
On the Generalization Properties of Adversarial Training

Yue Xing
Purdue University

Qifan Song
Purdue University

Guang Cheng
Purdue University

Abstract

Modern machine learning and deep learning models are shown to be vulnerable when testing data are slightly perturbed. Existing theoretical studies of adversarial training algorithms mostly focus on either adversarial training losses or local convergence properties. In contrast, this paper studies the generalization performance of a generic adversarial training algorithm. Specifically, we consider linear regression models and two-layer neural networks (with lazy training) using squared loss under low-dimensional and high-dimensional regimes. In the former regime, after overcoming the non-smoothness of adversarial training, the adversarial risk of the trained models can converge to the minimal adversarial risk. In the latter regime, we discover that data interpolation prevents the adversarially robust estimator from being consistent. Therefore, inspired by successes of the least absolute shrinkage and selection operator (LASSO), we incorporate the \mathcal{L}_1 penalty in the high dimensional adversarial learning and show that it leads to consistent adversarially robust estimation. A series of numerical studies are conducted to demonstrate how the smoothness and \mathcal{L}_1 penalization help improve the adversarial robustness of DNN models.

1 INTRODUCTION

Recent advances in deep learning and machine learning have led to breakthrough performance and are widely applied in practice. However, empirical experiments show that deep learning models can be fragile and vulnerable against adversarial input which is intentionally

perturbed (Biggio et al., 2013; Szegedy et al., 2014). For instance, in image recognition, a deep neural network will predict a wrong label when the testing image is slightly altered, while the change is not recognizable by the human eye (Papernot et al., 2016b). To ensure the reliability of machine learning and deep learning when facing real-world inputs, the demand for robustness is increasing. The related research efforts in adversarial learning include designing adversarial attacks in various applications (Papernot et al., 2016b,a; 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., 2018; Jalal et al., 2017; Balunovic and Vechev, 2020).

To introduce adversarial training, let l denote the loss function and $f_\theta(x) : \mathbb{R}^d \rightarrow \mathbb{R}$ be the model with parameter θ . The (population) adversarial loss is defined as $R_f(\theta, \epsilon) := \mathbb{E}[l(f_\theta[x + A_\epsilon(f_\theta, x, y)], y)]$, where A_ϵ is an attack of strength $\epsilon > 0$ and intends to deteriorate the loss in the following way

$$A_\epsilon(f_\theta, x, y) := \operatorname{argmax}_{z \in \mathcal{R}(0, \epsilon)} \{l(f_\theta(x + z), y)\}. \quad (1)$$

In the above, z is subject to the constraint $\mathcal{R}(0, \epsilon)$, i.e. an \mathcal{L}_2 ball centered at 0 with radius ϵ .

Given i.i.d. training samples $\{(x_i, y_i)\}_{i=1}^n$, the adversarial training aims to minimize an empirical version of $R_f(\theta, \epsilon)$ w.r.t. θ :

$$\widehat{R}_f(\theta, \epsilon) = \frac{1}{n} \sum_{i=1}^n l(f_\theta[x_i + A_\epsilon(f_\theta, x_i, y_i)], y_i), \quad (2)$$

and $\widehat{\theta} = \operatorname{argmin}_\theta \widehat{R}_f(\theta, \epsilon)$. The minimization in (2) is often implemented through an iterative two-step (min-max) update. In the t -th iteration, we first 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; see Algorithm 1. This generic algorithm and its variants have been studied in Shaham et al., 2015; Madry et al., 2018; Jalal et al., 2017; Balunovic and

Vechev, 2020; Sinha et al., 2018; Wang et al., 2019a among others. Note that for complex loss function l or model f_θ , there may not be an analytic form for A_ϵ (e.g. deep neural networks). In this case, an additional iterative optimization is needed to approximate A_ϵ at each step; see Wang et al., 2019a.

Algorithm 1 A General Form of Adversarial Training

Input: data $(x_1, y_1), \dots, (x_n, y_n)$, attack strength ϵ , number of steps T , initialization $\theta^{(0)}$, step size η .

for $t = 1$ **to** T **do**

for $i = 1$ **to** n **do**

 Calculate the attack for the i th sample and get

$$\tilde{x}_i^{(t-1)} = x_i + A_\epsilon(f_{\theta^{(t-1)}}, x_i, y_i).$$

end for

 Fixing $\tilde{x}_i^{(t-1)}$'s, update $\theta^{(t)}$ from $\theta^{(t-1)}$ through

$$\begin{aligned} \theta^{(t)} &= \theta^{(t-1)} - \eta \nabla_{\theta} \hat{R}_f(\theta^{(t-1)}, \epsilon) \\ &= \theta^{(t-1)} - \eta \nabla_{\theta} \left[\frac{1}{n} \sum_{i=1}^n l(f_{\theta^{(t-1)}}(\tilde{x}_i^{(t-1)}), y_i) \right]. \end{aligned} \quad (3)$$

end for

Output: $\theta^{(T)}$.

In the literature, there are three major strands of theoretical studies related to this work. The first strand focuses on the statistical properties or generalization performance of adversarially robust estimators without taking account of the role of optimization algorithms (Javanmard et al., 2020; Yin et al., 2019; Raghunathan et al., 2019; Schmidt et al., 2018; Najafi et al., 2019; Min et al., 2020; Zhai et al., 2019; Hendrycks et al., 2019; Chen et al., 2020). For instance, Javanmard et al. (2020) studied the statistical properties of $R_f(\hat{\theta}, \epsilon)$, without specifying how to obtain the exact/approximate global minimizer $\hat{f}_{\hat{\theta}}$. The second strand studies the adversarial training loss, i.e., the limiting behaviors of $\hat{R}_f(\theta^{(t)}, \epsilon)$ as a training algorithm iteration t grows. For instance, Gao et al. (2019); Zhang et al. (2020) showed that for overparameterized neural networks, the empirical adversarial loss could be arbitrarily close to the minimum value in a local region near initialization. The third strand studies the (local) convergence of adversarial training under certain convexity assumptions (e.g., Sinha et al., 2018; Wang et al., 2019a). Besides these theoretical results, there are a few empirical works as well (e.g., Wong et al., 2020; Wang et al., 2019b; Rice et al., 2020; Lee and Chandrakasan, 2020; Wu et al., 2020; Xie et al., 2020).

In this paper, we investigate the global convergence and the generalization ability (i.e., $R_f(\theta^{(T)}, \epsilon)$) of the adversarial training, for two models f_θ , linear regression and two-layer neural networks, with the squared

loss $l(f_\theta(x), y) = (f_\theta(x) - y)^2$ under \mathcal{L}_2 and \mathcal{L}_∞ attacks. Our theoretical contributions are summarized as follows.

First, the adversarial loss $R_f(\theta, \epsilon)$ suffers from non-differentiability even if both l and f_θ are smooth, thus the optimization, i.e., Algorithm 1, works poorly. This motivates us to introduce a surrogate attack to overcome the non-smoothness problem. Under low dimensional setup and \mathcal{L}_2 attack, we show that under proper conditions, the iterative estimate $\theta_\xi^{(T)}$ trained from the surrogate adversarial loss asymptotically achieves the minimum adversarial risk (Section 2).

Secondly, we observe the ‘‘data interpolation’’ behavior of \mathcal{L}_2 adversarial training under high dimensional setup ($d/n \rightarrow \infty$). More specifically, the training loss converges to zero, but the population loss $R_f(\theta_\xi^{(T)}, \epsilon)$ converges to a large constant which is the adversarial loss of null model. To remedy the poor generalization, we penalize the adversarial training loss using LASSO. The resulting adversarially robust estimator and adversarial risk are both consistent for under some sparsity assumption (Section 3).

Thirdly, we examine the differences between the \mathcal{L}_∞ and \mathcal{L}_2 adversarial training. One similarity with \mathcal{L}_2 case is that, when attack strength is small, data interpolation prevents \mathcal{L}_∞ adversarial training from achieving a consistent estimator as well, and the issue can be improved by the use of LASSO penalty. In terms of the differences, in general, it is harder to conduct adversarial training under \mathcal{L}_∞ attack, in the sense that a lower learning rate and more iterations are required to ensure the convergence of $\theta_\xi^{(T)}$ (Section 4).

It is worth mentioning that a recent work by Allen-Zhu and Li (2020) also conducted a similar analysis for the generalization ability of adversarial training for a two-layer ReLU network model. However, the focus of their work is to explain the feature purification effect of adversarial training in neural network and overlook the difficulties of adversarial training when ϵ does not asymptotic goes to 0.

For technical simplicity, throughout the paper, we assume the data are generated from linear regression:

$$y = \theta_0^\top x + \varepsilon, \quad (4)$$

where $x \in \mathbb{R}^d$ is Gaussian vector with mean 0 and variance Σ (θ_0 is not perpendicular to Σ), and ε is a Gaussian noise (independent of x) with variance $\sigma^2 < \infty$. As d diverges, we assume both maximum and minimum eigenvalues of Σ are finite and bounded away from 0. In addition, $\|\theta_0\|$ and σ^2 are allowed to increase in d , but the signal-to-noise ratio $\|\theta_0\|_\Sigma / \sigma$ is large, say bounded away from zero.

2 LOW DIMENSIONAL ASYMPTOTICS

This section considers linear regression models and two-layer neural network models (training the first layer weights) under the low dimensional scenario. In particular, we examine the landscape of the adversarial loss and further investigate the testing performance of the estimator $\theta_\xi^{(T)}$. In what follows, we rewrite R_f as R_L for linear models and as R_N for two-layer networks.

2.1 Linear regression model

Consider the linear regression model: $f_\theta(x) = \theta^\top x$. By the definition of $A_\epsilon(f_\theta, x, y)$ under the \mathcal{L}_2 ball constraint and the fact that x and ϵ are both Gaussian, R_L has an analytical form as

$$R_L(\theta, \epsilon) = \|\theta - \theta_0\|_\Sigma^2 + \sigma^2 + \epsilon^2 \|\theta\|^2 + 2\epsilon c_0 \|\theta\| \sqrt{\|\theta - \theta_0\|_\Sigma^2 + \sigma^2}, \quad (5)$$

where $\|a\|_\Sigma^2 := a^\top \Sigma a$, $\|a\| := \|a\|_2$ for any vector a , and $c_0 := \sqrt{2/\pi}$. Given any $\epsilon \geq 0$, define

$$\theta^*(\epsilon) := \underset{\theta}{\operatorname{argmin}} R_L(\theta, \epsilon) \text{ and } R^*(\epsilon) := \min_{\theta} R_L(\theta, \epsilon).$$

When no confusion arises, we will rewrite $\theta^*(\epsilon)$ and $R^*(\epsilon)$ as θ^* and R^* for simplicity.

We first analyze $R_L(\theta, \epsilon)$ based on the form (5).

Proposition 1. *Assume the attack strength $\epsilon > 0$ and the data generation follows (4). The adversarial loss R_L is differentiable w.r.t. θ if and only if (1) $\sigma^2 > 0$ and $\theta \neq 0$, or (2) $\sigma = 0$, $\theta \neq 0$, and $\theta \neq \theta_0$. The function R_L is convex w.r.t. θ , and there exists some constant c such that when $\epsilon \geq c$, $\theta^* = 0$.*

From Proposition 1, there always exists some θ where $R_L(\theta, \epsilon)$ is not differentiable (e.g., $\theta = 0$), and it is not avoidable when ϵ is large. Even if we smooth the standard loss as in Xie et al. (2020) to improve the quality of the gradients (or use smooth classifiers as in Salman et al., 2019), the non-differentiable issues remains for the adversarial loss. This makes it difficult to track the trajectory of gradient descent algorithms in the sense that the adversarial training loss will fluctuate during training rather than strictly decreases over iterations.

To solve this problem and improve the gradient quality, for both models under consideration, we introduce $A_{\epsilon, \xi}$ that is a surrogate for \mathcal{L}_2 attack A_ϵ :

$$A_{\epsilon, \xi}(f_\theta, x, y) = \frac{\|\partial l(f_\theta(x), y)/\partial x\|}{\sqrt{\|\partial l(f_\theta(x), y)/\partial x\|^2 + \xi^2}} A_\epsilon(f_\theta, x, y).$$

Lee and Chandrakasan (2020) showed that the instability of adversarial training when $\partial l(f_\theta(x), y)/\partial x$ is closed to zero. Therefore, the design of $A_{\epsilon, \xi}$ aims to impose shrinkage effect on A_ϵ when $\|\partial l(f_\theta(x), y)/\partial x\|$ is small, thus the training process will be more stable.

When $\xi = 0$, $A_{\epsilon, \xi}$ is reduced to A_ϵ . With the surrogate $A_{\epsilon, \xi}$, we define the empirical and population surrogate loss as:

$$\begin{aligned} \hat{R}_{L, \xi}(\theta, \epsilon) &:= \frac{1}{n} \sum_{i=1}^n l(f_\theta(x_i + A_{\epsilon, \xi}(f_\theta, x_i, y_i)), y_i), \\ R_{L, \xi}(\theta, \epsilon) &:= \mathbb{E} l(f_\theta(x_i + A_{\epsilon, \xi}(\theta, x_i, y_i)), y_i). \end{aligned}$$

One can show, $R_{L, \xi}(\theta, \epsilon)$ is smooth and convex everywhere. Accordingly, Algorithm 1 is modified by replacing A_ϵ with $A_{\epsilon, \xi}$. Note that a smaller value of ξ , which leads to a closer approximation to $R_{L, 0}$, makes $R_{L, \xi}$ less smooth at the origin and thus requires a lower learning rate η and more iteration steps T . Therefore, we require a fine-tuning of ξ in the sense that $\xi/\|\theta_0\|$ slowly decreases to 0. The surrogate loss for two-layer neural networks, i.e., $R_{N, \xi}$ and $\hat{R}_{N, \xi}$, is defined in the same way.

The use of surrogate attack improve the gradient quality of adversarial training and enforces the smoothness and convexity of surrogate loss $R_{L, \xi}(\theta, \epsilon)$. It facilitates the convergence and generalization analysis of adversarial training. We study the consistency of $\theta_\xi^{(T)}$ toward θ^* and $R_{L, \xi}(\theta_\xi^{(T)}, \epsilon)$ toward R^* , under suitable choices of (ξ, η, t) after accounting for the dimensional effect $v^2 := \|\theta_0\|_\Sigma^2 + \sigma^2$ (for simplicity assume v^2 bounded away from zero); see Theorem 2.

Theorem 2. *Assume the data generation follows (4) and let ϵ be a fixed constant. If there exists some constant B_0 such that $\|\theta_\xi^{(0)}\| \leq B_0 v$, and the dimension growth rate satisfies $\log n \sqrt{d^2/n} \rightarrow 0$, then with probability tending to 1, the surrogate loss $R_{L, \xi}(\theta_\xi^{(t)}, \epsilon)$ decreases in each iteration, and*

$$\frac{R_{L, \xi}(\theta_\xi^{(T)}, \epsilon) - R^*}{v^2} \rightarrow 0, \text{ and } \frac{\|\theta_\xi^{(T)} - \theta^*\|}{v} \rightarrow 0,$$

given $\eta = \xi/(v^2 L)$ for some large constant L , $T = (v^2 \log \log n)/\xi$ and $\xi = v^2 d/\sqrt{n} \log n$.

The proof of Theorem 2 is postponed to Appendix C. The results of Theorem 2 actually hold for a wide range of (ξ, η, T) (as elaborated in the Appendix C). In general, the choice of (η, T) can be invariant to v , while a larger v implies a wider range of possible ξ . In terms of tuning $\theta_\xi^{(0)}$ and ξ , one can estimate v^2 via the sample variance of y_1, \dots, y_n . Note that the above discussions apply to Theorems 3, 4, 5, 6, and 7 as well.




Figure 1: Adversarial training and testing loss in linear regression. Solid line: adversarial testing loss. Dashed line: adversarial training loss. The training process with $\xi = 0$ is more fluctuate.

Theorem 2 also confirms the phenomenon that “adversarial training hurts standard estimation” (Raghu-nathan et al., 2019). Denote $\hat{\theta}_{\text{OLS}}$ as the common least square estimator for un-corrupted data (i.e. without attack) and trivially $\hat{\theta}_{\text{OLS}} \rightarrow \theta_0$. By Theorem 2, $\theta_{\xi}^{(T)} \rightarrow \theta^* \neq \theta_0$, thus

$$\frac{R_L(\theta_{\xi}^{(T)}, 0) - R_L(\hat{\theta}_{\text{OLS}}, 0)}{v^2} \rightarrow \frac{R_L(\theta^*, 0) - R_L(\theta_0, 0)}{v^2} = c(\epsilon) > 0,$$

where the function $c(\epsilon)$ increases in ϵ , and converges to $\|\theta_0\|_{\Sigma}^2/v^2$ as ϵ diverges.

Remark 1. The Gaussian assumption of x is used in Theorem 2 for the convenience of showing the analytical form of R_L . All the results in this paper, except for Section 3.2, are still valid as long as the density of (x, y) is finite and $\exp(-t\|x\|^2) < \infty$ for some $t > 0$.

Numerical experiments

To demonstrate the necessity of smoothing attacks under large ϵ , a simple simulation is conducted. Let $d = 10$, $n = 1000$, $x \in \mathcal{R}^d$ and $\Sigma = I_d$. The response $y = \beta_0^\top x + \varepsilon$ with $\beta_0 = (1, \dots, 1)^\top$ and $\sigma^2 = 1$. The level of attack is taken as $\epsilon = 3$. We use the same training data set in the adversarial training for $\xi = 0$ and $\xi = 0.05$. Figure 1 displays the effectiveness of ξ : the loss with $\xi = 0$ always fluctuates, and after we smooth the training process through taking $\xi = 0.05$, (surrogate) training and (non-surrogate) testing loss smoothly decrease in each iteration.

2.2 Two-layer neural networks

We consider the two-layer neural network with an activation function ϕ , say

$$f_{\theta}(x) = \frac{1}{\sqrt{h}} \sum_{j=1}^h \phi(x^\top \theta_j) a_j, \quad (6)$$

where a_j ’s are known values and $\theta = (\theta_1, \dots, \theta_h) \in \mathbb{R}^{d \times h}$ is the parameter to be trained. This “lazy training” setup has been widely used in the literature (e.g. Du et al., 2018, 2019; Arora et al., 2019; Ba et al., 2020; Allen-Zhu and Li, 2020), it eases the theoretical analysis. Additional, we adopt a vanishing initialization scheme (Ba et al., 2020): $\theta_{\xi,j}^{(0)} \sim N(0, I_d/dh^{1+\delta})$ for some $\delta > 0$.

Theorem 3 shows that $R_{N,\xi}(\theta_{\xi}^{(T)}, \epsilon)$ converges to the same minimal loss R^* as in linear regression under proper choice of T and ξ , as n and h diverge.

Theorem 3. Under the generative model (4), assume the activation function ϕ in model (6) is twice continuously differentiable, $\phi'(0) \neq 0$, and $\phi(0) = 0$. If a and h satisfy $\|a\|_{\infty} = O(1)$, $\max |a_j|/(\min |a_j|) = \Theta(1)$, $(d \log n)\|a\|_{\infty} v/\sqrt{h} \rightarrow 0$, and $(d \log n)\|a\|_{\infty} v\sqrt{h}/\|a\|^2 \rightarrow 0$. When $\xi/v^2 \rightarrow 0$ where $\xi/v^2 = -\log \log n / \log(\sqrt{d^2 \log n/n} \vee (d \log n)\|a\|_{\infty}/\sqrt{h})$, and $\eta = \xi h/(v^2 L\|a\|^2)$ for some large constant L , with probability tending to 1, for $T = (v^2 \log \log n)/\xi$, if $\sqrt{d \log n}(1 + v^2 \eta\|a\|^3/(h^{3/2}\xi) + L\eta v)^T/h^{\delta/2} \rightarrow 0$, then

$$\frac{R_{N,\xi}(\theta_{\xi}^{(T)}, \epsilon) - R^*}{v^2} \rightarrow 0, \quad (7)$$

where R^* is the exactly the same as Theorem 2.

The detailed proof is postponed to Appendix D. The choice of ξ here depends on the weights a together with the number of hidden nodes h . Note that although both Theorems 2 and 3 establish the convergence of $R_{N,\xi}(\theta, \epsilon)$, Theorems 3 requires that ξ/v converges to zero in a slower speed, leading to a slower convergence rate for $R_{N,\xi}(\theta, \epsilon)$.

Remark 2. The proof of Theorem 3 is similar to Ba et al. (2020): as the number of hidden nodes h grows, the trajectories of optimization using linear network (with zero initialization) and nonlinear network (with vanishing initialization) are slightly different, while the convergence result of the former one can be simply extended from linear models. Different from Ba et al. (2020), we specify the learning rate as well as the number of iterations as functions of (d, a, n, h) , while Ba et al. (2020) utilized gradient flow, which is not applicable in our setup. In addition, compared with Ba et al. (2020), the relationship of (ξ, a, η) is revealed in our result when $\|\theta_0\| \rightarrow \infty$.

