

---

# Why Do Artificially Generated Data Help Adversarial Robustness?

---

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

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

**Guang Cheng**  
Department of Statistics  
University of California, Los Angeles  
guangcheng@ucla.edu

## Abstract

In the adversarial training framework of Carmon et al. (2019); Goyal et al. (2021), people use generated/real unlabeled data with pseudolabels to improve adversarial robustness. We provide statistical insights to explain why the artificially generated data improve adversarial training. In particular, we study how the attack strength and the quality of the unlabeled data affect adversarial robustness in this framework. Our results show that with a high-quality unlabeled data generator, adversarial training can benefit greatly from this framework under large attack strength, while a poor generator can still help to some extent. To make adaptations concerning the quality of generated data, we propose an algorithm that performs online adjustment to the weight between the labeled real data and the generated data, aiming to optimize the adversarial risk. Numerical studies are conducted to verify our theories and show the effectiveness of the proposed algorithm.

## 1 Introduction

Adversarial training is a popular and simple way to improve the adversarial robustness of modern machine learning models. There are fruitful results in the theoretical justification and methodology development. Among various research directions, one interesting aspect is to use extra unlabeled data to assist adversarial training. Recent works successfully demonstrate a great improvement in the adversarial robustness with additional unlabeled data. For example, Carmon et al. (2019); Xing et al. (2021b, 2022), show that additional external real data help improve adversarial robustness; Goyal et al. (2021) uses synthetically generated data to improve the adversarial robustness and achieves the highest 66% adversarial testing accuracy for CIFAR-10 dataset under AutoAttack (AA) in Croce et al. (2020)<sup>1</sup> in the literature. The algorithms utilizing real/generated unlabeled data in adversarial training in Carmon et al. (2019); Goyal et al. (2021) are summarized in Figure 1. Note that to unify these two algorithms, we view the unlabeled real data used in Carmon et al. (2019) as synthetic data generated from an *ideal generator*.

However, two fundamental questions about utilizing additional unlabeled data remain unclear.

First, as shown by Carmon et al. (2019), unlabeled data provide more information about the density of the data near the decision boundary, leading to a more robust adversarial estimator (e.g., it increases the robust test accuracy of CIFAR-10 from 53.08% to 59.53% for 8/255  $\mathcal{L}_\infty$  attack). However, the performance improvement of this training strategy is limited when the attack strength is small or zero. Therefore, it is natural to ask

*Q1: Compared to clean training, why and how can adversarial training significantly benefit from unlabeled data under the framework of Figure 1?*

---

<sup>1</sup><https://robustbench.github.io>

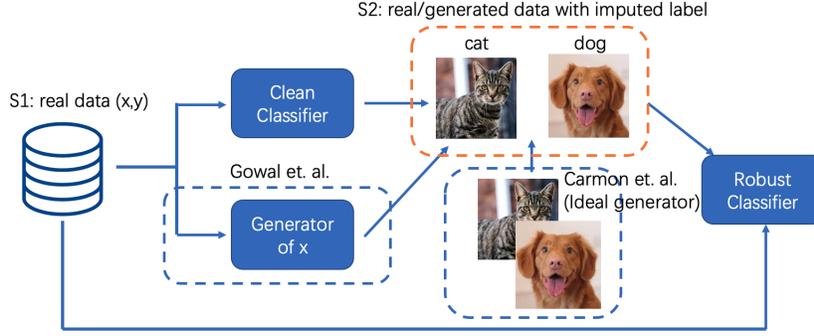


Figure 1: The procedure to do adversarial training with real/generated data. In Carmon et al. (2019), a clean classifier is trained to impute the label for additional real unlabeled data to form  $S_2$ . In Gowal et al. (2021), the clean classifier and the unlabeled data generator are trained via clean training to form  $S_2$ . Obtaining  $S_2$ , we use both  $S_1$  and  $S_2$  in the adversarial training. To unify the two algorithms, we call the unlabeled real data in Carmon et al. (2019) as data generated from *ideal generator*.

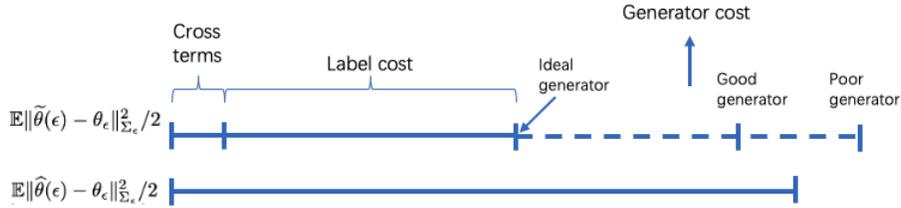


Figure 2: Exact risk decomposition of  $\tilde{\theta}(\epsilon)$  and  $\hat{\theta}(\epsilon)$  when  $n_2 \rightarrow \infty$ . The excess adversarial risk is  $\mathbb{E}R(\theta, \epsilon) - R(\theta_\epsilon, \epsilon) = \mathbb{E}\|\theta - \theta_\epsilon\|_{\Sigma_\epsilon}^2 / 2 + o$  for  $\theta \in \{\hat{\theta}(\epsilon), \tilde{\theta}(\epsilon)\}$  for some  $\Sigma_\epsilon$ . Ideal/good/poor generator refers to the unlabeled data generator ( $\mathcal{P}_a$ ). If  $n_2$  is finite, the label cost gets smaller because we use less noisy-label samples, but the generator cost gets larger because a smaller sample size gives larger estimation variance.

Second, besides the ideal generator in Carmon et al. (2019), Gowal et al. (2021) observe two other interesting phenomena regarding the quality of the generated unlabeled data. When the unlabeled data are generated from a generative model trained from original labeled data, although they introduce no more information beyond the original labeled data, the adversarial robustness gets enhanced. On the other hand, they also generate synthetic data from a multivariate Gaussian, which is much different from the real data. In this case, the generated data still slightly improve the robustness. Hence, it is important to investigate that

*Q2: How does the quality of the unlabeled data generator affect the adversarial robustness?*

This work aims to provide mathematical answers for the above two questions. Existing literature, (e.g., Alayrac et al., 2019; Uesato et al., 2019; Carmon et al., 2019; Zhai et al., 2019; Sehwag et al., 2021; Xing et al., 2021b, 2022) provide a justification that adversarial training potentially benefit from unlabeled/generated data and fail to explain that unlabeled data help little when attack strength is small. In contrast, we explicitly explain **why** and **how** adversarial training can be improved via an **exact** error decomposition rather than an upper bound statement.

For the convenience of our discussion below, we preliminarily introduce some notations, with detailed explanations given later. Denote  $\tilde{\theta}(\epsilon)$  as the robust estimator utilizing additional unlabeled data as in Figure 1, and  $\hat{\theta}(\epsilon)$  as the estimator via vanilla adversarial training. Denote the true distribution of  $X$  as  $\mathcal{P}_0$ , and the distribution of the generated unlabeled data as  $\mathcal{P}_a$ . The true conditional distribution of  $Y|X$  is  $\mathcal{P}_y$ , and  $Y|X - \mathbb{E}[Y|X]$  has a distribution  $\mathcal{P}_\epsilon$ . As defined in Figure 1,  $S_1$  is the set of original labeled training data with sample size  $n_1$ , and  $S_2$  is the set of generated data with size  $n_2$ . Denote  $\alpha = n_1/(n_1 + n_2)$  as the proportion of real data. Let  $l_\epsilon$  be the adversarial loss function w.r.t. attack strength  $\epsilon$ ,  $l_0$  be the clean loss (i.e.,  $\epsilon = 0$ ), and  $\theta_0$  be the parameter of the underlying true clean model.

**Contributions** We conduct theoretical analysis under some simple statistical setups. Through studying the convergence of  $\tilde{\theta}(\epsilon)$  and  $\hat{\theta}(\epsilon)$ , we decompose the excess adversarial risk of  $\tilde{\theta}(\epsilon)$  into two major components (and their cross product): label cost (related to the quality of the imputed labels of generated unlabeled data) and generator cost (related to the quality of unlabeled data generator) as in Figure 2. Using this decomposition, we can answer *Q1* and *Q2*.

First, unlike  $\hat{\theta}(\epsilon)$  whose risk noticeably affected as  $\epsilon$  increases, since the adversarial risk is stable to the imputation model, the label cost of  $\tilde{\theta}(\epsilon)$  is more insensitive w.r.t.  $\epsilon$ . In consequence, although in clean training, the label cost and the risk of  $\hat{\theta}(\epsilon)$  are comparable, the former one becomes relatively smaller as  $\epsilon$  increases. As in Figure 2, when  $\epsilon > 0$ , the label cost of  $\tilde{\theta}(\epsilon)$  is smaller than the risk of  $\hat{\theta}(\epsilon)$ . Using infinite data from the ideal unlabeled data generator, the generator cost becomes zero, and hence the overall cost of  $\tilde{\theta}(\epsilon)$  is smaller than  $\hat{\theta}(\epsilon)$ .

Second, we characterize how the quality of the unlabeled data generator affects the final adversarial robustness. In general, a better data generator is always preferred. For the ideal generator, it will always improve adversarial robustness **effectively**, i.e., the generator cost is negligible when  $n_2 \rightarrow \infty$ . If the generator is learned from  $S_1$  and captures extra information, it may **effectively** help as well when  $n_2 \rightarrow \infty$ . For a poor generator, it may **slightly** improve the performance for a large attack when  $n_2$  is small, and the generator cost is always not negligible.

Finally, following Goyal et al. (2021), we balance the weights of  $S_1$  and  $S_2$  in the adversarial training process in order to achieve the optimal adversarial robustness. We show that a larger  $n_2$  is always preferred with a proper choice of weight, even if generated from non-ideal generators. Tuning the best choice of weights via repeated trial-and-error runs can be infeasible due to the slow convergence of adversarial training. We propose an algorithm that dynamically adapts the weight during the training of neural networks and shows its promising performance in Section 5.

## 2 Adversarial Training

To formally introduce adversarial training, let  $l_0$  denote the loss function and  $f_\theta(x)$  be the model with parameter  $\theta$ . The adversarial loss function  $l_\epsilon$  and the (population) adversarial risk are defined as

$$R(\theta, \epsilon) := \mathbb{E}_{\mathcal{P}_0} [l_0(x + A_\epsilon(f_\theta, x, y), y, \theta)] := \mathbb{E}_{\mathcal{P}_0} [l_\epsilon(x, y, \theta)],$$

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

$$A_\epsilon(f_\theta, x, y) := \operatorname{argmax}_{z \in B_p(0, \epsilon)} \{l_0(x + z, y, \theta)\}, \quad (1)$$

where  $B_p(x, r)$  is a  $\mathcal{L}_p$  ball centering at  $x$  with radius  $r$ . Denote  $\theta_\epsilon = \operatorname{argmin}_\theta R(\theta, \epsilon)$ . Note that for simplicity, we only consider  $p = 2$  in our theorems and simulations, while in the numerical studies of real-data experiment, we follow the routine setup and set  $p = \infty$ .

Given the data set  $S_1$  consisting of  $n_1$  i.i.d. labeled samples, the estimator  $\hat{\theta}(\epsilon)$  from the vanilla adversarial training aims to minimize the empirical version of  $R(\theta, \epsilon)$ :

$$\hat{R}(\theta, \epsilon) = \frac{1}{n_1} \sum_{(x, y) \in S_1} l_\epsilon(x, y, \theta). \quad (2)$$

To take the advantage of the extra unlabeled data, we first assign pseudo-response variable  $\hat{y} = g(f_{\hat{\theta}(0)}(x), \epsilon)$  to them, where for regression,  $g(a, b) = a + b$  and  $\epsilon$  follows the true noise distribution; for classification,  $g(a, b) = 1(a \geq b)$  and  $\epsilon \sim \operatorname{Unif}[0, 1]$ . Then  $\tilde{\theta}(\epsilon)$  aims to minimize

$$\tilde{R}(\theta, \epsilon) = \frac{1}{n_1 + n_2} \left( \sum_{(x, y) \in S_1} l_\epsilon(x, y, \theta) + \sum_{(x, \hat{y}) \in S_2} l_\epsilon(x, \hat{y}, \theta) \right). \quad (3)$$

To evaluate the performance of an estimator  $\theta$ , we use the excess adversarial risk  $\mathbb{E}R(\theta, \epsilon) - R(\theta_\epsilon, \epsilon)$ .

Besides the pseudolabel  $\hat{y}$ , we define the the imaginary ‘‘true’’ response for the unlabeled data  $y = g(f_{\theta_0}(x), \epsilon)$ , where  $\epsilon$  is the same realization used in generating pseudolabels. These true responses are not observable but will be used when analyzing the convergence rate of  $\tilde{\theta}(\epsilon)$ . To evaluate the size of vectors and matrices, we denote  $\|\cdot\|$  as the  $\mathcal{L}_2$  norm of vectors, operator norm of matrices, and denote  $\|a\|_A^2 := a^\top A a$  for any vector  $a$  and positive definite matrix  $A$ .

### 3 Related Literature

Commonly used techniques (without extra data fed into training process) in adversarial training include adversarial regularization Goodfellow et al. (2015); Zhang et al. (2019); Wang et al. (2019b), curriculum-based algorithms Cai et al. (2018); Zhang et al. (2020a).

Theoretical investigations, besides the aforementioned works, have been conduct for adversarial training from different perspectives. For instance, Chen et al. (2020); Javanmard et al. (2020); Javanmard & Soltanolkotabi (2020); Taheri et al. (2021); Yin et al. (2018); Raghunathan et al. (2019); Najafi et al. (2019); Min et al. (2020); Hendrycks et al. (2019); Wu et al. (2020b); Xing et al. (2021a) studied the statistical properties of adversarial training, Sinha et al. (2017); Wang et al. (2019a) studied the optimization convergence of adversarial training, Zhang et al. (2020b); Allen-Zhu & Li (2020); Wu et al. (2020a); Xiao et al. (2021) studied theoretical issues related to adversarial training with Deep learning models.

### 4 Main Result

In this section, we present the main theorem, answer the two key questions, and show that balancing the weights of  $S_1$  and  $S_2$  leads to the better utility of the generated data. **Due to the page limit, we postpone the whole simulation study section to Appendix B.** Briefly speaking, our simulation results justify all the theoretical results below.

