# Trimmed Maximum Likelihood Estimation for Robust Learning in Generalized Linear Models 

Pranjal Awasthi<br>Google Research<br>pranjalawasthi@google.com

Abhimanyu Das<br>Google Research<br>abhidas@google.com

Weihao Kong<br>Google Research<br>weihaokong@google.com

Rajat Sen<br>Google Research<br>senrajat@google.com


#### Abstract

We study the problem of learning generalized linear models under adversarial corruptions. We analyze a classical heuristic called the iterative trimmed maximum likelihood estimator which is known to be effective against label corruptions in practice. Under label corruptions, we prove that this simple estimator achieves minimax near-optimal risk on a wide range of generalized linear models, including Gaussian regression, Poisson regression and Binomial regression. Finally, we extend the estimator to the more challenging setting of label and covariate corruptions and demonstrate its robustness and optimality in that setting as well.


## 1 Introduction

Generalized linear models (GLMs) are an elegant framework for statistical modeling of data and are widely used in many applications [NW72, DB18, MN19]. A generalized linear model captures the relationship of labels (or observations) $y$ to covariates $\mathbf{x}$ via the conditional density $f\left(y \mid \beta^{\top} \mathbf{x}\right) \propto$ $\exp \left(y \cdot \beta^{\top} \mathbf{x}-b\left(\beta^{\top} \mathbf{x}\right)\right)$. Here $\beta$ is the parameter of the model. Parameter estimation in GLMs is done via the standard maximum likelihood estimation (MLE) paradigm and has been extensively studied [NW72]. While GLMs offer a mathematically tractable formulation for statistical modeling, real data rarely satisfies the generative process of a GLM and as a result there has been considerable interest in developing robust learning algorithms for GLMs [LW11, NRWY12, PSBR18]. The predominant way to model misspecification or adversarial corruptions in the data is Huber's $\epsilon$-contamination model [Hub11] and its recent extensions that allow for stronger adversaries [DKK ${ }^{+}$19]. These models assume that a small $\epsilon$ fraction of the data is corrupted by an adversary. Furthermore, the adversary can be restricted to either only corrupting the labels $y$, or be allowed to corrupt both the covariates $\mathbf{x}$ and the labels $y$ for an $\epsilon$ fraction of the data.

Various algorithms have been proposed for robust estimation in generalized linear models. One line of work proposed computationally intractable algorithms that are based on non-convex M estimators [Hub11, LW11] or by running tournaments over an exponentially large search space [Yat85]. Another line of work proposes polynomial time algorithms that either achieve sub-optimal error rates [PSBR18] or only apply to restricted settings such as the noise being heavy tailed [ZZ21]. In practical settings a simple heuristic namely the iterative trimmed estimator has been shown to work well under settings where only the labels are corrupted [SS19b]. However, from a theoretical perspective the iterative trimmed MLE estimator has only been analyzed under restrictive settings such as when the underlying GLM is a Gaussian regression model.

Our key theoretical contribution in this work is a general analysis of the trimmed MLE estimator. In particular, we show that for a broad family of GLMs, and under adversarial corruptions of only the labels, not only does the iterative trimmed MLE estimator enjoy theoretical guarantees, it in fact nearly achieves the minimax error rate! Next, we also consider the more challenging setting of corruptions to both covariates and labels. In the setting where the covariance of the covariate is known, we leverage the same approach in [PJL20] by running the filtering algorithm [DHL19] as a prepossessing step to trim away the abnormal covariates before applying the iterative trimmed MLE estimator. This can simultaneously handle both covariate and label corruptions and nearly achieve minimax error rates. Below we state our main results.

Theorem 1.1 (Informal Theorem). Let $S_{\epsilon}=\left\{\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{n}, y_{n}\right)\right\}$ be independent and identically distributed samples generated by a generalized linear model with sub-Gaussian $x_{i}$. Let an $\epsilon$ fraction of the labels be adversarially corrupted. Then, with high probability, the iterative trimmed MLE estimator when given as input $S_{\epsilon}$ provides the following guarantees.

- $O(\sigma \epsilon \log (1 / \epsilon))$ parameter estimation error for the Gaussian regression model where $\sigma^{2}$ is the variance of Gaussian noise on $y$.
- $O(\epsilon \exp (\sqrt{\log (1 / \epsilon)}))$ parameter estimation error for Poisson regression model.
- $O\left(\frac{1}{\sqrt{m}} \epsilon \sqrt{\log (m / \varepsilon) \log (1 / \varepsilon)}\right)$ parameter estimation error for the Binomial regression model with $m$ trials.
- $O(\epsilon \log (1 / \epsilon))$ parameter estimation error for a general class of smooth and continuous GLMs (includes the Gaussian regression model).

Previous work [SS19b] analyze the iterative trimmed maximum likelihood estimator under the Gaussian regression settings (least square) where only the labels are corrupted, and proved an $O(\sigma)$ $\ell_{2}$-error bound for parameter estimation. Other iterative trimming approach such as [BJK15] also achieves only achieve $O(\sigma)$ error in this setting. Our error bound of $O(\sigma \varepsilon \log (1 / \varepsilon))$ significantly improved the dependency on the corruption level $\varepsilon$ and nearly matched the minimax lower bound of $\sigma \varepsilon$ [Gao20]. In all the generalized linear models studied in this work, we show that the iterative trimmed maximum likelihood estimator achieves $O\left(\varepsilon^{1-\delta}\right)$ error for any $\delta>0$, which matches the minimax lower bound $\Omega(\varepsilon)$ up to a sub-polynomial factor.

Next, we present our second main result that can simultaneously handle both covariate and label corruptions and nearly achieve minimax error rate.

Theorem 1.2 (Informal Theorem). Let $S_{\epsilon}=\left\{\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{n}, y_{n}\right)\right\}$ be independent and identically distributed samples generated by a generalized linear model with sub-Gaussian $x_{i}$ whose covariance is known. Let an $\epsilon$ fraction of the labels and covariates be adversarially corrupted. After a preprocessing step, the iterative trimmed MLE with high probability achieves the same parameter estimation recovery bounds as in Theorem 1.1 above.

Our algorithm requires the covariance matrix of the covariate distribution being identity or known. Thus the error rate we get is incomparable to the results in the general covariance settings.

Outline of the paper. In Section 2 we discuss related work. We define preliminaries in Section 3 followed by the iterative trimmed MLE algorithm, formal results, and proof sketches for the label corruption case in Section 4 In Section 55, we introduce our algorithm, results and proof sketches for the sample corruption case. We defer proofs to the Appendix.

## 2 Related Work

There is a vast amount of literature in statistics, machine learning and theoretical computer science on algorithms that are robust to adversarial corruptions and outliers. Classical works in the Huber's contamination model present algorithms for general robust estimation that obtain near optimal error rates [Hub11, Yat85]. Minimax optimal but computationally inefficient robust estimators have been established in the works of [Tuk75, Yat85, CGR15, Gao20] for a variety of problems such as mean and covariance estimation. In recent years there has also been a line of work in designing computationally efficient algorithms for handling adversarial corruptions [LRV16, DKK ${ }^{+}$19, CSV17, BJK15, CKMY22, $\mathrm{CAT}^{+}$20, DKSS21].

Several special cases of generalized linear model has been studied extensively from the robustness perspective. There is a long line of work on designing robust algorithms for the linear least squares problem. This corresponds to the special case of the GLM being a Gaussian regression model [SBRJ19, BJK15, BJKK17, KKM18, CCM13, BP21, PJL20, DKS19]. When the covariate follows from a sub-Gaussian distribution, previous best result PJL20 achieves $\ell_{2}$ error $\sigma \varepsilon \sqrt{\log (1 / \varepsilon)}$ using Huber regression while we show iterative thresholding achieves $\sigma \varepsilon \log (1 / \varepsilon)$ error. Another special case of GLMs that has been studied from a robustness perspective is the logistic regression model [FXMY14, PSBR18, CKMY20]. Our analysis of the iterative trimmed MLE estimator on binomial regression matches the best known error guarantees for this setting.
The work of PSBR18] proposes a general procedure for robust gradient descent and shows that one can use this to design robust estimation algorithms for a abroad class of GLMs. Building upon these works, the authors in [JLST21] present a nearly linear time algorithm for GLMs. Compared to our work, [PSBR18] assumes $x$ has bounded 8th moment, [JLST21] assumes $x$ is 2-4 hypercontractive, and both papers achieve an $O(\sqrt{\epsilon})$ error guarantee with $\varepsilon$-fraction of corruptions while our algorithm achieves a better $O(\varepsilon)$ guarantee under the stronger sub-Gaussian assumptions. In addition, for GLMs, both papers assume a uniform upper bound (and lower bound) on the second order derivative of function $b(\cdot)$, which is not satisfied by the widely used Poisson and Binomial regression studied in this paper. In a similar setting, [ZZ21] proposed a reweighted MLE estimator for dealing with settings where the covariates are heavy tailed but not corrupted by an adversary.

Iterative thresholding is a longstanding heuristic for robust linear regression that dates back to Legendre [LS59]. It's theoretical property in the non-asymptotic regime is first studied in [BJK15], which shows a $O(\sigma)$ error bound for sufficiently small label corruption level when the label noise is $N\left(0, \sigma^{2}\right)$. The iterative thresholding algorithm is later extended and analyzed in the oblivious label corruption setting [BJKK17, SBRJ19], and is shown to provide consistent estimate even when the corruption level goes to 1 . PJL20] extended the iterative thresholding algorithm to the heavy-tailed covariate setting and can simultaneously handle both labels and covariate corruptions. In addition, [PJL20] adapted the implicit result in [BJKK17] to show that iterative thresholding algorithm achieves $O(\sigma \sqrt{\varepsilon})$ error rate for sub-Gaussian covariate. [SS19b] empirically demonstrated the effectiveness of the trimmed loss estimator under labels corruptions, and also proved theoretical guarantees in a special class of GLMs, which, however, also has $O(\sigma)$ error when specialized to the Gaussian noise setting. [SS19a] studied the trimmed loss estimator in the mixed linear regression setting. Finally, [CKMY22] proposed an alternating minimization algorithm for the fixed design linear regression setting with Huber contamination on the labels, which is different from our strong contamination model (with adaptive replacement) for generalized linear model. Their algorithm incorporates a semidefinite programming in the set selection step, and achieves a near-optimal $\tilde{O}(\sigma \varepsilon)$ error. On a very high level, our proof follows from a similar framework as in [CKMY22].
[PJL20] first proposed a generic approach to modify an estimator which is robust against label corruption into one that is robust against simultaneous label and covariate corruptions by running a covariate filter algorithm [DKP20, DK19] as a preprocessing step. In the Gaussian regression setting, our Algorithm 2] is identical to Algorithm 3 of [PJL20] and the difference is in the improved error rate. In particular, Lemma 4.1 in [PJL20] implies a $O(\sigma \sqrt{\varepsilon})$ error bound while we proved a $O(\sigma \varepsilon \log (1 / \varepsilon))$ error bound.

## 3 Preliminaries

In this section, we formally introduce the robust generalized linear model studied in this paper. First we define the classical generalized linear model as follows.
Definition 3.1 (Generalized linear model). We say that $(\mathrm{x}, y)$ follows from a generalized linear model if there exist function $b(\cdot)$, and function $c(\cdot)$ such that the probability density function of $y$ equals

$$
f\left(y \mid \beta^{\top} \mathbf{x}\right)=c(y) \exp \left(y \cdot \beta^{\top} \mathbf{x}-b\left(\beta^{\top} \mathbf{x}\right)\right)
$$

The derivative $b^{\prime}(\cdot)$ is called mean function as $\mathbb{E}\left[y \mid \beta^{\top} \mathbf{x}\right]=b^{\prime}\left(\beta^{\top} \mathbf{x}\right)$. The second order derivative $b^{\prime \prime}(\cdot)$ is called variance function as $\operatorname{Var}\left[y \mid \beta^{\top} \mathbf{x}\right]=b^{\prime \prime}\left(\beta^{\top} \mathbf{x}\right)$

Here we provide three commonly used examples of generalized linear models studied in this paper.
Definition 3.2. Commonly used examples of generalized linear model

- Gaussian regression: $b(\theta)=\frac{1}{2} \theta^{2}, c(y)=\frac{\exp \left(-y^{2} / 2\right)}{\sqrt{2 \pi}}$
- Poisson regression: $b(\theta)=\exp (\theta), c(y)=\frac{1}{y!}$
- Binomial regression: Let $m$ be the number of trials of the binomial distribution. $b(\theta)=$ $m \log (1+\exp (\theta)), c(y)=\binom{m}{y}$

We consider the sub-Gaussian random design setting in this work, i.e., each covariate $\mathbf{x}_{i}$ is drawn i.i.d. from a sub-Gaussian distribution with zero mean and covariance $\Sigma$.

Definition 3.3 (Sub-Gaussian design). We assume that each covariate $\mathbf{x}_{i}$ is drawn independently from a zero mean, covariance $\Sigma$ sub-Gaussian distribution with sub-Guassian norm 1, namely, $\mathbb{E}\left[\mathbf{x}_{i}\right]=0, \mathbb{E}\left[\mathbf{x}_{i} \mathbf{x}_{i}^{\top}\right]=\Sigma$ and for all $\mathbf{v} \in \mathbb{R}^{d}$,

$$
\operatorname{Pr}\left(\mathbf{v}^{\top} \mathbf{x}_{i} \geq t\right) \leq \exp \left(-t^{2}\right)
$$

Having introduced the generation model of the good data, now we describe the corruption model where the adversary is allowed to corrupt a small fraction of the data points.
Definition 3.4 (Corruption model). We consider two different data generation model with $\varepsilon$ fraction of adversarial corruptions:

- Label corruption model: Given $n$ i.i.d. sample $\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}_{i=1}^{n}$ generated by a generalized linear model. The adversary is allowed to inspect the sample, and replace a total of $\varepsilon \cdot n$ labels $y_{i}$ with arbitrary values.
- Sample corruption model: Given $n$ i.i.d. sample $\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}_{i=1}^{n}$ generated by a generalized linear model. The adversary is allowed to inspect the sample, and replace a total of $\varepsilon \cdot n$ data points $\left(\mathbf{x}_{i}, y_{i}\right)$ with arbitrary values.

We call the corrupted dataset $S=T \cup E$ where $T$ contains the set of remaining uncorrupted data points, and $E$ contains the set of data points that is controlled by the adversary.

The goal of our algorithm is recovering the underlying regression coefficient $\beta^{*}$ from a set of examples with $\varepsilon$ fraction corrupted under $\ell_{2}$ error metric, i.e., $\left\|\hat{\beta}-\beta^{*}\right\|$. We assume $\left\|\beta^{*}\right\| \leq R$ for a constant $R$, and $\varepsilon<=c$ for a sufficiently small constant $c$. For simplicity of the presentation, throughout this paper, we assume $\Sigma=I_{d}$. We remark that our algorithm for label corruption model applies to the general covariance setting and is able to achieve small estimation error in terms of $\left\|\Sigma^{1 / 2}\left(\hat{\beta}-\beta^{*}\right)\right\|_{2}$ since we can always (implicitly) whiten the data to apply our analysis (see Section Cor more details). On the other hand, the algorithm for sample corruption model only works with the knowledge of $\Sigma$.

## 4 Label Corruption

In this section, we formally describe our algorithm and proof-sketch for the label corruption setting.

