that Assumption B.1 is satisfied, where $\sigma \stackrel{\text{def}}{=} \sigma_0 + \sigma_1 \|w^* - w\|_2$. Then SEVER applied to f_1, \dots, f_n, σ returns a point $w \in \mathcal{H}$ that, with probability at least $9/10$, is a $(\gamma + O(\sigma\sqrt{\varepsilon}))$ -approximate critical point of \bar{f} .

Observe that the above theorem holds quite generally; in particular, it holds for non-convex functions. As a corollary of this theorem, in Section B.2 we show that this immediately implies that SEVER robustly minimizes convex functions, if Assumption B.1 holds:

Corollary B.3. *For functions $f_1, \dots, f_n : \mathcal{H} \rightarrow \mathbb{R}$, suppose that Assumption B.1 holds and that \mathcal{H} is convex. Then, with probability at least $9/10$, for some universal constant ε_0 , if $\varepsilon < \varepsilon_0$, the output of SEVER satisfies the following:*

- (i) *If \bar{f} is convex, the algorithm finds a $w \in \mathcal{H}$ such that $\bar{f}(w) - \bar{f}(w^*) = O((\sigma_0 r + \sigma_1 r^2)\sqrt{\varepsilon} + \gamma r)$.*
- (ii) *If \bar{f} is ξ -strongly convex, the algorithm finds a $w \in \mathcal{H}$ such that*

$$\bar{f}(w) - \bar{f}(w^*) = O\left(\frac{\varepsilon}{\xi}(\sigma_0 + \sigma_1 r)^2 + \frac{\gamma^2}{\xi}\right).$$

In the strongly convex case and when $\sigma_1 > 0$, we can remove the dependence on σ_1 and r in the above by repeatedly applying SEVER with decreasing r :

Corollary B.4. *For functions $f_1, \dots, f_n : \mathcal{H} \rightarrow \mathbb{R}$, suppose that Assumption B.1 holds, that \mathcal{H} is convex and that \bar{f} is ξ -strongly convex for $\xi \geq C\sigma_1\sqrt{\varepsilon}$ for some absolute constant C . Then, with probability at least $9/10$, for some universal constant ε_0 , if $\varepsilon < \varepsilon_0$, we can find a \hat{w} with*

$$\bar{f}(\hat{w}) - \bar{f}(w^*) = O\left(\frac{\varepsilon\sigma_0^2 + \gamma^2}{\xi}\right).$$

and

$$\|\hat{w} - w^*\|_2 = O\left(\frac{\sqrt{\varepsilon}\sigma_0 + \gamma}{\xi}\right)$$

using at most $O(\log(r\xi/(\gamma + \sigma_0\sqrt{\varepsilon})))$ calls to SEVER.

To concretely use Theorem B.2, Corollary B.3, and Corollary B.4, in Section B.4 we show that the Assumption B.1 is satisfied with high probability under mild conditions on the distribution over the functions, after drawing polynomially many samples:

Proposition B.5. *Let $\mathcal{H} \subset \mathbb{R}^d$ be a closed bounded set with diameter at most r . Let p^* be a distribution over functions $f : \mathcal{H} \rightarrow \mathbb{R}$ with $\bar{f} = \mathbb{E}_{f \sim p^*}[f]$ so that $f - \bar{f}$ is L -Lipschitz and β -smooth almost surely. Assume furthermore that for each $w \in \mathcal{H}$ and unit vector v that $\mathbb{E}_{f \sim p^*}[(v \cdot (\nabla f(w) - \bar{f}(w)))^2] \leq \sigma^2/2$. Then for*

$$n = \Omega\left(\frac{d^2 L^2 \log(r\beta L/\sigma^2\varepsilon)}{\sigma^2\varepsilon}\right),$$

an ε -corrupted set of points f_1, \dots, f_n with high probability satisfy Assumption B.1.

The remaining subsections are dedicated to the proofs of Theorem B.2, Corollary B.3, Corollary B.4, and Proposition B.5.

B.1 Proof of Theorem B.2

Throughout this proof we let I_{good} be as in Assumption B.1. We require the following two lemmata. Roughly speaking, the first states that on average, we remove more corrupted points than uncorrupted points, and the second states that at termination, and if we have not removed too many points, then we have reached a point at which the empirical gradient is close to the true gradient. Formally:

Lemma B.6. *If the samples satisfy (1) of Assumption B.1, and if $|S| \geq 2n/3$ then if S' is the output of $\text{FILTER}(S, \tau, \sigma)$, we have that*

$$\mathbb{E}[|I_{\text{good}} \cap (S \setminus S')|] \leq \mathbb{E}[|([n] \setminus I_{\text{good}}) \cap (S \setminus S')|].$$

Lemma B.7. *If the samples satisfy Assumption B.1, $\text{FILTER}(S, \tau, \sigma) = S$, and $n - |S| \leq 11\epsilon n$, then*

$$\left\| \nabla \bar{f}(w) - \frac{1}{|I_{\text{good}}|} \sum_{i \in S} \nabla f_i(w) \right\|_2 \leq O(\sigma \sqrt{\epsilon})$$

Before we prove these lemmata, we show how together they imply Theorem B.2.

Proof of Theorem B.2 assuming Lemma B.6 and Lemma B.7. First, we note that the algorithm must terminate in at most n iterations. This is easy to see as each iteration of the main loop except for the last must decrease the size of S by at least 1.

It thus suffices to prove correctness. Note that Lemma B.6 says that each iteration will on average throw out as many elements not in I_{good} from S as elements in I_{good} . In particular, this means that $|([n] \setminus I_{\text{good}}) \cap S| + |I_{\text{good}} \setminus S|$ is a supermartingale. Since its initial size is at most ϵn , with probability at least $9/10$, it never exceeds $10\epsilon n$, and therefore at the end of the algorithm, we must have that $n - |S| \leq \epsilon n + |I_{\text{good}} \setminus S| \leq 11\epsilon n$. This will allow us to apply Lemma B.7 to complete the proof, using the fact that w is a γ -approximate critical point of $\frac{1}{|I_{\text{good}}|} \sum_{i \in S} \nabla f_i(w)$. \square

Thus it suffices to prove these two lemmata. We first prove Lemma B.6:

Proof of Lemma B.6. Let $S_{\text{good}} = S \cap I_{\text{good}}$ and $S_{\text{bad}} = S \setminus I_{\text{good}}$. We wish to show that the expected number of elements thrown out of S_{bad} is at least the expected number thrown out of S_{good} . We note that our result holds trivially if $\text{FILTER}(S, \tau, \sigma) = S$. Thus, we can assume that $\mathbb{E}_{i \in S}[\tau_i] \geq 12\sigma$.

It is easy to see that the expected number of elements thrown out of S_{bad} is proportional to $\sum_{i \in S_{\text{bad}}} \tau_i$, while the number removed from S_{good} is proportional to $\sum_{i \in S_{\text{good}}} \tau_i$ (with the same proportionality). Hence, it suffices to show that $\sum_{i \in S_{\text{bad}}} \tau_i \geq \sum_{i \in S_{\text{good}}} \tau_i$.

We first note that since $\text{Cov}_{i \in I_{\text{good}}}[\nabla f_i(w)] \preceq \sigma^2 I$, we have that

$$\begin{aligned} \text{Cov}_{i \in S_{\text{good}}}[v \cdot \nabla f_i(w)] &\stackrel{(a)}{\leq} \frac{3}{2} \text{Cov}_{i \in I_{\text{good}}}[v \cdot \nabla f_i(w)] \\ &= \frac{3}{2} \cdot v^\top \text{Cov}_{i \in I_{\text{good}}}[\nabla f_i(w)] v \leq 2\sigma^2, \end{aligned}$$

where (a) follows since $|S_{\text{good}}| \geq \frac{3}{2} |I_{\text{good}}|$.

Let $\mu_{\text{good}} = \mathbb{E}_{i \in S_{\text{good}}}[v \cdot \nabla f_i(w)]$ and $\mu = \mathbb{E}_{i \in S}[v \cdot \nabla f_i(w)]$. Note that

$$\mathbb{E}_{i \in S_{\text{good}}}[\tau_i] = \text{Cov}_{i \in S_{\text{good}}}[v \cdot \nabla f_i(w)] + (\mu - \mu_{\text{good}})^2 \leq 2\sigma + (\mu - \mu_{\text{good}})^2.$$

We now split into two cases.

Firstly, if $(\mu - \mu_{\text{good}})^2 \geq 4\sigma^2$, we let $\mu_{\text{bad}} = \mathbb{E}_{i \in S_{\text{bad}}}[v \cdot \nabla f_i(w)]$, and note that $|\mu - \mu_{\text{bad}}| |S_{\text{bad}}| = |\mu - \mu_{\text{good}}| |S_{\text{good}}|$. We then have that

$$\begin{aligned} \mathbb{E}_{i \in S_{\text{bad}}}[\tau_i] &\geq (\mu - \mu_{\text{bad}})^2 \\ &\geq (\mu - \mu_{\text{good}})^2 \left(\frac{|S_{\text{good}}|}{|S_{\text{bad}}|} \right)^2 \\ &\geq 2 \left(\frac{|S_{\text{good}}|}{|S_{\text{bad}}|} \right) (\mu - \mu_{\text{good}})^2 \\ &\geq \left(\frac{|S_{\text{good}}|}{|S_{\text{bad}}|} \right) \mathbb{E}_{i \in S_{\text{good}}}[\tau_i]. \end{aligned}$$

Hence, $\sum_{i \in S_{\text{bad}}} \tau_i \geq \sum_{i \in S_{\text{good}}} \tau_i$.

On the other hand, if $(\mu - \mu_{\text{good}})^2 \leq 4\sigma^2$, then $\mathbb{E}_{i \in S_{\text{good}}}[\tau_i] \leq 6\sigma^2 \leq \mathbb{E}_{i \in S}[\tau_i]/2$. Therefore $\sum_{i \in S_{\text{bad}}} \tau_i \geq \sum_{i \in S_{\text{good}}} \tau_i$ once again. This completes our proof. \square

We now prove Lemma B.7.