#### 4.1 General Convergence Result

To present our theoretical investigations in  $\tilde{\theta}(\epsilon)$ , we first introduce some assumptions:

**Assumption 1.** *The loss function satisfies the following conditions:*

- A1. *The clean loss function  $l_0$  is the square loss or logistic loss. The attributes  $X \in \mathbb{R}^d$  follow a sub-Gaussian distribution. The true clean model  $\theta_0$  is defined as  $\mathbb{E}[Y|X = x] = x^\top \theta_0$  for regression and  $P(Y = 1|X = x) = 1/(1 + \exp(x^\top \theta_0))$  for classification. In addition, it satisfies that  $\|\theta_0\| \leq b_0$  for some constant  $b_0 > 0$  and the noise in regression has finite variance. The distribution  $\mathcal{P}_a$  is sub-Gaussian as well.*
- A2. *The distribution  $\mathcal{P}_a$  and  $n_2$  satisfy  $\|\mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon)\| = o((n_1 + n_2)/(n_2 \log n_1))$ , where  $\mathcal{P}_a \otimes \mathcal{P}_y$  denotes the joint distribution induced by a marginal distribution  $\mathcal{P}_a$  and a conditional distribution  $\mathcal{P}_y$ .*

Assumption A1 is for the simplicity of derivation. When doing Taylor expansion as in (4), the remainder term  $o$  (if exists) does not explode under A1. Under Assumption A2, if  $\mathcal{P}_a$  is of poor quality, i.e.,  $\|\mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon)\|$  is not vanishing, then one can only take a relatively small  $n_2$  compared to  $n_1$ . Vice versa, to allow a large  $n_2$ , the generator  $\mathcal{P}_a$  need to be almost ‘‘unbiased’’. It is possible to relax Assumptions A1 and A2 as discussed in Section E in the appendix. However, the relaxations are technical and tailored to the derivations of Theorem 1.

The following theorem is a general result to decompose  $\mathbb{E}R(\tilde{\theta}(\epsilon), \epsilon) - R(\theta_\epsilon, \epsilon)$ :

**Theorem 1.** *Under Assumption A1 to A2, assuming the unlabeled data generator is independent to  $S_1$ , when  $n_1 \rightarrow \infty$ , the excess adversarial risk using  $\tilde{\theta}(\epsilon)$  is*

$$\mathbb{E}R(\tilde{\theta}(\epsilon), \epsilon) - R(\theta_\epsilon, \epsilon) = \frac{1}{2} \mathbb{E} \left\| \tilde{\theta}(\epsilon) - \theta_\epsilon \right\|_{\Sigma_\epsilon}^2 + o, \quad (4)$$

where  $\theta_\epsilon - \tilde{\theta}(\epsilon)$  is dominated by

$$\underbrace{\frac{n_2 \Sigma_\epsilon^{-1}}{n_1 + n_2} \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \left( \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, \hat{Y}, \theta_\epsilon) - \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon) \right)}_{:= E_1} + \underbrace{\frac{\Sigma_\epsilon^{-1}}{n_1 + n_2} \left( \sum_{S_1, S_2} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta_\epsilon) \right)}_{:= E_2},$$

and

$$\Sigma_\epsilon = \frac{\partial^2}{\partial \theta_\epsilon^2} \mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta_\epsilon), \quad \tilde{\Sigma}_\epsilon = \frac{\partial^2}{\partial \theta_\epsilon \partial \theta_0} \mathbb{E}_{\mathcal{P}_a} \mathbb{E}_\epsilon l_\epsilon(X, g(f_{\theta_0}(X), \epsilon), \theta_\epsilon).$$

The term “ $o$ ” represents the remainder terms which are not dominant.

Therefore, taking the decomposition of  $\tilde{\theta}(0)$  into  $\mathbb{E}R(\tilde{\theta}(\epsilon), \epsilon)$ , the excess adversarial risk becomes

$$\mathbb{E}R(\tilde{\theta}(\epsilon), \epsilon) - R(\theta_\epsilon, \epsilon) = \frac{1}{2}\mathbb{E}_{\hat{\theta}(0)}\|E_1\|_{\Sigma_\epsilon}^2 + \frac{1}{2}\mathbb{E}_{S_1, S_2}\|E_2\|_{\Sigma_\epsilon}^2 + \mathbb{E}_{S_1, S_2}E_1^\top \Sigma_\epsilon E_2 + o,$$

where  $\mathbb{E}_{\hat{\theta}(0)}\|E_1\|_{\Sigma_\epsilon}^2/2$  quantifies how the quality of pseudolabel affects the excess risk (label cost), and  $\mathbb{E}_{S_1, S_2}\|E_2\|_{\Sigma_\epsilon}^2/2$  quantifies how the quality of  $\mathcal{P}_a$  affects the excess risk (generator cost).

Meanwhile, the vanilla adversarial training estimate  $\hat{\theta}(\epsilon)$  satisfies

$$\theta_\epsilon - \hat{\theta}(\epsilon) = \Sigma_\epsilon^{-1} \left[ \frac{1}{n_1} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \theta_\epsilon) \right) \right] + o.$$

A proof sketch can be found in Section 6, while the detailed proof is postponed to the appendix. For simplicity of the derivation,  $\mathcal{P}_a$  in Theorem 1 is independent to  $S_1$ . A case-by-case study is essential if  $\mathcal{P}_a$  is trained from  $S_1$ . In Example 2 later, we provide a case where  $\mathcal{P}_a$  is trained from  $S_1$ .

Based on Theorem 1, as shown in Figure 2 in the introduction, the error of  $\tilde{\theta}(\epsilon)$  consists of two parts and their cross terms:

**Label cost**,  $\mathbb{E}_{\hat{\theta}(0)}\|E_1\|_{\Sigma_\epsilon}^2/2$ : It is mainly due to the discrepancy between the pseudolabel  $\hat{y}$  and “true” label  $y$  of the generated data. The label cost depends on both clean training error  $\hat{\theta}(0) - \theta_0$  and the matrix  $\Sigma_\epsilon^{-1}\tilde{\Sigma}_\epsilon$ , and is maximized when  $n_2 \rightarrow \infty$ . It measures the quality of the pseudolabel.

**Generator cost**,  $\mathbb{E}_{S_1, S_2}\|E_2\|_{\Sigma_\epsilon}^2/2$ : By the definition of  $\theta_\epsilon$ , i.e. the global minima of the adversarial risk, the expectation of  $\frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta_\epsilon)$  under  $\mathcal{P}_0 \otimes \mathcal{P}_y$  is  $\mathbf{0}$ , thus  $\sum_{S_1} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta_\epsilon) = O_p(\sqrt{n_1})$ . Consequently, the generator cost mostly rely on the term  $\sum_{S_2} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta_\epsilon)$ , which depends on the quality of the unlabeled data generator  $\mathcal{P}_a$ . For ideal  $\mathcal{P}_a$ , i.e.,  $\mathcal{P}_a = \mathcal{P}_0$ ,  $\mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon) = \mathbf{0}$ , and the generator cost goes to zero when  $n_2 \rightarrow \infty$ . For non-ideal  $\mathcal{P}_a$ ,  $\|\mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon)\|$  may be vanishing, but not exact zero.

Besides Figure 2, we also provide graphical illustration on how the label cost and generator cost change along  $n_2/n_1$  in Section F to ease the understanding.

## 4.2 Discrepancy between Clean and Adversarial Training with Unlabeled Data

We provide some intuitions to explain the discrepancy in the algorithm performance between  $\epsilon = 0$  and  $\epsilon > 0$  from two aspects, the loss function aspect and information (minimax lower bound) aspect. The following discussion assumes that the ideal data generator is used.

**Loss function and label cost** We rewrite the adversarial loss as  $R_\epsilon(\theta, \theta_0) = \mathbb{E}l_\epsilon(X, g(f_{\theta_0}(X), \epsilon), \theta)$  to explicitly reveal its dependency on the true parameter  $\theta_0$ . A crucial difference between clean training and adversarial training is the roles of the  $\theta$  and  $\theta_0$  in the loss functions. For clean loss, e.g. linear regression, the loss function can be rewritten as  $(x^\top \theta - x^\top \theta_0 - \epsilon)^2$ , where  $\theta$  and  $\theta_0$  play symmetric roles. In contrast, for adversarial loss, the loss becomes  $((x + A_\epsilon(f_\theta, x, y))^\top \theta - x^\top \theta_0 - \epsilon)^2$ , thus  $\theta$  and  $\theta_0$  have asymmetric roles. The adversarial loss is more sensitive to  $\theta$  than  $\theta_0$ . Consequently, when  $\hat{\theta}(0)$  is treated as true parameter and used to impute labels for  $S_2$ , the impact of the error in  $\hat{\theta}(0)$  is much less influential for adversarial training than clean training.

In Theorem 1, the matrix  $\tilde{\Sigma}_\epsilon$  reflects the sensitivity mentioned in the above illustration. Assume  $\epsilon = 0$ , for both

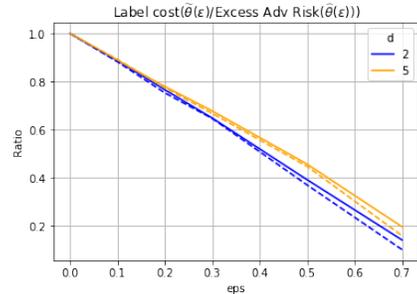


Figure 3: The ratio between label cost of  $\hat{\theta}(\epsilon)$  and  $\mathbb{E}\|\hat{\theta}(\epsilon) - \theta_\epsilon\|_{\Sigma_\epsilon}^2$  in linear regression,  $n_2 \rightarrow \infty$ . Solid line: theory. Dashed line: simulation. Derivations are in Appendix H.

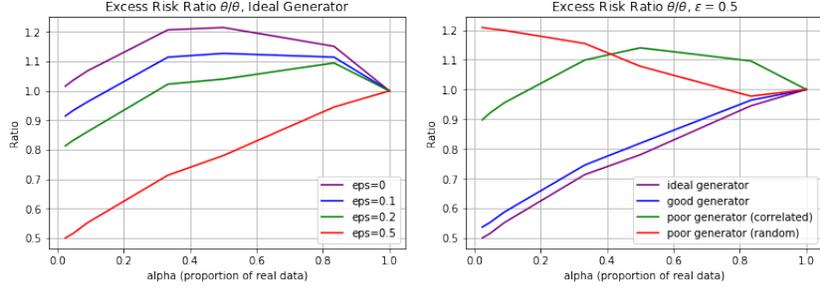


Figure 4: Simulation: how the attack strength ( $\epsilon$ ), quality of data generator, and the proportion of real data ( $\alpha = n_1/(n_1 + n_2)$ ) affect the convergence of the adversarial estimator. The y-axis represents the excess adversarial risk ratio of  $\tilde{\theta}(\epsilon)$  and  $\hat{\theta}(\epsilon)$ , i.e.,  $[\mathbb{E}R(\tilde{\theta}(\epsilon), \epsilon) - R(\theta_\epsilon, \epsilon)]/[\mathbb{E}R(\hat{\theta}(\epsilon), \epsilon) - R(\theta_\epsilon, \epsilon)]$ . A smaller-than-one ratio implies better performance of  $\tilde{\theta}(\epsilon)$ . Detailed model information is in Appendix B: the left panel corresponds to the left panel of Figure B.1, and the right panel is for  $\epsilon = 0.5$  from four different generators as in Figure B.1 and B.2.

linear regression and logistic regression, we have  $-\tilde{\Sigma}_0 = \Sigma_0$ , so when  $n_2 \rightarrow \infty$ ,

$$\tilde{\theta}(0) - \theta_0 = \hat{\theta}(0) - \theta_0 + o,$$

which means that clean training cannot be improved using  $S_2$  with pseudolabel.

On the other hand, for adversarial training, the matrix  $\Sigma_\epsilon^{-1}\tilde{\Sigma}_\epsilon$  is no longer  $-I_d$ , and the singular values get changed. As a consequence, the label cost is cheaper than the overall risk of  $\hat{\theta}(\epsilon)$  even if the label cost is maximized as  $n_2 \rightarrow \infty$ . A toy simulation in Figure 3 justifies this. In Figure 3, we calculate the theoretical values of the label cost of  $\tilde{\theta}(\epsilon)$  and the risk of  $\hat{\theta}(\epsilon)$  to compare them. When  $\epsilon$  gets larger, the label cost in  $\tilde{\theta}(\epsilon)$  becomes much smaller compared to the overall risk of  $\hat{\theta}(\epsilon)$ .

Together with an ideal  $\mathcal{P}_a$  that yields zero generator cost ( $n_2 \rightarrow \infty$ ), we obtain that  $\tilde{\theta}(\epsilon)$  is much better than  $\hat{\theta}(\epsilon)$ . The simulation results in the left panel of Figure 4 verifies this as well. All these theoretical insights and numerical observations under  $\epsilon > 0$ , are summarized in Figure 2. Figures 2 and 3 jointly are the main **answer towards Q1**. Note that we assert adversarial training benefits much more from unlabeled data than clean training, rather than that clean training cannot benefit from unlabeled data. See the remark below:

**Remark 1.** *It is possible to improve clean training using unlabeled data with pseudolabels. In Lee et al. (2013); Sun et al. (2017), in each training iteration, pseudolabels are updated based on the model of the current step. This is equivalent to a SGD with the gradient of  $\sum_{S_1} l(f_\theta(x), y)/n_1 + \mathcal{L}(\theta, \theta^{(t)})$  for some regularization term  $\mathcal{L}$  that stabilizes the training process, where  $\theta^{(t)}$  as the current model.*

**Minimax lower bound** Besides the exact decomposition of  $\tilde{\theta}(\epsilon)$  to answer Q1, we also provide minimax lower bound result to reveal another difference between clean and adversarial training.

Minimax rate refers to the best possible convergence rate that can be achieved by **any** estimator in the worst case given finite samples. As mentioned in Dan et al. (2020), it represents the information-theoretical limit given the data set.

In general, for the clean estimate without model misspecification, this information limit regards mostly the uncertainty about the conditional distribution of  $Y$  given  $X = x$ , which completely determines  $\theta_0$ . The additional information about the marginal distribution of  $X$  only affects the multiplicative constant of the bound but not the rate of convergence.

However, the true robust model  $\theta_\epsilon$  can depend on both  $Y|X$  and  $X$  for the adversarial estimate. The following example demonstrates the difference between clean and adversarial training:

**Example 1.** *Consider linear regression under Gaussian model  $Y = \theta_0^\top X + \varepsilon$  where  $X \sim N(\mathbf{0}, \Sigma)$  and  $\varepsilon \sim N(0, \sigma^2)$ . Then  $\theta_0$  is the best clean model, and  $\theta_\epsilon = (I + \lambda_\epsilon \Sigma^{-1})^{-1} \theta_0$  for some  $\lambda_\epsilon > 0$ .*

The derivation of Example 1 can be found in Appendix H. Based on Example 1, it is intuitive that the minimax lower bound for the adversarially robust estimate is contributed by two sources, one for the

uncertainty of  $\theta_0$ , and the other one for the distribution of  $X$ , where the additional unlabeled data can improve the latter term. The following proposition reveals this difference theoretically:

**Proposition 1.** *Under the setup of Example 1, assume  $\|\theta_0\| \leq b_0$  and there are  $n_2$  unlabeled real data. Assume  $R(\theta_\epsilon, \epsilon)$  is bounded from above and below, and  $\epsilon \in [0, \epsilon^*]$  with some constant  $\epsilon^*$  for  $\mathcal{L}_2$  attack. Assume  $\Sigma$  is unknown. Then for some positive constant  $c$ , for any estimator  $\hat{f}$ ,*

$$\inf_{\hat{f}} \sup_{\mathcal{P}_0 \otimes \mathcal{P}_y} \mathbb{E}[R(\hat{f}, \epsilon) - R(\theta_\epsilon, \epsilon)] \geq c \max \left( \inf_{\hat{\theta}_0} \sup_{\theta_0} \mathbb{E}\|\hat{\theta}_0 - \theta_0\|^2, \epsilon^2 b_0^2 \inf_{\hat{\Sigma}} \sup_{\Sigma} \mathbb{E}\|\hat{\Sigma} - \Sigma\|^2 \right).$$

### 4.3 Quality of Unlabeled Data Matters

We study how the quality of a non-ideal unlabeled data generator affects adversarial training. Unsurprisingly, we always prefer a high-quality/ideal  $\mathcal{P}_a$ . In addition, when there is some extra information of the marginal distribution of  $X$  (such as parametric modeling assumptions) which cannot be directly utilized in the vanilla adversarial training process, training  $\mathcal{P}_a$  from  $S_1$  can potentially give a high-quality adversarially robust model. On the other hand, given a poor  $\mathcal{P}_a$ , via bias-variance trade-off, it is still possible to reduce the error of  $\tilde{\theta}(\epsilon)$ . Note that the discussion in the section assumes that the unlabeled data generator ( $\mathcal{P}_a$ ) is independent to the training data  $S_1$ , as assumed in Theorem 1, unless stated otherwise.