### 4.1 Algorithm

We start with defining the trimmed maximum likelihood estimator, which is a simple and natural heuristic for robustly learning generalized linear model.
Definition 4.1 (Trimmed maximum likelihood estimator). Given a set of data points $S=$ $\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}_{i=1}^{n}$, define the trimmed maximum likelihood estimator as

$$
\hat{\beta}(S)=\min _{\beta} \min _{\hat{S} \subset S,|\hat{S}|=(1-\varepsilon) n} \sum_{\left(\mathbf{x}_{i}, y_{i}\right) \in \hat{S}}-\log f\left(y_{i} \mid \beta^{\top} \mathbf{x}_{i}\right)
$$

In the setting of generalized linear model, the objective of trimmed maximum likelihood estimator is a biconvex problem in $\hat{S}$ and $\beta$ but not jointly convex. The following alternating minimization algorithm is a simple heuristic to approximate the trimmed maximum likelihood estimator whose similar form has been studied in [BJK15, BJKK17, SS19b, CKMY22].

```
Algorithm 1: Alternating minimization of trimmed maximum likelihood estimator
Input: Set of examples \(S=\left\{\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{n}, y_{n}\right)\right\}, \varepsilon, \eta, R\)
Output: \(\hat{\beta}\)
\(S^{(0)} \leftarrow \arg \min _{T \subset[n]:|T|=(1-\varepsilon) n} \sum_{i \in T}\left|y_{i}\right| ;\)
\(\hat{\beta}^{(1)} \leftarrow 0 ;\)
for \(t=1\) to \(\infty d o\) do
    Choose \(\hat{S}^{(t)}=\arg \min _{T \subset S^{(0)}:|T|=(1-2 \varepsilon) n} \sum_{i \in T}-\log f\left(y_{i} \mid\left\langle\hat{\beta}^{(t)}, \mathbf{x}_{i}\right\rangle\right)\);
    Compute \(\hat{\beta}^{(t+1)}=\arg \min _{\beta,\|\beta\| \leq R} \sum_{i \in \hat{S}^{(t)}}-\log f\left(y_{i} \mid\left\langle\beta, \mathbf{x}_{i}\right\rangle\right)\);
    if \(\frac{1}{n} \sum_{i \in \hat{S}^{(t)}}-\log f\left(y_{i} \mid\left\langle\hat{\beta}^{(t+1)}, \mathbf{x}_{i}\right\rangle\right)>\frac{1}{n} \sum_{i \in \hat{S}^{(t)}}-\log f\left(y_{i} \mid\left\langle\hat{\beta}^{(t)}, \mathbf{x}_{i}\right\rangle\right)-\eta\) then
        Return \(\beta^{(t)}\)
```

The algorithm starts by naively pruning out $\varepsilon \cdot n$ data points whose labels have the largest magnitude. Then each round of the alternating minimization algorithm has two steps. In optimizing over set $S$, we find the set $\hat{S}^{(t)}$ of size $(1-\varepsilon) n$ with the best likelihood. In optimizing over $\beta$, we find regression coefficient $\beta^{(t)}$ which maximizes the likelihood on the current set of data $\hat{S}^{(t)}$. The algorithm terminates and outputs $\beta$ when the likelihood no longer improves by more than $\eta$. It is clear that the algorithm terminates in $O(1 / \eta)$ rounds when the log-likelihood is bounded. Since the original trimmed maximum likelihood estimator is a biconvex optimization problem in $S, \beta$ which is not jointly convex, our algorithm does not guarantee to return a global optimal solution. Nonetheless, as we showed, it does return a first order stationary point which will be close to the true coefficient $\beta^{*}$. It is worth noting that some recent papers on robust statistics [CDK ${ }^{+}$21, CDGS20, ZJS22] show similar nice statistical properties of an approximate first order stationary point for non-convex optimization problems.
We present the guarantee of our algorithm for Gaussian, Poisson, Binomial regression, and a broad class of generalized linear models.
Theorem 4.2 (Gaussian regression with label corruption). Let $S=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ be a set of data points generated by a Gaussian regression model with $y_{i}=\left\langle\mathbf{x}_{i}, \beta^{*}\right\rangle+\eta_{i}, \eta_{i} \sim N\left(0, \sigma^{2}\right)$, subGaussian design, with $\varepsilon_{c}$-fraction of label corruption and $n=\Omega\left(\frac{d+\log (1 / \delta)}{\varepsilon^{2}}\right)$. With probability $1-\delta$, Algorithm $\left\lceil 1\right.$ with parameters $\varepsilon=\varepsilon_{c}, \eta=\varepsilon_{c}^{2}, R=\infty$ terminate within $O\left(\frac{1}{\min \left(1, \sigma^{2}\right) \varepsilon_{c}^{2}}\right)$ iterations, and output an estimate $\hat{\beta}$ such that

$$
\left\|\hat{\beta}-\beta^{*}\right\|=O\left(\sigma \varepsilon_{c} \log \left(1 / \varepsilon_{c}\right)\right)
$$

Theorem 4.3 (Poisson regression with label corruption). Let $S=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ be a set of data points generated by a Poisson regression model with sub-Gaussian design, with $\varepsilon_{c}$-fraction of label corruption and $n=\Omega\left(\frac{d}{\varepsilon^{2}}\right)$. With probability 0.99, Algorithm 1$]$ with parameters $\varepsilon=2 \varepsilon_{c}, \eta=$ $\varepsilon_{c}^{2} /(d n)$, contant $R \geq\left\|\beta^{*}\right\|$ terminate within $d n / \varepsilon_{c}^{2}$ iterations, and output an estimate $\hat{\beta}$ such that

$$
\left\|\hat{\beta}-\beta^{*}\right\|=O\left(\varepsilon_{c} \exp \left(\sqrt{\log \left(1 / \varepsilon_{c}\right)}\right)\right)
$$

Theorem 4.4 (Binomial regression with label corruption). Let $S=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ be generated by a Binomial regression model with sub-Gaussian Design, with $\varepsilon_{c}$-fraction of label corruption and $n=\Omega\left(\frac{d+\log (1 / \delta)}{\varepsilon^{2}}\right)$. With probability $1-\delta$, Algorithm 1 with parameters $\varepsilon=\varepsilon_{c}, \eta=\varepsilon_{c}^{2} / m$, constant $R \geq\left\|\beta^{*}\right\|$ terminate within $m^{2} / \varepsilon_{c}^{2}$ iterations, and output an estimate $\hat{\beta}$ such that

$$
\left\|\hat{\beta}-\beta^{*}\right\|=O\left(\varepsilon_{c} \sqrt{\frac{\log \left(m / \varepsilon_{c}\right) \log \left(1 / \varepsilon_{c}\right)}{m}}\right)
$$

Theorem 4.5 (A class of generalized linear model with label corruption). Let $S=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ be generated by a generalized linear model with sub-Gaussian Design, with $\varepsilon_{c}$-fraction of label corruption and $n=\Omega\left(\frac{d+\log (1 / \delta)}{\varepsilon^{2}}\right)$. Assuming that $C_{0} \leq b^{\prime \prime}(\cdot) \leq C$ for non-zero constants $C_{0}, C$, $b(0)=0, b^{\prime}(0)=0$, and $\log (c(y))=O\left(\log \left(1 / \varepsilon_{c}\right)\right), \forall y \leq \Theta\left(\sqrt{\log \left(1 / \varepsilon_{c}\right)}\right)$, then with probability
$1-\delta$, Algorithm $\left[1\right.$ with parameters $\varepsilon=\varepsilon_{c}, \eta=\varepsilon_{c}^{2}, R=\infty$ terminates within $\log \left(1 / \varepsilon_{c}\right) / \varepsilon_{c}^{2}$ iterations, and output an estimate $\hat{\beta}$ such that

$$
\left\|\hat{\beta}-\beta^{*}\right\|=O\left(\varepsilon_{c} \log \left(1 / \varepsilon_{c}\right)\right)
$$

### 4.2 Proof Sketch

We provide an intuitive proof sketch for the above theorems (the full proofs are deferred to Section A in the Appendix). The high level proof framework is similar to [CKMY22], although the details are drastically different since the focus of our paper is on random design with strong contamination for a wide range of generalized linear model while [CKMY22] focuses on linear (least square) regression with Huber corruption on the labels.

The guarantee of our alternating minimization algorithm relies on two claims: First, the algorithm returns an approximate stationary point $\hat{\beta}$. Second, any approximate stationary point will be close to the true coefficient $\beta^{*}$. In this section, we will present high level intuition of the proof of the two claims.

Alternating minimization algorithm returns an approximate stationary point. First we define the first order approximate stationary point as follows. Let $\hat{\beta} \in \mathbb{R}^{d}$ be a regression coefficient vector and $\hat{S}$ contains the set of datapoints of size $(1-\varepsilon) n$ with the largest log-likelihood under $\hat{\beta}$. We call $\hat{\beta}$ a $\gamma$-approximate stationary point if

$$
\frac{1}{n} \sum_{i \in \hat{S}} \nabla_{\beta} \log f\left(y_{i} \mid\left\langle\hat{\beta}, \mathbf{x}_{i}\right\rangle\right)^{\top} \frac{\left(\beta^{*}-\hat{\beta}\right)}{\left\|\beta^{*}-\hat{\beta}\right\|} \leq \gamma
$$

i.e., the gradient of the log-likelihood projected along the $\beta^{*}-\hat{\beta}$ direction is small. Our goal is to show when the algorithm terminates, that is when $\hat{\beta}$ can not be improved by more than $\eta$, the gradient along the $\beta^{*}-\hat{\beta}$ must be small. This is clear where the empirical log-likelihood function is smooth, simply because if the gradient is large, one can improve the log-likelihood by more than $\eta$ which will result in a contradiction. The smoothness (norm of the Hessian matrix) of the empirical log-likelihood in generalized linear model is directly related to the range of $b^{\prime \prime}(\theta)$. In particular, $b^{\prime \prime}(\theta)$ is bounded for Gaussian regression and Binomial regression.

However, a problem arises for Poisson regression where $b^{\prime \prime}(\theta)=\exp (\theta)$ becomes extremely large for large $\theta$, which results in non-smooth curvature for empirical log-likelihood. We overcome this difficulty by leveraging the special property of function $b(\theta)=\exp (\theta)$ in the Poisson regression setting. Observing that the derivative $b^{\prime}(\theta)$ is equal to second order derivative $b^{\prime \prime}(\theta)$ for Poisson regression, the gradient along the $\beta^{*}-\hat{\beta}$ direction can not be small when the second order derivative along the $\beta^{*}-\hat{\beta}$ direction gets large, which will result in a more than $\eta$ improvement of the loglikelihood by moving toward $\beta^{*}$ and therefore a contradiction. Hence, the second order derivative along the $\beta^{*}-\hat{\beta}$ direction must be small, and we blue have the same argument as in the smooth objective function setting.
Any approximate stationary point will be close to the true coefficient. Let us first write down the $\gamma$-approximate stationary condition for generalized linear model as

$$
\frac{1}{n} \sum_{i \in \hat{S}} \nabla_{\beta} \log f\left(y_{i} \mid\left\langle\hat{\beta}, \mathbf{x}_{i}\right\rangle\right)^{\top}\left(\beta^{*}-\hat{\beta}\right)=\frac{1}{n} \sum_{i \in \hat{S}}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \leq \gamma\left\|\beta^{*}-\hat{\beta}\right\|
$$

Recall that $T$ contains the set of uncorrupted data points, and $E$ contains the set of data points that is controlled by the adversary. Split $\hat{S}$ into $\hat{S} \cap T$ and $\hat{S} \cap E$, and rearrange the terms we get

$$
\frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \leq-\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}+\gamma\left\|\beta^{*}-\hat{\beta}\right\| .
$$

To obtain an upper bound on $\left\|\beta^{*}-\hat{\beta}\right\|$, we will prove a lower bound in terms of $\left\|\beta^{*}-\hat{\beta}\right\|$ on the left hand side, and an upper bound in terms of $\left\|\beta^{*}-\hat{\beta}\right\|$ on the right hand side. Finally we will combine the upper and lower bound into an upper bound on $\left\|\beta^{*}-\hat{\beta}\right\|$.

Lower bound on the LHS. Note that $\hat{S} \cap T$ contains uncorrupted data points. The high level intuition is that since mean of $y_{i}$ is $b^{\prime}\left(\beta^{* \top} \mathbf{x}_{i}\right)$ and $b^{\prime}(\cdot)$ is monotone, $\left(y_{i}-b^{\prime}\left(\beta^{* \top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}$ should be roughly $O\left(\left(\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)^{2}\right)$, and $\left(\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)^{2}$ should be proportional to $\left\|\beta^{*}-\hat{\beta}\right\|^{2}$ given enough samples. More formally, we will decompose the LHS as

$$
\begin{aligned}
& \frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right) \mathbf{x}_{i} \\
= & \frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-b^{\prime}\left(\beta^{* \top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}+\frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(b^{\prime}\left(\beta^{* \top} \mathbf{x}_{i}\right)-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} .
\end{aligned}
$$

The first term contains a $(1-\varepsilon)$ fraction of uncorrupted random examples sampled from a zero mean distribution with certain tail bound, e.g. sub-exponential for Gaussian and Binomial regression, $k$-th moment bound for Poisson regression.
Therefore, we can apply resilience property to bound the first term. Overall, we heavily utilize the resilience property of the sample set that is drawn from "nice" distributions. Take sample mean as an example, resilience [SCV17, ZJS19] (also known as stability [DK19]) dictates that given a large enough sample set $S=\left\{\mathbf{x}_{i}\right\}_{i=1}^{n}$, the sample mean of any large enough subset of $S$ will be close to each other. We define mean resilience formally here:
Definition 4.6 (Resilience). Given a sample set $S=\left\{\mathbf{x}_{i}\right\}_{i=1}^{n}$, suppose for any $T \subset S,|T| \geq(1-\varepsilon) n$, it holds that $\left\|\frac{1}{|T|} \sum_{i \in T} \mathbf{x}_{i}-\frac{1}{|S|} \sum_{i \in S} \mathbf{x}_{i}\right\| \leq \tau$, then we call the set $S$ satisifes $(\varepsilon, \tau)$-resilience.

Specifically, under sub-Gaussian distribution, a set $S$ of i.i.d. samples with size $n=\Omega\left(d / \varepsilon^{2}\right)$ satisfies $(\varepsilon, \varepsilon \sqrt{\log (1 / \varepsilon)})$ resilience with high probability. Resilience property applies to sub-exponential and $k$-th moment bounded distribution as well, and this gives us a way to control the behavior of any subset of good data.
For the second term, we prove that $\frac{1}{n} \sum_{i \in \hat{S} \cap T} b\left(\beta^{\top} \mathbf{x}_{i}\right)$ is a strongly convex function again using the resilience property, which implies $\frac{1}{n} \sum_{i \in S \hat{S} \cap T}\left(b^{\prime}\left(\beta^{*}{ }^{\top} \mathbf{x}_{i}\right)-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}=\Omega\left(\| \beta^{*}-\right.$ $\hat{\beta} \|^{2}$ ).
Upper bound on the RHS. To upper bound $-\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}$, we will prove an upper bound on $\sqrt{\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)^{2}}$ and $\sqrt{\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)^{2}}$ separately, then apply Cauchy-Schwarz inequality. The key difficulty is bounding $\sqrt{\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)^{2}}$, as it contains corrupted data points controlled by an adversary, which does not follow any good property possessed by the good stochastic data. However, since $\hat{S}$ contains $(1-\varepsilon) n$ datapoints with the largest log-likelihood under $\hat{\beta}$, we can argue that

$$
\sum_{i \in \hat{S} \cap E}-\log f\left(y_{i} \mid\left\langle\hat{\beta}, \mathbf{x}_{i}\right\rangle\right) \leq \sum_{i \in T \backslash \hat{S}}-\log f\left(y_{i} \mid\left\langle\hat{\beta}, \mathbf{x}_{i}\right\rangle\right)
$$

or even

$$
\max _{i \in \hat{S} \cap E}-\log f\left(y_{i} \mid\left\langle\hat{\beta}, \mathbf{x}_{i}\right\rangle\right) \leq \min _{i \in T \backslash \hat{S}}-\log f\left(y_{i} \mid\left\langle\hat{\beta}, \mathbf{x}_{i}\right\rangle\right)
$$

since otherwise one can replace the data points in $\hat{S} \cap E$ by the ones in $T \backslash \hat{S}$ to form a new set with better likelihood than $\hat{S}$. This gives us an upper bound on the negative log-likelihood of $y_{i}$ in $\hat{S} \cap E$. Therefore we adopt a two step approach to upper bound the $\sqrt{\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)^{2}}$. First we prove an upper bound on the negative log-likelihood on $T \backslash \hat{S}$, which becomes a negative log-likelihood bound on $\hat{S} \cap E$ immediately. Second we turn the negative log-likelihood bound into a square error bound.
The two steps vary drastically for different regression models. For the first step of upper bounding the negative log-likelihood, in Gaussian regression we use the resilience property of the quadratic
form of sub-Gaussian random variable. In Poisson regression, we leverage the resilience property of $y_{i} \mathbf{x}_{i}$ which is heavy tailed. In Binomial regression, since the distribution only has support size $m$, we directly analyze the resilience of the negative log-likelihood. In general GLMs, due to the generality of the likelihood function, we have to again analyze the resilience of the negative log-likelihood directly. For the second step of turning the log-likelihood bound to a quadratic bound, we get the bound trivially in Gaussian regression since Gaussian likelihood is indeed quadratic. For Poisson and Binomial setting, we have to build a proxy function which lower bound the negative log-likelihood function $-\log f\left(y_{i} \mid\left\langle\hat{\beta}, \mathbf{x}_{i}\right\rangle\right)$ to connect it to quadratic function. For the general class of GLMs, we leverage the bounds on the $b^{\prime \prime}(\cdot)$ and $\log (c(y))$ to obtain a quadratic bound.

