# On the Algorithmic Stability of Adversarial Training

**Yue Xing**
Department of Statistics
Purdue University
xing49@purdue.edu

**Qifan Song**
Department of Statistics
Purdue University
qfsong@purdue.edu

**Guang Cheng**
Department of Statistics
Purdue University
chengg@purdue.edu

## Abstract

The adversarial training is a popular tool to remedy the vulnerability of deep learning models against adversarial attacks, and there is rich theoretical literature on the training loss of adversarial training algorithms. In contrast, this paper studies the algorithmic stability of a generic adversarial training algorithm, which can further help to establish an upper bound for generalization error. By figuring out the stability upper bound and lower bound, we argue that the non-differentiability issue of adversarial training causes worse algorithmic stability than their natural counterparts. To tackle this problem, we consider a noise injection method. While the non-differentiability problem seriously affects the stability of adversarial training, injecting noise enables the training trajectory to avoid the occurrence of non-differentiability with dominating probability, hence enhancing the stability performance of adversarial training. Our analysis also studies the relation between the algorithm stability and numerical approximation error of adversarial attacks.

## 1 Introduction

Successful machine learning algorithms require not only a good empirical performance but also generalizing well to unseen data. For the robustness towards unseen data, empirical experiments show that deep learning models can be fragile and vulnerable against adversarial input (Biggio et al., 2013; Szegedy et al., 2014). To set an example, in image recognition problems, a deep neural network will predict a wrong label when the testing image is slightly altered, while the change is barely recognizable by human eyes (Papernot et al., 2016a).

Related research efforts in adversarial learning include designing adversarial attacks in various applications (Papernot et al., 2016a,b; Moosavi-Dezfooli et al., 2016), detecting attacked samples (Tao et al., 2018; Ma and Liu, 2019), and modifications on the training process to obtain adversarially robust models, i.e., adversarial training (Shaham et al., 2015; Madry et al., 2017; Jalal et al., 2017).

However, although adversarial training improves the adversarial robustness during testing, its generalization performance is still poor. While Yin et al. (2018) presented that the adversarial Rademacher complexity is never smaller than its natural counterpart, Schmidt et al. (2018); Zhai et al. (2019) argued that a better adversarial generalization requires more labeled/unlabeled data.

In the literature of natural DNN optimization via iterative gradient moves, the empirical loss at each iteration can be characterized by convergence rate analysis, yet generalization properties are not well understood. To characterize the generalization error, one popular way is to study the algorithmic stability. Algorithmic stability is first considered by Kearns and Ron (1999); Bousquet and Elisseeff (2001, 2002). Later, Hardt et al. (2016) explored the connection between algorithmic stability and generalization performance of gradient-type optimization. Some follow-up research studies the stability for different classes of algorithms, or relax the definition of stability to generalize its usage, see Ramezani-Kebrya et al. (2018); Charles and Papailiopoulos (2018); Kuzborskij and Lampert (2018); Zhou et al. (2018); Lei and Ying (2020); Ho et al. (2020); Madden et al. (2020).

35th Conference on Neural Information Processing Systems (NeurIPS 2021).

In general, there are two ways to utilize algorithmic stability. On the one hand, as showed by Hardt et al. (2016), the algorithmic stability provides an upper bound for the generalization error; hence it will be useful when establishing the convergence of generalization error. On the other hand, the algorithmic stability itself is also a measure that evaluates the performance of an algorithm.

Our work extends algorithmic stability analysis to adversarial training. Our contributions are:

- Through figuring out the stability upper bound and lower bound, we argue that adversarial training leads to poor algorithmic stability even the clean loss is smooth. To solve this problem, we propose to inject noise into the adversarial training process. Although some existing works proposed the usage of noise injection, we highlight that such a method is more meaningful in adversarial training than its natural counterpart. Theoretical justification of the noise injection method is provided for a wide range of data-generating models in several tasks, including both linear regression and logistic classification.
- Noticing that, in practice, adversarial attacks are mostly approximated via numerical methods, e.g., fast gradient method (FGM) or projected gradient method (PGD), our theory investigates the role of accuracy of attack approximation for the stability of adversarial training algorithms. In short, a more accurate attack leads to better stability upper bound.
- The effectiveness of noise-injected adversarial training is further generalized to the $\mathcal{L}_\infty$ attack. Compared with $\mathcal{L}_2$, $\mathcal{L}_\infty$ training algorithm is generally less stable.
- Beyond the theoretical analysis under simple models, we provide a theory in two-layer ReLU network with lazy training (training the hidden layer) and observe the effectiveness of the noise injection method. We also obtain empirical evidence that for deep neural networks model, proper forms of noise injection and more accurate attack calculation (e.g., PGD-$k$ over FGM) improve the generalization error.

## 2 Related Works

**Theories in adversarial training** To theoretically understand how adversarial training works, Sinha et al. (2018); Wang et al. (2019a) investigated the convergence of adversarial training when the loss is strongly convex w.r.t. data attributes. In this case, it can be shown that the gradient of adversarial loss w.r.t. model parameters is Lipschitz, leading to good stability. However, when the loss is not strongly concave in data attributes, Xing et al. (2021a) figured out that adversarial training does not have a Lipschitz gradient even for linear regression. Some studies in deep neural networks (Gao et al., 2019; Zhang et al., 2020; Allen-Zhu and Li, 2020) studied the convergence of adversarial training loss, and Allen-Zhu and Li (2020) also provided a theoretical guarantee of the adversarial testing loss when attack strength is small enough. Some other literature in the generalization of adversarial training can be found in Khim and Loh (2018); Awasthi et al. (2020); Pinot et al. (2021); Xing et al. (2021b).

**Observations in deep learning** In terms of empirical studies, He et al. (2019); Zheltonozhskii et al. (2020); Xie et al. (2020); Lee and Chandrakasan (2020); Wu et al. (2020) focused on improving the performance of attack/adversarial robustness in deep learning.

In the literature, there are several ways to improve the adversarial training, including modifying the objective function to help the convergence of the training process (Zhang et al., 2019; Wang et al., 2019b), regularization (He et al., 2019; Zheltonozhskii et al., 2020; Wu et al., 2020), replacing non-smooth components (Lee and Chandrakasan, 2020; Xie et al., 2020), and handling over-fitting issue (Lee and Chandrakasan, 2020).

**Stability for non-smooth loss and min-max problem** Besides works in the algorithmic stability of first-order optimization methods on smooth loss, Bassily et al. (2020) studied the case when loss is convex but not smooth. In this scenario, the minimax lower bound and convergence upper bound of stability together imply that stochastic gradient descent[1] (SGD) and gradient descent (GD) have poor stability. It is recommended to run SGD/GD with an extremely small learning rate for a vast number of iterations, which implies that it is not practical to train a non-smooth model with good stability. Consequently, adaptations are essential for non-smooth models to improve the training process.

[1]We are considering sample-with-replacement SGD.

Another recent work, Farnia and Ozdaglar (2020), considered the algorithmic stability in the min-max problems for generative adversarial networks (GAN) to argue that simultaneous training in generator and discriminator leads to good stability. However, besides the strongly-convex-concave assumption in their loss, the "min-max" problem considered in GAN and adversarial training are not the same. These two differences lead to different conclusions between GAN and adversarial training.

# 3 Stability of adversarial training

In this section, we study the uniform argument stability (UAS) of adversarial training. Utilizing the notations introduced in Section 3.1, we present in Section 3.2 the upper and lower bounds of UAS. Section 3.3 studies the effect of attack error on stability.

## 3.1 Notations

**Adversarial training** To introduce adversarial training, let $l$ denote the loss function and $f_\theta(x)$ be the model with parameter $\theta$. The (population) adversarial loss is defined as

$$R(\theta, \epsilon) := \mathbb{E}\left[l\left(f_\theta[x + A_\epsilon(f_\theta, x, y)], y\right)\right],$$

where $A_\epsilon$ is an attack of strength $\epsilon > 0$ and intends to deteriorate the loss in the following way

$$A_\epsilon(f_\theta, x, y) := \underset{z \in B_p(0,\epsilon)}{\operatorname{argmax}} \{l(f_\theta(x+z), y)\}, \tag{1}$$

where $B_p(x, r)$ is a $\mathcal{L}_p$ ball centering at $x$ with radius $r$.

Given $n$ i.i.d. samples $S = \{(x_i, y_i)\}_{i=1}^n$, the adversarial training minimizes the sample version of $R(\theta, \epsilon)$ w.r.t. $\theta$:

$$R_S(\theta, \epsilon) = \frac{1}{n}\sum_{i=1}^{n} l\left(f_\theta[x_i + A_\epsilon(f_\theta, x_i, y_i)], y_i\right), \tag{2}$$

and the estimator $\widehat{\theta}$ aims to minimize $R_S(\theta, \epsilon)$. We rewrite $R_S(\theta, \epsilon)$ as $R_S(\theta)$ for simplicity when there is no confusion.

The minimization in (2) is often implemented through an iterative two-step (min-max) update. In the $t$-th iteration, we calculate the adversarial sample $\widetilde{x}_i^{(t)} = x_i + A_\epsilon(f_{\theta^{(t)}}, x_i, y_i)$ based on the current $\theta^{(t)}$, and then update $\theta^{(t+1)}$ based on the gradient of the adversarial training loss while fixing $\widetilde{x}_i^{(t)}$'s with learning rate $\eta_t$. The algorithm runs for $T$ iterations. A more detailed pseudocode is postponed to Algorithm 1 when introducing our adaptations. Note that for some loss function $l$ or model $f_\theta$, there may not be an analytic form for $A_\epsilon$ (e.g. deep neural networks), and numerical methods, e.g. FGM and PGD, are utilized to approximate $A_\epsilon$.

**Risk decomposition** Define $\theta_0$ and $\bar{\theta}$ as the minimizer of $R$ and $R_S$ respectively. Then for the algorithm output $\widehat{\theta}$, the excess risk can be decomposed into four parts as below:

$$R(\widehat{\theta}) - R(\theta_0) \quad = \quad \underbrace{R(\widehat{\theta}) - R_S(\widehat{\theta})}_{\mathcal{E}_{gen}} + \underbrace{R_S(\widehat{\theta}) - R_S(\bar{\theta})}_{\mathcal{E}_{opt}} + \underbrace{R_S(\bar{\theta}) - R_S(\theta_0)}_{\leq 0} + \underbrace{R_S(\theta_0) - R(\theta_0)}_{\mathbb{E}=0},$$

Since the last two parts are either negative or with zero expectation, we mainly focus on the first two parts, namely, generalization error $R(\widehat{\theta}) - R_S(\widehat{\theta})$ and optimization error $R_S(\widehat{\theta}) - R_S(\bar{\theta})$. Based on Hardt et al. (2016), $\mathcal{E}_{gen}$ is upper bounded by algorithmic stability.

**Uniform argument stability (UAS)** UAS aims to quantify the output sensitivity in $\mathcal{L}_2$ norm w.r.t an arbitrary change in a single data point. An algorithm is $\lambda$-UAS if for neighboring datasets $S_1 \sim S_2$ (i.e., $S_1$ and $S_2$ differ only in a single data point), it satisfies that

$$\sup_{S_1 \sim S_2} \|\widehat{\theta}(S_1) - \widehat{\theta}(S_2)\| := \sup_{S_1 \sim S_2} \lambda(S_1, S_2) \leq \lambda.$$

where $\|\cdot\|$ represents the $\mathcal{L}_2$ norm. Under proper conditions, the UAS bound implies a generalization error bound (Bassily et al., 2020): if $P(\|\lambda(S_1, S_2)\| \geq \gamma) \leq \kappa_0$ for any neighboring $(S_1, S_2)$, then for any $\kappa$,

$$P\left[|\mathcal{E}_{gen}| \geq c\left(\gamma(\log n)(\log(n/\kappa)) + \sqrt{\frac{\log(1/\kappa)}{n}}\right)\right] \leq \kappa + \kappa_0. \tag{3}$$

### 3.2 Upper and lower bound

This section presents the upper bound and lower bound of UAS of adversarial training when its natural counterpart is convex and smooth.

The upper bound of UAS of adversarial training can be directly extended from Bassily et al. (2020) as follows:

**Proposition 1.** *Assume $l(f_\theta(x), y)$ is L-Lipschitz and convex w.r.t. $\theta$, and $\theta \in B_2(0, r)$. The two models $\theta_1^{(t)}$ and $\theta_2^{(t)}$ are adversarial training estimators obtained using datasets $S_1, S_2$ respectively. For SGD,*

$$\sup_{S_1 \sim S_2} \mathbb{E}\left[\|\theta_1^{(T)} - \theta_2^{(T)}\|\right] = O\left(\min\left\{r, L\sqrt{\sum_{t=1}^{T}\eta_t^2} + L\frac{\sum_{t=1}^{T}\eta_t}{n}\right\}\right).$$

*The upper bound of GD is the same.*

The following theorem presents the lower bound of UAS. For simplicity, we consider the case of constant learning rate, i.e., $\eta_t = \eta$ for $t = 1, ..., T$.

**Theorem 1.** *Assume $\theta \in B_2(0, r)$. There exist some $\epsilon > 0$ and some loss function $l(f_\theta(x), y)$ which is differentiable and convex w.r.t. $\theta$, such that $\theta_1^{(t)}$ and $\theta_2^{(t)}$, which are SGD-based adversarial training estimators obtained using $S_1, S_2$ respectively under attack strength $\epsilon$, satisfies that*

$$\sup_{S_1 \sim S_2} \mathbb{E}\|\theta_1^{(T)} - \theta_2^{(T)}\| = \Omega\left(\min\left\{1, \frac{T}{n}\right\}\eta\sqrt{T} + \frac{\eta T}{n}\right).$$

*For GD, the lower bound is*

$$\sup_{S_1 \sim S_2} \|\theta_1^{(T)} - \theta_2^{(T)}\| = \Omega\left(\min\left\{1, \eta\sqrt{T} + \frac{\eta T}{n}\right\}\right).$$

To prove Theorem 1, similar to Bassily et al. (2020), we design a smooth clean loss function with two datasets $S_1 \sim S_2$ and study the exact change of the model parameters. The detailed proof is postponed to the Appendix D.

As discussed by Bassily et al. (2020), the non-smoothness of the loss is the main reason for poor stability. The presented low bounds match the result of Bassily et al. (2020), but it is worth mentioning that the two results are not directly comparable since Bassily et al. (2020) studied the UAS of clean training when the loss function $l(f_\theta(x), y)$ is non-smooth, while our work studies the UAS of adversarial training when the loss function is smooth. On the other hand, the UAS of cleaning training under smooth loss, implied by Theorem 3.8 of Hardt et al. (2016), is of order $O(\min\{r, L\sum_{t=t_0}^{T}\eta_t/n\})$. Therefore, we conclude that adversarial training has a worse stability than its natural counterpart.

To ensure the convergence of optimization (i.e., $\eta T$ is not so small) and the generalization performance (Proposition 1 and Theorem 1), one may take $T = n^2$ and $\eta = 1/n^{3/2}$. The resulting optimization error and stability then become $O(1/\sqrt{n})$, which matches the minimax lower bound of excess risk (Chen et al., 2018). However, such a choice of $(\eta, T)$ is impractical and needs to improve (refer to the discussion in Bassily et al., 2020).

### 3.3 The role of numerical attack error

In real-world applications, calculating the exact attack $A_\epsilon$ for general models is not easy, and usually, a numerical approximation $A'_\epsilon$ (e.g., by FGM or PGD) is used in the adversarial training algorithm.

Some recent literature start to aware the important impact of the numerical attack error (i.e., the difference between $A_\epsilon$ and $A'_\epsilon$). For example, Gao et al. (2019); Zhang et al. (2020) took account of the attack approximation method in the convergence analysis of adversarial training, and Deng et al. (2020) studied the convergence of PGD attack.

For algorithmic stability, extended from Proposition 1, the following result considers the effect of attack error. Comparing the upper bounds of Proposition 1 and Corollary 1, it suggests one to control the attack error carefully in the adversarial training.

**Corollary 1.** *Under the same conditions of Proposition 1, assume the algorithm uses an approximation $A'_\epsilon$ instead of the exact attack $A_\epsilon$ with attack error $\min \|A_\epsilon(x,y,w) - A'(x,y,w)\| \leq \Delta\varepsilon$ for any $(x,y,w)$, where the minimum is taken when the exact attack (i.e., (1)) is not unique. Assume $\nabla_\theta l(f_\theta(x),y)$ is $\kappa$-Lipschitz w.r.t. $x$. Then, for SGD*

$$\sup_{S_1 \sim S_2} \mathbb{E}\|\theta_1^{(T)} - \theta_2^{(T)}\| = O\left(\min\left\{r, L\sqrt{\sum_{t=1}^T \eta_t^2} + L\frac{\sum_{t=1}^T \eta_t}{n} + \kappa\Delta\varepsilon\sum_{t=1}^T \eta_t\right\}\right).$$

*The upper bound of GD is the same.*

Besides the convex case, some discussions for non-convex case can be found in Appendix A. The observations are similar.

# 4 Improving the stability

In this section, we show that injecting noise in adversarial training enhances the smoothness of adversarial loss, and hence improves the stability of adversarial training.

## 4.1 Source of non-smoothness

As mentioned after Theorem 1, the non-smoothness issue in adversarial training is the main cause of the poor stability. Summarizing from the related works, we identify two important sources of non-smoothness in adversarial training even when the standard loss is smooth: (1) when the data are overfitted, i.e., the training loss is almost 0 and $\nabla_{x_i} l(f_\theta(x_i), y_i) \approx 0$, the adversary has no preference on the attack direction at $x_i$, and the numerical estimation of $A_\epsilon$ is not stable, which possibly leads to an unstable update iteration of adversarial training; (2) there exists a certain set of $\theta$, such that the adversarial training loss is always non-differentiable regardless of the training data, even when its natural counterpart is smooth. For example, in linear regression, when $\theta_t$ is closed to the null model, the non-smoothness issue occurs (Xing et al., 2021a).

To tackle both non-smooth issues, we propose incorporating noise injection in the training process as described in the following section.

## 4.2 Injecting noise during training

In this section, we present the noise injection algorithm in adversarial training and provide some theoretical justifications.