**A high-quality  $\mathcal{P}_a$  is preferred** The term  $E_2$  in Theorem 1 satisfies

$$\begin{aligned} \mathbb{E}E_2E_2^\top &= \underbrace{\frac{n_2^2}{(n_1 + n_2)^2} \left[ \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon) \right] \left[ \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon) \right]^\top}_{\text{Square bias}} \quad (5) \\ &+ \underbrace{\frac{n_1}{(n_1 + n_2)^2} \text{Var}_{\mathcal{P}_0 \otimes \mathcal{P}_y} \left( \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon) \right) + \frac{n_2}{(n_1 + n_2)^2} \text{Var}_{\mathcal{P}_a \otimes \mathcal{P}_y} \left( \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon) \right)}_{\text{Variance}}. \end{aligned}$$

For the variance term in (5), when  $n_2 \rightarrow \infty$ , it is always negligible for reasonable  $\mathcal{P}_a$ .

For the square bias in (5), it is in general nonzero if  $\mathcal{P}_a \neq \mathcal{P}_0$ , and is similar to the maximum mean discrepancy (MMD, Gretton et al., 2012) that measures the dissimilarity between  $\mathcal{P}_a$  and  $\mathcal{P}_0$ . There exists some function class  $\mathcal{F}$  such that  $\partial l_\epsilon / \partial \theta_\epsilon \in \mathcal{F}$  and

$$\begin{aligned} \left\| \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon) \right\| &= \left\| \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon) - \underbrace{\mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon)}_{=0} \right\| \\ &\leq \sqrt{d} \text{MMD}(\mathcal{F}, \mathcal{P}_a, \mathcal{P}_0). \end{aligned}$$

Based on the above decomposition, a higher quality of the unlabeled data generator is preferred because it leads to small generator cost. And as shown in Figure 2, with the ideal  $\mathcal{P}_a$ , the generator cost can be zero when  $n_2 \rightarrow \infty$ . Using an ideal/good  $\mathcal{P}_a$ , we can have  $\mathbb{E}\|\tilde{\theta}(\epsilon) - \theta_\epsilon\|_{\Sigma_\epsilon}^2 < \mathbb{E}\|\hat{\theta}(\epsilon) - \theta_\epsilon\|_{\Sigma_\epsilon}^2$ .

A graphical illustration can be found in the right panel of Figure 4. When  $\epsilon$  is large, we know that the ideal  $\mathcal{P}_a$  will efficiently improve the adversarial robustness. When  $\mathcal{P}_a$  gets worse, there is less improvement in the adversarial robustness. This is the main **answer towards Q2**, and we illustrate some specific cases in the following discussions.

**Training data generator from  $S_1$**  For the vanilla adversarial training, there is no trivial way to incorporate the extra information of the marginal distribution of  $X$  into the optimization of (2). In contrast,  $\tilde{\theta}(\epsilon)$  can use the extra information to train a good unlabeled data generator from  $S_1$ .

**Example 2** (Sparse Covariance Matrix Estimate). *Assume  $X \sim N(\theta, \Sigma)$  and  $\Sigma$  is unknown.*

*We follow Cai et al. (2010) to consider a family of sparse covariance matrix as follows:*

$$\mathcal{F}_\beta = \left\{ \Sigma : \max_j \sum_{\{i: |i-j| > k\}} |\sigma_{ij}| \leq Mk^{-\beta} \forall k, \lambda_{\max}(\Sigma) \leq M_0, \lambda_{\min}(\Sigma) \geq m_0 > 0 \right\}.$$

*If we use the sparse estimator proposed in Cai et al. (2010), then  $\mathbb{E}\|\hat{\Sigma}_{\text{sparse}} - \Sigma\|^2 = O(n_1^{-2\beta/(2\beta+1)})$ .*

Besides, when  $\widehat{\Sigma}$  is either the sample covariance matrix or  $\widehat{\Sigma}_{\text{sparse}}$ , with probability tending to 1 over the generation of  $S_1$ , the decomposition of  $\tilde{\theta}(\epsilon)$  in Theorem 1 still holds, and the bias in (5) satisfies

$$\frac{n_2^2}{(n_1 + n_2)^2} \left\| \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon) \right\|^2 = O\left(\frac{n_2^2 \|\widehat{\Sigma} - \Sigma\|^2}{(n_1 + n_2)^2}\right). \quad (6)$$

If  $\Sigma \in \mathcal{F}_\beta$  and  $d \gg n^{1/(2\beta+1)}$ , when  $n_2 \rightarrow \infty$ ,  $\tilde{\theta}(\epsilon)$  with  $\widehat{\Sigma}_{\text{sparse}}$  leads to a rate of  $o(d/n_1)$  for the generator cost, which is negligible compared to the label cost.

Compared to Theorem 1, a difference in the assumption is that in Example 2, the generator is not independent to  $S_1$ . Therefore, the proof of (6), presented in Appendix G, is more involved.

**Bias-variance trade-off in poor  $\mathcal{P}_a$**  For a poor unlabeled data generator, when  $n_2 \rightarrow \infty$ , it leads to a poor  $\tilde{\theta}(\epsilon)$  as shown in Figure 2. However, based on the bias-variance trade-off of  $E_2$ , for a relatively  $n_2$ , it is still possible to obtain a  $\tilde{\theta}(\epsilon)$  better than  $\widehat{\theta}(\epsilon)$ .

**Proposition 2** (Bias-variance trade-off). *Assume the conditions of Theorem 1 hold,  $n_2 \ll n_1$  and the unlabeled data generator is independent to  $S_1$ , then*

$$\|\mathbb{E}\tilde{\theta}(\epsilon) - \theta_\epsilon\|_{\Sigma_\epsilon}^2 = O\left(\frac{dn_2^2}{(n_1 + n_2)^2}\right), \quad \text{tr}(\text{Var}[\Sigma_\epsilon^{1/2}(\tilde{\theta}(\epsilon) - \theta_\epsilon)]) = O\left(\frac{d}{n_1 + n_2}\right).$$

As shown by Proposition 2, there is a trade-off between bias and variance w.r.t. the size of  $S_2$ . This implies a U-shaped curve of the generator cost w.r.t.  $S_2$ , and there exists some small  $n_2$  that strikes the balance of this trade-off and can improve the adversarial robustness.

#### 4.4 Balancing the Weights between $S_1$ and $S_2$

Following Goyal et al. (2021) and Carmon et al. (2019), we try to improve the performance of  $\tilde{\theta}(\epsilon)$  via balancing the weights of  $S_1$  and  $S_2$  in (3), and consider the weighted minimization,

$$\tilde{R}(\theta, w, \epsilon) = \frac{1}{n_1 + wn_2} \left( \sum_{(x,y) \in S_1} l_\epsilon(x, y, \theta) + w \sum_{(x,\hat{y}) \in S_2} l_\epsilon(x, \hat{y}, \theta) \right). \quad (7)$$

When using the ideal  $\mathcal{P}_a$  and assuming  $\widehat{\theta}(0) \equiv \theta_0$  (i.e., we can generate extra independent labeled real data), trivially the optimal  $w$  minimizing the risk of  $\tilde{\theta}(\epsilon)$  is 1. When using non-ideal  $\mathcal{P}_a$  or  $\widehat{\theta}(0) \neq \theta_0$ , intuitively, one need to downweight the generated data and the best  $w < 1$ . We demonstrate this observation in the simulation in Figure B.3 as well.

In addition, as stated in the following proposition, through taking the optimal choice of  $w$ , for any  $\mathcal{P}_a$ , a larger  $n_2$  always leads to better  $\tilde{\theta}$  when  $\epsilon > 0$ .

**Proposition 3.** *We assume the assumptions of Theorem 1 hold, and consider a fixed  $(\epsilon > 0, n_1)$ . Denote  $\tilde{\theta}(w, n_2)$  as the minimizer of (7), and  $w^*(n_2)$  minimizes the excess adversarial risk of  $\tilde{\theta}(w, n_2)$  w.r.t.  $w$ . Then for  $n'_2 > n_2$ , taking  $w'$  such that  $w'n'_2 = w^*(n_2)n_2$ , the risk of  $\tilde{\theta}(w', n'_2)$  is smaller than  $\tilde{\theta}(w^*(n_2), n_2)$ , and the risk of  $\tilde{\theta}(w^*(n'_2), n'_2)$  is further smaller than  $\tilde{\theta}(w', n'_2)$ . As a result, a larger  $n_2$  always gives a better  $\tilde{\theta}$ .*

#### 4.5 Summary of Numerical Experiments

**Due to the space limit, we postpone simulations and most real experiments to Appendix B, C.** Below is a summary of simulation and empirical studies.

In Appendix B, we aim to use simple models to numerically verify: (1) given the ideal data generator, the performance of  $\tilde{\theta}(\epsilon)$  is better than  $\widehat{\theta}(\epsilon)$  when  $\epsilon$  deviates from zero; (2) the better quality of the data generator implies the better performance of  $\tilde{\theta}(\epsilon)$ ; and (3) balancing the weight between  $S_1$  and  $S_2$  improves the performance. We observe all (1) to (3) in the simulations.

In Appendix C, we aim to verify that the label cost and the generator cost are important factors in deep learning. We aim to show (1) adding more unlabeled samples from the ideal generator will improve adversarial robustness, and (2) adding unlabeled samples from a poor generator with a small  $n_2$  will slightly improve the performance. We perform an experiment on the CIFAR-10 data set. We take a part of the samples as labeled training data and remove the label of the other samples. Therefore, if we train a classifier to classify airplanes and cars, then (i) the unlabeled airplane and car pictures can be viewed as ideal unlabeled samples, and (ii) the unlabeled samples from other classes are viewed as from a poor generator. Our empirical results verify (1) and (2). Besides, we also find that the unlabeled data with the pseudo label can improve the adversarial training as much as labeled data of the same sample size (the middle panel of Figure C.1). This observation further verifies the effectiveness of unlabeled data and implies that the label cost is small in adversarial training.

## 5 How to Decide a Proper $w$

To determine a proper  $w$ , it is possible to conduct cross validation for simple models. However, due to the heavy computation of adversarial training under DNN models, this is infeasible for real applications. Our proposed algorithm in below aims to simplify the tuning process of  $w$ .

To design a proper metric to tune  $w$ , denote  $b$  as the average adversarial loss difference in  $S_1$  and  $S_2$ ,  $v_1$  and  $v_2$  as the variance of the adversarial loss for  $S_1$  and  $S_2$  respectively. Since in real practice one can often obtain a high-quality clean classifier (i.e., low label cost), we tune  $w$  based on the generator cost only. Inspired by the variance-bias trade-off in (5), we define a surrogate of the generator loss as

$$\kappa(b, v_1, v_2, n_1, n_2, w, w_{\text{bias}}) = w_{\text{bias}} \frac{b^2 (wn_2)^2}{(n_1 + wn_2)^2} + \frac{v_1 n_1 + w^2 v_1 n_2}{(n_1 + wn_2)^2}.$$

Due to over-parameterization, the mean and variance of the gradient are volatile in neural networks. Hence, instead of using the gradient of loss (i.e.,  $\partial l_\epsilon / \partial \theta_\epsilon$ ) as suggested by (5), we use the adversarial loss to construct  $\kappa$  and introduce  $w_{\text{bias}}$  to balance  $b$  and  $(v_1, v_2)$ . To control over-fitting problem, we tune  $w$  during the first 20% iterations in the experiment. The algorithm is shown in Algorithm 1.

---

### Algorithm 1 Select $\alpha$ during Training

---

**Input:** Training dataset  $S_1$  and  $S_2$  with size  $n_1$  and  $n_2$ , optimizer OPT, total number of iterations  $T$ , number of iteration per weight  $\tau$ , initial weight  $w_0$ , bias weight  $w_{\text{bias}}$ , decay factor  $\delta$ , number of epochs  $K$  to train weight.  
Initialize the model parameters  $\theta$ , take  $w = w_0$ .  
**for**  $k = 1, \dots, T/\tau$  **do**  
     $b = v_1 = v_2 = 0$ .  
    **for**  $t = 1, \dots, \tau$  **do**  
        Calculate the mean and variance of adversarial loss for the samples from  $S_1$  and  $S_2$  in this batch as  $b_{t1}, b_{t2}, v_{t1}, v_{t2}$ .  
        Update  $b = b * \delta + b_{t2} - b_{t1}$ ,  $v_1 = v_1 * \delta + v_{t1}$ ,  $v_2 = v_2 * \delta + v_{t2}$ .  
        Update  $\theta$  using OPT.  
    **end for**  
    Update  $b = b/(1 - \delta)$ ,  $v_1 = v_1/(1 - \delta)$ ,  $v_2 = v_2/(1 - \delta)$ .  
    **if**  $k \leq K$  **then**  
        Update  $w$  to reduce  $\kappa$ .  
    **end if**  
**end for**  
**Output:**  $\theta$ .

---

In the experiment, we consider binary classification for the CIFAR-10 dataset to classify airplane and car. The implementation for all real-data experiments is modified from Rice et al. (2020)<sup>2</sup>. We take 500 samples from each class as labeled data, i.e.,  $n_1 = 1,000$ . To form  $S_2$ , we sample  $n_2/2$  data from the other 9,000 airplane and car pictures, and  $n_2/2$  data from other classes, i.e., the unlabeled data generator generates both ideal and poor samples. The experiment setups are postponed to Appendix C. We repeat the experiment for 10 times to get the average robust testing accuracy and its standard error. The details of how to tune  $w_{\text{bias}}$  and how to update  $w$  using  $\kappa$  are in Appendix C.

<sup>2</sup>[https://github.com/locuslab/robust\\_overfitting](https://github.com/locuslab/robust_overfitting)

The results are summarized in Table 1. One can see that adjusting  $w$  using Algorithm 1 can always lead to a better performance than the unweighted training and is as good as the best fixed  $w$  obtained by grid search. We found that 75% of the total epochs are needed for each fixed  $w$  to reveal the performance difference among  $w$ 's (Figure C.3), leading to long computing time for grid search of  $w$ . In contrast, our proposed method only needs several epochs to tune  $w_{\text{bias}}$ .