Theorem 3 requires a continuous differentiable ϕ , and similar results can be established for ReLU activation function as well:

Theorem 4. Under the generative model (4), assume the activation function ϕ is ReLU function with zero initialization and no bias. Take $\xi/v^2 =$

$-\log \log n / \log(\sqrt{(d^2 \log n)/n})$. Set $\eta = \xi h / (v^2 L \|a\|^2)$ for some constant L . Denote a^+ as a vector such that $a_j^+ = a_j 1\{a_j > 0\}$, and similarly define a^- . If a satisfy $\|a^+\|/\|a^-\| = 1$ and $\|a\|_\infty = O(1)$, with probability tending to 1, for $T = (v^2 \log \log n)/\xi$, (7) holds.

Numerical experiments

A series of simulation studies of low dimensional linear regression and two-layer neural network model with lazy training are conducted. Due to page limit, these results are detailed in Appendix A and successfully validate our Theorems 2-4 on the convergence of adversarial training with surrogate loss.

Here we present another experiment that shows the improvement of predicting the performance of adversarial trained estimator via surrogate loss for complicated models beyond our theorems. We fit a ResNet-34 (WideResNet34-1) model for CIFAR-10 dataset. Since the ReLU function is not smooth, we utilize the training technique introduced by Xie et al. (2020): we use ReLU in the forward path and use Softplus($\beta = 10$) in the backward path to improve the gradient quality. The number of epochs is taken as 100. The initial learning rate is 0.1, and at the 75th and 90th epoch, it is multiplied by 0.1. The value of ξ is taken as 0.001 at the initial stage and multiplies 0.1 whenever the learning rate is changed. We repeat this experiment 10 times to obtain the mean and variance and conduct this experiment under various levels of attacks. We use \mathcal{L}_2 attack with strength $\epsilon = 3.0$ in this experiment. The results are summarized in Figure 2. It shows that using surrogate loss leads to slightly higher adversarial testing accuracy and much higher standard testing accuracy than the one with $\xi = 0$.

In addition, we conduct two experiments (in case that the adversarial training with $\xi = 0$ does not converges algorithmically in the above experiment): (1) with 200 epochs and (2) with initialization that is obtained by standard training as in Allen-Zhu and Li (2020). The results are postponed to the appendix. In short, a better performance is obtained under surrogate loss, and the initialization from standard training does not improve the performance.

Although our theory only reveals a single non-differentiable point, it is still important to handle this carefully in neural networks. The non-differentiable problem is partially due to that the adversary has no preference in the direction of attack. If we estimate the attack twice (from different initializations) and the difference between the two estimates will be large when the non-differentiable problem is severe. We conduct a small experiment to investigate the attack difference. The details are postponed to the appendix. The results justify the importance of accommodating this




Figure 2: Standard and adversarial testing accuracies for CIFAR-10 with ResNet34. Solid line: adversarial testing accuracy. Dashed line: standard testing accuracy. With $\xi = 0.001$, both adversarial testing accuracy and standard testing accuracy (with standard deviation as 0.458, 0.321) are higher than that for $\xi = 0$ (0.743, 0.219).

non-differentiability issue, especially for large ϵ .

3 HIGH DIMENSIONAL ASYMPTOTICS

In this section, we focus on the high dimensional regime where $d/n \rightarrow \infty$. It is first revealed that the adversarial training also suffers from the classical interpolation effect, i.e., near-zero (surrogate) adversarial training loss but high generalization error. As a potential remedy, we penalize the adversarial training loss using LASSO and show that the estimate is consistent when θ^* is sparse.

3.1 Effect of interpolation

It is well known that interpolation may occur under high dimensionality. For instance of linear regression, if a gradient descent with zero initialization is applied to minimize the squared loss when $d/n \rightarrow \infty$, then the solution converges to

$$\theta(\mathbf{y}) := \mathbf{X}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y},$$

where $\mathbf{X} = (x_1, x_2, \dots, x_n)^\top$ and $\mathbf{y} = (y_1, y_2, \dots, y_n)^\top$, given a sufficiently small learning rate. This perfectly interpolated estimator $\theta(\mathbf{y})$ is proven to be inconsistent to θ_0 and lead to a large generalization error (e.g., Hastie et al., 2019; Belkin et al., 2019). Note that this over-fitting scenario is different from the one in Rice et al. (2020), which is caused by over-parameterization in deep neural networks rather than high dimensionality of the input.

Our first result shows that $\theta(\mathbf{y})$ also induces the same effect of interpolation in adversarial learning for linear models.

Lemma 1. Assume data generation follows (4). When $\theta(\mathbf{y}) \neq 0$ and $d/n \rightarrow \infty$, we have $\|\theta(\mathbf{y})\|^2/v^2 = O_p(n/d)$, and with probability tending to 1, for any $\xi \geq 0$, it holds that

$$\frac{\hat{R}_{L,\xi}(\theta(\mathbf{y}), \epsilon)}{v^2} \rightarrow 0 \text{ and } \frac{R_{L,\xi}(\theta(\mathbf{y}), \epsilon)}{v^2} \rightarrow 1.$$

We next show that $\theta_\xi^{(T)}$ shares the same properties as $\theta(\mathbf{y})$. The core idea is that the training trajectory $\{\theta_\xi^{(t)}\}_{t=1}^T$ can be sufficiently close to that in the standard training, when both are initialized from zero. Since the latter converges to $\theta(\mathbf{y})$, the surrogate adversarial training loss and testing loss of $\theta_\xi^{(T)}$ act in a similar way as those for $\theta(\mathbf{y})$ respectively.

Theorem 5. Under the same assumptions as in Lemma 1, when $(\log n)\sqrt{n/d} \rightarrow 0$, take η small enough such that the largest eigenvalue of $\eta \mathbf{X}^\top \mathbf{X}$ is smaller than 1. Use zero initialization, and denote $T := \min\{t \in \mathbb{Z}^+ : \|\mathbf{X}\theta_\xi^{(t)} - \mathbf{y}\|_2/(v\sqrt{n}) < 1/\sqrt{\log n}\}$, then with probability tending to 1, for any $\xi > 0$, we have $T < \infty$ and

$$\frac{\hat{R}_{L,\xi}(\theta_\xi^{(T)}, \epsilon)}{v^2} \rightarrow 0, \text{ and } \frac{R_{L,\xi}(\theta_\xi^{(T)}, \epsilon)}{v^2} \rightarrow 1.$$

The proof of Theorem 5 is postponed to Appendix E. Compared with Theorem 2, Theorem 5 no longer requires ξ to be associated with (d, n) . A crucial reason for this difference is that under high dimensionality, when $t \leq T$, the smoothness of $\hat{R}_{L,\xi}$ along the training trajectory (i.e., the gradient of $\hat{R}_{L,\xi}(\theta_\xi^{(t)}, \epsilon)$) is always dominated by a term that is only determined by the eigenvalues of high dimensional design matrix \mathbf{X} , regardless of how small ξ is (refer to equation (13) in Appendix E for details). This is contrast to the low dimensional case.

Theorem 5 shows that $R_{L,\xi}(\theta_\xi^{(T)}, \epsilon)/v^2$ does not converge to R^*/v^2 . Similar results can be established for two-layer neural networks (with lazy training):

Theorem 6. For the two-layer neural network (6), under the same conditions on ϕ as in Theorem 3, $(\log n)\sqrt{d/n} \rightarrow \infty$, take zero/vanishing initialization and $\eta = \eta_{\text{linear}} h/\|a\|^2$. Assume $\|a\|_\infty = O(1)$, $(d \log n)\|a\|_\infty v/\sqrt{h} \rightarrow 0$, $\max |a_j|/(\min |a_j|) = \Theta(1)$, $\sqrt{dn}(\log n)^2\|a\|_\infty v\sqrt{h}/\|a\|^2 \rightarrow 0$, $\sqrt{d \log n}(1 + v^2\eta\|a\|^3/(h^{3/2}\xi) + L\eta v)^T/h^{\delta/2} \rightarrow 0$. Denote $T := \min\{t \in \mathbb{Z}^+ : \|\mathbf{X}\theta_\xi^{(t)} - \mathbf{y}\|_2/(v\sqrt{n}) < 1/\sqrt{\log n}\}$, then with probability tending to 1, for any $\xi > 0$, we have $T < \infty$ and

$$\frac{\hat{R}_{N,\xi}(\theta_\xi^{(T)}, \epsilon)}{v^2} \rightarrow 0, \text{ and } \frac{R_{N,\xi}(\theta_\xi^{(T)}, \epsilon)}{v^2} \rightarrow 1. \quad (8)$$

For ReLU network, we have the following result:

Theorem 7. For the two-layer neural network (6), $(\log n)\sqrt{d/n} \rightarrow \infty$, assume the activation function ϕ is ReLU function with zero initialization and no bias. If $\|a^+\|/\|a^-\| = 1$, $\|a\|_\infty = O(1)$. Denote $T := \min\{t \in \mathbb{Z}^+ : \|\mathbf{X}\theta_\xi^{(t)} - \mathbf{y}\|_2/(v\sqrt{n}) < 1/\sqrt{\log n}\}$, then for any $\xi > 0$, with probability tending to 1, we have $T < \infty$, and (8) holds.

Numerical experiment

A simulation is conducted to verify Theorem 5. We choose $n = 20$, $d = 1000$, and $\sigma^2 = 1$. The true underlying model θ_0 is all-zero except for its first 10 elements being 1. The attack intensity is $\epsilon = 0, 0.01, 0.1$. Learning rate is taken as 0.001 with zero initialization and $\xi = 0.5$. The curves in Figure 3 represent means of respective statistics, and the shaded areas represent mean \pm one standard deviation, based on 100 replications. Figure 3 shows that the surrogate adversarial training loss keeps decreasing to around zero for all the choices of ϵ , while the adversarial testing loss converges to some nonzero constant. Note that the three adversarial training loss curves in the left plot of Figure 3 overlap.

More experiments for larger d and $\xi = 0$ (with a change when $\theta = 0$) are postponed to Appendix A. Besides, we also postpone experiments for neural networks to Appendix A.




Figure 3: Adversarial training, high-dimension.

3.2 Improving adversarial robustness using LASSO under high dimensionality

In this section, we explore how incorporating LASSO improves adversarial learning under high dimensionality. In particular, we will present some theoretical justifications for linear models and conduct numerical exploration to evaluate the potential of LASSO in neural networks.

Intuitively, a sparse adversarial learning via LASSO makes sense only when the adversarial loss has a sparse global optimization, i.e., θ^* is a sparse vector. Therefore, certain investigation is necessary to understand the sparsity relation between θ_0 and θ^* .

Proposition 8. Under model (4), the optimal solution of $R_{L,\xi}(\theta, \epsilon)$ is of the form $\theta^* = (\Sigma + \kappa I)^{-1} \Sigma \theta_0$ for some κ as a function of $(\theta_0, \Sigma, \sigma^2, \xi)$. Assuming θ_0 is sparse, then whether the robust coefficient $((\Sigma + \kappa I)^{-1} \Sigma \theta_0)$ is sparse or not depends on Σ .

The following example illustrates that, based on Proposition 8, when the correlation between active set and inactive set is zero, the adversarially robust model θ^* is sparse as well.

Example 1. When θ_0 is sparse, and Σ can be represented as $\begin{Bmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{Bmatrix}$, where Σ_1 is the covariance of the active attributes, and Σ_2 is for the other attributes of x , the model θ^* will be sparse as well.

To simplify the derivation, we assume $\Sigma = I$ in the following result. Denote S as the active set of θ^* , and $s = |S|$ as the size of S . We consider applying LASSO in the adversarial training loss:

$$\frac{1}{n} \sum_{i=1}^n l(f_\theta(x_i + A_{\epsilon,\xi}(f_\theta, x_i, y_i)), y_i) + \lambda \|\theta\|_1.$$

The statistical property of $\hat{\theta}_\xi$, the minimizer of the above objective function, is as follows:

Theorem 9. Assume data generation follows (4), θ_0 is sparse and $\Sigma = I$. Take $\xi/v^2 \rightarrow 0$, $\lambda/v = o(1)$ and $\lambda/v \geq (c\sqrt{(s \log d)/n}) \vee (\xi a_n/v^2)$ for some large constant c and $a_n \rightarrow \infty$. If $\epsilon < \sqrt{\pi} \|\theta_0\|_2 / (\sqrt{2}v)$, then $\theta^* \neq 0$, and with probability tending to 1, we have

$$\frac{R_{L,\xi}(\hat{\theta}_\xi, \epsilon) - R^*}{v^2} \rightarrow 0, \text{ and } \frac{\|\hat{\theta}_\xi - \theta^*\|_1}{\lambda} < \infty.$$

The proof of Theorem 9 is similar to the traditional LASSO analysis as in Bickel et al. (2009); Belloni and Chernozhukov (2013) but with an important modification. In the literature, the Hessian of the standard training loss, i.e., $\mathbf{X}^\top \mathbf{X}/n$, is usually required to satisfy the so-called restricted eigenvalue condition. However, in adversarial setting, the Hessian changes as θ , so it takes more steps to verify the above condition.

Remark 3. Theorem 9 shows the effectiveness of LASSO in sparse linear model and it performs better than the case without LASSO. Note that the estimation consistency still holds for low-dimensional dense model, if d and n satisfies $(d \log d)/n \rightarrow 0$. But, to ensure that LASSO improves the performance in this case, λ should be carefully tuned.

We conduct some empirical study to explore the potential applications of LASSO in the adversarial training of neural networks. Similar experiments under large-sample regime can be found for adversarial training

Sinha et al. (2018); Wang et al. (2019a); Raghunathan et al. (2019), and pruning in adversarial training Ye et al. (2019); Li et al. (2020).

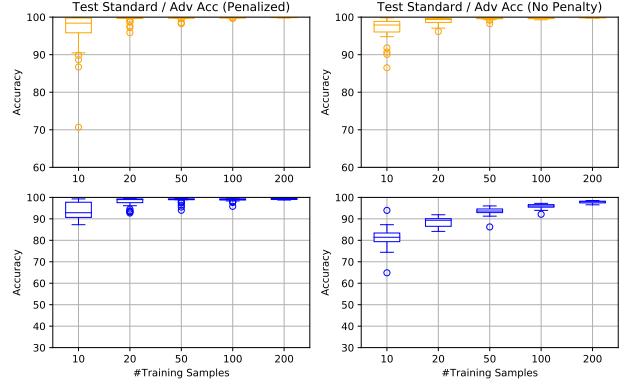


Figure 4: Comparison on standard (upper) / adversarial (lower) test accuracies between training with/without \mathcal{L}_1 penalty under \mathcal{L}_2 attack with $\xi = 10^{-4}$. Attack strength $\epsilon = 3$.

Numerical experiments

The program was modified from a repository in Github¹ and library `Advertorch`. A simple two-layer neural network is constructed with 1024 hidden nodes and ReLU as activation. We use MNIST dataset to distinguish between digits 0 and 1, and randomly select a small number of samples of 0 and 1 from the training dataset to create a high-dimension scenario. The \mathcal{L}_2 attack level is set to be 3. We trained 2000 epochs to ensure the convergence of the algorithms and repeat the experiment for 30 times to draw a box-plot. After training 2000 epochs, for both $\lambda = 0$ (No penalty) and $\lambda = 0.001$ (LASSO), the training accuracies for clean data and adversarial data both reach 100%. The penalty $\lambda = 0.001$ was chosen such that the magnitude of penalty is comparable with loss. The results are summarized in Figure 4.

For adversarial accuracy, as shown in Figure 4, the results with two different λ 's are significantly different, where the choice of $\lambda = 0.001$ improves the adversarial accuracy compared with $\lambda = 0$. As a reference, we also plot the standard accuracy (i.e. prediction accuracy for un-corrupted data), even though the objective function minimized is the (penalized) adversarial training loss. Figure 4 shows, it approaches 99% quickly in n for both adversarial training with and without LASSO.

We also observe similar results when $\xi = 0$, and the details are postponed to Appendix A.

We also tried on CIFAR-10 with WideResNet34-10.

¹<https://github.com/louis2889184/pytorch-adversarial-training>

The λ is chosen to ensure that the magnitude of cross entropy loss and the LASSO penalty are comparable. We use all data in the training dataset to conduct this experiment. The results are summarized in Table 1. Our theorem only concerns the high dimensional case, i.e., small- n -large- d , however, as showed in Table 1, both standard and adversarial testing accuracies are still enhanced when using LASSO, in this large- n application (Refer to Remark 3).

Method	std acc (%)	adv acc (%)
Benchmark	84.346(0.355)	61.760(0.204)
LASSO 10^{-5}	86.568(0.214)	63.072(0.292)

Table 1: Adversarial training in CIFAR-10 using WideResNet34-10 with \mathcal{L}_2 attack, $\epsilon = 0.5$.

4 \mathcal{L}_∞ ADVERSARIAL TRAINING

In this section, we discuss the adversarial loss and adversarial training under \mathcal{L}_∞ attack. Similar as stated in Chen et al. (2020), the adversarial risk of the linear model becomes

$$R_L^\infty(\theta, \epsilon) = \|\theta - \theta_0\|_\Sigma^2 + \sigma^2 + \epsilon^2 \|\theta\|_1^2 + 2\epsilon c_0 \|\theta\|_1 \sqrt{\|\theta - \theta_0\|_\Sigma^2 + \sigma^2}. \quad (9)$$

Below are discussions w.r.t \mathcal{L}_∞ adversarial training:

Harder to train

From (9), R_L^∞ is not differentiable when some element in θ is zero. Similar as for \mathcal{L}_2 attack, we propose to shrink the size of adversarial attack when R_L is not differentiable, while a difference is that the shrinkage is applied on each dimension of x : for $i = 1, \dots, d$,

$$[A_{\epsilon, \xi}^\infty(f_\theta, x, y)]_i = \frac{|\partial l / \partial x|_i}{|\partial l / \partial x|_i + \xi} [A_\epsilon^\infty(f_\theta, x, y)]_i.$$

A major difference between \mathcal{L}_∞ and \mathcal{L}_2 attacks is that, \mathcal{L}_∞ attack is more sensitive to ξ . For example, if $\theta = (1/d, \dots, 1/d)^\top$ and $\epsilon = 1/\sqrt{d}$, then A_ϵ^∞ becomes $(\epsilon, \dots, \epsilon)^\top$ whose \mathcal{L}_2 norm is 1, while $A_{\epsilon, \xi}^\infty$ is $(\epsilon/(1 + d\xi), \dots, \epsilon/(1 + d\xi))^\top$, whose \mathcal{L}_2 norm quickly shrinks to zero if $\xi d \rightarrow \infty$. As a result, it is necessary to require that $\xi = o(1/d)$ to avoid overshrinkage of the \mathcal{L}_∞ attack. However, as discussed in previous sections, a smaller ξ requires smaller learning rate and more training iterations, thus training under \mathcal{L}_∞ attack is more difficult.

Effect of interpolation

In high-dimensional case, adversarial training still suffers from data interpolation: when $\epsilon = O(1/\sqrt{d})$ and $d/n \rightarrow \infty$, the minimal adversarial training loss converges to zero, while the population adversarial loss

converges to v^2 (recall that $v^2 = \|\theta_0\|^2 + \sigma^2$). Similar as for \mathcal{L}_2 attack, we add LASSO in the \mathcal{L}_∞ adversarial training in MNIST and CIFAR-10. The results are summarized in Figure 15 and 16 in appendix, as well as Table 2. For both datasets, LASSO improves both standard and adversarial testing accuracies.

Method	std acc (%)	adv acc (%)
Benchmark	82.870(0.131)	50.338(0.315)
LASSO 10^{-5}	84.800(0.282)	54.260(0.376)

Table 2: Adversarial training in CIFAR-10 using WideResNet34-10 with \mathcal{L}_∞ attack, $\epsilon = 8/255$.

Remark 4. From the aspect of formulation, \mathcal{L}_∞ adversarial loss and LASSO has overlapped effect as both introduce \mathcal{L}_1 penalty effect into the loss function, see (9). However, LASSO and \mathcal{L}_∞ are designed for different purposes. From the aspect of loss landscape in deep learning, LASSO does not intend to change the loss landscape near the global minima as the penalty term goes to zero asymptotically, i.e., any global optimum for standard loss are optimum for LASSO problem given infinite training data. On the other hand, for adversarial robustness under \mathcal{L}_∞ attack, it aims to select the certain global minima such that the prediction is robust in the nearby region of training samples, and not all minimizers of standard loss are robust to adversarial attack. From this aspect, they can be applied simultaneously. We refer readers to Guo et al. (2020) for more discussion.

5 CONCLUSION AND FUTURE WORKS

This paper studies the convergence properties of adversarial training in linear models and two-layer neural networks (with lazy training). In the low-dimensional regime, using adversarial training with surrogate attack, the adversarial risk of the trained model converges to the minimal value. In a high-dimensional regime, data interpolation causes the adversarial training loss close enough to zero, while the generalization is poor. One potential solution is to add \mathcal{L}_1 penalty in the adversarial training, which results in both consistent adversarial estimate and risk in high dimensional sparse models.