Algorithm 1 below illustrates the details of the noise injection method. The basic idea behind this is that: the non-smoothness of adversarial loss occurs only when $\theta$ and $x_i$' belong to a certain special region (e.g., in linear regression, when $\theta$ is closed to either zero or when $\theta^\top x_i \equiv y_i$), thus injecting some small noise to both $\theta$ and $x$ helps them to escape from such region where non-smoothness occurs, which further ensures the Lipschitz continuity property.

**Remark 1.** *The Gaussian noise in Algorithm 1 is for proof simplicity. In general, it can be changed to other noise distributions if the tail is not heavy.*

In the literature, there have been some applications of noise injection. For example, He et al. (2019) considered injecting noise to the weights as a regularization method to improve the adversarial robustness. Besides literature in supervised learning (Weng et al. (2018); Wang et al. (2018); Ford et al. (2019), injecting noise in data was also considered to stabilize the training process of GAN (Arjovsky and Bottou, 2017; Jenni and Favaro, 2019).

**Algorithm 1** Add noise to weight and data

**Input:** data $\{(x_i, y_i)\}_{i=1}^n$, number of iterations $T$, learning rate $\{\eta_t\}_{t=1}^T$, attack strength $\epsilon$, noise size $(\xi_\theta, \xi_x)$, scale parameter $r$ , initialization $\theta^{(0)}$.
**for** $t = 1$ **to** $T$ **do**
  Add Gaussian noise with variance $\xi_\theta^2$ to each dimension of $\theta_t$ to form $\widetilde{\theta}^{(t)}$, and add Gaussian noise with variance $\xi_x^2$ to each dimension of $x_t$ to obtain $\widetilde{x}_{i_t}$.
  Calculate the attack (based on $\widetilde{x}_{i_t}$ and $\widetilde{\theta}^{(t)}$) as $\widehat{z}_{i_t}$.
  Take gradient w.r.t $\widetilde{\theta}^{(t)}$ on $l(f_{\widetilde{\theta}^{(t)}}(\widehat{z}_{i_t}), y_{i_t})$.
  Update $\theta^{(t)}$ to $\theta^{(t+1)}$ with rate $\eta_t$.
  Project $\theta^{(t+1)}$ onto $B_2(0, r)$.
**end for**
**Output:** $\theta^{(T)}$.

In the following theorems, we provide a theoretical justification for the stability and optimization when injecting noise into adversarial training for the following models:

- Linear regression: $l(f_\theta(x), y) = (x^\top\theta - y)^2$.
- Logistic regression: $l(f_\theta(x), y) = -\log^{1\{y=1\}}(p) - \log^{1\{y=-1\}}(1-p)$, where $p = p(x^\top\theta) = 1/(1+e^{-x^\top\theta})$.
- Smooth hinge loss: the hinge loss $\max\{0, 1-y(x^\top\theta)\}$ is not smooth at 0, hence is approximated by $l(f_\theta(x), y) = (1-y(x^\top\theta))H((1-y(x^\top\theta))/h)$, where $h > 0$ is a bandwidth parameter, and $H$ is a smooth approximation of the indicator function $I\{x \geq 0\}$. The detailed conditions on $H$ are postponed to Lemma 7 in the Appendix D.2.

The following assumption is imposed on the data:

**Assumption 1.** *The independent variable $x \in \mathbb{R}^d$ follows multivariate Gaussian distribution with zero-mean and $\Sigma$ whose eigenvalues are bounded and away from zero.*

*For regression, $\mathbb{E}|y|$ and $\mathbb{E}\|yx\|$ are finite. For some constant $C > 0$, any $\theta \in B_2(0, Cr)$ satisfies $P(|x^\top\theta - y| \in [\zeta_1 r, \zeta_2 r]) = O(\zeta_2 - \zeta_1)$ for $\zeta_1, \zeta_2 > 0$.*

*For classification, the label is $y \in \{\pm 1\}$. The upper bound $r$ satisfies $r/\max_{i=1,\ldots,n}\|x_i\| \to 0$.*

The Gaussian assumption in $x$ is merely for derivation simplicity. The assumptions w.r.t. regression avoids $|x^\top\theta - y|$ from clustering around zero when $\|\theta\|/r$ approaches zero. A linear model $\mathbb{E}[y|x] = \theta_0^\top x$ with Gaussian noise satisfies Assumption 1.

Given the above problems and data generating models, the following lemma studies the smoothness (i.e., the Lipschitz constant) of $\nabla_\theta l(f_\theta(x), y)$, and of the gradient of noise injected adversarial loss.

**Lemma 1** (Informal Statement for Lemma 3)**.** *Assume Assumption 1 holds. Denote $L$ as the Lipschitz constant of $l(f_\theta(x), y_j)$ w.r.t. $\theta$ for any $x \in B_2(x_j, 2\epsilon)$ and all $1 \leq j \leq n$. Then, in probability, $L$ is bounded by some finite $L^*$. Take the noise injected in data as zero-mean Gaussian with variance $(\xi_0^2/d)I_d$, and the noise injected in parameters is zero-mean Gaussian with variance $(\xi^2/d)I_d$ where $\xi = \xi_0 L^*$. Denote $E(\theta+\delta, \widetilde{x}, y)$ as the event that $\nabla_\theta l(f_{\theta+\delta}(\widetilde{x} + A_\epsilon(f_{\theta+\delta}, \widetilde{x}, y)), y)$ is $B^*/\zeta$-Lipschitz for some $B^* > 0$. There exists some choice of $(\xi, \zeta) \to 0$ such that with probability tending to one over the generation of $S$, uniformly for all $\theta \in B_2(0, r)$,*

$$P\left(E^c(\theta+\delta, \widetilde{x}, y) | (x, y) \in S\right) = o(1).$$

*The formal statement is postponed to Lemma 3 in the appendix.*

Let $P(E^c|S) := \sup_{\theta \in B_2(0,r),(x,y)} P\left(E^c(\theta+\delta, \widetilde{x}, y)|(x, y) \in S\right)$ in what follows, for notation simplicity.

**Remark 2.** *The terms $r$, $L$ are generic representations. For different loss functions and data dimension $d$, their values may change. In addition, the exact rate of $P\left(E^c|S\right)$ is affected by the value of $r, L$ as well as $\xi, \zeta$. We postpone the details to Appendix D.2 during the proof.*

The following lemma is an intermediate step in the derivation of Theorem 2 below, and reveals the important role played by $B^*/\zeta$. Note that since Lemma 1 holds over the randomness of $S$, instead of **uniform** argument stability, we turn to a bound similar to **hypothesis** stability (Bousquet and Elisseeff, 2002) for the following results. To simplify the representation, the values of $(B^*, L^*, r)$ are treated as constants in the following main text.

**Lemma 2.** *Under the same conditions as in Lemma 1, uniformly for all $i = 1, ..., n$, with probability tending to one over the generation of $S_1 \sim S_2$ where the $i$-th sample is replaced, for both GD and SGD, given $\|\theta_1^{(t-1)} - \theta_2^{(t-1)}\| = \Delta_{t-1}$, it follows that*

$$\mathbb{E}[\|\theta_1^{(t)} - \theta_2^{(t)}\|^2 | S_1, S_2, \Delta_{t-1}] \leq \left(1 + 2\eta_t^2 \frac{(B^*)^2}{\zeta^2} 1\{\eta_t \geq \frac{\zeta}{B^*}\}\right) \Delta_{t-1}^2 + reminder,$$

*where the detail of reminder term can be found in (8) in Appendix D.2. Note that the expectation taken on $\|\theta_1^{(t)} - \theta_1^{(t)}\|^2$ in GD is over the injected random noise, and the one for SGD is taken for both the sampling in SGD and the injected random noise.*

Lemma 2 illustrates the relationship between $\|\theta_1^{(t)} - \theta_2^{(t)}\|^2$ and $\Delta_{t-1}^2$. Recall that $B^*/\zeta$ is the Lipschitz constant of $\nabla_\theta l(f_{\theta+\delta}(\tilde{x} + A), y)$. When $\eta_t \geq \zeta/B^*$, a larger Lipschitz constant implies a larger upper bound of stability. When taking $\eta_t < \zeta/B^*$, we have the following result:

**Theorem 2.** *Under the same conditions as in Lemma 1, when taking $\eta_t \leq \zeta/B^*$, for both GD and SGD, with probability tending to one (where the probability refers to the generation measure of the $n + 1$ distinct independent samples in $S_1 \sim S_2$),*

$$\begin{aligned}\mathbb{E}[\|\theta_1^{(T)} - \theta_2^{(T)}\| | S_1, S_2] &= O\left(\left[\sqrt{P(E^c|S_1) + P(E^c|S_2)} + \sqrt{\frac{1}{n}}\right] \sqrt{\sum_{t=t_0}^{T} \eta_t^2}\right) \\ &\quad + O\left(\left[\Delta\varepsilon + \frac{1}{n} + P(E^c|S_1) + P(E^c|S_2)\right] \sum_{t=t_0}^{T} \eta_t\right).\end{aligned}$$

Furthermore, extending from Lemma 9 of Bousquet and Elisseeff (2002), the generalization gap is upper bounded using hypothesis stability as follows.

**Proposition 2.** *Assume $\theta \in B_2(0, r)$. Denote $\widehat{\theta}(S)$ as the model obtained based on dataset $S$ using some algorithm. Assume $l(f_\theta(x), y) \in [0, M]$ when $\|x\| \leq \sqrt{d \log n}$, we have for any $i = 1, ..., n$,*

$$\begin{aligned}\mathbb{E}\left[\left(R(\widehat{\theta}(S_1)) - R_{S_1}(\widehat{\theta}(S_1))\right)^2\right] &\leq \frac{M^2}{2n} + 4\mathbb{E}\left[\sup_{\theta \in B_2(0,r)} l^2(f_\theta(x + A_\epsilon), y) 1\{\|x\| \geq \sqrt{d \log n}\}\right] \\ &\quad + 3M\mathbb{E}\left[\left|l(f_{\widehat{\theta}(S_1)}(x_i + A_\epsilon), y_i) - l(f_{\widehat{\theta}(S_2^i)}(x_i + A_\epsilon), y_i)\right|\right],\end{aligned} \tag{4}$$

*where $S_2^i$ represents the neighboring dataset of $S_1$ whose $i$th sample is replaced. The notion $A_\epsilon$ is an abbreviation of the attack $A_\epsilon(f, x, y)$ or $A_\epsilon(f, x_i, y_i)$.*

Note that the last term on the RHS of (4) can be bounded according to the result of Theorem 2, under Lipschitz condition of the adversarial loss $l(f_\theta(x_i + A_\epsilon), y_j)$. The second term on the RHS of (4) can be bounded given some further conditions on the tail behavior of loss function.

Compared with the generalization upper bound obtained in Bousquet and Elisseeff (2002), in Proposition 2, there is an extra term corresponding to $\|x\| > \sqrt{d \log n}$. In Proposition 2, we only assume $l(f_\theta(x), y) \in [0, M]$ when $\|x\| \leq \sqrt{d \log n}$, which is weaker than the uniform bounded assumption in Bousquet and Elisseeff (2002).

Besides, we also establish the optimization error bound. The following theorem presents the convergence of noise-injected adversarial training when $\eta_t \equiv \eta$. For the proof, one can refer to the Appendix D.2.

**Theorem 3.** *Under the same conditions as in Lemma 1, for both GD and SGD,* $\bar{\theta} := \operatorname{argmin}_{\theta \in B_2(0,r)} R_S(\theta)$*, when* $\eta_t \equiv \eta$*,*

$$\mathbb{E}\left[\min_{t=1,\ldots,T} R_S(\theta^{(t)}) - \min_{\theta \in B_2(0,r)} R_S(\theta) \middle| S\right] \leq \frac{\mathbb{E}\|\theta^{(0)} - \bar{\theta}\|^2}{2\eta T} - \frac{\mathbb{E}[\|\theta^{(T)} - \bar{\theta}\|^2 | S]}{2\eta T} + \frac{\eta (L^*)^2}{2} + O(L^* \xi) + O(r\Delta\varepsilon).$$

Theorem 3 presents the convergence of adversarial training loss throughout the training. Under the boundedness of $\theta$ and $\xi = \xi_0 L^*$, when $(\eta, T, \xi_0)$ is chosen properly (e.g. $(\eta T) \to \infty$ and $\xi_0 \to 0$) and $\Delta\epsilon \to 0$, the adversarial training loss converges to its minimal asymptotically.

It is noteworthy that the general design of Algorithm 1 does not specify the noise distribution.

While in Section 4.3, we use Gaussian noise to justify our theorems under linear regression empirically, different forms of noise can be utilized for complex models, refer to our experiments on deep neural networks in Section C.2.

**Remark 3.** *If an intercept term exists in the loss function, e.g.,* $l = (x^\top \theta + b - y)^2$ *for linear regression, the analysis is similar to Lemma 1, leading to the same final conclusions as in Theorem 2 and 3.*

**Remark 4** ($\mathcal{L}_\infty$ Attack in Adversarial Training)**.** *In general, the stability of* $\mathcal{L}_\infty$ *adversarial training is worse. To set an example, we consider the linear regression setup. For* $\mathcal{L}_2$ *attack, the gradient of adversarial loss is not Lipschitz only when* $\theta$ *approaches zero or* $\theta^\top x$ *is closed to* $y$*. Under* $\mathcal{L}_\infty$ *attack, the adversarial loss becomes*

$$(x^\top \theta - y)^2 + \epsilon^2 \|\theta\|_1^2 + 2\epsilon \|\theta\|_1 |x^\top \theta - y|,$$

*indicating there is a much larger set where the* $\mathcal{L}_\infty$ *adversarial loss is not smooth.*

*Noise injection is still helpful to remedy the non-smooth issue for* $\mathcal{L}_\infty$ *adversarial training and leads to results similar to Theorems 2 and 3. However, one will derive a worse upper bound for the stability and optimization error. Refer to Appendix E for more detailed arguments.*

### 4.3 Numerical illustration

We use simulation to illustrate how noise-injected adversarial training affects performance. In short, the quality of the updating gradient is better after injecting noise. The Lipschitz constant of $\nabla_\theta l(f_\theta(x + A), y)$ in Lemma 2 is smaller.

We consider linear regression problem in this experiment. The data is generated using $y = x^\top \theta^* + \delta$ with $x \sim N(0, I_d)$ with $d = 10$ and $\delta \sim N(0, \sigma^2)$. The coefficient $\theta_0$ is taken as $\theta_i^* = 1/\sqrt{d}$ for $i = 1, ..., d$. The variance of noise is taken as $\sigma^2 = 4$ and attack strength is $\epsilon = 2$. We randomly generated $n = 1000$ samples.

To train the regression model, we train $T = 500$ epochs with learning rate $\eta = 0.01$ and initialization $\theta^{(0)} = \mathbf{0}$. In each iteration, we calculate $\|\nabla f_t(\theta^{(t)}) - \nabla f_{t-1}(\theta^{(t-1)})\| / \|\theta^{(t)} - \theta^{(t-1)}\|$ as an approximation for the Lipschitz constant of the gradient, where $\nabla f_t(\theta)$ is the averaged gradient of adversairal loss for the $t$th batch of data $S_t$. Based on Lemma 2, a larger Lipschitz constant ($B^*/\zeta$) indicates a worse stability, which is the right tail of the histogram. The results are summarized the histograms in Figure 1.

From the left three histograms in Figure 1, one can see that injecting noise on parameters and data (where we set $\xi_x = \xi_\theta = \xi$) leads to a smaller distribution of Lipschitz constant for $\nabla f_t(\theta^{(t)})$, in terms of right tail percentile. For the right two histograms in Figure 1, a smaller batch size implies a heavier tail in the distribution of Lipschitz constant due to larger stochastic noise in estimating $\nabla f_t(\theta^{(t)})$.

Besides the experiment showing how noise injection affects the training process, we also conduct a simulation to illustrate the effect of attack error on the generalization. Due to the space limit, the simulation is postponed to the Appendix B. To briefly summarize the observations, for all scenarios we consider, when there is an error when calculating the attack, the generalization gap becomes larger.

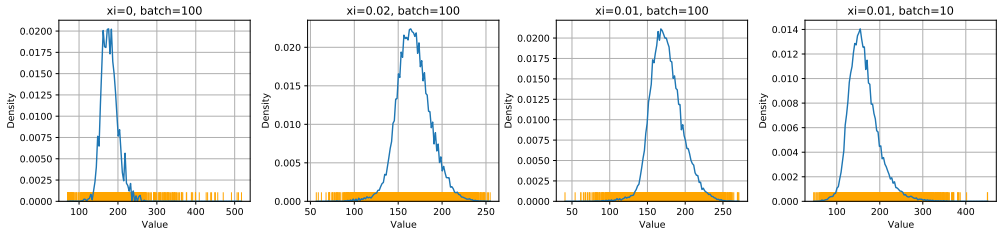


Figure 1: Density of $\|\nabla f_t(\theta^{(t)}) - \nabla f_{t-1}(\theta^{(t-1)})\|/\|\theta^{(t)} - \theta^{(t-1)}\|$. $\theta_{0,i} = 1/\sqrt{d}$ for $i = 1, ..., d$ with $d = 10$. $n = 1000$, $\sigma = 2$, $\epsilon = 2$. $\eta = 0.01$, $T = 500$. Vanishing initialization. A larger $\xi$ implies a smaller Lipschitz gradient of $\nabla f_t(\theta^{(t)})$. The (mean, sd, 99.9%-quantile) are (178.90, 22.34, 323.36), (167.72, 19.20, 234.43), (174.95, 20.71, 247.93), (163.05, 34.89, 337.07) for the above four histograms.

## 5 Exploration in Neural Networks

While our main contributions are for statistical models, we also provide some theoretical results and numerical experiments associated with neural networks.

It is still an open question how to connect existing algorithmic stability tools to neural networks. Since the number of parameters in neural networks is much larger than simple models, a simple bound on $\|\theta_1^{(T)} - \theta_2^{(T)}\|$ is not useful. Instead, we consider a two-layer neural network with lazy training and vanishing initialization in regression and provide a stability bound directly on the loss. In order to track the neural network parameters, we track both the convergence and the stability together. This is more restrictive than simple models.