Table 1: Adversarial test accuracy using unweighted training ( $w = 1$ ), Algorithm 1, and best fixed  $w$ . The latter two methods reaches similar performance.

$n_2$	200	2000	4000	18000
$w = 1$	0.7847(0.0041)	0.7612(0.0168)	0.8170(0.0096)	0.8622(0.0057)
Algorithm 1	0.7870(0.0090)	0.7942(0.0126)	0.8407(0.0070)	0.8765(0.0038)
Best fixed $w$	0.7868(0.0091)	0.7957(0.0128)	0.8429(0.0068)	0.8754(0.0027)

We also evaluate the performance of Algorithm 1 using the full CIFAR-10, CIFAR-100, and SVHN with the additional data generated by Gowal et al. (2021). The results are postponed to Appendix C.

## 6 Proof Sketch of Theorem 1

We assume  $\tilde{\theta}(\epsilon)$  is consistent to  $\theta_\epsilon$  and start from the first-order optimality condition of  $\tilde{\theta}(\epsilon)$  and obtain that

$$\begin{aligned}
\mathbf{0} &= \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) + \sum_{S_2} l_\epsilon(x, \hat{y}, \tilde{\theta}(\epsilon)) \right) \tag{8} \\
&= \underbrace{\frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right)}_{:=A_1} - \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \theta_\epsilon) \right) \\
&\quad + \underbrace{\frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, \hat{y}, \tilde{\theta}(\epsilon)) \right)}_{:=A_2} - \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right) \\
&\quad + \underbrace{\frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right)}_{:=A_3} - \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, y, \theta_\epsilon) \right) \\
&\quad + \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \theta_\epsilon) + \sum_{S_2} l_\epsilon(x, y, \theta_\epsilon) \right).
\end{aligned}$$

We observe that, when  $\tilde{\theta}(\epsilon)$  is consistent, i.e.,  $\|\tilde{\theta}(\epsilon) - \theta_\epsilon\| \xrightarrow{P} 0$ , with probability tending to 1,

$$\begin{aligned}
A_1 + A_3 &= \Sigma_\epsilon(\tilde{\theta}(\epsilon) - \theta_\epsilon) + o, \\
A_2 &= \frac{n_2}{n_1 + n_2} \tilde{\Sigma}_\epsilon(\hat{\theta}(0) - \theta_0) + o.
\end{aligned}$$

As a result, rearranging the terms in (8), we can prove Theorem 1. One the other hand, one can also show that  $\|\tilde{\theta}(\epsilon) - \theta_\epsilon\| \xrightarrow{P} 0$  because  $\tilde{R}(\theta, \epsilon) \rightarrow R(\theta, \epsilon)$  for any reasonable  $\theta$ .

## 7 Conclusion

This paper studies how adversarial training benefits from unlabeled (generated) data. We show that (i) the label cost of  $\tilde{\theta}(\epsilon)$  is small in adversarial training; (ii) with a high-quality data generator, a large  $S_2$  leads to a negligible generator cost. These two facts together indicate that adversarial training benefits a lot from unlabeled (generated) data. Motivated by these observations and our theory, we balance the weights between  $S_1$  and  $S_2$  to improve the performance and propose an algorithm to determine the weight automatically during training.

## Acknowledgements

This project is partially supported by NSF-SCALE MoDL (2134209) and ONR N00014-22-1-2680.

## References

- Alayrac, J.-B., Uesato, J., Huang, P.-S., Fawzi, A., Stanforth, R., and Kohli, P. Are labels required for improving adversarial robustness? *Advances in Neural Information Processing Systems*, 32, 2019.
- Allen-Zhu, Z. and Li, Y. Feature purification: How adversarial training performs robust deep learning. *arXiv preprint arXiv:2005.10190*, 2020.
- Cai, Q.-Z., Du, M., Liu, C., and Song, D. Curriculum adversarial training. *arXiv preprint arXiv:1805.04807*, 2018.
- Cai, T. T., Zhang, C.-H., and Zhou, H. H. Optimal rates of convergence for covariance matrix estimation. *The Annals of Statistics*, 38(4):2118–2144, 2010.
- Carmon, Y., Raghunathan, A., Schmidt, L., Duchi, J. C., and Liang, P. S. Unlabeled data improves adversarial robustness. In *Advances in Neural Information Processing Systems*, pp. 11192–11203, 2019.
- Chen, L., Min, Y., Zhang, M., and Karbasi, A. More data can expand the generalization gap between adversarially robust and standard models. In *International Conference on Machine Learning*, pp. 1670–1680. PMLR, 2020.
- Croce, F., Andriushchenko, M., Sehwag, V., Debenedetti, E., Flammarion, N., Chiang, M., Mittal, P., and Hein, M. Robustbench: a standardized adversarial robustness benchmark. *arXiv preprint arXiv:2010.09670*, 2020.
- Dan, C., Wei, Y., and Ravikumar, P. Sharp statistical guarantees for adversarially robust gaussian classification. In *International Conference on Machine Learning*, pp. 2345–2355. PMLR, 2020.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. In *3rd International Conference on Learning Representations*, 2015.
- Gowal, S., Rebuffi, S.-A., Wiles, O., Stimberg, F., Calian, D. A., and Mann, T. A. Improving robustness using generated data. *Advances in Neural Information Processing Systems*, 34, 2021.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- Hendrycks, D., Lee, K., and Mazeika, M. Using pre-training can improve model robustness and uncertainty. *arXiv preprint arXiv:1901.09960*, 2019.
- Javanmard, A. and Soltanolkotabi, M. Precise statistical analysis of classification accuracies for adversarial training. *arXiv preprint arXiv:2010.11213*, 2020.
- Javanmard, A., Soltanolkotabi, M., and Hassani, H. Precise tradeoffs in adversarial training for linear regression. In *Conference on Learning Theory*, pp. 2034–2078. PMLR, 2020.
- Koltchinskii, V. et al. Local rademacher complexities and oracle inequalities in risk minimization. *The Annals of Statistics*, 34(6):2593–2656, 2006.
- Lee, D.-H. et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, pp. 896, 2013.
- Min, Y., Chen, L., and Karbasi, A. The curious case of adversarially robust models: More data can help, double descend, or hurt generalization. *arXiv preprint arXiv:2002.11080*, 2020.
- Najafi, A., Maeda, S.-i., Koyama, M., and Miyato, T. Robustness to adversarial perturbations in learning from incomplete data. In *Advances in Neural Information Processing Systems*, pp. 5542–5552, 2019.

- Raghunathan, A., Xie, S. M., Yang, F., Duchi, J. C., and Liang, P. Adversarial training can hurt generalization. *arXiv preprint arXiv:1906.06032*, 2019.
- Rice, L., Wong, E., and Kolter, J. Z. Overfitting in adversarially robust deep learning. *arXiv preprint arXiv:2002.11569*, 2020.
- Sehwag, V., Mahloujifar, S., Handina, T., Dai, S., Xiang, C., Chiang, M., and Mittal, P. Improving adversarial robustness using proxy distributions. *arXiv preprint arXiv:2104.09425*, 2021.
- Sinha, A., Namkoong, H., Volpi, R., and Duchi, J. Certifying some distributional robustness with principled adversarial training. *arXiv preprint arXiv:1710.10571*, 2017.
- Sun, W., Tseng, T.-L. B., Zhang, J., and Qian, W. Enhancing deep convolutional neural network scheme for breast cancer diagnosis with unlabeled data. *Computerized Medical Imaging and Graphics*, 57:4–9, 2017.
- Taheri, M., Xie, F., and Lederer, J. Statistical guarantees for regularized neural networks. *Neural Networks*, 142:148–161, 2021.
- Uesato, J., Alayrac, J.-B., Huang, P.-S., Stanforth, R., Fawzi, A., and Kohli, P. Are labels required for improving adversarial robustness? *arXiv preprint arXiv:1905.13725*, 2019.
- Van Der Vaart, A. W. and Wellner, J. A. Weak convergence. In *Weak convergence and empirical processes*, pp. 16–28. Springer, 1996.
- Wang, Y., Ma, X., Bailey, J., Yi, J., Zhou, B., and Gu, Q. On the convergence and robustness of adversarial training. In *International Conference on Machine Learning*, pp. 6586–6595, 2019a.
- Wang, Y., Zou, D., Yi, J., Bailey, J., Ma, X., and Gu, Q. Improving adversarial robustness requires revisiting misclassified examples. In *International Conference on Learning Representations*, 2019b.
- Wu, B., Chen, J., Cai, D., He, X., and Gu, Q. Does network width really help adversarial robustness? *arXiv preprint arXiv:2010.01279*, 2020a.
- Wu, D., Wang, Y., and Xia, S.-t. Adversarial weight perturbation helps robust generalization. *arXiv preprint arXiv:2004.05884*, 2020b.
- Xiao, J., Fan, Y., Sun, R., and Luo, Z.-Q. Adversarial rademacher complexity of deep neural networks. 2021.
- Xie, C. and Yuille, A. Intriguing properties of adversarial training at scale. *arXiv preprint arXiv:1906.03787*, 2019.
- Xing, Y., Song, Q., and Cheng, G. On the generalization properties of adversarial training. In *International Conference on Artificial Intelligence and Statistics*, pp. 505–513. PMLR, 2021a.
- Xing, Y., Zhang, R., and Cheng, G. Adversarially robust estimate and risk analysis in linear regression. In *International Conference on Artificial Intelligence and Statistics*, pp. 514–522. PMLR, 2021b.
- Xing, Y., Song, Q., and Cheng, G. Unlabeled data helps: minimax analysis and adversarial robustness. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022.
- Yin, D., Ramchandran, K., and Bartlett, P. Rademacher complexity for adversarially robust generalization. *arXiv preprint arXiv:1810.11914*, 2018.
- Zhai, R., Cai, T., He, D., Dan, C., He, K., Hopcroft, J., and Wang, L. Adversarially robust generalization just requires more unlabeled data. *arXiv preprint arXiv:1906.00555*, 2019.
- Zhang, H., Yu, Y., Jiao, J., Xing, E. P., Ghaoui, L. E., and Jordan, M. I. Theoretically principled trade-off between robustness and accuracy. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 7472–7482. PMLR, 2019.

Zhang, J., Xu, X., Han, B., Niu, G., Cui, L., Sugiyama, M., and Kankanhalli, M. Attacks which do not kill training make adversarial learning stronger. In *International Conference on Machine Learning*, pp. 11278–11287. PMLR, 2020a.

Zhang, Y., Plevrakis, O., Du, S. S., Li, X., Song, Z., and Arora, S. Over-parameterized adversarial training: An analysis overcoming the curse of dimensionality. *arXiv preprint arXiv:2002.06668*, 2020b.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
  - (b) Did you describe the limitations of your work? [Yes]
  - (c) Did you discuss any potential negative societal impacts of your work? [No] Our work is a theoretical study, and it studies the existing adversarial training method other than introducing new applications.
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
  - (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] We are using the code of Rice et al. (2020) to do experiments, and we also provide the code for CIFAR-10 experiment with data from Goyal et al. (2021).
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] In Section C.
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] In Section C.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? [Yes]
  - (b) Did you mention the license of the assets? [N/A]
  - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
  - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A]
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

The following is the list of contents for the appendix:

- Section A: limitations of this work.
- Section B: simulation study.
- Section C: deep learning experiments.
- Section D: some basic results in empirical process.
- Section E: discussions about Assumptions A1 and A2.
- Section F: discussions about the effect of  $n_2/n_1$  on the label cost and the generator cost.
- Section G: proof of Theorem 1 and other examples/propositions.
- Section H: derivations for linear regression under Gaussian design.

## A Limitations of this Work

In this study, we aim to provide theoretical justification on why adversarial training can benefit more from generated data than clean training, which potentially improves the adversarial robustness of machine learning methods. We mainly focus on the intuitions, theories in simple statistical models, and numerical verification.

Major limitations of the scope of this work include: (1) We focus on justifying how the quality of the data generator affects the final adversarial robustness, rather than how to design better generators, while designing better generators is an interesting and important practical problem; (2) Our theoretical investigation only focuses on simple models. For deep neural networks that are used in real data sciences, we only provide empirical justifications, and an exact theory of how neural networks balances the label cost and generator cost is not currently available.

Regarding social impacts, our paper mainly focuses on the theories and uses public data sets in our real-data experiments. We are not aware of any direct negative social impact by our results. However, we still suggest people be cautious when applying any machine learning algorithm to real applications, e.g., auto-drive, in which safety is essential. Based on our knowledge, there is only limited literature interpreting how machine learning algorithms work.

## B Simulation Study

In the simulation, we aim to numerically verify: (1) given the ideal data generator, the performance of  $\tilde{\theta}(\epsilon)$  is better than  $\hat{\theta}(\epsilon)$  when  $\epsilon$  deviates from zero; (2) the better quality of the data generator implies the better performance of  $\tilde{\theta}(\epsilon)$ ; and (3) balancing the weight between  $S_1$  and  $S_2$  improves the performance.

We first present the results and analysis under  $\mathcal{L}_2$  attack, and then present the corresponding  $\mathcal{L}_\infty$  attack version in a later section.

### B.1 Simulations under $\mathcal{L}_2$ Attack

The simulation model is as follows. Assume  $X \sim N(0, I_{10})$ ,  $\theta_0 = [1/\sqrt{10}, \dots, 1/\sqrt{10}]^\top$ , and the response is generated via  $Y = \theta_0^\top X + \varepsilon$  where the noise  $\varepsilon$  follows  $N(0, 0.1)$ . The unlabeled data are generated from  $N(0, \Sigma)$ , where  $\Sigma$  can be  $I_{10}$ ,  $3I_{10}$ ,  $\hat{\Sigma}_{\text{sparse}}$  in Example 2 or sample covariance matrix  $\hat{\Sigma}$ . We take  $\mathcal{L}_2$  attack,  $n_1 = 50$ , and repeat the experiment for 100 times to obtain the MSE of  $\tilde{\theta}(\epsilon)$ .

We first verify (1) and (2) when the data generator is not related to  $S_1$ . The results are shown in Figure B.1. In the left panel, we use the ideal data generator and change the value of  $\epsilon$ . With a larger attack, it is more likely that adding more data will help reduce the MSE of  $\tilde{\theta}(\epsilon)$ . In the right panel, when using a poor generator  $N(0, 3I_d)$  to generate  $S_2$ , the reduction of MSE of  $\tilde{\theta}(\epsilon)$  is limited. In addition, when  $\epsilon = 0.5$ , the best performance is obtained when  $n_2$  is a finite value. This verifies the observation in Proposition 2.

For the case where the data generator is learned from  $S_1$ , the simulation results are summarized in Figure B.2. As discussed in Example 2, the performance of  $\hat{\Sigma}_{\text{sparse}}$  should be theoretically better than

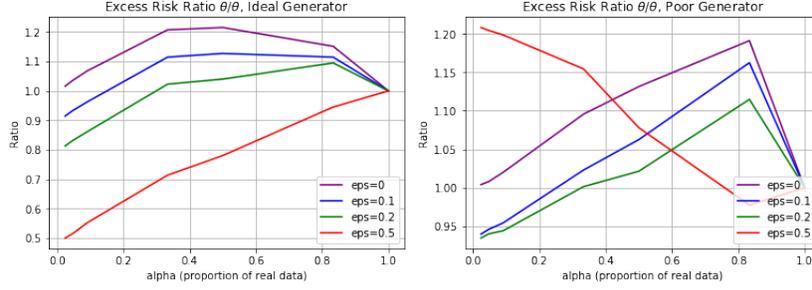


Figure B.1: Simulation: how the attack strength, quality of data generator, and the proportion of real data ( $\alpha$ ) affect the convergence of the adversarial estimator. Left: ideal generator. Right: other generator.