### 4.3 Proof for Poisson Regression

As an illustrative example, we show how the above proof sketch can be used for Poisson regression to formally prove Theorem 4.3 (Proofs for the other GLMs are deferred to Section A in the Appendix)
Lemma 4.7 (Approximate stationary point close to $\beta^{*}$ for Poisson regression). Given a set of datapoints $S=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ generated by a Poisson regression model with $\varepsilon$-fraction of label corruption, and the largest $\varepsilon$ labels removed. Let $\hat{\beta}$ be $a \max \left(\varepsilon, \varepsilon^{2} /\left\|\beta^{*}-\hat{\beta}\right\|\right)$-stationary point and $\|\hat{\beta}\| \leq R$. Given that $n=\Omega\left(\frac{d}{\varepsilon^{2}}\right)$, with probability 0.99 , it holds that

$$
\left\|\hat{\beta}-\beta^{*}\right\|=O(\varepsilon \exp (\Theta(\sqrt{\log (1 / \varepsilon)})))
$$

Since $b^{\prime \prime}(\theta)=\exp (\theta)$ is unbounded for Poisson regression, the following lemma (proved in the appendix) shows that alternating minimization algorithm still return an approximate stationary point
Lemma 4.8 (Algorithm 1 finds an approximate stationary point for Poisson regression). Given a set of datapoints $S=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ generated by a Poisson model with $\varepsilon_{c}$-fraction of corruption. Assuming that $n=\Omega\left(\frac{d+\log (1 / \delta)}{\varepsilon^{2}}\right)$, then with probability $1-\delta$, the output of Algorithm 1 with input parameters $\varepsilon=2 \varepsilon_{c}, R \geq\left\|\beta^{*}\right\|, \eta=\varepsilon^{2} /(d n)$, is $\max \left(\varepsilon, \frac{2 \varepsilon^{2}}{\left\|\beta^{*}-\hat{\beta}\right\|}\right)$-approximate stationary point.

Proof of Theorem 4.3. Lemma 4.8 implies the output of Algorithm 1 is a $\max \left(\varepsilon, \frac{\varepsilon^{2}}{\left\|\beta^{*}-\hat{\beta}\right\|}\right)$ approximate stationary point. Lemma 4.7 then implies that $\left\|\hat{\beta}-\beta^{*}\right\|=O(\varepsilon \exp (\sqrt{\log (1 / \varepsilon)}))$. To bound the number of iterations, we need an upper bound on the negative log-likelihood on $\beta=0$, and a uniform lower bound on the negative log-likelihood. The initial negative log-likelihood is upper bounded by $\frac{1}{n} \sum_{i \in \hat{S}^{(1)}} \log \left(y_{i}!\right)+1 \leq O\left(\mathbb{E}\left[y_{i}^{2}\right]+1\right)=O(1)$ where $S^{(1)}$ contains the smallest $(1-\varepsilon) n$ labels. Trivially, there is a 0 lower bound on the negative log-likelihood for Poisson distribution. Therefore, the algorithm will terminate in $d n / \varepsilon_{c}^{2}$ iterations.

## 5 Result for Sample Corruption Model

The learning problem becomes much harder in the presence of label and covariate corruption, since it is hard to tell whether a data point is corrupted by simply looking at the likelihood of $y_{i}$. From a technical level, the resilience condition we leveraged on covariate $\mathbf{x}_{i}$ in set $E$ breaks down when there is covariate corruption. Luckily, we are able to restore the resilience property by first running the filtering algorithm for robust mean estimation [DHL19]. Specifically, if the covariate distribution has identity (or known) covariance and sub-Gaussian tail, one can apply the filtering algorithm to the $\varepsilon$ corrupted data set, and the resulting data set $\left\{w_{i} \mathbf{x}_{i}\right\}_{i=1}^{n}$ will have close to identity covariance and the same resilient condition as an uncorrupted data set. This prepossessing step only takes nearly linear time. This approach is firstly proposed in [PJL20] as a general method to make an algorithm robust against covariate-corruptions.

### 5.1 Algorithm

The guarantee of Algorithm 2 is formalized in the following theorems.
Theorem 5.1 (Gaussian regression with sample corruption). Given a set of datapoints $S=$ $\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ generated by a Gaussian regression model with $y_{i}=\left\langle\mathbf{x}_{i}, \beta^{*}\right\rangle+\eta_{i}, \eta_{i} \sim N\left(0, \sigma^{2}\right)$, sub-Gaussian design, $\varepsilon_{c}$-fraction of sample corruption and $n=\Omega\left(\frac{d+\log (1 / \delta)}{\varepsilon^{2}}\right)$. With probability

```
Algorithm 2: Alternating minimization of trimmed maximum likelihood estimator in sample
corruption model
Input: Set of examples \(S=\left\{\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{n}, y_{n}\right)\right\}, \Sigma, \varepsilon, \eta, R\)
Output: \(\hat{\beta}\)
\(S_{0} \leftarrow\left\{\left(\Sigma^{-1 / 2} \mathbf{x}_{1}, y_{1}\right), \ldots,\left(\Sigma^{-1 / 2} \mathbf{x}_{n}, y_{n}\right)\right\} / /\) Whiten the covariates.
\(S^{\prime} \leftarrow \operatorname{Filtering}(S, \varepsilon) / /\) Algorithm 4 in [DHL19]
\(\hat{\beta} \leftarrow\) Algorithm \(1\left(S^{\prime}, \varepsilon, \eta, R\right) ;\)
Return \(\hat{\beta}\);
```

$1-\delta$, Algorithm 2 with parameters $\varepsilon=\varepsilon_{c}, \eta=\varepsilon_{c}^{2}, R=\infty$ terminate within $O\left(\frac{1}{\min \left(1, \sigma^{2}\right) \varepsilon_{c}^{2}}\right)$ iterations, and output an estimate $\hat{\beta}$ such that

$$
\left\|\hat{\beta}-\beta^{*}\right\|=O\left(\sigma \varepsilon_{c} \log \left(1 / \varepsilon_{c}\right)\right)
$$

Theorem 5.2 (Poisson regression with sample corruption). Given a set of datapoints $S=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ generated by a Poisson regression model with $\varepsilon_{c}$-fraction of label corruption and $n=\Omega\left(\frac{d}{\varepsilon^{2}}\right)$. With probability 0.99, Algorithm 2 with parameters $\varepsilon=2 \varepsilon_{c}, \eta=\varepsilon_{c}^{2} /(d n), R \geq\left\|\beta^{*}\right\|$ terminate within $d n / \varepsilon_{c}^{2}$ iterations, and output an estimate $\hat{\beta}$ such that

$$
\left\|\hat{\beta}-\beta^{*}\right\|=O\left(\varepsilon_{c} \exp \left(\sqrt{\log \left(1 / \varepsilon_{c}\right)}\right)\right)
$$

Theorem 5.3 (Binomial regression with sample corruption). Given a set of datapoints $S=$ $\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ generated by a Binomial regression model with $\varepsilon_{c}$-fraction of sample corruption and $n=\Omega\left(\frac{d+\log (1 / \delta)}{\varepsilon^{2}}\right)$. With probability $1-\delta$, Algorithm 2 with parameters $\varepsilon=\varepsilon_{c}, \eta=\varepsilon_{c}^{2}, R \geq\left\|\beta^{*}\right\|$ terminate within $m / \varepsilon_{c}^{2}$ iterations, and output an estimate $\hat{\beta}$ such that

$$
\left\|\hat{\beta}-\beta^{*}\right\|=O\left(\varepsilon_{c} \sqrt{\frac{\log \left(m / \varepsilon_{c}\right) \log \left(1 / \varepsilon_{c}\right)}{m}}\right)
$$

Theorem 5.4 (A class of generalized linear model with sample corruption). Let $S=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ be generated by a generalized linear model with sub-Gaussian Design, with $\varepsilon_{c}$-fraction of sample corruption and $n=\Omega\left(\frac{d+\log (1 / \delta)}{\varepsilon^{2}}\right)$. Assuming that $C_{0} \leq b^{\prime \prime}(\cdot) \leq C$ for non-zero constants $C_{0}, C$, $b(0)=0, b^{\prime}(0)=0$, and $\log (c(y))=O\left(\log \left(1 / \varepsilon_{c}\right)\right), \forall y \leq \Theta\left(\sqrt{\log \left(1 / \varepsilon_{c}\right)}\right)$ With probability $1-\delta$, Algorithm 2 with parameters $\varepsilon=\varepsilon_{c}, \eta=\varepsilon_{c}^{2}, R=\infty$ terminate within $\log \left(1 / \varepsilon_{c}\right) / \varepsilon_{c}^{2}$ iterations, and output an estimate $\hat{\beta}$ such that

$$
\left\|\hat{\beta}-\beta^{*}\right\|=O\left(\varepsilon_{c} \log \left(1 / \varepsilon_{c}\right)\right)
$$

The proof is the same compared to the label corruption setting except that since $\mathbf{x}_{i}, i \in E$ in now controlled by the adversary, we can no longer bound $\sqrt{\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)^{2}}$ by the resilience property of (uncorrupted) sub-Gaussian samples. Instead, we will leverage the fact that corrupted sample with small covariance is also resilient.

## 6 Conclusion

In this paper, we provided a general theoretical analysis showing that a simple and practical heuristic namely the iterative trimmed MLE estimator achieves minimax optimal error rates upto a logarithmic factor under adversarial corruptions for a wide class of generalized linear models (GLMs). It would also be interesting to study whether our techniques can be extended to design robust algorithms for more general exponential families beyond GLMs.

## References

[Ahl22] Thomas D Ahle. Sharp and simple bounds for the raw moments of the binomial and poisson distributions. Statistics \& Probability Letters, 182:109306, 2022.
[BJK15] Kush Bhatia, Prateek Jain, and Purushottam Kar. Robust regression via hard thresholding. Advances in neural information processing systems, 28, 2015.
[BJKK17] Kush Bhatia, Prateek Jain, Parameswaran Kamalaruban, and Purushottam Kar. Consistent robust regression. Advances in Neural Information Processing Systems, 30, 2017.
[BP21] Ainesh Bakshi and Adarsh Prasad. Robust linear regression: Optimal rates in polynomial time. In Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, pages 102-115, 2021.
$\left[\mathrm{CAT}^{+} 20\right]$ Yeshwanth Cherapanamjeri, Efe Aras, Nilesh Tripuraneni, Michael I Jordan, Nicolas Flammarion, and Peter L Bartlett. Optimal robust linear regression in nearly linear time. arXiv preprint arXiv:2007.08137, 2020.
[CCM13] Yudong Chen, Constantine Caramanis, and Shie Mannor. Robust sparse regression under adversarial corruption. In International Conference on Machine Learning, pages 774-782. PMLR, 2013.
[CDGS20] Yu Cheng, Ilias Diakonikolas, Rong Ge, and Mahdi Soltanolkotabi. High-dimensional robust mean estimation via gradient descent. In International Conference on Machine Learning, pages 1768-1778. PMLR, 2020.
[CDK $\left.{ }^{+} 21\right]$ Yu Cheng, Ilias Diakonikolas, Daniel M Kane, Rong Ge, Shivam Gupta, and Mahdi Soltanolkotabi. Outlier-robust sparse estimation via non-convex optimization. arXiv preprint arXiv:2109.11515, 2021.
[CGR15] Mengjie Chen, Chao Gao, and Zhao Ren. Robust covariance matrix estimation via matrix depth. arXiv preprint arXiv:1506.00691, 2015.
[CKMY20] Sitan Chen, Frederic Koehler, Ankur Moitra, and Morris Yau. Classification under misspecification: Halfspaces, generalized linear models, and evolvability. Advances in Neural Information Processing Systems, 33:8391-8403, 2020.
[CKMY22] Sitan Chen, Frederic Koehler, Ankur Moitra, and Morris Yau. Online and distributionfree robustness: Regression and contextual bandits with huber contamination. In 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS), pages 684-695. IEEE, 2022.
[CSV17] Moses Charikar, Jacob Steinhardt, and Gregory Valiant. Learning from untrusted data. In Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, pages 47-60, 2017.
[DB18] Annette J Dobson and Adrian G Barnett. An introduction to generalized linear models. Chapman and Hall/CRC, 2018.
[DHL19] Yihe Dong, Samuel Hopkins, and Jerry Li. Quantum entropy scoring for fast robust mean estimation and improved outlier detection. Advances in Neural Information Processing Systems, 32, 2019.
[DK19] Ilias Diakonikolas and Daniel M Kane. Recent advances in algorithmic high-dimensional robust statistics. arXiv preprint arXiv:1911.05911, 2019.
[DKK ${ }^{+}$19] Ilias Diakonikolas, Gautam Kamath, Daniel Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robust estimators in high-dimensions without the computational intractability. SIAM Journal on Computing, 48(2):742-864, 2019.
[DKP20] Ilias Diakonikolas, Daniel M Kane, and Ankit Pensia. Outlier robust mean estimation with subgaussian rates via stability. Advances in Neural Information Processing Systems, 33:1830-1840, 2020.
[DKS19] Ilias Diakonikolas, Weihao Kong, and Alistair Stewart. Efficient algorithms and lower bounds for robust linear regression. In Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 2745-2754. SIAM, 2019.
[DKSS21] Ilias Diakonikolas, Daniel M Kane, Alistair Stewart, and Yuxin Sun. Outlier-robust learning of ising models under dobrushin's condition. In Conference on Learning Theory, pages 1645-1682. PMLR, 2021.
[FXMY14] Jiashi Feng, Huan Xu, Shie Mannor, and Shuicheng Yan. Robust logistic regression and classification. Advances in neural information processing systems, 27, 2014.
[Gao20] Chao Gao. Robust regression via mutivariate regression depth. Bernoulli, 26(2):11391170, 2020.
[Hub11] Peter J Huber. Robust statistics. In International encyclopedia of statistical science, pages 1248-1251. Springer, 2011.
[JLST21] Arun Jambulapati, Jerry Li, Tselil Schramm, and Kevin Tian. Robust regression revisited: Acceleration and improved estimation rates. Advances in Neural Information Processing Systems, 34, 2021.
[JLT20] Arun Jambulapati, Jerry Li, and Kevin Tian. Robust sub-gaussian principal component analysis and width-independent schatten packing. Advances in Neural Information Processing Systems, 33:15689-15701, 2020.
[KKM18] Adam Klivans, Pravesh K Kothari, and Raghu Meka. Efficient algorithms for outlierrobust regression. In Conference On Learning Theory, pages 1420-1430. PMLR, 2018.
[LRV16] Kevin A Lai, Anup B Rao, and Santosh Vempala. Agnostic estimation of mean and covariance. In 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS), pages 665-674. IEEE, 2016.
[LS59] Adrien-Marie Legendre and DE Smith. On the method of least squares. A Source Book in Mathemathics, Ed. DE Smith (originally published in 1805), pages 576-579, 1959.
[LW11] Po-Ling Loh and Martin J Wainwright. High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity. Advances in Neural Information Processing Systems, 24, 2011.
[MN19] Peter McCullagh and John A Nelder. Generalized linear models. Routledge, 2019.
[NRWY12] Sahand N Negahban, Pradeep Ravikumar, Martin J Wainwright, and Bin Yu. A unified framework for high-dimensional analysis of $m$-estimators with decomposable regularizers. Statistical science, 27(4):538-557, 2012.
[NW72] John Ashworth Nelder and Robert WM Wedderburn. Generalized linear models. Journal of the Royal Statistical Society: Series A (General), 135(3):370-384, 1972.
[PJL20] Ankit Pensia, Varun Jog, and Po-Ling Loh. Robust regression with covariate filtering: Heavy tails and adversarial contamination. arXiv preprint arXiv:2009.12976, 2020.
[PSBR18] Adarsh Prasad, Arun Sai Suggala, Sivaraman Balakrishnan, and Pradeep Ravikumar. Robust estimation via robust gradient estimation. arXiv preprint arXiv:1802.06485, 2018.
[SBRJ19] Arun Sai Suggala, Kush Bhatia, Pradeep Ravikumar, and Prateek Jain. Adaptive hard thresholding for near-optimal consistent robust regression. In Conference on Learning Theory, pages 2892-2897. PMLR, 2019.
[SCV17] Jacob Steinhardt, Moses Charikar, and Gregory Valiant. Resilience: A criterion for learning in the presence of arbitrary outliers. arXiv preprint arXiv:1703.04940, 2017.
[SS19a] Yanyao Shen and Sujay Sanghavi. Iterative least trimmed squares for mixed linear regression. Advances in Neural Information Processing Systems, 32, 2019.
[SS19b] Yanyao Shen and Sujay Sanghavi. Learning with bad training data via iterative trimmed loss minimization. In International Conference on Machine Learning, pages 5739-5748. PMLR, 2019.
[Tuk75] John W Tukey. Mathematics and the picturing of data. In Proceedings of the International Congress of Mathematicians, Vancouver, 1975, volume 2, pages 523-531, 1975.
[Yat85] Yannis G Yatracos. Rates of convergence of minimum distance estimators and kolmogorov's entropy. The Annals of Statistics, 13(2):768-774, 1985.
[ZJS19] Banghua Zhu, Jiantao Jiao, and Jacob Steinhardt. Generalized resilience and robust statistics. arXiv preprint arXiv:1909.08755, 2019.
[ZJS22] Banghua Zhu, Jiantao Jiao, and Jacob Steinhardt. Robust estimation via generalized quasi-gradients. Information and Inference: A Journal of the IMA, 11(2):581-636, 2022.
[ZZ21] Ziwei Zhu and Wenjing Zhou. Taming heavy-tailed features by shrinkage. In International Conference on Artificial Intelligence and Statistics, pages 3268-3276. PMLR, 2021.

