

# Adversarial Fine-tune with Dynamically Regulated Adversary

Pengyue Hou , Ming Zhou , Jie Han , Petr Musilek , and Xingyu Li

*Department of Electrical and Computer Engineering, University of Alberta*  
{pengyue, mzhou4, jhan8, pmusilek, xingyu}@ualberta.ca

**Abstract**—Adversarial training is an effective method to boost model robustness to malicious, adversarial attacks. However, such improvement in model robustness often leads to a significant sacrifice of standard performance on clean images. In many real-world applications such as health diagnosis and autonomous surgical robotics, the standard performance is more valued over model robustness against such extremely malicious attacks. This leads us to the question: to what extent can we improve the robustness of the model without sacrificing standard performance? This work tackles this problem and proposes a simple yet effective transfer learning based adversarial training strategy that disentangles the negative effects of adversarial samples on model’s standard performance. In addition, we introduce a training-friendly adversarial attack algorithm, which facilitates the boost of adversarial robustness without introducing significant training complexity. Extensive experiments show that the proposed approach outperforms previous adversarial training algorithms with the following objective: to improve the robustness of the model while preserving model’s standard accuracy on clean data.

## I. INTRODUCTION

Deep learning has been widely applied to real-world applications with excellent performance, such as autonomous driving, efficient web search and facial recognition. However, recent works have shown that most deep learning models, including state-of-the-art deep neural networks (DNNs), are vulnerable to adversarial attacks [1–4]. These adversarial samples are carefully designed by adding small, human-imperceptible noises to original images, but can lead a well-trained classifier to wrong predictions. To improve model adversarial robustness, many defense strategies, such as data augmentation, gradient masking, adversarial example detection, adversarial training, *etc.*, have been proposed with the aim of finding countermeasures to protect DNNs [4–7]. Particularly, adversarial training is widely accepted as the most effective solution. It incorporates adversarial data in model training and helps build model robustness to adversarial attacks.

Despite its success in improving model robustness to adversarial data, state-of-the-art (SOTA) adversarial training strategies are observed to cause model performance degradation [7–10]. For instance, SOTA adversarial training methods such as TRADES [11] loses about 10% standard accuracy for a 50% adversarial robustness improvement on CIFAR-10 image set. This observation has led to a discussion of the relationship between adversarial robustness and standard generalization (e.g., classification accuracy), with a central debate on whether

accuracy and robustness are intrinsically in conflict. Some studies, for example, the works by Tsipras *et al.* [12] and Raghunathan *et al.* [13], claim that the trade-off of model accuracy and adversarial robustness is unavoidable due to the nature of DNNs. In contrast, Raghunathan *et al.* [14, 15] argue that the robustness-accuracy trade-off could disappear with unlimited data. Yang *et al.* [16] present a theoretical analysis, as well as a proof-of-concept example, showing that this trade-off is not inherent and arguing that the observed accuracy-robustness trade-off is introduced by limitations in current adversary training methods. More recently, Xie *et al.* [17] leverage adversarial samples to improve model accuracy by introducing an auxiliary batch normalization layer particularly designed for adversarial samples. However, this study does not discuss if the model robustness is improved or not. To summarize, although adversarial training improves model robustness against adversarial attacks, how to achieve this goal without trading off model accuracy on clean data in adversarial training is still an open question and remains under explored.

It should be noted that a big loss of standard performance on clean data is unacceptable in many applications that might cause severe consequences. Instead, the standard performance is more valued than model robustness against the extremely malicious adversarial attacks. Such applications include medical diagnosis, autonomous surgical robotics, *etc.* This leads to the question: To what extent can we boost model robustness without sacrificing standard performance? Unlike prior adversarial training works that allow model performance degradation on clean data, we investigate if it is feasible to improve adversarial performance without any standard performance loss. Specifically, this paper proposes a pre-train based adversarial training strategy, where adversarial training is applied to clean-data (vanilla-trained) models. To prevent model catastrophic forgetting in model refining, we follow the replay-based strategy and maintain the 1:1 clean-adversarial data ratio in model refinement. To further reduce adversarial samples’ negative impacts on standard accuracy, we incorporate novel dynamically regulated adversarial (DRA) samples in model adversarial refinement. Unlike most adversarial attacks that add adversarial noise to every image pixel, DRA searches for highly stimulated adversarial features and generates adversarial samples accordingly. Such strategy in adversarial sample generation enforces the model refinement to learn descriptive but non-robust features. Extensive experi-

mentation shows that the proposed adversarial training strategy improves model adversarial robustness with a large margin, but without standard performance loss. The contributions of this study are summarized as follows.

