

# Failure Cases Are Better Learned But Boundary Says Sorry: Facilitating Smooth Perception Change for Accuracy-Robustness Trade-Off in Adversarial Training

Yanyun Wang Li Liu\*

The Hong Kong University of Science and Technology (Guangzhou)

## Abstract

Adversarial Training (AT) is one of the most effective methods to train robust Deep Neural Networks (DNNs). However, AT creates an inherent trade-off between clean accuracy and adversarial robustness, which is commonly attributed to the more complicated decision boundary caused by the insufficient learning of hard adversarial samples. In this work, we reveal a counterintuitive fact for the first time: **From the perspective of perception consistency, hard adversarial samples that can still attack the robust model after AT are already learned better than those successfully defended.** Thus, different from previous views, we argue that it is rather the over-sufficient learning of hard adversarial samples that degrades the decision boundary and contributes to the trade-off problem. Specifically, the excessive pursuit of perception consistency would force the model to view the perturbations as noise and ignore the information within them, which should have been utilized to induce a smoother perception transition towards the decision boundary to support its establishment to an appropriate location. In response, we define a new AT objective named **Robust Perception**, encouraging the model perception to change smoothly with input perturbations, based on which we propose a novel **Robust Perception Adversarial Training (RPAT)** method, effectively mitigating the current accuracy-robustness trade-off. Experiments on CIFAR-10, CIFAR-100, and Tiny-ImageNet with ResNet-18, PreActResNet-18, and WideResNet-34-10 demonstrate the effectiveness of our method beyond four common baselines and 12 state-of-the-art (SOTA) works. The code is available at <https://github.com/FlaAI/RPAT>.

## 1. Introduction

Deep neural networks (DNNs) have demonstrated impressive performance across various real-world applications [24, 26, 60]. However, they are inherently vulnerable to adversarial attacks [4, 39]. Specifically, even impercep-

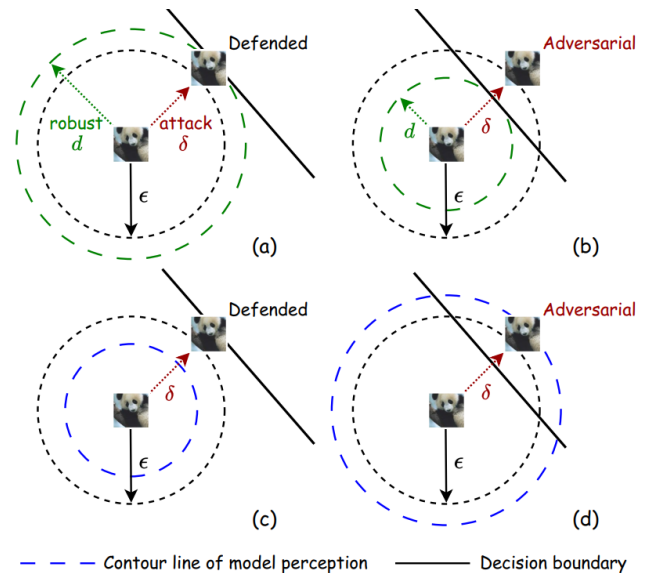


Figure 1. Success and failure cases of AT defense from the perspectives of prediction and perception, respectively. (a) and (b): AT works through building a robust ball with radius  $d$  such that the model prediction remain unchanged within it. Then if the robust  $d$ -ball fails to contain the  $\epsilon$ -ball, there would still be failure cases beating the robust model under perturbation  $\delta$  within the  $\epsilon$ -ball. (c) and (d): In this work, we reveal that the smaller  $d$ -ball in the failure cases is not a result of more vulnerable model perception as previous intuition. In fact, as demonstrated by the contour line, which shows the perturbation ball to the same degree of perception change, the failure cases already achieve less perception change under similar  $\delta$  within the  $\epsilon$ -ball. This implies that the issue of these failure cases is not failing to learn perception consistency following the current AT objective, but failing to support the establishment of decision boundary to an appropriate location.

tible perturbations to inputs can cause substantial changes in the outputs [12, 17]. As this problem raises significant security concerns [45], previously, several defense strategies, such as Defense Distillation [32], Feature Squeezing [52], Randomization [51], and Input Denoising [14, 25], have been proposed to improve the adversarial robustness of DNNs. However, most of these approaches have proven

\*Corresponding Author: avrillliu@hkust-gz.edu.cn

ineffective against more advanced adaptive attacks [2, 42], such as Auto-Attack [7]. Currently, Adversarial Training (AT) [12, 29] is regarded as one of the most effective methods for training robust DNNs [2, 9].

However, an **accuracy-robustness trade-off** problem have been widely identified for AT in recent studies [3, 28, 34–36, 43, 44, 48, 54, 58]. Compared with standard training, the improvement in adversarial robustness acquired by AT is usually at the expense of the reduction in clean accuracy, which is afraid to significantly degrade the user experience in clean context and hinder the application of AT in real-world practices. The current consensus is that AT aggravates the accuracy-robustness trade-off by creating overly complex decision boundaries, primarily because models struggle to sufficiently learn from hard adversarial samples during training [6, 8, 10, 19, 27, 28, 34, 48, 53, 54]. Specifically, as the conventional AT objective encourages the model to perceive any adversarial inputs within the  $\epsilon$ -ball similar to the corresponding benign one to produce consistent predictions for them, previous works empirically assume that the wrong predictions of the hard adversarial samples are attributed to the insufficiently optimized perception of them under the current objective. Therefore, although defining “hard sample” from different perspectives (*e.g.*, classification result and distance to decision boundary), previous works agree to further optimize the perception consistency on them, ending up with various learning strategies specific for them, such as data reweighting [27, 59], adaptive perturbation radius [8, 34], and non-one-hot supervised signal [10, 19].

However, in this work, we reveal that for the robust model trained by convention AT, **the statistical difference in model perception from benign to hard adversarial samples causing failure defense is already smaller than to others (*i.e.*, those successfully defended)**. It implies that the model has even more sufficiently learned them under the current AT objective, and should have successfully defended them if the above assumption holds. Thus, as demonstrated in Figure 1, we suggest that the key factor here is no longer the optimization of model perception but the location of the decision boundary. In that case, keeping on pursuing over-sufficient small perception change would not only help little in further reducing the prediction error, but also force the model to view the perturbations as noise and ignore the information within them, which impairs their ability in claiming robust balls with reasonable gradients to push the decision boundary to an appropriate location. As a consequence, such sample-level imbalance in boundary establishment would naturally increase its complexity, harming generalization ability and degrading the trade-off between accuracy and robustness.

To mitigate the over-sufficient learning of model perception on the failure cases and retain effective gradient sig-

nals *w.r.t.* them to facilitate the establishment of the decision boundary to an appropriate location, **firstly**, we define a new AT objective named **Robust Perception** in addition to the conventional one, encouraging smoother perception changes towards the decision boundary along with the input perturbations. Rigorous theoretical derivations also demonstrate its effectiveness in smoothening the decision boundary and releasing the current trade-off problem from the perspectives of local linearity [12] and *Lipschitz* regularization [1, 11, 21, 33]. **Secondly**, based on our new objective, we propose a novel **Robust Perception Adversarial Training (RPAT)** method, concurrently improving clean accuracy and adversarial robustness in AT. Extensive experiments on CIFAR-10, CIFAR-100, and Tiny-ImageNet with ResNet-18, PreActResNet-18, and WideResNet-34-10 demonstrate the advancement of RPAT beyond four common benchmarks and 12 state-of-the-art (SOTA) works.