There are several future directions. First, we may focus on classification tasks as a future work. In regression, the adversarially robust model generally outputs smaller-in-magnitude predictions, which is not practical in classification. One may be interested in how adversarial training works in classification. Second, the scenarios we consider are $d/n \rightarrow 0$ and ∞ , and one can consider the linear dimensionality case, i.e.

$d/n \rightarrow c$, as a future direction. Finally, the non-smoothness issue happens to the adversarial loss and the penalty term (e.g., LASSO, Wang et al., 2019b; Wu et al., 2020), so there is potential to improve further the gradient quality of penalized adversarial training via smoothing the penalty function.

Acknowledgements

Dr. Song’s research activities are partially supported by National Science Foundation DMS-1811812.

References

- Allen-Zhu, Z. and Li, Y. (2020), “Feature Purification: How Adversarial Training Performs Robust Deep Learning,” *arXiv preprint arXiv:2005.10190*.
- Arora, S., Du, S. S., Hu, W., Li, Z., and Wang, R. (2019), “Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks,” in *Proceedings of the 36th International Conference on Machine Learning*, PMLR, vol. 97 of *Proceedings of Machine Learning Research*, pp. 322–332.
- Ba, J., Erdogdu, M., Suzuki, T., Wu, D., and Zhang, T. (2020), “Generalization of two-layer neural networks: an asymptotic viewpoint,” in *8th International Conference on Learning Representations*.
- Bai, Z.-D. and Yin, Y.-Q. (2008), “Limit of the smallest eigenvalue of a large dimensional sample covariance matrix,” in *Advances In Statistics*, World Scientific, pp. 108–127.
- Balunovic, M. and Vechev, M. (2020), “Adversarial training and provable defenses: bridging the gap,” in *8th International Conference on Learning Representations*.
- Belkin, M., Hsu, D., and Xu, J. (2019), “Two models of double descent for weak features,” *arXiv preprint arXiv:1903.07571*.
- Belloni, A. and Chernozhukov, V. (2013), “Least squares after model selection in high-dimensional sparse models,” *Bernoulli*, 19, 521–547.
- Bickel, P. J., Ritov, Y., and Tsybakov, A. B. (2009), “Simultaneous analysis of Lasso and Dantzig selector,” *The Annals of Statistics*, 37, 1705–1732.
- 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.
- Chen, L., Min, Y., Zhang, M., and Karbasi, A. (2020), “More data can expand the generalization gap between adversarially robust and standard models,” *arXiv preprint arXiv:2002.04725*.
- Du, S. S., Lee, J. D., Li, H., Wang, L., and Zhai, X. (2019), “Gradient descent finds global minima of deep neural networks,” in *Proceedings of the 36th International Conference on Machine Learning*, PMLR, vol. 97 of *Proceedings of Machine Learning Research*, pp. 1675–1685.
- Du, S. S., Zhai, X., Poczos, B., and Singh, A. (2018), “Gradient descent provably optimizes over-parameterized neural networks,” *arXiv preprint arXiv:1810.02054*.
- Gao, R., Cai, T., Li, H., Hsieh, C., Wang, L., and Lee, J. D. (2019), “Convergence of adversarial training in overparametrized neural networks,” in *Advances in Neural Information Processing Systems*, pp. 13009–13020.
- Guo, Y., Chen, L., Chen, Y., and Zhang, C. (2020), “On connections between regularizations for improving dnn robustness,” *IEEE transactions on pattern analysis and machine intelligence*.
- Hastie, T., Montanari, A., Rosset, S., and Tibshirani, R. J. (2019), “Surprises in high-dimensional ridgeless least squares interpolation,” *arXiv preprint arXiv:1903.08560*.
- Hendrycks, D., Lee, K., and Mazeika, M. (2019), “Using pre-training can improve model robustness and uncertainty,” in *Proceedings of the 36th International Conference on Machine Learning*, vol. 97 of *Proceedings of Machine Learning Research*, pp. 2712–2721.
- Ing, C.-K. and Lai, T. L. (2011), “A stepwise regression method and consistent model selection for high-dimensional sparse linear models,” *Statistica Sinica*, 21, 1473–1513.
- Jalal, A., Ilyas, A., Daskalakis, C., and Dimakis, A. G. (2017), “The robust manifold defense: Adversarial training using generative models,” .
- Javanmard, A., Soltanolkotabi, M., and Hassani, H. (2020), “Precise tradeoffs in adversarial training for linear regression,” *arXiv preprint arXiv:2002.10477*.
- Lee, K. and Chandrakasan, A. P. (2020), “Rethinking Empirical Evaluation of Adversarial Robustness Using First-Order Attack Methods,” *arXiv preprint arXiv:2006.01304*.
- Li, B., Wang, S., Jia, Y., Lu, Y., Zhong, Z., Carin, L., and Jana, S. (2020), “Towards Practical Lottery Ticket Hypothesis for Adversarial Training,” *arXiv preprint arXiv:2003.05733*.

- 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*.
- Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. (2018), “Towards deep learning models resistant to adversarial attacks,” in *6th International Conference on Learning Representations*.
- Min, Y., Chen, L., and Karbasi, A. (2020), “The curious case of adversarially robust models: More data can help, double descend, or hurt generalization,” *arXiv preprint arXiv:2002.11080*.
- 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.
- Najafi, A., Maeda, S.-i., Koyama, M., and Miyato, T. (2019), “Robustness to adversarial perturbations in learning from incomplete data,” in *Advances in Neural Information Processing Systems*, pp. 5542–5552.
- Papernot, N., McDaniel, P., Swami, A., and Harang, R. (2016a), “Crafting adversarial input sequences for recurrent neural networks,” in *Military Communications Conference, MILCOM 2016-2016 IEEE*, IEEE, pp. 49–54.
- Papernot, N., McDaniel, P. D., Jha, S., Fredrikson, M., Celik, Z. B., and Swami, A. (2016b), “The limitations of deep learning in adversarial settings,” in *IEEE European Symposium on Security and Privacy, EuroS&P 2016, Saarbrücken, Germany, March 21-24, 2016*, IEEE, pp. 372–387.
- Raghunathan, A., Xie, S. M., Yang, F., Duchi, J. C., and Liang, P. (2019), “Adversarial training can hurt generalization,” *arXiv preprint arXiv:1906.06032*.
- Rice, L., Wong, E., and Kolter, J. Z. (2020), “Overfitting in adversarially robust deep learning,” *arXiv preprint arXiv:2002.11569*.
- Salman, H., Li, J., Razenshteyn, I., Zhang, P., Zhang, H., Bubeck, S., and Yang, G. (2019), “Provably robust deep learning via adversarially trained smoothed classifiers,” in *Advances in Neural Information Processing Systems*, pp. 11292–11303.
- 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*.
- Sinha, A., Namkoong, H., and Duchi, J. C. (2018), “Certifiable distributional robustness with principled adversarial training,” in *6th International Conference on Learning Representations*.
- 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, 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*.
- Wong, E., Rice, L., and Kolter, J. Z. (2020), “Fast is better than free: Revisiting adversarial training,” *arXiv preprint arXiv:2001.03994*.
- Wu, D., Wang, Y., and Xia, S.-t. (2020), “Revisiting Loss Landscape for Adversarial Robustness,” *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*.
- Ye, S., Xu, K., Liu, S., Cheng, H., Lambrechts, J.-H., Zhang, H., Zhou, A., Ma, K., Wang, Y., and Lin, X. (2019), “Adversarial robustness vs. model compression, or both,” in *The IEEE International Conference on Computer Vision (ICCV)*, vol. 2.
- Yin, D., Ramchandran, K., and Bartlett, P. L. (2019), “Rademacher complexity for adversarially robust generalization,” 97, 7085–7094.
- 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, 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*.

The structure of appendix is as follows. In Section A, we provide more numerical experiments. Section B presents the proof of Theorem 2. Section C presents the proof of Theorem 3, 4, 6 and 7. Section D shows the proof for Lemma 1 and Theorem 5. And finally Section E is for high-dimensional sparse model (Theorem 9).

A More numerical results

A.1 Low-dimensional linear models

To verify Theorem 2 and the statement that “adversarial training hurts standard testing performance”, we run a linear model this experiment. The model is set to be $d = 10$ with $\theta_{0,i} = 1$ for $i = 1, \dots, 10$. The covariance Σ is I , and for noise, $\sigma^2 = 1$.

For adversarial training, we use zero initialization, $\xi = 0.1$, and $\eta = 0.01$. We repeat 100 times to get mean and standard deviation. The results are summarized in Figure 5. From Figure 5, one can find that the adversarial testing loss is closed to $R^*(\epsilon)$, while the standard testing loss is away from 1 when $\epsilon = 0.5$.




Figure 5: Adversarial training in linear regression under low data dimension. Dashed line in the middle panel: $R^*(\epsilon)$.

A.2 Low-dimensional two-Layer networks

Here we present a numerical experiment to verify our results in Theorems 3, and 4 on lazy training.

We take $d = 3$, $\theta_0 = (1/\sqrt{3}, 1/\sqrt{3}, 1/\sqrt{3})$, $\Sigma = I$, $\sigma^2 = 1$ for the data generation model, and $n = 100$. For the two-layer neural network, we take $h = 50$, and $a_j \sim \text{Unif}[-1, 1]$. For a network in Theorem 3, we take $\phi(x) = 1/(1 + e^{-x}) - 1/2$ and $\eta = 0.2$. To match the same $\phi'(0)$ and learning rate for all three models, we take $\phi(x) = x/4$ and $\eta = 0.2$ for linear network, and $\phi(x) = x1\{x > 0\}/4$ and $\eta = 0.8$ for ReLU network. For ReLU, we adjust negative a_j 's so that $\|a^+\| = \|a^-\|$.

For initialization, we take $\delta = 0.5$. For nonlinear networks, we use fast gradient method to approximate $A_{\epsilon, \xi}$ for both training and testing, and for surrogate loss, we take $\xi = 0.01$. We run the optimization for 4000 iterations, and repeat 50 times to get mean and standard deviation for the (population) adversarial risk. To estimate the adversarial risk, we randomly simulate 10000 samples and calculate the sample adversarial loss.

The results are shown in Figure 6. Since we match $\phi'(0)$ and learning rate for all the three networks, the adversarial risks decrease in the same speed and all converges to $R^*(\epsilon)$. However, due to the existence of ξ , linear network cannot reach an adversarial testing loss as $R^*(\epsilon)$. For finite h , sigmoid networks and ReLU networks have higher adversarial testing loss than linear networks.

A.3 CIFAR-10

As mentioned in the main text, we conduct some additional experiments. In the first experiment, we run 200 epochs of adversarial training with \mathcal{L}_2 attack $\epsilon = 3.0$. The initial learning rate is 0.1, and multiplies 0.1 at the 100th and 150th epoch. The value of ξ is initialized as 0.001 and changes according to the learning rate. The first 74 epochs are the same as those in Figure 2, and we display the remaining epochs in Figure 7. From Figure 7, the performance of $\xi > 0$ is still better than the case of $\xi = 0$ for both adversarial and standard testing accuracy.




Figure 6: Adversarial training in three two-Layer neural networks (with lazy training) under low data dimension. Dashed line : $R^*(\epsilon)$.




Figure 7: Adversarial training in CIFAR-10.

In the second experiment, besides the 100 epochs of adversarial training, we use standard training to first train 50 epochs. The final standard/adversarial testing accuracy for $\xi = 0$ is 33.347(1.257)/17.285(0.241), and the one w.r.t. $\xi = 0.001$ becomes 10.313(0.532)/29.762(2.767), while training only in 100 epochs of adversarial training with $\xi = 0$ is 33.637(0.724)/17.806(0.241). To conclude, using additional standard training at the beginning may lead to a volatile training process when ϵ is large.

In the last experiment, we evaluate the attack difference. We use WideResNet34-1 in CIFAR-10 for \mathcal{L}_2 PGD-5 attack of strength 0.5 and 3.0 and calculate the attack difference. Attack step size is taken as $2\epsilon/k$ for PGD- k . The results are shown in Figure 8. When $\epsilon = 0.5$, the attack difference is around 0.16 to 0.17. When $\epsilon = 3.0$, the attack difference slowly increases to 1.5 in the end. As a result, a larger attack leads to a larger attack difference, which indicates that potential improvements should be considered to stabilize the training process when ϵ is large. A similar observation can be found when using PGD-20. Despite of that it is difficult to mathematically characterize all non-differentiable points for DNN adversarial loss, the above simulations justify the importance of accommodating this non-differentiability issue, especially for large ϵ .

A.4 High-dimensional dense models

A.4.1 Linear model

Besides the experiment in Figure 3, we further run some experiments with larger d and $\xi = 0$. Figure 9 shows the experiment with the same setting as in Figure 3 but with $\xi = 0$. When $\theta = 0$, since the adversarial training loss is not differentiable, we does not impose attack. From Figure 9, the adversarial training / testing loss have similar performance as when $\xi = 0.1$, while the difference between the gradients of adversarial training and standard training becomes larger than the case when $\xi > 0$. To explain this, since $\|\theta_\xi^{(t)}\| \rightarrow 0$ when $d/n \rightarrow \infty$, the introduction of positive ξ will leads to a surrogate attack with strength almost zero.

Figures 10 and 11 show the experiment with the same setting as in Figure 3 but with $\eta = 0.0002$ and $d = 5000$. In addition to the adversarial training / testing loss, from Figures 10 and 11, one can observe that the difference between the gradients of adversarial training and standard training gets smaller.




Figure 8: Attack difference in \mathcal{L}_2 adversarial training. Both the attack difference and “relative” attack difference (dividing ϵ) for $\epsilon = 3.0$ are larger than the ones in $\epsilon = 0.5$.




Figure 9: Adversarial training in high-dimensional setup. $\xi = 0$




Figure 10: Adversarial training in high-dimensional setup. $\xi = 0.5$, $d = 5000$




Figure 11: Adversarial training in high-dimensional setup. $\xi = 0$, $d = 5000$

A.4.2 Two-layer neural networks

Similar as in Section A.2, we conduct experiment on neural network with three different activation functions. The data generation follows those for Figure 3 with $\sigma^2 = 0.1$, $d = 1000$, $n = 20$. For neural networks, we take $h = 50$, and $a_j \sim \text{Unif}[-1, 1]$. For ReLU, we adjust negative a_j 's so that $\|a^+\| = \|a^-\|$. Initialization takes δ such that $h^\delta = d^{0.6}$. The learning rate is taken as 0.16 for linear and sigmoid networks, and 0.64 for ReLU network so

that the convergence pattern is clear in the first 100 iterations. For adversarial surrogate loss, we take $\xi = 0.1$. The results are summarized in Figure 12 for adversarial training loss and 13 for adversarial testing loss. For all three neural networks, the adversarial training loss decreases as fast as standard loss, while the adversarial testing loss are as higher than 1.




Figure 12: Adversarial training loss of adversarial training in three two-Layer neural networks (with lazy training) under high data dimension.




Figure 13: Adversarial testing loss of adversarial training in three two-Layer neural networks (with lazy training) under high data dimension.

A.5 High-dimensional sparse models

In addition to the \mathcal{L}_2 attack with $\xi = 0$ as in Figure 4, we also conduct experiments of \mathcal{L}_∞ attack with $\xi = 0$ and the two attacks with $\xi = 1e - 04$ below. When $\xi > 0$, the adversarial testing accuracy is based on the original attack, i.e. not the surrogate one. As mentioned in Figure 9, under high-dimensional setup, a positive ξ leads to the adversarial training getting closer to standard training when $n/d \rightarrow 0$. As a result, we choose a small enough ξ such that the adversarial testing performance is closed to the case when $\xi = 0$. Similar as Figure 4, all the experiments in Figures 14, 15, and 16 shows a better performance using LASSO.

B Proofs for low dimension linear model

In the proof of Theorem 2, we assume v^2 is a constant number, which implies that $\|\theta^*\| \leq B_0$ for some B_0 . After investigating the results for bounded v , we then use a trick to extend to the case when v is changing.

Lemma 2. *Under the model in (4), when $\|\theta_0\| \leq B_0$ for some constant $B_0 > 0$, there exists some function $g_1(\delta, \xi, d, n)$ such that*

$$P \left(\sup_{\|\theta\| \leq B_0} \left\| \nabla \widehat{R}_{L,\xi}(\theta, \epsilon) - \nabla R_{L,\xi}(\theta, \epsilon) \right\| > \delta \right) \leq e^{-g_1(\delta, \xi, d, n)},$$

where $\nabla \widehat{R}_{L,\xi}(\theta, \epsilon)$ and $\nabla R_{L,\xi}(\theta, \epsilon)$ denote the gradient of θ **after fixing** $A_{\epsilon,\xi}$ (so it is not $\partial \widehat{R}_{L,\xi} / \partial \theta$ or $\partial R_{L,\xi} / \partial \theta$ if $\xi > 0$).

Proof of Lemma 2. Assume $\theta \in B(0, B_0)$. For any sample i , the gradient of θ on surrogate adversarial loss is




Figure 14: Comparison on standard (upper)/adversarial (lower) test accuracies between training with/without \mathcal{L}_1 penalty under \mathcal{L}_2 attack. Attack strength $\epsilon = 3$.

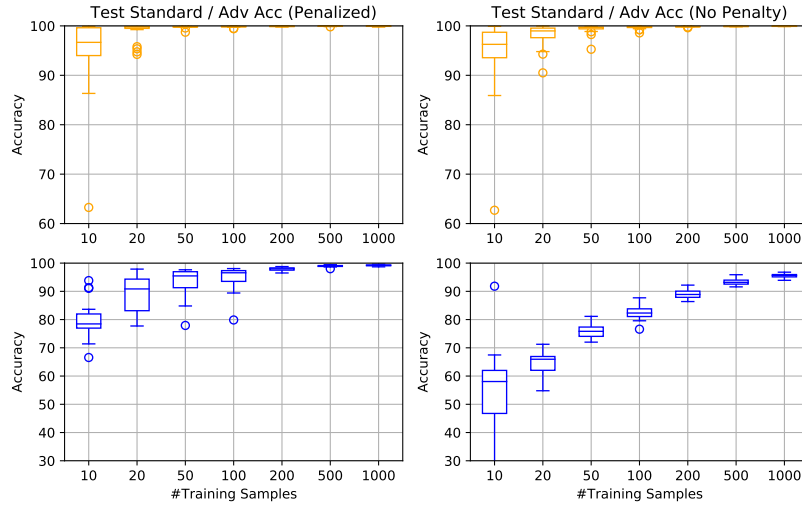


Figure 15: Comparison on standard (upper)/adversarial (lower) test accuracies between training with/without \mathcal{L}_1 penalty under \mathcal{L}_∞ attack with $\xi = 1e - 04$. Attack strength $\epsilon = 0.3$.

(we ignore the constant multiplier of $\xi^2/|x_i^\top \theta - y_i|^2$)

$$\begin{aligned}
 & 2x_i(x_i^\top \theta - y_i) + 2\epsilon \frac{\theta}{\sqrt{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2}} |x_i^\top \theta - y_i| \\
 & + 2\epsilon x_i \frac{\|\theta\|^2}{\sqrt{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2}} \text{sgn}(x_i^\top \theta - y_i) + 2\epsilon^2 \theta \frac{\|\theta\|^2}{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2}.
 \end{aligned}$$

Therefore, by Bernstein inequality, for any fixed θ ,

$$P\left(\left\|\nabla \widehat{R}_{L,\xi}(\theta, \epsilon) - \nabla R_{L,\xi}(\theta, \epsilon)\right\| > \delta\right) \leq e^{-c_1 n \delta^2/d}.$$




Figure 16: Comparison on standard (upper)/adversarial (lower) test accuracy between training with/without \mathcal{L}_1 penalty under \mathcal{L}_∞ attack.