Proof of Lemma B.7. We need to show that

$$\delta := \left\| \sum_{i \in S} (\nabla f_i(w) - \nabla \bar{f}(w)) \right\|_2 = O(n\sigma\sqrt{\varepsilon}).$$

We note that

$$\begin{aligned} & \left\| \sum_{i \in S} (\nabla f_i(w) - \nabla \bar{f}(w)) \right\|_2 \\ & \leq \left\| \sum_{i \in I_{\text{good}}} (\nabla f_i(w) - \nabla \bar{f}(w)) \right\|_2 + \left\| \sum_{i \in (I_{\text{good}} \setminus S)} (\nabla f_i(w) - \nabla \bar{f}(w)) \right\|_2 + \left\| \sum_{i \in (S \setminus I_{\text{good}})} (\nabla f_i(w) - \nabla \bar{f}(w)) \right\|_2 \\ & = \left\| \sum_{i \in (I_{\text{good}} \setminus S)} (\nabla f_i(w) - \nabla \bar{f}(w)) \right\|_2 + \left\| \sum_{i \in (S \setminus I_{\text{good}})} (\nabla f_i(w) - \nabla \bar{f}(w)) \right\|_2 + O(n\sqrt{\sigma^2\varepsilon}). \end{aligned}$$

First we analyze

$$\left\| \sum_{i \in (I_{\text{good}} \setminus S)} (\nabla f_i(w) - \nabla \bar{f}(w)) \right\|_2.$$

This is the supremum over unit vectors v of

$$\sum_{i \in (I_{\text{good}} \setminus S)} v \cdot (\nabla f_i(w) - \nabla \bar{f}(w)).$$

However, we note that

$$\sum_{i \in I_{\text{good}}} (v \cdot (\nabla f_i(w) - \nabla \bar{f}(w)))^2 = O(n\sigma^2).$$

Since $|I_{\text{good}} \setminus S| = O(n\varepsilon)$, we have by Cauchy-Schwarz that

$$\sum_{i \in (I_{\text{good}} \setminus S)} v \cdot (\nabla f_i(w) - \nabla \bar{f}(w)) = O(\sqrt{(n\sigma^2)(n\varepsilon)}) = O(n\sqrt{\sigma^2\varepsilon}),$$

as desired.

We note that since for any such v that

$$\sum_{i \in S} (v \cdot (\nabla f_i(w) - \nabla \bar{f}(w)))^2 = \sum_{i \in S} (v \cdot (\nabla f_i(w) - \nabla \hat{f}(w)))^2 + \delta^2 = O(n\sigma^2) + \delta^2$$

(or otherwise our filter would have removed elements) and since $|S \setminus I_{\text{good}}| = O(n\varepsilon)$, and so we have similarly that

$$\left\| \sum_{i \in (S \setminus I_{\text{good}})} \nabla f_i(w) - \nabla \bar{f}(w) \right\|_2 = O(n\sigma\sqrt{\varepsilon} + \delta\sqrt{n\varepsilon}).$$

Combining with the above we have that

$$\delta = O(\sigma\sqrt{\varepsilon} + \delta\sqrt{\varepsilon/n}),$$

and therefore, $\delta = O(\sigma\sqrt{\varepsilon})$ as desired. \square

B.2 Proof of Corollary B.3

In this section, we show that the SEVER algorithm finds an approximate global optimum for convex optimization in various settings, under Assumption B.1. We do so by simply applying the guarantees of Theorem B.2 in a fairly black box manner.

Before we proceed with the proof of Corollary B.3, we record a simple lemma that allows us to translate an approximate critical point guarantee to an approximate global optimum guarantee:

Lemma B.8. *Let $f : \mathcal{H} \rightarrow \mathbb{R}$ be a convex function and let $x \neq y \in \mathcal{H}$. Let $v = (y - x) / \|y - x\|_2$ be the unit vector in the direction of $y - x$. Suppose that for some δ that $v \cdot (\nabla f(x)) \geq -\delta$ and $-v \cdot (\nabla f(y)) \geq -\delta$. Then we have that:*

1. $|f(x) - f(y)| \leq \|x - y\|_2 \delta$.
2. If f is ξ -strongly convex, then $|f(x) - f(y)| \leq 2\delta^2/\xi$ and $\|x - y\|_2 \leq 2\delta/\xi$.

Proof. Let $r = \|x - y\|_2 > 0$ and $g(t) = f(x + tv)$. We have that $g(0) = f(x)$, $g(r) = f(y)$ and that g is convex (or ξ -strongly convex) with $g'(0) \geq -\delta$ and $g'(r) \leq \delta$. By convexity, the derivative of g is increasing on $[0, r]$ and therefore $|g'(t)| \leq \delta$ for all $t \in [0, r]$. This implies that

$$|f(x) - f(y)| = |g(r) - g(0)| = \left| \int_0^r g'(t) dt \right| \leq r\delta.$$

To show the second part of the lemma, we note that if g is ξ -strongly convex that $g''(t) \geq \xi$ for all t . This implies that $g'(r) > g'(0) + \xi r$. Since $g'(r) - g'(0) \leq 2\delta$, we obtain that $r \leq 2\delta/\xi$, from which the second statement follows. \square

Proof of Corollary B.3. By applying the algorithm of Theorem B.2, we can find a point w that is a $\gamma' \stackrel{\text{def}}{=} (\gamma + O(\sigma\sqrt{\varepsilon}))$ -approximate critical point of \bar{f} , where $\sigma \stackrel{\text{def}}{=} \sigma_0 + \sigma_1 \|w^* - w\|_2$. That is, for any unit vector v pointing towards the interior of \mathcal{H} , we have that $v \cdot \nabla \bar{f}(w) \geq -\gamma'$.

To prove (i), we apply Lemma B.8 to \bar{f} at w which gives that

$$|\bar{f}(w) - \bar{f}(w^*)| \leq r \cdot \gamma'.$$

To prove (ii), we apply Lemma B.8 to \bar{f} at w which gives that

$$|\bar{f}(w) - \bar{f}(w^*)| \leq 2\gamma'^2/\xi.$$

Plugging in parameters appropriately then immediately gives the desired bound. \square

B.3 Proof of Corollary B.4

We apply SEVER iteratively starting with a domain $\mathcal{H}_1 = \mathcal{H}$ and radius $r_1 = r$. After each iteration, we know the resulting point is close to w^* will be able to reduce the search radius.

At step i , we have a domain of radius r_i . As in the proof of Corollary B.3 above, we apply algorithm of Theorem B.2, we can find a point w_i that is a $\gamma'_i \stackrel{\text{def}}{=} (\gamma + O(\sigma'_i\sqrt{\varepsilon}))$ -approximate critical point of \bar{f} , where $\sigma'_i \stackrel{\text{def}}{=} \sigma_0 + \sigma_1 r_i$. Then using Lemma B.8, we obtain that $\|w_i - w^*\|_2 \leq 2\gamma'_i/\xi$.

Now we can define \mathcal{H}_{i+1} as the intersection of \mathcal{H} and the ball of radius $r_{i+1} = 2\gamma'_i/\xi$ around w_i and repeat using this domain. We have that $r_{i+1} = 2\gamma'_i/\xi = 2\gamma/\xi + O(\sigma_0\sqrt{\varepsilon}/\xi + \sigma_1\sqrt{\varepsilon}r_i/\xi)$. Now if we choose the constant C such that the constant in this $O()$ is $C/4$, then using our assumption that $\xi \geq 2\sigma_1\sqrt{\varepsilon}$, we obtain that

$$r_{i+1} \leq 2\gamma/\xi + C\sigma_0\sqrt{\varepsilon}/4\xi + C\sigma_1\sqrt{\varepsilon}r_i/4\xi \leq 2\gamma/\xi + C\sigma_0\sqrt{\varepsilon}/4 + r_i/4$$

Now if $r_i \geq 8\gamma/\xi + 2C\sigma_0\sqrt{\varepsilon}/\xi$, then we have $r_{i+1} \leq r_i/2$ and if $r_i \leq 8\gamma/\xi + 2C\sigma_0\sqrt{\varepsilon}/\xi$ then we also have $r_{i+1} \leq 8\gamma/\xi + 2C\sigma_0\sqrt{\varepsilon}/\xi$. When r_i is smaller than this we stop and output w_i . Thus we stop in at most

$O(\log(r) - \log(8\gamma/\xi + 2C\sigma_0\sqrt{\varepsilon}/\xi)) = O(\log(r\xi/(\gamma + \sigma_0\sqrt{\varepsilon})))$ iterations and have $r_i = O(\gamma/\xi + C\sigma_0\sqrt{\varepsilon})$. But then $\gamma'_i = \gamma + O(\sigma'_i\sqrt{\varepsilon}) \leq \gamma + C(\sigma_0 + \sigma_1 r'_i)\sqrt{\varepsilon}/8 = O(\gamma + \sigma_0\sqrt{\varepsilon})$. Using Lemma B.8 we obtain that

$$|\bar{f}(w_i) - \bar{f}(w^*)| \leq 2\gamma'_i{}^2/\xi = O(\gamma^2/\xi + \sigma_0^2\varepsilon/\xi).$$

as required. The bound on $\|\hat{w} - w^*\|_2$ follows similarly.

Remark B.1. While we don't give explicit bounds on the number of calls to the approximate learner needed by SEVER, such bounds can be straightforwardly obtained under appropriate assumptions on the f_i (see, e.g., the following subsection). Two remarks are in order. First, in this case we cannot take advantage of assumptions that only hold at \bar{f} but might not on the corrupted average f . Second, our algorithm can take advantage of a closed form for the minimum. For example, for the case of linear regression considered in Section E, f_i is not Lipschitz with a small constant if x_i is far from the mean, but there is a simple closed form for the minimum of the least squares loss.