## Checklist

1. For all authors...
(a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
(b) Did you describe the limitations of your work? [Yes] (See Section6)
(c) Did you discuss any potential negative societal impacts of your work? [N/A]
(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
(a) Did you state the full set of assumptions of all theoretical results? [Yes]
(b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
(a) If your work uses existing assets, did you cite the creators? [N/A]
(b) Did you mention the license of the assets? [N/A]
(c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

## A Label Corruption Proofs

## A. 1 Gaussian

Proof of Theorem 4.2. Applying Lemma A.12 with $\eta=\varepsilon^{2}, C=1 / \sigma^{2}$ implies the output of Algorithm 1 is a $O(\varepsilon / \sigma)$ approximate stationary point. Lemma A.1 then implies that $\left\|\hat{\beta}-\beta^{*}\right\|=$ $O(\sigma \varepsilon \log (1 / \varepsilon))$. To bound the number of iterations, we need an upper bound on the negative loglikelihood on $\beta=0$, and a uniform lower bound on the negative log-likelihood. The initial negative log-likelihood is upper bounded by $\frac{1}{n} \sum_{i \in \hat{S}^{(1)}} y_{i}^{2} / \sigma^{2}$ where $S^{(1)}$ contains the smallest $(1-\varepsilon) n$ labels, which is bounded by $O\left(\max \left(1,1 / \sigma^{2}\right)\right)$. Trivially, there is a lower bound on the negative $\log$-likelihood for Gaussian. Therefore, the algorithm will terminate in $\frac{1}{\min \left(1, \sigma^{2}\right) \varepsilon_{c}^{2}}$ iterations.

Lemma A. 1 (Approximate stationary point close to $\beta^{*}$ for Gaussian regression). Given a set of datapoints $S=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ generated by a Gaussian regression model with $\varepsilon$-fraction of label corruption. Let $\hat{\beta}$ be a $\varepsilon / \sigma$-stationary point defined in Definition D.4. Given that $n=\Omega\left(\frac{d+\log (1 / \delta)}{\varepsilon^{2}}\right)$, with probability $1-\delta$, it holds that

$$
\left\|\hat{\beta}-\beta^{*}\right\|=O(\sigma \varepsilon \log (1 / \varepsilon))
$$

Proof. Let $\hat{S}$ be the set defined in Definition D.4. The first order stationary property guarantees

$$
\frac{1}{n} \sum_{i \in \hat{S}} \frac{1}{\sigma^{2}}\left(y_{i}-\hat{\beta}^{\top} \mathbf{x}_{i}\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \leq \frac{\varepsilon}{\sigma}\left\|\beta^{*}-\hat{\beta}\right\|
$$

Denote $T=G \backslash L$ as the uncorrupted set of data points. Then we get

$$
\begin{equation*}
\frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-\hat{\beta}^{\top} \mathbf{x}_{i}\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \leq-\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-\hat{\beta}^{\top} \mathbf{x}_{i}\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}+\sigma \varepsilon\left\|\beta^{*}-\hat{\beta}\right\| \tag{1}
\end{equation*}
$$

## Lower bound on the LHS

We will first establish a lower bound on the LHS of Equation 1, which contains terms from $\hat{S} \cap T$. Note that

$$
\begin{align*}
& \frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-\hat{\beta}^{\top} \mathbf{x}_{i}\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \\
= & \frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(\eta_{i}+\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \\
= & \frac{1}{n} \sum_{i \in \hat{S} \cap T} \eta_{i}\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}+\left(\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)^{2} \\
\gtrsim & -\sigma \varepsilon \log (1 / \varepsilon)\left\|\beta^{*}-\hat{\beta}\right\|+\left\|\beta^{*}-\hat{\beta}\right\|^{2}(1-\varepsilon \log (1 / \varepsilon)), \tag{2}
\end{align*}
$$

where we have leveraged the following resilience property in Equation 2

$$
\begin{array}{r}
\left\|\frac{1}{n} \sum_{i \in \hat{S} \cap T} \eta_{i} \mathbf{x}_{i}\right\| \lesssim \sigma \varepsilon \log (1 / \varepsilon) \\
\left\|\frac{1}{n} \sum_{i \in \hat{S} \cap T} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}-I\right\| \lesssim \varepsilon \log (1 / \varepsilon)
\end{array}
$$

Upper bound on the RHS

Then we establish an upper bound on the RHS of Equation 1. which contains terms from $\hat{S} \cap E$

$$
\begin{gathered}
\quad-\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-\hat{\beta}^{\top} \mathbf{x}_{i}\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \\
\stackrel{\text { Cauchy-Schwarz }}{\lesssim}\left(\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-\hat{\beta}^{\top} \mathbf{x}_{i}\right)^{2}\right)^{1 / 2}\left(\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)^{2}\right)^{1 / 2} \\
\stackrel{\text { resilience }}{\lesssim}\left(\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-\hat{\beta}^{\top} \mathbf{x}_{i}\right)^{2}\right)^{1 / 2} \varepsilon^{1 / 2} \log (1 / \varepsilon)^{1 / 2}\left\|\beta^{*}-\hat{\beta}\right\|
\end{gathered}
$$

Observe that since

$$
\left(\frac{1}{n} \sum_{i \in \hat{S}}\left(y_{i}-\hat{\beta}^{\top} \mathbf{x}_{i}\right)^{2}\right) \stackrel{\substack{\text { optimality of } \hat{S}}}{\leq}\left(\frac{1}{n} \sum_{i \in T}\left(y_{i}-\hat{\beta}^{\top} \mathbf{x}_{i}\right)^{2}\right)
$$

it holds that

$$
\begin{aligned}
\left(\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-\hat{\beta}^{\top} \mathbf{x}_{i}\right)^{2}\right) & \leq\left(\frac{1}{n} \sum_{i \in T}\left(y_{i}-\hat{\beta}^{\top} \mathbf{x}_{i}\right)^{2}\right)-\left(\frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-\hat{\beta}^{\top} \mathbf{x}_{i}\right)^{2}\right) \\
& =\left(\frac{1}{n} \sum_{i \in T \backslash \hat{S}}\left(y_{i}-\hat{\beta}^{\top} \mathbf{x}_{i}\right)^{2}\right) \\
& =\left(\frac{1}{n} \sum_{i \in T \backslash \hat{S}}\left(\eta_{i}+\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)^{2}\right) \\
& \stackrel{\text { resilience }}{=} O\left(\sigma^{2} \varepsilon \log (1 / \varepsilon)+\left\|\beta^{*}-\hat{\beta}\right\|^{2} \varepsilon \log (1 / \varepsilon)\right)
\end{aligned}
$$

Combining we get

$$
\begin{equation*}
\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-\hat{\beta}^{\top} \mathbf{x}_{i}\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \leq \varepsilon \log (1 / \varepsilon)\left(\sigma\left\|\beta^{*}-\hat{\beta}\right\|+\left\|\beta^{*}-\hat{\beta}\right\|^{2}\right) \tag{3}
\end{equation*}
$$

Plugging in Equation 2 and 3 into Equation 1, we we have

$$
\begin{array}{r}
-\sigma \varepsilon \log (1 / \varepsilon)\left\|\beta^{*}-\hat{\beta}\right\|+\left\|\beta^{*}-\hat{\beta}\right\|^{2}(1-\varepsilon \log (1 / \varepsilon)) \\
\lesssim \varepsilon \log (1 / \varepsilon)\left(\sigma\left\|\beta^{*}-\hat{\beta}\right\|+\left\|\beta^{*}-\hat{\beta}\right\|^{2}\right)+\varepsilon \sigma\left\|\beta^{*}-\hat{\beta}\right\| \\
\Longrightarrow \varepsilon \log (1 / \varepsilon)\left(\sigma\left\|\beta^{*}-\hat{\beta}\right\|+\left\|\beta^{*}-\hat{\beta}\right\|^{2}\right) \gtrsim\left\|\beta^{*}-\hat{\beta}\right\|^{2} \\
\Longrightarrow\left\|\beta^{*}-\hat{\beta}\right\| \lesssim \sigma \varepsilon \log (1 / \varepsilon)
\end{array}
$$

## A. 2 Poisson

Lemma A. 2 (Resilience condition for Poisson regression). With probability 0.99 it holds that for all $Q \subset T$ with $|Q| \geq(1-2 \varepsilon) n$,

$$
\frac{1}{n}\left\|\sum_{i \in Q}\left(y_{i}-\exp \left(\beta^{* \top} \mathbf{x}_{i}\right)\right) \mathbf{x}_{i}\right\| \lesssim \varepsilon \exp (\Theta(\sqrt{\log (1 / \varepsilon)}))
$$

and for all $Q \subset T$ with $Q \leq \varepsilon n$,

$$
\begin{aligned}
\frac{1}{n} \sum_{i \in Q} y_{i} & \lesssim \varepsilon \exp (\Theta(\sqrt{\log (1 / \varepsilon)})) \\
\frac{1}{n}\left\|\sum_{i \in Q} y_{i} \mathbf{x}_{i}\right\| & \lesssim \varepsilon \exp (\Theta(\sqrt{\log (1 / \varepsilon)}))
\end{aligned}
$$

Proof.
Proof of the first statement
We first prove the distribution of $\left(y-\exp \left(\beta^{* \top} \mathbf{x}\right)\right) \mathbf{x}$ is $k$-th moment bounded. Note that the $k$-th moment along direction $v$ can be written as

$$
\begin{gathered}
\mathbb{E}\left[\left(y-\exp \left(\beta^{*^{\top}} \mathbf{x}\right)\right)^{k}\left(\mathbf{v}^{\top} \mathbf{x}\right)^{k}\right] \\
\stackrel{\text { Cauchy-Schwarz }}{\leq} \sqrt{\mathbb{E}\left[\left(y-\exp \left(\beta^{*}{ }^{\top} \mathbf{x}\right)\right)^{2 k}\right] \sqrt{\mathbb{E}\left[\left(\mathbf{v}^{\top} \mathbf{x}\right)^{2 k}\right]}}
\end{gathered}
$$

Applying the Poisson $k$-th moment bound from Fact D.5, and the $k$-th moment bound of sub-Gaussian random variable yields

$$
\begin{align*}
& \leq \sqrt{\mathbb{E}\left[\max \left\{\exp \left(2 k\left(\beta^{*}{ }^{\top} \mathbf{x}\right)\right), 1\right\} 2 k^{2 k}\right]} \sqrt{(2 k)^{k}} \\
& \leq \exp \left(\Theta\left(k^{2}\left\|\beta^{*}\right\|^{2}\right)\right) k^{\Theta(k)}  \tag{4}\\
& =\exp \left(\Theta\left(k^{2}\right)\right)
\end{align*}
$$

where we use the fact that $\mathbb{E}[\exp (\lambda x)] \leq \exp \left(\kappa^{2} \lambda^{2}\right)$ for $\kappa$-sub-Gaussian random variable in Equation 4. Applying Corollary G. 1 in [ZJS19], we have that with probability $0.99, \forall Q \subset T, \mid Q \geq$ $(1-2 \varepsilon) n$,

$$
\frac{1}{n}\left\|\sum_{i \in Q}\left(y_{i}-\exp \left(\left\langle\beta^{*}, \mathbf{x}_{i}\right\rangle\right)\right) \mathbf{x}_{i}\right\| \leq \exp (C k)\left(\varepsilon^{1-1 / k}+\sqrt{d / n}\right)
$$

Setting $k=\sqrt{\log (1 / \varepsilon)} /$ in the above upper bound yields an upper bound of

$$
\varepsilon \exp (\Theta(\sqrt{\log (1 / \varepsilon)}))
$$

## Proof of the second statement

Since $\mathbb{E}\left[y^{k}\right]^{1 / k}=\exp (\Theta(k))$ Applying Corollary G. 1 in [ZJS19] and setting $k=\sqrt{\log (1 / \varepsilon)}$, we have that with probability $0.99, \forall Q \subset T, \mid Q \geq(1-\varepsilon) n$,

$$
\frac{1}{n}\left|\sum_{i \in Q}\left(y_{i}-\mathbb{E}[y]\right)\right| \leq \varepsilon \exp (\Theta(\sqrt{\log (1 / \varepsilon)}))
$$

This implies for all $Q \subset T$ with $|Q| \leq \varepsilon n$,

$$
\begin{array}{r}
\frac{1}{n}\left|\sum_{i \in Q} y_{i}-\mathbb{E}[y]\right| \leq \varepsilon \exp (\Theta(\sqrt{\log (1 / \varepsilon)})) \\
\mathbb{E}[y]=O(1) \\
\Longrightarrow
\end{array} \frac{1}{n}\left|\sum_{i \in Q} y_{i}\right| \leq \varepsilon \exp (\Theta(\sqrt{\log (1 / \varepsilon)}))
$$

Setting $k=\sqrt{\log (1 / \varepsilon)} /$ in the above upper bound yields an upper bound of $\varepsilon \exp (\Theta(\sqrt{\log (1 / \varepsilon)}))$
Proof of the third statement
Since we have the same $k$-th moment bound on $y$ and $y-\mathbb{E}[y]$. The bound $\frac{1}{n}\left\|\sum_{i \in Q} y_{i} \mathbf{x}_{i}\right\| \leq$ $\varepsilon \exp (\Theta(\sqrt{\log (1 / \varepsilon)}))$ can be proved similarly.