Our aim is to figure out a bound for $\sup_{\|\theta\| \leq B_0} \left\| \nabla \hat{R}_{L,\xi}(\theta, \epsilon) - \nabla R_{L,\xi}(\theta, \epsilon) \right\|$, thus we consider the following decomposition:

$$\begin{aligned} \left\| \nabla \hat{R}_{L,\xi}(\theta, \epsilon) - \nabla R_{L,\xi}(\theta, \epsilon) \right\| &\leq \left\| \nabla \hat{R}_{L,\xi}(\theta, \epsilon) - \nabla \hat{R}_{L,\xi}(\theta_k, \epsilon) \right\| \\ &\quad + \left\| \nabla R_{L,\xi}(\theta, \epsilon) - \nabla R_{L,\xi}(\theta_k, \epsilon) \right\| \\ &\quad + \left\| \nabla \hat{R}_{L,\xi}(\theta_k, \epsilon) - \nabla R_{L,\xi}(\theta_k, \epsilon) \right\|, \end{aligned}$$

where θ_k is an element in the fixed sequence $\{\theta_j\}_{j=1,\dots}$. Now we first introduce how to design $\{\theta_j\}_{j=1,\dots}$. For the ball $B(0, B_0)$, we use balls with radius $1/M$ to cover it. Then there are total $c_d B_0^d M^d$ balls for some constant c_d which only depends on d . Denote θ_k as the center of the k th ball and thus we obtain $\{\theta_j\}_{j=1,\dots}$. The worst case among the $c_d M^d$ centers of balls satisfies

$$P \left(\sup_k \left\| \nabla \hat{R}_{L,\xi}(\theta_k, \epsilon) - \nabla R_{L,\xi}(\theta_k, \epsilon) \right\| > \delta \right) \leq c_d B_0^d M^d e^{-c_1 n \delta^2 / d}.$$

For any $\theta \in B(0, B_0)$, the distance from θ to its nearest θ_k is at most $1/M$, thus there exists some constant c_2 such that

$$\inf_k \left\| \nabla R_{L,\xi}(\theta_k, \epsilon) - \nabla R_{L,\xi}(\theta, \epsilon) \right\| \leq c_2 / M \xi.$$

In terms of $\nabla \hat{R}_{L,\xi}(\theta_k, \epsilon) - \nabla \hat{R}_{L,\xi}(\theta, \epsilon)$, we have

$$\begin{aligned} &\nabla \hat{R}_{L,\xi}(\theta_k, \epsilon) - \nabla \hat{R}_{L,\xi}(\theta, \epsilon) \\ &= \frac{2}{n} \sum_{i=1}^n x_i x_i^\top (\theta_k - \theta) + \frac{2\epsilon}{n} \sum_{i=1}^n \left(\frac{\theta_k}{\sqrt{\|\theta_k\|^2 + \xi^2 / |x_i^\top \theta_k - y_i|^2}} |x_i^\top \theta_k - y_i| - \frac{\theta}{\sqrt{\|\theta\|^2 + \xi^2 / |x_i^\top \theta - y_i|^2}} |x_i^\top \theta - y_i| \right) \\ &\quad + \frac{2\epsilon}{n} \sum_{i=1}^n \left(\frac{\|\theta_k\|^2}{\sqrt{\|\theta_k\|^2 + \xi^2 / |x_i^\top \theta_k - y_i|^2}} x_i \operatorname{sgn}(x_i^\top \theta_k - y_i) - \frac{\|\theta\|^2}{\sqrt{\|\theta\|^2 + \xi^2 / |x_i^\top \theta - y_i|^2}} x_i \operatorname{sgn}(x_i^\top \theta - y_i) \right) \\ &\quad + 2\epsilon^2 \frac{1}{n} \sum_{i=1}^n \theta_k \frac{\|\theta_k\|^2}{\|\theta_k\|^2 + \xi^2 / |x_i^\top \theta_k - y_i|^2} - \theta \frac{\|\theta\|^2}{\|\theta\|^2 + \xi^2 / |x_i^\top \theta - y_i|^2}, \end{aligned}$$

thus taking $\widehat{\Sigma} = \frac{1}{n} \sum_{i=1}^n x_i x_i^\top$, and denoting $\|A\|$ as the operator norm of matrix A ,

$$\begin{aligned}
 & \left\| \nabla \widehat{R}_{L,\xi}(\theta_k, \epsilon) - \nabla \widehat{R}_{L,\xi}(\theta, \epsilon) \right\| \\
 \leq & 2 \left\| \widehat{\Sigma} - \Sigma \right\| \left\| \theta_k - \theta \right\| + \left\| \frac{2\epsilon}{n} \sum_{i=1}^n \left(\frac{\theta_k}{\sqrt{\|\theta_k\|^2 + \xi^2/|x_i^\top \theta_k - y_i|^2}} |x_i^\top \theta_k - y_i| - \frac{\theta}{\sqrt{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2}} |x_i^\top \theta_k - y_i| \right) \right\| \\
 & + \left\| \frac{2\epsilon}{n} \sum_{i=1}^n \left(\frac{\theta}{\sqrt{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2}} |x_i^\top \theta_k - y_i| - \frac{\theta}{\sqrt{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2}} |x_i^\top \theta - y_i| \right) \right\| \\
 & + \left\| \frac{2\epsilon}{n} \sum_{i=1}^n \left(\frac{\|\theta_k\|^2}{\sqrt{\|\theta_k\|^2 + \xi^2/|x_i^\top \theta_k - y_i|^2}} x_i \operatorname{sgn}(x_i^\top \theta_k - y_i) - \frac{\|\theta\|^2}{\sqrt{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2}} x_i \operatorname{sgn}(x_i^\top \theta_k - y_i) \right) \right\| \\
 & + \left\| \frac{2\epsilon}{n} \sum_{i=1}^n \left(\frac{\|\theta\|^2}{\sqrt{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2}} x_i \operatorname{sgn}(x_i^\top \theta_k - y_i) - \frac{\|\theta\|^2}{\sqrt{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2}} x_i \operatorname{sgn}(x_i^\top \theta - y_i) \right) \right\| \\
 & + \left\| 2\epsilon^2 \frac{1}{n} \sum_{i=1}^n \theta_k \frac{\|\theta_k\|^2}{\|\theta_k\|^2 + \xi^2/|x_i^\top \theta_k - y_i|^2} - \theta \frac{\|\theta\|^2}{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2} \right\| \\
 := & A_1 + A_2 + A_3 + A_4 + A_5 + A_6.
 \end{aligned}$$

Following Lemma A.2 of [Ing and Lai \(2011\)](#) to quantify $\|\widehat{\Sigma} - \Sigma\|$, there is some large enough c'_2 such that with probability tending to 1,

$$A_1 \leq 2c'_2 \sqrt{\frac{d \log n}{n}} \|\theta_k - \theta\|.$$

For A_2 ,

$$\begin{aligned}
 A_2 &= \left\| \frac{2\epsilon}{n} \sum_{i=1}^n \left(\frac{\theta_k}{\sqrt{\|\theta_k\|^2 + \xi^2/|x_i^\top \theta_k - y_i|^2}} - \frac{\theta}{\sqrt{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2}} \right) |x_i^\top \theta_k - y_i| \right\| \\
 &\leq \left\| \frac{2\epsilon}{n} \sum_{i=1}^n \left(\frac{\theta_k - \theta}{\sqrt{\|\theta_k\|^2 + \xi^2/|x_i^\top \theta_k - y_i|^2}} \right) |x_i^\top \theta_k - y_i| \right\| \\
 &\quad + \left\| \frac{2\epsilon}{n} \sum_{i=1}^n \left(\frac{\theta}{\sqrt{\|\theta_k\|^2 + \xi^2/|x_i^\top \theta_k - y_i|^2}} - \frac{\theta}{\sqrt{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2}} \right) |x_i^\top \theta_k - y_i| \right\| \\
 &\leq \frac{1}{M} \frac{2\epsilon}{n} \sum_{i=1}^n \frac{|x_i^\top \theta_k - y_i|}{\sqrt{\|\theta_k\|^2 + \xi^2/|x_i^\top \theta_k - y_i|^2}} \\
 &\quad + \left\| \frac{2\epsilon}{n} \sum_{i=1}^n \left(\frac{\theta}{\sqrt{\|\theta_k\|^2 + \xi^2/|x_i^\top \theta_k - y_i|^2}} - \frac{\theta}{\sqrt{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2}} \right) |x_i^\top \theta_k - y_i| \right\|.
 \end{aligned}$$

Using Bernstein inequality, for some constant c'_1 ,

$$P \left(\sup_k \left| \frac{2\epsilon}{n} \sum_{i=1}^n \frac{|x_i^\top \theta_k - y_i|}{\sqrt{\|\theta_k\|^2 + \xi^2/|x_i^\top \theta_k - y_i|^2}} - \mathbb{E} \frac{2\epsilon |x^\top \theta_k - y|}{\sqrt{\|\theta_k\|^2 + \xi^2/|x^\top \theta_k - y|^2}} \right| \geq \delta \right) \leq c_d B_0^d M^d e^{-c_1 n \delta^2 \xi^2}.$$

In addition, one can take c'_2 large enough such that

$$\begin{aligned}
 & \frac{2\epsilon}{n} \sum_{i=1}^n \left| \frac{1}{\sqrt{\|\theta_k\|^2 + \xi^2/|x_i^\top \theta_k - y_i|^2}} - \frac{1}{\sqrt{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2}} \right| |x_i^\top \theta_k - y_i| \\
 = & \frac{2\epsilon}{n} \sum_{i=1}^n \left| \frac{|x_i^\top \theta_k - y_i|}{\sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2}} - \frac{|x_i^\top \theta - y_i|}{\sqrt{\|\theta\|^2 |x_i^\top \theta - y_i|^2 + \xi^2}} \right| |x_i^\top \theta_k - y_i| \\
 = & \frac{2\epsilon}{n} \sum_{i=1}^n \left| \frac{|x_i^\top \theta_k - y_i| \sqrt{\|\theta\|^2 |x_i^\top \theta - y_i|^2 + \xi^2} - |x_i^\top \theta - y_i| \sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2}}{\sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2} \sqrt{\|\theta\|^2 |x_i^\top \theta - y_i|^2 + \xi^2}} \right| |x_i^\top \theta_k - y_i|.
 \end{aligned}$$

Furthermore,

$$\begin{aligned}
 & \left| |x_i^\top \theta_k - y_i| \sqrt{\|\theta\|^2 |x_i^\top \theta - y_i|^2 + \xi^2} - |x_i^\top \theta - y_i| \sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2} \right| \\
 & \leq \left| |x_i^\top \theta_k - y_i| \sqrt{\|\theta\|^2 |x_i^\top \theta - y_i|^2 + \xi^2} - |x_i^\top \theta_k - y_i| \sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2} \right| \\
 & \quad + \left| |x_i^\top \theta_k - y_i| \sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2} - |x_i^\top \theta - y_i| \sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2} \right| \\
 & \leq |x_i^\top \theta_k - y_i| \left| \sqrt{\|\theta\|^2 |x_i^\top \theta - y_i|^2 + \xi^2} - \sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2} \right| \\
 & \quad + \|x_i\| \|\theta_k - \theta\| \sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2}.
 \end{aligned}$$

As a result, we have

$$\begin{aligned}
 & \left\| \frac{2\epsilon}{n} \sum_{i=1}^n \left(\frac{\theta}{\sqrt{\|\theta_k\|^2 + \xi^2 / |x_i^\top \theta_k - y_i|^2}} - \frac{\theta}{\sqrt{\|\theta\|^2 + \xi^2 / |x_i^\top \theta - y_i|^2}} \right) |x_i^\top \theta_k - y_i| \right\| \\
 & \leq \|\theta\| \frac{2\epsilon}{n} \sum_{i=1}^n \left| \frac{1}{\sqrt{\|\theta_k\|^2 + \xi^2 / |x_i^\top \theta_k - y_i|^2}} - \frac{1}{\sqrt{\|\theta\|^2 + \xi^2 / |x_i^\top \theta - y_i|^2}} \right| |x_i^\top \theta_k - y_i| \\
 & \leq \|\theta\| \frac{2\epsilon}{n} \sum_{i=1}^n \left| \frac{\|x_i\| \|\theta_k - \theta\| \sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2}}{\sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2} \sqrt{\|\theta\|^2 |x_i^\top \theta - y_i|^2 + \xi^2}} \right| |x_i^\top \theta_k - y_i| \\
 & \quad + \|\theta\| \frac{2\epsilon}{n} \sum_{i=1}^n \left| \frac{\sqrt{\|\theta\|^2 |x_i^\top \theta - y_i|^2 + \xi^2} - \sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2}}{\sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2} \sqrt{\|\theta\|^2 |x_i^\top \theta - y_i|^2 + \xi^2}} \right| |x_i^\top \theta_k - y_i|^2 \\
 & \leq \frac{2\epsilon B_0}{nM} \sum_{i=1}^n \frac{\|x_i\|}{\xi} |x_i^\top \theta_k - y_i| \\
 & \quad + \frac{2\epsilon}{n} \sum_{i=1}^n \left| \frac{\|\theta\|}{\sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2}} - \frac{\|\theta\|}{\sqrt{\|\theta\|^2 |x_i^\top \theta - y_i|^2 + \xi^2}} \right| |x_i^\top \theta_k - y_i|^2 \\
 & = \frac{2\epsilon B_0}{nM} \sum_{i=1}^n \frac{\|x_i\|}{\xi} |x_i^\top \theta_k - y_i| \\
 & \quad + \frac{2\epsilon}{n} \sum_{i=1}^n \left| \frac{\|\theta_k\| |x_i^\top \theta_k - y_i| - \|\theta\| |x_i^\top \theta_k - y_i|}{\sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2}} \right| |x_i^\top \theta_k - y_i| \\
 & \leq \frac{2\epsilon B_0}{nM} \sum_{i=1}^n \frac{\|x_i\|}{\xi} |x_i^\top \theta_k - y_i| \\
 & \quad + \frac{2\epsilon}{n} \sum_{i=1}^n \left| \frac{\|\theta_k\| |x_i^\top \theta_k - y_i| - \|\theta_k\| |x_i^\top \theta - y_i|}{\sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2}} \right| |x_i^\top \theta_k - y_i| + \frac{2\epsilon}{n} \sum_{i=1}^n \left| \frac{\|\theta_k\| |x_i^\top \theta - y_i| - \|\theta\| |x_i^\top \theta - y_i|}{\sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2}} \right| |x_i^\top \theta_k - y_i| \\
 & \leq \frac{2\epsilon B_0}{nM} \sum_{i=1}^n \frac{\|x_i\|}{\xi} |x_i^\top \theta_k - y_i| \\
 & \quad + B_0 \frac{2\epsilon}{n} \sum_{i=1}^n \frac{\|x_i\|}{M\xi} |x_i^\top \theta_k - y_i| + \frac{2\epsilon}{n} \sum_{i=1}^n \left| \frac{\|\theta_k - \theta\| |x_i^\top \theta_k - y_i| + \|\theta_k - \theta\|^2 \|x_i\|}{\sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2}} \right| |x_i^\top \theta_k - y_i| \\
 & \leq \frac{2\epsilon B_0}{nM} \sum_{i=1}^n \frac{\|x_i\|}{\xi} |x_i^\top \theta_k - y_i| + B_0 \frac{2\epsilon}{n} \sum_{i=1}^n \frac{\|x_i\|}{M\xi} |x_i^\top \theta_k - y_i| + \frac{2\epsilon}{n} \frac{1}{M\xi} \sum_{i=1}^n |x_i^\top \theta_k - y_i|^2 + \frac{2\epsilon}{n} \frac{1}{M^2\xi} \sum_{i=1}^n |x_i^\top \theta_k - y_i| \|x_i\|,
 \end{aligned}$$

which can be bounded using Bernstein inequality as well.

On the other hand,

$$\begin{aligned} A_3 &= \left\| \frac{2\epsilon}{n} \sum_{i=1}^n \frac{\theta}{\sqrt{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2}} (|x_i^\top \theta_k - y_i| - |x_i^\top \theta - y_i|) \right\| \\ &\leq \left| \frac{2\epsilon}{n} \sum_{i=1}^n (|x_i^\top \theta_k - y_i| - |x_i^\top \theta - y_i|) \right| \leq 2\epsilon \|\theta_k - \theta\| (\sup_i \|x_i\|). \end{aligned}$$

Further, assume $\|\theta_k\| \leq \|\theta\|$,

$$\begin{aligned} A_4 &= \frac{2\epsilon}{n} \sum_{i=1}^n \left| \frac{\|\theta_k\|^2}{\sqrt{\|\theta_k\|^2 + \xi^2/|x_i^\top \theta_k - y_i|^2}} - \frac{\|\theta\|^2}{\sqrt{\|\theta\|^2 + \xi^2/|x_i^\top \theta - y_i|^2}} \right| \|x_i\| \\ &\leq \frac{2\epsilon}{n} \sum_{i=1}^n \left| \frac{\|\theta_k\|^2}{\sqrt{\|\theta_k\|^2 + \xi^2/|x_i^\top \theta_k - y_i|^2}} - \|\theta\| + \|\theta_k\| - \|\theta_k\| \right| \|x_i\| \\ &\leq \frac{2\epsilon}{n} \sum_{i=1}^n \|\theta_k - \theta\| \|x_i\| \\ &\quad + \frac{2\epsilon}{n} \sum_{i=1}^n \left| \frac{\|\theta_k\|^2 - \|\theta_k\| \sqrt{\|\theta_k\|^2 + \xi^2/|x_i^\top \theta_k - y_i|^2}}{\sqrt{\|\theta_k\|^2 + \xi^2/|x_i^\top \theta_k - y_i|^2}} \right| \|x_i\| \\ &\leq \frac{2\epsilon}{n} \sum_{i=1}^n \|\theta_k - \theta\| \|x_i\| + \frac{2\epsilon}{n} \sum_{i=1}^n \frac{\|\theta_k\| \xi}{\sqrt{\|\theta_k\|^2 |x_i^\top \theta_k - y_i|^2 + \xi^2}} \|x_i\|, \end{aligned}$$

where both the terms can be bounded using Bernstein inequality.

For A_5 ,

$$\begin{aligned} A_5 &\leq B_0 \left\| \frac{2\epsilon}{n} \sum_{i=1}^n x_i \operatorname{sgn}(x_i^\top \theta_k - y_i) - x_i \operatorname{sgn}(x_i^\top \theta - y_i) \right\| \\ &\leq B_0 \left(\frac{2\epsilon}{n} \sum_{i=1}^n \|x_i\| 1\{|x_i^\top \theta_k - y_i| \leq \|x_i\| \|\theta_k - \theta\|\} \right). \end{aligned}$$

For any θ_k , $x_i^\top \theta_k - y_i$ follows a Gaussian distribution with some mean and the variance is finite. As a result, for some $c_3 > 0$, we have

$$P \left(\left| \frac{1}{n} \sum_{i=1}^n \|x_i\| 1\{|x_i^\top \theta_k - y_i| \leq \kappa\} - \mathbb{E}(\|x\| 1\{|x^\top \theta_k - y| \leq \kappa\}) \right| > \delta \right) \leq e^{-nc_3 \delta^2 / (d\kappa)}.$$

Note that $\mathbb{E}(\|x\| 1\{|x^\top \theta_k - y| \leq \kappa\}) = O(\kappa \sqrt{d})$. Finally,

$$A_6 \leq c'_2 \|\theta - \theta_k\|.$$

Putting $\|\nabla \widehat{R}_{L,\xi}(\theta, \epsilon) - \widehat{\nabla} R_{L,\xi}(\theta_k, \epsilon)\|$, $\|\nabla R_{L,\xi}(\theta, \epsilon) - \nabla R_{L,\xi}(\theta_k, \epsilon)\|$, and $\|\nabla \widehat{R}_{L,\xi}(\theta_k, \epsilon) - \nabla R_{L,\xi}(\theta_k, \epsilon)\|$ together, there exists some function g_1 such that

$$P \left(\sup_{\|\theta\| \leq B_0} \|\nabla \widehat{R}_{L,\xi}(\theta, \epsilon) - \nabla R_{L,\xi}(\theta, \epsilon)\| > \delta \right) \leq e^{-g_1(\delta, \xi, d, n)}.$$

When using this lemma, instead of figuring out the complicated form of g_1 , one can directly use (ξ, δ, d, n) to determine a suitable M , and then use the upper bound of $\|\nabla R_{L,\xi}(\theta, \epsilon) - \nabla R_{L,\xi}(\theta_k, \epsilon)\|$, $\|\nabla \widehat{R}_{L,\xi}(\theta_k, \epsilon) - \nabla R_{L,\xi}(\theta_k, \epsilon)\|$ and A_1 to A_5 to obtain a probability bound. \square

Proof of Theorem 2. Rewrite $\theta_\xi^{(t)}$ as $\theta^{(t)}$ for simplicity. After we defining $\nabla R_{L,\xi}$, unlike $\nabla R_{L,0}$, we may not find a function whose gradient is $\nabla R_{L,\xi}$. As a result, we propose another function $\tilde{R}_{L,\xi}(\theta, \epsilon)$ and figure out its gradient $\partial \tilde{R}_{L,\xi}(\theta, \epsilon)/\partial \theta$, then bound the difference between $\partial \tilde{R}_{L,\xi}(\theta, \epsilon)/\partial \theta$ and $\nabla R_{L,\xi}$.