B.4 Proof of Proposition B.5

We let I_{good} be the set of uncorrupted functions f_i . It is then the case that $|I_{\text{good}}| \geq (1 - \varepsilon)n$. We need to show that for each $w \in \mathcal{H}$ that

$$\text{Cov}_{i \in I_{\text{good}}}[\nabla f_i(w)] \leq 3\sigma^2 I/4 \quad (3)$$

and

$$\left\| \nabla \bar{f}(w) - \frac{1}{|I_{\text{good}}|} \sum_{i \in I_{\text{good}}} \nabla f_i(w) \right\|_2 \leq O(\sigma^2 \sqrt{\varepsilon}). \quad (4)$$

We will proceed by a cover argument. First we claim that for each $w \in \mathcal{H}$ that (3) and (4) hold with high probability. For Equation (3), it suffices to show that for each unit vector v in a cover \mathcal{N} of size $2^{O(d)}$ of the sphere that

$$\mathbb{E}_{i \in I_{\text{good}}}[(v \cdot (\nabla f_i(w) - \bar{f}))^2] \leq 2\sigma^2/3. \quad (5)$$

However, we note that

$$\mathbb{E}_{p^*}[(v \cdot (\nabla f(w) - \bar{f}))^2] \leq \sigma^2/2.$$

Since $|v \cdot (\nabla f(w) - \bar{f})|$ is always bounded by L , Equation (5) holds for each v, w with probability at least $1 - \exp(-\Omega(n\sigma^2/L^2))$ by a Chernoff bound (noting that the removal of an ε -fraction of points cannot increase this by much). Similarly, to show Equation 4, it suffices to show that for each such v that

$$\mathbb{E}_{i \in I_{\text{good}}}[(v \cdot (\nabla f_i(w) - \bar{f}))] \leq O(\sigma\sqrt{\varepsilon}). \quad (6)$$

Noting that

$$\mathbb{E}_{p^*}[(v \cdot (\nabla f(w) - \bar{f}))] = 0$$

A Chernoff bound implies that with probability $1 - \exp(-\Omega(n\sigma^2\varepsilon/L^2))$ that the average over our original set of f 's of $(v \cdot (\nabla f(w) - \bar{f}))$ is $O(\sigma\sqrt{\varepsilon})$. Assuming that Equation (5) holds, removing an ε -fraction of these f 's cannot change this value by more than $O(\sigma\sqrt{\varepsilon})$. By union bounding over \mathcal{N} and standard net arguments, this implies that Equations (3) and (4) hold with probability $1 - \exp(-\Omega(dn\sigma^2\varepsilon/L^2))$ for any given w .

To show that our conditions hold for all $w \in \mathcal{H}$, we note that by β -smoothness, if Equation (4) holds for some w , it holds for all other w' in a ball of radius $\sqrt{\sigma^2\varepsilon}/\beta$ (up to a constant multiplicative loss). Similarly, if Equation (3) holds at some w , it holds with bound $\sigma^2 I$ for all w' in a ball of radius $\sigma^2/(2L\beta)$. Therefore, if Equations (3) and (4) hold for all w in a $\min(\sqrt{\sigma^2\varepsilon}/\beta, \sigma/(2L\beta))$ -cover of \mathcal{H} , the assumptions of Theorem B.2 will hold everywhere. Since we have such covers of size $\exp(O(d \log(r\beta L/(\sigma^2\varepsilon))))$, by a union bound, this holds with high probability if

$$n = \Omega\left(\frac{d^2 L^2 \log(r\beta L/\sigma^2\varepsilon)}{\sigma^2\varepsilon}\right),$$

as claimed.

C Analysis of Sever for GLMs

A case of particular interest is that of Generalized Linear Models (GLMs):

Definition C.1. Let $\mathcal{H} \subseteq \mathbb{R}^d$ and \mathcal{Y} be an arbitrary set. Let D_{xy} be a distribution over $\mathcal{H} \times \mathcal{Y}$. For each $Y \in \mathcal{Y}$, let $\sigma_Y : \mathbb{R} \rightarrow \mathbb{R}$ be a convex function. The *generalized linear model* (GLM) over $\mathcal{H} \times \mathcal{Y}$ with distribution D_{xy} and link functions σ_Y is the function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ defined by $\bar{f}(w) = \mathbb{E}_{X,Y}[f_{X,Y}(w)]$, where

$$f_{X,Y}(w) := \sigma_Y(w \cdot X) .$$

A *sample* from this GLM is given by $f_{X,Y}(w)$ where $(X, Y) \sim D_{xy}$.

Our goal, as usual, is to approximately minimize \bar{f} given ε -corrupted samples from D_{xy} . Throughout this section we assume that \mathcal{H} is contained in the ball of radius r around 0, i.e. $\mathcal{H} \subseteq B(0, r)$. Moreover, we will let $w^* = \arg \min_{w \in \mathcal{H}} \bar{f}(w)$ be a minimizer of \bar{f} in \mathcal{H} .

This case covers a number of interesting applications, including SVMs and logistic regression. Unfortunately, the tools developed in Appendix B do not seem to be able to cover this case in a simple manner. In particular, it is unclear how to demonstrate that Assumption B.1 holds after taking polynomially many samples from a GLM. To rectify this, in this section, we demonstrate a different deterministic regularity condition under which we show SEVER succeeds, and we show that this condition holds after polynomially many samples from a GLM. Specifically, we will show that SEVER succeeds under the following deterministic condition:

Assumption C.1. Fix $0 < \varepsilon < 1/2$. There exists an unknown set $I_{\text{good}} \subseteq [n]$ with $|I_{\text{good}}| \geq (1 - \varepsilon)n$ of “good” functions $\{f_i\}_{i \in I_{\text{good}}}$ and parameters $\sigma_0, \sigma_2 \in \mathbb{R}_+$ such that the following conditions simultaneously hold:

- Equation (1) holds with $\sigma_1 = 0$ and the same σ_0 , and
- The following equations hold:

$$\|\nabla \hat{f}(w^*) - \nabla \bar{f}(w^*)\|_2 \leq \sigma_0 \sqrt{\varepsilon} \quad , \text{ and} \tag{7}$$

$$|\hat{f}(w) - \bar{f}(w)| \leq \sigma_2 \sqrt{\varepsilon}, \text{ for all } w \in \mathcal{H} \quad , \tag{8}$$

where $\hat{f} \stackrel{\text{def}}{=} \frac{1}{|I_{\text{good}}|} \sum_{i \in I_{\text{good}}} f_i$.

In this section, we will show the following two statements. The first demonstrates that Assumption C.1 implies that SEVER succeeds, and the second shows that Assumption C.1 holds after polynomially many samples from a GLM. Formally:

Theorem C.2. For functions $f_1, \dots, f_n : \mathcal{H} \rightarrow \mathbb{R}$, suppose that Assumption C.1 holds and that \mathcal{H} is convex. Then, for some universal constant ε_0 , if $\varepsilon < \varepsilon_0$, there is an algorithm which, with probability at least $9/10$, finds a $w \in \mathcal{H}$ such that

$$\bar{f}(w) - \bar{f}(w^*) = r(\gamma + O(\sigma_0 \sqrt{\varepsilon})) + O(\sigma_2 \sqrt{\varepsilon}) .$$

If the link functions are ξ -strongly convex, the algorithm finds a $w \in \mathcal{H}$ such that

$$\bar{f}(w) - \bar{f}(w^*) = 2 \frac{(\gamma + O(\sigma_0 \sqrt{\varepsilon}))^2}{\xi} + O(\sigma_2 \sqrt{\varepsilon}) .$$

Proposition C.3. Let $\mathcal{H} \subseteq \mathbb{R}^d$ and let \mathcal{Y} be an arbitrary set. Let f_1, \dots, f_n be obtained by picking f_i i.i.d. at random from a GLM \bar{f} over $\mathcal{H} \times \mathcal{Y}$ with distribution D_{xy} and link functions σ_Y , where

$$n = \Omega \left(\frac{d \log(dr/\varepsilon)}{\varepsilon} \right) .$$

Suppose moreover that the following conditions all hold:

1. $E_{X \sim D_{xy}}[XX^T] \preceq I$,
2. $|\sigma'_Y(t)| \leq 1$ for all $Y \in \mathcal{Y}$ and $t \in \mathbb{R}$, and
3. $|\sigma_Y(0)| \leq 1$ for all $Y \in \mathcal{Y}$.

Then with probability at least $9/10$ over the original set of samples, there is a set of $(1 - \varepsilon)n$ of the f_i that satisfy Assumption C.1 on \mathcal{H} with $\sigma_0 = 2$, $\sigma_1 = 0$ and $\sigma_2 = 1 + r$. and $\sigma_2 = 1 + r$.

C.1 Proof of Theorem C.2

As before, since SEVER either terminates or throws away at least one sample, clearly it cannot run for more than n iterations. Thus the runtime bound is simple, and it suffices to show correctness.

We first prove the following lemma:

Lemma C.4. *Let f_1, \dots, f_n satisfy Assumption C.1. Then with probability at least $9/10$, SEVER applied to $f_1, \dots, f_n, \sigma_0$ returns a point $w \in \mathcal{H}$ which is a $(\gamma + O(\sigma_0\sqrt{\varepsilon}))$ -approximate critical point of \hat{f} .*

Proof. We claim that the empirical distribution over f_1, \dots, f_n satisfies Assumption B.1 for the function \hat{f} with σ_0 as stated and $\sigma_1 = 0$, with the I_{good} in Assumption B.1 being the same as in the definition of Assumption C.1. Clearly these functions satisfy (2) (since the LHS is zero), so it suffices to show that they satisfy (1) Indeed, we have that for all $w \in \mathcal{H}$,

$$\mathbb{E}_{I_{\text{good}}}[(\nabla f_i(w) - \nabla \hat{f}(w))(\nabla f_i(w) - \nabla \hat{f}(w))^\top] \preceq \mathbb{E}_{I_{\text{good}}}[(\nabla f_i(w) - \nabla \bar{f}(w))(\nabla f_i(w) - \nabla \bar{f}(w))^\top],$$

so they satisfy (1), since the RHS is bounded by Assumption C.1. Thus this lemma follows from an application of Theorem B.2. \square

With this critical lemma in place, we can now prove Theorem C.2:

Proof of Theorem C.2. Condition on the event that Lemma C.4 holds, and let $w \in \mathcal{H}$ be the output of SEVER. By Assumption C.1, we know that $\hat{f}(w^*) \geq \bar{f}(w^*) - \sigma_2\sqrt{\varepsilon}$, and moreover, w^* is a $\gamma + \sigma_0\sqrt{\varepsilon}$ -approximate critical point of \hat{f} .

Since each link function is convex, so is \hat{f} . Hence, by Lemma B.8, since w is a $(\gamma + O(\sigma_0\sqrt{\varepsilon}))$ -approximate critical point of \hat{f} , we have $\hat{f}(w) - \hat{f}(w^*) \leq r(\gamma + O(\sigma_0\sqrt{\varepsilon}))$. By Assumption B.1, this immediately implies that $\bar{f}(w) - \bar{f}(w^*) \leq r(\gamma + O(\sigma_0\sqrt{\varepsilon})) + O(\sigma_2\sqrt{\varepsilon})$, as claimed.