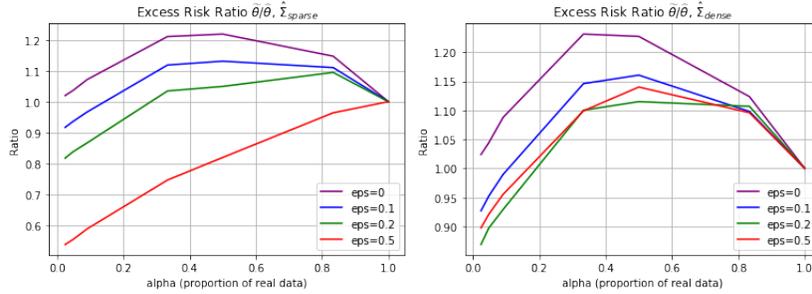


Figure B.2: Simulation: how the attack strength, quality of data generator, and the proportion of real data ( $\alpha$ ) affect the convergence of the adversarial estimator.

$\hat{\Sigma}$ . Figure B.2 and Figure B.1 together show that the ideal generator leads the MSE largest reduction, the generator with  $\hat{\Sigma}_{\text{sparse}}$  has a slightly worse performance, and the generator with  $\hat{\Sigma}$  cannot work well. These observations match our theoretical investigations.

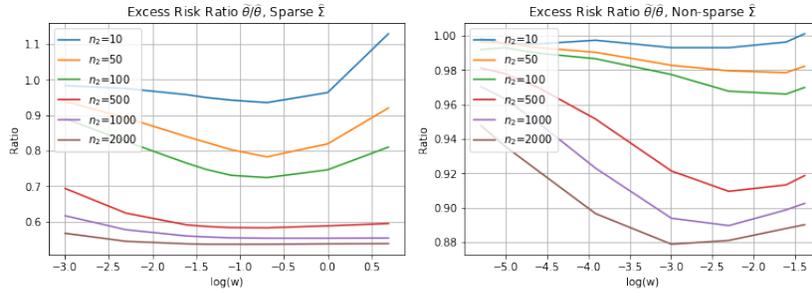


Figure B.3: Simulation: how the weight and quality of data generator, and  $n_2$  affect the convergence of the adversarial estimator. We take  $\epsilon = 0.5$ .

In terms of (3) adjusting the weight, we fix  $\epsilon = 0.5$  and compare the performance while varying  $w$ . The results are summarized in Figure B.3. When the data generator is of high quality (the left panel), the best  $w$  is around 1. When the data generator is not good (the right panel), the best  $w$  has larger deviation from 1. In addition, from both cases in Figure B.3, the minimal MSE ratio is smaller than 1 and keeps decreasing when enlarging  $n_2$ , which verifies Proposition 3.

## B.2 Simulations under $\mathcal{L}_\infty$ Attack

We repeat the experiments in Section B.1 using  $\mathcal{L}_\infty$  attack. Since  $\mathcal{L}_\infty$  attack is stronger, we take slightly smaller  $\epsilon$ . From Figures B.4, B.5, B.6, all the observations are the same as in Section B.1.

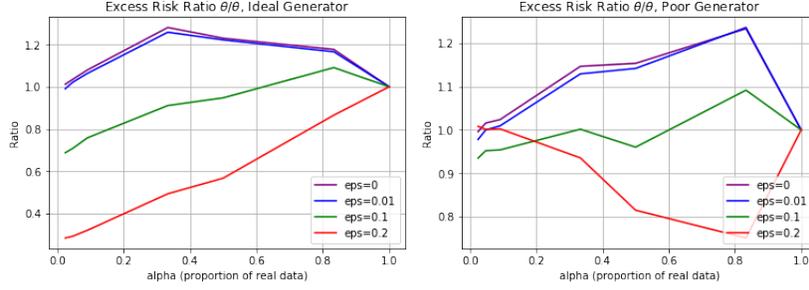


Figure B.4:  $\mathcal{L}_\infty$  version of Figure B.1: how the attack strength, quality of data generator, and the proportion of real data ( $\alpha$ ) affect the convergence of the adversarial estimator. Left: ideal generator. Right: other generator.

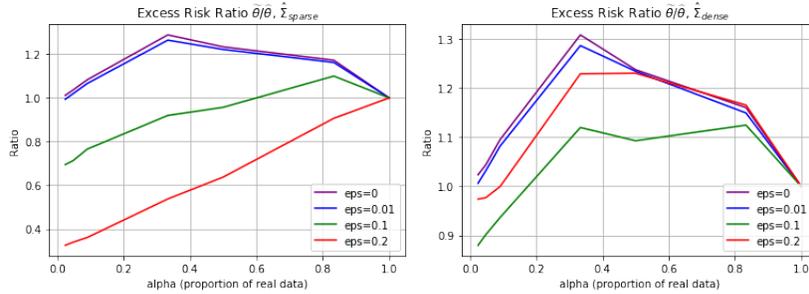


Figure B.5:  $\mathcal{L}_\infty$  version of Figure B.2: how the attack strength, quality of data generator, and the proportion of real data ( $\alpha$ ) affect the convergence of the adversarial estimator.

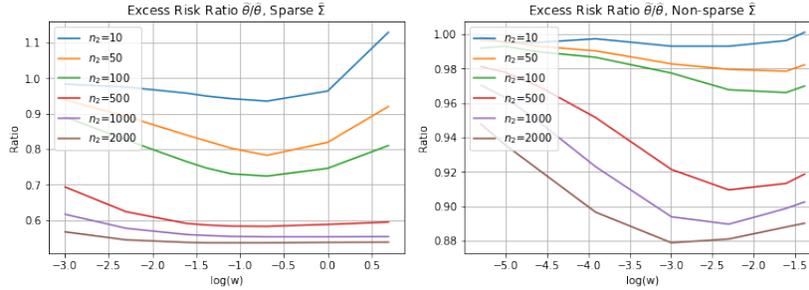


Figure B.6:  $\mathcal{L}_\infty$  version of Figure B.3: how the weight and quality of data generator, and  $n_2$  affect the convergence of the adversarial estimator. We take  $\epsilon = 0.2$ .

## C Deep Learning Experiments

### C.1 General Configurations

We use the CIFAR-10 dataset for this experiment. We pick two classes, airplanes and cars, to do a binary classification. Each class in CIFAR-10 contains 5,000 training samples. We randomly pick 500 from each class as labeled training data. For unlabeled data, we (i) sample from remaining 4,500\*2 samples of airplanes and cars as an ideal data generator or (ii) sample from the images of the other 8 classes as a poor generator. We repeat this experiment 10 times to obtain a boxplot for each setup. We use a single NVIDIA Tesla V100 GPU to do the training.

## C.2 Quality of Data Generator Matters

In Section 4.1, we decompose the error of  $\tilde{\theta}(\epsilon)$  into a label cost and a generator cost. We want to examine whether these two errors are the key factors in deep learning.

We vary the value of  $n_2$  and expect that (1) adding more unlabeled samples from the idea generator will help improve adversarial robustness; (2) when adding unlabeled sample from the poor generator, a small  $n_2$  will lead to slight improvement. For both adversarial training and clean training, the learning rate is divided by ten at the 50% and 75% epochs. For each  $(n_1, n_2)$ , we take the total number of iterations as  $T = 100(n_1 + n_2)/128$ . Although small  $n_2$  leads to fewer training iterations, some simple trials indicate that this choice of  $T$  is sufficient to get the best performance.

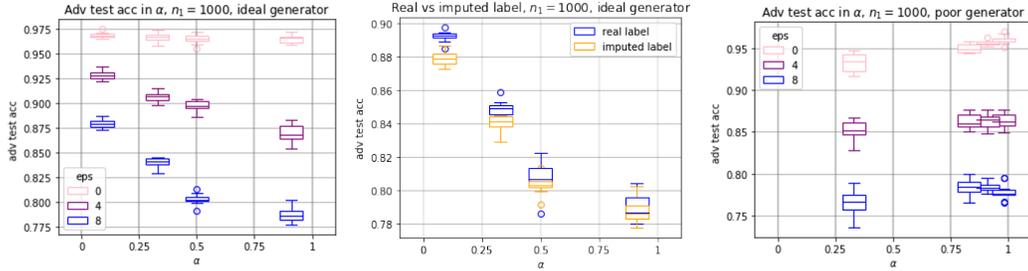


Figure C.1: Adversarial testing accuracy of the neural network trained from  $500 \times 2$  labeled samples with additional unlabeled data. The y-axis refers to adversarial testing accuracy for  $\epsilon = 0$  for clean training,  $4/255$  or  $8/255$  for the corresponding adversarial training. Left: ideal generator. Middle: ideal generator using real/imputed label for  $S_2$ ,  $8/255 \mathcal{L}_\infty$  attack. Right: poor generator.

The results are shown in Figure C.1. For (1), from the left figure of Figure C.1, when  $n_2$  gets larger, i.e.  $\alpha$  gets smaller, for adversarial training with  $\epsilon = 4/255, 8/255$ , the adversarial robustness of the model gets improved. Clean training also benefits from  $S_2$ , but not as much as adversarial training.

In addition, the middle figure in Figure C.1 compares the performance of using  $n_1$  labeled samples and  $n_2$  unlabeled samples, against the performance of using  $n_1 + n_2$  labeled samples (we still take  $\alpha = n_1/(n_1 + n_2)$  to match the x-axis). Using extra unlabeled data improves the performance significantly as  $n_2$  increases. Knowing the true labels of the unlabeled data can further improve the performance, but this additional improvement is only marginal. These observations also verify the hypothesis in Gowal et al. (2021) that it is critical to complement the original training data set with additional (unlabeled) data in adversarial training.

In terms of (2), from the right figure of Figure C.1, for clean training and  $4/255 \mathcal{L}_\infty$  adversarial training, the testing performance does not improve. Only when  $\epsilon = 8/255$ , the adversarial robustness slightly gets better when  $\alpha \in (0.75, 1)$ . Based on our theory, adversarial training is more likely to benefit from the poor generator when  $\epsilon$  is large and  $n_2$  is small, and our experiment observations match this result.

## C.3 Weighted Sampling

We examine the performance of weighted sampling in adversarial training when  $n_1 = 500$ . For  $S_2$ , we randomly select  $n_2/2$  from the remaining samples whose label is airplane/car, and  $n_2/2$  from those with other labels.

**How to implement weighted sampling** Instead of performing a weighted optimization (7), equivalently, we use stratified sampling to sample a batch of  $S_1$  and  $S_2$ , such that the proportion of  $S_1$  samples in the batch is  $n_1/(n_1 + wn_2)$ , and perform SGD for the unweighted objective over the random batch. This helps reduce the variance of the updating gradient. We take  $n_2 \in \{200, 2000, 4000, 18000\}$  and try different  $w$  for each  $n_2$ . The results are shown in Figure C.2. When changing  $w$ ,  $S_1$  is larger, leading to more batches in each epoch. We keep the total number of iterations unchanged as  $T = 100(n_1 + n_2)/128$ .

**Performance** Based on the left panel of Figure C.2, through tuning the weight  $w$ , one can obtain a better adversarial test accuracy.

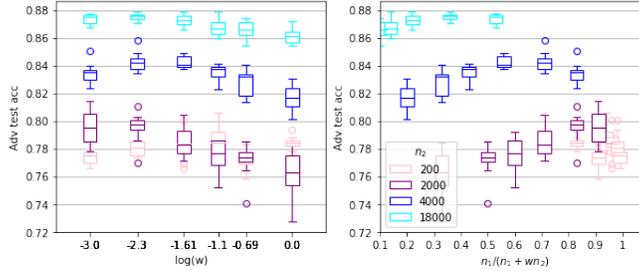


Figure C.2: How  $w$  affects the testing robust accuracy. The generator outputs half true samples and half samples from the other 8 classes in CIFAR-10. The x-axis of the left panel is  $\log(w)$ . The x-axis of the right panel is  $n_1/(n_1 + wn_2)$ , corresponding to the definition of  $\alpha$  in Gowal et al. (2021).

**Configurations in Algorithm 1** We initialize  $w$  as  $w_0 = n_1/n_2$  and take  $\delta = 0.9$  in the experiments. We check  $\kappa$  after every epoch and update  $1/w$  with a magnitude of  $1/(4w_0)$ , i.e., in the next epoch, we take  $1/w_{\text{next}}$  as  $1/w$ ,  $1/w + 1/(4w_0)$ , or  $\max(0, 1/w - 1/(4w_0))$  depending on which one gives the smallest  $\kappa$ . Table 2 is the  $w_{\text{bias}}$  we take in experiments. To choose  $w_{\text{bias}}$ , we try five to six different candidates to pick the one that is most stable in the first 20% iterations.

Table 2: The choice of  $w_{\text{bias}}$  and the resulting final  $w$  for the experiment w.r.t Algorithm 1.

$n_2$	200	2000	4000	18000
$w_{\text{bias}}$	5	5	20	50

#### C.4 Full Real Datasets

For the generated data, we download from Gowal et al. (2021)<sup>3</sup>. The implementations of the training, attack, and neural network architectures are the same as Rice et al. (2020). Detailed configurations can be found in Table 4.

The results are summarized in Table 3. For all the data sets and neural network architectures, the proposed algorithm leads to better performance than the vanilla adversarial training and  $w = 1$ .

On the other hand, for CIFAR-10, compared to the  $>10\%$  improvement for binary classification, there is only 3% to 6% improvement in the robust test accuracy. There are two possible explanations. First, the model capacity of the neural network architectures may not be sufficient to learn the ten classes. In our experiment as well as Gowal et al. (2021); Xie & Yuille (2019), we all observe that

<sup>3</sup>[https://github.com/deepmind/deepmind-research/tree/master/adversarial\\_robustness](https://github.com/deepmind/deepmind-research/tree/master/adversarial_robustness)

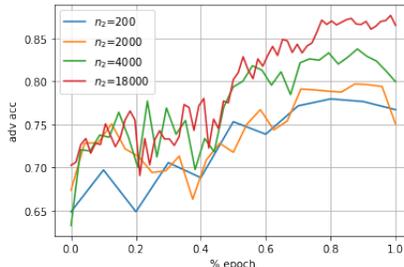


Figure C.3: Adversarial testing accuracy of one repetition when  $w = 1/5$ . The accuracy grows to the highest level at around 75% total iterations for all choices of  $n_2$ .

wider neural networks lead to better robustness. Second, we are using ideal (real) data in the binary classification but generated data in the ten-class classification task. The quality of generated data deteriorates the performance.