Denote $\tilde{R}_{L,\xi}(\theta, \epsilon)$ as

$$\tilde{R}_{L,\xi}(\theta, \epsilon) = \mathbb{E} \left[(x^\top \theta - y)^2 + 2\epsilon \sqrt{\|\theta\|^2 + \xi^2/|x^\top \theta - y|^2} |x^\top \theta - y| + \epsilon^2 \|\theta\|^2 \right],$$

then taking gradient on \tilde{R}_ξ w.r.t θ , it becomes

$$\begin{aligned} & \nabla \tilde{R}_{L,\xi}(\theta, \epsilon) \\ = & \mathbb{E} \left[2x(x^\top \theta - y) + 2\epsilon \frac{\theta}{\sqrt{\|\theta\|^2 + \xi^2/|x^\top \theta - y|^2}} |x^\top \theta - y| - \frac{2\epsilon \xi^2 x}{\sqrt{\|\theta\|^2 + \xi^2/|x^\top \theta - y|^2}} \frac{\text{sgn}(x^\top \theta - y)}{(x^\top \theta - y)^2} \right. \\ & \left. + 2\epsilon x \sqrt{\|\theta\|^2 + \xi^2/|x^\top \theta - y|^2} \text{sgn}(x^\top \theta - y) + 2\epsilon^2 \theta \right] \end{aligned}$$

For $\nabla R_{L,\xi}(\theta, \epsilon)$, we have

$$\begin{aligned} & \nabla R_{L,\xi}(\theta, \epsilon) \\ = & \mathbb{E} \left[2x(x^\top \theta - y) + 2\epsilon \frac{\theta}{\sqrt{\|\theta\|^2 + \xi^2/|x^\top \theta - y|^2}} |x^\top \theta - y| \right. \\ & \left. + 2\epsilon x \frac{\|\theta\|^2}{\sqrt{\|\theta\|^2 + \xi^2/|x^\top \theta - y|^2}} \text{sgn}(x^\top \theta - y) + 2\epsilon^2 \theta \frac{\|\theta\|^2}{\|\theta\|^2 + \xi^2/|x^\top \theta - y|^2} \right] \end{aligned}$$

We assume $\theta \in B(0, B_1)$, then there exists B_2 such that both $\|\nabla R_{L,\xi}(\epsilon, \theta)\|$ and $\|\partial \tilde{R}_{L,\xi}(\theta, \epsilon)/\partial \theta\|$ are bounded by some large constant $B_2 > 0$.

To compare the difference between $\frac{\partial \tilde{R}_{L,\xi}}{\partial \theta}$ and $\nabla R_{L,\xi}(\theta, \epsilon)$, we have

$$\left\| \nabla R_{L,\xi}(\theta, \epsilon) - \frac{\partial \tilde{R}_{L,\xi}}{\partial \theta} \right\| = 2\epsilon \left\| \theta \mathbb{E} \left(\frac{\xi^2/|x^\top \theta - y|^2}{\|\theta\|^2 + \xi^2/|x^\top \theta - y|^2} \right) \right\|.$$

As a result, there exists some $B_3 > 0$ such that

$$\left\| \nabla R_\xi(\theta, \epsilon) - \frac{\partial \tilde{R}_{L,\xi}}{\partial \theta} \right\| \leq \frac{\xi}{B_3} \left\| \frac{\partial \tilde{R}_{L,\xi}}{\partial \theta} \right\|.$$

Therefore the gradient of $\tilde{R}_{L,\xi}$ is the dominant term when updating θ in each iteration.

Next we show that $\tilde{R}_{L,\xi}(\theta^{(t+1)}, \epsilon)$ is smaller than $\tilde{R}_{L,\xi}(\theta^{(t)}, \epsilon)$ in probability. From the definition of \tilde{R} , similar with Proposition 1, there exists L such that

$$\tilde{R}_{L,\xi}(\theta^{(t+1)}, \epsilon) \leq \tilde{R}_{L,\xi}(\theta^{(t)}, \epsilon) + \left(\frac{\partial \tilde{R}_{L,\xi}}{\partial \theta^{(t)}} \right)^\top (\theta^{(t+1)} - \theta^{(t)}) + \frac{L}{\xi} \|\theta^{(t+1)} - \theta^{(t)}\|^2.$$

From Lemma 2, with probability at least $1 - e^{-g_1(\delta, \xi, d, n)}$,

$$\left\| \nabla \hat{R}_{L,\xi}(\theta, \epsilon) - \nabla R_{L,\xi}(\theta, \epsilon) \right\| < \delta,$$

thus

$$\left(\frac{\partial \tilde{R}_{L,\xi}}{\partial \theta^{(t)}} \right)^\top (\theta^{(t+1)} - \theta^{(t)}) = -\eta \left(\frac{\partial \tilde{R}_{L,\xi}}{\partial \theta^{(t)}} \right)^\top \nabla \hat{R}_{L,\xi}(\theta^{(t)}, \epsilon) \leq -\eta \left\| \frac{\partial \tilde{R}_{L,\xi}}{\partial \theta^{(t)}} \right\|^2 (1 - \xi/B_3) + \eta \delta B_2,$$

and

$$\begin{aligned}\|\theta^{(t+1)} - \theta^{(t)}\|^2 &= \eta^2 \|\nabla \widehat{R}_{L,\xi}(\theta^{(t)}, \epsilon)\|^2 \leq \eta^2 \|\nabla R_{L,\xi}(\theta^{(t)}, \epsilon)\|^2 + 2\eta^2 \delta B_2 + \eta^2 B_2^2 \delta^2 \\ &\leq \eta^2 \left\| \frac{\partial \widetilde{R}_{L,\xi}}{\partial \theta^{(t)}} \right\|^2 (1 + \xi/B_3)^2 + 2\eta^2 \delta B_2 + o,\end{aligned}$$

where $A + o$ represents $A + a$ where $a/A \rightarrow 0$ for expressions a and A .

Therefore,

$$\widetilde{R}_{L,\xi}(\theta^{(t+1)}, \epsilon) \leq \widetilde{R}_{L,\xi}(\theta^{(t)}, \epsilon) - \eta \left\| \frac{\partial \widetilde{R}_{L,\xi}}{\partial \theta^{(t)}} \right\|^2 (1 - \xi/B_3) + \eta \delta B_2 + \frac{L\eta^2}{\xi} \left\| \frac{\partial \widetilde{R}_{L,\xi}}{\partial \theta^{(t)}} \right\|^2 (1 + 2\xi/B_3) + o.$$

When taking $\eta = \frac{\xi}{2L}$,

$$\widetilde{R}_{L,\xi}(\theta^{(t+1)}, \epsilon) \leq \widetilde{R}_{L,\xi}(\theta^{(t)}, \epsilon) - \frac{\xi}{4L} \left\| \frac{\partial \widetilde{R}_{L,\xi}}{\partial \theta^{(t)}} \right\|^2 + \frac{\xi^2 B_2^2}{LB_3} + \frac{\xi \delta B_2}{2L} + o. \quad (10)$$

From (10), when taking $\xi \rightarrow 0$ and $\delta \rightarrow 0$, we know that $\widetilde{R}_{L,\xi}(\theta^{(t)}, \epsilon)$ will decrease in t at the beginning of training; thus based on the shape of $\widetilde{R}_{L,\xi}$ and $\|\theta^{(0)}\| \leq B_0$, one can use induction to show the existence of B_1 , i.e. $\theta^{(t)} \in B(0, B_1)$ for any t .

To bound the difference $\widetilde{R}_{L,\xi}(\theta^{(t+1)}, \epsilon)$ and $\widetilde{R}_{L,\xi}(\theta^*, \epsilon)$, since $\widetilde{R}_{L,\xi}$ is a convex function, we have

$$\widetilde{R}_{L,\xi}(\theta^{(t)}, \epsilon) \leq \widetilde{R}_{L,\xi}(\theta^*, \epsilon) + \left(\frac{\partial \widetilde{R}_{L,\xi}}{\partial \theta^{(t)}} \right)^\top (\theta^{(t)} - \theta^*).$$

Thus

$$\begin{aligned}\widetilde{R}_{L,\xi}(\theta^{(t+1)}, \epsilon) &\leq \widetilde{R}_{L,\xi}(\theta^*, \epsilon) + \left(\frac{\partial \widetilde{R}_{L,\xi}}{\partial \theta^{(t)}} \right)^\top (\theta^{(t)} - \theta^*) - \frac{\xi}{4L} \left\| \frac{\partial \widetilde{R}_{L,\xi}}{\partial \theta^{(t)}} \right\|^2 + \frac{\xi^2 B_2^2}{LB_3} + \frac{\xi \delta B_2}{2L} + o \\ &\leq \widetilde{R}_{L,\xi}(\theta^*, \epsilon) + \nabla \widehat{R}_{L,\xi}(\theta^{(t)}, \epsilon)^\top (\theta^{(t)} - \theta^*) - \frac{\xi}{4L} \left\| \nabla \widehat{R}_{L,\xi}(\theta^{(t)}, \epsilon) \right\|^2 + \frac{\xi B_1 B_2}{B_3} + \delta B_2 + o \\ &= \widetilde{R}_{L,\xi}(\theta^*, \epsilon) - \frac{L}{\xi} \left\| \theta^{(t+1)} - \theta^* \right\|^2 + \frac{L}{\xi} \left\| \theta^{(t)} - \theta^* \right\|^2 + \frac{\xi B_1 B_2}{B_3} + \delta B_2 + o,\end{aligned}$$

where the last step is obtained from the updating rule $\theta^{(t+1)} - \theta^{(t)} = -\eta \nabla \widehat{R}_{L,\xi}(\theta^{(t)}, \epsilon)$.

As a result, summing up from $t = 1$ to T ,

$$\sum_{t=1}^T \widetilde{R}_{L,\xi}(\theta^{(t)}, \epsilon) - \widetilde{R}_{L,\xi}(\theta^*, \epsilon) \leq \frac{L}{\xi} \|\theta^{(0)} - \theta^*\|^2 + \frac{T\xi B_1 B_2}{B_3} + T\delta B_2 + o.$$

If $\theta^{(t)}$ is away from θ^* , then $\widetilde{R}_{L,\xi}(\theta^{(t)}, \epsilon)$ reduces in every step.

When $\theta^{(t)}$ is close to θ^* , since $\xi \rightarrow 0$, we have $\widetilde{R}_{L,\xi}(\theta, \epsilon) - R_{L,\xi}(\theta, \epsilon) \rightarrow 0$, thus through taking suitable choice of (ξ, t, η) , we can get $R_{L,\xi}(\theta^{(T)}, \epsilon) \rightarrow R_{L,0}(\theta^*, \epsilon)$.

The proof of $\widehat{R}_{L,\xi}(\theta^{(T)}, \epsilon) - R_{L,\xi}(\theta^{(T)}, \epsilon) \rightarrow 0$ in probability follows a concentration bound similar as in Lemma 2.

Finally we relax the condition of $\|\theta_0\| \leq B_0$. For $y_i = \theta_0^\top x_i + \varepsilon_i$ with $\|\theta_0\|_\Sigma^2 + \sigma^2 = 1$, $\sigma > \xi/v$, denote $y'_i := v y_i = (v\theta_0)^\top x_i + v\varepsilon_i := (\theta'_0)^\top x_i + \varepsilon'_i$ for some $v \rightarrow \infty$. In this case, if the initialization satisfy $(\theta^{(0)}) = v\theta^{(0)}$ and $\xi' = v^2 \xi$, then for any $t = 1, \dots, T$, we always have $(\theta^{(t)}) = v\theta^{(t)}$. Therefore, since $\theta^{(T)} \rightarrow \theta^*$, we also have $\theta^{(T)}/v \rightarrow \theta^*$, which is just the minimizer of population adversarial loss for y'_i .

□

C Proofs for low-dimensional nonlinear network

Proof of Theorem 3. In the proof, we first consider taking gradient on loss w.r.t x to get the attack direction, i.e. fast gradient method to obtain the attack. And after the main proof, we discuss how to adapt the true attack.

We consider three optimization problems: (1) a linear network using zero initialization, (2) a nonlinear network using zero initialization, and (3) a nonlinear network with vanishing initialization. Extending from Theorem 2, we know how (1) works, then we bound the difference between (1), (2), and (3).

We first bound the difference among (1) and (2) assuming ξ and t are any arbitrary number. And finally choose some suitable ξ and t to ensure the consistency.

Denote

$$A_{\epsilon, \xi, OP1}(f, x, y) = \epsilon \operatorname{sgn}(f(x) - y) \frac{\frac{1}{\sqrt{h}} \sum_{j=1}^h a_j \phi'(0) \theta_j}{\sqrt{\left\| \frac{1}{\sqrt{h}} \sum_{j=1}^h a_j \phi'(0) \theta_j \right\|^2 + \xi^2 / |f_\theta(x) - y|^2}},$$

and for the weight θ_j for the j th node,

$$g_{j, \xi}^{OP1}(\theta, \epsilon) = \frac{2}{n} \sum_{i=1}^n \frac{\phi(0)' a_j}{\sqrt{h}} [x_i + A_{\epsilon, \xi, OP1}(f, x_i, y_i)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h (\phi(0)' x_i^\top \theta_k a_k + \phi'(0) A_{\epsilon, \xi, OP1}(f, x_i, y_i)^\top \theta_k a_k) - y_i \right].$$

$$g_{j, \xi}^{OP2}(\theta, \epsilon) = \frac{2}{n} \sum_{i=1}^n \frac{\phi'((x + A_{\epsilon, \xi}(f, x, y))^\top \theta_j) a_j}{\sqrt{h}} [x_i + A_{\epsilon, \xi}(f, x_i, y_i)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h \phi((x + A_{\epsilon, \xi}(f, x, y))^\top \theta_k) a_k - y_i \right].$$

Denote $\theta_j^{OP1}(t)$ as the weight obtained using linear network with zero initialization, i.e. for each hidden node $j = 1, \dots, h$,

$$\begin{aligned} \theta_j^{OP1}(0) &= 0, \\ \theta_j^{OP1}(t+1) &= \theta_j^{OP1}(t) - \eta g_{j, \xi}^{OP1}(\theta_j^{OP1}(t), \epsilon). \end{aligned}$$

Also denote $\theta_j^{OP2}(t)$ as the weight obtained using nonlinear network with nonzero initialization, i.e.

$$\begin{aligned} \theta_j^{OP2}(0) &= 0, \\ \theta_j^{OP2}(t+1) &= \theta_j^{OP2}(t) - \eta g_{j, \xi}^{OP2}(\theta_j^{OP2}(t), \epsilon). \end{aligned}$$

For the original problem we consider (i.e. nonlinear network with vanishing initialization), define

$$\begin{aligned} \theta_j^{OP3}(0) &\sim N\left(0, \frac{1}{dh^{1+\delta'}} I\right), \\ \theta_j^{OP3}(t+1) &= \theta_j^{OP3}(t) - \eta g_{j, \xi}^{OP2}(\theta_j^{OP3}(t), \epsilon). \end{aligned}$$

Difference between (1) and (2) At $(t+1)$ th step, we have

$$\begin{aligned} \theta_j^{OP1}(t+1) - \theta_j^{OP2}(t+1) &= \theta_j^{OP1}(t) - \theta_j^{OP2}(t) - \eta g_{j, \xi}^{OP1}(\theta_j^{OP1}(t), \epsilon) + \eta g_{j, \xi}^{OP2}(\theta_j^{OP2}(t), \epsilon) \\ &= \theta_j^{OP1}(t) - \theta_j^{OP2}(t) - \eta [g_{j, \xi}^{OP1}(\theta_j^{OP1}(t), \epsilon) - g_{j, \xi}^{OP1}(\theta_j^{OP2}(t), \epsilon)] \\ &\quad - \eta [g_{j, \xi}^{OP1}(\theta_j^{OP2}(t), \epsilon) - g_{j, \xi}^{OP2}(\theta_j^{OP2}(t), \epsilon)]. \end{aligned}$$

Denote

$$\Delta_1(\theta, x, y) = A_{\epsilon, \xi}(f, x, y) - A_{\epsilon, \xi, OP1}(f, x, y),$$

$$\begin{aligned} \Delta_2(\theta, x, y) &= \frac{1}{\sqrt{h}} \sum_{j=1}^h \phi((x + A_{\epsilon, \xi}(f, x, y))^\top \theta_j) a_j \\ &\quad - \frac{1}{\sqrt{h}} \sum_{j=1}^h \phi(0)' x^\top \theta_j a_j + \phi'(0) A_{\epsilon, \xi, OP1}(f, x, y)^\top \theta_j a_j, \end{aligned}$$

and

$$\Delta_{3,j}(\theta, x, y) = \phi' \left((x + A_{\epsilon,\xi}(f, x, y))^\top \theta_j \right) - \phi'(0).$$

Assume all $\Delta_1(\theta, x, y)$, $\Delta_2(\theta, x, y)$, $\Delta_{3,j}(\theta, x, y)$ converges to zero for any j and $(x, y) \in \{(x_i, y_i)\}_{i=1,\dots,n}$ (we will later go back to validate this assumption), then

$$\begin{aligned} & g_{j,\xi}^{OP1}(\theta, \epsilon) - g_{j,\xi}^{OP2}(\theta, \epsilon) \\ = & -\frac{2}{n} \sum_{i=1}^n \frac{\phi' \left((x + A_{\epsilon,\xi}(f, x, y))^\top \theta_j \right) a_j}{\sqrt{h}} [x_i + A_{\epsilon,\xi}(f, x_i, y_i)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h \phi \left((x_i + A_{\epsilon,\xi}(f, x_i, y_i))^\top \theta_j \right) a_k - y_i \right] \\ & + \frac{2}{n} \sum_{i=1}^n \frac{\phi(0)' a_j}{\sqrt{h}} [x_i + A_{\epsilon,\xi,OP1}(f, x_i, y_i)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h (\phi(0)' x_i^\top \theta_k a_k + \phi'(0) A_{\epsilon,\xi,OP1}(f, x_i, y_i)^\top \theta_k a_k) - y_i \right] \\ = & -\frac{2}{n} \sum_{i=1}^n \frac{a_j \Delta_{3,j}(\theta, x_i, y_i)}{\sqrt{h}} [x_i + A_{\epsilon,\xi,OP1}(f, x_i, y_i)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h (\phi(0)' x_i^\top \theta_k a_k + \phi'(0) A_{\epsilon,\xi,OP1}(f, x_i, y_i)^\top \theta_k a_k) - y_i \right] \\ & - \frac{2}{n} \sum_{i=1}^n \frac{\phi(0)' a_j}{\sqrt{h}} \Delta_1(\theta, x_i, y_i) \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h (\phi(0)' x_i^\top \theta_k a_k + \phi'(0) A_{\epsilon,\xi,OP1}(f, x_i, y_i)^\top \theta_k a_k) - y_i \right] \\ & - \frac{2}{n} \sum_{i=1}^n \frac{\phi(0)' a_j}{\sqrt{h}} [x_i + A_{\epsilon,\xi,OP1}(f, x_i, y_i)] \Delta_2(\theta, x_i, y_i) + o, \end{aligned}$$

for some remainder term o , thus

$$\begin{aligned} & \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j,\xi}^{OP1}(\theta, \epsilon) - \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j,\xi}^{OP2}(\theta, \epsilon) \\ = & -\frac{2}{n} \sum_{i=1}^n \frac{\sum_{j=1}^h a_j^2 \Delta_{3,j}(\theta, x_i, y_i)}{h} [x_i + A_{\epsilon,\xi,OP1}(f, x_i, y_i)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h (\phi(0)' x_i^\top \theta_k a_k + \phi'(0) A_{\epsilon,\xi,OP1}(f, x_i, y_i)^\top \theta_k a_k) - y_i \right] \\ & - \frac{2}{n} \sum_{i=1}^n \frac{\phi(0)' \|a\|^2}{h} \Delta_1(\theta, x_i, y_i) \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h (\phi(0)' x_i^\top \theta_k a_k + \phi'(0) A_{\epsilon,\xi,OP1}(f, x_i, y_i)^\top \theta_k a_k) - y_i \right] \\ & - \frac{2}{n} \sum_{i=1}^n \frac{\phi(0)' \|a\|^2}{h} [x_i + A_{\epsilon,\xi,OP1}(f, x_i, y_i)] \Delta_2(\theta, x_i, y_i) + o. \end{aligned}$$

When $\|\theta^\top a / \sqrt{h}\| \leq B_1 v$, we have

$$\begin{aligned} & \left\| \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j,\xi}^{OP1}(\theta, \epsilon) - \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j,\xi}^{OP2}(\theta, \epsilon) \right\| \\ = & O \left(v \max_i \frac{|\sum_{j=1}^h a_j^2 \Delta_{3,j}(\theta, x_i, y_i)|}{h} \right) \\ & + \left\| \frac{2}{n} \sum_{i=1}^n \frac{\phi(0)' \|a\|^2}{h} \Delta_1(\theta, x_i, y_i) \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h (\phi(0)' x_i^\top \theta_k a_k + \phi'(0) A_{\epsilon,\xi,OP1}(f, x_i, y_i)^\top \theta_k a_k) - y_i \right] \right\| \\ & + \left\| \frac{2}{n} \sum_{i=1}^n \frac{\phi(0)' \|a\|^2}{h} [x_i + A_{\epsilon,\xi,OP1}(f, x_i, y_i)] \Delta_2(\theta, x_i, y_i) \right\| + o. \end{aligned}$$

We know that with probability at least $1 - e^{-g_1(\delta, \xi, d, n)}$,

$$\begin{aligned}
 & \left\| \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j,\xi}^{OP1}(\theta, \epsilon) - \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j,\xi}^{OP1}(\theta', \epsilon) \right\| \\
 & \leq \frac{\|a\|^2}{h} \|\nabla \widehat{R}_{L,\xi}(\theta^\top a / \sqrt{h}, \epsilon) - \nabla \widehat{R}_{L,\xi}((\theta')^\top a / \sqrt{h}, \epsilon)\| \quad (\text{from linear network to linear model}) \\
 & \leq 2\delta + \frac{\|a\|^2}{h} \|\nabla R_{L,\xi}(\theta^\top a / \sqrt{h}, \epsilon) - \nabla R_{L,\xi}((\theta')^\top a / \sqrt{h}, \epsilon)\| \\
 & \leq 2\delta + \frac{\|a\|^2}{h} \frac{Lv}{\sqrt{h}\xi} \|\theta^\top a - (\theta')^\top a\|.
 \end{aligned}$$

As a result, with probability at least $1 - e^{-g_1(\delta, \xi, d, n)}$,

$$\begin{aligned}
 & \left\| \frac{1}{\sqrt{h}} \theta^{OP1}(t+1)^\top a - \frac{1}{\sqrt{h}} \theta^{OP2}(t+1)^\top a \right\| \tag{11} \\
 & \leq \left\| \frac{1}{\sqrt{h}} \theta^{OP1}(t)^\top a - \frac{1}{\sqrt{h}} \theta^{OP2}(t)^\top a \right\| + \eta \left\| \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j,\xi}^{OP1}(\theta^{OP1}(t), \epsilon) - \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j,\xi}^{OP1}(\theta^{OP2}(t), \epsilon) \right\| \\
 & \quad + \eta \left\| \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j,\xi}^{OP1}(\theta^{OP2}(t), \epsilon) - \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j,\xi}^{OP2}(\theta^{OP2}(t), \epsilon) \right\|. \tag{12}
 \end{aligned}$$

As will be mentioned later, due to zero initialization, one can track $\theta_j^{OP1}(t)$ once given $\frac{1}{\sqrt{h}} \theta^{OP1}(t)^\top a / \sqrt{h}$. Similar property holds for $\theta_j^{OP2}(t)$. Therefore we bound $\frac{1}{\sqrt{h}} \theta^{OP1}(t+1)^\top a - \frac{1}{\sqrt{h}} \theta^{OP2}(t+1)^\top a$, which further enable us to bound $\|\theta_j^{OP1}(t) - \theta_j^{OP2}(t)\|$.