The bound for strongly convex functions follows from the exact argument, except using the statement in Lemma B.8 pertaining to strongly convex functions. \square

C.2 Proof of Proposition C.3

Proof. We first note that $\nabla f_{X,Y}(w) = X\sigma'_Y(w \cdot X)$. Thus, under Assumption C.1, we have for any v that

$$\mathbb{E}_i[(v \cdot (\nabla f_i(w) - \nabla \bar{f}(w)))^2] \ll \mathbb{E}_i[(v \cdot \nabla f_i(w))^2] + 1 \ll \mathbb{E}_i[(v \cdot X_i)^2] + 1.$$

In particular, since this last expression is independent of w , we only need to check this single matrix bound.

We let our good set be the set of samples with $|X| \leq 80\sqrt{d/\varepsilon}$ that were not corrupted. We use Lemma A.18 of [DKK⁺17]. This shows that with 90% probability that the non-good samples make up at most an $\varepsilon/2 + \varepsilon/160$ -fraction of the original samples, and that $\mathbb{E}[XX^T]$ over the good samples is at most $2I$. This proves that the spectral bound holds everywhere. Applying it to the $\nabla f_{X,Y}(w^*)$, we find also with 90% probability that the expectation over all samples of $\nabla f_{X,Y}(w^*)$ is within $\sqrt{\varepsilon}/3$ of $\nabla \bar{f}(w^*)$. Additionally, throwing away the samples with $|\nabla f_{X,Y}(w^*) - \nabla \bar{f}(w^*)| > 80\sqrt{d/\varepsilon}$ changes this by at most $\sqrt{\varepsilon}/2$. Finally, it also implies that the variance of $\nabla f_{X,Y}(w^*)$ is at most $3/2I$, and therefore, throwing away any other ε -fraction of the samples changes it by at most an additional $\sqrt{3\varepsilon/2}$.

We only need to show that $|\mathbb{E}_{i \text{ good}}[f_i(w)] - \mathbb{E}_X[f_X(w)]| \leq \sqrt{\varepsilon}$ for all $w \in \mathcal{H}$. For this we note that since the f_X and f_i are all 1-Lipschitz, it suffices to show that $|\mathbb{E}_{i \text{ good}}[f_i(w)] - \mathbb{E}_X[f_X(w)]| \leq (1 + |w|)\sqrt{\varepsilon}/2$ on an $\varepsilon/2$ -cover of \mathcal{H} . For this it suffices to show that the bound will hold pointwise except with probability $\exp(-\Omega(d \log(r/\varepsilon)))$. We will want to bound this using pointwise concentration and union bounds, but this runs into technical problems since very large values of $X \cdot w$ can lead to large values of f , so we will need to make use of the condition above that the average of $X_i X_i^T$ over our good samples is bounded by $2I$. In particular, this implies that the contribution to the average of $f_i(w)$ over the good i coming from samples where $|X_i \cdot w| \geq 10|w|/\sqrt{\varepsilon}$ is at most $\sqrt{\varepsilon}(1 + |w|)/10$. We consider the average of $f_i(w)$ over the remaining i . Note that these values are uniform random samples from $f_X(w)$ conditioned on $|X| \leq 80\sqrt{d}/\varepsilon$ and $|X_i \cdot w| < 10|w|/\sqrt{\varepsilon}$. It will suffice to show that taking n samples from this distribution has average within $(1 + |w|)\sqrt{\varepsilon}/2$ of the mean with high probability. However, since $|f_X(w)| \leq O(1 + |X \cdot w|)$, we have that over this distribution $|f_X(w)|$ is always $O(1 + |w|/\sqrt{\varepsilon})$, and has variance at most $O(1 + |w|)^2$. Therefore, by Bernstein's Inequality, the probability that n random samples from $f_X(w)$ (with the above conditions on X) differ from their mean by more than $(1 + |w|)\sqrt{\varepsilon}/2$ is

$$\exp(-\Omega(n^2(1 + |w|)^2\varepsilon/((1 + |w|)^2 + n(1 + |w|)^2))) = \exp(-\Omega(n\varepsilon)).$$

Thus, for n at least a sufficiently large multiple of $d \log(dr/\varepsilon)/\varepsilon$, this holds for all w in our cover of \mathcal{H} with high probability. This completes the proof. \square

D An Alternative Algorithm: Robust Filtering in Each Iteration

In this section, we describe another algorithm for robust stochastic optimization. This algorithm uses standard robust mean estimation techniques to compute approximate gradients pointwise, which it then feeds into a standard projective gradient descent algorithm. This algorithm in practice turns out to be somewhat slower than the one employed in the rest of this paper, because it employs a filtering algorithm at every step of the projective gradient descent, and does not remember which points were filtered between iterations. On the other hand, we present this algorithm for two reasons. Firstly, because it is a conceptually simpler interpretation of the main ideas of this paper, and secondly, because the algorithm works under somewhat more general assumptions. In particular, this algorithm only requires that for each $w \in \mathcal{H}$ that there is a corresponding good set of functions, rather than that there exists a single good set that works simultaneously for all w .

In particular, we can make do with the following somewhat weaker assumption:

Assumption D.1. Fix $0 < \varepsilon < 1/2$ and parameter $\sigma \in \mathbb{R}_+$. For each $w \in \mathcal{H}$, there exists an unknown set $I_{\text{good}} = I_{\text{good}}(w) \subseteq [n]$ with $|I_{\text{good}}| \geq (1 - \varepsilon)n$ of “good” functions $\{f_i\}_{i \in I_{\text{good}}}$ such that:

$$\left\| \mathbb{E}_{I_{\text{good}}} [(\nabla f_i(w) - \nabla \bar{f}(w))(\nabla f_i(w) - \nabla \bar{f}(w))^T] \right\|_2 \leq (\sigma)^2, \quad (9)$$

and

$$\|\nabla \hat{f}(w) - \nabla \bar{f}(w)\|_2 \leq \sigma\sqrt{\varepsilon}, \text{ where } \hat{f} \stackrel{\text{def}}{=} \frac{1}{|I_{\text{good}}|} \sum_{i \in I_{\text{good}}} f_i. \quad (10)$$

We make essential use of the following result, which appears in both [DKK⁺17, SCV18]:

Theorem D.2. Let $\mu \in \mathbb{R}^d$ and a collection of points $x_i \in \mathbb{R}^d$, $i \in [n]$ and $\sigma > 0$. Suppose that there exists $I_{\text{good}} \subseteq [n]$ with $|I_{\text{good}}| \geq (1 - \varepsilon)n$ satisfying the following:

$$\frac{1}{|I_{\text{good}}|} \sum_{i \in I_{\text{good}}} (x_i - \mu)(x_i - \mu)^\top \preceq \sigma^2 I \text{ and } \left\| \frac{1}{|I_{\text{good}}|} \sum_{i \in I_{\text{good}}} (x_i - \mu) \right\|_2 \leq \sigma\sqrt{\varepsilon}. \quad (11)$$

Then, if $\varepsilon < \varepsilon_0$ for some universal constant ε_0 , there is an efficient algorithm, Algorithm \mathcal{A} , which outputs an estimate $\hat{\mu} \in \mathbb{R}^d$ such that $\|\hat{\mu} - \mu\|_2 = O(\sigma\sqrt{\varepsilon})$.

Our general robust algorithm for stochastic optimization will make calls to Algorithm \mathcal{A} in a black-box manner, as well as to the projection operator onto \mathcal{H} . We will measure the cost of our algorithm by the total number of such calls.

Remark D.1. While it is not needed for the theoretical results established in this subsection, we note that the robust mean estimation algorithm of [DKK⁺17] relies on an iterative outlier removal method only requiring basic eigenvalue computations (SVD), while the [SCV18] algorithm employs semidefinite programming. In our experiments, we use the algorithm in [DKK⁺17] and variants thereof.

Using the above black-box, together with known results on convex optimization with errors, we obtain the following meta-theorem:

Theorem D.3. *For functions $f_1, \dots, f_n : \mathcal{H} \rightarrow \mathbb{R}$, bounded below on a closed domain \mathcal{H} , suppose that either Assumption D.1 is satisfied with some parameters $\varepsilon, \sigma > 0$. Then there exists an efficient algorithm that finds an $O(\sigma\sqrt{\varepsilon})$ -approximate critical point of \bar{f} .*

Proof. We note that by applying Algorithm A on $\{\nabla f_i(w)\}$, we can find an approximation to $\nabla \bar{f}(w)$ with error $O(\sigma\sqrt{\varepsilon})$. We note that standard projective gradient descent algorithms can be made to run efficiently even if the gradients given are only approximate, and this can be used to find our $O(\sigma\sqrt{\varepsilon})$ -approximate critical point. \square

E Applications of the General Algorithm

In this section, we present three concrete applications of our general robust algorithm. In particular, we describe how to robustly optimize models for linear regression, support vector machines, and logistic regression, in Sections E.1, E.2, E.3, respectively.

E.1 Linear Regression

In this section, we demonstrate how our results apply to linear regression. We are given pairs $(X_i, Y_i) \in \mathbb{R}^d \times \mathbb{R}$ for $i \in [n]$. The X_i 's are drawn i.i.d. from a distribution D_x , and $Y_i = \langle w^*, X_i \rangle + e_i$, for some unknown $w^* \in \mathbb{R}^d$ and the noise random variables e_i 's are drawn i.i.d. from some distribution D_e . Given $(X_i, Y_i) \sim D_{xy}$, the joint distribution induced by this process, let $f_i(w) = (Y_i - \langle w, X_i \rangle)^2$. The goal is then to find a \hat{w} approximately minimizing the objective function

$$\bar{f}(w) = \mathbb{E}_{(X,Y) \sim D_{xy}} [(Y - \langle w, X \rangle)^2].$$

We work with the following assumptions:

Assumption E.1. Given the model for linear regression described above, assume the following conditions for D_e and D_x :

- $\mathbb{E}_{e \sim D_e} [e] = 0$;
- $\text{Var}_{e \sim D_e} [e] \leq \xi$;
- $\mathbb{E}_{X \sim D_x} [XX^T] \preceq \sigma^2 I$ for some $\sigma > 0$;
- There is a constant $C > 0$, such that for all unit vectors v , $\mathbb{E}_{X \sim D_x} [\langle v, X \rangle^4] \leq C\sigma^4$.