The following (informal) statement presents the stability of two-layer nonlinear networks with lazying training in adversarial training setup. The formal statement of the theorem is postponed to Appendix C.1. Based on the following theorem, under proper configurations, the noise-injected training in neural networks improve the stability:

**Theorem 4** (Informal Statement)**.** *For two-layer nonlinear (including ReLU) networks, with proper initialization and training configurations, training only on the hidden layer with proper noise injection, it satisfies that*

$$\begin{aligned} & \mathbb{E}_{S_1 \sim S_2} \left| l(f_{\theta_1^{(T)}}[x + A_\epsilon(f_{\theta_1^{(T)}}, x_i, y_i)], y_i) - l(f_{\theta_2^{(T)}}[x + A_\epsilon(f_{\theta_2^{(T)}}, x_i, y_i)], y_i) \right| \\ = \quad & O\left( \left[ L\sqrt{P(E^c)} + \sqrt{\frac{L^2}{n}} \right] \eta\sqrt{T} + \left[ \frac{L}{n} + LP(E^c) \right] \eta T \right) + rem, \end{aligned}$$

*where $rem = o(1)$ and is not the dominant term.*

There are two differences between Theorem 4 and the results in simple models. First, it is not useful to directly assume the weights of the neural network parameters within a large ball and put this large number into the bound, thus we simultaneously study the convergence and stability of the neural network to tighten the stability bound. Second, instead of bounding the stability of the parameters, we turn to bound the stability of the loss, which is more meaningful to this over-parameterized method.

Besides the results in two-layer networks, we also numerically study the generalization gap using deep neural networks with CIFAR10 dataset. Due to space limit, we postpone the experiments to Appendix C.2. The observations from numerical experiments are (1) injecting noise can reduce the generalization gap between training and testing performance, and (2) improving the accuracy of attack also improves the quality of the adversarial training. Both of the observations are similar to those in simulations.

## 6 Conclusion

In this paper, we evaluate the algorithmic stability of the adversarial training method. Based on the lower bound and upper bound of UAS, we reveal that the naive adversarial training is not as stable

as its natural counterpart. To improve the stability, we argue that it is helpful to inject noise into model parameters and input data. Our theory verifies the effectiveness of noise injection in some simple models. Besides, our theory also considers the effect of attack error and indicates that the upper bound of UAS is smaller when the attack error is smaller.

The above theoretical investigations emphasize the usage of noise injection and controlling numerical attack error during the adversarial training. These theoretical insights are well validated by our simulations under simple regression models.

There are two future research directions motivated by this study. Although we observe a similar phenomenon in deep neural networks as our theory in simple models, there is a gap between the exact algorithmic stability of deep neural networks and the UAS bounds in simple models. Our analysis in the two-layer neural networks is a trail in this area, but a more comprehensive study study on algorithmic stability of the deep neural network is wanted. Second, as we mentioned in the numerical experiments, a wider neural network has a poor attack. Corresponding theoretical explanation is also an interesting topic.

## 7 Acknowledgements

This project is partially supported by NSF-SCALE MoDL(2134209) and NSF-DMS-1811812.

## References

Allen-Zhu, Z. and Li, Y. (2020), "Feature Purification: How Adversarial Training Performs Robust Deep Learning," *arXiv preprint arXiv:2005.10190*.

Arjovsky, M. and Bottou, L. (2017), "Towards principled methods for training generative adversarial networks," *arXiv preprint arXiv:1701.04862*.

Awasthi, P., Frank, N., and Mohri, M. (2020), "Adversarial learning guarantees for linear hypotheses and neural networks," in *International Conference on Machine Learning*, PMLR, pp. 431–441.

Bassily, R., Feldman, V., Guzmán, C., and Talwar, K. (2020), "Stability of Stochastic Gradient Descent on Nonsmooth Convex Losses," *arXiv preprint arXiv:2006.06914*.

Biggio, B., Corona, I., Maiorca, D., Nelson, B., Šrndić, N., Laskov, P., Giacinto, G., and Roli, F. (2013), "Evasion attacks against machine learning at test time," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, Springer, pp. 387–402.

Bousquet, O. and Elisseeff, A. (2001), "Algorithmic stability and generalization performance," in *Advances in Neural Information Processing Systems*, pp. 196–202.

— (2002), "Stability and generalization," *Journal of machine learning research*, 2, 499–526.

Charles, Z. and Papailiopoulos, D. (2018), "Stability and generalization of learning algorithms that converge to global optima," in *International Conference on Machine Learning*, pp. 745–754.

Chen, Y., Jin, C., and Yu, B. (2018), "Stability and convergence trade-off of iterative optimization algorithms," *arXiv preprint arXiv:1804.01619*.

Deng, Z., He, H., Huang, J., and Su, W. J. (2020), "Towards Understanding the Dynamics of the First-Order Adversaries," .

Ding, G. W., Wang, L., and Jin, X. (2019), "AdverTorch v0.1: An Adversarial Robustness Toolbox based on PyTorch," *arXiv preprint arXiv:1902.07623*.

Farnia, F. and Ozdaglar, A. (2020), "Train simultaneously, generalize better: Stability of gradient-based minimax learners," *arXiv preprint arXiv:2010.12561*.

Ford, N., Gilmer, J., Carlini, N., and Cubuk, D. (2019), "Adversarial examples are a natural consequence of test error in noise," *arXiv preprint arXiv:1901.10513*.

Gao, R., Cai, T., Li, H., Wang, L., Hsieh, C.-J., and Lee, J. D. (2019), "Convergence of adversarial training in overparametrized networks," *arXiv preprint arXiv:1906.07916*.

Hardt, M., Recht, B., and Singer, Y. (2016), "Train faster, generalize better: Stability of stochastic gradient descent," in *International Conference on Machine Learning*, pp. 1225–1234.

He, Z., Rakin, A. S., and Fan, D. (2019), "Parametric noise injection: Trainable randomness to improve deep neural network robustness against adversarial attack," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 588–597.

Ho, N., Khamaru, K., Dwivedi, R., Wainwright, M. J., Jordan, M. I., and Yu, B. (2020), "Instability, Computational Efficiency and Statistical Accuracy," *arXiv preprint arXiv:2005.11411*.

Jalal, A., Ilyas, A., Daskalakis, C., and Dimakis, A. G. (2017), "The robust manifold defense: Adversarial training using generative models," .

Jenni, S. and Favaro, P. (2019), "On stabilizing generative adversarial training with noise," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 12145–12153.

Kearns, M. and Ron, D. (1999), "Algorithmic stability and sanity-check bounds for leave-one-out cross-validation," *Neural computation*, 11, 1427–1453.

Khim, J. and Loh, P.-L. (2018), "Adversarial risk bounds via function transformation," *arXiv preprint arXiv:1810.09519*.

Kuzborskij, I. and Lampert, C. (2018), "Data-dependent stability of stochastic gradient descent," in *International Conference on Machine Learning*, pp. 2815–2824.

Lee, K. and Chandrakasan, A. P. (2020), "Rethinking Empirical Evaluation of Adversarial Robustness Using First-Order Attack Methods," *arXiv preprint arXiv:2006.01304*.

Lei, Y. and Ying, Y. (2020), "Fine-Grained Analysis of Stability and Generalization for Stochastic Gradient Descent," *arXiv preprint arXiv:2006.08157*.

Ma, S. and Liu, Y. (2019), "Nic: Detecting adversarial samples with neural network invariant checking," in *Proceedings of the 26th Network and Distributed System Security Symposium (NDSS 2019)*.

Madden, L., Dall'Anese, E., and Becker, S. (2020), "High probability convergence and uniform stability bounds for nonconvex stochastic gradient descent," *arXiv preprint arXiv:2006.05610*.

Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. (2017), "Towards deep learning models resistant to adversarial attacks," *arXiv preprint arXiv:1706.06083*.

Moosavi-Dezfooli, S.-M., Fawzi, A., and Frossard, P. (2016), "Deepfool: a simple and accurate method to fool deep neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2574–2582.

Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, Z. B., and Swami, A. (2016a), "The limitations of deep learning in adversarial settings," in *Security and Privacy (EuroS&P), 2016 IEEE European Symposium on*, IEEE, pp. 372–387.

Papernot, N., McDaniel, P., Swami, A., and Harang, R. (2016b), "Crafting adversarial input sequences for recurrent neural networks," in *Military Communications Conference, MILCOM 2016-2016 IEEE*, IEEE, pp. 49–54.

Pinot, R., Meunier, L., Yger, F., Gouy-Pailler, C., Chevaleyre, Y., and Atif, J. (2021), "On the robustness of randomized classifiers to adversarial examples," *arXiv preprint arXiv:2102.10875*.

Ramezani-Kebrya, A., Khisti, A., and Liang, B. (2018), "On the Stability and Convergence of Stochastic Gradient Descent with Momentum," *arXiv preprint arXiv:1809.04564*.

Rice, L., Wong, E., and Kolter, J. Z. (2020), "Overfitting in adversarially robust deep learning," *arXiv preprint arXiv:2002.11569*.

Schmidt, L., Santurkar, S., Tsipras, D., Talwar, K., and Madry, A. (2018), "Adversarially robust generalization requires more data," in *Advances in Neural Information Processing Systems*, pp. 5014–5026.

Shaham, U., Yamada, Y., and Negahban, S. (2015), "Understanding adversarial training: Increasing local stability of neural nets through robust optimization," *arXiv preprint arXiv:1511.05432*.

Shorten, C. and Khoshgoftaar, T. M. (2019), "A survey on image data augmentation for deep learning," *Journal of Big Data*, 6, 1–48.

Sinha, A., Namkoong, H., and Duchi, J. (2018), "Certifying some distributional robustness with principled adversarial training," .

Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. J., and Fergus, R. (2014), "Intriguing properties of neural networks," in *2nd International Conference on Learning Representations*.

Tao, G., Ma, S., Liu, Y., and Zhang, X. (2018), "Attacks meet interpretability: Attribute-steered detection of adversarial samples," in *Advances in Neural Information Processing Systems*, pp. 7717–7728.

Wang, B., Yuan, B., Shi, Z., and Osher, S. J. (2018), "Resnets ensemble via the feynman-kac formalism to improve natural and robust accuracies," *arXiv preprint arXiv:1811.10745*.

Wang, Y., Ma, X., Bailey, J., Yi, J., Zhou, B., and Gu, Q. (2019a), "On the convergence and robustness of adversarial training," in *International Conference on Machine Learning*, pp. 6586–6595.

Wang, Y., Zou, D., Yi, J., Bailey, J., Ma, X., and Gu, Q. (2019b), "Improving adversarial robustness requires revisiting misclassified examples," in *International Conference on Learning Representations*.

Weng, T.-W., Zhang, H., Chen, P.-Y., Yi, J., Su, D., Gao, Y., Hsieh, C.-J., and Daniel, L. (2018), "Evaluating the robustness of neural networks: An extreme value theory approach," *arXiv preprint arXiv:1801.10578*.

Wu, D., Wang, Y., and Xia, S.-t. (2020), "Adversarial Weight Perturbation Helps Robust Generalization," *arXiv preprint arXiv:2004.05884*.

Xie, C., Tan, M., Gong, B., Yuille, A., and Le, Q. V. (2020), "Smooth Adversarial Training," *arXiv preprint arXiv:2006.14536*.

Xing, Y., Song, Q., and Cheng, G. (2021a), "On the generalization properties of adversarial training," in *International Conference on Artificial Intelligence and Statistics*, PMLR, pp. 505–513.

Xing, Y., Zhang, R., and Cheng, G. (2021b), "Adversarially Robust Estimate and Risk Analysis in Linear Regression," in *International Conference on Artificial Intelligence and Statistics*, PMLR, pp. 514–522.

Yin, D., Ramchandran, K., and Bartlett, P. (2018), "Rademacher complexity for adversarially robust generalization," *arXiv preprint arXiv:1810.11914*.

Zhai, R., Cai, T., He, D., Dan, C., He, K., Hopcroft, J., and Wang, L. (2019), "Adversarially robust generalization just requires more unlabeled data," *arXiv preprint arXiv:1906.00555*.

Zhang, H., Yu, Y., Jiao, J., Xing, E. P., Ghaoui, L. E., and Jordan, M. I. (2019), "Theoretically Principled Trade-off between Robustness and Accuracy," in *Proceedings of the 36th International Conference on Machine Learning*, PMLR, vol. 97 of *Proceedings of Machine Learning Research*, pp. 7472–7482.

Zhang, Y., Plevrakis, O., Du, S. S., Li, X., Song, Z., and Arora, S. (2020), "Over-parameterized Adversarial Training: An Analysis Overcoming the Curse of Dimensionality," *arXiv preprint arXiv:2002.06668*.

Zheltonozhskii, E., Baskin, C., Nemcovsky, Y., Chmiel, B., Mendelson, A., and Bronstein, A. M. (2020), "Colored Noise Injection for Training Adversarially Robust Neural Networks," *arXiv preprint arXiv:2003.02188*.

Zhou, Y., Liang, Y., and Zhang, H. (2018), "Generalization error bounds with probabilistic guarantee for sgd in nonconvex optimization," *arXiv preprint arXiv:1802.06903*.

# 8 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. We mention some limitations when discussing potential extensions in Conclusion.*
   (c) Did you discuss any potential negative societal impacts of your work? *No. This paper studies theories and only uses CIFAR-10 for real-data experiment.*
   (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)? *No. We mention in the main text that we are using some implementation from other papers shared in Github.*
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? *Yes.*
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? *Yes.*
   (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)? *No. From the design of the algorithms, they are similar to the original adversarial training algorithm, so the computation complexity does not change a lot.*

The appendix is arranged as follows: Section A presents the analysis on non-convex loss. Some additional simulation results are presented in Section B. Section C.1 presents some theoretical results in neural networks, and Section C.2 is a numerical study using CIFAR10. Section D is a collection of proofs for theorems. Section E displays some theoretical results w.r.t $\mathcal{L}_\infty$ attack in regression task.

# A General loss

The following theorem provides a more general upper bound for the algorithm stability beyond convex $l(f_\theta(x, y))$.

**Theorem 5.** *Assume $\|\theta\|$, $\|x\|$ are both bounded, and we choose $\eta_t \leq c/t$ for some $c > 0$. If , If $l(f_\theta(x, y))$ is L-Lipschitz w.r.t. $\theta$, $(\nabla_\theta l(f_\theta(x), y), \nabla_x l(f_\theta(x), y))$ is $\kappa$-Lipschitz in $(\theta, x)$, and the attack error is smaller than $\Delta\varepsilon$, then the UAS upper bound for SGD becomes*

$$\sup_{S_1 \sim S_2} \mathbb{E}\|\theta_1^{(T)} - \theta_2^{(T)}\| = O\left((\kappa\Delta\varepsilon + L)\left(\frac{T}{n}\right)^{\frac{c\kappa}{c\kappa+1}}\right),$$

*and the one for GD is*

$$\sup_{S_1 \sim S_2} \|\theta_1^{(T)} - \theta_2^{(T)}\| = O\left((\kappa\Delta\varepsilon + L)\, T^{c\kappa}\right).$$

By Theorem 5, to obtain better stability, it suffices to control the number of iterations and the attack error. To ensure a diminishing UAS upper bound, the choice of $T$ should be much smaller than $n$. Note that SGD only has a small probability of encountering the exact different data points between the two datasets at the very first iteration, leading to a smaller UAS upper bound than GD due to the diminishing learning rate.

**Remark 5.** *Compared with the UAS bound $O(T^{\frac{c\kappa}{c\kappa+1}}/n)$ for standard training (Hardt et al., 2016), Theorem 5 also suggests that adversarial training prefers a smaller number of steps to reduce the corresponding UAS upper bound. This echos the observations in Rice et al. (2020) that early stopping is necessary in adversarial training.*

**Remark 6.** *The class of functions considered in Proposition 1 (smooth convex function) is not a special case of those in Theorem 5, so results of Proposition 1 and Theorem 5 are not directly comparable.*

# B Additional Experimental Results

Below is the simulation study mentioned in Section 4.3 on the effect of attack error.

In this experiment, we follow the data generating model as for Figure 1 and take $\epsilon = 0.2$. Unlike noise injection which changes the Lipschitz constant $B^*/\zeta$, from Lemma 2 and Theorem 2, attack error has little effect on the Lipschitz constant, and directly deteriorates the stability. Consequently, instead of showing the Lipschitz constant (Figure 1), we present the generalization error.

To create error in the attack, for each data, after we obtain the true attack using analytical solution, denoting as $\widehat{z}_i$, we create random noise $\delta_{z_i} \sim N(0, (\sigma_\epsilon^2/d)I_d)$, and finally project $\delta_{z_i} + \widehat{z}_i$ onto $B_2(x_i, \epsilon)$ to obtain the corrupted attack. There is no noise injection for parameters and $x_i$ in this experiment, i.e. $\delta_\xi = \delta_x = 0$.

We take $(\epsilon, \delta_\epsilon)$ as (0.2,0), (0.2,0.1), (0.4,0), (0.4,0.2). For each setup, the experiment repeats 100 times to obtain a mean and variance. The results are summarized in Figure 2 for the whole training process and Table 1 for the detailed values at the 500th iteration. Compared with $\delta_\epsilon = 0$, when $\delta_\epsilon > 0$, both adversarial testing loss and generalization error increases. For the increase in adversarial training loss, our conjecture is that the attack error perturbs the training process, so the trained model is slight away from its optimum.