- We propose a simple, yet generic adversarial training strategy that improves adversarial robustness by a large margin without sacrificing standard accuracy. We argue that the proposed method is significant for applications where standard performance is more valued than adversarial robustness.
- We further introduce an unbounded adversarial attack method, namely DRA. It introduces smaller image distortions and facilitates the adversarial refinement to focus on the learning of most descriptive but non-robust features.
- We show that our adversarially trained models exhibit robustness not only to adversarial samples; but also against naturally corrupted images, which suggests its potential for real-world applications.

The rest of this paper is organized as follows. Section II presents a brief review of related works. Section III formulates the target problem and specifies our motivations. The technical details of the proposed method are elaborated in Section IV. Section V presents extensive experiments and discussions, then followed by our conclusions in section VI.

## II. RELATED WORKS

**Adversarial attacks** aim to add small, carefully-designed, human-imperceptible noises to clean data to fool a machine learning model. Depending on the availability of prior knowledge of the target, adversarial attacks can be categorized into two types: white box attacks and black box attacks. Since adversarial samples are transferable, many white box attacks follow the paradigm that exploits a surrogate white box model to generate attack samples [6, 18, 19]. This section briefly reviews the first-order white-box adversarial attacks, due to their close relevance to this study.

Many white box attacks belong to first-order adversarial methods. Fast Gradient Method (FGM) is an intuitive back-propagation based method to generate adversarial samples [2, 20] and evolves into a better constraint method, namely Fast Gradient Sign Method (FGSM) [3]. FGSM searches for adversarial samples following the direction of the gradients of current parameter values and projects them on images.

$$x_{adv} = x + \epsilon \cdot \text{sign}(\nabla_x L(x, y; \theta)), \quad (1)$$

where  $\epsilon$  is the noise constraint for projecting the attack noises onto  $L_\infty$  ball,  $x$  and  $y$  are the clean input samples and their labels in dataset  $D$  and  $L(\cdot)$  is the loss function.

Iterative Fast Gradient Sign Method (I-FGSM), also called Basic Iterative Method (BIM) [21], is a multi-step variant of FGSM. It provides stronger adversarial examples but with more computational cost. Projected Gradient Descent (PGD) [4] is another iterative variant of FGSM with a random starting point.

$$x_{adv}^t = \Pi_S[x_{adv}^{t-1} + \epsilon \cdot \text{sign}(\nabla_x L(x_{adv}^{t-1}, y; \theta))], \quad (2)$$

where  $t$  denotes the number of iteration and  $\Pi_S$  is the clipping function that forces the value to reside in a predetermined range. PGD is widely agreed to be one of the strongest attacks and often used to benchmark model robustness.

**Adversarial training** incorporates adversarial samples in model optimization and is considered a very effective method to boost model robustness. Particularly, Madry *et al.* [7] formulate adversarial training as a MinMax optimization problem,

$$\mathcal{L}(\theta) = \mathbb{E}_{(x,y) \sim D} [\max_{\delta \in S} L(x + \delta, y; \theta)], \quad (3)$$

Where  $\delta$  is the perturbation superimposed on the input, and  $\theta$  represents network parameters. The inner maximization tries to find adversarial samples that maximize the loss within the allowed permutation range, while the outer minimization tries to minimize the adversarial loss.

Since Madry's MinMax optimization is prone to over-fit adversarial samples, many adversarial training methods mix clean data and adversarial data in training.

$$\mathcal{L}(\theta) = \mathbb{E}_{(x,y) \sim D} [L(x, y; \theta) + \max_{\delta \in S} L(x + \delta, y; \theta)]. \quad (4)$$

Note, in (4), adversarial examples is leveraged to regularize the vanilla training on clean data. Recently, Raghunathan *et al.* [14] introduce robust self-training (RST) to balance the vanilla and adversarial loss by a regularization parameter  $\beta > 0$ .

$$\mathcal{L}(\theta, \beta) = \mathbb{E}_{(x,y) \sim D} [L(x, y; \theta) + \beta \max_{\delta \in S} L(x + \delta, y; \theta)]. \quad (5)$$

Another regularized adversarial training strategy, TRADES [11], is proposed to boost model robustness following the Locally-Lipschitz smoothness constraint.