## 2. Background and Related Works

### 2.1. Robustness and Adversarial Training

Adversarial robustness reflects model vulnerability to malicious perturbations, which is commonly quantified by test accuracy under adversarial attacks [3, 39, 46]. Representatively, the Fast Gradient Sign Method (FGSM) [12] perturbs input samples along the gradient of the loss function, whereas the Projected Gradient Descent (PGD) attack [29] applies FGSM iteratively and projects the perturbation onto the  $\epsilon$ -ball. Auto-Attack [7] introduces and combines several parameter-free attacks to mitigate practical issues such as hyper-parameter tuning and gradient masking. Due to their effectiveness, PGD and Auto-Attack are widely used to assess adversarial robustness, and so does this work.

Previously, although many defense strategies including Defense Distillation [32], Feature Squeezing [52], Input Denoising [14, 25] and Randomization [51] have been developed, most rely on obfuscated gradients [2] and can be evaded by advanced adaptive attacks [42]. Nowadays, Adversarial Training [12, 29] is recognized as the most effective method to enhance adversarial robustness [2, 9, 41]. Unlike clean training on benign samples, AT directly learns on adversarially augmented data. Its objective can be formulated as the following min-max problem [29]:

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \max_{\|\mathbf{x}'_i - \mathbf{x}_i\|_p \leq \epsilon} \mathcal{L}(f_{\theta}(\mathbf{x}'_i), y_i), \quad (1)$$

where  $f_{\theta}(\mathbf{x}_i) = \arg \max_{c=1, \dots, C} \mathbf{p}_c(\mathbf{x}_i, \theta)$  with  $\mathbf{p}_c$  denoting the *softmax* probability of class  $c$  within the prediction of  $\mathbf{x}_i$ ,  $n$  denotes the number of training samples, and  $\mathbf{x}'_i$  is generated by adding the strongest  $L_p$ -norm perturbation within the  $\epsilon$ -ball to the original sample  $\mathbf{x}_i$ . This objective aims at prediction consistency within the  $\epsilon$ -ball, encouraging the model to also produce similar perception there.

PGD-AT [29] is the first attempt at solving this min-max optimization iteratively by approximating the inner maximization via PGD to craft adversarial samples and then performing stochastic gradient descent (SGD) on these samples for the outer minimization. TRADES [58] first considers balancing clean accuracy and adversarial robustness in AT, which decomposes robust error into natural error and boundary error, with the latter occurring only if specific data points are sufficiently close to the decision boundary. MART [45] further refines AT by distinguishing misclassified from correctly classified samples, which suggests appending boundary error only for misclassified samples to the basic AT objective in PGD-AT. Consistency-AT [40] encourages the prediction between adversarial samples derived from different augmentations of the same benign sample to be similar, which explicitly enhances the prediction consistency further.

## 2.2. Accuracy-Robustness Trade-Off in AT

An important problem identified for AT is the trade-off between clean accuracy and adversarial robustness [3, 28, 34–36, 43, 44, 54, 58], which means the robustness gained by AT often comes at the cost of reduced model accuracy compared with clean training. Consequently, AT can degrade user experience under benign conditions, which may seriously delay its real-world deployment. A widely acknowledged cause of this problem is that AT tends to learn more complicated decision boundaries than clean training [6, 10, 28, 34, 53, 54], thereby harming the ability of the robust model to generalize to unseen data. A number of previous works attribute this to the insufficient learning of hard adversarial samples [8, 10, 19, 27, 34, 59]. Though defining “hard adversarial samples” from different perspectives, such as the classification result and the distance to decision boundary, they basically reached a consensus that the model perception *w.r.t.* hard adversarial samples is not well-optimized, and as a consequence, fails to produce the consistent final prediction. In response, these works propose various solutions to help the model learn better prediction consistency for those hard adversarial samples.

For instance, GAIRAT [59] proposes that hard samples nearer to the decision boundary are more critical but less robust and should be assigned a larger weight. MAIL [27] further proposes three types of probabilistic margins to quantify the closeness for better reweighting adversarial data accordingly. In contrast, MMA [8] suggests learning hard adversarial samples with reduced perturbations, which employs an adaptive  $\epsilon$  for perturbations to maximize the margin between data and the decision boundary. EWAT [20] weighs the loss for each adversarial sample proportionally to the entropy of its prediction distribution during AT to focus on those with more uncertain labels. TE [10] reveals that the model may attempt to memorize certain hard ad-

versarial samples during AT, as it is difficult to assign high-confidence one-hot labels for them [6, 38]. SOVR [19] suggests increasing the margins of *logits* for certain hard samples by switching to a one-vs-the-rest loss.

Additionally, several works also study the trade-off problem from other perspectives. AWP [49] proposes a double perturbation mechanism that flattens the loss landscape by weight perturbation to improve robust generalization. KD+SWA [5] injects more smoothing during AT by leveraging knowledge distillation to smooth the logits and performing stochastic weight averaging [18] to smooth the weights. ADR [50] employs soft labels as a guidance mechanism, accurately reflecting the distribution shift under attack during AT. CURE [13] selectively updates specific layers while preserving others to enhance the learning capacity of the model. ReBAT [47] views AT as a dynamic min-max game between the model trainer and the attacker, rebalancing the two players by either constraining the trainer’s capacity or intensifying the attack.

## 3. Motivation: AT Over-Learns Failure Cases

It is basically common sense for DNN-based classification that the model learns to perceive input better and is thus able to predict it more precisely (especially when the model perception is indicated by its *logits*). Therefore, it seems reasonable at first glance that the previous works assume that the wrong prediction of the hard adversarial samples is due to the insufficient model perception of them, and propose various solutions upon this assumption in response.

However, to the best of our knowledge, there is no previous discussion on a simple but crucial question: **What if the assumption does not hold?** If the model perception on the hard adversarial samples is already relatively close to the corresponding benign ones (*e.g.*, at a similar or even better level than non-hard ones), then it would be no longer effective enough to further optimize the model on such samples through an objective simply pursuing prediction consistency as the conventional AT. Also, in turn, this can no longer effectively benefit the further reduction of the prediction error. Surprisingly, our following proof-of-concept experiments reveal that this is exactly the case.

The proof-of-concept experiments are conducted on CIFAR-10 dataset with ResNet-18 architecture, with the results illustrated in Figure 2. As suggested by the x-axis, we involve three cases with different pairs of training and evaluation methods, which are respectively 1) trained on clean data and evaluated with random perturbations, 2) trained and evaluated with randomly perturbed data, and 3) trained with PGD-AT and evaluated with PGD-20 adversary. The budget of the random perturbations  $\epsilon = 8/255$  is assigned the same as PGD, following our experimental setup detailed in Appendix A, so do other training and evaluation settings. Given that the conventional AT objective in Equation (1)

encourages the model to produce a similar perception for benign and adversarial samples, we utilize the similarity in model perception between the corresponding benign and perturbed samples to estimate the learning sufficiency of the perturbed samples in AT. Specifically, the model perception is represented by the *logits*, and the similarity is measured through the *mean squared error (MSE)* score. We calculate the score respectively for success and failure cases of defense, which are divided according to the correctness of the final classification results of the perturbed samples. In addition, the model accuracy on clean and perturbed data is also recorded and illustrated together.