Theorem A. 3 (Poisson regression with label corruption (Restatement of Theorem4.3). Let $S=$ $\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ be a set of data points generated by a Poisson regression model with sub-Gaussian design, with $\varepsilon_{c}$-fraction of label corruption and $n=\Omega\left(\frac{d}{\varepsilon_{c}^{2}}\right)$. With probability 0.99, Algorithm 1 with parameters $\varepsilon=2 \varepsilon_{c}, \eta=\varepsilon_{c}^{2} /(d n), R \geq\left\|\beta^{*}\right\|$ terminate within $d n / \varepsilon_{c}^{2}$ iterations, and output an estimate $\hat{\beta}$ such that

$$
\left\|\hat{\beta}-\beta^{*}\right\| \lesssim \varepsilon_{c} \exp \left(\Theta\left(\sqrt{\log \left(1 / \varepsilon_{c}\right)}\right)\right)
$$

Proof of Theorem A.3. Lemma A. 13 implies the output of Algorithm 1 is a $\max \left(\varepsilon, \frac{\varepsilon^{2}}{\left\|\beta^{*}-\hat{\beta}\right\|}\right)$ approximate stationary point. Lemma A. 4 then implies that $\left\|\hat{\beta}-\beta^{*}\right\|=O(\varepsilon \exp (\sqrt{\log (1 / \varepsilon)}))$. To bound the number of iterations, we need an upper bound on the negative log-likelihood on $\beta=0$, and a uniform lower bound on the negative log-likelihood. The initial negative log-likelihood is upper bounded by $\frac{1}{n} \sum_{i \in \hat{S}^{(1)}} \log \left(y_{i}!\right)+1 \leq O\left(\mathbb{E}\left[y_{i}^{2}\right]+1\right)=O(1)$ where $S^{(1)}$ contains the smallest $(1-\varepsilon) n$ labels. Trivially, there is a 0 lower bound on the negative log-likelihood for Poisson distribution. Therefore, the algorithm will terminate in $d n / \varepsilon_{c}^{2}$ iterations.
Lemma A. 4 (Approximate stationary point close to $\beta^{*}$ for Poisson regression (Restatement of Lemma 4.7). Given a set of datapoints $S=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ generated by a Poisson regression model with $\varepsilon$-fraction of label corruption, and the largest $\varepsilon$ labels removed. Let $\hat{\beta}$ be a $\max \left(\varepsilon, \varepsilon^{2} / \| \beta^{*}-\right.$ $\hat{\beta} \|)$-stationary point defined in Definition D.4 and $\|\hat{\beta}\| \leq R$. Given that $n=\Omega\left(\frac{d}{\varepsilon^{2}}\right)$, with probability 0.99, it holds that

$$
\left\|\hat{\beta}-\beta^{*}\right\|=O(\varepsilon \exp (\Theta(\sqrt{\log (1 / \varepsilon)})))
$$

Proof. Recall that $y_{i} \sim \operatorname{Poi}\left(\exp \left(\beta^{* \top} \mathbf{x}_{i}\right)\right)$. Poisson regression log-likelihood:

$$
\log \operatorname{Pr}\left(y_{i} \mid\left\langle\beta^{*}, \mathbf{x}_{i}\right\rangle\right)=y_{i}\left(\beta^{* \top} \mathbf{x}_{i}\right)-\exp \left(\beta^{* \top} \mathbf{x}_{i}\right)-\log y_{i}!
$$

If $\left\|\beta^{*}-\hat{\beta}\right\| \leq \varepsilon$, the proof is done. When $\left\|\beta^{*}-\hat{\beta}\right\| \geq \varepsilon$, the first order approximate stationary property guarantees

$$
\begin{aligned}
& \frac{1}{n} \sum_{i \in S}\left(y_{i}-\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \leq \varepsilon\left\|\beta^{*}-\hat{\beta}\right\| \\
\Longrightarrow & \frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \\
\leq & -\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}+\varepsilon\left\|\beta^{*}-\hat{\beta}\right\|
\end{aligned}
$$

## Lower bound on the LHS

We will first establish a lower bound on the LHS, which contains terms from $\hat{S} \cap T$. Note that

$$
\begin{align*}
& \frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \\
= & \frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-\exp \left(\beta^{* \top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}  \tag{5}\\
& +\frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(\exp \left(\beta^{* \top} \mathbf{x}_{i}\right)-\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}, \tag{6}
\end{align*}
$$

and we bound the two terms separately.
Lemma A. 2 implies

$$
\begin{equation*}
\frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-\exp \left(\beta^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \leq\left\|\hat{\beta}-\beta^{*}\right\| \varepsilon \exp (\Theta(\sqrt{\log (1 / \varepsilon)})) \tag{7}
\end{equation*}
$$

Now we bound the second term in Equation 6. By resilience property (Proposition D.1), the set $L_{\beta}=\left\{\beta^{\top} \mathbf{x}_{i}<-C\|\beta\| \log (1 / \gamma)\right\}$ satisfies $\left|L_{\beta}\right| \leq \gamma n$ for any $\gamma>\varepsilon$, and also it is clear that $\frac{1}{n} \sum_{i \in(\hat{S} \cap T) \backslash L_{\beta}} \mathbf{x}_{i} \mathbf{x}_{i}^{\top} \succeq(1-C \gamma \log (1 / \gamma)) \cdot I$, hence

$$
\frac{1}{n} \sum_{i \in(\hat{S} \cap T) \backslash L_{\beta}} \exp \left(\beta^{\top} \mathbf{x}_{i}\right) \mathbf{x}_{i} \mathbf{x}_{i}^{\top} \succeq \exp (-C\|\beta\| \log (1 / \gamma))(1-C \gamma \log (1 / \gamma)) \cdot I
$$

Let $\gamma$ be a small constant yields

$$
\frac{1}{n} \sum_{i \in \hat{S} \cap T} \exp \left(\beta^{\top} \mathbf{x}_{i}\right) \mathbf{x}_{i} \mathbf{x}_{i}^{\top} \succeq e^{-O(\max (\|\beta\|, 1))} I
$$

This implies that $\sum_{i \in \hat{S} \cap T} \exp \left(\beta^{\top} \mathbf{x}_{i}\right)$ is a strongly convex function in $\beta$ in the range of $\|\beta\|=O(1)$. Since both $\left\|\beta^{*}\right\|=O(1)$ and $\|\hat{\beta}\|=O(1)$, by the definition of strongly convex function

$$
\begin{aligned}
& \frac{1}{n}\left(\nabla \sum_{i \in \hat{S} \cap T} \exp \left(\beta^{* \top} \mathbf{x}_{i}\right)-\nabla \sum_{i \in \hat{S} \cap T} \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right) \geq \Omega\left(\left\|\beta^{*}-\hat{\beta}\right\|^{2}\right) \\
\Longrightarrow & \frac{1}{n}\left(\sum_{i \in \hat{S} \cap T} \exp \left(\beta^{* \top} \mathbf{x}_{i}\right)-\sum_{i \in \hat{S} \cap T} \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \geq \Omega\left(\left\|\beta^{*}-\hat{\beta}\right\|^{2}\right)
\end{aligned}
$$

Combining this term with Equation 7, we have shown that

$$
\begin{equation*}
\frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)(\beta-\hat{\beta})^{\top} \mathbf{x}_{i} \geq C_{1}\left\|\beta^{*}-\hat{\beta}\right\|^{2}-\left\|\hat{\beta}-\beta^{*}\right\| \varepsilon \exp \left(C_{2} \sqrt{\log (1 / \varepsilon)}\right) \tag{8}
\end{equation*}
$$

## Upper bound on the RHS

1. Upper bound on the negative log-likelihood. Recall that $|\hat{S}|=(1-2 \varepsilon) n,|T|=(1-\varepsilon) n$, and the negative log-likelihood of Poisson regression is

$$
\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)+\log y_{i}!.
$$

By the optimality of $\hat{S}$, it must hold that

$$
\begin{aligned}
& \max _{i \in \hat{S} \cap E} \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)+\log y_{i}! \\
\leq & \min _{i \in T \backslash \hat{S}} \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)+\log y_{i}
\end{aligned}
$$

since otherwise one can replace a data point in $\hat{S} \cap E$ by one in $T \backslash \hat{S}$. Since $|T \backslash \hat{S}| \geq \varepsilon n$

$$
\begin{aligned}
& \min _{i \in T \backslash \hat{S}} \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)+\log y_{i} \\
\leq & \max _{Q \subset T,|Q|=\varepsilon n} \min _{i \in Q} \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)+\log y_{i}! \\
\leq & \left(\max _{Q \subset T,|Q|=\varepsilon n / 3} \min _{i \in Q} \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)+\max _{Q \subset T,|Q|=\varepsilon n / 3} \min _{i \in Q}-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)+\max _{Q \subset T,|Q|=\varepsilon n / 3} \min _{i \in Q} \log y_{i}!\right)
\end{aligned}
$$

Applying Lemma A. 2 and Proposition D. 1 gives

$$
\begin{aligned}
& \leq\left(\exp (\Theta(\sqrt{\log (1 / \varepsilon)}))+\max _{Q \subset T,|Q|=\varepsilon n / 3} \min _{i \in Q} y_{i} \log y_{i}\right) \\
& \leq(\exp (\Theta(\sqrt{\log (1 / \varepsilon)}))+\exp (\Theta(\sqrt{\log (1 / \varepsilon)})) \Theta(\sqrt{\log (1 / \varepsilon)})) \\
& \leq(\exp (\Theta(\sqrt{\log (1 / \varepsilon)})))
\end{aligned}
$$

2. Turn likelihood bound into square error bound. Now we have an upper bound on the negative log-likelihood of a data point in $\hat{S} \cap E$, next step we will turn the log-likelihood bound to a squared error bound. Define proxy function $g_{\hat{\beta}^{\top} \mathbf{x}_{i}}\left(y_{i}\right)$ as

$$
g_{\hat{\beta}^{\top} \mathbf{x}_{i}}\left(y_{i}\right)=\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)+y_{i} \log y_{i}-y_{i}
$$

which, by Fact D.6 is always smaller the negative log-likelihood function

$$
-\log f\left(y_{i} \mid\left\langle\hat{\beta}, \mathbf{x}_{i}\right\rangle\right)=\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)+\log y_{i}!
$$

and hence $g_{\hat{\beta}^{\top} \mathbf{x}_{i}}\left(y_{i}\right) \leq \exp (\Theta(\sqrt{\log (1 / \varepsilon)}))$.

Note that

$$
\begin{array}{r}
g_{\hat{\beta}^{\top} \mathbf{x}_{i}}\left(\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)=0 \\
g_{\hat{\beta}^{\top} \mathbf{x}_{i}}^{\prime}\left(\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)=0 \\
g_{\hat{\beta}^{\top} \mathbf{x}_{i}}^{\prime}\left(y_{i}\right)=\frac{1}{y_{i}} .
\end{array}
$$

Therefore, we can lower bound $g_{\hat{\beta}^{\top} \mathbf{x}_{i}}^{\prime}\left(y_{i}\right)$ as

$$
g_{\hat{\beta}^{\top} \mathbf{x}_{i}}\left(y_{i}\right) \geq \min \left(\frac{1}{y_{i}}, \frac{1}{\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)}\right)\left(y_{i}-\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)^{2} .
$$

Now there are two cases to consider: 1) if $y_{i} \geq \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right) / 2$, it holds that

$$
\left(y_{i}-\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)^{2} \leq \sqrt{y_{i} g_{\hat{\beta}^{\top} \mathbf{x}_{i}}\left(y_{i}\right)} \leq 2 \exp (\Theta(\sqrt{\log (1 / \varepsilon)}))
$$

where we leveraged the fact that $g_{\hat{\beta}^{\top} \mathbf{x}_{i}}\left(y_{i}\right) \leq \exp (\Theta(\sqrt{\log (1 / \varepsilon)}))$ and $y_{i} \leq \exp (\Theta(\sqrt{\log (1 / \varepsilon)}))$ after throwing away the largest $\varepsilon n y_{i}$ s in the beginning of the algorithm. 2) if $y<\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right) / 2$, we have

$$
\begin{array}{r}
g_{\hat{\beta}^{\top} \mathbf{x}_{i}}\left(y_{i}\right) \geq \frac{1}{\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)}\left(y_{i}-\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)^{2} \geq \frac{1}{2}\left|y_{i}-\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right| \\
\Longrightarrow\left(y_{i}-\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)^{2} \leq \exp (\Theta(\sqrt{\log (1 / \varepsilon)}))
\end{array}
$$

By Cauchy Schwarz

$$
\begin{align*}
& -\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \\
\leq & \sqrt{\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)^{2}} \sqrt{\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)^{2}} \\
\leq & \varepsilon^{1 / 2} \exp (\Theta(\sqrt{\log (1 / \varepsilon)})) \varepsilon^{1 / 2} \log ^{1 / 2}(1 / \varepsilon)\left\|\beta^{*}-\hat{\beta}\right\| \\
= & \varepsilon \exp \left(C_{3}(\sqrt{\log (1 / \varepsilon)})\right)\left\|\beta^{*}-\hat{\beta}\right\| \tag{9}
\end{align*}
$$

Combining LHS and RHS. Combining Equation 8 and Equation 9 yields

$$
\begin{array}{r}
C_{1}\left\|\beta^{*}-\hat{\beta}\right\|^{2}-\left\|\hat{\beta}-\beta^{*}\right\| \varepsilon \exp \left(C_{2}(\sqrt{\log (1 / \varepsilon)})\right) \\
\leq \varepsilon \exp \left(C_{3}(\sqrt{\log (1 / \varepsilon)})\right)\left\|\beta^{*}-\hat{\beta}\right\|+\varepsilon\left\|\beta^{*}-\hat{\beta}\right\| \\
\Longrightarrow\left\|\beta^{*}-\hat{\beta}\right\| \leq \varepsilon \exp \left(C_{4}(\sqrt{\log (1 / \varepsilon)})\right)
\end{array}
$$

## A. 3 Binomial

Proof of Theorem 4.4. Lemma A. 12 implies the output of Algorithm 1 is a $\max \left(\varepsilon / \sqrt{m}, \frac{\varepsilon^{2}}{m\left\|\beta^{*}-\hat{\beta}\right\|}\right)$ approximate stationary point. Lemma A. 5 then implies that $\left\|\hat{\beta}-\beta^{*}\right\|=O\left(\varepsilon \sqrt{\frac{\log (m / \varepsilon) \log (1 / \varepsilon)}{m}}\right)$. The initial negative log-likelihood is upper bounded by $m$, and trivially, there is a 0 lower bound on the negative log-likelihood. Therefore, the algorithm will terminate in $m^{2} / \varepsilon_{c}^{2}$ iterations.
Lemma A. 5 (Approximate stationary point close to $\beta^{*}$ for Binomial regression). Given a set of datapoints $S=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ generated by a Binomial regression model with $m$ trials and $\varepsilon$-fraction of label corruption. Let $\hat{\beta}$ be an $\max \left(\varepsilon / \sqrt{m}, \frac{\varepsilon^{2}}{m\left\|\beta^{*}-\hat{\beta}\right\|}\right)$-stationary point defined in Definition D. 4 with $\|\hat{\beta}\| \leq R$. Given that $n=\Omega\left(\frac{d+\log (1 / \delta)}{\varepsilon^{2}}\right)$, with probability $1-\delta$, it holds that

$$
\left\|\beta^{*}-\hat{\beta}\right\| \leq O\left(\varepsilon \sqrt{\frac{\log (m / \varepsilon) \log (1 / \varepsilon)}{m}}\right)
$$

Proof. Recall that the log-likelihood of Binomial regression is

$$
\log \operatorname{Pr}(y \mid\langle\beta, \mathbf{x}\rangle)=y \beta^{\top} \mathbf{x}-m \log \left(1+\exp \left(\beta^{\top} \mathbf{x}\right)\right)+\log \binom{m}{y}
$$

The following holds for fucntion $b\left(\beta^{\top} \mathbf{x}\right)$ in the Binomial regression.

$$
\begin{aligned}
b\left(\beta^{\top} \mathbf{x}\right) & =m \log \left(1+\exp \left(\beta^{\top} \mathbf{x}\right)\right) \\
\text { Mean function: } b^{\prime}\left(\beta^{\top} \mathbf{x}\right) & =m \frac{1}{1+\exp \left(-\beta^{\top} \mathbf{x}\right)} \\
\text { Variance function: } \quad b^{\prime \prime}\left(\beta^{\top} \mathbf{x}\right) & =m \frac{1}{\left(1+\exp \left(-\beta^{\top} \mathbf{x}\right)\right)\left(1+\exp \left(\beta^{\top} \mathbf{x}\right)\right)}
\end{aligned}
$$