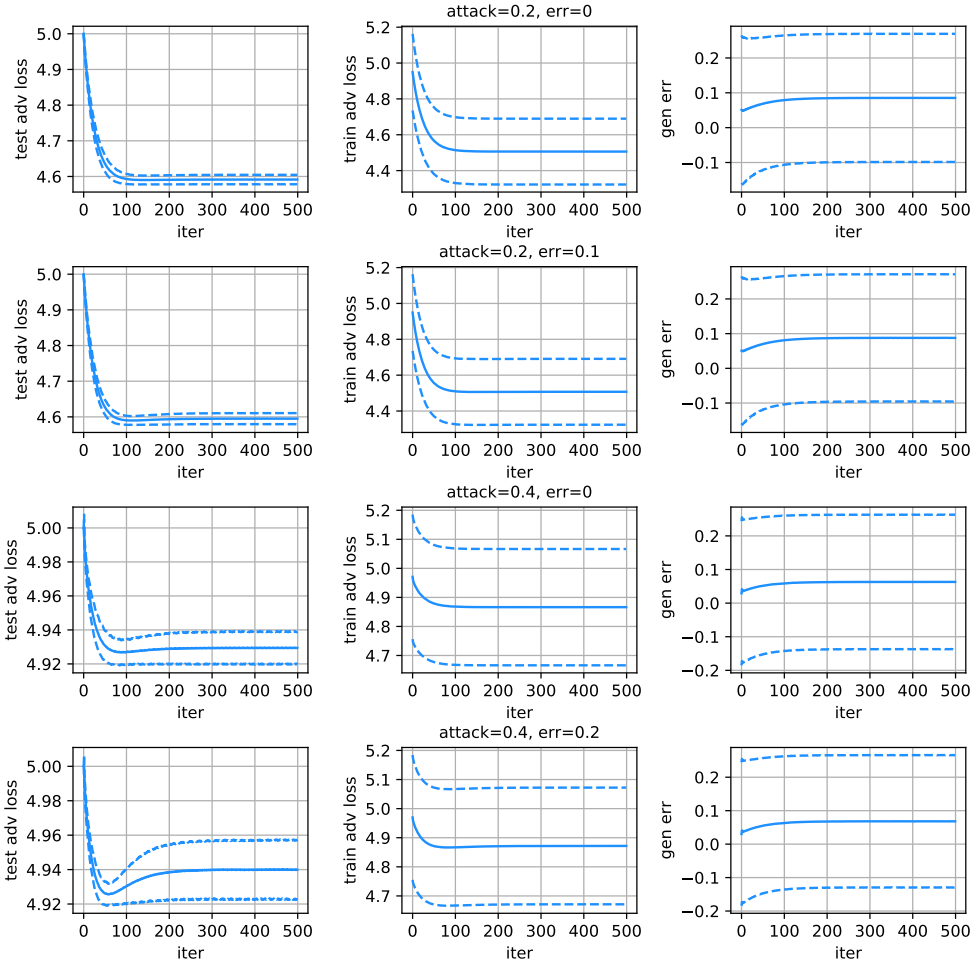


Figure 2: Mean and standard deviation of adversarial testing loss (left), adversarial training loss (middle), generalization error (right) for four groups of $(\epsilon, \sigma_\epsilon)$. After injecting error into attack (2nd and 4th row), all the adversarial testing loss, adversarial training loss, and generalization error increase (compared with 1st and 3rd row).

| $\epsilon$ | $\sigma_\epsilon$ | Test adv loss | Train adv loss | Gen err |
|---|---|---|---|---|
| 0.2 | 0.0 | 4.5914 (0.0134) | 4.5061 (0.1836) | 0.0854 (0.1838) |
| 0.2 | 0.1 | 4.5950 (0.0157) | 4.5073 (0.1837) | 0.0877 (0.1833) |
| 0.4 | 0.0 | 4.9294 (0.0096) | 4.8664 (0.2003) | 0.0631 (0.1998) |
| 0.4 | 0.2 | 4.9399 (0.0175) | 4.8718 (0.2005) | 0.0681 (0.1977) |

Table 1: Mean and standard deviation of adversarial testing loss, adversarial training loss, generalization error for the 500th iteration. After injecting error into attack, the adversarial testing loss, adversarial training loss, and generalization error increase.

# C Results in Neural Networks

## C.1 Two-Layer Neural Network

This section provides some results w.r.t two-layer neural network with lazy training, i.e. training the first hidden layer, and vanishing initialization.

Denote $h$ as the number of hidden nodes, $\theta = [\theta_1 \mid \theta_2 \mid ... \mid \theta_h]$ are weights for each node. The neural network outputs value as follows

$$f(x, \theta) = \frac{1}{\sqrt{h}} \sum_{j=1}^{h} a_j \phi(\theta_j^\top x),$$

where $\phi$ is a smooth activation function which is either (1) twice differentiable and $\phi(0) = 0$, or (2) ReLU activation function. The value of $a_j$'s are unchanged after the initialization, and we randomly generate them from $\{\pm 1\}$. The initial value of $\theta_j$ ($\theta_j^{(0)}$) is generated from $N(0, I_d/(dh^{\delta+1}))$, i.e. vanishing initialization. The adversarial training will update $\theta_j^{(t)}$ throughout training.

We consider regression task, and to simplify the analysis, we use the following data generation model:

$$y = \theta_0^\top x + \omega,$$

where $x$ follows $N(0, I_d)$, $\|\theta_0\|$ is a constant, and $\omega$ is a Gaussian noise with zero mean and finite variance.

Also denote adversarial risk minimizor and the minimal adversarial risk for linear model as follows:

$$\theta^* = \operatorname*{argmin}_\theta \mathbb{E}(y - \theta^\top (x + A_\epsilon(f_\theta, x, y)))^2, \; R^* = \min_\theta \mathbb{E}(y - \theta^\top (x + A_\epsilon(f_\theta, x, y)))^2.$$

We have the following theorem:

**Theorem 6.** *Assume $\Delta\epsilon = 0$. Under the model setup of neural network and data generating model in this section, assume $\log n \sqrt{d^2/n} \to 0$, $(d \log n)/\sqrt{h} \to 0$ and $\sqrt{d \log n}(1 + D)^T / h^{\delta/2} \to 0$ for some large constant $D$. Take $L = \Theta(\sqrt{d \log n})$. If $\eta = \zeta / B^*$, $T = (\log \log n)/\eta$, and $\sqrt{d \log n \log(hT)}\xi \to 0$, then using GD, taking $(x, y)$ as a random testing data,*

$$\begin{aligned} & \mathbb{E}_{S_1 \sim S_2} \left| l(f_{\theta_1^{(T)}}[x + A_\epsilon(f_{\theta_1^{(T)}}, x_i, y_i)], y_i) - l(f_{\theta_2^{(T)}}[x + A_\epsilon(f_{\theta_2^{(T)}}, x_i, y_i)], y_i) \right| \\ = \; & O\left( \left[ L\sqrt{P(E^c)} + \sqrt{\frac{L^2}{n}} \right] \eta \sqrt{T} + \left[ \frac{L}{n} + LP(E^c) \right] \eta T \right) + rem, \end{aligned} \tag{5}$$

*where $rem = o(1)$ and is not the dominant term when $h$ and $\delta$ are large enough. The value $B^*$ is defined similarly as in Lemma 5 for linear regression.*

Theorem 6 illustrates how the hypothesis stability changes throughout the training. This stability result can be trivially used to bound the third term on the RHS of generalization inequality 4 in Proposition 2.

To prove Theorem 6, it is harder than simple models since $\|\theta_1^{(T)} - \theta_2^{(T)}\|$ is not so meaningful in neural networks. Starting from the vanishing initialization, we track the change in each node one by one. We only consider GD because GD ensures the convergence of neural network.

For the two terms in (5), the first term is obtained when bounding the difference between $\theta_1^{(T)}$ and $\theta_2^{(T)}$; the second term counts for (1) the rare events on $S_1, S_2$ (we mention "with probability tending to 1 over $S_1 \sim S_2$" in Lemma 2) ;(2) the difference when we use linear network to approximate nonlinear network.

*Proof.* We provide the proof for the first term in (5) for smooth activation function. The idea is similar to the arguments in Xing et al. (2021a). We first consider how linear network with zero initialization could act, compared to a linear model with zero initialization under noise injection.

Then we bound the difference between noise-injected linear network with zero initialization and noise-injected nonlinear network with vanishing initialization.

First, we consider a linear network

$$f_L(x,\theta)=\frac{1}{\sqrt{h}}\sum_{j=1}^{h}\phi'(0)a_j\theta_j^\top x=\left(\frac{1}{\sqrt{h}}\sum_{j=1}^{h}\phi'(0)a_j\theta_j\right)^\top x.$$

When injecting noise $\delta_j$ into $\theta_j$, it becomes

$$f_L(x,\theta+\delta)=\frac{1}{\sqrt{h}}\sum_{j=1}^{h}\phi'(0)a_j(\theta_j+\delta_j)^\top x=\left(\frac{1}{\sqrt{h}}\sum_{j=1}^{h}\phi'(0)a_j\theta_j\right)^\top x+\left(\frac{1}{\sqrt{h}}\sum_{j=1}^{h}\phi'(0)a_j\delta_j\right)^\top x,$$

where $\frac{1}{\sqrt{h}}\sum_{j=1}^{k}\phi'(0)a_j\delta_j$ is a random vector with zero mean and covariance $N(0,\phi'(0)^2\xi^2 I_d/d)$.

Denote the parameters trained from linear model with zero initialization as follows: for dataset $S_k$,

$$\begin{aligned}\theta_{j,k}^{OP1}(0) &= \mathbf{0},\\ \theta_{j,k}^{OP1}(t+1) &= \theta_{j,k}^{OP1}(t)-\eta\left(\frac{2}{n}\sum_{(x,y)\in S_k}\frac{a_j\phi'(0)}{\sqrt{h}}(x+A_\epsilon(f,x,y))\left(\frac{\phi'(0)}{\sqrt{h}}\sum_{m=1}^{h}a_j(\theta_{m,k}^{OP1}(t)+\delta_m)^\top x-y\right)\right).\end{aligned}$$

As a result, injecting noise $\delta_j\sim N(0,\xi^2 I_d/d)$ into a linear network is equivalent to injecting noise $N(0,\phi'(0)^2\xi^2 I_d/d)$ into a linear model. So one may use arguments in Lemma 2 and Theorem 2 to study

$$\left\|\frac{1}{\sqrt{h}}\sum_{j=1}^{h}\phi'(0)a_j(\theta_{j,1}^{OP1}(t)-\theta_{j,2}^{OP1}(t))\right\|$$

given $\left\|\frac{1}{\sqrt{h}}\sum_{j=1}^{h}\phi'(0)a_j(\theta_{j,1}^{OP1}(t-1)-\theta_{j,2}^{OP1}(t-1))\right\|$.

On the other hand, from the updating rule of $\theta_{j,k}^{OP1}(t)$, one can also see that $\theta_{j,k}^{OP1}(t)\equiv\theta_{l,k}^{OP1}(t)$ if $a_j=a_l$. Therefore, besides $\frac{1}{\sqrt{h}}\sum_{j=1}^{h}\phi'(0)a_j\theta_{j,k}^{OP1}(T)\to\theta^*$ (since we are using GD and $n\to\infty$), one can also solve $\theta_{j,k}^{OP1}(T)$ for each $j$.

Denote the parameters trained from nonlinear model with vanishing initialization as follows: for dataset $S_k$,

$$\begin{aligned}&\theta_{j,k}^{OP2}(t+1)\\ =\;&\theta_{j,k}^{OP2}(t)-\eta\left(\frac{2}{n}\sum_{(x,y)\in S_k}\frac{a_j\phi'((\theta_{m,k}^{OP2}(t)+\delta_m)^\top x)}{\sqrt{h}}(x+A_\epsilon(f,x,y))\left(\frac{1}{\sqrt{h}}\sum_{m=1}^{h}a_j\phi((\theta_{m,k}^{OP2}(t)+\delta_m)^\top x)-y\right)\right).\end{aligned}$$

When we do not inject noise to network parameters, the difference between $\theta_{j,k}^{OP2}(t)$ and $\theta_{j,k}^{OP1}(t)$ can be neglected when $(d\log n)/\sqrt{h}\to 0$ and $\sqrt{d\log n}(1+D)^T/h^{\delta/2}\to 0$ (Theorem 3 of Xing et al. (2021a)).

When injecting noise to network parameters, we further want $\max_{i,k}|\delta_k^\top x_i|\to 0$ in probability so that $\phi((\theta_j+\delta_j)^\top x)=\phi(0)+\phi'(0)(\theta_j+\delta_j)^\top x+O(((\theta_j+\delta_j)^\top x)^2)$ and the second order term is a remainder. Note that since $\delta_k$ and $x_i$ are generated from Gaussian, we have $\|\delta_k\|/\xi$ is in $O(\sqrt{\log(hT)})$ (among the $T$ iterations and $h$ nodes in each iteration) and $\|x_i\|$ is in $O(\sqrt{d\log n})$. As a result, $\max_{i,k}|\delta_k^\top x_i|=O(\xi\sqrt{d\log n\log(hT)})$, which by assumption is a vanishing term. Further injecting noise in data has little impact on the difference between $\theta_{j,k}^{OP2}(t)$ and $\theta_{j,k}^{OP1}(t)$ since $\xi_0\to 0$, and we skip this part.

The proofs for ReLU networks follows similar arguments with assistance of Theorem 4 of Xing et al. (2021a).

□

### C.2 Exploration in deep learning

Section 4.2 and 3.3 suggest that (1) the noise injection in model parameters and input data, and (2) better accuracy on the approximation of attack $A_\epsilon$, both lead to a better stability of the algorithm. Although these theoretical insights are derived under certain simple statistical models, we conjecture that they apply to modern complex models. Hence, in this section, we assess their effects on the stability of DNNs. Inspired by Proposition 3, we use the difference between adversarial training accuracy and adversarial testing accuracy to measure algorithm stability.

We use SGD with batch size 128 and weight decay 0.0002 as the optimizer. The learning rate is taken as 0.1 at the beginning and multiplies 0.1 at the 75th and 90th epoch. The total number of epochs is 100. To overcome the non-smoothness from ReLU activation, if not specified, we use WideResNet34-1 as the network structure with replacing backward update of ReLU into Softmax(10) using the `BPDA` in AdverTorch[2] (Ding et al. 2019).

#### C.2.1 Noise injection

Theorem 2 indicates that injecting noise in adversarial training improves algorithmic stability. This section examines the effect of noise injection in adversarial training under a deep learning setup.

In this experiment, we use CIFAR10 and compare the adversarial testing accuracy before/after injecting noise using the implementation of TRADES[3] in Zhang et al. (2019). For the model parameters, the noise for each element is generated from a zero-mean normal distribution with a standard deviation equal to $\alpha_{h,t}\sigma_{h,t}$, where $\sigma^2_{h,t}$ is the variance of the parameters in $h$-th layer at $t$-th iteration and $\alpha_{h,t}$ is a trainable parameter initialized as 0.1. The implementation follows the one in He et al. (2019)[4]. We consider data augmentation (Shorten and Khoshgoftaar, 2019) as a form of noise injection to the data, and compare generalization performance with/without data augmentation during the training. For data augmentation method, we follow Zhang et al. (2019); Wang et al. (2019b); He et al. (2019) to include `Randomop(32, padding=4)` and `RandomHorizontalFlip()`. Each setup is repeated for five time to obtain a mean and variance. The results are summarized in Table 2. AT represents the vanilla adversarial training.

| # | Method | Noise | Aug | $\epsilon$ | Adv Train Acc | Adv Test Acc | Gen Gap | Std(Gen Gap) |
|---|---|---|---|---|---|---|---|---|
| 1 | AT | No | No | $\mathcal{L}_2$ 0.5 | 97.1775 | 55.4525 | 41.725 | 0.3414 |
| 2 | AT | No | Yes | $\mathcal{L}_2$ 0.5 | 72.542 | 52.698 | 19.844 | 1.302 |
| 3 | AT | Yes | Yes | $\mathcal{L}_2$ 0.5 | 69.27 | 54.72 | 14.56 | 1.412 |
| 4 | AT | No | No | $\mathcal{L}_\infty$ 8/255 | 75.525 | 37.055 | 38.47 | 0.2493 |
| 5 | AT | No | Yes | $\mathcal{L}_\infty$ 8/255 | 50.486 | 36.162 | 14.324 | 0.5118 |
| 6 | AT | Yes | Yes | $\mathcal{L}_\infty$ 8/255 | 47.94 | 37.64 | 10.30 | 1.252 |
| 7 | TRADES | No | No | $\mathcal{L}_2$ 0.5 | 91.42 | 51.555 | 39.865 | 0.781 |
| 8 | TRADES | No | Yes | $\mathcal{L}_2$ 0.5 | 67.632 | 57.0 | 10.632 | 0.7764 |
| 9 | TRADES | Yes | Yes | $\mathcal{L}_2$ 0.5 | 65.904 | 61.508 | 4.396 | 0.4583 |

Table 2: Effect of different methods on generalization gap under $\mathcal{L}_2$/$\mathcal{L}_\infty$ attack.

In Table 2, when training using adversarial training for (#1, #2, #3), both data augmentation and noise injection in model parameters reduce the accuracy difference a lot. The vanilla adversarial training/testing gap is around 42%. After introducing data augmentation, the gap reduces to 20%. Finally, after further injecting noise into model parameters, the gap gets sown to only 15%. We also compare the adversarial testing accuracy before/after injecting noise for $\mathcal{L}_\infty$ attack. The observations are similar. The observation from TRADES are similar to AT. It still suffers from the stability issue, and injecting noise can reduce the generalization gap.

It is worth emphasizing that the aim of injection noise is not to improve the final testing performance, but to reduce the generalization gap (without sacrificing the testing performance). It has been shown in Rice et al. (2020) that, the adversarial training tends to overfit, i.e., the high training

[2] `https://github.com/BorealisAI/advertorch`

[3] `https://github.com/yaodongyu/TRADES`

[4] `https://github.com/elliothe/CVPR_2019_PNI/blob/master/code/models/noise_layer.py`

acc is deceptive and not reliable. Therefore, the training loss shall decrease if a stable algorithm is implemented.

**Remark 7.** *Although our main target of this experiment is to examine the reduction of generalization gap after noise injection, one may also find some other observations from the experiment results. For example, with noise injection, TRADES gets a better adversarial testing accuracy, while AT does not have significant change in this, which implies some potential differences between AT and TRADES to help people better understand why AT does not perform well.*

### C.2.2 Improving the attack error

From Corollary 1, an accurate attack leads to better algorithmic stability. To explore this in neural networks, since PGD-5 is more accurate than FGM, we use $\mathcal{L}_2$ adversarial training under different choices of $\epsilon$, and compare the adversarial training accuracy against the generalization gap for these two choices of attack method (i.e., FGM vs. PGD-5). To ensure the comparison is fair, we use FGM in testing data if the training uses FGM and use PGD-5 in testing data if the training uses PGD-5. The value of $\epsilon$ ranges from 0.25 to 4.0 to achieve different levels of adversarial training accuracy.

As shown in Figure 3, when the same level of adversarial training accuracy (70%$\sim$80%) is achieved, the generalization gap of FGM, in a worse-case scenario, can be much larger than PGD.