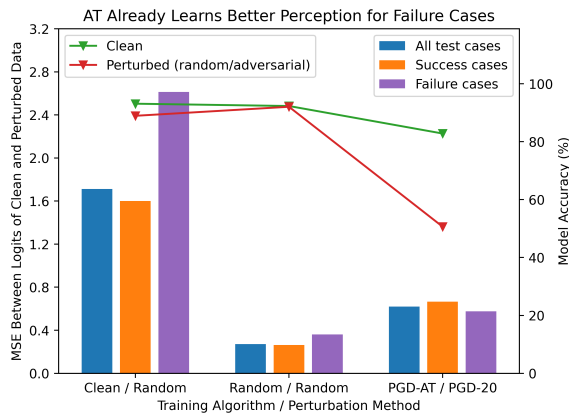


Figure 2. Proof-of-concept experiments for the different learning sufficiency in model perception for success and failure cases on CIFAR-10 with ResNet-18. In the x-axis, the “Random” for “Training Algorithm” means the training data is randomly perturbed within the budget  $\epsilon$ , while that for “Perturbation Method” indicates the attack used for evaluation on test data is random perturbation. The bars demonstrate the *MSE* score on *logits* between each pair of benign and perturbed samples, with the lines additionally showing the model accuracy on clean and perturbed data. Our results reveal a surprising phenomenon that, different from the clean training (with or without random noise), where failure cases are worse learned as we intuitively expected, AT already learns a better model perception for failure cases.

Surprisingly, we experimentally found that **the failure cases already learn better in model perception than the success cases**. First, for the clean trained model, the perception on failure cases is significantly worse than on the success ones, just as we expected. Although the difference between success and failure cases has decreased for the randomly trained model, the failure cases remain worse, which suggests that non-adversarial perturbations are influential but not enough to change the trend. However, unlike the previous intuition, the model trained by PGD-AT can already perceive failure cases better than success ones. Our finding updates the current knowledge that, unlike those in clean training, the failure cases in AT do not fail simply because they are not sufficiently learned. On the contrary,

the model even learns them harder to try to achieve the prediction and perception consistency. The smaller perception change under similar perturbations is direct evidence for this point of view.

Then, as a straightforward deduction, now that we have success cases with worse perception but failure cases with better perception, there should be a statistical difference between them in the maximum perception change they can tolerate while still maintaining the consistent prediction, just as demonstrated in Figure 1. In other words, the decision boundary is established relatively closer to the failure cases, which is instead the issue we should focus on, because such sample-level imbalance in boundary establishment can naturally complicate it and thus degrade the accuracy-robustness trade-off as mentioned in Section 2.2. In that case, the over-consistent model perception between benign and adversarial samples in the failure cases rather harms the final robustness, because this is expected to force the model to view the perturbations as noise and ignore the information within them, which should have been utilized to induce a smoother perception transition towards the decision boundary to support its establishment, thus impairing the ability of such cases to claim a reasonable robust distance during AT compared with success cases. Therefore, in this work, we argue that **it is not insufficient but over-sufficient learning of failure cases that contributes to the more complicated decision boundary and finally results in the current trade-off problem**.

## 4. Proposed Method

### 4.1. Robust Perception: A New AT Objective

In Section 3, the too-small perception change under perturbations caused by the over-sufficient learning of failure cases has been identified as an important factor contributing to the current trade-off problem, which in turn questions the effectiveness of the conventional AT objective on such failure cases.

Formally, under Equation (1) which implicitly optimize the perception distance  $\|h_{\theta}(\mathbf{x}') - h_{\theta}(\mathbf{x})\| \rightarrow 0$  between corresponding benign sample  $\mathbf{x}$  and adversarial sample  $\mathbf{x}'$ , the gradient of the conventional AT loss *w.r.t.* the perturbation,  $\nabla_{\|\mathbf{x}' - \mathbf{x}\|} \mathcal{L}$ , would also approach 0. In that case, the model cannot further optimize the decision boundary using the perturbation information. Thus, a straightforward idea to mitigate this is to encourage smoother perception change towards the decision boundary along with the input perturbation, such that  $\|h_{\theta}(\mathbf{x}') - h_{\theta}(\mathbf{x})\| \propto \|\mathbf{x}' - \mathbf{x}\|$ , as long as the final prediction remains unchanged. This is expected to retain an effective gradient signal in the direction of the perturbation, which enables the model to utilize this information to guide the decision boundary to a more reasonable location. Accordingly, in this section, we propose *Robust*

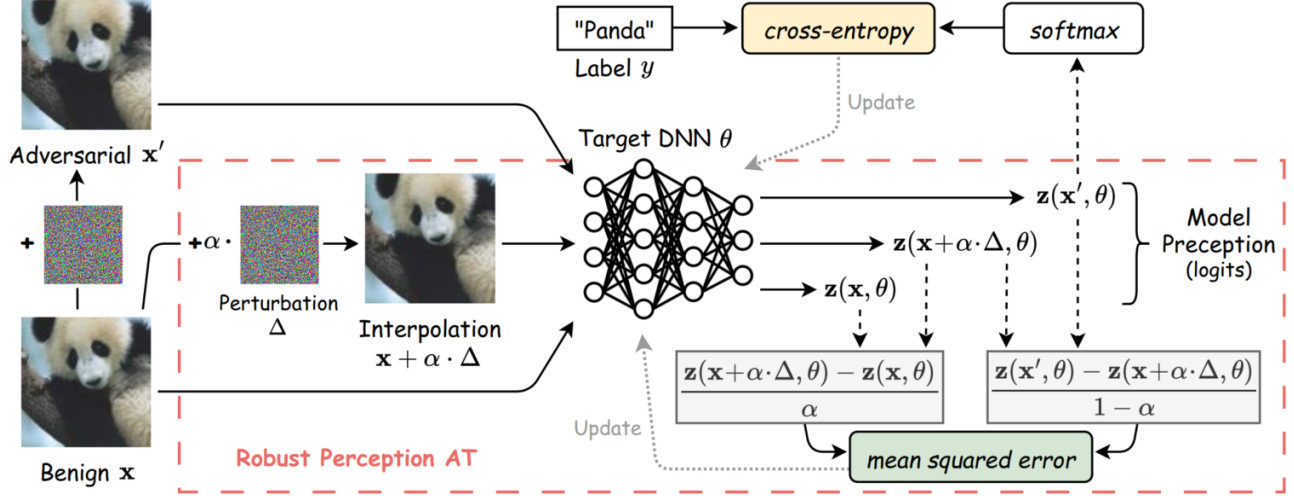


Figure 3. Illustration of the proposed RPAT method. In addition to the conventional AT objective, which learns adversarial samples with the original supervised labels, RPAT introduces a novel regularization term to achieve the newly defined **Robust Perception** objective. Given benign sample  $\mathbf{x}$  and adversarial samples  $\mathbf{x}' = \mathbf{x} + \Delta$ , for any interpolation samples  $\mathbf{x} + \lambda \cdot \Delta$  with  $\lambda \in [0, 1]$ , it encourages the ratio of the perception changes,  $z_{\theta}(\mathbf{x} + \lambda \cdot \Delta) - z_{\theta}(\mathbf{x})$  and  $z_{\theta}(\mathbf{x}') - z_{\theta}(\mathbf{x} + \lambda \cdot \Delta)$ , respectively from  $\mathbf{x}$  to  $\mathbf{x} + \lambda \cdot \Delta$  and from  $\mathbf{x} + \lambda \cdot \Delta$  to  $\mathbf{x}'$ , to approach the ratio indicated by  $\lambda$ , such that the model perception would change smoothly along with the input perturbations from  $\mathbf{x}$  to  $\mathbf{x}'$ .