The first order approximate stationary property guarantees

$$
\begin{aligned}
& \frac{1}{n} \sum_{i}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \leq \frac{\varepsilon}{\sqrt{m}}\left\|\beta^{*}-\hat{\beta}\right\| \\
& \Longrightarrow \frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \\
& \quad \leq-\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}+\frac{\varepsilon}{\sqrt{m}}\left\|\beta^{*}-\hat{\beta}\right\|
\end{aligned}
$$

## Lower bound on the LHS

We will first establish a lower bound on the LHS, which contains terms from $\hat{S} \cap T$. Note that

$$
\begin{array}{r}
\sum_{i \in \hat{S} \cap T}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right) \mathbf{x}_{i} \\
=\sum_{i \in \hat{S} \cap T}\left(y_{i}-b^{\prime}\left(\beta^{* \top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}+\sum_{i \in \hat{S} \cap T}\left(b^{\prime}\left(\beta^{* \top} \mathbf{x}_{i}\right)-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}, \tag{10}
\end{array}
$$

and we bound the two terms separately. Note that $\frac{y_{i}-b^{\prime}\left(\beta^{\top} \mathbf{x}_{i}\right)}{\sqrt{m}}$ has sub-Gaussian norm 1 , $\mathbf{x}_{i}$ has sub-Gaussian norm 1. Let

$$
\tilde{x}_{i}=\left[\begin{array}{c}
\mathbf{x}_{i} \\
\left(y_{i}-b^{\prime}\left(\beta^{\top} \mathbf{x}_{i}\right)\right) / \sqrt{m}
\end{array}\right]
$$

which is a 1 sub-Gaussian random vector with mean 0 and covariance $\tilde{\Sigma}=\left[\begin{array}{cc}I_{d} & 0 \\ 0 & c\end{array}\right]$ for a constant $c$. By Proposition D. 1 we have that

$$
\left\|\frac{1}{n} \sum_{i \in \hat{S} \cap T} \tilde{\mathbf{x}}_{i} \tilde{\mathbf{x}}_{i}^{\top}-\tilde{\Sigma}\right\| \leq \varepsilon \log (1 / \varepsilon)
$$

which implies for $\mathbf{u}=[0, \ldots, 1]^{\top} \in \mathbb{R}^{d+1}$ and any $\mathbf{w}=[\mathbf{v}, 0]^{\top},\|\mathbf{v}\|=1, \mathbf{v} \in \mathbb{R}^{d}$

$$
\begin{aligned}
& \frac{1}{n \sqrt{m}} \sum_{i \in \hat{S} \cap T}\left(y_{i}-b^{\prime}\left(\beta^{* \top} \mathbf{x}_{i}\right)\right) \mathbf{v}^{\top} \mathbf{x}_{i} \\
& =\frac{1}{n} \sum_{i \in \hat{S} \cap T} \mathbf{u}^{\top} \tilde{\mathbf{x}}_{i} \tilde{\mathbf{x}}_{i}^{\top} \mathbf{w} \leq \varepsilon \log (1 / \varepsilon)
\end{aligned}
$$

Hence,

$$
\begin{equation*}
\frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-b^{\prime}\left(\beta^{* \top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \leq \sqrt{m} \varepsilon \log (1 / \varepsilon)\left\|\beta^{*}-\hat{\beta}\right\| \tag{11}
\end{equation*}
$$

Now we bound the second term in Equation 10 By resilience property of sub-Gaussian distribution, for any $\|\beta\|=O(1)$, there is a constant $C$ such that the set $L_{\beta}=\left\{\left|\beta^{\top} \mathbf{x}_{i}\right|>C \log (1 / \gamma)\right\}$ satisfies $\left|L_{\beta}\right| \leq \gamma n$ for any $\gamma>\varepsilon$, and also $\sum_{i \in(\hat{S} \cap T) \backslash L_{\beta}} \mathbf{x}_{i} \mathbf{x}_{i}^{\top} \succeq(1-C \gamma \log (1 / \gamma)) \cdot I$. Hence

$$
\sum_{i \in(\hat{S} \cap T) \backslash L_{\beta}} b^{\prime \prime}\left(\beta^{\top} \mathbf{x}_{i}\right) \mathbf{x}_{i} \mathbf{x}_{i}^{\top} \succeq m \frac{1}{\left(1+1 / \gamma^{C}\right)\left(1+\gamma^{C}\right)}(1-C \gamma \log (1 / \gamma)) I \succeq \Theta(m \cdot I)
$$

by Setting $\gamma$ as a small constant. This implies $\sum_{i \in(\hat{S} \cap T) \backslash L_{\beta}} b\left(\beta^{\top} \mathbf{x}_{i}\right)$ is a strongly convex function with parameter $\Theta(m)$ when $\|\beta\|=O(1)$. By the definition of strongly convex function, we have

$$
\sum_{i \in \hat{S} \cap T}\left(b^{\prime}\left(\beta^{\top} \mathbf{x}_{i}\right)-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)(\beta-\hat{\beta})^{\top} \mathbf{x}_{i} \geq \Omega\left(m\|\beta-\hat{\beta}\|^{2}\right)
$$

Combining the two terms, we have shown that

$$
\begin{equation*}
\frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \geq C_{1}\left(m\left\|\beta^{*}-\hat{\beta}\right\|^{2}\right)-C_{2}\left(\sqrt{m}\left\|\hat{\beta}-\beta^{*}\right\| \varepsilon \log (1 / \varepsilon)\right) \tag{12}
\end{equation*}
$$

## Upper bound on the RHS

1. Upper bound on the negative log-likelihood. By the optimality of $\hat{S}$, it must hold that

$$
\begin{aligned}
& \sum_{i \in \hat{S} \cap E} b\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-\log \binom{m}{y_{i}} \leq \sum_{i \in T \backslash \hat{S}} b\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-\log \binom{m}{y_{i}} \\
& \leq \sum_{i \in T \backslash \hat{S}} b\left(\beta^{* \top} \mathbf{x}_{i}\right)-y_{i}\left(\beta^{* \top} \mathbf{x}_{i}\right)-\log \binom{m}{y_{i}} \\
& +\sum_{i \in T \backslash \hat{S}}\left(b\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)-\left(b\left(\beta^{* \top} \mathbf{x}_{i}\right)-y_{i}\left(\beta^{* \top} \mathbf{x}_{i}\right)\right)
\end{aligned}
$$

The first summation corresponds to the negative log-likelihood of good data under the right model $\beta^{*}$. The second summation corresponds to the shift in the likelihood from $\beta^{*}$ to $\hat{\beta}$. We first bound the first summation.
Define random variable $z_{i}=-\log f\left(y_{i} \mid\left\langle\beta^{*}, \mathbf{x}_{i}\right\rangle\right)=b\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-\log \binom{m}{y_{i}}$. Notice that condition on $\mathbf{x}_{i}, z_{i}$ can only take $m$ values, hence it is easy to see that $\forall \delta, \operatorname{Pr}\left(z_{i} \geq \log (1 / \delta) \mid \mathbf{x}_{i}\right) \leq m \delta$. Since this is true regardless of $\mathbf{x}_{i}$, we have

$$
\begin{aligned}
\operatorname{Pr}\left(z_{i} \geq \log (1 / \delta)\right) & \leq m \delta \\
\Longrightarrow \operatorname{Pr}\left(z_{i}-\log m \geq t\right) & \leq e^{-t}
\end{aligned}
$$

Hence $z_{i}-\log m$ is a 1 sub-exponential random variable. By the resilience property (Corollary D.2, we have that

$$
\frac{1}{n} \sum_{i \in T \backslash \hat{S}} b\left(\beta^{* \top} \mathbf{x}_{i}\right)-y_{i}\left(\beta^{* \top} \mathbf{x}_{i}\right)-\log \binom{m}{y_{i}}=\frac{1}{n} \sum_{i \in T \backslash \hat{S}} z_{i} \lesssim \varepsilon \log (m / \varepsilon)
$$

Now we bound the second summation. Note that

$$
\begin{aligned}
& b\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right) \\
\leq & b\left(\beta^{*}{ }^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\beta^{* \top} \mathbf{x}_{i}\right)+\left(b^{\prime}\left(\beta^{* \top} \mathbf{x}_{i}\right)-y_{i}\right)\left(\hat{\beta}-\beta^{*}\right)^{\top} \mathbf{x}_{i}+\max b^{\prime \prime}\left(\beta^{\top} \mathbf{x}_{i}\right) \cdot\left(\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)^{2}
\end{aligned}
$$

Leveraging Equation 11 , resilience of sub-Gaussian sample (Proposition D.1, and $b^{\prime \prime}\left(\beta^{\top} \mathbf{x}_{i}\right) \leq m$, we get

$$
\begin{aligned}
& \left.\frac{1}{n} \sum_{i \in T \backslash \hat{S}} b\left(\hat{\beta}^{\top}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-\sum_{i \in T \backslash \hat{S}} b\left(\beta^{* \top} \mathbf{x}_{i}\right)-y_{i}\left(\beta^{* \top} \mathbf{x}_{i}\right)\right) \\
& \quad \lesssim \sqrt{m} \varepsilon \log (1 / \varepsilon)\left\|\beta^{*}-\hat{\beta}\right\|+m \varepsilon \log (1 / \varepsilon)\left\|\beta^{*}-\hat{\beta}\right\|^{2}
\end{aligned}
$$

Combining the two summations yields

$$
\begin{array}{r}
\frac{1}{n} \sum_{i \in \hat{S} \cap E} b\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-\log \binom{m}{y_{i}} \\
\leq \frac{1}{n} \sum_{i \in T \backslash \hat{S}} b\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-\log \binom{m}{y_{i}} \\
\lesssim \sqrt{m} \varepsilon \log (1 / \varepsilon)\left\|\beta^{*}-\hat{\beta}\right\|+m \varepsilon \log (1 / \varepsilon)\left\|\beta^{*}-\hat{\beta}\right\|^{2}+\varepsilon \log (m / \varepsilon)
\end{array}
$$

2. Turn likelihood bound into square error bound. From Fact D.7. define proxy function

$$
\begin{array}{r}
g_{\hat{\beta}^{\top} \mathbf{x}_{i}}\left(y_{i}\right)=b\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)+y_{i} \log \left(\frac{y_{i}}{m}\right)+\left(m-y_{i}\right) \log \left(\frac{m-y_{i}}{m}\right)+\log \frac{y(m-y)}{m}+C \\
\leq b\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-\log \binom{m}{y_{i}}
\end{array}
$$

Note that

$$
\begin{array}{r}
g_{\hat{\beta}^{\top} \mathbf{x}_{i}}\left(b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)=C \\
g_{\hat{\beta}^{\top} \mathbf{x}_{i}}^{\prime}\left(b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)=0 \\
g_{\hat{\beta}^{\top} \mathbf{x}_{i}}^{\prime}\left(y_{i}\right)=\frac{1}{y_{i}}+\frac{1}{m-y_{i}} \geq 4 / m
\end{array}
$$

Combining this with the likelihood bound we get

$$
\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)^{2} \lesssim \varepsilon\left(m^{2}\left(\left\|\beta^{*}-\hat{\beta}\right\|^{2}+\sqrt{1 / m}\left\|\beta^{*}-\hat{\beta}\right\|\right) \log (1 / \varepsilon)+(\log m / \varepsilon) m\right)
$$

By Cauchy Schwaz

$$
\begin{array}{r}
-\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)(\beta-\hat{\beta})^{\top} \mathbf{x}_{i} \lesssim \sqrt{\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-b\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)^{2}} \sqrt{\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left((\beta-\hat{\beta})^{\top} \mathbf{x}_{i}\right)^{2}} \\
\lesssim \varepsilon m\left(\sqrt{\left\|\beta^{*}-\hat{\beta}\right\|^{2}+\frac{1}{\sqrt{m}}\left\|\beta^{*}-\hat{\beta}\right\|} \log (1 / \varepsilon)+\sqrt{\frac{\log (m / \varepsilon) \log (1 / \varepsilon)}{m}}\right) \cdot\left\|\beta^{*}-\hat{\beta}\right\| \tag{13}
\end{array}
$$

Combining Equation 12 and Equation 13 yields

$$
\left\|\beta^{*}-\hat{\beta}\right\| \leq O\left(\varepsilon \sqrt{\frac{\log (m / \varepsilon) \log (1 / \varepsilon)}{m}}\right)
$$

## A. 4 Generalized Linear Model

Theorem A.6 (Generalized linear model with label corruption (Restatement of Theorem4.5). Let $S=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ be generated by a generalized linear model with sub-Gaussian Design, with $\varepsilon_{c^{-}}$ fraction of label corruption and $n=\Omega\left(\frac{d+\log (1 / \delta)}{\varepsilon^{2}}\right)$. Assuming that $C_{0} \leq b^{\prime \prime}(\cdot) \leq C$ for non-zero constants $C_{0}, C, b(0)=0, b^{\prime}(0)=0$, and $\log (c(y))=O(\log (1 / \varepsilon)), \forall y \leq \Theta(\sqrt{\log (1 / \varepsilon)})$ With probability $1-\delta$, Algorithm $\left\lceil\right.$ with parameters $\varepsilon=\varepsilon_{c}, \eta=\varepsilon_{c}^{2}, R=\infty$ terminate within $\log \left(1 / \varepsilon_{c}\right) / \varepsilon_{c}^{2}$ iterations, and output an estimate $\hat{\beta}$ such that

$$
\left\|\hat{\beta}-\beta^{*}\right\|=O\left(\varepsilon_{c} \log \left(1 / \varepsilon_{c}\right)\right)
$$

Proof. Lemma A. 12 implies the output of $\operatorname{Algorithm~} 1$ is a $\max \left(\varepsilon_{c}, \frac{\varepsilon_{c}^{2}}{\left\|\beta^{*}-\hat{\beta}\right\|}\right)$ approximate stationary point. Lemma A. 8 then implies that $\left\|\hat{\beta}-\beta^{*}\right\|=O\left(\varepsilon_{c} \log \left(1 / \varepsilon_{c}\right)\right)$. Now we analyze the
iteration complexity. Note that Equation 17 shows that the $\varepsilon$ quantile of negative log-likelihood $-\log f\left(y_{i} \mid\left\langle\beta^{*}, \mathbf{x}_{i}\right\rangle\right)$ is upper bounded by $\log (1 / \varepsilon)$. Since the algorithm start $\hat{\beta}$ from 0 , using the fact that $b(0)=0$, the initial negative log-likelihood is upper bounded by

$$
\begin{array}{r}
-\log f\left(y_{i} \mid\left\langle\beta^{*}, \mathbf{x}_{i}\right\rangle\right)-b\left(\left\langle\beta^{*}, \mathbf{x}_{i}\right\rangle\right)+y_{i}\left\langle\beta^{*}, \mathbf{x}_{i}\right\rangle \\
\leq-\log f\left(y_{i} \mid\left\langle\beta^{*}, \mathbf{x}_{i}\right\rangle\right)+y_{i}\left\langle\beta^{*}, \mathbf{x}_{i}\right\rangle .
\end{array}
$$