Our main result for linear regression is the following:

Theorem E.2. *Let $\varepsilon > 0$, and let D_{xy} be a distribution over pairs (X, Y) which satisfies the conditions of Assumption E.1. Suppose we are given $O\left(\frac{d^5}{\varepsilon^2}\right)$ ε -noisy samples from D_{xy} . Then in either of the following two cases, there exists an algorithm that, with probability at least 9/10, produces a \hat{w} with the following guarantees:*

1. If $\mathbb{E}_{X \sim D_x}[XX^T] \succeq \gamma I$ for $\gamma = \Omega(\sqrt{\varepsilon})$, then $\bar{f}(\hat{w}) \leq \bar{f}(w^*) + O\left(\frac{(\xi + \varepsilon)\varepsilon}{\gamma}\right)$ and $\|\hat{w} - w^*\|_2 = O\left(\frac{\sqrt{\xi\varepsilon + \varepsilon}}{\gamma}\right)$.
2. If $\|w^*\|_2 \leq r$, then $\bar{f}(\hat{w}) \leq \bar{f}(w^*) + O((\sqrt{\xi} + \sqrt{\varepsilon})r + \sqrt{Cr^2}\sqrt{\varepsilon})$.

The proof will follow from two lemmas (proved in Section E.1.1 and E.1.2, respectively). First, we will bound the covariance of the gradient, in Lemma E.3:

Lemma E.3. *Suppose D_{xy} satisfies the conditions of Assumption E.1. Then for all unit vectors $v \in \mathbb{R}^d$, we have*

$$v^\top \text{Cov}_{(X,Y) \sim D_{xy}}[\nabla f_i(w, (X, Y))] v \leq 4\sigma^2\xi + 4C\sigma^4\|w^* - w\|_2^2.$$

With this in hand, we can prove Lemma E.4, giving us a polynomial sample complexity which is sufficient to satisfy the conditions of Assumption B.1.

Lemma E.4. *Suppose D_{xy} satisfies the conditions of Assumption E.1. Given $O(d^5/\varepsilon^2)$ ε -noisy samples from D_{xy} , then with probability at least 9/10, they satisfy Assumption B.1 with parameters $\sigma_0 = 30\sqrt{\xi} + \sqrt{\varepsilon}$ and $\sigma_1 = 18\sqrt{C + 1}$.*

The proof concludes by applying Corollary B.4 or case (i) of Corollary B.3 for the first and second cases respectively.

E.1.1 Proof of Lemma E.3

Note that for this setting we have that $f(w, z) = f(w, x, y) = (y - \langle w, x \rangle)^2$. We then have that $\nabla_w f(w, z) = -2(\langle w^* - w, x \rangle + e)x$. Our main claim is the following:

Claim E.5. We have that $\text{Cov}[\nabla_w f(w, z)] = 4\mathbb{E}_{X \sim D}[\langle w^* - w, x \rangle^2 (xx^T)] + 4\text{Var}[E]\Sigma - 4\Sigma(w^* - w)(w^* - w)^T\Sigma$.

Proof. Let us use the notation $A = \nabla_w f(w, z)$ and $\mu = \mathbb{E}[A]$. By definition, we have that $\text{Cov}[A] = \mathbb{E}[AA^T] - \mu\mu^T$.

Note that $\mu = \mathbb{E}_z[\nabla_w f(w, z)] = \mathbb{E}_z[(-2\langle w^* - w, x \rangle + e)x] = -2\Sigma(w^* - w)$, where we use the fact that $\mathbb{E}_z[e] = 0$ and e is independent of x . Therefore, $\mu\mu^T = 4\Sigma(w^* - w)(w^* - w)^T\Sigma$.

To calculate $\mathbb{E}[AA^T]$, note that $A = \nabla_w f(w, z) = -2(\langle w^* - w, x \rangle + e)x$, and $A^T = -2(\langle w^* - w, x \rangle + e)x^T$. Therefore, $AA^T = 4(\langle w^* - w, x \rangle^2 + e^2 + 2\langle w^* - w, x \rangle e)(xx^T)$ and

$$\mathbb{E}_z[AA^T] = 4\mathbb{E}_x[\langle w^* - w, x \rangle^2 (xx^T)] + 4\text{Var}[e]\Sigma + 0,$$

where we again used the fact that the noise e is independent of x and its expectation is zero.

By gathering terms, we get that

$$\text{Cov}[\nabla_w f(w, z)] = 4\mathbb{E}_x[\langle w^* - w, x \rangle^2 (xx^T)] + 4\text{Var}[e]\Sigma - 4\Sigma(w^* - w)(w^* - w)^T\Sigma.$$

This completes the proof. \square

Given the above claim, we can bound from above the spectral norm of the covariance matrix of the gradients as follows: Specifically, for a unit vector v , the quantity $v^T \text{Cov}[\nabla_w f(w, z)]v$ is bounded from above by a constant times the following quantities:

- The first term is $v^T \mathbb{E}_x[\langle w^* - w, x \rangle^2 (xx^T)]v = \mathbb{E}_x[\langle w^* - w, x \rangle^2 \cdot \langle v, x \rangle^2]$. By Cauchy-Schwarz and our 4th moment bound, this is at most $C\sigma^4\|w^* - w\|_2^2$, where $\Sigma \preceq \sigma^2 I$.
- The second term is at most the upper bound of the variance of the noise ξ times σ^2 .
- The third term is at most $v^T \Sigma(w^* - w)(w^* - w)^T \Sigma v$, which by our bounded covariance assumption is at most $\sigma^4\|w^* - w\|_2^2$.

This gives the parameters in the meta-theorem.

E.1.2 Proof of Lemma E.4

Let S be the set of uncorrupted samples and I be the subset of S with $\|X\|_2 \leq 2\sqrt{d}/\varepsilon^{1/4}$. We will take I_{good} to be the subset of I that are not corrupted.

Firstly, we show that with probability at least $39/40$, at most an $\varepsilon/2$ -fraction of points in S have $\|X\|_2 > 2\sqrt{d}/\varepsilon^{1/4}$, and so $|I_{\text{good}}| \geq (1-\varepsilon)|S|$. Note that $\mathbb{E}_D[\|X\|_2^4] = \mathbb{E}_D[(\sum_{j=1}^d X_j^2)^2] \leq \sum_{j=1}^d \sum_{k=1}^d \sqrt{\mathbb{E}_D[X_j^2] \mathbb{E}_D[X_k^2]} \leq Cd^2$, since $\mathbb{E}_{D_x}[XX^T] \preceq I$. Thus, by Markov's inequality, $\Pr_D[\|X\|_2 > 2\sqrt{d}(C/\varepsilon)^{1/4}] = \Pr_D[\|X\|_2^4 > 16d^2/\varepsilon] \leq \varepsilon/16$. By a Chernoff bound, since $n \geq 10\varepsilon^2$ this probability is at most $\varepsilon/2$ for the uncorrupted samples with probability at least $39/40$.

Next, we show that (1) holds with probability at least $39/40$. To do this, we will apply Lemma E.3 to I_{good} . Since S consists of independent samples, the variance over the randomness of S of $|S| \mathbb{E}_S[e^2]$ is at most $|S|\xi$. By Chebyshev's inequality, except with probability $1/99$, we have that $\mathbb{E}_S[e^2] \leq 99\xi$ and since $I_{\text{good}} \subset S$, $\mathbb{E}_{I_{\text{good}}}[e^2] \leq |S| \mathbb{E}_S[e^2]/|I| \leq 100\xi$. This is condition (i) of Lemma E.3.

We note that I consists of $\Omega(d^5/\varepsilon^2)$ independent samples from D conditioned on $\|X\|_2 < 2\sqrt{d}/\varepsilon^{1/4}$, a distribution that we will call D' . Since the VC-dimension of all halfspaces in \mathbb{R}^d is $d+1$, by the VC inequality, we have that, except with probability $1/80$, for any unit vector v and $T \in \mathbb{R}$ that $|\Pr_I[v \cdot X > T] - \Pr_{D'}[v \cdot X > T]| \leq \varepsilon/d^2$. Note that for unit vector v and positive integer m , $\mathbb{E}[(v \cdot X)^m] = \int_0^\infty m(v \cdot X)^{m-1} \Pr[v \cdot X > T] dT$. Thus we have that

$$\begin{aligned} \mathbb{E}_I[(v \cdot X)^m] &= \int_0^\infty m(v \cdot X)^{m-1} \Pr[v \cdot X > T] dT \\ &\leq \int_0^{2d^{1/2}(C/\varepsilon)^{1/4}} m(v \cdot X)^{m-1} (\Pr_{D'}[v \cdot X > T] + \varepsilon/d^2) dT \\ &= \mathbb{E}_{D'}[(v \cdot X)^m] + (2d^{1/2}(C/\varepsilon)^{1/4})^m (\varepsilon/d^2) \\ &\leq (1 + \varepsilon) \mathbb{E}_D[(v \cdot X)^m] + 2^m C^{m/4} (\varepsilon/d^2)^{1-m/4}. \end{aligned}$$

Applying this for $m = 2$ gives $\mathbb{E}_I[XX^T] \preceq (1 + \varepsilon + 4\sqrt{C}\varepsilon/d^2)I \preceq 2I$ and with $m = 4$ gives $\mathbb{E}_I[(v \cdot X)^4] \leq (1 + \varepsilon)C + 16C$. Similar bounds apply to I_{good} , with an additional $1 + \varepsilon$ factor.

Thus, with probability at least $39/40$, I_{good} satisfies the conditions of Lemma E.3 with $\xi := 100\xi$, $\sigma^2 := 2$ and $C := 5C$. Hence, it satisfies (1) with $\sigma_0 = 20\sqrt{\xi}$ and $\sigma_1 = 18\sqrt{C+1}$.

For (2), note that $\nabla_w f_i(w) = (w \cdot x_i - y_i)x_i = ((w - w^*) \cdot x_i)x_i - e_i x_i$. We will separately bound $\|\mathbb{E}_{I_{\text{good}}}[(w - w^*) \cdot X]X - \mathbb{E}_D[(w - w^*) \cdot X]X\|_2$ and $\|\mathbb{E}_{I_{\text{good}}}[eX] - \mathbb{E}_D[eX]\|_2$.