*Perception* as a new AT objective in addition to the conventional one to realize this idea.

**Definition 1 (Robust Perception).** *Provided an AT task with target model  $\theta$ , let  $\mathbf{x}$  and  $\mathbf{x}'$  be any pairs of corresponding benign and adversarial samples with  $\Delta = \mathbf{x}' - \mathbf{x}$ , for any hidden representation  $h_{\theta}(\mathbf{x})$  denoting the model perception, our additional AT objective can be formulated as:*

$$\forall \alpha \in [0, 1], \|h_{\theta}(\mathbf{x} + \alpha \cdot \Delta) - h_{\theta}(\mathbf{x})\| = \alpha \cdot \|h_{\theta}(\mathbf{x}') - h_{\theta}(\mathbf{x})\|.$$

To facilitate the understanding of the proposed *Robust Perception*, in addition to the most straightforward perspective that it encourages the change in model perception along with the adversarial perturbations, the effectiveness of this new objective can also be demonstrated from other two perspectives, local linearity [12] and *Lipschitz* regularization [1, 11, 21, 33], both of which also help smooth the decision boundary and mitigate the current trade-off problem. Below, we further detail them theoretically.

**Theorem 1.** *Let  $H$  be the Hessian Matrix such that  $H_{h_{\theta}}(\mathbf{x}) = \nabla_{\mathbf{x}}^2 h_{\theta}(\mathbf{x})$ , then with the new optimization objective of *Robust Perception*, we have:*

$$\forall \Delta, \Delta^{\top} \cdot H_{h_{\theta}}(\mathbf{x}) \cdot \Delta \rightarrow 0.$$

*Proof.* We defer the proof to Appendix B.1.

This means *Robust Perception* limits the second-order (and higher-order as well) nonlinear effects within the adversarial perturbation to the model perception, such that the perception would change mainly along with the linear

term of the perturbation, and thus be smoother. This also aligns with the local linearity hypothesis of adversarial samples [12], which assumes that linear approximations dominate the model behavior under adversarial perturbations.

**Theorem 2.** *Let  $J$  be the Jacobian Matrix such that  $J_{h_{\theta}}(\mathbf{x}) = \nabla_{\mathbf{x}} h_{\theta}(\mathbf{x})$ , then with the new optimization objective of *Robust Perception*, we have:*

$$\forall \alpha \in [0, 1], J_{h_{\theta}}(\mathbf{x} + \alpha \cdot \Delta) \rightarrow J_{h_{\theta}}(\mathbf{x}),$$

then with  $\|\cdot\|_{\text{spec}}$  denoting the Spectral Norm and  $\gamma$  being any micro value, given  $\|J_{h_{\theta}}(\mathbf{x} + \alpha \cdot \Delta) - J_{h_{\theta}}(\mathbf{x})\|_{\text{spec}} \leq \gamma$ , the function  $h_{\theta}(\mathbf{x})$  can be referred to as *K-Lipschitz* with the Lipschitz constant  $K$  upper-bounded by:

$$K \leq \sup_{\mathbf{x}} \|J_{h_{\theta}}(\mathbf{x})\|_{\text{spec}} + \gamma.$$

*Proof.* We defer the proof to Appendix B.2.

This theorem tells us that *Robust Perception* also regularizes the change  $\gamma$  in Jacobian along with the adversarial perturbation, which can be interpreted as limiting the increase of the global Lipschitz constant by  $\gamma$  under perturbation. Previous works have demonstrated that this is expected to help achieve smoother decision boundaries and improve the accuracy-robustness trade-off [1, 11, 21, 33].

## 4.2. Robust Perception Adversarial Training

Based on the above considerations, in this section, we propose a novel AT method named **Robust Perception Adversarial Training (RPAT)**. It utilizes a new additional

regularization term appended to the conventional AT objective to achieve the *Robust Perception*. Formally, given a  $C$ -class classification task ( $C \geq 2$ ) and a dataset  $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1, \dots, n}$  with benign sample  $\mathbf{x}_i \in \mathbb{R}^d$  and supervised label  $y_i \in \{1, \dots, C\}$ , let  $\hat{\mathbf{x}}'_i$  denotes the adversarial sample corresponding to  $\mathbf{x}_i$ , which is generated as  $\hat{\mathbf{x}}'_i = \arg \max_{\|\mathbf{x}'_i - \mathbf{x}_i\|_p \leq \epsilon} \mathbb{1}(\arg \max_{k=1, \dots, C} h_{\theta}(\mathbf{x}'_i)^{(k)} \neq y_i)$ , then based on the Definition 1 in Section 4.1, the adversarial risk to be minimized for the proposed RPAT *w.r.t.*  $\mathbf{x}_i$  can be formulated with the  $0$ -1 loss [58] as:

$$\mathcal{R}(\theta, \mathbf{x}_i) := \mathbb{1}(f_{\theta}(\hat{\mathbf{x}}'_i) \neq y_i) + \mathbb{1}\left(\frac{h_{\theta}(\tilde{\mathbf{x}}_i) - h_{\theta}(\mathbf{x}_i)}{\alpha} \neq \frac{h_{\theta}(\hat{\mathbf{x}}'_i) - h_{\theta}(\tilde{\mathbf{x}}_i)}{1 - \alpha}\right), \quad (2)$$

where  $\tilde{\mathbf{x}}_i = \mathbf{x}_i + \alpha \cdot (\hat{\mathbf{x}}'_i - \mathbf{x}_i)$ .

While the proposed *Robust Perception* aims to improve learning efficacy for the failure cases, the above risk does not explicitly distinguish success and failure cases. This is because *Robust Perception* is designed to mitigate the gap of the failure cases in the effectiveness of learning compared with the success ones, rather than require any new characteristic beyond the success cases. In other words, the success cases are expected to also satisfy *Robust Perception* at the beginning. Therefore, considering the computational efficiency, it is not necessary to specifically exclude the success cases for the *Robust Perception* term.