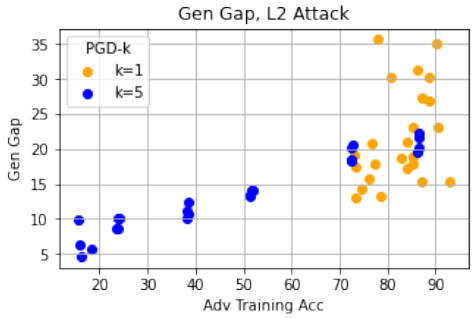


Figure 3: Comparison between FGM (i.e., PGD-1) and PGD-5 in adversarial training accuracy against generalization gap. The generalization gap using FGM is poor. Different values are obtained under different attack strength.

## D Proofs

### D.1 Adversarial Training without Noise Injection

*Proof of Theorem 1.* The proof mainly follows Bassily et al. (2020) and we will transfer their worst-case scenario into the format of adversarial training.

Let $D = \min\{t, 1/\eta^2\} \leq d$, and $\nu > 0$, $K \geq \sqrt{D}$. Take linear loss functions and $y \equiv 0$. Denote $h$ as a smooth activation function to approximate ReLU, with $h = 0$ when $z < -\zeta$, and $h(z) = z$ when $z > \zeta$. Define $f$ as

$$f_\theta(x, z) = \begin{cases} \|h(\theta - x)\|_H & z < 0.5 - \lambda \\ r^\top \theta / K & z \geq 0.5 + \lambda \\ \frac{z - 0.5 + \lambda}{2\lambda} \|h(\theta - x)\|_H + \frac{0.5 + \lambda - z}{2\lambda} r^\top \theta / K & \text{otherwise} \end{cases},$$

where $\|h(\theta - x)\|_H$ is a smooth approximation of $\max\{0, \theta - x\}$, and $r$ is a vector with first $D$ elements as -1 and others as 0. The tuple $(x, z)$ represents the independent variables in the data.

Consider attack strength $\epsilon < 0.5$, then taking $z \in \{0, 1\}$, this attack strength is not strong enough to change whether $z$ is greater than 0.5 or not, and the attack will only change $x$.

In the first dataset $S_1$, $(x_1, z_1) = (\nu, ..., \nu, 1)$, and the others are $(x_i, z_i) = (\nu, ..., \nu, 0)$. In the second dataset $S_2$, all the samples are $(x_i, z_i) = (\nu, ..., \nu, 0)$.

Take initialization $\theta_1^{(0)} = \theta_2^{(0)} = \mathbf{0}$. If attack strength is small enough such that $\epsilon < \nu$, then $\theta_2^{(t)}$ is always $\mathbf{0}$.

To analyze $\theta_1^{(t)}$, for SGD, denote $i_t$ as the sample index at $t$th iteration. The value of $\theta_1^{(t)}$ keeps $\mathbf{0}$ before the first $(x_1, z_1)$ appears. For the first $t$ (denote as $t_0$) such that $i_t = 1$, since the function $f_\theta$ is not related to $x$ when $z = 1$, we have $\theta_1^{(t_0)} = -\eta r/K$. Taking $\nu < \eta/K$, for the next step $t_0 + 1$, if $i_{t_0+1}$ is not 1, then the attack will randomly select an element of $x$ to result in

$$\max_{x' \in B_2(x_{i_{t_0+1}}, \epsilon)} f_{\theta_1^{(t_0)}}(x', z_{i_{t_0+1}}) = \eta/K - \nu + \epsilon + o,$$

where we chose $h$ and $H$ properly so that its gradient has only minor difference with $\max\{0, \theta - x\}$, which is represented as $o$. The update of $\theta_1$ will be in the corresponding dimension chosen by attack, which has the same outcome as the nonsmooth function considered in Bassily et al. (2020) for clean training. Then the remaining proof follows the one in Bassily et al. (2020) directly.

The proof of GD is similar as SGD, as we take sufficiently large $K$ such that $\nu < \eta\|r\|/K$.

□

*Proof of Theorem 5.* In the proof, we slightly change the assumption on $\kappa$ to

$$\|\nabla_\theta l(f_{\theta_1}(x_1), y) - \nabla_\theta l(f_{\theta_2}(x_2), y)\|^2 + \|\nabla_x l(f_{\theta_1}(x_1), y) - \nabla_x l(f_{\theta_2}(x_2), y)\|^2 \le \kappa^2(\|\theta_1 - \theta_2\|^2 + \|x_1 - x_2\|^2).$$

Define $z_i^1$ and $z_i^2$ as the correct attack of sample $(x_i, y_i)$ given the models $\theta_1^{(t)}$ and $\theta_2^{(t)}$. For SGD, we have

$$\begin{aligned}
\Delta_{t+1} &\le \left\|\theta_1^{(t)} - \theta_2^{(t)} - \eta_t \left(\nabla_\theta l(f_{\theta_1^{(t)}}(\widehat{z}_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t)}}(\widehat{z}_{i_t}^2), y_{i_t}^2)\right)\right\| \\
&\le \Delta_t + \eta_t \left\|\nabla_\theta l(f_{\theta_1^{(t)}}(\widehat{z}_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t)}}(\widehat{z}_{i_t}^2), y_{i_t}^2)\right\| \\
&= \Delta_t + \eta_t \left\|\nabla_\theta l(f_{\theta_1^{(t)}}(\widehat{z}_{i_t}^1 - z_{i_t}^1 + z_{i_t}^1 - x_{i_t}^1 + x_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t)}}(\widehat{z}_{i_t}^2 - z_{i_t}^2 + z_{i_t}^2 - x_{i_t}^2 + x_{i_t}^2), y_{i_t}^2)\right\| \\
&\le \Delta_t + 2\eta_t l\Delta\varepsilon + 2\eta_t L + \eta_t \left\|\nabla_\theta l(f_{\theta_1^{(t)}}(x_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t)}}(x_{i_t}^2), y_{i_t}^2)\right\|.
\end{aligned}$$

Therefore, given $\Delta_t$,

$$\begin{aligned}
\mathbb{E}[\Delta_{t+1}|\Delta_t] &\le \Delta_t + 2\eta_t(\kappa\Delta\varepsilon + L) + \eta_t \mathbb{E}\left[\left\|\nabla_\theta l(f_{\theta_1^{(t)}}(x_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t)}}(x_{i_t}^2), y_{i_t}^2)\right\| 1\{(x_{i_t}^1, y_{i_t}^1) = (x_{i_t}^2, y_{i_t}^2)\}\right] \\
&\quad + \eta_t \mathbb{E}\left[\left\|\nabla_\theta l(f_{\theta_1^{(t)}}(x_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t)}}(x_{i_t}^2), y_{i_t}^2)\right\| 1\{(x_{i_t}^1, y_{i_t}^1) \ne (x_{i_t}^2, y_{i_t}^2)\}\right] \\
&\le \Delta_t + 2\eta_t(\kappa\Delta\varepsilon + L) + \eta_t\kappa\Delta_t \frac{n-1}{n} + \frac{\eta_t L}{n}.
\end{aligned}$$

Since $\eta_t \le c/t$, denoting $t_0$ as the known first time that the $i_t$th sample in the two datasets are difference, we have

$$\begin{aligned}
\mathbb{E}[\Delta_t|t_0] &\le \left(1 + \frac{c\kappa}{t}\right)\mathbb{E}[\Delta_t|t_0] + \frac{c}{t}\left(2\kappa\Delta\varepsilon + 2L + \frac{L}{n}\right) \\
&\le \sum_{\tau=t_0+1}^{t} \prod_{k=\tau+1}^{t} \left(1 + \frac{c\kappa}{k}\right) \frac{c}{\tau}\left(2\kappa\Delta\varepsilon + 2L + \frac{L}{n}\right) \\
&\le \sum_{\tau=t_0+1}^{t} \prod_{k=\tau+1}^{t} \exp\left(\frac{c\kappa}{k}\right) \frac{c}{\tau}\left(2\kappa\Delta\varepsilon + 2L + \frac{L}{n}\right) \\
&= \sum_{\tau=t_0+1}^{t} \exp\left(\sum_{k=\tau+1}^{t} \frac{c\kappa}{k}\right) \frac{c}{\tau}\left(2\kappa\Delta\varepsilon + 2L + \frac{L}{n}\right) \\
&\le \sum_{\tau=t_0+1}^{t} \exp\left(c\kappa\log(t/\tau)\right) \frac{c}{\tau}\left(2\kappa\Delta\varepsilon + 2L + \frac{L}{n}\right) \\
&= ct^{c\kappa}\left(2\kappa\Delta\varepsilon + 2L + \frac{L}{n}\right) \sum_{\tau=t_0+1}^{t} \tau^{-c\kappa-1} \\
&\le c\left(2\kappa\Delta\varepsilon + 2L + \frac{L}{n}\right)\left(\frac{t}{t_0}\right)^{c\kappa}.
\end{aligned}$$

As a result, taking expectation w.r.t $t_0$, we have

$$\begin{aligned}\mathbb{E}[\Delta_t] &\leq c\left(2\kappa\Delta\varepsilon + 2L + \frac{L}{n}\right)\mathbb{E}\left[\left(\frac{t}{t_0}\right)^{c\kappa} 1\{t_0 \geq t_1\}\right] + 2rP\left(t_0 < t_1\right) \\ &\leq c\left(2\kappa\Delta\varepsilon + 2L + \frac{L}{n}\right)\left(\frac{t}{t_1}\right)^{c\kappa} P\left(t_0 \geq t_1\right) + 2rP\left(t_0 < t_1\right).\end{aligned}$$

Taking $t_1 = \Theta((nt^{c\kappa})^{\frac{1}{1+c\kappa}})$, it becomes

$$\mathbb{E}[\Delta_t] = O\left(\left(2\kappa\Delta\varepsilon + 2L + \frac{L}{n}\right)\left(\frac{t}{n}\right)^{\frac{c\kappa}{1+c\kappa}}\right).$$

For GD, since the different sample appears in the first iteration, we directly take $t_0 = 1$ in (6) and obtain the result. □

### D.2 Adversarial Training with Noise Injection

We first present the formal statement of Lemma 1 as follows:

**Lemma 3.** *Assume Assumption 1 holds. Denote $L$ as the Lipschitz constant of $l(f_\theta(x), y_j)$ w.r.t. $\theta$ for any $x \in B_2(x_j, 2\epsilon)$ and all $1 \leq j \leq n$, and $\kappa$ as the Lipschitz constant of $\nabla_\theta l(f_\theta(x), y)$ w.r.t. $x$. Take $B$ as some function (specified later) of $(L, n, d, \kappa)$. Then, $(B, L, \kappa)$ is bounded by some finite $(B^*, L^*, \kappa^*)$ with probability tending 1, where the probability refers to the generation measure of $S = \{x_j, y_j\}_{j=1}^n$.*

*Assume the noise injected in data is zero-mean Gaussian with variance $(\xi_0^2/d)I_d$, and the noise injected in parameters is zero-mean Gaussian with variance $(\xi^2/d)I_d$ with $\xi = \xi_0 L^*$ and $\xi(d\log n) \to 0$. Denote $E(\theta + \delta, \widetilde{x}, y)$ as the event that $\nabla_\theta l(f_{\theta+\delta}(\widetilde{x} + A_\epsilon(f_{\theta+\delta}, \widetilde{x}, y)), y)$ is $B^*/\zeta$-Lipschitz. Then for the regression and classification tasks, there exists some $\zeta \ll \xi \to 0$ in $n$ such that, with probability tending to one over the generation of $S$, uniformly for all $\theta \in B_2(0, r)$,*

$$P\left(E^c(\theta + \delta, \widetilde{x}, y)|(x, y) \in S\right) = o(1).$$

Let $P(E^c|S) := \sup_{\theta \in B_2(0,r),(x,y)} P\left(E^c(\theta + \delta, \widetilde{x}, y)|(x, y) \in S\right)$ in what follows, for notation simplicity.

**Remark 8.** *The terms $r$, $L$, $\kappa$ are generic representations. For different loss functions and data dimension $d$, their values may change. In addition, the exact rate of $P\left(E^c|S\right)$ is affected by the value of $r, L, \kappa$ as well as $\xi_0, \zeta_0$. We postpone the details to the proof.*

In the following proofs regarding to Theorem 2 and 3, we use linear regression as an example. To be more specific, the three lemmas to be used in the main proof, Lemma 4, Lemma 5 and 6, provide some results w.r.t $E$ and $\mathbb{E}g$ for linear regression model. The proof for Theorem 2 and 3 directly utilize the results on $E$ and $\mathbb{E}g$ instead of any specific model. We provide the results of $E$ and $\mathbb{E}g$ for other models in the next section. Theorem 2 and 3 also hold after replacing Lemma 5 and 6 by these lemmas.

In terms of Lemma 3, it is a summary of results of $L$ and $P(E^c|S)$ over different models.

**Lemma 4.** *For linear regression, there exists some $(L^*, \kappa^*)$ such that, with probability tending to one over the choice of $S$, $L \leq L^*$ and $\kappa \leq \kappa^*$.*

*Proof of Lemma 4.* The gradient can be written as

$$\frac{1}{2}\nabla_\theta l(f_\theta(x), y) = x(x^\top\theta - y).$$

Then from the definition of the Lipschitz constant $L$, when taking $\delta_x$ such that $\delta_x \in B_2(0, 2\epsilon)$,

$$\frac{1}{2}L = \max_{\theta \in B_2(0,r), i \in [n], \delta_x} \|x_i + \delta_x\||(x_i + \delta_x)^\top\theta - y_i| \leq (\max_i \|x_i\| + 2\epsilon)^2 r + (\max_i \|x_i\||y_i| + 2\epsilon|y_i|).$$

In addition,

$$\begin{aligned}
& \frac{1}{2}\left\|\nabla_\theta l(f_\theta(x),y)-\nabla_\theta l(f_\theta(x+\delta_x),y)\right\| \\
= & \frac{1}{2}\left\|x(x^\top\theta-y)-x((x+\delta_x)^\top\theta-y)-\delta_x((x+\delta_x)^\top\theta-y)\right\| \\
= & \frac{1}{2}\left\|x\delta_x^\top\theta-\delta_x((x+\delta_x)^\top\theta-y)\right\| \\
= & \frac{1}{2}\left\|x\delta_x^\top\theta-\delta_x(x^\top\theta-y)-\delta_x\delta_x^\top\theta\right\| \\
\le & \frac{1}{2}\left(\|x\|\|\theta\|\|\delta_x\|+\|\delta_x\||x^\top\theta-y|+\|\delta_x\|^2\|\theta\|\right) \\
\le & \frac{1}{2}\left(\|x\|\|\theta\|\|\delta_x\|+\|\delta_x\||x^\top\theta-y|+\|\delta_x\|2\epsilon\|\theta\|\right).
\end{aligned}$$

Thus for a given set of data $S$,

$$\frac{1}{2}\kappa=\max_{\theta\in B_2(0,r),i\in[n]}\left[(\|x_i\|+2\epsilon)\|\theta\|+|x_i^\top\theta-y_i|\right]\le 2(\max_i\|x_i\|+\epsilon)r+\max_i|y_i|.$$

From the distribution of $x$, we know that $\max_i\|x_i\|=O(\sqrt{d\log n})$ almost surely. In addition, $\mathbb{E}\|x\||y|$ and $\mathbb{E}|y|$ are finite, thus $\max_i\|x_i\||y_i|$ and $\max_i|y_i|$ are some functions of $n$ as well.

□

**Lemma 5.** *For linear regression, denote $\zeta=L\zeta_0$ for some $\zeta_0/\xi_0\to 0$. Denote $E(\theta,\delta,\widetilde{x},y)=1\{\|\theta+\delta\|\ge\zeta,|\widetilde{x}^\top(\theta+\delta)-y|\ge\zeta_0 r(d\log n)\}$, then $E=1$ implies that $\nabla_\theta l(f_{\theta+\delta}(\widetilde{x}),y)$ is $B/\zeta_0$-Lipschitz. Uniformly for all $\theta$, with probability tending to one over the $n$ random samples, we have*

$$P(E^c(\theta,\delta,\widetilde{x},y)|S)=o(1).$$

*Proof of Lemma 5.* We show that $E=1$ implies that $\nabla_\theta l(f_{\theta+\delta}(\widetilde{x}),y)$ is $B/\zeta_0$-Lipschitz. The gradient of adversarial loss is

$$\frac{1}{2}g(\widetilde{x},y,\theta)=\widetilde{x}(\widetilde{x}^\top(\theta+\delta)-y)+\epsilon^2(\theta+\delta)+\epsilon\frac{(\theta+\delta)}{\|\theta+\delta\|}|y-\widetilde{x}^\top(\theta+\delta)|-\epsilon\widetilde{x}\|(\theta+\delta)\|\text{sgn}(y-\widetilde{x}^\top(\theta+\delta)).$$

When $\|\theta+\delta\|\ge\zeta$, we have for any $\theta'$,

$$\frac{1}{\|\theta+\delta-\theta'\|^2}\left\|\frac{\theta'}{\|\theta'\|}-\frac{\theta+\delta}{\|\theta+\delta\|}\right\|^2=\frac{2}{\|\theta+\delta-\theta'\|^2}-\frac{2}{\|\theta+\delta-\theta'\|^2}\frac{(\theta+\delta)^\top\theta'}{\|\theta'\|\|\theta+\delta\|}.$$

Taking $\theta'\propto-(\theta+\delta)$, the above quantity is maximized. Therefore, taking $\theta'=-\alpha(\theta+\delta)$ for $\alpha>0$,

$$\begin{aligned}
\frac{1}{\|\theta+\delta-\theta'\|^2}\left\|\frac{\theta'}{\|\theta'\|}-\frac{\theta+\delta}{\|\theta+\delta\|}\right\|^2 & \le \frac{4}{\|\theta+\delta+\alpha(\theta+\delta)\|^2} \\
& \le \lim_{\alpha\to 0^+}\frac{4}{\|\theta+\delta+\alpha(\theta+\delta)\|^2} \\
& \le \frac{4}{\zeta^2}.
\end{aligned}$$

When $|y - \widetilde{x}^\top(\theta+\delta)| \geq \gamma$ for some $\gamma$, this implies that the nearest $\theta'$ such that $\text{sgn}(y - \widetilde{x}^\top(\theta+\delta))$ gets changed satisfies $\|\theta' - (\theta+\delta)\| = \gamma/\|\widetilde{x}\|$. As a result,

$$\begin{aligned}
& \frac{1}{\|\theta+\delta-\theta'\|}\left\|\widetilde{x}\|\theta+\delta\|\text{sgn}(y-\widetilde{x}^\top(\theta+\delta)) - \widetilde{x}\|\theta'\|\text{sgn}(y-\widetilde{x}^\top\theta')\right\| \\
\leq\ & \frac{1}{\|\theta+\delta-\theta'\|}\left\|\widetilde{x}\|\theta+\delta\|\text{sgn}(y-\widetilde{x}^\top\theta') - \widetilde{x}\|\theta'\|\text{sgn}(y-\widetilde{x}^\top\theta')\right\| \\
& + \frac{1}{\|\theta+\delta-\theta'\|}\left\|\widetilde{x}\|\theta+\delta\|\text{sgn}(y-\widetilde{x}^\top(\theta+\delta)) - \widetilde{x}\|\theta+\delta\|\text{sgn}(y-\widetilde{x}^\top\theta')\right\| \\
\leq\ & \frac{\|\widetilde{x}\|\|\theta+\delta-\theta'\|}{\|\theta+\delta-\theta'\|} + \frac{\|\widetilde{x}\|\|\theta+\delta\|}{\|\theta+\delta-\theta'\|}\left|\text{sgn}(y-\widetilde{x}^\top(\theta+\delta)) - \text{sgn}(y-\widetilde{x}^\top\theta')\right| \\
\leq\ & \|\widetilde{x}\| + \frac{2\|\widetilde{x}\|^2 r}{\gamma}.
\end{aligned}$$

Take $\gamma = \zeta_0 r\|\widetilde{x}\|^2$ in the above inequality to obtain $\sqrt{d\log n} + 2/\zeta_0$-Lipschitz.