We will repeatedly make use of the following, which bounds how much removing points or probability mass affects an expectation in terms of its variance:

Claim E.6. For a mixture of distributions $P = (1 - \delta)Q + \delta R$ for distributions P, Q, R and a real valued function f , we have that $|\mathbb{E}_{X \sim P}[f(X)] - \mathbb{E}_{X \sim Q}[f(X)]| \leq 2\sqrt{\delta \mathbb{E}_{X \sim P}[f(X)^2]}/(1 - \delta)$

Proof. By Cauchy-Schwarz $|\mathbb{E}_{X \sim R}[f(X)]| \leq \sqrt{\mathbb{E}_{X \sim R}[f(X)^2]} \leq \sqrt{\mathbb{E}_{X \sim P}[f(X)^2]}/\delta$. Since $\mathbb{E}_{X \sim P}[f(X)] = (1 - \delta) \mathbb{E}_{X \sim Q}[f(X)] + \delta \mathbb{E}_{X \sim R}[f(X)]$, this implies that $|\mathbb{E}_{X \sim P}[f(X)]/(1 - \delta) - \mathbb{E}_{X \sim Q}[f(X)]| \leq \sqrt{\delta \mathbb{E}_{X \sim P}[f(X)^2]}/(1 - \delta)$. However $|\mathbb{E}_{X \sim P}[f(X)]/(1 - \delta) - \mathbb{E}_{X \sim P}[f(X)]| = (\delta/(1 - \delta))|\mathbb{E}_{X \sim P}[f(X)]| \leq \sqrt{\delta \mathbb{E}_{X \sim P}[f(X)^2]}/(1 - \delta)$ and the triangle inequality gives the result. \square

We can apply this to $P = I$ and $Q = I_{\text{good}}$ with $\delta = \varepsilon/2$ and also to $P = D$ and $Q = D'$ with $\delta = \varepsilon/16$, with error $2\sqrt{\delta}/(1 - \delta) \leq 2\sqrt{\varepsilon}$ in either case.

For the first of term we wanted to bound, we have $\|\mathbb{E}_{I_{\text{good}}}[(w - w^*) \cdot X]X - \mathbb{E}_D[(w - w^*) \cdot X]X\|_2 = \|(w - w^*)^T (\mathbb{E}_{I_{\text{good}}}[XX^T] - \mathbb{E}_D[XX^T])\|_2 \leq \|w - w^*\|_2 \|\mathbb{E}_{I_{\text{good}}}[XX^T] - \mathbb{E}_D[XX^T]\|_2$. For any unit vector v , the VC dimension argument above gave that $|\mathbb{E}_I[(v \cdot X)^2] - \mathbb{E}_{D'}[(v \cdot X)^2]| \leq 4\sqrt{C}\varepsilon/d^2$ and Claim E.6 both gives that $|\mathbb{E}_I[(v \cdot X)^2] - \mathbb{E}_{I_{\text{good}}}[(v \cdot X)^2]| \leq 2\sqrt{\varepsilon} \mathbb{E}_I[(v \cdot X)^4] \leq 10\sqrt{C}\varepsilon$ and that $|\mathbb{E}_D[(v \cdot X)^2] - \mathbb{E}_{D'}[(v \cdot X)^2]| \leq 2\sqrt{\varepsilon} \mathbb{E}_D[(v \cdot X)^4] \leq 2\sqrt{C}\varepsilon$. By the triangle inequality, we have that $|\mathbb{E}_D[(v \cdot X)^2] - \mathbb{E}_{I_{\text{good}}}[(v \cdot X)^2]| \leq 16\sqrt{C}\varepsilon$. Since this holds for all unit v and the matrices involved are symmetric, we have that $\|\mathbb{E}_{I_{\text{good}}}[XX^T] -$

$\mathbb{E}_D[XX^T]\|_2 \leq 16\sqrt{C\varepsilon}$. The overall first term is bounded by $\|\mathbb{E}_{I_{\text{good}}}[(w - w^*) \cdot X]X - \mathbb{E}_D[(w - w^*) \cdot X]X\|_2 \leq 16\sqrt{C\varepsilon}\|w - w^*\|_2$.

Now we want to bound the second term, $\|\mathbb{E}_{I_{\text{good}}}[eX] - \mathbb{E}_D[eX]\|_2$. Note that $\mathbb{E}_D[eX] = \mathbb{E}_D[e]\mathbb{E}_D[X] = 0$. So we need to bound $\mathbb{E}_{I_{\text{good}}}[eX]$. First we bound the expectation and variance on D' using Claim E.6. It yields that, for any unit vector v , $|\mathbb{E}_{D'}[e(v \cdot X)]| \leq 2\sqrt{\varepsilon\mathbb{E}_D[e^2(v \cdot X)^2]} \leq 2\sqrt{\varepsilon\xi}$.

Next we bound the expectation on I . Since I consists of independent samples from D' , the covariance matrix over the randomness on I of $|I|\mathbb{E}_I[eX - \mathbb{E}_{D'}[eX]]$ is $|I|\mathbb{E}_{D'}[(eX - \mathbb{E}_{D'}[eX])(eX - \mathbb{E}_{D'}[eX])^T] \leq |I|\mathbb{E}_{D'}[XX^T] \leq |I|(1 + \varepsilon)I$ and its expectation is 0. Thus the expectation over the randomness of I of $(|I|^2\|\mathbb{E}_I[eX] - \mathbb{E}_{D'}[eX]\|_2)^2$ is $\text{Tr}(|I|\mathbb{E}_{D'}[(eX - \mathbb{E}_{D'}[eX])(eX - \mathbb{E}_{D'}[eX])^T]) \leq |I|(1 + \varepsilon + 4\xi\varepsilon|I|)d$. By Markov's inequality, except with probability $1/40$, $\Pr[\|\mathbb{E}_I[eX]\|_2 \geq 2\sqrt{\xi\varepsilon + \varepsilon}] \leq d/|I|\varepsilon^2$. Since $|I| \geq 40d/\varepsilon^2$. This happens with probability at least $1/40$.

Next we bound the expectation on I_{good} which follows by a slight variation of Claim E.6. Let $J = I - I_{\text{good}}$. Then, for any v , $\mathbb{E}_J[e(v \cdot X)] \leq \sqrt{\mathbb{E}_J[e^2]\mathbb{E}_J[(v \cdot X)^2]} \leq \sqrt{\mathbb{E}_S[e^2]\mathbb{E}_I[(v \cdot X)^2]}|J|/\sqrt{|S||I|} \leq \sqrt{100\xi(1 + \varepsilon + 4\sqrt{C\varepsilon}/d^2)}|J|/\sqrt{|S||I|} \leq 20|J|\sqrt{\xi/|S||I|}$ by bounds we obtained earlier. Now $\|\mathbb{E}_{I_{\text{good}}}[eX]\|_2 = \|(|I|/|I_{\text{good}}|)\mathbb{E}_I[eX] - (|J|/|I_{\text{good}}|)\mathbb{E}_J[eX]\|_2 \leq 20\sqrt{\xi\varepsilon} + \varepsilon + (1 + \varepsilon)\sqrt{\xi\varepsilon/16} \leq 30\sqrt{\xi\varepsilon} + \varepsilon$.

We can thus take $\sigma_0 = 30\sqrt{\xi\varepsilon} + \sqrt{\varepsilon}$ and $\sigma_1 = 18\sqrt{C + 1} \geq 16\sqrt{C}$ to get (2).

To get both (2) and (1) hold with $\sigma_0 = 30\sqrt{\xi\varepsilon} + \sqrt{\varepsilon}$ and $\sigma_1 = 18\sqrt{C + 1}$. This happens with probability at least $9/10$ by a union bound on the probabilistic assumptions above.

E.2 Support Vector Machines

In this section, we demonstrate how our results apply to learning support vector machines (i.e., halfspaces under hinge loss). In particular, we describe how SVMs fit into the GLM framework described in Section C.

We are given pairs $(X_i, Y_i) \in \mathbb{R}^d \times \{\pm 1\}$ for $i \in [n]$, which are drawn from some distribution D_{xy} . Let $L(w, (x, y)) = \max\{0, 1 - y(w \cdot x)\}$, and $f_i(w) = L(w, (x_i, y_i))$. The goal is to find a \hat{w} approximately minimizing the objective function

$$\bar{f}(w) = \mathbb{E}_{(X, Y) \sim D_{xy}}[L(w, (X, Y))].$$

One technical point is that f_i does not have a gradient everywhere – instead, we will be concerned with the sub-gradients of the f_i 's. All our results which operate on the gradients also work for sub-gradients. To be precise, we will take the sub-gradient to be 0 when the gradient is undefined:

Definition E.1. Let ∇f_i be the *sub-gradient* of $f_i(w)$ with respect to w , where $\nabla f_i = -y_i x_i$ if $y_i(w \cdot x_i) < 1$, and 0 otherwise.

To get a bound on the error of hinge loss, we will need to assume the marginal distribution D_x is anti-concentrated.

Definition E.2. A distribution is δ -*anticoncentrated* if at most an $O(\delta)$ -fraction of its probability mass is within Euclidean distance δ of any hyperplane.

We work with the following assumptions:

Assumption E.7. Given the model for SVMs as described above, assume the following conditions for the marginal distribution D_x :

- $\mathbb{E}_{X \sim D_x}[XX^T] \preceq I$;
- D_x is $\varepsilon^{1/4}$ -anticoncentrated.

Our main result on SVMs is the following:

Theorem E.8. Let $\varepsilon > 0$, and let D_{xy} be a distribution over pairs (X, Y) , where the marginal distribution D_x satisfies the conditions of Assumption E.7. Then there exists an algorithm that with probability $9/10$, given $O(d \log(d/\varepsilon)/\varepsilon)$ ε -noisy samples from D_{xy} , returns a \hat{w} such that for any w^* ,

$$\mathbb{E}_{(X,Y) \sim D_{xy}}[L(\hat{w}, (X, Y))] \leq \mathbb{E}_{(X,Y) \sim D_{xy}}[L(w^*, (X, Y))] + O(\varepsilon^{1/4}).$$

Our approach will be to fit this problem into the GLM framework developed in Section C. First, we will restrict our search over w to \mathcal{H} , a ball of radius $r = \varepsilon^{-1/4}$. As we argue in Lemma E.9, this restriction comes at a cost of at most $O(\varepsilon^{1/4})$ in our algorithm's loss. With this restriction, we will argue that the problem satisfies the conditions of Proposition C.3. This allows us to argue that, with a polynomial number of samples, we can obtain a set of f_i 's satisfying the conditions of Assumption C.1. This will allow us to apply Theorem C.2, concluding the proof.

We start by showing that, due to anticoncentration of D , there is a $w' \in \mathcal{H}$ with loss close to w^* :

Lemma E.9. Let w' be a rescaling of w^* , such that $\|w'\|_2 \leq \varepsilon^{-1/4}$ (i.e. $w' = \min\{1, \varepsilon^{-1/4}/\|w^*\|_2\}w^*$). Then $\mathbb{E}_{(X,Y) \sim D_{xy}}[L(w', (X, Y))] \leq \mathbb{E}_{(X,Y) \sim D_{xy}}[L(w^*, (X, Y))] + O(\varepsilon^{1/4})$.