Finally, by introducing practical surrogate losses to optimize the adversarial risk in Equation (2), we end up with our new RPAT method. To be specific, for the supervised learning term for adversarial samples, we adopt *cross-entropy* (CE) just as the conventional AT by default. Then, for the *Robust Perception* term, as it is formulated in a format of similarity regularization, various criterion that measures the similarity distance between the two input items are feasible options, such as *mean squared error* (MSE), *Kullback–Leibler divergence* (KL), *Jensen-Shannon divergence* (JS), and *Cosine similarity*. We simply adopt *MSE loss* for the basic version of RPAT here, and would also demonstrate the effectiveness of *KL loss* later in Section 5.3. With the logits  $\mathbf{z}(\mathbf{x}_i, \theta)$  utilized as the hidden representation  $h_{\theta}(\cdot)$  to denote the model perception, and  $\mathbf{p}(\mathbf{x}_i, \theta)$  being the *softmax* output as given in Equation (1), the **overall objective** of RPAT can be formulated as:

$$\mathcal{L}^{\text{RPAT}}(\theta, \mathcal{D}, \lambda, \alpha) := \frac{1}{n} \sum_{i=1}^n \left( \mathcal{L}^{\text{CE}}(\mathbf{p}(\hat{\mathbf{x}}'_i, \theta), y_i) + \lambda \cdot \mathcal{L}^{\text{MSE}}\left(\frac{\mathbf{z}(\tilde{\mathbf{x}}_i, \theta) - \mathbf{z}(\mathbf{x}_i, \theta)}{\alpha} \parallel \frac{\mathbf{z}(\hat{\mathbf{x}}'_i, \theta) - \mathbf{z}(\tilde{\mathbf{x}}_i, \theta)}{1 - \alpha}\right) \right), \quad (3)$$

where  $\lambda$  is a weight hyper-parameter to balance the conventional AT objective and our new *Robust Perception* objective. An intuitive illustration of the proposed RPAT method is provided in Figure 3.

## 5. Experiments

Following previous AT works, the effectiveness of the proposed RPAT is evaluated on CIFAR-10, CIFAR-100 [22] and Tiny-ImageNet [23] datasets. The main experiments can be divided into two parts. First, based on Tack et al. [40], we implement RPAT upon three common AT baselines, namely PGD-AT, TRADES, and MART, as well an advanced method, Consistency-AT, with ResNet-18 [15] architecture. This is to show the general effectiveness of the proposed *Robust Perception* in improving the current trade-off. Second, by integrating RPAT into one of the current SOTA methods, ReBAT [47], with PreActResNet-18 [16] and WideResNet-34-10 [56] architectures, we achieve new SOTA performance beyond 12 SOTA methods on the accuracy-robustness trade-off problem. This is to demonstrate the advancement achieved by the proposed RPAT in the AT area. The detailed experimental setup is provided in Appendix A.

### 5.1. Evaluation Measures

There are five measures in total adopted to evaluate the experimental methods from three different perspectives. Firstly, we report clean accuracy on benign test data, which reflects the ability of the robust models to maintain performance in benign cases. Secondly, we evaluate adversarial robustness under two adversarial attack methods, namely PGD-20 [29] and Auto-Attack (AA) [7], with default settings aligning with Appendix A and random starts. Finally, to test the performance *w.r.t.* the current accuracy-robustness trade-off problem in AT, on the one hand, we report the mean of the clean accuracy and the robust score under Auto-Attack acquired above, and on the other hand, we introduce the Natural-Robustness Ratio (NRR) [13]:

$$\text{NRR} = \frac{2 \times \text{Clean Accuracy} \times \text{Adversarial Robustness}}{\text{Clean Accuracy} + \text{Adversarial Robustness}},$$

which emphasizes how well the robust model strikes a balance between clean accuracy and robustness. A lower NRR value suggests that the model might prioritize one side over the other. All the results reported are averages of three runs, with the “best” performance of each run recorded on the best checkpoint achieving the highest PGD-20 accuracy, which ensures that the training process is isolated from AA, the test-time adversary to evaluate the robustness.

### 5.2. Effectiveness in Improving Baselines

The proposed RPAT can be easily implemented with existing AT methods by simply appending the *Robust Perception* term. In this section, we illustrate the improvements that RPAT can bring to our four experimental benchmarks, namely PGD-AT, TRADES, MART, and Consistency-AT, on CIFAR-10, CIFAR-100, and Tiny-ImageNet with

Table 1. Improvements that the proposed RPAT achieved upon four benchmarks. The results are acquired with ResNet-18 and  $\ell_\infty$  norm.

Dataset	Method	Benchmarks					+RPAT (Ours)							
		Clean	PGD-20	AA	Mean	NRR	Clean		PGD-20		AA		Mean	NRR
		best	best	best			best	final	best	final	best	final		
CIFAR-10	PGD-AT	82.92	50.61	46.74	64.830	59.782	83.20	82.99	51.29	50.97	48.00	47.60	65.600	60.878
	TRADES	79.67	51.91	<u>47.62</u>	63.645	<u>59.610</u>	80.02	80.02	51.21	51.21	47.47	47.45	63.745	59.590
	MART	<u>77.93</u>	53.00	46.70	<u>62.315</u>	<u>58.402</u>	75.44	75.95	52.88	52.87	47.03	47.30	61.235	57.940
	Consistency-AT	83.42	51.96	47.72	65.570	60.711	<b>84.12</b>	84.04	52.33	52.14	<b>48.98</b>	49.06	<b>66.550</b>	<b>61.911</b>
CIFAR-100	PGD-AT	56.56	28.80	<u>25.02</u>	40.790	34.693	58.22	58.58	29.16	28.57	24.88	24.23	41.550	34.862
	TRADES	55.39	29.36	24.51	39.950	33.983	57.50	57.25	29.42	29.00	25.05	24.45	41.275	34.897
	MART	49.83	30.38	25.00	37.415	33.295	50.72	50.72	30.33	30.33	25.34	25.06	38.030	33.796
	Consistency-AT	58.53	29.28	25.39	41.960	35.417	<b>60.33</b>	60.51	29.97	29.64	<b>26.31</b>	25.97	<b>43.320</b>	<b>36.641</b>
Tiny-ImageNet	PGD-AT	46.32	21.58	17.07	31.695	24.947	47.68	48.32	22.13	22.00	17.77	17.91	32.725	25.891
	TRADES	46.75	21.52	16.60	31.675	24.500	48.77	48.85	21.98	21.93	16.92	16.71	32.845	25.124
	MART	39.70	22.90	17.18	28.440	23.982	41.76	41.66	23.40	23.29	17.79	17.78	29.775	24.951
	Consistency-AT	46.54	21.89	17.60	32.070	25.541	<b>49.74</b>	49.85	23.16	22.84	<b>18.84</b>	18.74	<b>34.290</b>	<b>27.329</b>

Table 2. Improvements that the proposed RPAT achieved upon four benchmarks. The results are acquired with ResNet-18 and  $\ell_2$  norm.