Therefore the overall gradient is Lipschitz with

$$\kappa + 2\epsilon^2 + 8\epsilon/\zeta_0 + 2\epsilon\sqrt{d\log n} \tag{6}$$

which can be rewritten as $B/\zeta_0$ for some $B$.

Now we turn to bound the probability of $E^c$.

$$P(E^c(\theta,\delta,\widetilde{x},y)|S) \leq P(\|\theta+\delta\| < \zeta) + P(|\widetilde{x}^\top(\theta+\delta) - y| < \zeta_0 r(d\log n)|S).$$

For any $\theta$, based on the distribution of $\delta$, we have

$$P(\|\theta+\delta\| < \zeta|\theta) = O\left(\left(\frac{\zeta}{\xi}\right)^d\right).$$

On the other hand,

$$P(|\widetilde{x}^\top(\theta+\delta) - y| < \zeta_0 r(d\log n)|S) = P(|(\widetilde{x}-x)^\top\theta + (\widetilde{x}-x)^\top\delta + (x^\top\theta - y) + x^\top\delta| < \zeta_0 r(d\log n)|S).$$

When $\|\theta\| > Cr$, from the distribution of $x^\top\delta$, $(\widetilde{x}-x)^\top\delta$, $(x^\top\theta - y)$, and $(\widetilde{x}-x)^\top\theta$, we have for any $(x,y,\theta)$,

$$\begin{aligned}
& P\left(|(\widetilde{x}-x)^\top\theta + (\widetilde{x}-x)^\top\delta + (x^\top\theta - y) + x^\top\delta| < \zeta_0 r(d\log n)\middle| x,y,\theta\right) \\
=\ & O\left(P\left(|x^\top\delta + (\widetilde{x}-x)^\top\theta| < \zeta_0 r(d\log n)\middle| x\right)\right) \\
=\ & O\left(\min\left(\frac{\zeta_0 r(d\log n)}{\|x\|\xi/\sqrt{d}}, \frac{\zeta_0 r(d\log n)}{\xi_0 r}, 1\right)\right).
\end{aligned}$$

From the distribution of $x$, with probability tending to one over the choice of $S$,

$$\begin{aligned}
& \mathbb{E}\left[\min\left(\frac{\zeta_0 r(d\log n)}{\|x\|\xi/\sqrt{d}}, \frac{\zeta_0 r(d\log n)}{\xi_0 r}, 1\right)\middle| S\right] \\
\leq\ & \mathbb{E}\left[\frac{\zeta_0 r(d\log n)}{\xi_0 r}1\{\|x\| \leq \zeta_0'\}\middle| S\right] + \mathbb{E}\left[\frac{\zeta_0 r(d\log n)1\{\|x\| > \zeta_0'\}}{\|x\|\xi/\sqrt{d}}\middle| S\right] \\
=\ & O\left(\frac{\zeta_0 r(d\log n)}{\xi_0 r}\left(\zeta_0'\right)^d + \frac{\zeta_0 r(d\log n)}{\zeta_0'\xi/\sqrt{d}}\right),
\end{aligned}$$

and take $\zeta_0' = (\xi_0 r\sqrt{d}/\xi)^{1/(d+1)}$ to reach the minimized upper bound as $O(\zeta_0 r(d\log n)(\xi_0 r\sqrt{d}/\xi)^{d/(d+1)})$.

When $\|\theta\| \leq Cr$, we first assume that $P(|x^\top\theta - y| \leq zr|S) = O(z)\ \forall z$ is correct to finish the main proof, and finally provide the proof of itself. From Assumption 1, we have $P(|x^\top\theta - y| \leq zr) = O(z)$. Since $\max\|x_i\| = O(\sqrt{d\log n})$ almost surely,

$$\begin{aligned}
& P\left(|(\widetilde{x}-x)^\top\theta + (\widetilde{x}-x)^\top\delta + (x^\top\theta - y) + x^\top\delta| < \zeta_0 r(d\log n)\middle|S\right) \\
\leq\ & P\left(|x^\top\theta - y| < \zeta_0 r(d\log n) + |(\widetilde{x}-x)^\top\theta + (\widetilde{x}-x)^\top\delta + x^\top\delta|\middle|S\right) \\
=\ & O\left(\frac{\zeta_0 r(d\log n) + L\xi_0 + \sqrt{d\log n}\xi}{r}\right) = O\left(\zeta_0(d\log n) + \xi_0\frac{L}{r} + \frac{\sqrt{d\log n}\xi_0 L}{r}\right).
\end{aligned}$$

To conclude, with probability tending to one over the generation of $S$, we have

$$P(E^c|S) = O\left(\left(\frac{\zeta_0}{\xi_0}\right)^d + \zeta_0(d\log n)\left(\frac{r\sqrt{d}}{L}\right)^{\frac{d}{d+1}} + \xi_0\frac{L}{r} + \frac{\sqrt{d\log n}\xi_0 L}{r}\right).$$

The last thing is to verify $P(|x^\top\theta - y| \leq \zeta_0 r|S) = O(\zeta_0)$ for any $\|\theta\| \leq r$. Using Bernstein inequality, we know that for any fixed $\theta$,

$$P\left(\sum_{i=1}^n 1\{|x_i^\top\theta - y_i| \leq \zeta_0 r\} - nP(|x^\top\theta - y| \leq \zeta_0 r) \geq t\right) \leq e^{-\frac{t^2}{n+t}}. \tag{7}$$

We construct some intervals and design a series of points in $B_2(0,r)$. For the interval $[-r,r]$, we equally divide it into $n^m$ sub-intervals and repeat this procedure on all the $d$ dimensions. Through this construction, there are $\Theta(n^{md})$ points in $B_2(0,r)$. Denote these points as $\alpha_i$ for $i = 1, ..., K$. A consequence of this construction is that, for any $\theta \in B_2(0,r)$, the nearest $\alpha_j$ to $\theta$ has distance less than $D = 2\sqrt{d}r/n^m$.

Taking $\{\alpha_j\}_{j=1,...,K}$ into (7), we obtain

$$P\left(\sup_{j\in[K]}\sum_{i=1}^n 1\{|x_i^\top\alpha_j - y_i| \leq \zeta_0 r\} - nP(|x^\top\alpha_j - y| \leq \zeta_0 r) \geq t\right) \leq Ke^{-\frac{t^2}{n+t}}.$$

For any $\theta$, denote $\alpha_k$ as the one in $\{\alpha_j\}_{j=1,...,K}$ such that $\|\theta - \alpha_k\|$ is minimized, then for a sample $(x_i, y_i)$,

$$\begin{aligned}
& |1\{|x_i^\top\theta - y_i| \leq \zeta_0 r\} - 1\{|x_i^\top\alpha_k - y_i| \leq \zeta_0 r\}| \\
=\ & 1\{|x_i^\top\theta - y_i| \leq \zeta_0 r, |x_i^\top\alpha_k - y_i| > \zeta_0 r\} + 1\{|x_i^\top\theta - y_i| > \zeta_0 r, |x_i^\top\alpha_k - y_i| \leq \zeta_0 r\} \\
\leq\ & 1\{|x_i^\top\alpha_k - y_i| - \|x_i\|\|\theta - \alpha_k\| \leq \zeta_0 r, |x_i^\top\alpha_k - y_i| > \zeta_0 r\} \\
& +1\{|x_i^\top\alpha_k - y_i| + \|x_i\|\|\theta - \alpha_k\| > \zeta_0 r, |x_i^\top\alpha_k - y_i| \leq \zeta_0 r\} \\
\leq\ & 1\{|x_i^\top\alpha_k - y_i| - \|x_i\|D \leq \zeta_0 r, |x_i^\top\alpha_k - y_i| > \zeta_0 r\} \\
& +1\{|x_i^\top\alpha_k - y_i| + \|x_i\|D > \zeta_0 r, |x_i^\top\alpha_k - y_i| \leq \zeta_0 r\} \\
\leq\ & 1\{\zeta_0 r - \|x_i\|D \leq |x_i^\top\alpha_k - y_i| \leq \zeta_0 r + \|x_i\|D\}.
\end{aligned}$$

Since $\max_i\|x_i\| = O(\sqrt{d\log n})$ almost surely, we can further expand the above formula into

$$\begin{aligned}
& 1\{\zeta_0 r - \|x_i\|D \leq |x_i^\top\alpha_k - y_i| \leq \zeta_0 r + \|x_i\|D\} \\
\leq\ & 1\{\zeta_0 r - Dc\sqrt{d\log n} \leq |x_i^\top\alpha_k - y_i| \leq \zeta_0 r + Dc\sqrt{d\log n}, \|x_i\| \leq c\sqrt{d\log n}\} + 1\{\|x_i\| > c\sqrt{d\log n}\} \\
\leq\ & 1\{\zeta_0 r - Dc\sqrt{d\log n} \leq |x_i^\top\alpha_k - y_i| \leq \zeta_0 r + Dc\sqrt{d\log n}\} + 1\{\|x_i\| > c\sqrt{d\log n}\}.
\end{aligned}$$

As a result,

$$\begin{aligned}
& \left|\sum_{i=1}^n 1\{|x_i^\top\theta - y_i| \leq \zeta_0 r\} - \sum_{i=1}^n 1\{|x_i^\top\alpha_k - y_i| \leq \zeta_0 r\}\right| \\
\leq\ & \sum_{i=1}^n 1\{\zeta_0 r - Dc\sqrt{d\log n} \leq |x_i^\top\alpha_k - y_i| \leq \zeta_0 r + Dc\sqrt{d\log n}\} + n1\{\max_i\|x_i\| > c\sqrt{d\log n}\},
\end{aligned}$$

where $n1\{\max_i \|x_i\| > c\sqrt{d\log n}\} = 0$ almost surely, and

$$P\left(\sup_j \sum_{i=1}^n 1\{|x_i^\top \alpha_j - y_i| \in \zeta_0 r \pm Dc\sqrt{d\log n}\} - nP\{|x^\top \alpha_j - y| \in \zeta_0 r \pm Dc\sqrt{d\log n}\} \geq t\right) \leq Ke^{-\frac{t^2}{n+t}}.$$

Consequently, rewrite $k$ as $k(\theta)$, we have

$$\begin{aligned}
& P\left(\sup_\theta \sum_{i=1}^n 1\{|x_i^\top \theta - y_i| \leq \zeta_0 r\} - nP(|x^\top \theta - y| \leq \zeta_0 r) \geq t\right) \\
\leq\ & P\Bigg(\sup_\theta \left[\sum_{i=1}^n 1\{|x_i^\top \alpha_{k(\theta)} - y_i| \leq \zeta_0 r\} - nP(|x^\top \alpha_{k(\theta)} - y| \leq \zeta_0 r)\right] \\
& \quad + \left[\sum_{i=1}^n 1\{|x_i^\top \alpha_{k(\theta)} - y_i| \in \zeta_0 r \pm Dc\sqrt{d\log n}\} - nP\{|x^\top \alpha_j - y| \in \zeta_0 r \pm Dc\sqrt{d\log n}\}\right] \\
& \quad + nP\{|x^\top \alpha_{k(\theta)} - y| \in \zeta_0 r \pm Dc\sqrt{d\log n}\} \\
& \quad + nP(|x^\top \alpha_{k(\theta)} - y| \leq \zeta_0 r) - nP(|x^\top \theta - y| \leq \zeta_0 r) \geq t\Bigg) \\
\leq\ & P\Bigg(\sup_j \left[\sum_{i=1}^n 1\{|x_i^\top \alpha_j - y_i| \leq \zeta_0 r\} - nP(|x^\top \alpha_j - y| \leq \zeta_0 r)\right] \\
& \quad + \left[\sum_{i=1}^n 1\{|x_i^\top \alpha_j - y_i| \in \zeta_0 r \pm Dc\sqrt{d\log n}\} - nP\{|x^\top \alpha_j - y| \in \zeta_0 r \pm Dc\sqrt{d\log n}\}\right] \\
& \quad + nP\{|x^\top \alpha_j - y| \in \zeta_0 r \pm Dc\sqrt{d\log n}\} + nP(|x^\top \alpha_j - y| \leq \zeta_0 r) \geq t\Bigg) \\
\leq\ & P\Bigg(\sup_j \left[\sum_{i=1}^n 1\{|x_i^\top \alpha_j - y_i| \leq \zeta_0 r\} - nP(|x^\top \alpha_j - y| \leq \zeta_0 r)\right] \\
& \quad + nP\{|x^\top \alpha_j - y| \in \zeta_0 r \pm Dc\sqrt{d\log n}\} + nP(|x^\top \alpha_j - y| \leq \zeta_0 r) \geq \frac{t}{2}\Bigg) \\
& + P\Bigg(\sup_j \left[\sum_{i=1}^n 1\{|x_i^\top \alpha_j - y_i| \in \zeta_0 r \pm Dc\sqrt{d\log n}\} - nP\{|x^\top \alpha_j - y| \in \zeta_0 r \pm Dc\sqrt{d\log n}\}\right] \\
& \quad + nP\{|x^\top \alpha_j - y| \in \zeta_0 r \pm Dc\sqrt{d\log n}\} + nP(|x^\top \alpha_j - y| \leq \zeta_0 r) \geq \frac{t}{2}\Bigg).
\end{aligned}$$

Denote $\gamma = \sup_j P\{|x^\top \alpha_j - y| \in \zeta_0 r \pm Dc\sqrt{d\log n}\} + P(|x^\top \alpha_j - y| \leq \zeta_0 r)$, then

$$P\left(\sup_{\theta \in B_2(0,r)} \frac{1}{n}\sum_{i=1}^n 1\{|x_i^\top \theta - y_i| \leq \zeta_0 r\} - P(|x^\top \theta - y| \leq \zeta_0 r) \geq t\right) \leq 2K\exp\left\{-\frac{n[(t/2-\gamma)^+]^2}{1+(t/2-\gamma)^+}\right\}.$$

Recall that $K = \Theta(n^{md})$ and $D = 2\sqrt{d}r/n^m$, thus $\gamma = O\left(\sqrt{d}/n^m + \zeta_0\right)$. Taking $m$ as a constant such that $\zeta_0 \gg \sqrt{d}/n^m$, and $n\zeta_0$ grows polynomially in $n$, we have with probability tending to one over the generation of $S$, for any $\theta \in B_2(0,r)$

$$\frac{1}{n}\sum_{i=1}^n 1\{|x_i^\top \theta - y_i| \leq \zeta_0 r\} = O(\zeta_0).$$

□

**Lemma 6.** *Under the same conditions as Lemma 5, with probability tending to one over the choice of $S$,*

$$\mathbb{E}[g(\widetilde{x}, y, \theta + \delta)^\top (\theta - \bar{\theta})|S] \geq R_S(\theta) - R_S(\bar{\theta}) + O(\xi L^*).$$

*Proof of Lemma 6.* Since the adversarial loss is a convex function in both $\theta$ and $x$, and is smooth in $x$, we have

$$\mathbb{E}[g(\widetilde{x},y,\theta+\delta)^\top(\theta-\bar{\theta})|S] \geq \mathbb{E}[l(f_{\theta+\delta}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)),y)|S] - \mathbb{E}[l(f_{\bar{\theta}}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)),y)|S].$$

To quantify the error introduced by $\widetilde{x}$, we have

$$\begin{aligned}
& \mathbb{E}[l(f_{\theta+\delta}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)),y)|\delta,S] \\
= \;& \mathbb{E}\left[(y-\widetilde{x}^\top(\theta+\delta))^2+\epsilon^2\|\theta+\delta\|^2+2\epsilon\|\theta+\delta\||y-\widetilde{x}^\top(\theta+\delta)|\middle|\delta,S\right] \\
\geq \;& \mathbb{E}\left[(y-x^\top(\theta+\delta))^2+((\widetilde{x}-x)(\theta+\delta))^2+\epsilon^2\|\theta+\delta\|^2+2\epsilon\|\theta+\delta\||y-x^\top(\theta+\delta)|-2\epsilon\|\theta+\delta\||(\widetilde{x}-x)^\top(\theta+\delta)|\right. \\
= \;& \mathbb{E}[l(f_{\theta+\delta}(x+A_\epsilon(f,x,y)),y)|\delta,S]+O(\xi_0^2r^2)+O(\xi_0r^2).
\end{aligned}$$

Similarly we can obtain a bound for $\mathbb{E}[l(f_{\bar{\theta}}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)),y)|S]$.

Finally, from the distribution of $\delta$, we have

$$\mathbb{E}[l(f_{\theta+\delta}(x+A_\epsilon(f,x,y)),y)|S] = \mathbb{E}[l(f_\theta(x+A_\epsilon(f,x,y)),y)|S]+O(\xi L).$$

From the definition of $L^*$, we know that $r=O(L^*)$ and $L=O(L^*)$.