$$\begin{aligned} \mathcal{L}(\theta, \lambda) &= \mathbb{E}_{(x,y) \sim D} [L(f_\theta(x), y) \\ &+ 1/\lambda \cdot \max_{\delta \in S} L(f_\theta(x + \delta), f_\theta(x))], \end{aligned} \quad (6)$$

where  $f_\theta$  denotes the training model parameterized by  $\theta$ . Unlike RST computing an adversarial loss between the prediction  $f_\theta(x + \delta)$  and label  $y$  as the regularization term in (5), TRADES regularizes the training by calculating  $L(f_\theta(x + \delta), f_\theta(x))$  from a pair of clean sample  $x$  and its adversarial version  $x + \delta$ . The regularization parameter  $\lambda$  determines the trade-off between accuracy and robustness in the overall optimization. It is worth notice that small regularization parameter helps the model emphasize more on accuracy over robustness which closely align to our objective of adversarial training. However, we show later in Section V that even very small value of  $\lambda^{-1}$  cannot guarantee models to achieve comparable accuracy to vanilla trained models.

Pre-training is a popular training framework that can help reduce training time or improve accuracy performance for fine-tuning downstream tasks. Jeddi *et al.* [22] start with a clean data pre-trained model and fine-tune with PGD adversarial training with the aim of reducing the time cost and overfitting issue [23]. Hendrycks *et al.* [24] adversarially pre-train their model on a downsampled ImageNet and apply adversarial fine-tuning which can significantly improve model robustness

**TABLE I:** Performance of adversarial training methods on MNIST and CIFAR-10. Adversarial accuracy is evaluated with PGD attacks.

	MNIST ( $\epsilon = 0.3$ )		CIFAR-10 ( $\epsilon = 8/255$ )	
	$\mathcal{A}_{std}$	$\mathcal{A}_{rob}$	$\mathcal{A}_{std}$	$\mathcal{A}_{rob}$
Vanilla	99.3%	0.3%	93.0%	0.0%
Madry's [7]	99.2%	95.6%	87.3%	47.0%
Trades( $1/\lambda = 1$ ) [11]	99.3%	94.1%	86.6%	44.3%
Trades( $1/\lambda = 6$ ) [11]	99.3%	96.0%	81.2%	53.5%
MART [28]	99.1%	96.2%	83.4%	52.8%

on CIFAR datasets compare with adversarial training from scratch. Chen *et al.* [25] show that self-supervised pre-training such as Selfie [26], Jigsaw [27], also lead to better robustness than traditional adversarial training. In contrast to previous adversarial pre-training strategies, our approach mainly focuses on maintaining accuracy and treats adversarial robustness as an added bonus. More specifically, we incorporate a novel adversary generating method (DRA) to facilitate our goal by reducing the learning complexity of adversarial training.

### III. PRIMITIVES

#### A. Problem Formulation

In this study, we focus on improving model robustness to imperceptible, in-distribution adversarial samples defined by Fawzi *et al.* [29]. Briefly, assume that clean data follows a distributing  $D$ , in-distribution adversarial samples ( $x'$ ) can be also roughly described within  $D$ .

To answer the question: to what extent can we boost model robustness without sacrificing standard performance? We formulate the problem as


$$\mathcal{L}(\theta) = \mathbb{E}_{(x,y) \sim D} [\max_{\delta \in S} L(x + \delta, y; \theta)] \quad (7)$$

$$\text{s.t. } \mathbb{E}_{(x,y) \sim D} [L(x, y; \theta)] \leq \mathbb{E}_{(x,y) \sim D} [L(x, y; \theta_{std})],$$

where  $\theta_{std} = \underset{\theta}{\operatorname{argmin}} \mathbb{E}_{(x,y) \sim D} [L(x, y; \theta)]$  represents the model with parameter  $\theta_{std}$  that yields the minimized standard loss. Comparing the new problem in (7) with previous adversarial training methods in (3-6), the new regularization term in (7) explicitly defines the behavior of the model: to improve adversarial robustness without the loss of model's standard performance.