Dataset	Method	Benchmarks					+RPAT (Ours)							
		Clean	PGD-20	AA	Mean	NRR	Clean		PGD-20		AA		Mean	NRR
		best	best	best			best	final	best	final	best	final		
CIFAR-10	PGD-AT	87.76	67.92	66.36	77.060	75.574	88.20	88.45	68.56	68.52	67.63	67.62	77.915	76.557
	TRADES	83.99	68.60	65.93	74.960	73.872	85.17	84.80	68.95	68.90	67.67	67.37	76.420	75.418
	MART	84.09	68.32	66.28	75.185	74.130	84.65	84.65	68.60	68.60	66.69	66.70	75.670	74.604
	Consistency-AT	88.76	69.35	67.46	78.110	76.658	<b>89.38</b>	89.52	70.41	70.36	<b>69.44</b>	69.55	<b>79.410</b>	<b>78.158</b>
CIFAR-100	PGD-AT	65.00	41.54	<u>39.27</u>	52.135	48.960	65.14	65.44	41.61	41.28	39.23	39.20	52.185	48.969
	TRADES	61.25	43.05	40.15	50.70	48.505	62.41	62.64	43.56	43.10	<b>40.36</b>	40.12	51.385	49.020
	MART	60.08	43.42	39.85	49.965	47.92	60.57	60.57	43.92	43.90	<b>40.36</b>	40.36	50.465	48.442
	Consistency-AT	65.14	42.28	39.92	52.530	49.503	<b>65.54</b>	66.81	42.81	42.51	40.26	40.61	<b>52.900</b>	<b>49.880</b>
Tiny-ImageNet	PGD-AT	59.36	43.12	40.90	50.130	48.431	60.28	60.28	44.06	44.06	41.89	41.87	51.085	49.430
	TRADES	57.61	44.70	42.33	49.970	48.802	58.46	58.50	44.87	44.65	42.42	42.44	50.440	49.165
	MART	56.47	44.84	42.22	49.345	48.316	57.17	56.95	45.02	44.67	42.51	42.16	49.840	48.762
	Consistency-AT	61.49	44.92	42.59	52.040	50.324	<b>62.19</b>	62.64	45.28	44.87	<b>43.17</b>	42.85	<b>52.680</b>	<b>50.963</b>

ResNet-18, demonstrating the wide effectiveness of our method. The results under  $\ell_\infty$  and  $\ell_2$  threat models are shown in Table 1 and Table 2, respectively, where the best clean accuracy, AA robustness, as well mean and NRR for their trade-off on each of the datasets are marked in bold, while individual scores not improved are marked with underline. Note that PGD-20 is the indicator for selecting the “best” checkpoint as explained in Section 5.1, and the “final” scores are additionally provided to show that RPAT does not suffer from serious robust-overfitting, which is also evidence for appropriate algorithm design. Thus, these fields are not involved in the comparison. It can be found that RPAT achieves significant improvement for most scores among the tables, where Consistency-AT+RPAT is basically the most competitive for all the experimental groups, which demonstrates the general effectiveness of our method in mitigating the current trade-off problem.

### 5.3. Advancement beyond 12 SOTAs

To demonstrate the advancement that the proposed RPAT contributes, we further integrate RPAT into ReBAT [47],


one of the current SOTA AT methods on the accuracy-robustness trade-off problem. It implements WA [18], averaging multiple points along the SGD trajectory for better generalization, through the exponential moving average (EMA) [37] strategy with a decay rate of 0.999. Thus, we also involve WA itself (*i.e.*, with PGD-AT), as well as KD+SWA [5] which shares a similar idea, for comparison. Other differences in ReBAT include *KL loss* and a reduced learning rate decay factor from 0.1 to 0.5. We refer to our method upon ReBAT as RPAT<sup>++</sup>, and compare it with 12 previous SOTAs introduced in Section 2.2, respectively with PreActResNet-18 and WideResNet-34-10 as shown in Table 3 and Table 4. The clean and AA scores of the SOTAs are from the best checkpoint among their original papers, Kanai et al. [19], Wang et al. [47], and our reproduction if any (so due to the lack of previous record, ADR and CURE are excluded for PreActResNet-18). Beyond the basic CIFAR-10 results under  $\ell_\infty$  norm, Table 3 involves CIFAR-100 and  $\ell_2$  norm, and Table 4 further consider the *Cutmix* [55] data augmentation strategy as marked by “\*”, which ensures the comprehensiveness of our comparison.

Table 3. Comparison of the proposed RPAT<sup>++</sup> with current SOTAs on accuracy-robustness trade-off with PreActResNet-18. The best results are marked in bold.

Norm	Method	CIFAR-10				CIFAR-100				
		Clean	AA	Mean	NRR	Clean	AA	Mean	NRR	
$\ell_\infty$	WA	UAI'18	83.50	49.89	66.695	62.461	57.26	25.83	41.545	35.601
	MMA	ICLR'20	<b>85.50</b>	37.20	61.350	51.844	<b>60.60</b>	18.40	39.500	28.229
	AWP	NeurIPS'20	81.11	50.09	65.600	61.933	54.10	25.16	39.630	34.347
	GAIRAT	ICLR'21	78.70	37.70	58.200	50.979	52.00	19.80	35.900	28.680
	KD+SWA	ICLR'21	84.06	49.82	<b>66.940</b>	62.562	57.17	25.66	41.415	35.422
	EWAT	ICML'21	82.80	48.20	65.500	60.931	54.20	23.52	38.860	32.805
	MAIL	NeurIPS'21	79.50	39.60	59.550	52.867	46.50	16.70	31.600	24.574
	TE	ICLR'22	82.04	50.12	66.080	62.225	56.41	25.84	41.125	35.444
	SOVR	ICML'23	81.90	49.40	65.650	61.628	52.10	24.30	38.200	33.142
	ReBAT	NeurIPS'23	82.09	50.72	66.405	62.700	56.13	27.60	41.865	37.004
	<b>RPAT<sup>++</sup></b>	<b>Ours</b>	82.63	<b>51.00</b>	66.815	<b>63.072</b>	56.84	<b>27.68</b>	<b>42.260</b>	<b>37.230</b>
$\ell_2$	ReBAT	NeurIPS'23	88.79	71.00	79.895	78.905	65.58	42.67	54.125	51.701
	<b>RPAT<sup>++</sup></b>	<b>Ours</b>	<b>89.06</b>	<b>71.26</b>	<b>80.160</b>	<b>79.172</b>	<b>65.63</b>	<b>42.85</b>	<b>54.240</b>	<b>51.848</b>

Table 4. Comparison of RPAT<sup>++</sup> with SOTAs on CIFAR-10 with WideResNet-34-10 and  $\ell_\infty$  norm.

Method		Clean	AA	Mean	NRR
WA	UAI'18	87.66	52.65	70.155	65.787
MMA	ICLR'20	<b>87.80</b>	43.10	65.450	57.818
AWP	NeurIPS'20	85.63	53.32	69.475	65.718
GAIRAT	ICLR'21	83.00	41.80	62.400	55.599
KD+SWA	ICLR'21	87.45	53.59	70.520	66.456
EWAT	ICML'21	86.00	51.60	68.800	64.500
MAIL	NeurIPS'21	82.20	43.30	62.750	56.721
TE	ICLR'22	85.97	52.88	69.425	65.482
SOVR	ICML'23	85.00	53.10	69.050	65.366
ReBAT	NeurIPS'23	85.25	54.78	70.015	66.700
ADR	ICLR'24	84.67	53.25	68.960	65.381
CURE	ICLR'24	87.05	52.10	69.575	65.186
<b>RPAT<sup>++</sup></b>	<b>Ours</b>	86.76	<b>54.97</b>	<b>70.865</b>	<b>67.300</b>
ReBAT*	NeurIPS'23	86.66	55.64	71.150	67.769
<b>RPAT<sup>++</sup>*</b>	<b>Ours</b>	<b>87.57</b>	<b>55.79</b>	<b>71.680</b>	<b>68.158</b>



(a) Impact of loss function. (b) Impact of hyper-parameter  $\alpha$ .