Consequently, aggregating all the above results, we have

$$\begin{aligned}
\mathbb{E}[g(\widetilde{x},y,\theta+\delta)^\top(\theta-\bar{\theta})|S] \;&\geq\; \mathbb{E}[l(f_{\theta+\delta}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)),y)|S]-\mathbb{E}[l(f_{\bar{\theta}}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)),y)|S] \\
&=\; \mathbb{E}[l(f_{\theta+\delta}(x+A_\epsilon(f,x,y)),y)|\delta,S]-\mathbb{E}[l(f_{\bar{\theta}}(x+A_\epsilon(f,x,y)),y)|\delta,S]+O(\xi L^*) \\
&=\; \mathbb{E}[l(f_\theta(x+A_\epsilon(f,x,y)),y)|S]-\mathbb{E}[l(f_{\bar{\theta}}(x+A_\epsilon(f,x,y)),y)|\delta,S]+O(\xi L^*) \\
&=\; R_S(\theta)-R_S(\bar{\theta})+O(\xi L^*).
\end{aligned}$$

□

*Proof of Lemma 2 and Theorem 2.* We use $(L,B,l)$ rather than $(L^*,B^*,l^*)$. The latter can is just an simple upper bound after obtain results regarding to the former one.

To show the stability of SGD, denoting $\Delta_t=\theta_1^{(t)}-\theta_2^{(t)}$, we have

$$\begin{aligned}
\|\Delta_t\|^2 \;&\leq\; \left\|\theta_1^{(t-1)}-\theta_2^{(t-1)}-\eta_t\left(\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(\widehat{z}_{i_t}^1),y_{i_t}^1)-\nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(\widehat{z}_{i_t}^2),y_{i_t}^2)\right)\right\|^2 \\
&=\; \|\Delta_{t-1}\|^2+\eta_t^2\left\|\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(\widehat{z}_{i_t}^1),y_{i_t}^1)-\nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(\widehat{z}_{i_t}^2),y_{i_t}^2)\right\|^2 \\
&\quad -2\eta_t\Delta_{t-1}^\top\left(\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(\widehat{z}_{i_t}^1),y_{i_t}^1)-\nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(\widehat{z}_{i_t}^2),y_{i_t}^2)\right).
\end{aligned}$$

Further,

$$\begin{aligned}
& \eta_t^2\left\|\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(\widehat{z}_{i_t}^1-z_{i_t}^1+z_{i_t}^1,y_{i_t}^1)-\nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(\widehat{z}_{i_t}^2-z_{i_t}^2+z_{i_t}^2,y_{i_t}^2)\right\|^2 \\
\leq\; & \eta_t^2\left(2\kappa\Delta\varepsilon+\left\|\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(z_{i_t}^1),y_{i_t}^1)-\nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(z_{i_t}^2),y_{i_t}^2)\right\|\right)^2, \\
\leq\; & 2\eta_t^2\left(4\kappa^2\Delta\varepsilon^2+\left\|\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(z_{i_t}^1),y_{i_t}^1)-\nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(z_{i_t}^2),y_{i_t}^2)\right\|^2\right)
\end{aligned}$$

and

$$\begin{aligned}
& -2\eta_t\Delta_{t-1}^\top\left(\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(\widehat{z}_{i_t}^1),y_{i_t}^1)-\nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(\widehat{z}_{i_t}^2),y_{i_t}^2)\right) \\
\leq\; & 4\eta_t\|\Delta_{t-1}\|\kappa\Delta\varepsilon-2\eta_t\Delta_{t-1}^\top\left(\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(z_{i_t}^1),y_{i_t}^1)-\nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(z_{i_t}^2),y_{i_t}^2)\right).
\end{aligned}$$

Therefore, taking conditional expectation, we obtain

$$
\begin{aligned}
& \mathbb{E}\left[\eta_t^2 \left\|\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(\widehat{z}_{i_t}^1 - z_{i_t}^1 + z_{i_t}^1, y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(\widehat{z}_{i_t}^2 - z_{i_t}^2 + z_{i_t}^2, y_{i_t}^2)\right\|^2 \middle| \Delta_{t-1}\right] \\
& -\mathbb{E}\left[2\eta_t \Delta_{t-1}^\top \left(\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(\widehat{z}_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(\widehat{z}_{i_t}^2), y_{i_t}^2)\right) \middle| \Delta_{t-1}\right] \\
\leq\ & 8\eta_t^2\kappa^2\Delta\varepsilon^2 + 4\eta_t\|\Delta_{t-1}\|\kappa\Delta\varepsilon + 2\eta_t^2\mathbb{E}\left[\left\|\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(z_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(z_{i_t}^2), y_{i_t}^2)\right\|^2 \middle| \Delta_{t-1}\right] \\
& -2\eta_t\Delta_{t-1}^\top\mathbb{E}\left[\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(z_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(z_{i_t}^2), y_{i_t}^2)\middle| \Delta_{t-1}\right] \\
\leq\ & 8\eta_t^2\kappa^2\Delta\varepsilon^2 + 4\eta_t\|\Delta_{t-1}\|\kappa\Delta\varepsilon + 2\eta_t^2(2L)^2\frac{1}{n} \\
& +2\eta_t^2\mathbb{E}\left[\left\|\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(z_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(z_{i_t}^2), y_{i_t}^2)\right\|^2 \middle| \Delta_{t-1}, (x_{i_t}^1, y_{i_t}^1) = (x_{i_t}^2, y_{i_t}^2)\right] \\
& +2\eta_t\|\Delta_{t-1}\|(2L)\frac{1}{n} \\
& -2\eta_t\Delta_{t-1}^\top\mathbb{E}\left[\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(z_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(z_{i_t}^2), y_{i_t}^2)\middle| \Delta_{t-1}, (x_{i_t}^1, y_{i_t}^1) = (x_{i_t}^2, y_{i_t}^2)\right] \\
\leq\ & 8\eta_t^2\kappa^2\Delta\varepsilon^2 + 4\eta_t\|\Delta_{t-1}\|\kappa\Delta\varepsilon + 2\eta_t^2(2L)^2\frac{1}{n} + 2\eta_t^2(2L)^2P(E^c|S) \\
& +2\eta_t^2\mathbb{E}\left[\left\|\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(z_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(z_{i_t}^2), y_{i_t}^2)\right\|^2 \middle| \Delta_{t-1}, (x_{i_t}^1, y_{i_t}^1) = (x_{i_t}^2, y_{i_t}^2), E\right] \\
& +2\eta_t\|\Delta_{t-1}\|(2L)\frac{1}{n} + 2\eta_t\|\Delta_{t-1}\|(2L)P(E^c|S) \\
& -2\eta_t\Delta_{t-1}^\top\mathbb{E}\left[\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(z_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(z_{i_t}^2), y_{i_t}^2)\middle| \Delta_{t-1}, (x_{i_t}^1, y_{i_t}^1) = (x_{i_t}^2, y_{i_t}^2), E\right].
\end{aligned}
$$

Under $E$, since $l(f_{\theta_1^{(t-1)+\delta}}(z_{i_t}^1), y_{i_t}^1)$ is convex, following (A.1) of Hardt et al. (2016), we have

$$
\begin{aligned}
& -2\eta_t\Delta_{t-1}^\top\mathbb{E}\left[\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(z_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(z_{i_t}^2), y_{i_t}^2)\middle| \Delta_{t-1}, (x_{i_t}^1, y_{i_t}^1) = (x_{i_t}^2, y_{i_t}^2), E\right] \\
\leq\ & -2\eta_t\frac{\zeta}{B}\mathbb{E}\left[\left\|\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(z_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(z_{i_t}^2), y_{i_t}^2)\right\|^2 \middle| \Delta_{t-1}, (x_{i_t}^1, y_{i_t}^1) = (x_{i_t}^2, y_{i_t}^2), E\right],
\end{aligned}
$$

thus we obtain Lemma 2 for general choices of $\eta_t$.

$$
\begin{aligned}
& \mathbb{E}[\|\theta_1^{(t)} - \theta_1^{(t)}\|^2|S] \\
\leq\ & \left(1 + 2\eta_t^2\frac{B^2}{\zeta^2}1\{\eta_t \geq \frac{\zeta}{B}\}\right)\|\theta_1^{(t-1)} - \theta_2^{(t-1)}\|^2 + 8\eta_t^2\kappa^2\Delta\varepsilon^2 + 2\eta_t^2(2L)^2\frac{1}{n} + 2\eta_t^2(2L)^2P(E^c|S) \\
& +4\eta_t\|\Delta_{t-1}\|\kappa\Delta\varepsilon + 2\eta_t\|\Delta_{t-1}\|(2L)\frac{1}{n} + 2\eta_t\|\Delta_{t-1}\|(2L)P(E^c|S).
\end{aligned} \tag{8}
$$

When taking $\eta_t \leq \zeta/B$, we have

$$
\begin{aligned}
& 2\eta_t^2\mathbb{E}\left[\left\|\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(z_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(z_{i_t}^2), y_{i_t}^2)\right\|^2 \middle| \Delta_{t-1}, (x_{i_t}^1, y_{i_t}^1) = (x_{i_t}^2, y_{i_t}^2), E\right] \\
& -2\eta_t\Delta_{t-1}^\top\mathbb{E}\left[\nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(z_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(z_{i_t}^2), y_{i_t}^2)\middle| \Delta_{t-1}, (x_{i_t}^1, y_{i_t}^1) = (x_{i_t}^2, y_{i_t}^2), E\right] \\
\leq\ & 0.
\end{aligned}
$$

Therefore,

$$\begin{aligned}
&\mathbb{E}\left[\eta_t^2\left\|\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(\widehat{z}_{i_t}^1 - z_{i_t}^1 + z_{i_t}^1, y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(\widehat{z}_{i_t}^2 - z_{i_t}^2 + z_{i_t}^2, y_{i_t}^2)\right\|^2 \middle| \Delta_{t-1}\right] \\
&-\mathbb{E}\left[2\eta_t\Delta_{t-1}^\top\left(\nabla_\theta l(f_{\theta_1^{(t-1)}+\delta}(\widehat{z}_{i_t}^1), y_{i_t}^1) - \nabla_\theta l(f_{\theta_2^{(t-1)}+\delta}(\widehat{z}_{i_t}^2), y_{i_t}^2)\right)\middle|\Delta_{t-1}\right] \\
\leq\ & 8\eta_t^2\kappa^2\Delta\varepsilon^2 + 2\eta_t^2(2L)^2\frac{1}{n} + 2\eta_t^2(2L)^2P(E^c) \\
&+4\eta_t\|\Delta_{t-1}\|\kappa\Delta\varepsilon + 2\eta_t\|\Delta_{t-1}\|(2L)\frac{1}{n} + 2\eta_t\|\Delta_{t-1}\|(2L)P(E^c),
\end{aligned}$$

and

$$\begin{aligned}
\mathbb{E}[\|\Delta_t\|^2|\Delta_{t-1}] \leq\ & \|\Delta_{t-1}\|^2 + 8\eta_t^2\kappa^2\Delta\varepsilon^2 + 2\eta_t^2(2L)^2\frac{1}{n} + 2\eta_t^2(2L)^2P(E^c) \\
&+4\eta_t\|\Delta_{t-1}\|\kappa\Delta\varepsilon + 2\eta_t\|\Delta_{t-1}\|(2L)\frac{1}{n} + 2\eta_t\|\Delta_{t-1}\|(2L)P(E^c),
\end{aligned}$$

which leads to

$$\mathbb{E}^2\|\Delta_T\| \leq \mathbb{E}\|\Delta_T\|\sum_{t=t_0}^{T}\left[4\eta_t\kappa\Delta\varepsilon + \frac{4L\eta_t}{n} + 4\eta_t LP(E^c)\right] + \sum_{t=t_0}^{T} 8\eta_t^2\kappa^2\Delta\varepsilon^2 + 8\frac{L^2\eta_t^2}{n} + 8\eta_t^2L^2P(E^c).$$

Reordering some terms in the above inequality, we get

$$\left(\mathbb{E}\|\Delta_T\| - \sum_{t=t_0}^{T}\left[4\eta_t\kappa\Delta\varepsilon + \frac{4L\eta_t}{n} + 4\eta_t LP(E^c)\right]\right)^2 \leq \sum_{t=t_0}^{T} 8\eta_t^2\kappa^2\Delta\varepsilon^2 + 8\frac{L^2\eta_t^2}{n} + 8\eta_t^2L^2P(E^c),$$

so finally we obtain

$$\begin{aligned}
\mathbb{E}\|\Delta_T\| &= O\left(\left[\Delta\varepsilon + \sqrt{P(E^c)} + \sqrt{\frac{1}{n}}\right]\sqrt{\sum_{t=t_0}^{T}\eta_t^2}\right) + O\left(\left[\Delta\varepsilon + \frac{1}{n} + P(E^c)\right]\sum_{t=t_0}^{T}\eta_t\right) \\
&= O\left(\left[\sqrt{P(E^c)} + \sqrt{\frac{1}{n}}\right]\sqrt{\sum_{t=t_0}^{T}\eta_t^2}\right) + O\left(\left[\Delta\varepsilon + \frac{1}{n} + P(E^c)\right]\sum_{t=t_0}^{T}\eta_t\right).
\end{aligned}$$

The proof for GD is similar. □

*Proof of Theorem 3.* We first assume the analytical solution of attack exists and there is no attack error in linear regression problem, then discussing how to consider the attack error and for the other loss functions.

The updating rule of SGD leads to

$$\|\theta^{(t)} - \bar{\theta}\|^2 \leq \|\theta^{(t-1)} - \bar{\theta} - \eta_t g_t\|^2 \leq \|\theta^{(t-1)} - \bar{\theta}\|^2 - 2\eta_t g_t^\top(\theta^{(t-1)} - \bar{\theta}) + \eta_t^2L^2.$$

Taking expectation and move some terms, it becomes

$$\mathbb{E}g_t^\top(\theta^{(t-1)} - \bar{\theta}) \leq \frac{1}{2\eta_t}\mathbb{E}\|\theta^{(t-1)} - \bar{\theta}\|^2 - \frac{1}{2\eta_t}\mathbb{E}\|\theta^{(t)} - \bar{\theta}\|^2 + \frac{1}{2}\eta_t L^2.$$

Taking average over $t = 1$ to $T$, we have

$$\begin{aligned}
\frac{1}{T}\mathbb{E}\left[\sum_{t=1}^{T} g_t^\top(\theta^{(t-1)} - \bar{\theta})\right] &\leq \frac{1}{2T}\mathbb{E}\left[\sum_{t=1}^{T}\frac{1}{\eta_t}\|\theta^{(t-1)} - \bar{\theta}\|^2 - \frac{1}{\eta_t}\|\theta^{(t)} - \bar{\theta}\|^2\right] + \frac{L^2}{2T}\sum_{t=1}^{T}\eta_t \\
&= \frac{\mathbb{E}\|\theta^{(0)} - \bar{\theta}\|^2}{2\eta_1 T} - \frac{\mathbb{E}\|\theta^{(T)} - \bar{\theta}\|^2}{2\eta_T T} \\
&\quad + \frac{1}{2T}\mathbb{E}\left[\sum_{t=1}^{T-1}\left(\frac{1}{\theta_{t+1}} - \frac{1}{\theta_t}\right)\|\theta^{(t)} - \bar{\theta}\|^2\right] + \frac{L^2}{2T}\sum_{t=1}^{T}\eta_t.
\end{aligned}$$

Finally, since $R_S$ is a convex function, based on Lemma 6, we have

$$\frac{1}{T}\mathbb{E}\left[\sum_{t=1}^{T} g_t^\top(\theta^{(t-1)}-\bar{\theta})\right] \geq \frac{1}{T}\mathbb{E}\sum_{t=1}^{T} R_S(\theta^{(t)}) - R_S(\bar{\theta}) + O(\xi L^*) \geq \mathbb{E}\left[\min_{t=1,\ldots,T} R_S(\theta^{(t)}) - R_S(\bar{\theta})\right] + O(\xi L^*).$$

The above proof assumes that attack has no error. To count for the attack error, since $\triangledown R_S$ is Lipschitz in $x$, denoting $\widehat{g}_t$ as the gradient approximated, then $\widehat{g}_t = g_t + O(\kappa\Delta\varepsilon)$. So an additional $O(\kappa r\Delta\varepsilon)$ is introduced.