Since $y_{i} \leq \sqrt{\log \left(1 / \varepsilon_{c}\right)}$ and the $\varepsilon_{c}$-quantile of $\left\langle\beta^{*}, \mathbf{x}\right\rangle$ is bounded by $\sqrt{\log \left(1 / \varepsilon_{c}\right)}$. We get that the negative $\log$-likelihood is upper bounded by $\Theta\left(\log \left(1 / \varepsilon_{C}\right)\right)$. From Equation 18 , we know that there is $\mathrm{a}-\log \left(1 / \varepsilon_{c}\right)$ lower bound on the negative log-likelihood. Therefore, the algorithm will terminate in $\log \left(1 / \varepsilon_{c}\right) / \varepsilon_{c}^{2}$ iterations.
Remark A.7. The assumption that $b^{\prime \prime}(\cdot) \leq C$ makes sure the variance of $y$ is bounded, and $b^{\prime \prime}(\cdot) \geq C_{0}$ makes sure the distribution of $y$ never degenerate to a singular point. Without these conditions, in the minimax sense, learning $\beta^{*}$ with finite sample is impossible. The additional condition on $c(y)$ exists to rule out having a single $y_{i}$ with extremely large density. This would be problematic since the adversary can inject $y_{i}$ at the point and the trimmed maximum likelihood estimator will not be able to remove these datapoints. Note that this is a mild condition since this essentially only requires $c(y) \leq \exp \left(y^{2}\right)$, and it is trivially true for probability mass function. All these assumptions are satisfied by the Gaussian regression model studied in this work.
Lemma A. 8 (Approximate stationary point close to $\beta^{*}$ for generalized linear regression). Given a set of datapoints $S=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ generated by a generalized linear model with $\varepsilon$-fraction of label corruption. Let $\hat{\beta}$ be an $\max \left(\varepsilon, \varepsilon^{2} /\left\|\beta^{*}-\hat{\beta}\right\|\right)$-stationary point defined in Definition D. 4 . Assuming that $C_{0} \leq b^{\prime \prime}(\cdot) \leq C$ for non-zero constants $C_{0}, C, b(0)=0, b^{\prime}(0)=0$, and $\log (c(y))=$ $O(\log (1 / \varepsilon)), \forall y \leq \Theta(\sqrt{\log (1 / \varepsilon)})$. Given that $n=\Omega\left(\frac{d+\log (1 / \delta)}{\varepsilon^{2}}\right)$, with probability $1-\delta$, it holds that

$$
\left\|\beta^{*}-\hat{\beta}\right\| \leq O(\varepsilon \log (1 / \varepsilon))
$$

Proof. The first order approximate stationary property guarantees

$$
\begin{aligned}
& \frac{1}{n} \sum_{i}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \leq \varepsilon\left\|\beta^{*}-\hat{\beta}\right\| \\
\Longrightarrow & \frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \\
& \leq-\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}+\varepsilon\left\|\beta^{*}-\hat{\beta}\right\|
\end{aligned}
$$

## Lower bound on the LHS

We first show $y$ is a sub-Gaussian random variable under our assumptions.
Proposition A.9. Suppose a generalized linear model satisfies $b^{\prime \prime}(\theta) \leq C$ for all $\theta \in \mathbb{R}$, then $(y-\mathbb{E}[y]) \mid\left\langle\beta^{*}, \mathbf{x}\right\rangle$ has sub-Gaussian norm $\sqrt{C}$ for any $\mathbf{x}$

Proof. First note that since a probability density function must sum to one, it hold that for any $\theta \in R$,

$$
\begin{array}{r}
\int c(y) \exp (\theta y-b(\theta)) d y=1 \\
\Longrightarrow \quad \int c(y) \exp (\theta y) d y=\exp (b(\theta)) .
\end{array}
$$

Also recall that $\mathbb{E}[y \mid \theta]=b^{\prime}(\theta)$. Fix $\theta$, the moment generating function of $y-\mathbb{E}[y]$ can be written as

$$
\begin{array}{r}
\mathbb{E}\left[e^{t\left(y-b^{\prime}(\theta)\right)}\right]=\int \exp \left(t\left(y-b^{\prime}(\theta)\right)\right) \cdot c(y) \exp (\theta \cdot y-b(\theta)) d y \\
=\int c(y)
\end{array} \begin{array}{r}
\exp \left((\theta+t) y-b(\theta)-t b^{\prime}(\theta)\right) d y \\
=\exp \left(b(\theta+t)-b(\theta)-t b^{\prime}(\theta)\right)
\end{array}
$$

Note that since $b^{\prime \prime}(\theta) \leq C$, it holds that $b(\theta+t) \leq b(\theta)+b^{\prime}(\theta) t+C t^{2}$, and therefore

$$
\begin{array}{r}
\exp \left(b(\theta+t)-b(\theta)-t b^{\prime}(\theta)\right) \\
\leq \exp \left(C t^{2}\right) \tag{15}
\end{array}
$$

This implies $y-\mathbb{E}[y] \mid\left\langle\beta^{*}, \mathbf{x}\right\rangle$ is $\sqrt{C}$-sub-Gaussian random variable.
Note that $\mathbb{E}[y \mid \mathbf{x}]=b^{\prime}\left(\left\langle\beta^{*}, \mathbf{x}\right\rangle\right) \leq C\left\langle\beta^{*}, \mathbf{x}\right\rangle+b^{\prime}(0)$. Since $\left\langle\beta^{*}, \mathbf{x}\right\rangle$ has sub-Gaussian norm 1, $\mathbb{E}[y \mid \mathbf{x}]-b^{\prime}(0)$ is a $C$-sub-Gaussian random variable. Therefore $y-b^{\prime}(0)$ is has sub-Gaussian norm $C+1$. Therefore, the following resilience property holds just like in the Gaussian case.
Proposition A. 10 (Resilience condition for generalized linear regression). With probability $1-\delta$ it holds that for all $Q \subset T$ with $|Q| \geq(1-2 \varepsilon) n$,

$$
\frac{1}{n}\left\|\sum_{i \in Q}\left(y_{i}-b^{\prime}\left(\beta^{* \top} \mathbf{x}_{i}\right)\right) \mathbf{x}_{i}\right\| \leq \Theta(\varepsilon \log (1 / \varepsilon))
$$

and for all $Q \subset T$ with $Q \leq \varepsilon n$,

$$
\begin{aligned}
\frac{1}{n} \sum_{i \in Q} y_{i} & \leq \Theta(\varepsilon \sqrt{\log (1 / \varepsilon)}) \\
\frac{1}{n}\left\|\sum_{i \in Q} y_{i} \mathbf{x}_{i}\right\| & \leq \Theta(\varepsilon \log (1 / \varepsilon))
\end{aligned}
$$

We conclude that

$$
\begin{equation*}
\frac{1}{n} \sum_{i \in \hat{S} \cap T}\left(y_{i}-b^{\prime}\left(\beta^{* \top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \leq \varepsilon \log (1 / \varepsilon)\left\|\beta^{*}-\hat{\beta}\right\| \tag{16}
\end{equation*}
$$

Next, note that by the lower bound on $b^{\prime \prime}(\cdot)$, it holds that

$$
\sum_{i \in(\hat{S} \cap T)} b^{\prime \prime}\left(\beta^{\top} \mathbf{x}_{i}\right) \mathbf{x}_{i} \mathbf{x}_{i}^{\top} \succeq \Theta(I)
$$

and hence by strong convexity

$$
\sum_{i \in \hat{S} \cap T}\left(b^{\prime}\left(\beta^{* \top} \mathbf{x}_{i}\right)-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} \geq \Theta\left(\left\|\beta^{*}-\hat{\beta}\right\|^{2}\right)
$$

Together we get

$$
\sum_{i \in \hat{S} \cap T}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)(\beta-\hat{\beta})^{\top} \mathbf{x}_{i} \geq \Theta\left(\left\|\beta^{*}-\hat{\beta}\right\|^{2}\right)+\Theta\left(\varepsilon \log (1 / \varepsilon)\left\|\beta^{*}-\hat{\beta}\right\|\right)
$$

## Upper bound on the RHS

## 1. Upper bound on the negative log-likelihood.

Define random variable $z_{i}=-\log f\left(y_{i} \mid\left\langle\beta^{*}, \mathbf{x}_{i}\right\rangle\right)$. Since $y_{i}-\mathbb{E}\left[y_{i} \mid \mathbf{x}_{i}\right]$ is a sub-Gaussian random variable, for a given $\mathbf{x}_{i}, \operatorname{Pr}\left(y_{i}-\mathbb{E}\left[y_{i}\right] \geq \sqrt{\log (1 / \delta)}\right) \leq \delta$, thus

$$
\begin{aligned}
& \operatorname{Pr}\left(z_{i} \geq \log (1 / \delta)\right)=\operatorname{Pr}\left(z_{i} \geq \log (1 / \delta) \mid y_{i} \leq \sqrt{\log (1 / \delta)}\right)+\operatorname{Pr}\left(y_{i} \geq \sqrt{\log (1 / \delta)}\right) \\
& \leq \\
& \hline \log (1 / \delta) \\
& \log \delta \\
& \Longrightarrow \operatorname{Pr}\left(z_{i} \geq t\right) \leq e^{-t} \sqrt{t} \leq e^{-t / 2}
\end{aligned}
$$

Since this is true regardless of $\mathbf{x}_{i}$, we have $z_{i}$ is a sub-exponential random variable. By the resilience property (Corollary D.2), we have that

$$
\begin{equation*}
\frac{1}{n} \sum_{i \in T \backslash \hat{S}}-\log f\left(y_{i} \mid\left\langle\beta^{*}, \mathbf{x}_{i}\right\rangle\right)=\frac{1}{n} \sum_{i \in T \backslash \hat{S}} z_{i} \leq \Omega(\varepsilon \log (1 / \varepsilon)) \tag{17}
\end{equation*}
$$

2. Turn likelihood bound into square error bound. Let $\theta=\left\langle\beta^{*}, \mathbf{x}\right\rangle$, and $g(\cdot)=b^{-1}(\cdot)$. Taking derivative of the negative likelihood function over the mean yields $\frac{b^{\prime}(\theta)-y}{b^{\prime \prime}(\theta)}$. This implies for each $y$, $b^{\prime}(\theta)=y$ minimize the negative log-likelihood function, and

$$
(-\log f(y \mid \theta)) \geq(-\log f(y \mid g(y)))+\left(b^{\prime}(\theta)-y\right)^{2} / C
$$

Next we lower bound the likelihood of the minimizer

$$
-\log f(y \mid g(y))=b(g(y))-y g(y)-\log (c(y))
$$

Since $b(0)=0$, and note that since $b^{\prime \prime}(\theta) \leq C$

$$
b(0) \leq b(\theta)-b^{\prime}(\theta) \theta+C \theta^{2}
$$

This implies

$$
-\log f(y \mid g(y)) \geq-C g(y)^{2}-\log (c(y))
$$

Since $b^{\prime}(x) \geq C_{0} x$, it holds that $g(y) \leq y / C_{0}$. This implies

$$
\begin{equation*}
-\log f(y \mid g(y)) \geq-C y^{2} / C_{0}-\log (c(y)) \tag{18}
\end{equation*}
$$

Since after truncation $y_{i}$ are all bounded by $\Theta(\sqrt{\log (1 / \varepsilon)})$ from Proposition A. 10 We have

$$
\frac{1}{n} \sum_{i \in \hat{S} \cap E}-\log f\left(y_{i} \mid\left\langle\beta^{*}, \mathbf{x}_{i}\right\rangle\right) \geq \Theta\left(\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-b^{\prime}\left(\beta^{* \top} \mathbf{x}_{i}\right)\right)^{2}\right)-\Theta(\varepsilon \log (1 / \varepsilon)) .
$$

Combining this with the negative log-likelihood upper bound proved for $T \backslash \hat{S}$, we have

$$
\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-b^{\prime}\left(\beta^{* \top} \mathbf{x}_{i}\right)\right)^{2} \leq \Theta(\varepsilon \log (1 / \varepsilon))
$$

By Cauchy Schwaz

$$
\begin{align*}
-\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-b^{\prime}\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i} & \leq \sqrt{\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(y_{i}-b\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\right)^{2}} \sqrt{\frac{1}{n} \sum_{i \in \hat{S} \cap E}\left(\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)^{2}} \\
& \lesssim \varepsilon \log (1 / \varepsilon)\left\|\beta^{*}-\hat{\beta}\right\| \tag{19}
\end{align*}
$$

Combining Equation 19 with the lower bound yields

$$
\left\|\beta^{*}-\hat{\beta}\right\|=O(\varepsilon \log (1 / \varepsilon))
$$

## A. 5 Convergence analysis of the alternating minimization algorithm

Proposition A.11. Suppose $\left\langle\nabla f(\hat{\beta}), \frac{\beta^{*}-\hat{\beta}}{\left\|\beta^{*}-\hat{\beta}\right\|}\right\rangle=\Delta$, and $\forall \beta, \frac{1}{\left\|\beta^{*}-\hat{\beta}\right\|^{2}}\left(\beta^{*}-\hat{\beta}\right)^{\top} \nabla^{2} f(\beta)\left(\beta^{*}-\hat{\beta}\right) \leq$ $H$. There exists a point $\beta$ such that

$$
f(\beta) \leq f(\hat{\beta})-\frac{\Delta^{2}}{4 H}
$$

Lemma $\mathbf{A . 1 2}$ (Algorithm 1 finds an approximate stationary point for generalized linear model). Given a set of datapoints $S=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ generated by a generalized linear model with $\varepsilon_{c}$-fraction of corruption. Assuming that $b^{\prime \prime}(\mu) \leq C$ for the generalized linear model and $n=\Omega\left(\frac{d+\log (1 / \delta)}{\varepsilon^{2}}\right)$, then with probability $1-\delta$, the output of Algorithm $1, \hat{\beta}$, is $\operatorname{arax}\left(2 \sqrt{C \eta}, \frac{2 \eta}{\left\|\beta^{*}-\hat{\beta}\right\|}\right)$-approximate stationary point with input parameters $R \geq\left\|\beta^{*}\right\|, \varepsilon=\varepsilon_{c}$. In particular, when $R=\infty$, the algorithm returns a $2 \sqrt{C \eta}$-approximate stationary point.

Proof. Given that Algorithm 1 returns $\hat{\beta}$, and the correponding set is $\hat{S}$. Let $f(\beta)$ be the negative log likelihood function. Then $\nabla^{2} f(\beta)=\frac{1}{n} \sum_{i \in \hat{S}} b^{\prime \prime}\left(\beta^{\top} x_{i}\right) \mathbf{x}_{i} \mathbf{x}_{i}^{\top} \preceq C \cdot I$. Therefore, $\frac{1}{\left\|\beta^{*}-\hat{\beta}\right\|^{2}}\left(\beta^{*}-\right.$ $\hat{\beta})^{\top} \nabla^{2} f(\beta)\left(\beta^{*}-\hat{\beta}\right) \leq C$. Applying Proposition A.11. it holds that there exists $\beta=\frac{\Delta}{2 C} \frac{\beta^{*}-\hat{\beta}}{\left\|\beta^{*}-\hat{\beta}\right\|}+\hat{\beta}$ such that

$$
f(\beta) \leq f(\hat{\beta})-\frac{\Delta^{2}}{4 C}
$$

When $R=\infty$, since the algorithm terminate when it is impossible to make an $\eta$ improvement over the current point $\hat{\beta}$, it holds that

$$
|\Delta| \leq 2 \sqrt{C \eta}
$$

and $\hat{\beta}$ is a $\sqrt{2 C \eta}$-approximate stationary point.
When $R$ only satisfies $\left\|\beta^{*}\right\| \leq R$, we need to make sure $\beta$ is a convex combination of $\beta^{*}$ and $\hat{\beta}$ to make sure it is a valid solution for the optimization problem. Therefore, if $\Delta / 2 C \geq\left\|\beta^{*}-\hat{\beta}\right\|$, we have

$$
f\left(\beta^{*}\right) \leq f(\hat{\beta})-\Delta\left\|\beta^{*}-\hat{\beta}\right\| / 2
$$

which implies

$$
\Delta \leq \frac{2 \eta}{\left\|\beta^{*}-\hat{\beta}\right\|}
$$

Since $b^{\prime \prime}(\theta)=\exp (\theta)$ is unbounded for Poisson regression, the above analysis does not apply, and here we prove our alternating minimization algorithm still returns an approximate stationary point
Lemma A. 13 (Algorithm 1 finds an approximate stationary point for Poisson regression (Restatement of Lemma 4.8). Given a set of datapoints $S=\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$ generated by a Poisson model with $\varepsilon_{c}$-fraction of corruption. Assuming that $n=\Omega\left(\frac{d+\log (1 / \delta)}{\varepsilon^{2}}\right)$, then with probability $1-\delta$, the output of Algorithm $\left\lfloor\right.$ with input parameters $\varepsilon=2 \varepsilon_{c}, R \geq\left\|\beta^{*}\right\|, \eta=\varepsilon^{2} /(d n)$, is a $\max \left(\varepsilon, \frac{2 \varepsilon^{2}}{\left\|\beta^{*}-\hat{\beta}\right\|}\right)$ approximate stationary point.