Table 3: How unlabeled data help adversarial training. Vanilla: vanilla adversarial training( $\hat{\theta}(\epsilon)$ ) without generated data, copied from Rice et al. (2020). Algorithm 1 gives better performance than the vanilla adversarial training. Similar to Goyal et al. (2021), taking  $w = 1$  possibly leads to poor performance.

Dataset	Architecture	Method	Adv Acc
CIFAR-10	PreActResNet18	Vanilla	0.5349
		$w = 1$	0.5549
		Algorithm 1	<b>0.5625</b>
CIFAR-10	WideResNet34-10	Vanilla	0.5680
		$w = 1$	0.6029
		Algorithm 1	<b>0.6203</b>
SVHN	PreActResNet18	Vanilla	0.6104
		$w = 1$	0.6147
		Algorithm 1	<b>0.6179</b>
CIFAR-100	PreActResNet18	Vanilla	0.2780
		$w = 1$	0.2793
		Algorithm 1	<b>0.3071</b>

Table 4: Training configurations corresponding to Table 3.

Dataset	Batch size	Clean lr	Clean iteration	Adv lr	Adv iteration	$w_{bias}$
CIFAR-10	128	0.1	200*50000/128	0.1	same	1000
SVHN	128	0.1	100*73257/128	0.01	same	2000
CIFAR-100	128	0.1	200*50000/128	0.1	same	20000

## D Results on Empirical Process Theory

In this section, we present some results in Koltchinskii et al. (2006) which will be used in our proof.

Let  $\sigma^2 = \sup_{f \in \mathcal{F}} \mathbb{E} f^2$ , and  $|f(x)| \leq F(x) \leq U$ . Define  $P_n$  as the empirical measure and  $\|P_n - P\|_{\mathcal{F}} = \sup_{f \in \mathcal{F}} |P_n f - P f|$ . If for some  $A > 0, V > 0$ ,

$$\forall \epsilon > 0, \quad N(\epsilon, \|\cdot\|_{\mathcal{L}_2(P_n)}, \mathcal{F}) \leq \left( \frac{A \|F\|_{\mathcal{L}_2(P_n)}}{\epsilon} \right)^V,$$

then (2.2) in Koltchinskii et al. (2006) states that with some universal constant  $C > 0$ ,

$$\mathbb{E} \|P_n - P\|_{\mathcal{F}} \leq C \left[ \left( \frac{V}{n} \right)^{1/2} \sigma \left( \log \frac{A \|F\|_{\mathcal{L}_2(P)}}{\sigma} \right)^{1/2} \vee \frac{VU}{n} \log \frac{A \|F\|_{\mathcal{L}_2(P)}}{\sigma} \right]. \quad (9)$$

Furthermore, a refined version of Talagrand’s concentration inequality states that for any countable class of measurable functions  $\mathcal{F}$  with elements mapping into some  $[-M, M]$ . Combining Bousquet’s bound and Klein’s bound together,

$$P \left\{ \left| \|P_n - P\|_{\mathcal{F}} - \mathbb{E} \|P_n - P\|_{\mathcal{F}} \right| \geq \left( \frac{2t}{n} (\sigma^2 + 2\mathbb{E} \|P_n - P\|_{\mathcal{F}}) \right)^{1/2} + \frac{8t}{3n} \right\} \leq e^{-t}. \quad (10)$$

## E How Does Assumptions A1 and A2 Affect the Theory

Assumptions A1 and A2 are imposed to regulate the behavior of the loss function, gradient, and Hessian. The following discussion provides some details on how to possibly relax the conditions and what are the potential observations if some conditions are violated.

## E.1 Assumption A1

Assumption A1 is a sufficient condition for Theorem 1, Lemma 1, 2, 3, and Example 2. It can be relaxed to general distributions if the following conditions are satisfied. For some  $\epsilon_0 > 0$ ,  $r > 0$ , for any  $\epsilon \in [0, \epsilon_0]$  and  $\theta \in B_2(0, r)$ ,

- With probability one, the loss function  $l_\epsilon$  is smooth, and the gradient and Hessian exist.
- The loss function  $l_\epsilon$  is convex.
- The loss function  $l_\epsilon$  has a finite expectation and finite variance.
- The gradient  $\partial l_\epsilon(X, Y, \theta)/\partial\theta$  has finite expectation and finite variance.
- The eigenvalues of the expected Hessian  $\mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} \partial^2 l_\epsilon(X, Y, \theta)/\partial\theta^2$  are finite and bounded away from zero.
- Considering the expectation over  $Y$  as a function of  $\theta_0$ ,  $\mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta)$  is smooth in  $\theta_0$ , the gradient  $\partial \mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta)/\partial\theta_0$  is finite, and  $\partial^2 \mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta)/\partial\theta_0 \partial\theta$  has non-zero singular values.

## E.2 Assumption A2

Assumption A2 is essential because the proof in Theorem 1 is based on the fact that  $\tilde{\theta}(\epsilon)$  is consistent to  $\theta_\epsilon$ .

## E.3 If the Assumptions Do Not Hold

**Multi-class classification, or other generalized linear models.** In general, as long as the relationship between  $Y$  and  $X$  is well specified, and the loss function is convex and smooth enough, Theorem 1 is valid. In addition, the intuitions of label cost and generator cost in Sections 4.2 and 4.3 still apply.

**The distribution  $\mathcal{P}_0$  has a heavy tail.** It depends on the design of the loss function. If  $l_\epsilon$  and its gradient and Hessian are still well-behaved, we still have the convergence. Otherwise,  $\hat{\theta}(0)$  may not be consistent, i.e., it is hard even for the clean model estimation.

**The parameterization between  $Y$  and  $X$  is misspecified.** In this case,  $\theta_0$  is a biased model to describe  $P(Y|X = x)$ . If  $l_\epsilon$  is smooth enough, the results in Theorem 1 will still hold. However, it is more important to deal with the bias in this case.

**Assumption A2 is violated.** In this case, Theorem 1 may fail because  $\tilde{\theta}(\epsilon)$  may not be consistent to  $\theta_\epsilon$ .

## F How $n_2/n_1$ Affects the Performance

In Figure 2, we explain how the quality of the generator affects the generator cost when  $n_2 \rightarrow \infty$ . In this section, we discuss how  $n_2/n_1$  affects the label cost and generator cost.

Based on the formula of  $E_1$  and  $E_2$ , when  $n_2$  increases,  $\mathbb{E}E_1^2$  is a monotone increasing function, and  $\mathbb{E}E_2^2$  could even decrease or increase depending on the quality of the generator. Figures F.1 and F.2 show how the costs are changed with an ideal generator or a poor generator.

## G Proofs

### G.1 Theorem 1

Since  $\tilde{\theta}(\epsilon)$  minimizes  $\tilde{R}(\theta, \epsilon)$  as in (3), we have

$$\frac{1}{n_1 + n_2} \frac{\partial}{\partial\theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) + \sum_{S_2} l_\epsilon(x, \hat{y}, \tilde{\theta}(\epsilon)) \right) = \mathbf{0},$$

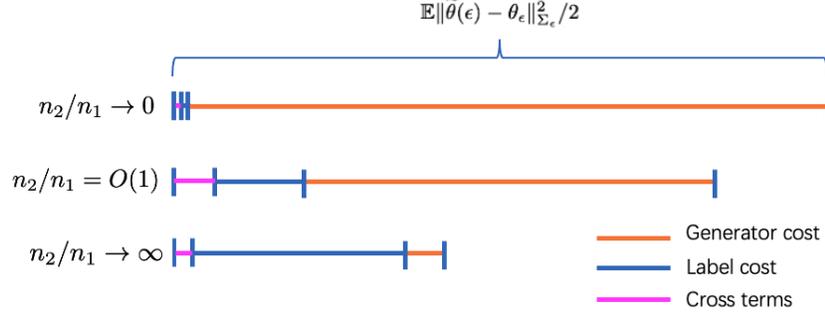


Figure F.1: How  $n_2/n_1$  changes generator cost and label cost when using an **ideal generator**. With  $n_2/n_1$  increases, the generator cost becomes smaller and smaller because more samples lead to a smaller variance in  $E_2$ . The label cost will slightly increase.

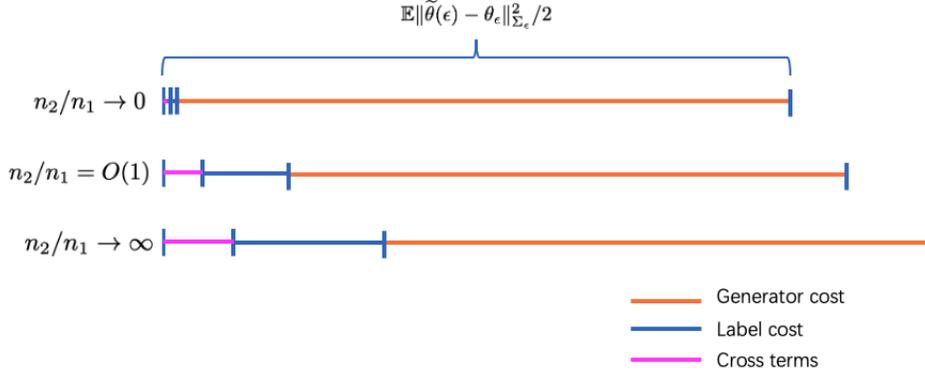


Figure F.2: How  $n_2/n_1$  changes generator cost and label cost when using a **poor generator**. With  $n_2/n_1$  increases, since  $E_2$  is biased away from zero, the generator cost cannot be efficiently reduced even if large  $n_2$  leads to a small variance in  $E_2$ .

using which we can get the following decomposition,

$$\begin{aligned}
\mathbf{0} &= \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) + \sum_{S_2} l_\epsilon(x, \hat{y}, \tilde{\theta}(\epsilon)) \right) \tag{11} \\
&= \underbrace{\frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right)}_{:=A_1} - \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \theta_\epsilon) \right) \\
&\quad + \underbrace{\frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, \hat{y}, \tilde{\theta}(\epsilon)) \right)}_{:=A_2} - \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right) \\
&\quad + \underbrace{\frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right)}_{:=A_3} - \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, y, \theta_\epsilon) \right) \\
&\quad + \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \theta_\epsilon) + \sum_{S_2} l_\epsilon(x, y, \theta_\epsilon) \right).
\end{aligned}$$

In the following, we present Lemma 1 to show the consistency of  $\tilde{\theta}(\epsilon)$ , then show Lemma 2 and 3 to bound some terms in  $A_1$  to  $A_3$ , and finally present the main proof for Theorem 1.

**Lemma 1** (Consistency of  $\tilde{\theta}(\epsilon)$ ). *Under the same assumptions as Theorem 1,  $\tilde{\theta}(\epsilon)$  converges to  $\theta_\epsilon$  in probability.*

*Proof of Lemma 1.* Denote  $\bar{\theta}_\epsilon$  as the minimizer of  $n_1 \mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta) + n_2 \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta)$ . Denote  $\mathcal{P}$  as the weighted average of  $\mathcal{P}_0$  and  $\mathcal{P}_a$ . Under Assumption A2, we know that  $\|\theta_\epsilon - \bar{\theta}_\epsilon\| \rightarrow 0$ .

For both linear regression and logistic regression, one can check that  $l_\epsilon$  is strongly convex in  $\theta$ . Therefore, we have for some  $C > 0$ , such that

$$\inf_{\|\theta - \bar{\theta}_\epsilon\| > \xi} (\mathbb{E}_{\mathcal{P} \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta) - \mathbb{E}_{\mathcal{P} \otimes \mathcal{P}_y} l_\epsilon(X, Y, \bar{\theta}_\epsilon)) \geq C\xi^2. \quad (12)$$

Furthermore, when  $\|\theta\| \leq b_0$  for some large constant  $B$ , following the same arguments in Lemma 2, there exists some  $C_0 > 0$ ,

$$\begin{aligned} P \left( \sup_{\theta \in B(\mathbf{0}, b_0)} \left| \mathbb{E}_{\mathcal{P} \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta) - \frac{1}{n_1 + n_2} \sum_{S_1, S_2} l_\epsilon(x, y, \theta) \right| \right. \\ \left. \geq C_0 \left( \sqrt{\frac{t}{n_1 + n_2}} + \frac{t}{n_1 + n_2} \right) \right) \leq e^{-t}, \quad (13) \end{aligned}$$

$$P \left( \sup_{\theta \in B(\mathbf{0}, b_0)} \left| \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} l_\epsilon(X, \hat{Y}, \theta) - \frac{1}{n_2} \sum_{S_2} l_\epsilon(x, \hat{y}, \theta) \right| \geq C_0 \left( \sqrt{\frac{t}{n_2}} + \frac{t}{n_2} \right) \right) \leq e^{-t}, \quad (14)$$

and

$$P \left( \sup_{\theta \in B(\mathbf{0}, b_0)} \left| \mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta) - \frac{1}{n_1} \sum_{S_1} l_\epsilon(x, y, \theta) \right| \geq C_0 \left( \sqrt{\frac{t}{n_1}} + \frac{t}{n_1} \right) \right) \leq e^{-t}. \quad (15)$$

From the definition of  $\tilde{\theta}(\epsilon)$ , we know that

$$\frac{1}{n_1 + n_2} \left[ \sum_{S_1} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) + \sum_{S_2} l_\epsilon(x, \hat{y}, \tilde{\theta}(\epsilon)) \right] \leq \frac{1}{n_1 + n_2} \left[ \sum_{S_1} l_\epsilon(x, y, \bar{\theta}_\epsilon) + \sum_{S_2} l_\epsilon(x, \hat{y}, \bar{\theta}_\epsilon) \right],$$

then combining with (14), (15) and the fact that  $\mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} l_\epsilon(X, \hat{Y}, \theta) \xrightarrow{P} \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta)$  ( $\hat{\theta}(0) \xrightarrow{P} \theta_0$ ), we have

$$\begin{aligned} & \frac{1}{n_1 + n_2} \left[ \sum_{S_1} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) + \sum_{S_2} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right] \\ & \leq \frac{1}{n_1 + n_2} \left[ \sum_{S_1} l_\epsilon(x, y, \bar{\theta}_\epsilon) + \sum_{S_2} l_\epsilon(x, y, \bar{\theta}_\epsilon) \right] + o. \quad (16) \end{aligned}$$

Combining (16) and (13), with probability tending to 1,

$$\mathbb{E}_{\mathcal{P} \otimes \mathcal{P}_y} l_\epsilon(X, Y, \tilde{\theta}(\epsilon)) \leq \mathbb{E}_{\mathcal{P} \otimes \mathcal{P}_y} l_\epsilon(X, Y, \bar{\theta}_\epsilon) + o.$$

On the other hand, from the definition of  $\tilde{\theta}_\epsilon$ , we have

$$\mathbb{E}_{\mathcal{P} \otimes \mathcal{P}_y} l_\epsilon(X, Y, \tilde{\theta}(\epsilon)) \geq \mathbb{E}_{\mathcal{P} \otimes \mathcal{P}_y} l_\epsilon(X, Y, \bar{\theta}_\epsilon),$$

thus one can conclude that  $\mathbb{E}_{\mathcal{P} \otimes \mathcal{P}_y} l_\epsilon(X, Y, \tilde{\theta}(\epsilon)) - \mathbb{E}_{\mathcal{P} \otimes \mathcal{P}_y} l_\epsilon(X, Y, \bar{\theta}_\epsilon) \rightarrow 0$ .