Since fixing the $n$ samples, taking expectation in SGD is reduced to GD, the above result also holds. □

### D.3 Lemmas for Other Loss

**Lemma 7** (Smoothed Hinge Loss)**.** *For smoothed hinge loss, $H(x)$ is defined as a strictly monotone function in $x$ with $H(x) = 1$ when $x \geq 1$ and $H(x) = 0$ when $x \leq -1$, and $xH(x)$ is convex. The derivative $H'$ satisfies $H'(-1) = H'(1) = 0$, and $H''$ is finite. define $E(\theta, \delta, \widetilde{x}, y) = 1\{\|\theta+\delta\| \geq \zeta\}$, then $E = 1$ implies that $g_t$ is Lipschitz with $B/\max(h, \zeta)$. Assume $h$ is a fixed constant. In addition, when $r/(\sqrt{d\log n}) \to 0$, with probability tending to one over the choice of $S$,*

$$\mathbb{E}[g(\widetilde{x}, y, \theta+\delta)^\top(\theta-\bar{\theta})|S] \geq R_S(\theta) - R_S(\bar{\theta}) + O(\xi L^*).$$

*Proof of Lemma 7.* Given $(\theta, x, y)$, we have

$$l(f_\theta(x), y) = (1 - y(x^\top\theta))H\left(\frac{1-yx^\top\theta}{h}\right),$$

for $y \in \{\pm 1\}$. Thus the attack is

$$A = \begin{cases} -y\epsilon\frac{\theta}{\|\theta\|} & \text{if } 1 - y(x^\top\theta) > -\epsilon \\ \text{any } z \in B_2(x,\epsilon) & \text{otherwise} \end{cases}.$$

The adversarial risk becomes

$$l(f_\theta(x+A), y) = (1 - y((x+A)^\top\theta))H\left(\frac{1-y(x+A)^\top\theta}{h}\right),$$

and the gradient becomes

$$g(x,y,\theta) \quad = \quad -y(x+A)\left[H\left(1-y(x^\top\theta)+\epsilon\|\theta\|\right) + \frac{(1-y(x^\top\theta)+\epsilon\|\theta\|)}{h}H'\left(\frac{1-y(x^\top\theta)+\epsilon\|\theta\|}{h}\right)\right]$$

Since $H$ and $H'$ are differentiable, for any $g$, when $\|\theta\| \geq \zeta$, for any other $\theta' \neq \mathbf{0}$.

$$\|g(x,y,\theta) - g(x,y,\theta')\| \leq \frac{B}{\zeta}\|\theta - \theta'\|.$$

In terms of the expectation of $g(\widetilde{x}, y, \theta+\delta)$, since $l(f_{\theta+\delta}(\widetilde{x} + A_\epsilon(f, \widetilde{x}, y)), y)$ is convex, we have

$$\mathbb{E}g(\widetilde{x}, y, \theta+\delta)^\top(\theta-\bar{\theta}) \geq \mathbb{E}l(f_{\theta+\delta}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)), y) - \mathbb{E}l(f_{\bar{\theta}}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)), y).$$

Further, for any $(x, y, \theta)$,

$$\begin{aligned}
& \mathbb{E}l(f_{\theta+\delta}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)), y) \\
= \quad & \mathbb{E}\left[(1-y(\widetilde{x}^\top(\theta+\delta)) + \epsilon\|\theta+\delta\|)H\left(\frac{(1-y(\widetilde{x}^\top(\theta+\delta)) + \epsilon\|\theta+\delta\|)}{h}\right)\right] \\
= \quad & (1-y(x^\top\theta)+\epsilon\|\theta\|)H\left(\frac{(1-y(x^\top(\theta+\delta)) + \epsilon\|\theta\|)}{h}\right) \\
& +\mathbb{E}\left[(1-y(\widetilde{x}^\top(\theta+\delta)) + \epsilon\|\theta+\delta\|)\left(H\left(\frac{(1-y(\widetilde{x}^\top(\theta+\delta)) + \epsilon\|\theta+\delta\|)}{h}\right) - H\left(\frac{(1-y(x^\top\theta)+\epsilon\|\theta\|)}{h}\right)\right)\right] \\
= \quad & l(f_\theta(x+A_\epsilon(f,x,y)), y) \\
& +\mathbb{E}\left[(1-y(\widetilde{x}^\top(\theta+\delta)) + \epsilon\|\theta+\delta\|)\left(H\left(\frac{(1-y(\widetilde{x}^\top(\theta+\delta)) + \epsilon\|\theta+\delta\|)}{h}\right) - H\left(\frac{(1-y(x^\top\theta)+\epsilon\|\theta\|)}{h}\right)\right)\right].
\end{aligned}$$

Furthermore,

$$\begin{aligned}\left|(1-y(\widetilde{x}^\top(\theta+\delta))+\epsilon\|\theta+\delta\|)\right| &\leq \left|(1-y(x^\top(\theta))+\epsilon\|\theta\|)\right|+\epsilon\|\delta\|+\left|y((\widetilde{x}-x)^\top(\theta+\delta))\right|+\left|y(x^\top\delta)\right| \\ &\leq L^*+\epsilon\|\delta\|+\left|(\widetilde{x}-x)^\top(\theta+\delta)\right|+\left|x^\top\delta\right| \\ &= L^*(1+o_p(1)).\end{aligned}$$

From the definition of $H$, we have

$$\begin{aligned}& H\left(\frac{(1-y(\widetilde{x}^\top(\theta+\delta))+\epsilon\|\theta+\delta\|)}{h}\right)-H\left(\frac{(1-y(x^\top\theta)+\epsilon\|\theta\|)}{h}\right) \\ =\ & \left[\frac{-y\widetilde{x}^\top(\theta+\delta)+yx^\top\theta+\epsilon\|\theta+\delta\|-\epsilon\|\theta\|}{h}\right]H'\left(\frac{(1-y(x^\top\theta)+\epsilon\|\theta\|)}{h}\right) \\ & +\Bigg\{H\left(\frac{(1-y(\widetilde{x}^\top(\theta+\delta))+\epsilon\|\theta+\delta\|)}{h}\right)-H\left(\frac{(1-y(x^\top\theta)+\epsilon\|\theta\|)}{h}\right) \\ & -\left[\frac{-y\widetilde{x}^\top(\theta+\delta)+yx^\top\theta+\epsilon\|\theta+\delta\|-\epsilon\|\theta\|}{h}\right]H'\left(\frac{(1-y(x^\top\theta)+\epsilon\|\theta\|)}{h}\right)\Bigg\} \\ =\ & O_p(\xi+r\xi_0).\end{aligned}$$

Consequently, when $r/\max_i\|x_i\|\to 0$, $r\xi_0=O(\xi L^*)$, and

$$\mathbb{E}l(f_{\theta+\delta}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)),y)=l(f_\theta(x+A_\epsilon(f,x,y)),y)+O(\xi L^*). \tag{9}$$

□

**Lemma 8.** *For Logistic regression, denote $E(\theta,\delta,\widetilde{x},y)=1\{\|\theta+\delta\|\geq\zeta\}$. Then $E=1$ implies that $\triangledown_{\theta+\delta}l(f_{\theta+\delta}(\widetilde{x}+A),y)$ is $1/\zeta_0$-Lipschitz.*

*In addition,*

$$\mathbb{E}g(\widetilde{x},y,\theta+\delta)^\top(\theta-\bar{\theta})\geq R_S(\theta)-R_S(\bar{\theta})+O(\xi L^*).$$

*Proof of Lemma 8.* For each data $(x,y)$, $l(f(x,\theta),y)=-1\{y=1\}\log(p(x^\top\theta))-1\{y=-1\}\log(1-p(x^\top\theta))$, where

$$p(x^\top\theta)=\frac{1}{1+e^{-x^\top\theta}}.$$

Taking gradient w.r.t $\theta$, we obtain

$$\triangledown_\theta l(f_\theta(x),y)=-1\{y=1\}\frac{xp'(x^\top\theta)}{p(x^\top\theta)}+1\{y=-1\}\frac{xp'(x^\top\theta)}{1-p(x^\top\theta)},$$

where

$$p'(x^\top\theta)=\frac{e^{-x^\top\theta}}{(1+e^{-x^\top\theta})^2}.$$

When $\theta\neq\mathbf{0}$, the attack is

$$A=\begin{cases}-\epsilon\frac{\theta}{\|\theta\|} & y=1\\ \epsilon\frac{\theta}{\|\theta\|} & y=-1\end{cases}, \tag{10}$$

thus

$$\triangledown_\theta l(f_\theta(x+A),y)=-1\{y=1\}\frac{(x-\frac{\epsilon\theta}{\|\theta\|})p'(x^\top\theta-\epsilon\|\theta\|)}{p(x^\top\theta-\epsilon\|\theta\|)}+1\{y=-1\}\frac{(x+\frac{\epsilon\theta}{\|\theta\|})p'(x^\top\theta+\epsilon\|\theta\|)}{1-p(x^\top\theta+\epsilon\|\theta\|)}.$$

The above representation indicates that $\|\theta\|\geq\zeta$ implies $\triangledown_\theta l(f_\theta(x+A),y)$ is $B/\zeta$-Lipschitz.

In terms of $\mathbb{E}g(\widetilde{x},y,\theta+\delta)^\top(\theta-\bar{\theta})$, since $l(f_{\theta+\delta}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)),y)$ is convex, we have

$$\mathbb{E}g(\widetilde{x},y,\theta+\delta)^\top(\theta-\bar{\theta})\geq\mathbb{E}l(f_{\theta+\delta}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)),y)-\mathbb{E}l(f_{\bar{\theta}}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)),y).$$

The remaining proof is similar as in Lemma 6 and Lemma 7. □

## E $\mathcal{L}_\infty$ Attack in Adversarial Training

As mentioned in the proof of Theorem 2 and 3, it only requires some conditions w.r.t Lipschitz continuous of the gradient as well as $\mathbb{E}g^\top(\theta-\widetilde{\theta})$. Therefore, we provide results of $\mathcal{L}_\infty$ adversarial training for linear regression setup.

In addition, based on the following lemmas, the value of $(L, L^*, \kappa, \kappa^*)$ is different from those in $\mathcal{L}_2$ adversarial training. Furthermore, the value of $B$ and $B^*$ are also enlarged for $\mathcal{L}_\infty$ adversarial training.

**Lemma 9.** *For linear regression, there exists some $(L^*, \kappa^*)$ such that, with probability tending to one over the choice of $S$, $L \le L^*$ and $\kappa \le \kappa^*$.*

*Proof of Lemma 9.*

$$\frac{1}{2}\nabla_\theta l(f_\theta(x), y) = x(x^\top\theta - y).$$

Then from the definition of $L$, when $\|\delta_x\|_\infty \le 2\epsilon$, $\|\delta_x\| \le 2\sqrt{d}\epsilon$.

$$\frac{1}{2}L = \max_{\theta\in B_2(0,r), i\in[n], \delta_x} \|x_i+\delta_x\||(x_i+\delta_x)^\top\theta - y_i| \le (\max_i \|x_i\| + 2\sqrt{d}\epsilon)^2 r + (\max_i \|x_i\||y_i| + 2\sqrt{d}\epsilon|y_i|).$$

In addition,

$$\begin{aligned}
& \frac{1}{2}\|\nabla_\theta l(f_\theta(x), y) - \nabla_\theta l(f_\theta(x+\delta_x), y)\| \\
= & \frac{1}{2}\left\|x(x^\top\theta - y) - x((x+\delta_x)^\top\theta - y) - \delta_x((x+\delta_x)^\top\theta - y)\right\| \\
= & \frac{1}{2}\left\|x\delta_x^\top\theta - \delta_x((x+\delta_x)^\top\theta - y)\right\| \\
= & \frac{1}{2}\left\|x\delta_x^\top\theta - \delta_x(x^\top\theta - y) - \delta_x\delta_x^\top\theta\right\| \\
\le & \frac{1}{2}\left(\|x\|\|\theta\|\|\delta_x\| + \|\delta_x\||x^\top\theta - y| + \|\delta_x\|^2\|\theta\|\right) \\
\le & \frac{1}{2}\left(\|x\|\|\theta\|\|\delta_x\| + \|\delta_x\||x^\top\theta - y| + \|\delta_x\|2\sqrt{d}\epsilon\|\theta\|\right).
\end{aligned}$$

Thus for a given set of data $S$,

$$\frac{1}{2}\kappa = \max_{\theta\in B_2(0,r), i\in[n]}\left[(\|x_i\| + 2\sqrt{d}\epsilon)\|\theta\| + |x_i^\top\theta - y_i|\right] \le 2(\max_i\|x_i\| + \sqrt{d}\epsilon)r + \max_i|y_i|.$$

From the distribution of $x$, we know that $\max_i\|x_i\| = O(\sqrt{d\log n})$ almost surely. In addition, $\mathbb{E}\|x\||y|$ and $\mathbb{E}|y|$ are finite, thus $\max_i\|x_i\||y_i|$ and $\max_i|y_i|$ are some functions of $n$ as well. □

**Lemma 10.** *For linear regression, denote $\zeta = L\zeta_0$ for some $\zeta_0/\xi_0 \to 0$. Denote $E(\theta, \delta, \widetilde{x}, y) = 1\{\min_j|\theta+\delta|_j \ge \zeta/\sqrt{d}, |\widetilde{x}^\top(\theta+\delta) - y| \ge \zeta_0 r(d\log n)\}$, then $E = 1$ implies that $\nabla_\theta l(f_{\theta+\delta}(\widetilde{x}), y)$ is $\sqrt{d}/\zeta_0$-Lipschitz. Then uniformly for all $\theta$, with probability tending to one over the $n$ random samples, we have*

$$P(E^c(\theta, \delta, \widetilde{x}, y)|S) = o(1).$$

*Proof of Lemma 10.* We show that $E = 1$ implies that $\nabla_\theta l(f_{\theta+\delta}(\widetilde{x}), y)$ is $1/\zeta_0$-Lipschitz. The gradient of adversarial loss is

$$\begin{aligned}
\frac{1}{2}g(\widetilde{x}, y, \theta) = & \widetilde{x}(\widetilde{x}^\top(\theta+\delta) - y) + \epsilon^2\|\theta+\delta\|_1\text{sgn}(\theta+\delta) + \epsilon\text{sgn}(\theta+\delta)|y - \widetilde{x}^\top(\theta+\delta)| \\
& -\epsilon\widetilde{x}\|(\theta+\delta)\|_1\text{sgn}(y - \widetilde{x}^\top(\theta+\delta)).
\end{aligned}$$

When $\min_j|\theta+\delta|_j \ge \zeta/\sqrt{d}$, we have for any $\theta'$,

$$\frac{1}{\|\theta+\delta-\theta'\|^2}\|\text{sgn}(\theta+\delta) - \text{sgn}(\theta')\|^2 \le \lim_{\alpha\to 0^+}\frac{d}{\|\theta+\delta+\alpha(\theta+\delta)\|^2} \le \frac{d}{\zeta^2}.$$

When $|y-\widetilde{x}^\top(\theta+\delta)| \geq \gamma$ for some $\gamma$, this implies that the nearest $\theta'$ such that $\text{sgn}(y-\widetilde{x}^\top(\theta+\delta))$ gets changed satisfies $\|\theta'-(\theta+\delta)\| = \gamma/\|\widetilde{x}\|$. As a result,

$$\begin{aligned}
& \frac{1}{\|\theta+\delta-\theta'\|}\left\|\widetilde{x}\|\theta+\delta\|\text{sgn}(y-\widetilde{x}^\top(\theta+\delta)) - \widetilde{x}\|\theta'\|\text{sgn}(y-\widetilde{x}^\top\theta')\right\| \\
\leq\ & \frac{1}{\|\theta+\delta-\theta'\|}\left\|\widetilde{x}\|\theta+\delta\|_1\text{sgn}(y-\widetilde{x}^\top\theta') - \widetilde{x}\|\theta'\|_1\text{sgn}(y-\widetilde{x}^\top\theta')\right\| \\
& + \frac{1}{\|\theta+\delta-\theta'\|}\left\|\widetilde{x}\|\theta+\delta\|_1\text{sgn}(y-\widetilde{x}^\top(\theta+\delta)) - \widetilde{x}\|\theta+\delta\|_1\text{sgn}(y-\widetilde{x}^\top\theta')\right\| \\
\leq\ & \frac{\|\widetilde{x}\|\sqrt{d}\|\theta+\delta-\theta'\|}{\|\theta+\delta-\theta'\|} + \frac{\|\widetilde{x}\|\|\theta+\delta\|_1}{\|\theta+\delta-\theta'\|}\left|\text{sgn}(y-\widetilde{x}^\top(\theta+\delta)) - \text{sgn}(y-\widetilde{x}^\top\theta')\right| \\
\leq\ & \|\widetilde{x}\| + \frac{2\|\widetilde{x}\|^2\sqrt{d}}{\gamma}.
\end{aligned}$$

Take $\gamma = \zeta_0 r\|\widetilde{x}\|^2$ in the above inequality to obtain $\|\tilde{x}\| + 2\sqrt{d}/\zeta_0$-Lipschitz.

Now we turn to bound the probability of $E^c$.

$$P(E^c(\theta,\delta,\widetilde{x},y)|S) \leq P(\min_j|\theta+\delta|_j < \zeta/\sqrt{d}) + P(|\widetilde{x}^\top(\theta+\delta)-y| < \zeta_0 r(d\log n)|S)$$

For any $\theta$, we have

$$P(\min_j|\theta+\delta|_j < \zeta/\sqrt{d}|\theta) = O\left(1-\left(1-\frac{\zeta}{\xi}\right)^d\right) = O\left(\frac{d\zeta}{\xi}\right).$$

The remaining steps follows the same as in Lemma 5.

□

**Lemma 11.** *Under the same conditions as Lemma 5, with probability tending to one over the set of $n$ random samples,*

$$\mathbb{E}g(\widetilde{x},y,\theta+\delta)^\top(\theta-\bar{\theta}) \geq R_S(\theta) - R_S(\bar{\theta}) + O(\xi L^*).$$

*Proof of Lemma 11.* Since the adversarial loss is a convex function in both $\theta$ and $x$, and is smooth in $x$, we have

$$\mathbb{E}g(\widetilde{x},y,\theta+\delta)^\top(\theta-\bar{\theta}) \geq \mathbb{E}l(f_{\theta+\delta}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)),y) - \mathbb{E}l(f_{\bar{\theta}}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)),y).$$

To quantify the error introduced by $\widetilde{x}$, we have

$$\begin{aligned}
& \mathbb{E}[l(f_{\theta+\delta}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)),y)|\delta] \\
=\ & \mathbb{E}\left[(y-\widetilde{x}^\top(\theta+\delta))^2 + \epsilon^2\|\theta+\delta\|_1^2 + 2\epsilon\|\theta+\delta\|_1|y-\widetilde{x}^\top(\theta+\delta)|\middle|\delta\right] \\
\geq\ & \mathbb{E}\Big[(y-x^\top(\theta+\delta))^2 + ((\widetilde{x}-x)(\theta+\delta))^2 + \epsilon^2\|\theta+\delta\|_1^2 + 2\epsilon\|\theta+\delta\|_1|y-x^\top(\theta+\delta)| - 2\epsilon\|\theta+\delta\|_1|(\widetilde{x}-x)^\top(\theta+\delta)\Big] \\
=\ & \mathbb{E}[l(f_{\theta+\delta}(x+A_\epsilon(f,x,y)),y)|\delta] + O(\xi_0^2 r^2) + O(\xi_0 r^2\sqrt{d}).
\end{aligned}$$

Since $l$ is square loss and $f_\theta$ is the linear model, we have $r = O(L^*/\sqrt{d})$, thus

$$\mathbb{E}[l(f_{\theta+\delta}(\widetilde{x}+A_\epsilon(f,\widetilde{x},y)),y)|\delta] = \mathbb{E}[l(f_{\theta+\delta}(x+A_\epsilon(f,x,y)),y)|\delta] + O(\xi_0(L^*)^2/\sqrt{d}).$$

Finally, from the definition of $L$, we have

$$\mathbb{E}l(f_{\theta+\delta}(x+A_\epsilon(f,x,y)),y) = \mathbb{E}l(f_\theta(x+A_\epsilon(f,x,y)),y) + O(\xi L^*).$$

□