Proof. If $w' = w^*$, then $\mathbb{E}_{(X,Y) \sim D_{xy}}[L(w', (X, Y))] = \mathbb{E}_{(X,Y) \sim D_{xy}}[L(w^*, (X, Y))]$.

Otherwise, we break into case analysis, based on the value of (x, y) :

- $|w' \cdot x| > 1$: If $y(w' \cdot x) > 1$, then $L(w', (x, y)) = L(w^*, (x, y)) = 0$. If $y(w' \cdot x) < -1$, then $L(w', (x, y)) = 1 - y(w' \cdot x) \leq 1 - y(w^* \cdot x) = L(w^*, (x, y))$. Both cases use the fact that $\|w'\|_2 < \|w^*\|_2$.
- $|w' \cdot x| \leq 1$: In this case, we have that $L(w', (x, y)) \leq 2$. Since $L(w^*, (x, y)) \geq 0$, we have that $L(w', (x, y)) \leq L(w^*, (x, y)) + 2$.

Note that if $|w' \cdot x| \leq 1$, then x is within $1/\|w'\|_2 = \varepsilon^{1/4}$ of the hyperplane defined by the normal vector w' . Since D_x is $\varepsilon^{1/4}$ -anticoncentrated, we have that $\Pr_{X \sim D_x}[|w' \cdot X| \leq 1] \leq \varepsilon^{1/4}$. Thus, we have that $\mathbb{E}_{(X,Y) \sim D_{xy}}[L(w', (X, Y))] \leq \mathbb{E}_{(X,Y) \sim D_{xy}}[L(w^*, (X, Y))] + 2 \cdot \mathbb{1}(|w' \cdot X| \leq 1) \leq \mathbb{E}_{(X,Y) \sim D_{xy}}[L(w^*, (X, Y))] + O(\varepsilon^{1/4})$. \square

Proof of Theorem E.8. We first show that this problem fits into the GLM framework, in particular, satisfying the conditions of Proposition C.3. The link function is $\sigma_y(t) = \max\{0, 1 - yt\}$, giving us the loss function $L(w, (x, y)) = \sigma_y(w \cdot x)$. We let \mathcal{H} be the set $\|w\|_2 \leq \varepsilon^{-1/4}$, giving us the parameter $r = \varepsilon^{-1/4}$. Condition 1 is satisfied by Assumption E.7. For $y \in \{-1, 1\}$, $\sigma'_y(t) = 0$ for $yt \geq 1$ and $\sigma'_y(t) = -y$ for $yt < 1$. Thus we have that $|\sigma'_1(t)| \leq 1$ for all t and y , satisfying Condition 2. Finally, one can observe that $\sigma_y(0) = 1$ for all y , satisfying Condition 3. Thus we can apply Proposition C.3: if we take $O(d \log(dr/\varepsilon)/\varepsilon)$ ε -corrupted samples, then they satisfy Assumption C.1 on \mathcal{H} with $\sigma_0 = 2$, $\sigma_1 = 0$ and $\sigma_2 = 1 + \varepsilon^{-1/4}$, with probability $9/10$.

Now we can apply the algorithm of Theorem C.2. Since the loss is convex, we get a vector \hat{w} with $\bar{f}(\hat{w}) - \bar{f}(w^{*'}) = O((\sigma_0 r + \sigma_1 r^2 + \sigma_2)\sqrt{\varepsilon}) = O((2\varepsilon^{-1/4} + \varepsilon^{-1/4})\sqrt{\varepsilon}) = O(\varepsilon^{1/4})$ where $w^{*'}$ is the minimizer of \bar{f} on \mathcal{H} .

We thus have that $\bar{f}(\hat{w}) \leq \bar{f}(w^{*'}) + O(\varepsilon^{1/4}) \leq \bar{f}(w') + O(\varepsilon^{1/4}) \leq \bar{f}(w^*) + O(\varepsilon^{1/4})$. The second inequality follows because $w^{*'}$ is the minimizer of \bar{f} on \mathcal{H} , and the third inequality follows from Lemma E.9. \square

E.3 Logistic Regression

In this section, we demonstrate how our results apply to logistic regression. In particular, we describe how logistic regression fits into the GLM framework described in Section C.

We are given pairs $(X_i, Y_i) \in \mathbb{R}^d \times \{\pm 1\}$ for $i \in [n]$, which are drawn from some distribution D_{xy} . Let $\phi(t) = \frac{1}{1 + \exp(-t)}$. Logistic regression is the model where $y = 1$ with probability $\phi(w \cdot x)$, and $y = -1$ with probability $\phi(-w \cdot x)$. We define the loss function to be the log-likelihood of y given x . More precisely, we let $f_i(w, (x_i, y_i)) = L(w, (x_i, y_i))$, which is defined as follows:

$$L(w, (x, y)) = \frac{1+y}{2} \ln \left(\frac{1}{\phi(w \cdot x)} \right) + \frac{1-y}{2} \ln \left(\frac{1}{\phi(-w \cdot x)} \right) = \frac{1}{2} (-\ln(\phi(w \cdot x)\phi(-w \cdot x)) - y(w \cdot x)).$$

The gradient of this function is $\nabla L(w, (x, y)) = \frac{1}{2}(\phi(w \cdot x) - \phi(-w \cdot x) - y)x$. The goal is to find a \hat{w} approximately minimizing the objective function

$$\bar{f}(w) = \mathbb{E}_{(X, Y) \sim D_{xy}} [L(w, (X, Y))].$$

We work with the following assumptions:

Assumption E.10. Given the model for logistic regression as described above, assume the following conditions for the marginal distribution D_x :

- $\mathbb{E}_{X \sim D_x} [XX^T] \preceq I$;
- D_x is $\varepsilon^{1/4} \sqrt{\log(1/\varepsilon)}$ -anticoncentrated.

We can get a similar result to that for hinge loss for logistic regression:

Theorem E.11. *Let $\varepsilon > 0$, and let D_{xy} be a distribution over pairs (X, Y) , where the marginal distribution D_x satisfies the conditions of Assumption E.10. Then there exists an algorithm that with probability $9/10$, given $O(d \log(d/\varepsilon)/\varepsilon)$ ε -noisy samples from D_{xy} , returns a \hat{w} such that for any w^* ,*

$$\mathbb{E}_{(X, Y) \sim D_{xy}} [L(\hat{w}, (X, Y))] \leq \mathbb{E}_{(X, Y) \sim D_{xy}} [L(w^*, (X, Y))] + O(\varepsilon^{1/4} \sqrt{\log(1/\varepsilon)}).$$

The approach is very similar to that of Theorem E.8, which we repeat here for clarity. First, we will restrict our search over w to \mathcal{H} , a ball of radius $r = \varepsilon^{-1/4} \sqrt{\log(1/\varepsilon)}$. As we argue in Lemma E.12, this restriction comes at a cost of at most $O(\varepsilon^{1/4} \sqrt{\log(1/\varepsilon)})$ in our algorithm's loss. With this restriction, we will argue that the problem satisfies the conditions of Proposition C.3. This allows us to argue that, with a polynomial number of samples, we can obtain a set of f_i 's satisfying the conditions of Assumption C.1. This will allow us to apply Theorem C.2, concluding the proof.

We start by showing that, due to anticoncentration of D , there is a $w' \in \mathcal{H}$ with loss close to w^* :

Lemma E.12. *Let w' be a rescaling of w^* , such that $\|w'\|_2 \leq \varepsilon^{-1/4} \sqrt{\ln(1/\varepsilon)}$ (i.e. $w' = \min\{1, \varepsilon^{-1/4} \sqrt{\ln(1/\varepsilon)} / \|w^*\|_2\} w^*$). Then $\mathbb{E}_{(X, Y) \sim D_{xy}} [L(w', (X, Y))] \leq \mathbb{E}_{(X, Y) \sim D_{xy}} [L(w^*, (X, Y))] + O(\varepsilon^{1/4} \sqrt{\ln(1/\varepsilon)})$.*

Proof. We need the following claim:

Claim E.13.

$$|t| \leq -\ln(\phi(t)\phi(-t)) \leq |t| + 3 \exp(-|t|)$$

Proof. Recalling that $\phi = 1/(1 + \exp(-t))$, we have that $-\ln(\phi(t)\phi(-t)) = \ln(\exp(t) + \exp(-t) + 2)$. Since $\exp(t) + \exp(-t) + 2 \geq \exp(|t|)$, we have $|t| \leq -\ln(\phi(t)\phi(-t))$. On the other hand, $\ln(\exp(t) + \exp(-t) + 2) = |t| + \ln(1 + 2 \exp(-|t|) + \exp(-2|t|)) \leq |t| + \ln(1 + 3 \exp(-|t|)) \leq |t| + 3 \exp(-|t|)$. \square

For any $x \in \mathbb{R}^d$, we have that:

$$\begin{aligned} -\ln(\phi(w' \cdot x)\phi(-w' \cdot x)) - y(w' \cdot x) - 3 \exp(-3|w' \cdot x|) &\leq |w' \cdot x| - y(w' \cdot x) \\ &\leq |w^* \cdot x| - y(w^* \cdot x) \\ &\leq -\ln(\phi(w^* \cdot x)\phi(-w^* \cdot x)) - y(w^* \cdot x) \end{aligned}$$

The first and last inequality hold by Claim E.13. For the second inequality, we do a case analysis on y . When $y = \text{sign}(w' \cdot x) = \text{sign}(w^* \cdot x)$, then both sides of the inequality are 0. When $y = -\text{sign}(w' \cdot x) = -\text{sign}(w^* \cdot x)$, then the inequality becomes $2|w' \cdot x| \leq 2|w^* \cdot x|$, which holds since $\|w'\|_2 \leq \|w^*\|_2$. We thus have that for any $y \in \{\pm 1\}$, $L(w', (x, y)) \leq L(w^*, (x, y)) + \frac{3}{2} \exp(-3|w' \cdot x|)$. If $|w' \cdot x| \leq \frac{1}{3} \ln(1/\varepsilon)$, then $L(w', (x, y)) \leq L(w^*, (x, y)) + \frac{3}{2}$. If $|w' \cdot x| \geq \frac{1}{3} \ln(1/\varepsilon)$, then $L(w', (x, y)) \leq L(w^*, (x, y)) + \frac{3}{2} \varepsilon$. Since $\|w'\|_2 \leq \varepsilon^{-1/4} \sqrt{\ln(1/\varepsilon)}$ and D_x is $\varepsilon^{1/4} \sqrt{\ln(1/\varepsilon)}$ -anticoncentrated, we have that $\Pr_{D_x} [|w' \cdot x| \leq \frac{1}{3} \ln(1/\varepsilon)] \leq O(\varepsilon^{1/4} \sqrt{\ln(1/\varepsilon)})$. Thus, $\mathbb{E}_{(X, Y) \sim D_{xy}} [L(w', (X, Y))] \leq \mathbb{E}_{(X, Y) \sim D_{xy}} [L(w^*, (X, Y))] + O(\varepsilon^{1/4} \sqrt{\ln(1/\varepsilon)})$, as desired. \square

With this in hand, we can conclude with the proof of Theorem E.11.