Further using (12), we get  $\|\tilde{\theta}(\epsilon) - \bar{\theta}_\epsilon\| \rightarrow 0$  as well, and thus  $\tilde{\theta}(\epsilon)$  converges to  $\theta_\epsilon$ .  $\square$

**Lemma 2.** Denoting

$$\delta_1(x, y, \theta) = \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta) - \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta_\epsilon), \quad (17)$$

under the same assumptions as in Theorem 1, there exists some  $\xi = o(1)$  such that

$$P \left( \sup_{\theta \in B(\theta_\epsilon, \xi)} \left\| \frac{1}{n_1} \sum_{S_1} \delta_1(x, y, \theta) - \mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} \delta_1(X, Y, \theta) \right\| \geq C_{12} \left( \sqrt{\frac{t}{n_1}} \xi + \frac{t}{n_1} \right) \right) \leq e^{-t}.$$

*Proof of Lemma 2.* Since  $\partial l_\epsilon / \partial \theta$  is square integrable, based on Theorem 2.6.7 of Van Der Vaart & Wellner (1996), taking  $\mathcal{F}_1$  as the class of function of  $\delta_1$  for  $\theta \in B(\theta_\theta, \xi)$ , we have for some  $C_{11} > 0$ ,  $V > 0$ ,

$$\forall \epsilon > 0, \quad N(\epsilon, \|\cdot\|_{\mathcal{L}_2(P)}, \mathcal{F}_1) \leq C_{11} \left( \frac{1}{\epsilon} \right)^V.$$

Based on Lemma 1, we know that there exists some  $\xi \rightarrow 0$  such that  $\|\tilde{\theta}(\epsilon) - \theta_\epsilon\| \leq \xi$  in probability. As a result, with probability tending to 1,  $\sup_{\|\theta - \theta_\epsilon\| \leq \xi} \mathbb{E} \|\partial l_\epsilon / \partial \theta\|^2 = o(\xi^2)$ . Following (9) and (10), we know that there exists some positive constant  $C_{12}$  such that

$$P \left( \sup_{\theta \in B(\theta_\epsilon, \xi)} \left\| \frac{1}{n_1} \sum_{S_1} \delta_1(x, y, \theta) - \mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} \delta_1(X, Y, \theta) \right\| \geq C_{12} \left( \sqrt{\frac{t}{n_1}} \xi + \frac{t}{n_1} \right) \right) \leq e^{-t}. \quad \square$$

**Lemma 3.** Define  $r_0$  such that  $r_0 = o(1)$  and  $\|\hat{\theta}(0) - \theta_0\| \leq r_0$  with probability tending to 1. Denote

$$\delta_2(x, \theta, \hat{\theta}(0)) = \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, \hat{y}, \theta) - \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta),$$

then under the same assumptions as Theorem 1, there exists some positive constants  $C_{21}, C_{22}$  such that

$$\begin{aligned} P \left( \sup_{\theta \in B(\theta_\epsilon, \xi), \theta' \in B(\theta_0, r_0)} \left\| \frac{1}{n_2} \sum_{S_2} \delta_2(x, \theta, \theta') - \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \delta_2(X, \theta, \theta') \right\| \right. \\ \left. \geq C_{22} \left( \sqrt{\frac{t}{n_2}} \max(\xi, r_0) + \frac{t}{n_2} \right) \right) \leq e^{-t}. \end{aligned}$$

*Proof of Lemma 3.* The proof of Lemma 3 is the same as Lemma 2. Note that we do not assume  $S_2$  and  $\hat{\theta}(0)$  are independent. Instead, we consider changing both  $\theta$  and  $\hat{\theta}(0)$  in  $\delta_2$  to form the class of functions.  $\square$

*Proof of Theorem 1.* We start from studying  $A_1$  to  $A_3$ . For  $A_1$ ,

$$\begin{aligned} A_1 &= \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right) - \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \theta_\epsilon) \right) \\ &= \underbrace{\frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right) - \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \theta_\epsilon) \right)}_{:=B_{11}} \\ &\quad - \underbrace{\mathbb{E} \left[ \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right) - \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \theta_\epsilon) \right) \right]}_{:=B_{12}} \\ &\quad + \mathbb{E} \left[ \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right) - \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_1} l_\epsilon(x, y, \theta_\epsilon) \right) \right]. \end{aligned}$$

From the definition of  $\Sigma_\epsilon$ , we obtain that

$$B_{12} = \frac{n_1}{n_1 + n_2} \Sigma_\epsilon(\tilde{\theta}(\epsilon) - \theta_\epsilon),$$

and based on Lemma 2, (taking  $t = 2 \log n_1$  in (17)) with probability tending to 1,  $\|B_{11} - B_{12}\| = o(\|B_{12}\|)$ .

In terms of  $A_2$ , we have

$$\begin{aligned} A_2 &= \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, \hat{y}, \theta) \right) - \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right) \\ &= \underbrace{\frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, \hat{y}, \tilde{\theta}(\epsilon)) \right) - \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right)}_{:=B_{21}} \\ &\quad - \underbrace{\mathbb{E} \left[ \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, \hat{y}, \tilde{\theta}(\epsilon)) \right) - \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right) \right]}_{:=B_{22}} \\ &\quad + \mathbb{E} \left[ \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, \hat{y}, \tilde{\theta}(\epsilon)) \right) - \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right) \right]. \end{aligned}$$

Similar to  $A_1$ , from the definition of  $\tilde{\Sigma}_\epsilon$ ,

$$\begin{aligned} B_{22} &= \frac{n_2}{n_1 + n_2} \left[ \frac{\partial^2}{\partial \theta_\epsilon \partial \theta_0} \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_\epsilon} l_\epsilon(X, Y, \tilde{\theta}(\epsilon)) \right] (\hat{\theta}(0) - \theta_0) + o \\ &= \frac{n_2}{n_1 + n_2} \left[ \frac{\partial^2}{\partial \theta_\epsilon \partial \theta_0} \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_\epsilon} l_\epsilon(X, Y, \theta_\epsilon) \right] (\hat{\theta}(0) - \theta_0) + o \\ &= \frac{n_2}{n_1 + n_2} \tilde{\Sigma}_\epsilon(\hat{\theta}(0) - \theta_0) + o, \end{aligned}$$

and from Lemma 3, with probability tending to 1,  $\|B_{21} - B_{22}\| = o(\|B_{22}\|)$ .

For  $A_3$ , similar to  $A_1$ , based on Lemma 2,

$$\begin{aligned} A_3 &= \mathbb{E} \left[ \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, y, \tilde{\theta}(\epsilon)) \right) - \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \left( \sum_{S_2} l_\epsilon(x, y, \theta_\epsilon) \right) \right] + o \\ &= \frac{n_2}{n_1 + n_2} \left[ \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \frac{\partial^2 l_\epsilon(X, Y, \theta_\epsilon)}{\partial \theta_\epsilon^2} \right] (\tilde{\theta}(\epsilon) - \theta_\epsilon) + o. \end{aligned}$$

Under Assumption A2, although  $\mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \partial^2 l_\epsilon(X, Y, \theta_\epsilon) / \partial \theta_\epsilon^2$  is different from  $\Sigma_\epsilon$ , we still have

$$A_1 + A_3 = \Sigma_\epsilon(\tilde{\theta}(\epsilon) - \theta_\epsilon) + o.$$

Finally, taking  $A_1$  to  $A_3$  back to the decomposition of  $\tilde{\theta}(\epsilon) - \theta_\epsilon$ , we conclude that with probability tending to 1,

$$\tilde{\theta}(\epsilon) - \theta_\epsilon = -\Sigma_\epsilon^{-1} \left[ \frac{n_2}{n_1 + n_2} \tilde{\Sigma}_\epsilon(\hat{\theta}(0) - \theta_0) + \frac{1}{n_1 + n_2} \frac{\partial}{\partial \theta_\epsilon} \sum_{S_1, S_2} l_\epsilon(x, y, \theta_\epsilon) \right] + o.$$

□

**Remark 2.** A converge-in-probability result and a Bahadur representation does not directly imply  $\mathcal{L}_2$  convergence, i.e.,  $\tilde{\theta}(\epsilon) \xrightarrow{P} \theta_\epsilon$  does not imply  $\mathbb{E}\|\tilde{\theta}(\epsilon) - \theta_\epsilon\|^2 \rightarrow 0$ . However, based on assumption A1 and A2, due to the choice of  $\mathcal{P}_0$  and the loss functions we take, one can obtain  $\mathbb{E}\|\tilde{\theta}(\epsilon) - \theta_\epsilon\|^2 \rightarrow 0$ . As a result, we can use the terms in the Bahadur representation to discuss the label cost and the generator cost. One can replace A1 and A2 with more general conditions in other models with proper conditions.

## G.2 Proposition 1

The proof mainly follows Xing et al. (2021b). The definition of adversarial risk in Xing et al. (2021b) is different to our scenario. Using our definition, in linear regression setup, we know that as long as  $\sigma^2 (\text{Var}(Y|X))$  and  $\epsilon$  are nonzero, then  $\theta_\epsilon$  will deviate from  $\theta_0$ . And one can obtain that,

$$\theta_\epsilon = (\lambda I_d + \Sigma)^{-1} \Sigma \theta_0,$$

and  $\lambda = \Theta(\epsilon)$  when  $\epsilon \in [0, \epsilon^*]$ . One can take this rate of  $\lambda$  into Lemma 5 of Xing et al. (2021b) to obtain the rate  $\Omega(\epsilon^2 d R^2 / (n_1 + n_2))$ . The other term in the lower bound is obtained from Lemma 4 of Xing et al. (2021b) without any changes.

## G.3 Proposition 2

From Theorem 1,  $\theta_\epsilon - \tilde{\theta}(\epsilon)$  is decomposed into

$$\theta_\epsilon - \tilde{\theta}(\epsilon) = E_1 + E_2 + o,$$

where

$$E_1 = \frac{n_2}{n_1 + n_2} \Sigma_\epsilon^{-1} \tilde{\Sigma}_\epsilon (\hat{\theta}(0) - \theta_0) + o, \quad \text{and } E_2 = \frac{\Sigma_\epsilon^{-1}}{n_1 + n_2} \sum_{S_1, S_2} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta_\epsilon).$$

In terms of the variance of  $\theta_\epsilon - \tilde{\theta}(\epsilon)$ , we know that  $\text{tr}(\text{Var}(E_2)) = O(d/(n_1 + n_2))$ , and  $\text{tr}(\text{Var}(E_1)) = O(d n_2^2 / (n_1 + n_2)^2 n_1)$ . Therefore, the dominant part is  $\text{tr}(\text{Var}(E_2))$ .

In terms of the expectation of  $\theta_\epsilon - \tilde{\theta}(\epsilon)$ , since  $\theta_\epsilon$  is the minimizer of  $\mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta)$ , we have

$$\mathbb{E} \sum_{S_1} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta_\epsilon) = \mathbf{0}.$$

As a result, the bias in  $E_2$  are all contributed by the samples in  $S_2$ . When  $\|\mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon)\| = O(\sqrt{d})$ , we have  $\|\mathbb{E} E_2\| = O(\sqrt{d} n_2 / (n_1 + n_2))$ . Meanwhile, since  $\hat{\theta}(0)$  is unbiased to  $\theta_0$ ,  $\mathbb{E} E_1$  is at most in  $O(\sqrt{d} n_2 / (n_1 + n_2))$ .

## G.4 Proposition 3

The statement of Proposition 3 itself can be viewed as a simple sketch of the proof. We provide some more details below:

Adding  $w$  into Theorem 1, we have

$$\begin{aligned} \theta_\epsilon - \tilde{\theta}(w, n_2) &= \underbrace{\frac{w n_2}{n_1 + w n_2} \Sigma_\epsilon^{-1} \tilde{\Sigma}_\epsilon (\hat{\theta}(0) - \theta_0)}_{:=A(w, n_2)} + \underbrace{\frac{\Sigma_\epsilon^{-1}}{n_1 + w n_2} \left( \sum_{S_1} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta_\epsilon) \right)}_{:=B(w, n_2)} \\ &\quad + \underbrace{\frac{w \Sigma_\epsilon^{-1}}{n_1 + w n_2} \left( \sum_{S_2} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta_\epsilon) \right)}_{:=C(w, n_2)} + o. \end{aligned} \quad (18)$$

$$\begin{aligned} \|\theta_\epsilon - \tilde{\theta}(w, n_2)\|_{\Sigma_\epsilon}^2 &= \|A(w, n_2) + B(w, n_2)\|_{\Sigma_\epsilon}^2 + \|C(w, n_2)\|_{\Sigma_\epsilon}^2 \\ &\quad + \underbrace{2(A(w, n_2) + B(w, n_2))^\top \Sigma_\epsilon C(w, n_2)}_{:=D(w, n_2)} + o. \end{aligned}$$

Denote  $\tilde{\theta}(w, n_2)$  as the minimizer of (7), and  $w^*(n_2)$  minimizes the excess adversarial risk of  $\tilde{\theta}(w, n_2)$  w.r.t.  $w$ .

For  $n'_2 > n_2$ , take  $w'$  such that  $w' n'_2 = w^*(n_2) n_2$ , then

- In (18),  $A(w^*(n_2), n_2) = A(w', n'_2)$  and  $B(w^*(n_2), n_2) = B(w', n'_2)$ .
- Since  $n'_2 > n_2$ , we have  $w' < w^*(n_2)$ , and  $\mathbb{E}\|C(w^*(n_2), n_2)\|_{\Sigma_\epsilon}^2 > \mathbb{E}\|C(w', n'_2)\|_{\Sigma_\epsilon}^2$ .
- Since  $\widehat{\theta}(0)$  is consistent, and data in  $S_1$  are real data, we have  $\mathbb{E}A(w, n_2) + B(w, n_2)$  is negligible, and thus  $\mathbb{E}D(w, n_2)$  is negligible.

As a result, the risk of  $\widetilde{\theta}(w', n'_2)$  is smaller than  $\widetilde{\theta}(w^*(n_2), n_2)$ .

On the other hand, by the definition of  $w^*(\cdot)$ , the risk of  $\widetilde{\theta}(w^*(n'_2), n'_2)$  is further smaller than  $\widetilde{\theta}(w', n'_2)$ .

Eventually, we obtain that a larger  $n_2$  always gives a better  $\widetilde{\theta}$ .

## G.5 Derivation of Example 2

In Theorem 1, we assume the data generator and  $S_1$  are independent, which is different this scenario where  $\widehat{\Sigma}$  is learnt from  $S_1$ . To show that the estimate in this example still converges, there are several steps in the proof.