Figure 4. Two ablation studies of the proposed RPAT on CIFAR-10 with ResNet-18 under  $\ell_\infty$  norm. The *MSE* loss in (a) and the points with  $\alpha = 0.5$  in (b) correspond to the record of “PGD-AT+RPAT” on CIFAR-10 in Table 1.

The results show that the proposed RPAT<sup>++</sup> achieves new SOTA performance for most of the measures beyond previous SOTAs on the accuracy-robustness trade-off problem, which strongly supports the contribution of this work.

## 5.4. Ablation Study

As the proposed RPAT can be implemented upon existing AT methods by simply appending the *Robust Perception* term into their original objectives, the comparison in performance between and after adding RPAT in Section 5.2 and Section 5.3 have achieved a comprehensive ablation study on the effectiveness of RPAT. Therefore, in this section, we mainly investigate the impact of two setting details, respectively the loss function adopted for the *Robust Perception* term (Figure 4 (a)) and the specific value for the hyper-parameter  $\alpha$  (Figure 4 (b)). Besides, we further discuss an optional trick suggested by ReBAT but excluded in the above experiments in Appendix 5.3.

Specifically, just as we expected, different loss functions to calculate the similarity between model perceptions, including *MSE*, *KL*, *JS* and *Cosine*, would not significantly

impact the final performance of RPAT, which suggests that the effectiveness of our method is from the successful conceptual design of *Robust Perception*, instead of relying on specific implementation of the regularization term. Then, for  $\alpha$ , we try 0.2 and 0.8 besides the default value 0.5 on the one hand, and also introduce *beta distribution*  $\beta \sim \text{Beta}(0.75, 0.75)$  as suggested by Zhang et al. [57] on the other hand. Specifically, we try  $\alpha = \min(\beta, (1 - \beta))$  and  $\alpha = \max(\beta, (1 - \beta))$ , namely indicated as “*Beta-*” and “*Beta+*” in (Figure 4 (b)), as they respectively approach adversarial and benign samples more. The results show that the default  $\alpha = 0.5$  is an appropriate choice, which is reasonable because the midpoint is expected to be more representative for the ideal target that *Robust Perception* holds for all  $\alpha \in [0, 1]$ .

## 6. Conclusion

This work studies the current accuracy-robustness trade-off problem in AT. To the best of our knowledge, this is the first work that explicitly attributes the trade-off problem to the over-sufficient (rather than insufficient) learning of hard adversarial samples. We surprisingly reveal that such hard samples are already learned better than those successfully defended in model perception, which means the conventional AT objective is no longer effective to further optimize the model *w.r.t.* these hard samples. Instead, it forces the model to ignore the perturbations, in turn degrading their effectiveness in supporting the establishment of decision boundary to an appropriate location. To mitigate this, our new AT objective, *Robust Perception*, encourages smoother perception transition towards the decision boundary along with perturbations, based on which the proposed RPAT method achieves new SOTA performance. Still, exploring better methods to approximate this ideal objective would be a promising future direction.

## Acknowledgements

This work was supported by the National Natural Science Foundation of China (No. 62471420), Guangdong Basic and Applied Basic Research Foundation (2025A1515012296), and CCF-Tencent Rhino-Bird Open Research Fund.

## References

- [1] Abulikemu Abuduweili and Changliu Liu. Estimating neural network robustness via lipschitz constant and architecture sensitivity. In *Safe and Robust Robot Learning for Operation in the Real World Workshop @ CoRL*, 2024. 2, 5, 11
- [2] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *ICML*, 2018. 2
- [3] Tao Bai, Jinqi Luo, Jun Zhao, Bihan Wen, and Qian Wang. Recent advances in adversarial training for adversarial robustness. In *IJCAI*, 2021. 2, 3
- [4] Battista Biggio, Igino Corona, Davide Maiorca, Blaine Nelson, Nedim Šrndić, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. Evasion attacks against machine learning at test time. In *ECML-PKDD*, 2013. 1
- [5] Tianlong Chen, Zhenyu Zhang, Sijia Liu, Shiyu Chang, and Zhangyang Wang. Robust overfitting may be mitigated by properly learned smoothening. In *ICLR*, 2021. 3, 7
- [6] Minhao Cheng, Qi Lei, Pin-Yu Chen, Inderjit Dhillon, and Cho-Jui Hsieh. Cat: Customized adversarial training for improved robustness. In *IJCAI*, 2022. 2, 3
- [7] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *ICML*, 2020. 2, 6
- [8] Gavin Weiguang Ding, Yash Sharma, Kry Yik Chau Lui, and Ruitong Huang. Mma training: Direct input space margin maximization through adversarial training. In *ICLR*, 2020. 2, 3
- [9] Yinpeng Dong, Qi-An Fu, Xiao Yang, Tianyu Pang, Hang Su, Zihao Xiao, and Jun Zhu. Benchmarking adversarial robustness on image classification. In *CVPR*, 2020. 2
- [10] Yinpeng Dong, Ke Xu, Xiao Yang, Tianyu Pang, Zhijie Deng, Hang Su, and Jun Zhu. Exploring memorization in adversarial training. In *ICLR*, 2022. 2, 3
- [11] Mahyar Fazlyab, Taha Entesari, Aniket Roy, and Rama Chellappa. Certified robustness via dynamic margin maximization and improved lipschitz regularization. In *NeurIPS*, 2023. 2, 5, 11
- [12] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015. 1, 2, 5, 11
- [13] Shruthi Gowda, Bahram Zonooz, and Elahe Arani. Conserve-update-revise to cure generalization and robustness trade-off in adversarial training. In *ICLR*, 2024. 3, 6
- [14] Chuan Guo, Mayank Rana, Moustapha Cisse, and Laurens van der Maaten. Countering adversarial images using input transformations. In *ICLR*, 2018. 1, 2
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *ECCV*, 2016. 6
- [17] Xiaowei Huang, Daniel Kroening, Wenjie Ruan, James Sharp, Youcheng Sun, Emese Thamo, Min Wu, and Xinpeng Yi. A survey of safety and trustworthiness of deep neural networks: Verification, testing, adversarial attack and defence, and interpretability. *Computer Science Review*, 2020. 1
- [18] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. In *UAI*, 2018. 3, 7
- [19] Sekitoshi Kanai, Shin'ya Yamaguchi, Masanori Yamada, Hiroshi Takahashi, Kentaro Ohno, and Yasutoshi Ida. One-vs-the-rest loss to focus on important samples in adversarial training. In *ICML*, 2023. 2, 3, 7
- [20] Minseon Kim, Jihoon Tack, Jinwoo Shin, and Sung Ju Hwang. Entropy weighted adversarial training. In *Adversarial Machine Learning Workshop @ ICML*, 2021. 3
- [21] Vishaal Krishnan, Al Makdah, Abed AlRahman, and Fabio Pasqualetti. Lipschitz bounds and provably robust training by laplacian smoothing. In *NeurIPS*, 2020. 2, 5, 11
- [22] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. 2009. 6
- [23] Fei-Fei Li, Andrej Karpathy, and Justin Johnson. Tiny imagenet dataset. 2015. 6
- [24] Zi Liang, Pinghui Wang, Ruofei Zhang, Nuo Xu, Shuo Zhang, Lifeng Xing, Haitao Bai, and Ziyang Zhou. Merge: Fast private text generation. In *AAAI*, 2024. 1
- [25] Fangzhou Liao, Ming Liang, Yinpeng Dong, Tianyu Pang, Xiaolin Hu, and Jun Zhu. Defense against adversarial attacks using high-level representation guided denoiser. In *CVPR*, 2018. 1, 2
- [26] Weilin Lin, Nanjun Zhou, Yanyun Wang, Jianze Li, Hui Xiong, and Li Liu. Backdoordm: A comprehensive benchmark for backdoor learning in diffusion model. *arXiv preprint arXiv:2502.11798*, 2025. 1
- [27] Feng Liu, Bo Han, Tongliang Liu, Chen Gong, Gang Niu, Mingyuan Zhou, Masashi Sugiyama, et al. Probabilistic margins for instance reweighting in adversarial training. In *NeurIPS*, 2021. 2, 3
- [28] Xin Liu, Yichen Yang, Kun He, and John E Hopcroft. Parameter interpolation adversarial training for robust image classification. *IEEE Transactions on Information Forensics and Security (TIFS)*, 2025. 2, 3
- [29] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018. 2, 3, 6
- [30] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018. 12
- [31] Tianyu Pang, Xiao Yang, Yinpeng Dong, Hang Su, and Jun Zhu. Bag of tricks for adversarial training. In *ICLR*, 2021. 11
- [32] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practi-