Proof of Theorem E.11. We first show that this problem fits into the GLM framework, in particular, satisfying the conditions of Proposition C.3. The link function is $\sigma_y(t) = \frac{1}{2}(-\ln(\phi(t)\phi(-t)) - yt)$, giving us the loss function $L(w, (x, y)) = \sigma_y(w \cdot x)$. We let \mathcal{H} be the set $\|w\|_2 \leq \varepsilon^{-1/4} \sqrt{\ln(1/\varepsilon)}$, giving us the parameter $r = \varepsilon^{-1/4} \sqrt{\ln(1/\varepsilon)}$. Condition 1 is satisfied by Assumption E.10. For $y \in \{-1, 1\}$, $\sigma'_y(t) = \frac{1}{2}(\phi(t) - \phi(-t) - y)$, which gives that $|\sigma'_y(t)| \leq 1$ for all t and y , satisfying Condition 2. Finally, $\sigma_y(0) = \ln 2 < 1$ for all y , satisfying Condition 3. Thus we can apply Proposition C.3: if we take $O(d \log(dr/\varepsilon)/\varepsilon)$ ε -corrupted samples, then they satisfy Assumption C.1 on \mathcal{H} with $\sigma_0 = 2$, $\sigma_1 = 0$ and $\sigma_2 = 1 + \varepsilon^{-1/4} \sqrt{\ln(1/\varepsilon)}$, with probability 9/10.

Now we can apply the algorithm of Theorem C.2. Since the loss is convex, we get a vector \hat{w} with $\bar{f}(\hat{w}) - \bar{f}(w^{*'}) = O((\sigma_0 r + \sigma_1 r^2 + \sigma_2) \sqrt{\varepsilon}) = O((2\varepsilon^{-1/4} \sqrt{\ln(1/\varepsilon)} + \varepsilon^{-1/4} \sqrt{\ln(1/\varepsilon)}) \sqrt{\varepsilon}) = O(\varepsilon^{1/4} \sqrt{\ln(1/\varepsilon)})$ where $w^{*'}$ is the minimizer of \bar{f} on \mathcal{H} .

We thus have that $\bar{f}(\hat{w}) \leq \bar{f}(w^{*'}) + O(\varepsilon^{1/4} \sqrt{\ln(1/\varepsilon)}) \leq \bar{f}(w^*) + O(\varepsilon^{1/4} \sqrt{\ln(1/\varepsilon)}) \leq \bar{f}(w^*) + O(\varepsilon^{1/4} \sqrt{\ln(1/\varepsilon)})$. The second inequality follows because $w^{*'}$ is the minimizer of \bar{f} on \mathcal{H} , and the third inequality follows from Lemma E.12. \square

F Additional Experimental Results

In this section, we provide additional plots of our experimental results, comparing with all baselines considered.

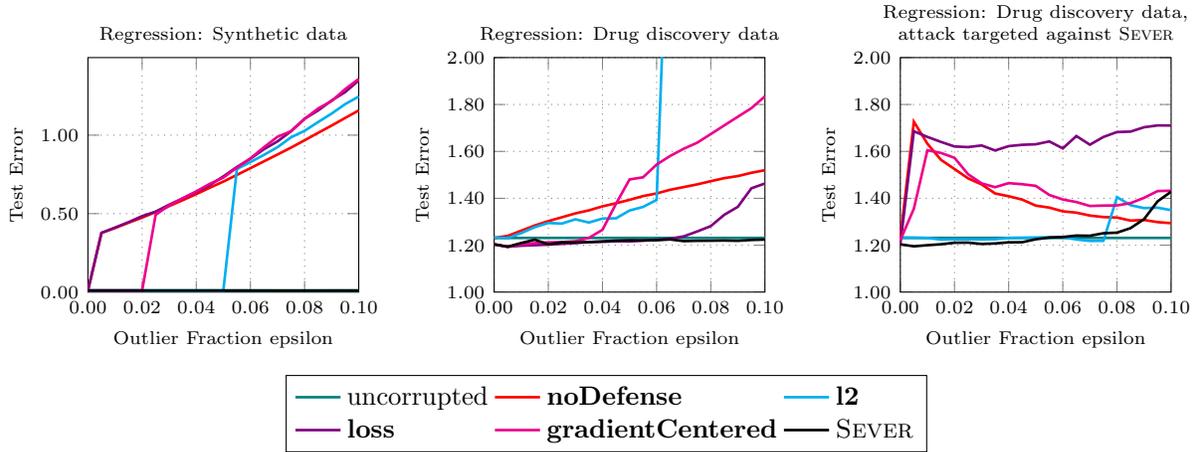


Figure 7: ε vs test error for baselines and SEVER on synthetic data and the drug discovery dataset. The left and middle figures show that SEVER continues to maintain statistical accuracy against our attacks which are able to defeat previous baselines. The right figure shows an attack with parameters chosen to increase the test error SEVER on the drug discovery dataset as much as possible. Despite this, SEVER still has relatively small test error.

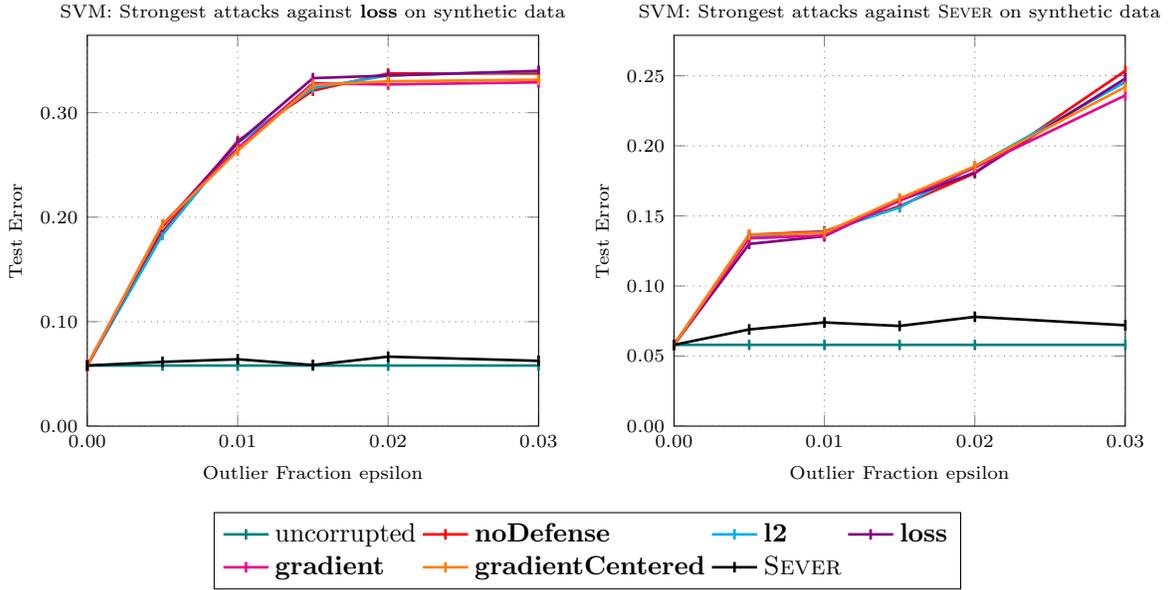


Figure 8: ϵ vs test error for baselines and SEVER on synthetic data. The left figure demonstrates that SEVER is accurate when outliers manage to defeat previous baselines. The right figure shows the result of attacks which increased the test error the most against SEVER. Even in this case, SEVER performs much better than the baselines.

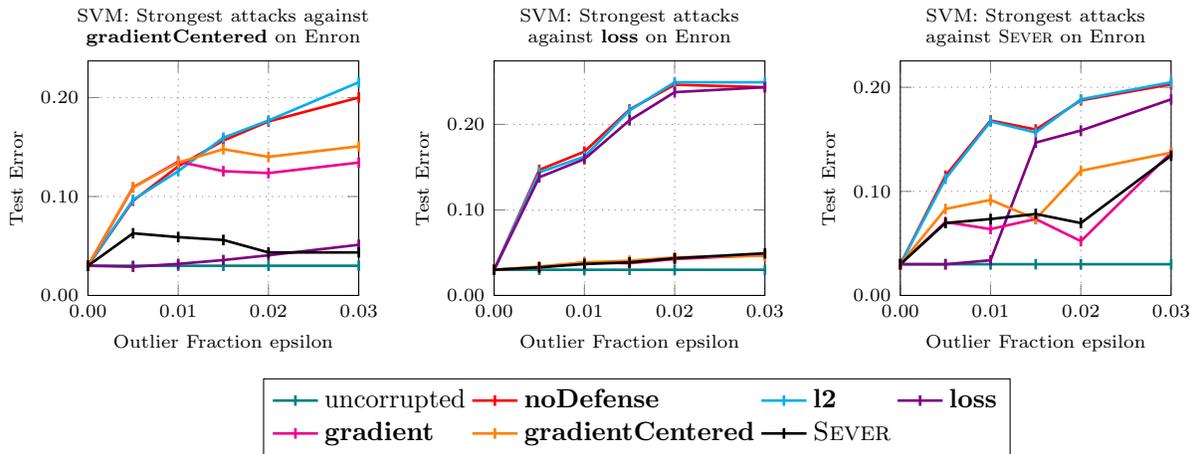


Figure 9: ϵ versus test error for baselines and SEVER on the Enron spam corpus. The left and middle figures are the attacks which perform best against two baselines, while the right figure performs best against SEVER. Though other baselines may perform well in certain cases, only SEVER is consistently accurate. The exception is for certain attacks at $\epsilon = 0.03$, which, as shown in Figure 6, require three rounds of outlier removal for any method to obtain ϵ reasonable test error – in these plots, our defenses perform only two rounds.