Proof. Define

$$
H=\frac{1}{n\left\|\beta^{*}-\hat{\beta}\right\|^{2}} \sum_{i \in \hat{S}} \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\left(\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)^{2}
$$

to be the second order derivative along the $\beta^{*}-\hat{\beta}$ direction. For every point $\beta=(1-\lambda) \hat{\beta}+\lambda \beta^{*}, 0 \leq$ $\lambda \leq 1$, the second order derivative is

$$
\begin{array}{r}
\frac{1}{n\left\|\beta^{*}-\hat{\beta}\right\|^{2}} \sum_{i \in \hat{S}} \exp \left(\left((1-\lambda) \hat{\beta}+\lambda \beta^{*}\right)^{\top} \mathbf{x}_{i}\right)\left(\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)^{2} \\
\leq \frac{1}{n\left\|\beta^{*}-\hat{\beta}\right\|^{2}}\left(\sum_{i \in \hat{S}} \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\left(\left(\hat{\beta}-\beta^{*}\right)^{\top} \mathbf{x}_{i}\right)^{2}+\sum_{i \in \hat{S}} \exp \left(\beta^{* \top} \mathbf{x}_{i}\right)\left(\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)^{2}\right) \\
\leq H+\frac{1}{n\left\|\beta^{*}-\hat{\beta}\right\|^{2}} \sum_{i \in \hat{S}} \exp \left(\beta^{* \top} \mathbf{x}_{i}\right)\left(\left(\beta^{*}-\hat{\beta}\right)^{\top} \mathbf{x}_{i}\right)^{2}
\end{array}
$$

Since $\beta^{* \top} \mathbf{x}_{i}$ is sub-Gaussian, it holds that $\max _{i \in \hat{S}} \beta^{* \top} \mathbf{x}_{i} \leq \sqrt{\log (n)}$ with probability 0.99 . Hence with probability 0.99 , we have

$$
\leq H+\exp (\sqrt{\log (n)})
$$

Consider two scenarios for the value of $H$.

1. $H \leq C \cdot d n$.

We can apply the same argument as in Lemma A. 12 and obtain that $\hat{\beta}$ is a $\max \left(2 \sqrt{d n \eta}, \frac{2 \eta}{\left\|\beta^{*}-\hat{\beta}\right\|}\right)$ approximate stationary point. Plugging in that $\eta=\varepsilon^{2} /(d n)$, we get that $\hat{\beta}$ is a $\max \left(\varepsilon, \frac{\varepsilon^{2}}{\left\|\beta^{*}-\hat{\beta}\right\|}\right)$ approximate stationary point.
2. $H \geq C \cdot d n$.

We will show this can not happen due to the termination condition of our algorithm. Due to the norm concentration of sub-Gaussian random vector, with probability 0.99 , it holds that $\max _{i \in \hat{S}}\left\|\mathbf{x}_{i}\right\|^{2} \leq$ $d+\Theta(\sqrt{d})$. Denote $D=\sum_{i} \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)$. Then

$$
D=\frac{1}{n} \sum_{i \in \hat{S}} \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right) \geq H / \max _{i \in \hat{S}}\left\|\mathbf{x}_{i}\right\|^{2} \geq \Theta(H / d)
$$

The following inequality always holds (log sum inequality)

$$
\frac{1}{n} \sum_{i \in \hat{S}} \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\left(\hat{\beta}^{\top} \mathbf{x}_{i}\right) \geq D \log (D)
$$

The following inequality holds since $\beta^{* \top} \mathbf{x}_{i} \leq \sqrt{\log (n)}$

$$
\frac{1}{n} \sum_{i} \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\left(\beta^{* \top} \mathbf{x}_{i}\right) \leq D \sqrt{\log (n)}
$$

Therefore,

$$
\frac{1}{n} \sum_{i} \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\left(\hat{\beta}^{\top}-\beta^{*}\right)^{\top} \mathbf{x}_{i} \geq D(\log (D)-\sqrt{\log n}) \geq \frac{H}{\Theta(d)}(\log (H / \Theta(d))-\sqrt{\log n})
$$

Since $H \geq \Theta(d n)$, it holds that

$$
\Delta:=\frac{1}{n} \sum_{i} \exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)\left(\hat{\beta}^{\top}-\beta^{*}\right)^{\top} \mathbf{x}_{i} \geq \Theta\left(\frac{H \sqrt{\log n}}{d}\right)
$$

and that the gradient satisfies of the

$$
\begin{aligned}
& \frac{1}{n} \sum_{i \in \hat{S}}\left(\exp \left(\hat{\beta}^{\top} \mathbf{x}_{i}\right)-y_{i}\right)\left(\hat{\beta}^{\top}-\beta^{*}\right)^{\top} \mathbf{x}_{i} \\
& \geq \frac{H}{d}-\exp (\sqrt{\log (1 / \varepsilon)}) \sqrt{d} \\
& \geq \Theta(H / d)
\end{aligned}
$$

where the first inequality holds since $y_{i} \leq \exp (\sqrt{\log (1 / \varepsilon)}),\left\|\mathbf{x}_{i}\right\|=\Theta(\sqrt{d})$, and the second inequality holds since $n \geq 2 \sqrt{d} / \varepsilon \geq 2 \exp (\sqrt{\log (1 / \varepsilon)}) \sqrt{d}$.
Note that in this regime the second order derivative along the $\beta^{*}-\hat{\beta}$ direction is upper bounded by $2 H$ for every point $\beta=(1-\lambda) \hat{\beta}+\lambda \beta^{*}, 0 \leq \lambda \leq 1$. Therefore, for every point in $\beta_{\lambda}=(1-\lambda) \hat{\beta}+\lambda \beta^{*}$, the value of the likelihood function is bounded as

$$
-\frac{1}{n} \sum_{i \in \hat{S}} \log f\left(y_{i} \mid\left\langle\beta, \mathbf{x}_{i}\right\rangle\right) \leq-\frac{1}{n} \sum_{i \in \hat{S}} \log f\left(y_{i} \mid\left\langle\hat{\beta}, \mathbf{x}_{i}\right\rangle\right)-\lambda \Delta+H \lambda^{2}
$$

If $\Delta \leq 2 H$, setting $\lambda=\Delta / 2 H$ yield

$$
-\frac{1}{n} \sum_{i \in \hat{S}} \log f\left(y_{i} \mid\left\langle\beta, \mathbf{x}_{i}\right\rangle\right) \leq-\frac{1}{n} \sum_{i \in \hat{S}} \log f\left(y_{i} \mid\left\langle\hat{\beta}, \mathbf{x}_{i}\right\rangle\right)-\Delta^{2} / 4 H
$$

Since $\Delta \gtrsim H / d$, and $H \geq n d$, the drop in objective value $\Delta^{2} / 4 H=\Omega(n / d) \geq \Omega\left(1 / \varepsilon^{2}\right)$.
If $\Delta \geq 2 H$, setting $\lambda=1$ yield

$$
-\frac{1}{n} \sum_{i \in \hat{S}} \log f\left(y_{i} \mid\left\langle\beta, \mathbf{x}_{i}\right\rangle\right) \leq-\frac{1}{n} \sum_{i \in \hat{S}} \log f\left(y_{i} \mid\left\langle\hat{\beta}, \mathbf{x}_{i}\right\rangle\right)-H
$$

and the drop in objective value is $H \geq n d \gg 1 / \varepsilon^{2}$. Since $\eta \leq 1 / \varepsilon^{2}$, we conclude that $H$ must be smaller than $C \cdot d n$ when the algorithm terminate.

## B Handling sample corruption model

This section we discuss the algorithm for the sample corruption model, where the adversary is allowed to change the covariate $\mathbf{x}_{i}$ in addition to label $y_{i}$. At the end of Algorithm 4 in [DHL19], with high probability, it will return a set of $\varepsilon$ corrupted data points that satisfies $\left\|\frac{1}{n} \sum_{i \in \hat{S}} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}-I\right\|=$ $O(\varepsilon \log (\varepsilon))$. Corollary D.3 then implies that $\left\|\frac{1}{n} \sum_{i \in \hat{S} \cap E} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}\right\|=O(\varepsilon \log (\varepsilon))$. The rest of the proof proceeds as in the label corruption setting.

## C Non-identity covariance

Our result of Theorem 4.2. Theorem 4.3, Theorem4.4 apply to the case where the covariance of $\mathbf{x}_{i}$ is identity. Here we argue that the result holds for general covariance $\Sigma$ where the guarantee is in terms of $\left\|\hat{\beta}-\beta^{*}\right\|_{\Sigma}$, which is the root-mean-square error in the Gaussian setting. First note that Algorithm 1 on input $\left(S=\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{n}, y_{n}\right), \varepsilon, \eta, R\right)$ with non-identity covariance $\Sigma$ output $\hat{\beta}$ if and only if running it on input $S=\left(\Sigma^{-1 / 2} \mathbf{x}_{1}, y_{1}\right), \ldots,\left(\Sigma^{-1 / 2} \mathbf{x}_{n}, y_{n}\right), \varepsilon, \eta$ with the constraint $\left\|\Sigma^{-1 / 2} \beta\right\|_{2} \leq R$ output $\Sigma^{1 / 2} \hat{\beta}$.
Although the constraint $\left\|\Sigma^{-1 / 2} \beta\right\|_{2} \leq R$ is different from the $\ell_{2}$ constraint in the analysis of Algorithm 1, with slightly different arguments as in Theorem 4.2. Theorem 4.3. Theorem4.4, we can show that running Algorithm 1 on $S=\left(\Sigma^{-1 / 2} \mathbf{x}_{1}, y_{1}\right), \ldots,\left(\Sigma^{-1 / 2} \mathbf{x}_{n}, y_{n}\right), \varepsilon, \eta$ with the constraint $\left\|\Sigma^{-1 / 2} \beta\right\|_{2} \leq R$ will output $\hat{\beta}^{\prime}$ such that $\left\|\hat{\beta}^{\prime}-\Sigma^{1 / 2} \beta\right\|_{2}$ has the desired error rate. This implies Algorithm 1 running on $\left(S=\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{n}, y_{n}\right), \varepsilon, \eta, R\right)$ will output $\hat{\beta}=\Sigma^{-1 / 2} \hat{\beta}^{\prime}$ such that $\|\hat{\beta}-\beta\|_{\Sigma}$ has the desired error rate.

## D Auxiliary Lemmas

Proposition D. 1 (Resilience of sub-Gaussian sample [JLT20, Corollary 4]). Let $G=\left\{\mathbf{x}_{i} \in \mathbb{R}^{d}\right\}_{i=1}^{n}$ be a dataset satisfies Assumption 3.3 . For any $\varepsilon \in[0,1 / 2]$. If

$$
\begin{equation*}
n=\Omega\left(\frac{d+\log (1 / \delta)}{\varepsilon^{2}}\right), \tag{20}
\end{equation*}
$$

then with probability at least $1-\delta$ there exists an absolute constant $C>0$ such that for any subset $T \subset G$ and $|T| \geq(1-\varepsilon) n$, we have

$$
\begin{array}{r}
\left\|\frac{1}{|T|} \sum_{i \in T} \mathbf{x}_{i}\right\| \leq C \varepsilon \sqrt{\log (1 / \varepsilon)} \\
\left\|\left(\frac{1}{|T|} \sum_{i \in T} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}\right)-I\right\| \leq C \varepsilon \log (1 / \varepsilon) \tag{22}
\end{array}
$$

and that for any subset $T \subset G$ and $|T| \geq \varepsilon n$, we have

$$
\begin{array}{r}
\left\|\frac{1}{|T|} \sum_{i \in T} \mathbf{x}_{i}\right\| \leq C \sqrt{\log (1 / \varepsilon)} \\
\left\|\left(\frac{1}{|T|} \sum_{i \in T} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}\right)-I\right\| \leq C \log (1 / \varepsilon) \tag{24}
\end{array}
$$

Corollary D. 2 (Resilience of 1-d sub-exponential sample). Let $G=\left\{x_{i} \in \mathbb{R}^{d}\right\}_{i=1}^{n}$ contains i.i.d. samples drawn from 1 sub-exponential distribution with zero mean. For any $\varepsilon \in[0,1 / 2]$. If

$$
\begin{equation*}
n=\Omega\left(\frac{\log (1 / \delta)}{\varepsilon^{2}}\right) \tag{25}
\end{equation*}
$$

then with probability at least $1-\delta$ there exists an absolute constant $C>0$ such that for any subset $T \subset G$ and $|T| \geq \varepsilon n$, we have

$$
\begin{equation*}
\left|\frac{1}{|T|} \sum_{i \in T} \mathbf{x}_{i}\right| \leq C \log (1 / \varepsilon) \tag{26}
\end{equation*}
$$

Corollary D. 3 (Resilience of corrupted sample with small covariance). Let $G=\left\{\mathbf{x}_{i} \in \mathbb{R}^{d}\right\}_{i=1}^{n}$ be a dataset satisfies the condition in Proposition $\overline{D .1}$ with $\Sigma=I$ and parameter $\varepsilon$. Let $S=G \backslash L \cup E=$ $T \cup E$ be obtained from set $G$ by corrupting $\varepsilon$ an samples arbitrarily. There exists an absolute constant $C>0$ such that if

$$
\left\|\frac{1}{n} \sum_{i \in \hat{S}} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}-I\right\|=O(\varepsilon \log (\varepsilon))
$$

for $\hat{S}$ with $\hat{S} \geq(1-\varepsilon) n$, we have

$$
\begin{array}{r}
\frac{1}{n}\left\|\sum_{i \in \hat{S} \cap T} \mathbf{x}_{i}\right\| \leq C \varepsilon \sqrt{\log (1 / \varepsilon)} \\
\left\|\left(\frac{1}{n} \sum_{i \in \hat{S} \cap T} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}\right)-I\right\| \leq C \varepsilon \log (1 / \varepsilon)
\end{array}
$$

and

$$
\begin{array}{r}
\frac{1}{n}\left\|\sum_{i \in \hat{S} \cap E} \mathbf{x}_{i}\right\| \leq C \varepsilon \sqrt{\log (1 / \varepsilon)} \\
\left\|\left(\frac{1}{n} \sum_{i \in \hat{S} \cap E} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}\right)\right\| \leq C \varepsilon \log (1 / \varepsilon)
\end{array}
$$

Proof. The statement about $T$ follows directly from resilience of sub-Gaussian sample. The statement about $E$ follows from the covariance bound and the condition of $T$.

Definition D. 4 (Approximate stationary point). Given a set of datapoints $\left\{\mathbf{x}_{i}, y_{i}\right\}_{i=1}^{n}$, probability density function $f(\cdot \mid \cdot)$, and $\beta^{*} \in \mathbb{R}^{d}, \hat{\beta} \in \mathbb{R}^{d}$. Let $\hat{S}=\arg \min _{S:|S|=(1-\varepsilon) n} \sum_{i \in S}-\log f\left(y_{i} \mid\left\langle\hat{\beta}, \mathbf{x}_{i}\right\rangle\right)$. We call $\hat{\beta}$ a $\eta$-stationary point if

$$
\left(\nabla_{\beta} \sum_{i \in \hat{S}}-\log f\left(y_{i} \mid\left\langle\hat{\beta}, \mathbf{x}_{i}\right\rangle\right)\right)^{\top} \frac{\left(\hat{\beta}-\beta^{*}\right)}{\left\|\hat{\beta}-\beta^{*}\right\|} \leq \eta
$$

Fact D. 5 ( $k$-th moment bound of Poisson Ah122]). Assuming that $y \sim \operatorname{Poi}(\lambda)$, then $\mathbb{E}\left[y^{k}\right]=$ $\max \left(k^{k}, \lambda^{k}\right)$ and $\mathbb{E}\left[(y-\lambda)^{k}\right]=\max \left(k^{k}, \lambda^{k}\right)$
Fact D.6. For all $n>0, n \log n-n+1 \leq \log n!\leq(n+1) \log n-n+1$
Fact D.7.

$$
\log \binom{m}{y} \leq-y \log (y / m)-(m-y) \log ((m-y) / m)-\log \frac{y(m-y)}{m}+C_{2}
$$