Linear network in (1) To further figure out the difference, we need some knowledge on $\theta^{OP1}(t)$, i.e. how $\theta^{OP1}(t)$ affects Δ_1 , Δ_2 , and Δ_3 .

Observe that $\frac{\phi'(0)}{\sqrt{h}} \theta^{OP1}(T)^\top a$ is just the coefficients of a linear model, so using zero initialization on $\theta^{OP1}(t)$, we have for any j, k in $1, \dots, h$,

$$\frac{\|\theta_j^{OP1}(t)\|}{\|\theta_k^{OP1}(t)\|} = \frac{|a_j|}{|a_k|}.$$

Now we study the sufficient conditions on (a, h) which enable us to bound the difference between $\theta^{OP1}(t)$ and $\theta^{OP2}(t)$. In general, we want the Taylor approximation of $\phi(x_i^\top \theta_j)$ valid, and $\left\| \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j,\xi}^{OP1}(\theta, \epsilon) - \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j,\xi}^{OP2}(\theta, \epsilon) \right\|$ goes to zero.

To validate the approximation of $\phi(x_i^\top \theta_k) = \phi(0) + \phi'(0)(x_i^\top \theta_k) + o$, we require $\max_{i,j} |\theta_j^{OP1}(t)^\top x_i| \rightarrow 0$. Since $\max_i \|x_i\| = O(\sqrt{d \log n})$ and $\|\theta_j^{OP1}(t)\| \leq \sqrt{h} |a_j| \|\theta_0\| / \|a\|^2$, a sufficient condition becomes

$$\sqrt{d \log n} \frac{\|a\|_\infty}{\|a\|^2} \|\theta_0\| \sqrt{h} \rightarrow 0.$$

(We will use a stronger version $(d \log n) \|a\|_\infty v \sqrt{h} / \|a\|^2 \rightarrow 0$.)

Next we study the bound of $\left\| \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j,\xi}^{OP1}(\theta, \epsilon) - \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j,\xi}^{OP2}(\theta, \epsilon) \right\|$. Recall that θ_0 is the true model (i.e. a vector, not a matrix). By the definition of $\Delta_{3,j}$, we have

$$\begin{aligned}
 |\Delta_{3,j}(\theta, x, y)| &= \left| \phi' \left((x + A_{\epsilon,\xi}(f, x, y))^\top \theta_j \right) - \phi'(0) \right| \\
 &= O(\phi''(0) |(x + A_{\epsilon,\xi}(f, x, y))^\top \theta_j|) \\
 &= O(\phi''(0) \|\theta_j\| (\|x\| + \epsilon)).
 \end{aligned}$$

As a result, since $\max_i \|x_i\| = O(\sqrt{d \log n})$ in probability, we have

$$\begin{aligned} \max_i \frac{|\sum_{j=1}^h a_j^2 \Delta_{3,j}(\theta^{OP1}(t), x_i, y_i)|}{\|a\|^2} &= O\left(\sqrt{d \log n} \max_j \|\theta^{OP1}(t)_j\|\right) \\ &= O\left(\sqrt{d \log n} \frac{\|a\|_\infty}{\|a\|^2} \sqrt{h} \|\theta_0\|\right). \end{aligned}$$

Similarly, for Δ_1 , if the sign of $\frac{\phi'(0)}{\sqrt{h}} x^\top \theta^\top a - y$ and $f_\theta(x) - y$ are the same, then

$$\|A_{\epsilon,\xi}(f, x, y) - A_{\epsilon,\xi,OP1}(f, x, y)\| = O\left(\frac{\|x\| \sum_{j=1}^h |a_j| \|\theta_j\|^2}{\sum_{j=1}^h a_j \theta_j}\right) = O\left(\sqrt{d \log n} \frac{\|a\|_\infty}{\|a\|^2} \sqrt{h} \|\theta_0\|\right).$$

If the signs are different, then since

$$\frac{\phi'(0)}{\sqrt{h}} x_i^\top \theta^\top a - f_\theta(x_i) = O\left(d \log n \sqrt{h} \|\theta_0\|^2 \frac{\|a\|_\infty}{\|a\|^2}\right),$$

there is a proportion of at most $O\left(d \log n \sqrt{h} \|\theta_0\|^2 \frac{\|a\|_\infty}{\|a\|^2} / v\right)$ of samples such that $A_{\epsilon,\xi}$ and $A_{\epsilon,\xi,OP1}$ are far away from each other. Denote b_i as the indicator that the sign of $\frac{\phi'(0)}{\sqrt{h}} x^\top \theta^\top a - y$ and $f_\theta(x) - y$ are the same, then

$$\begin{aligned} &\left\| \frac{2}{n} \sum_{i=1}^n \frac{\phi(0)' \|a\|^2}{h} \Delta_1(\theta, x_i, y_i) \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h (\phi(0)' x_i^\top \theta_k a_k + \phi'(0) A_{\epsilon,\xi,OP1}(f, x_i, y_i)^\top \theta_k a_k) - y_i \right] \right\| \\ &\leq \left\| \frac{2}{n} \sum_{i=1}^n \frac{\phi(0)' \|a\|^2}{h} \Delta_1(\theta, x_i, y_i) \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h (\phi(0)' x_i^\top \theta_k a_k + \phi'(0) A_{\epsilon,\xi,OP1}(f, x_i, y_i)^\top \theta_k a_k) - y_i \right] 1\{b_i = 1\} \right\| \\ &\quad + \left\| \frac{2}{n} \sum_{i=1}^n \frac{\phi(0)' \|a\|^2}{h} \Delta_1(\theta, x_i, y_i) \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h (\phi(0)' x_i^\top \theta_k a_k + \phi'(0) A_{\epsilon,\xi,OP1}(f, x_i, y_i)^\top \theta_k a_k) - y_i \right] 1\{b_i = 0\} \right\| \\ &= O\left((d \log n) \|a\|_\infty \frac{\|\theta_0\|^2}{\sqrt{h}}\right). \end{aligned}$$

Based on Δ_1 , for Δ_2 , when $\sqrt{dh \log n} \|a\|_\infty \|\theta_0\| / \|a\|^2 \rightarrow 0$, i.e. $\max_{i,j} |\theta_j^{OP1}(t)^\top x_i| \rightarrow 0$, we have

$$\begin{aligned} |\Delta_2(\theta, x, y)| &= \left| \frac{1}{\sqrt{h}} \sum_{j=1}^h \phi\left((x + A_{\epsilon,\xi}(f, x, y))^\top \theta_j\right) a_j - \frac{1}{\sqrt{h}} \sum_{j=1}^h \phi(0)' x^\top \theta_j a_j + \phi'(0) A_{\epsilon,\xi,OP1}(f, x, y)^\top \theta_j a_j \right| \\ &= \left| \frac{1}{\sqrt{h}} \sum_{j=1}^h \phi(0)' x^\top \theta_j a_j + \phi'(0) A_{\epsilon,\xi}(f, x, y)^\top \theta_j a_j + o - \frac{1}{\sqrt{h}} \sum_{j=1}^h \phi(0)' x^\top \theta_j a_j + \phi'(0) A_{\epsilon,\xi,OP1}(f, x, y)^\top \theta_j a_j \right| \\ &= \left| \frac{1}{\sqrt{h}} \sum_{j=1}^h \phi'(0) [A_{\epsilon,\xi}(f, x, y) - A_{\epsilon,\xi,OP1}(f, x, y)]^\top \theta_j a_j \right| \\ &\leq \frac{1}{\sqrt{h}} \sum_{j=1}^h \phi'(0) \Delta_1(\theta, x, y) |a_j| \|\theta_j\| + o \\ &\leq \frac{1}{\sqrt{h}} \sum_{j=1}^h \phi'(0) \Delta_1(\theta, x, y) \frac{\sqrt{h} a_j^2}{\|a\|^2} \|\theta_0\| + o \\ &= \phi'(0) \Delta_1(\theta, x, y) \|\theta_0\| + o. \end{aligned}$$

Thus with probability tending to 1,

$$\begin{aligned}
 & \left\| \frac{2}{n} \sum_{i=1}^n \frac{\phi(0)' \|a\|^2}{h} [x_i + A_{\epsilon, \xi, OP1}(f, x_i, y_i)] \Delta_2(\theta, x_i, y_i) \right\| \\
 & \leq 2 \frac{\phi'(0) \|a\|^2}{h} (\max_i \|x_i\| + \epsilon) \phi'(0) \|\theta_0\| \left(\frac{1}{n} \sum_{i=1}^n |\Delta_1(f, x_i, y_i)| \right) + o \\
 & = O \left((d \log n) \|a\|_\infty \frac{\|\theta_0\|^2}{\sqrt{h}} \right).
 \end{aligned}$$

To summarize, with probability tending to 1,

$$\left\| \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j, \xi}^{OP1}(\theta, \epsilon) - \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j, \xi}^{OP2}(\theta, \epsilon) \right\| = O \left((d \log n) \|a\|_\infty \frac{\|\theta_0\|^2}{\sqrt{h}} \right)$$

when $(d \log n) \|a\|_\infty v^2 \sqrt{h} / \|a\|^2 \rightarrow 0$.

Return to the difference between (1) and (2) Now we use the property that $\theta^{OP2}(0) = 0$. Similar as θ^{OP1} , if $a_j a_k > 0$, then

$$\frac{\|\theta_j^{OP2}(t)\|}{\|\theta_k^{OP2}(t)\|} \approx \frac{|a_j|}{|a_k|}.$$

Therefore, we define $\theta_+^{OP2}(t) = \sum_{a_j > 0} a_j \theta_j^{OP2}(t)$, and similarly define $\theta_-^{OP2}(t)$, $\theta_+^{OP1}(t)$, $\theta_-^{OP1}(t)$. One can obtain similar result for (11) to (12) when considering $\theta_+^{OP1}(t) - \theta_+^{OP2}(t)$ and $\theta_-^{OP1}(t) - \theta_-^{OP2}(t)$.

When $\eta(\delta + (d \log n) \|a\|_\infty v^2 / \sqrt{h})(1 + \|a\|^2 L v \eta / (h \xi))^T \rightarrow 0$, one can use induction to bound $\|\theta^{OP1}(T)^\top a / \sqrt{h} - \theta^{OP2}(T)^\top a / \sqrt{h}\| / \|\theta_0\|$: recall that

$$\begin{aligned}
 \Delta_{t+1} &:= \left\| \frac{1}{\sqrt{h}} \theta^{OP1}(t+1)^\top a - \frac{1}{\sqrt{h}} \theta^{OP2}(t+1)^\top a \right\| \\
 &\leq \left\| \frac{1}{\sqrt{h}} \theta^{OP1}(t)^\top a - \frac{1}{\sqrt{h}} \theta^{OP2}(t)^\top a \right\| + \eta \left\| \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j, \xi}^{OP1}(\theta^{OP1}(t), \epsilon) - \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j, \xi}^{OP1}(\theta^{OP2}(t), \epsilon) \right\| \\
 &\quad + \eta \left\| \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j, \xi}^{OP1}(\theta^{OP2}(t), \epsilon) - \frac{1}{\sqrt{h}} \sum_{i=1}^h a_j g_{j, \xi}^{OP2}(\theta^{OP2}(t), \epsilon) \right\| \\
 &:= \Delta_t + \eta \Delta_{1,t} + \eta \Delta_{2,t}(\theta^{OP1}(t)).
 \end{aligned}$$

When studying the linear network $\theta^{OP1}(t)$, we argue that $\Delta_{2,t}(\theta^{OP1}(t)) \rightarrow 0$ under proper choice of (a, h) . If $\|\theta_j^{OP2}(t) - \theta_j^{OP1}(t)\| / \|\theta_j^{OP1}(\infty)\| \rightarrow 0$, then using $\theta^{OP2}(t)$ we can also obtain $\Delta_{2,t}(\theta^{OP2}(t)) \rightarrow 0$. In addition, $\Delta_{1,t}$ is a linear function of Δ_t plus some error δ . Furthermore, after obtaining Δ_{t+1} , one can further figure out $\|\theta_j^{OP2}(t+1) - \theta_j^{OP1}(t+1)\|$. To conclude, using induction, we have with probability $1 - e^{-g_1(\delta, \xi, d, n)}$, $\|\theta^{OP1}(T)^\top a / \sqrt{h} - \theta^{OP2}(T)^\top a / \sqrt{h}\| / \|\theta_0\|$ converges to zero.

Difference between (2) and (3) Denote $A_{2,i}$ and $A_{3,i}$ as the attack for i th data w.r.t. $\theta^{OP2}(t)$ and $\theta^{OP3}(t)$, then we have

$$\begin{aligned}
 & g_{j,\xi}^{OP2}(\theta^{OP2}(t), \epsilon) - g_{j,\xi}^{OP2}(\theta^{OP3}(t), \epsilon) \\
 = & \frac{2}{n} \sum_{i=1}^n \frac{\phi' \left((x + A_{2,i})^\top \theta_j^{OP2}(t) \right) a_j}{\sqrt{h}} [x_i + A_{2,i}(t)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h \phi \left((x_i + A_{2,i}(t))^\top \theta_k^{OP2}(t) \right) a_k - y_i \right] \\
 & - \frac{2}{n} \sum_{i=1}^n \frac{\phi' \left((x + A_{3,i})^\top \theta_j^{OP3}(t) \right) a_j}{\sqrt{h}} [x_i + A_{3,i}(t)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h \phi \left((x_i + A_{3,i}(t))^\top \theta_k^{OP3}(t) \right) a_k - y_i \right] \\
 = & \frac{2}{n} \sum_{i=1}^n \frac{\phi' \left((x + A_{2,i})^\top \theta_j^{OP2}(t) \right) a_j}{\sqrt{h}} [x_i + A_{2,i}(t)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h \phi \left((x_i + A_{2,i}(t))^\top \theta_k^{OP2}(t) \right) a_k - y_i \right] \\
 & - \frac{2}{n} \sum_{i=1}^n \frac{\phi' \left((x + A_{2,i})^\top \theta_j^{OP3}(t) \right) a_j}{\sqrt{h}} [x_i + A_{2,i}(t)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h \phi \left((x_i + A_{2,i}(t))^\top \theta_k^{OP3}(t) \right) a_k - y_i \right] \\
 & + \frac{2}{n} \sum_{i=1}^n \frac{\phi' \left((x + A_{2,i})^\top \theta_j^{OP3}(t) \right) a_j}{\sqrt{h}} [x_i + A_{2,i}(t)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h \phi \left((x_i + A_{2,i}(t))^\top \theta_k^{OP3}(t) \right) a_k - y_i \right] \\
 & - \frac{2}{n} \sum_{i=1}^n \frac{\phi' \left((x + A_{3,i})^\top \theta_j^{OP3}(t) \right) a_j}{\sqrt{h}} [x_i + A_{3,i}(t)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h \phi \left((x_i + A_{3,i}(t))^\top \theta_k^{OP3}(t) \right) a_k - y_i \right] \\
 := & (B_{1,j} - B_{2,j}) + (B_{2,j} - B_{3,j}).
 \end{aligned}$$

The term $B_{1,j} - B_{2,j}$ follows the same as C.8.3 in Ba et al. (2020). For $B_{2,j} - B_{3,j}$, we know that

$$\|A_{2,i} - A_{3,i}\| = O \left(v \left\| \frac{\sum_{j=1}^h a_j (\phi'(x^\top \theta_j^{OP2}(t)) \theta_j^{OP2}(t) - \phi'(x^\top \theta_j^{OP3}(t)) \theta_j^{OP3}(t))}{\sqrt{h}\xi} \right\| \right).$$

Since

$$\|\theta^{OP3}(T)a - \theta^{OP2}(T)a\| \leq \|a\|_\infty \|\theta^{OP3}(T) - \theta^{OP2}(T)\|_F,$$

when $f_{\theta^{OP2}(t)}(x_i) - y_i$ and $f_{\theta^{OP3}(t)}(x_i) - y_i$ have the same sign, assuming $|x_i \theta_j^{OP3}(t)| \rightarrow 0$,

$$\begin{aligned}
 \|A_{2,i} - A_{3,i}\| &= O \left(v \left\| \frac{\sum_{j=1}^h a_j (\phi'(0)(\theta_j^{OP2}(t) - \theta_j^{OP3}(t)) + o)}{\sqrt{h}\xi} \right\| \right) \\
 &= O \left(\frac{v \|a\|_\infty \|\theta^{OP3}(t) - \theta^{OP2}(t)\|_F}{\sqrt{h}\xi} \right),
 \end{aligned}$$

and there is a proportion at most $O(\sqrt{d \log n} \|(\theta^{OP2}(t) - \theta^{OP3}(t))^\top a / \sqrt{h}\| / v)$ samples whose $f_{\theta^{OP2}(t)}(x_i) - y_i$ and $f_{\theta^{OP3}(t)}(x_i) - y_i$ have different signs. As a result, denote b'_i as the indicator that $f_{\theta^{OP2}(t)}(x_i) - y_i$ and

$f_{\theta^{OP3}(t)}(x_i) - y_i$ have the same sign, then assuming $|x_i \theta_j^{OP3}(t)| \rightarrow 0$,

$$\begin{aligned}
 & \|g_{j,\xi}^{OP2}(\theta^{OP2}(t), \epsilon) - g_{j,\xi}^{OP2}(\theta^{OP3}(t), \epsilon)\| \\
 \leq & \|B_{1,j} - B_{2,j}\| \\
 & + \left\| \frac{2}{n} \sum_{i=1}^n \frac{\phi'((x + A_{2,i})^\top \theta_j^{OP3}(t)) a_j}{\sqrt{h}} [x_i + A_{2,i}(t)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h \phi((x + A_{2,i}(t))^\top \theta_k^{OP3}(t)) a_k - y_i \right] \right. \\
 & \quad \left. - \frac{2}{n} \sum_{i=1}^n \frac{\phi'((x + A_{3,i})^\top \theta_j^{OP3}(t)) a_j}{\sqrt{h}} [x_i + A_{3,i}(t)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h \phi((x + A_{3,i}(t))^\top \theta_k^{OP3}(t)) a_k - y_i \right] \right\| \\
 = & \|B_{1,j} - B_{2,j}\| \\
 & + \left\| \frac{2}{n} \sum_{i=1}^n \frac{\phi'(0) a_j}{\sqrt{h}} [x_i + A_{2,i}(t)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h \phi'(0) (x_i + A_{2,i})^\top \theta_k^{OP3}(t) a_k - y_i \right] \right. \\
 & \quad \left. - \frac{2}{n} \sum_{i=1}^n \frac{\phi'(0) a_j}{\sqrt{h}} [x_i + A_{3,i}(t)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h \phi'(0) (x_i + A_{3,i})^\top \theta_k^{OP3}(t) a_k - y_i \right] \right\| + o \\
 \leq & \|B_{1,j} - B_{2,j}\| \\
 & + \left\| \frac{2}{n} \sum_{i=1}^n \frac{\phi'(0) a_j}{\sqrt{h}} [A_{3,i}(t) - A_{2,i}(t)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h \phi'(0) (x_i + A_{2,i})^\top \theta_k^{OP2}(t) a_k - y_i \right] \right\| \\
 & + \left\| \frac{2}{n} \sum_{i=1}^n \frac{\phi'(0) a_j}{\sqrt{h}} [A_{3,i}(t)] \left[\frac{1}{\sqrt{h}} \sum_{k=1}^h \phi'(0) (A_{3,i} - A_{2,i})^\top \theta_k^{OP2}(t) a_k - y_i \right] \right\| + o \\
 = & \|B_{1,j} - B_{2,j}\| + O \left(\frac{|a_j| v L \|\theta_0\| \|a\|_\infty \|\theta^{OP3}(t) - \theta^{OP2}(t)\|_F}{\sqrt{h} \xi} \right).
 \end{aligned}$$

Taking $\xi \gg \sqrt{d \log n} / (\sqrt{h} v)$, we have

$$\begin{aligned}
 & \|\theta^{OP3}(t+1) - \theta^{OP2}(t+1)\|_F \\
 \leq & \|\theta^{OP3}(t) - \theta^{OP2}(t)\|_F + \eta \sqrt{\sum_{j=1}^h \left\| g_{j,\xi}^{OP2}(\theta^{OP3}(t), \epsilon) - g_{j,\xi}^{OP2}(\theta^{OP2}(t), \epsilon) \right\|^2} \\
 = & \|\theta^{OP3}(t) - \theta^{OP2}(t)\|_F O \left(1 + \frac{v L \eta \|\theta_0\| \|a\|_\infty \|a\|^2}{h \xi} + L v \eta \right).
 \end{aligned}$$

Therefore,

$$\begin{aligned}
 \|\theta^{OP3}(T) - \theta^{OP2}(T)\|_F &= O \left(\|\theta^{OP3}(0)\|_F \left(1 + \frac{v L \eta \|\theta_0\| \|a\|_\infty \|a\|^2}{h \xi} + L v \eta \right)^T \right) \\
 &= O \left(\frac{1}{h^{\delta'/2}} \left(1 + \frac{v L \eta \|\theta_0\| \|a\|_\infty \|a\|^2}{h \xi} + L v \eta \right)^T \right).
 \end{aligned}$$

As a result, we require $\|\theta^{OP3}(T) - \theta^{OP2}(T)\|_F \sqrt{d \log n} \rightarrow 0$ so that with probability tending to 1, for all nodes, $|x_i^\top \theta_j^{OP3}(T)| \rightarrow 0$.