- cal black-box attacks against machine learning. In *AsiaCCS*, 2017. 1, 2
- [33] Patricia Pauli, Anne Koch, Julian Berberich, Paul Kohler, and Frank Allgöwer. Training robust neural networks using lipschitz bounds. *IEEE Control Systems Letters*, 2021. 2, 5, 11
- [34] Rahul Rade and Seyed-Mohsen Moosavi-Dezfooli. Reducing excessive margin to achieve a better accuracy vs. robustness trade-off. In *ICLR*, 2022. 2, 3
- [35] Aditi Raghunathan, Sang Michael Xie, Fanny Yang, John Duchi, and Percy Liang. Adversarial training can hurt generalization. In *Identifying and Understanding Deep Learning Phenomena Workshop @ ICML*, 2019.
- [36] Aditi Raghunathan, Sang Michael Xie, Fanny Yang, John Duchi, and Percy Liang. Understanding and mitigating the tradeoff between robustness and accuracy. In *ICML*, 2020. 2, 3
- [37] Sylvestre-Alvise Rebuffi, Sven Gowal, Dan Andrei Calian, Florian Stimberg, Olivia Wiles, and Timothy A Mann. Data augmentation can improve robustness. In *NeurIPS*, 2021. 7
- [38] David Stutz, Matthias Hein, and Bernt Schiele. Confidence-calibrated adversarial training: Generalizing to unseen attacks. In *ICML*, 2020. 3
- [39] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *ICLR*, 2014. 1, 2
- [40] Jihoon Tack, Sihyun Yu, Jongheon Jeong, Minseon Kim, Sung Ju Hwang, and Jinwoo Shin. Consistency regularization for adversarial robustness. In *AAAI*, 2022. 3, 6
- [41] Lue Tao, Lei Feng, Jinfeng Yi, Sheng-Jun Huang, and Songcan Chen. Better safe than sorry: Preventing delusive adversaries with adversarial training. In *NeurIPS*, 2021. 2
- [42] Florian Tramer, Nicholas Carlini, Wieland Brendel, and Aleksander Madry. On adaptive attacks to adversarial example defenses. In *NeurIPS*, 2020. 2
- [43] Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. In *ICLR*, 2019. 2, 3
- [44] Haotao Wang, Tianlong Chen, Shupeng Gui, Tingkui Hu, Ji Liu, and Zhangyang Wang. Once-for-all adversarial training: In-situ tradeoff between robustness and accuracy for free. In *NeurIPS*, 2020. 2, 3
- [45] Yisen Wang, Difan Zou, Jinfeng Yi, James Bailey, Xingjun Ma, and Quanquan Gu. Improving adversarial robustness requires revisiting misclassified examples. In *ICLR*, 2020. 1, 3
- [46] Yanyun Wang, Dehui Du, Haibo Hu, Zi Liang, and Yuanhao Liu. Tsfool: Crafting highly-imperceptible adversarial time series through multi-objective attack. *arXiv preprint arXiv:2209.06388*, 2022. 2
- [47] Yifei Wang, Liangchen Li, Jiansheng Yang, Zhouchen Lin, and Yisen Wang. Balance, imbalance, and rebalance: Understanding robust overfitting from a minimax game perspective. In *NeurIPS*, 2023. 3, 6, 7
- [48] Yanyun Wang, Li Liu, Zi Liang, Qingqing Ye, and Haibo Hu. New paradigm of adversarial training: Releasing accuracy-robustness trade-off via dummy class. *arXiv preprint arXiv:2410.12671*, 2024. 2
- [49] Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. In *NeurIPS*, 2020. 3
- [50] Yu-Yu Wu, Hung-Jui Wang, and Shang-Tse Chen. Annealing self-distillation rectification improves adversarial training. In *ICLR*, 2024. 3
- [51] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, and Alan Yuille. Mitigating adversarial effects through randomization. In *ICLR*, 2018. 1, 2
- [52] Weilin Xu, David Evans, and Yanjun Qi. Feature squeezing: Detecting adversarial examples in deep neural networks. In *NDSS*, 2018. 1, 2
- [53] Shuo Yang and Chang Xu. One size does not fit all: Data-adaptive adversarial training. In *ECCV*, 2022. 2, 3
- [54] Jia-Li Yin, Bin Chen, Wanqing Zhu, Bo-Hao Chen, and Ximeng Liu. Push stricter to decide better: A class-conditional feature adaptive framework for improving adversarial robustness. *IEEE Transactions on Information Forensics and Security (TIFS)*, 2023. 2, 3
- [55] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019. 7
- [56] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016. 6
- [57] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. Mixup: Beyond empirical risk minimization. In *ICLR*, 2019. 8
- [58] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In *ICML*, 2019. 2, 3, 6, 11
- [59] Jingfeng Zhang, Jianing Zhu, Gang Niu, Bo Han, Masashi Sugiyama, and Mohan Kankanhalli. Geometry-aware instance-reweighted adversarial training. In *ICLR*, 2021. 2, 3
- [60] Ziyu Zhou, Gengyu Lyu, Yiming Huang, Zihao Wang, Ziyu Jia, and Zhen Yang. Sdformer: Transformer with spectral filter and dynamic attention for multivariate time series long-term forecasting. In *IJCAI*, 2024. 1