We firstly bound the difference between  $\mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta_\epsilon)$  and  $\mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta_\epsilon)$ . Denote  $f_A$  as the density of  $(x, y)$  given a covariance matrix  $A$ , then given  $\widehat{\Sigma}$ ,

$$\begin{aligned}
& \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon) \\
&= \int f_{\widehat{\Sigma}}(x, y) \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta_\epsilon) d(x, y) \\
&= \underbrace{\int f_{\widehat{\Sigma}}(x, y) \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta_\epsilon) d(x, y)}_{=0} + \int (f_{\widehat{\Sigma}} - f_{\Sigma})(x, y) \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta_\epsilon) d(x, y) \\
&= \int \left\langle \widehat{\Sigma} - \Sigma, \frac{\partial f_{\Sigma}(x, y)}{\partial \Sigma} \right\rangle \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(x, y, \theta_\epsilon) d(x, y) \\
&= O(\|\widehat{\Sigma} - \Sigma\|).
\end{aligned} \tag{19}$$

Since  $\widehat{\Sigma}$  is consistent,  $\mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon)$  converges to  $\mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta_\epsilon)$  as well, where the latter one is a just zero vector.

Second, to show that  $\widetilde{\theta}(\epsilon)$  is consistent as in Lemma 1, for some constant  $c_0$ , with probability tending to 1 over the generation of  $S_1$ ,

$$P \left( \sup_{\theta \in B(\mathbf{0}, b_0)} \left| \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta) - \frac{1}{n_2} \sum_{S_2} l_\epsilon(x, y, \theta) \right| \geq c_0 \left( \sqrt{\frac{t}{n_2}} + \frac{t}{n_2} \right) \middle| S_1 \right) \leq e^{-t}.$$

Since  $\widehat{\theta}(0) \xrightarrow{P} \theta_0$ , following the same arguments as in Lemma 1, we have

$$\frac{1}{n_1 + n_2} \sum_{S_2} l_\epsilon(x, \widehat{y}, \theta) \rightarrow \frac{1}{n_1 + n_2} \sum_{S_2} l_\epsilon(x, y, \theta) + o.$$

Furthermore,

$$\begin{aligned}
& \left| \mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta) - \frac{1}{n_1 + n_2} \sum_{S_1, S_2} l_\epsilon(x, y, \theta) \right| \\
& \leq \left| \mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta) - \frac{1}{n_1 + n_2} \sum_{S_1} l_\epsilon(x, y, \theta) - \frac{n_2}{n_1 + n_2} \mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta) \right| \\
& \quad + \frac{n_2}{n_1 + n_2} \left| \frac{1}{n_2} \sum_{S_2} l_\epsilon(x, y, \theta) - \mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta) \right| \\
& \leq \underbrace{\left| \mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta) - \frac{1}{n_1 + n_2} \sum_{S_1} l_\epsilon(x, y, \theta) - \frac{n_2}{n_1 + n_2} \mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta) \right|}_{\xrightarrow{P} 0} \\
& \quad + \frac{n_2}{n_1 + n_2} \underbrace{\left| \mathbb{E}_{\mathcal{P}_0 \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta) - \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta) \right|}_{\xrightarrow{P} 0} \\
& \quad + \frac{n_2}{n_1 + n_2} \underbrace{\left| \frac{1}{n_2} \sum_{S_2} l_\epsilon(x, y, \theta) - \mathbb{E}_{\mathcal{P}_a \otimes \mathcal{P}_y} l_\epsilon(X, Y, \theta) \right|}_{\xrightarrow{P} 0} \\
& \xrightarrow{P} 0.
\end{aligned}$$

Thus following the same arguments as in Lemma 1, we get  $\tilde{\theta}(\epsilon) \rightarrow \theta_\epsilon$  in probability.

In terms of the Bahadur representation of  $\tilde{\theta}(\epsilon)$ , from the decomposition in (11), we always deal with  $S_1$  and  $S_2$  separately. The clean estimate  $\hat{\theta}(0)$  is related to  $\mathcal{P}_a$ , but this does not affect Lemma 2 and 3 because their bounds are valid uniformly for proper  $\theta$ . As a result, we obtain the same Bahadur representation for  $\tilde{\theta}(\epsilon) - \theta_\epsilon$ .

## H Linear Regression with Gaussian Design

In this section, we present the derivation of label cost in  $\tilde{\theta}(\epsilon)$  and excess adversarial risk of  $\hat{\theta}(\epsilon)$  in linear regression with Gaussian design.

For Figure 3, we consider  $Y = X^\top \theta_0 + \varepsilon$ ,  $\theta_0 = \mathbf{1}_d / \sqrt{d}$ ,  $X \sim N(\mathbf{0}, I_d)$ ,  $\varepsilon \sim N(0, 0.1^2)$ .

The following is a list of contents in this section:

- Paragraph “formula of  $\theta_\epsilon$ ” describes how  $\theta_\epsilon$  is related to  $\Sigma$  and  $\theta_0$ .
- Paragraph “label cost” and “vanilla adversarial training” present the formula of these two costs under Gaussian design, using which we obtain the theoretical curves in Figure 3.
- Paragraph “simulation” presents the simulation setups to get the dashed curves in Figure 3.

**Formula of  $R(\theta, \epsilon)$**  Assume  $X \sim N(\mathbf{0}, \Sigma)$  and  $Var(\varepsilon) = \sigma^2$ , then in the adversarial loss:

$$l_\epsilon(x, y, \theta) = (x^\top \theta - x^\top \theta_0 - \varepsilon)^2 + \epsilon^2 \|\theta\|^2 + 2\epsilon \|\theta\| |x^\top \theta - x^\top \theta_0 - \varepsilon|.$$

We know that  $X^\top \theta - x^\top \theta_0 - \varepsilon$  is also a Gaussian distribution with zero mean and variance as  $\|\theta - \theta_0\|_\Sigma^2 + \sigma^2$ , so the expectation of its absolute value is  $c_0 \sqrt{\|\theta - \theta_0\|_\Sigma^2 + \sigma^2}$  where  $c_0 = \sqrt{2/\pi}$ . As a result, the population adversarial risk is

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

and its gradient becomes

$$\begin{aligned}\frac{\partial R}{\partial \theta_\epsilon} &= 2\Sigma(\theta - \theta_0) + 2\epsilon^2\theta + 2\epsilon c_0 \frac{\theta}{\|\theta\|} \sqrt{\|\theta - \theta_0\|_\Sigma^2 + \sigma^2} + 2\epsilon c_0 \|\theta\| \frac{\Sigma(\theta - \theta_0)}{\sqrt{\|\theta - \theta_0\|_\Sigma^2 + \sigma^2}} \\ &:= 2(\Sigma(\theta - \theta_0) + \epsilon^2\theta + \delta_1\theta + \delta_2\Sigma(\theta - \theta_0)).\end{aligned}$$

One can solve  $\theta_\epsilon$  numerically via solving  $\partial R/\partial \theta = \mathbf{0}$  or minimizing  $R(\theta, \epsilon)$ .

**Formula of  $\theta_\epsilon$**  Under first-order optimality condition, we have

$$\mathbf{0} = \Sigma(\theta - \theta_0) + \epsilon^2\theta + \delta_1\theta + \delta_2\Sigma(\theta - \theta_0) = (\Sigma + \delta_2\Sigma)(\theta - \theta_0) + (\epsilon^2 + \delta_1)\theta,$$

which implies that

$$\theta_\epsilon = (\Sigma + \delta_2\Sigma + \epsilon^2I + \delta_1I)^{-1}(\Sigma + \delta_2\Sigma)\theta_0 := (\Sigma + \lambda_\epsilon I)^{-1}\Sigma\theta_0.$$

**Remark 3.** As in Javanmard et al. (2020); Xing et al. (2021b), when  $\Sigma = I_d$ ,  $\lambda_\epsilon$  is only a function of  $(\|\theta_0\|, \sigma^2, \epsilon)$ , regardless of  $d$ . However, in terms of the excess adversarial risk, it is related to  $d$ .

**Label cost** For  $E_1$ ,

$$\begin{aligned}\frac{\partial^2 R}{\partial \theta \partial \theta_0} &= -2\Sigma - 2\epsilon c_0 \frac{\theta(\theta - \theta_0)^\top \Sigma}{\|\theta\| \sqrt{\|\theta - \theta_0\|_\Sigma^2 + \sigma^2}} - 2\epsilon c_0 \|\theta\| \frac{\Sigma}{\sqrt{\|\theta - \theta_0\|_\Sigma^2 + \sigma^2}} \\ &\quad + 2\epsilon c_0 \|\theta\| \frac{\Sigma(\theta - \theta_0)(\theta - \theta_0)^\top \Sigma}{\sqrt{\|\theta - \theta_0\|_\Sigma^2 + \sigma^2}^3}.\end{aligned}$$

For the term  $\widehat{\theta}(0) - \theta_0$  in  $E_1$ , we know that,

$$\widehat{\theta}(0) - \theta_0 = \frac{\Sigma_0^{-1}}{n_1} \sum_{S_1} \frac{\partial}{\partial \theta_\epsilon} l_0(x, y, \theta).$$

When  $\theta = \theta_0$ ,

$$\text{Var}\left(\frac{\partial l_0}{\partial \theta_\epsilon}\right) = 4\sigma^2\Sigma, \quad \Sigma_0 = 2\Sigma.$$

As a result, the label cost becomes

$$\begin{aligned}\mathbb{E}\|E_1\|_{\Sigma_\epsilon}^2 &= \frac{1}{n_1} \text{tr} \left( \Sigma_\epsilon^{-1} \frac{\partial^2 R}{\partial \theta \partial \theta_0} \Sigma_0^{-1} (4\sigma^2\Sigma) \Sigma_0^{-1} \left( \frac{\partial^2 R}{\partial \theta \partial \theta_0} \right)^\top \right) \\ &= \frac{1}{n_1} \text{tr} \left( \Sigma_\epsilon^{-1} \frac{\partial^2 R}{\partial \theta \partial \theta_0} (\sigma^2\Sigma^{-1}) \left( \frac{\partial^2 R}{\partial \theta \partial \theta_0} \right)^\top \right).\end{aligned}$$

The matrix  $\Sigma_\epsilon$  has an analytical form as

$$\begin{aligned}\Sigma_\epsilon &= 2\Sigma + 2\epsilon^2I_d + 2\epsilon c_0 \left( \frac{I_d}{\|\theta\|} - \frac{\theta\theta^\top}{\|\theta\|^3} \right) \sqrt{\|\theta - \theta_0\|_\Sigma^2 + \sigma^2} \\ &\quad + 2\epsilon c_0 \frac{\theta}{\|\theta\|} \frac{(\Sigma(\theta - \theta_0))^\top}{\sqrt{\|\theta - \theta_0\|_\Sigma^2 + \sigma^2}} + 2\epsilon c_0 \frac{\Sigma(\theta - \theta_0)}{\sqrt{\|\theta - \theta_0\|_\Sigma^2 + \sigma^2}} \frac{\theta^\top}{\|\theta\|} \\ &\quad + 2\epsilon c_0 \|\theta\| \frac{\Sigma}{\sqrt{\|\theta - \theta_0\|_\Sigma^2 + \sigma^2}} + 2\epsilon c_0 \|\theta\| \frac{(\Sigma(\theta - \theta_0))(\Sigma(\theta - \theta_0))^\top}{\sqrt{\|\theta - \theta_0\|_\Sigma^2 + \sigma^2}}.\end{aligned}$$

**Vanilla adversarial training** For the excess adversarial risk of  $\widehat{\theta}(\epsilon)$ , we have

$$\mathbb{E}\|\widehat{\theta}(\epsilon) - \theta_\epsilon\|_{\Sigma_\epsilon}^2 = \frac{1}{n_1} \text{tr} \left( \Sigma_\epsilon^{-1} \text{Var} \left( \frac{\partial}{\partial \theta_\epsilon} l_\epsilon(X, Y, \theta) \right) \right),$$

where

$$\begin{aligned} \frac{\partial l_\epsilon}{\partial \theta} &= \underbrace{2x(x^\top \theta - x^\top \theta_0 - \epsilon)}_{:=\xi_1} + \underbrace{2\epsilon^2 \theta}_{:=\xi_2} + \underbrace{2\epsilon \operatorname{sgn}(x^\top \theta - x^\top \theta_0 - \epsilon) \|\theta\| x}_{:=\xi_3} \\ &\quad + \underbrace{2\epsilon \frac{\theta}{\|\theta\|} |x^\top \theta - x^\top \theta_0 - \epsilon|}_{:=\xi_4}. \end{aligned}$$

Since  $\mathbb{E} \partial l_\epsilon / \partial \theta = \mathbf{0}$ , we have

$$\begin{aligned} e_2 &= \mathbb{E} \left( \frac{\partial}{\partial \theta} l_\epsilon(X, Y, \theta) \right) \left( \frac{\partial}{\partial \theta} l_\epsilon(X, Y, \theta) \right)^\top \\ &= \mathbb{E}(\xi_1 + \xi_2 + \xi_3 + \xi_4)(\xi_1 + \xi_2 + \xi_3 + \xi_4)^\top, \end{aligned}$$

where

$$\begin{aligned} \mathbb{E} \xi_1 \xi_1^\top &= 4\mathbb{E}[X X^\top (\theta - \theta_0)^\top X X^\top (\theta - \theta_0)] + 4\sigma^2 \Sigma, \\ \mathbb{E} \xi_2 \xi_2^\top &= 4\epsilon^4 \theta \theta^\top, \\ \mathbb{E} \xi_3 \xi_3^\top &= 4\epsilon^2 \|\theta\|^2 \Sigma, \\ \mathbb{E} \xi_4 \xi_4^\top &= 4\epsilon^2 \frac{\theta \theta^\top}{\|\theta\|^2} (\|\theta - \theta_0\|_\Sigma^2 + \sigma^2). \end{aligned}$$

And some cross terms can be calculated as

$$\begin{aligned} \mathbb{E} \xi_1 \xi_2^\top + \mathbb{E} \xi_2 \xi_1^\top &= 4\epsilon^2 (\Sigma(\theta - \theta_0) \theta^\top + \theta(\theta - \theta_0)^\top \Sigma), \\ \mathbb{E} \xi_2 \xi_4^\top + \mathbb{E} \xi_4 \xi_2^\top &= 8c_0 \epsilon^3 \frac{\theta \theta^\top}{\|\theta\|} \sqrt{\|\theta - \theta_0\|_\Sigma^2 + \sigma^2} \\ \mathbb{E} \xi_3 \xi_4^\top + \mathbb{E} \xi_4 \xi_3^\top &= 4\epsilon^2 (\Sigma(\theta - \theta_0) \theta^\top + \theta(\theta - \theta_0)^\top \Sigma). \end{aligned}$$

For the other cross terms,  $\xi_1 \xi_3^\top$ ,  $\xi_1 \xi_4^\top$ ,  $\xi_2 \xi_3^\top$ , there is no closed-form formula for their expectation, and we use simulation to approximate them.

**Simulation** To obtain an  $\tilde{\theta}(\epsilon)$  which does not involve generator cost, the procedure is slightly different to other simulation studies in this paper. Rewrite  $R(\theta, \epsilon) = R_\epsilon(\theta, \theta_0)$ . We take  $n_1 = 1,000$  to obtain  $\hat{\theta}(0)$ , and directly numerically minimize  $R_\epsilon(\theta, \hat{\theta}(0))$  to get  $\tilde{\theta}(\epsilon)$ . We repeat this procedure for 1,000 times to get a stable estimate of the excess risk ratio.