Deciding proper choice of ξ In Theorem 1, since $\frac{\phi'(0)}{\sqrt{h}} \theta^{OP1}(t) a$ is a linear model, denote η_{linear} as the learning rate for linear model, then the corresponding η in linear network is $\eta = \eta_{linear} h / (\|a\|^2 \phi'(0)^2)$. We require $T \eta_{linear} \rightarrow \infty$, $\xi / v^2 \rightarrow 0$, $\eta_{linear} = \xi / (2 L v^2)$, and $\delta v / \xi \rightarrow 0$ in Theorem 2.

Now we list all the assumptions we made on (ξ, η, T, δ) during derivation.

- Difference between $\theta^{OP1}(t)$ and $\theta^{OP2}(t)$: $e^{-g_1(\delta, \xi, d, n)} \rightarrow 0$, $(d \log n) \|a\|_\infty v / \sqrt{h} \rightarrow 0$, $(d \log n) \frac{\|a\|_\infty}{\|a\|^2} \|\theta_0\|^2 / v \rightarrow 0$ (the gradient of linear network is in $\Theta(v)$ when it is far from θ^* , thus we divide v here), and $\eta(\delta + (d \log n) \|a\|_\infty v^2 / \sqrt{h})(1 + \|a\|^2 L v^2 \eta / (h \xi) + L)^T / v \rightarrow 0$.

- Difference between $\theta^{OP2}(t)$ and $\theta^{OP3}(t)$: $\frac{\sqrt{d \log n}}{h^{\delta'/2}} \left(1 + \frac{vL\eta\|\theta_0\|\|a\|_\infty\|a\|^2}{h\xi} + Lv\eta \right)^T / v \rightarrow 0$.

So we take $\delta = v\sqrt{(d^2 \log n)/n}$, and

$$\xi/v^2 = -\log \log n / \log \left(\sqrt{\frac{d^2 \log n}{n}} \vee (d \log n) \frac{\|a\|_\infty}{\sqrt{h}} \right).$$

Finally, we adapt the above proof for the true attack. From Taylor expansion, we have when $x^\top \theta_j \rightarrow 0$,

$$\phi(z^\top \theta_j) \approx \phi(x^\top \theta_j \phi'(0)(z - x)^\top \theta_j + O(\|\theta_j\|^2)),$$

which implies that the fast gradient attack used in the above proof converges to the true attack. Recall that we require $\Delta_3 \rightarrow 0$, thus for $\theta^{OP2}(t)$, the arguments also holds when replacing fast gradient attack by true attack. The same argument for $\theta^{OP3}(t)$ as well. \square

Proof of Theorem 6. The proof is similar as the one for Theorem 4 through replacing θ_0 with $\theta(\mathbf{y})$. Based on Lemma 1, we know that $\|\theta(\mathbf{y})\|/v = O_p(n/d)$, thus in the step “Linear Network in (1)”, to ensure $\|f_{\theta_\xi^{(t)}}(x_i) - \mathbf{y}\|$ decreases in each iteration, we require $\sqrt{d \log n} \sqrt{n \log n} \|a\|_\infty v / \sqrt{h} \rightarrow 0$. On the other hand, for the step “Difference between (2) and (3)”, we take a and δ' such that $\sqrt{d \log n} \|\theta^{OP3}(t)a - \theta^{OP2}(t)a\| = o(1)$. \square

Proof of Theorem 4 and 7. Assume $X = [X_0, -X_0]^\top$. Since $\|a^+\| = \|a^-\|$, one can check that all Δ_1, Δ_2 , and Δ_3 are zeros. Take $\phi'(0) = 1$. Denote $S^+(\theta)$ as $\{j \mid f(x + A_{\epsilon, \xi}(f, x_i, y_i)) > 0\}$, and similarly denote $S^-(\theta)$. Then one can observe that, using zero initialization,

$$\frac{1}{\sqrt{h}} \sum_{j \in S^+(\theta)} a_j g_{j, \xi}^{OP2}(\theta, \epsilon) + \frac{1}{\sqrt{h}} \sum_{j \in S^-(\theta)} a_j g_{j, \xi}^{OP2}(\theta, \epsilon) = 0,$$

and hence ReLU-activated neural networks with zero initialization perform the same as linear networks when the learning rate for ReLU-activated neural is twice as the one for linear networks. \square

D Proofs for high-dimensional dense model

Proof of Lemma 1. Denote $\theta^+ = \theta(\mathbf{y})$. When v is a constant, $\|\theta^+\|$ satisfies

$$\begin{aligned} \|\theta^+\|^2 &= \mathbf{y}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{y} \\ &= \theta_0^\top \mathbf{X}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{X} \theta_0 + \epsilon^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \epsilon \\ &\rightarrow \theta_0^\top \mathbf{X}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{X} \theta_0 + \sigma^2 \text{tr}((\mathbf{X}\mathbf{X}^\top)^{-1}). \end{aligned}$$

Furthermore, similar as Belkin et al. (2019), we have

$$(\theta_0^\top \Sigma^{1/2}) \Sigma^{-1/2} \mathbf{X}^\top (\mathbf{X} \Sigma^{-1} \mathbf{X}^\top)^{-1} \mathbf{X} \Sigma^{-1/2} (\Sigma^{1/2} \theta_0) \rightarrow \frac{n}{d} \|\theta_0 \Sigma^{1/2}\|^2.$$

Therefore, since $\mathbf{X} \Sigma^{-1} \mathbf{X}^\top - \frac{2}{\lambda_{\min}} \mathbf{X} \mathbf{X}^\top$ is negative definite for the smallest eigenvalue λ_{\min} of Σ and , we have

$$\begin{aligned} \theta_0^\top \mathbf{X}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{X} \theta_0 &\leq \frac{2}{\lambda_{\min}} (\theta_0^\top \Sigma^{1/2}) \Sigma^{-1/2} \mathbf{X}^\top (\mathbf{X} \Sigma^{-1} \mathbf{X}^\top)^{-1} \mathbf{X} \Sigma^{-1/2} (\Sigma^{1/2} \theta_0) \\ &\rightarrow \frac{n}{d} \|\theta_0 \Sigma^{1/2}\|^2. \end{aligned}$$

Finally, if $v \rightarrow \infty$, we obtain that $\theta_0^\top \mathbf{X}^\top (\mathbf{X}\mathbf{X}^\top)^{-1} \mathbf{X} \theta_0 / v^2 \rightarrow 0$. \square

Lemma 3. When $(\log n)\sqrt{n/d} \rightarrow 0$, with probability tending to 1, the smallest eigenvalue of $\mathbf{X}\mathbf{X}^\top$ is in $\Theta(d)$.

Proof of Lemma 3. Assume $\Sigma = I$ for simplicity. Denote $b \in (0, 1)$. Since $(\log n)\sqrt{n/d} \rightarrow 0$, we append $(bd - n)$ i.i.d samples of x after \mathbf{X} and denote the new data matrix as \mathbf{Z} . Based on [Bai and Yin \(2008\)](#), the smallest eigenvalue of $\mathbf{Z}\mathbf{Z}^\top/d$ converges to $(1 - \sqrt{b})^2$, and the largest eigenvalue converges to $(1 + \sqrt{b})^2$. Since $\lambda_{\min}(\mathbf{X}\mathbf{X}^\top/d) \geq \lambda_{\min}(\mathbf{Z}\mathbf{Z}^\top/d)$, and $\lambda_{\min}(\mathbf{X}\mathbf{X}^\top/d) \leq \lambda_{\max}(\mathbf{X}\mathbf{X}^\top/d) \leq \lambda_{\max}(\mathbf{Z}\mathbf{Z}^\top/d)$, we conclude that $\lambda_{\min} = \Theta(d)$ in probability. \square

Proof of Theorem 5. Assume v is constant for simplicity and we take $\xi = 0$. Denote $\theta^{OP1}(t)$ and $\theta^{OP2}(t)$ satisfy

$$\begin{aligned}\theta^{OP1}(0) &= 0, \\ \theta^{OP1}(t+1) &= \theta^{OP1}(t) - \eta \nabla \hat{R}_\xi(\theta^{OP1}(t), 0),\end{aligned}$$

and

$$\begin{aligned}\theta^{OP2}(0) &= 0, \\ \theta^{OP2}(t+1) &= \theta^{OP2}(t) - \eta \nabla \hat{R}_\xi(\theta^{OP2}(t), \epsilon).\end{aligned}$$

Proof sketch The proof idea is that, since [Ba et al. \(2020\)](#) has studied $\theta^{OP1}(t)$, we want know how θ^{OP2} is closed to θ^{OP1} . We show that when $\|\theta\|$ satisfies some certain condition \mathbf{M} , $\hat{R}_\xi(\theta, 0)$ will get decreased in the next update in adversarial training. As a result, when θ^{OP2} satisfies these condition \mathbf{M} , $\hat{R}_\xi(\theta^{OP2}(t), 0)$ is always decreasing. On the hand, when θ^{OP2} satisfies condition \mathbf{M} , one can also show that $\|\nabla \hat{R}_\xi(\theta^{OP2}(t), \epsilon) - \nabla \hat{R}_\xi(\theta^{OP2}(t), 0)\| = o(\|\nabla \hat{R}_\xi(\theta^{OP2}(t), 0)\|)$. Since zero initialization satisfies condition \mathbf{M} and clean training satisfies condition \mathbf{M} , one can use induction to show that adversarial training is dominated by clean training. Finally, we obtain $\hat{R}_\xi(\theta^{OP2}(T), 0) \rightarrow 0$ and $\|\theta^{OP2}(T)\| \rightarrow 0$ if η and T are chosen properly.

Condition M $\|\theta\| = O(\sqrt{n/d})$, $\|\mathbf{z}\| = O(\sqrt{n})$, and $\|\mathbf{z} - \mathbf{y}\| > \sqrt{n}/\sqrt{\log n}$.

Now we begin our proof. For $\theta^{OP2}(t)$, the change on $\hat{R}_\xi(\theta, 0)$ is

$$\hat{R}_\xi(\theta^{OP2}(t+1), 0) = \hat{R}_\xi(\theta^{OP2}(t), 0) - \frac{2\eta}{n} \nabla \hat{R}_\xi(\theta^{OP2}(t), \epsilon)^\top \mathbf{X}^\top (\mathbf{X}\theta^{OP2}(t) - \mathbf{y}) + \eta^2 \|\mathbf{X} \nabla \hat{R}_\xi(\theta^{OP2}(t), \epsilon)\|^2.$$

Rewrite $\mathbf{X}\theta = \mathbf{z}$ for simplicity, then

$$\mathbf{X} \nabla \hat{R}_\xi(\theta, \epsilon) = \frac{2\mathbf{X}\mathbf{X}^\top(\mathbf{z} - \mathbf{y})}{n} + \frac{2\epsilon}{n} \|\theta\| \mathbf{X}\mathbf{X}^\top \text{sgn}(\mathbf{z} - \mathbf{y}) + \frac{2\epsilon}{n} \frac{\mathbf{z}}{\|\theta\|} \|\mathbf{z} - \mathbf{y}\|_1 + 2\epsilon^2 \mathbf{z}.$$

Since with probability tending to 1, $\|\mathbf{X}\mathbf{X}^\top \text{sgn}(\mathbf{z} - \mathbf{y})\| = O(d\sqrt{n})$, thus when $\|\mathbf{X}\mathbf{X}^\top(\mathbf{z} - \mathbf{y})\| = \Omega(d\|\mathbf{z} - \mathbf{y}\|)$, with probability tending to 1, we have

$$\left\| \frac{2\epsilon}{n} \frac{\mathbf{z}}{\|\theta\|} \|\mathbf{z} - \mathbf{y}\|_1 \right\| \leq \frac{2\epsilon}{\sqrt{n}} \frac{\|\mathbf{X}\theta\|}{\|\theta\|} \|\mathbf{z} - \mathbf{y}\|_2 = o\left(\left\| \frac{\mathbf{X}\mathbf{X}^\top(\mathbf{z} - \mathbf{y})}{n} \right\|\right),$$

and

$$\left\| \frac{2\epsilon}{n} \|\theta\| \mathbf{X}\mathbf{X}^\top \text{sgn}(\mathbf{z} - \mathbf{y}) + 2\epsilon^2 \mathbf{z} \right\| = \left\| \frac{\mathbf{X}\mathbf{X}^\top(\mathbf{z} - \mathbf{y})}{n} \right\| O\left(\left(\|\theta\| + \frac{\sqrt{n}\|\mathbf{z}\|}{d}\right) \frac{\sqrt{n}}{\|\mathbf{z} - \mathbf{y}\|}\right).$$

The statement $\|\mathbf{X}\mathbf{X}^\top(\mathbf{z} - \mathbf{y})\| = \Omega(d\|\mathbf{z} - \mathbf{y}\|)$ for any \mathbf{z} holds based on [Lemma 3](#).

Consequently, under condition \mathbf{M} ,

$$\hat{R}_\xi(\theta^{OP2}(t+1), 0) = \hat{R}_\xi(\theta^{OP2}(t), 0) - \frac{4\eta}{n^2} (\mathbf{z} - \mathbf{y})^\top \mathbf{X}\mathbf{X}^\top (\mathbf{z} - \mathbf{y}) + o,$$

which implies that the decrease in $\hat{R}_\xi(\theta^{OP2}(t), 0)$ is almost the same as $\hat{R}_\xi(\theta^{OP1}(t), 0)$.

Now our aim becomes to figure out $\|\theta^{OP2}(t)\|$ when $\|\mathbf{z} - \mathbf{y}\| > \sqrt{n}/\sqrt{\log n}$.

For $\theta^{OP1}(t)$, i.e. standard training, when η is small enough such that the largest eigenvalue of $\eta \mathbf{X}^\top \mathbf{X}$ is smaller than 1, then

$$\begin{aligned}\theta^{OP1}(t) &= \mathbf{X}^\top (I - (I - \eta \mathbf{X} \mathbf{X}^\top / n)^t) (\mathbf{X} \mathbf{X}^\top)^{-1} \mathbf{y}, \\ \mathbf{X} \theta^{OP1}(t) &= \mathbf{y} - (I - \eta \mathbf{X} \mathbf{X}^\top / n)^t \mathbf{y},\end{aligned}$$

which means that $\|\mathbf{X} \theta^{OP1}(t) - \mathbf{y}\|$ monotonly decreases in t . Solving $\|(I - \eta \mathbf{X} \mathbf{X}^\top / n)^T \mathbf{y}\|/\sqrt{n} = 1/\sqrt{\log n}$, we obtain $\eta T = O(\frac{n}{d} \log \log n)$.

Observe that

$$\begin{aligned}& \|\theta^{OP1}(t+1) - \theta^{OP2}(t+1)\| \\ &= \left\| (\theta^{OP1}(t) - \theta^{OP2}(t)) - \eta \left(\nabla \hat{R}_\xi(\theta^{OP1}(t), 0) - \nabla \hat{R}_\xi(\theta^{OP2}(t), 0) + \nabla \hat{R}_\xi(\theta^{OP2}(t), 0) - \nabla \hat{R}_\xi(\theta^{OP2}(t), \epsilon) \right) \right\| \\ &\leq \|\theta^{OP1}(t) - \theta^{OP2}(t)\| \left(1 + O\left(\frac{\eta d}{n}\right) \right) + \eta \left\| \nabla \hat{R}_\xi(\theta^{OP2}(t), 0) - \nabla \hat{R}_\xi(\theta^{OP2}(t), \epsilon) \right\|.\end{aligned}$$

When $\theta = \theta^{OP1}(t)$ for some $t > 0$ and $\|\theta\| = O(\sqrt{n/d})$. Since $\mathbf{z} - \mathbf{y} = (I - \eta \mathbf{X} \mathbf{X}^\top / n)^t \mathbf{y}$, $\|\mathbf{X}^\top (\mathbf{z} - \mathbf{y})\| \rightarrow c(d, n, t) \sqrt{d} \|\mathbf{z} - \mathbf{y}\|$ for some function $c(d, n, t)$ which is finite and bounded away from zero. As a result, with probability tending to 1, for the difference between $\nabla \hat{R}_\xi(\theta, \epsilon)$ and $\nabla \hat{R}_\xi(\theta, 0)$, it becomes

$$\begin{aligned}\left\| \nabla \hat{R}_\xi(\theta, \epsilon) - \nabla \hat{R}_\xi(\theta, 0) \right\| &\leq \left\| \frac{2\epsilon}{n} \|\theta\| \mathbf{X}^\top \text{sgn}(\mathbf{z} - \mathbf{y}) \right\| + \left\| \frac{2\epsilon}{n} \frac{\theta}{\|\theta\|} \|\mathbf{z} - \mathbf{y}\|_1 \right\| + \|2\epsilon^2 \theta\| \\ &= \left\| \nabla \hat{R}_\xi(\theta, 0) \right\| O\left(\frac{n/\sqrt{d}}{\|\mathbf{z} - \mathbf{y}\|} + \sqrt{\frac{n}{d}} + \frac{n^{3/2}/d}{\|\mathbf{z} - \mathbf{y}\|} \right) \\ &= \left\| \nabla \hat{R}_\xi(\theta, 0) \right\| O\left(\frac{n/\sqrt{d}}{\|\mathbf{z} - \mathbf{y}\|} \right).\end{aligned}\tag{13}$$

When $\|\theta - \theta^{OP1}(t)\| = o(\|\theta^{OP1}(t)\|)$, a similar result can be obtained.

As a result, one can use induction to show that $\|\theta^{OP2}(T)\| = O(\|\theta^{OP1}(T)\| + \eta T) = o(\sqrt{n/d})$, $\|\mathbf{X} \theta^{OP2}(t)\| = O(\sqrt{n})$ and $\|\mathbf{X} \theta^{OP2}(t) - \mathbf{y}\|$ decreases in t , i.e. condition **M** holds when taking $\eta T = O(\frac{n}{d} \log \log n)$. So the final conclusion holds. \square

E Proofs for high-dimensional sparse model

Proof of Theorem 9. Assume $v = O(1)$ first. For simplicity we assume $\xi = 0$. Denote

$$\frac{1}{n} \sum_{i=1}^n l(f_\theta(x_i + A_\epsilon(f_\theta, x_i, y_i)), y_i) + \lambda \|\theta\|_1 := \frac{1}{n} \sum_{i=1}^n l_\epsilon(\theta, x_i, y_i) + \lambda \|\theta\|_1.$$

Since $\hat{\theta}$ minimizes the empirical penalized loss function, take $\Delta = \hat{\theta} - \theta^*$, we have

$$\frac{1}{n} \sum_{i=1}^n l_\epsilon(\hat{\theta}, x_i, y_i) - \frac{1}{n} \sum_{i=1}^n l_\epsilon(\theta^*, x_i, y_i) \leq \lambda \|\theta^*\|_1 - \lambda \|\hat{\theta}\|_1 \leq \lambda (\|\Delta_S\|_1 - \|\Delta_{S^c}\|_1).$$

Moreover, the structure of l implies that it is a convex function, thus

$$\Delta^\top \frac{1}{n} \sum_{i=1}^n l'_\epsilon(\theta^*, x_i, y_i) \leq \lambda (\|\Delta_S\|_1 - \|\Delta_{S^c}\|_1).$$

Further,

$$\lambda(\|\Delta_S\|_1 - \|\Delta_{S^c}\|_1) \geq -\left|\Delta^\top \frac{1}{n} \sum_{i=1}^n l'_\epsilon(\theta^*, x_i, y_i)\right| \geq -\|\Delta\|_1 \left\|\frac{1}{n} \sum_{i=1}^n l'_\epsilon(\theta^*, x_i, y_i)\right\|_\infty.$$

Since θ^* is fixed, we can figure out that with probability tending to 1,

$$\left\|\frac{1}{n} \sum_{i=1}^n l'_\epsilon(\theta^*, x_i, y_i)\right\|_\infty \leq c_1 \sqrt{\frac{s \log d}{n}}$$

for some constant $c_1 > 0$. Consequently, as our choice of M is large enough, with probability tending to 1,

$$\|\Delta_S\|_1 \geq \frac{\lambda - c_1 \sqrt{\frac{s \log d}{n}}}{\lambda + c_1 \sqrt{\frac{s \log d}{n}}} \|\Delta_{S^c}\|_1.$$

As a result, from Lemma 4, we know that with probability tending to 1, for some constant $c_2 > 0$,

$$\lambda \|\Delta_S\|_1 \geq \lambda \|\Delta_S\|_1 - \lambda \|\Delta_{S^c}\|_1 \geq \Delta^\top \frac{1}{n} \sum_{i=1}^n l'_\epsilon(\theta^*, x_i, y_i) - c_2 \|\Delta\|_1 \sqrt{\frac{s \log d}{n}} + \epsilon^2 \left(1 - \frac{2}{\pi}\right) \|\Delta\|^2.$$

Therefore,

$$\begin{aligned} \lambda \|\Delta_S\|_1 + \left(c_1 \sqrt{\frac{s \log d}{n}} + c_2 \sqrt{\frac{s \log d}{n}}\right) \|\Delta\|_1 &\geq \epsilon^2 \left(1 - \frac{2}{\pi}\right) \|\Delta\|^2 \\ &\geq \epsilon^2 \left(1 - \frac{2}{\pi}\right) \frac{\|\Delta_S\|_1^2}{s}, \end{aligned}$$

hence with probability tending to 1,

$$\|\Delta\|_1 = O(\lambda), \quad \|\Delta\|_2 = O(\lambda).$$

Finally, using Lemma 5, we have

$$\begin{aligned} &l_\epsilon(\hat{\theta}, x_i, y_i) - l_\epsilon(\theta^*, x_i, y_i) \\ &\leq \Delta^\top \frac{\partial l_\epsilon(\theta^*, x_i, y_i)}{\partial \theta^*} + (\Delta^\top x_i)^2 + \epsilon^2 \|\Delta\|^2 + 4\epsilon \|\theta^* + \Delta\| |\Delta^\top x_i| \\ &\quad + 2\epsilon (\|\theta^* + \Delta\| - \|\theta^*\|) \Delta^\top x_i + 2\epsilon \frac{\|\theta^* + \Delta\| \|\theta^*\| - (\theta^* + \Delta)^\top \theta^*}{\|\theta^*\|} |y_i - x_i^\top \theta^*|. \end{aligned}$$

Taking expectation on (x_i, y_i) , it becomes

$$\begin{aligned} &\mathbb{E} \left(l_\epsilon(\hat{\theta}, x_i, y_i) - l_\epsilon(\theta^*, x_i, y_i) \right) \\ &\leq (1 + \epsilon^2) \|\Delta\|^2 + O(\|\theta^* + \Delta_S\| \|\theta^*\| - (\theta^* + \Delta)^\top \theta^*) + O(\|\Delta\|) = O(\|\Delta\|). \end{aligned}$$

Therefore, if $\lambda \rightarrow 0$, with probability tending to 1,

$$R_0(\hat{\theta}, \epsilon) - R_0(\theta^*, \epsilon) \rightarrow 0.$$

To extend for general v , similar with Theorem 2, the change on v does not affect the convergence property of $\hat{\theta}$ after adjustment w.r.t. v . \square

Lemma 4. *Under the conditions in Theorem 9, when $\|\Delta_{S^c}\|_1 \leq c \|\Delta_S\|_1$, with probability tending to 1, for some $c_3 > 0$,*

$$\begin{aligned} &\frac{1}{n} \sum_{i=1}^n l_\epsilon(\hat{\theta}, x_i, y_i) - l(\theta^*, x_i, y_i) \\ &\geq \frac{1}{n} \sum_{i=1}^n \Delta^\top \frac{\partial l_\epsilon}{\partial \theta} + (\Delta^\top x_i)^2 - \|\Delta\|^2 + \epsilon^2 (1 - 2/\pi) \|\Delta\|^2 - c_3 \|\Delta\|_1 \sqrt{\frac{s \log d}{n}}, \end{aligned}$$

and

$$\frac{1}{\|\Delta_S\|_1^2} \left| \frac{1}{n} \sum_{i=1}^n (\Delta^\top x_i)^2 - \|\Delta\|^2 \right| \rightarrow 0.$$

Proof. Assume $x^\top \theta_2 - y > 0$, then for any θ_1 ,

$$\begin{aligned} & (x^\top \theta_1 - y)^2 + 2\epsilon \|\theta_1\| |x^\top \theta_1 - y| + \epsilon^2 \|\theta_1\|^2 - (x^\top \theta_2 - y)^2 - 2\epsilon \|\theta_2\| (x^\top \theta_2 - y) - \epsilon^2 \|\theta_2\|^2 \\ \geq & (x^\top \theta_1 - y)^2 + 2\epsilon \|\theta_1\| (x^\top \theta_1 - y) + \epsilon^2 \|\theta_1\|^2 - (x^\top \theta_2 - y)^2 - 2\epsilon \|\theta_2\| (x^\top \theta_2 - y) - \epsilon^2 \|\theta_2\|^2 \\ = & 2(\theta_1 - \theta_2)^\top x (x^\top \theta_2 - y) + (\theta_1 - \theta_2)^\top x x^\top (\theta_1 - \theta_2) + 2\epsilon^2 (\theta_1 - \theta_2)^\top \theta_2 + \epsilon^2 \|\theta_1 - \theta_2\|^2 \\ & + 2\epsilon \|\theta_1\| (x^\top \theta_1 - y) - 2\epsilon \|\theta_2\| (x^\top \theta_2 - y) - 2\epsilon (\theta_1 - \theta_2) \left[\frac{\theta_2}{\|\theta_2\|} (x^\top \theta_2 - y) + \|\theta_2\| x \right] \\ & + 2\epsilon (\theta_1 - \theta_2) \left[\frac{\theta_2}{\|\theta_2\|} (x^\top \theta_2 - y) + \|\theta_2\| x \right], \end{aligned}$$

where

$$\begin{aligned} & \|\theta_1\| (x^\top \theta_1 - y) - \|\theta_2\| (x^\top \theta_2 - y) - (\theta_1 - \theta_2) \left[\frac{\theta_2}{\|\theta_2\|} (x^\top \theta_2 - y) + \|\theta_2\| x \right] \\ = & \|\theta_1\| (x^\top \theta_2 - y) + \|\theta_1\| (\theta_1 - \theta_2)^\top x - \theta_1 \left[\frac{\theta_2}{\|\theta_2\|} (x^\top \theta_2 - y) + \|\theta_2\| x \right] + \theta_2^\top x \|\theta_2\| \\ = & \frac{\|\theta_1\| \|\theta_2\| - \theta_1^\top \theta_2}{\|\theta_2\|} (x^\top \theta_2 - y) + (\|\theta_1\| - \|\theta_2\|) (\theta_1 - \theta_2)^\top x \\ \geq & (\|\theta_1\| - \|\theta_2\|) (\theta_1 - \theta_2)^\top x \operatorname{sgn}(x^\top \theta_2 - y) \end{aligned}$$

Thus if $x_i^\top \theta^* - y_i > 0$,

$$l_\epsilon(\hat{\theta}, x_i, y_i) - l_\epsilon(\theta^*, x_i, y_i) \geq \Delta^\top \frac{\partial l_\epsilon(\theta^*, x_i, y_i)}{\partial \theta^*} + (\Delta^\top x_i)^2 + \epsilon^2 \|\Delta\|^2 + 2\epsilon(\|\theta^* + \Delta\| - \|\theta^*\|) \Delta^\top x_i \quad (14)$$

When $x^\top \theta_2 - y < 0$,

$$\begin{aligned} & \|\theta_1\| (y - x^\top \theta_1) - \|\theta_2\| (y - x^\top \theta_2) - (\theta_1 - \theta_2) \left[\frac{\theta_2}{\|\theta_2\|} (y - x^\top \theta_2) + \|\theta_2\| x \right] \\ = & \|\theta_1\| (y - x^\top \theta_2) - \|\theta_1\| (\theta_1 - \theta_2)^\top x - \theta_1 \left[\frac{\theta_2}{\|\theta_2\|} (y - x^\top \theta_2) + \|\theta_2\| x \right] + \theta_2^\top x \|\theta_2\| \\ = & \frac{\|\theta_1\| \|\theta_2\| - \theta_1^\top \theta_2}{\|\theta_2\|} (y - x^\top \theta_2) - (\|\theta_1\| - \|\theta_2\|) (\theta_1 - \theta_2)^\top x \\ \geq & -(\|\theta_1\| - \|\theta_2\|) (\theta_1 - \theta_2)^\top x \\ = & (\|\theta_1\| - \|\theta_2\|) (\theta_1 - \theta_2)^\top x \operatorname{sgn}(x^\top \theta_2 - y). \end{aligned}$$

As a result,

$$\begin{aligned}
 & \frac{1}{n} \sum_{i=1}^n l_\epsilon(\hat{\theta}, x_i, y_i) - l(\theta^*, x_i, y_i) \\
 & \geq \frac{1}{n} \sum_{i=1}^n \Delta^\top \frac{\partial l_\epsilon}{\partial \theta} + (\Delta^\top x_i)^2 + \epsilon^2 \|\Delta\|^2 + 2\epsilon(\|\theta^* + \Delta\| - \|\theta^*\|) \Delta^\top x_i \operatorname{sgn}(x_i^\top \theta^* - y_i) \\
 & = \frac{1}{n} \sum_{i=1}^n \Delta^\top \frac{\partial l_\epsilon}{\partial \theta} + (\Delta^\top x_i)^2 + \epsilon^2 \|\Delta\|^2 + 2\epsilon(\|\theta^* + \Delta\| - \|\theta^*\|) (\Delta_S^\top x_{i,S} + \Delta_{S^c}^\top x_{i,S^c}) \operatorname{sgn}(x_i^\top \theta^* - y_i) \\
 & \geq \frac{1}{n} \sum_{i=1}^n \Delta^\top \frac{\partial l_\epsilon}{\partial \theta} + (\Delta^\top x_i)^2 + \epsilon^2 \|\Delta\|^2 \\
 & \quad - 2\epsilon \|\|\theta^* + \Delta\| - \|\theta^*\|\| \left(\|\Delta_{S^c}\|_1 \left\| \frac{1}{n} \sum_{i=1}^n x_{i,S^c} \operatorname{sgn}(x_i^\top \theta^* - y_i) \right\|_\infty + O_p \left(\|\Delta_S\|_1 \sqrt{\frac{s \log s}{n}} \right) \right) \\
 & \quad + 2\epsilon(\|\theta^* + \Delta\| - \|\theta^*\|) \mathbb{E} \Delta_S^\top x_{i,S} \operatorname{sgn}(x_i^\top \theta^* - y_i) \\
 & \geq \frac{1}{n} \sum_{i=1}^n \Delta^\top \frac{\partial l_\epsilon}{\partial \theta} + (\Delta^\top x_i)^2 + \epsilon^2 \|\Delta\|^2 \\
 & \quad - 2\epsilon \|\|\theta^* + \Delta\| - \|\theta^*\|\| \left(\|\Delta_{S^c}\|_1 \left\| \frac{1}{n} \sum_{i=1}^n x_{i,S^c} \operatorname{sgn}(x_i^\top \theta^* - y_i) \right\|_\infty + O_p \left(\|\Delta_S\|_1 \sqrt{\frac{s \log s}{n}} \right) \right) \quad (15)
 \end{aligned}$$

$$\begin{aligned}
 & \quad - 2\epsilon \|\Delta\| \sqrt{\frac{2}{\pi}} \left| \Delta^\top \frac{\partial}{\partial \theta} \sqrt{\|\theta - \theta_0\|^2 + \sigma^2} \right| \\
 & \geq \frac{1}{n} \sum_{i=1}^n \Delta^\top \frac{\partial l_\epsilon}{\partial \theta} + (\Delta^\top x_i)^2 - \|\Delta\|^2 + \epsilon^2 (1 - 2/\pi) \|\Delta\|^2 - O_p \left(\|\Delta\|_1 \sqrt{\frac{s \log d}{n}} \right). \quad (16)
 \end{aligned}$$

From (15) to (16),

$$\begin{aligned}
 & \|\Delta\|^2 - 2\epsilon \|\Delta\| \sqrt{\frac{2}{\pi}} \left| \Delta^\top \frac{\partial}{\partial \theta} \sqrt{\|\theta - \theta_0\|^2 + \sigma^2} \right| \\
 & \geq \left(\|\Delta\| - \epsilon \sqrt{\frac{2}{\pi}} \left| \Delta^\top \frac{\partial}{\partial \theta} \sqrt{\|\theta - \theta_0\|^2 + \sigma^2} \right| \right)^2 - \epsilon^2 \frac{2}{\pi} \left| \Delta^\top \frac{\partial}{\partial \theta} \sqrt{\|\theta - \theta_0\|^2 + \sigma^2} \right|^2 \\
 & \geq -\epsilon^2 \frac{2}{\pi} \|\Delta\|^2.
 \end{aligned}$$

When $\|\Delta_{S^c}\|_1 \leq c\|\Delta_S\|_1$, since $\varepsilon_{d,n} := \max_{i,j} |\hat{\Sigma}_{i,j} - I_{i,j}| \rightarrow 0$ in probability, we have with probability tending to 1,

$$\frac{1}{\|\Delta_S\|_1^2} \left| \frac{1}{n} \sum_{i=1}^n (\Delta^\top x_i)^2 - \|\Delta\|^2 \right| = \frac{1}{\|\Delta_S\|_1^2} \left| \Delta^\top (I - \hat{\Sigma}) \Delta \right| \leq \frac{1}{\|\Delta_S\|_1^2} \|\Delta\|_1^2 \varepsilon_{d,n} \rightarrow 0.$$

□

Lemma 5.

$$\begin{aligned}
 & l_\epsilon(\hat{\theta}, x_i, y_i) - l_\epsilon(\theta^*, x_i, y_i) \\
 & \leq \Delta^\top \frac{\partial l_\epsilon(\theta^*, x_i, y_i)}{\partial \theta^*} + (\Delta^\top x_i)^2 + \epsilon^2 \|\Delta\|^2 + 4\epsilon \|\theta^* + \Delta\| |\Delta^\top x_i| \\
 & \quad + 2\epsilon(\|\theta^* + \Delta\| - \|\theta^*\|) \Delta^\top x_i + 2\epsilon \frac{\|\theta^* + \Delta\| \|\theta^*\| - (\theta^* + \Delta)^\top \theta^*}{\|\theta^*\|} |y_i - x_i^\top \theta^*|.
 \end{aligned}$$

Proof. Assume $x^\top \theta_2 - y > 0$,

$$\begin{aligned}
 & (x^\top \theta_1 - y)^2 + 2\epsilon \|\theta_1\| |x^\top \theta_1 - y| + \epsilon^2 \|\theta_1\|^2 - (x^\top \theta_2 - y)^2 - 2\epsilon \|\theta_2\| (x^\top \theta_2 - y) - \epsilon^2 \|\theta_2\|^2 \\
 = & (x^\top \theta_1 - y)^2 + 2\epsilon \|\theta_1\| (x^\top \theta_1 - y) + \epsilon^2 \|\theta_1\|^2 - (x^\top \theta_2 - y)^2 - 2\epsilon \|\theta_2\| (x^\top \theta_2 - y) - \epsilon^2 \|\theta_2\|^2 \\
 & + 2\epsilon \|\theta_1\| |x^\top \theta_1 - y| - 2\epsilon \|\theta_1\| (x^\top \theta_1 - y) \\
 = & 2(\theta_1 - \theta_2)^\top x (x^\top \theta_2 - y) + (\theta_1 - \theta_2)^\top x x^\top (\theta_1 - \theta_2) + 2\epsilon^2 (\theta_1 - \theta_2)^\top \theta_2 + \epsilon^2 \|\theta_1 - \theta_2\|^2 \\
 & + 2\epsilon \|\theta_1\| (x^\top \theta_1 - y) - 2\epsilon \|\theta_2\| (x^\top \theta_2 - y) - 2\epsilon (\theta_1 - \theta_2) \left[\frac{\theta_2}{\|\theta_2\|} (x^\top \theta_2 - y) + \|\theta_2\| x \right] \\
 & + 2\epsilon (\theta_1 - \theta_2) \left[\frac{\theta_2}{\|\theta_2\|} (x^\top \theta_2 - y) + \|\theta_2\| x \right] + 2\epsilon \|\theta_1\| |x^\top \theta_1 - y| - 2\epsilon \|\theta_1\| (x^\top \theta_1 - y) \\
 = & \Delta^\top \frac{\partial l_\epsilon}{\partial \theta_2} + (\Delta^\top x)^2 + \epsilon^2 \|\Delta\|^2 \\
 & + 2\epsilon \|\theta_1\| (x^\top \theta_1 - y) - 2\epsilon \|\theta_2\| (x^\top \theta_2 - y) - 2\epsilon (\theta_1 - \theta_2) \left[\frac{\theta_2}{\|\theta_2\|} (x^\top \theta_2 - y) + \|\theta_2\| x \right] \\
 & + 2\epsilon \|\theta_1\| |x^\top \theta_1 - y| - 2\epsilon \|\theta_1\| (x^\top \theta_1 - y).
 \end{aligned}$$

From Lemma 4, we know that

$$\begin{aligned}
 & \|\theta_1\| (x^\top \theta_1 - y) - \|\theta_2\| (x^\top \theta_2 - y) - (\theta_1 - \theta_2) \left[\frac{\theta_2}{\|\theta_2\|} (x^\top \theta_2 - y) + \|\theta_2\| x \right] \\
 = & \frac{\|\theta_1\| \|\theta_2\| - \theta_1^\top \theta_2}{\|\theta_2\|} (x^\top \theta_2 - y) + (\|\theta_1\| - \|\theta_2\|) (\theta_1 - \theta_2)^\top x.
 \end{aligned}$$

For $2\epsilon \|\theta_1\| |x^\top \theta_1 - y| - 2\epsilon \|\theta_1\| (x^\top \theta_1 - y)$, since $x^\top \theta_2 - y > 0$, it becomes

$$\begin{aligned}
 & 2\epsilon \|\theta_1\| |x^\top \theta_1 - y| - 2\epsilon \|\theta_1\| (x^\top \theta_1 - y) \\
 = & 2\epsilon \|\theta_1\| |x^\top (\theta_2 + \Delta) - y| - 2\epsilon \|\theta_1\| (x^\top \theta_1 - y) \\
 = & 2\epsilon \|\theta_1\| |x^\top (\theta_2 + \Delta) - y| - 2\epsilon \|\theta_1\| (x^\top (\theta_2 + \Delta) - y) \\
 \leq & 2\epsilon \|\theta_1\| (x^\top \theta_2 - y + |\Delta^\top x|) - 2\epsilon \|\theta_1\| (x^\top (\theta_2 + \Delta) - y) \\
 = & 2\epsilon \|\theta_1\| |\Delta^\top x| - 2\epsilon \|\theta_1\| x^\top \Delta \\
 = & 4\epsilon \|\theta_1\| |\Delta^\top x|.
 \end{aligned}$$

The case of $x^\top \theta_2 - y < 0$ has the same result. □