#### B. Motivation

To tackle the question in (7), we re-examine SOTA adversarial training strategies and obtain an interesting observation. Yang *et al.* [16] show that many real image sets, such as MNIST, CIFAR-10, SVHN and Restricted ImageNet, are  $r$ -separated, with the smallest inter-category sample distance being no smaller than  $2r$ . Furthermore, their empirical separation distance is  $3x-7x$  larger than the typical adversarial perturbation constraint  $\epsilon$  adopted in prior arts, i.e.  $\epsilon < r$ . In theory, any  $r$ -separated dataset has more than one classifier that are both accurate and robust up to perturbations of size  $r$ . This  $r$ -separated claim aligns well with adversarial robustness experiments on the MNIST dataset in previous adversarial training studies [7, 11, 28]. However, as summarized in Table I



**Fig. 1:** A conceptual demonstration of our adversarial fine-tune method. Classifier 1 represents the clean data pre-trained classifier which is accurate but not robust to adversarial samples. Our adversarial fine-tune method seeking for both accurate and robust classifier by pushing the classifier 2 out of the yellow adversarial regions defined by  $\delta \in S$ .

where the standard performance and adversarial performance are denoted by  $\mathcal{A}_{std}$  and  $\mathcal{A}_{rob}$  respectively, on CIFAR-10 which is a more complicated dataset, previous adversarial training often leads to around 10% standard accuracy drop for models to gain desired adversarial robustness.


From the above observation, we hypothesize that the observed trade-off between adversarial robustness and standard accuracy is due to the limitation of model capacity. The problem of image classification on the MNIST dataset is relatively easy and the models adopted in previous adversarial training studies have enough capacity to jointly benefit from standard and adversarial samples following (3-6). However, for complex problems such as classification on CIFAR-10 dataset, conventional adversarial training in (3-6) increases the difficulty in model optimization due to the noisy representation of adversarial perturbations. In another word, learning both standard features and adversarial features (i.e. those non-robust, yet highly predictive patterns [30]) is beyond the capacity of those models (e.g. ResNet-18, ResNet-50, etc.).

Under the hypothesis that model capacity is not enough to simultaneously learn standard and adversarial features, we value clean-data accuracy over adversarial robustness. Thus, we made two modifications to existing adversarial training strategies to tackle the problem formulated in (7). Briefly, we propose a transfer learning based adversarial training strategy that starts with a clean data pre-trained model which already has strong generalizability on clean data. To further reduce the learning complexity, we incorporate easy-to-learn adversarial samples in the adversarial fine-tuning. Please refer to the next section for details of the proposed method.

## IV. METHODOLOGY

### A. Transfer Adversarial Training

SOTA adversarial training methods usually train a robust model from scratch using either adversarial data only in (3) or a combination of clean and adversarial data following (4-6). However, none of them guarantee that the standard performance will be preserved. Indeed, due to the highly non-



**Fig. 2:** Systematic diagram of the proposed adversarial transfer adversarial training strategy, where we use various ResNets as the backbone. From left to right: DRA, a proposed method to generate adversarial samples. It filters out negligible adversarial noises and reduces adversarial training complexity. The orange block and green block represent standard training and robust training, respectively. In standard training, we train a model with  $\theta_{std}$  on clean data that yields high standard performance. Then the robust training aims to find a better  $\theta$  in the vicinity of  $\theta_{std}$  to boost adversarial robustness without standard performance loss.

convex loss surface in model optimization, optimizing both targets simultaneously may be at odds with each other [12].

To solve the primal conditioned optimization problem in (7), we apply the Karush-Kuhn-Tucker (KKT) approach and obtain a dual unconstrained optimization problem by introducing a KKT coefficient  $\lambda$ ,

$$\begin{aligned} \mathcal{L}(\theta, \lambda) &= \mathbb{E}_{(x,y) \sim D} [\max_{\delta \in S} L(x + \delta, y; \theta)] \\ &+ \lambda \mathbb{E}_{(x,y) \sim D} [L(x, y; \theta) - L(x, y; \theta_{std})]. \end{aligned} \quad (8)$$

Therefore, the solution to the dual problem

$$\min_{\theta} \max_{\lambda, \lambda > 0} \mathcal{L}(\theta, \lambda) \quad (9)$$

is identical to the solution of the primal problem in (7). Note that to solve the problem in (9), we need the margin value  $\mathbb{E}_{(x,y) \sim D} [L(x, y; \theta_{std})]$  which is fixed before the adversarial training described in (9). That is, a model with  $\theta_{std}$  is already obtained for high standard performance. Therefore, instead of training a robust model with  $\theta$  from scratch, we introduce a transfer learning strategy to solve (9) and propose to search a new  $\theta$  from  $\theta_{std}$ , as the standard performance is unlikely to change severely in the small vicinity of  $\theta_{std}$ . Fig. 1 visualizes the conceptual idea of the proposed adversarial training strategy on a  $r$ -separated dataset. Classifier 1 is vanilla-trained over clean data. It is accurate but not adversarial robustness. By pushing the classification boundary out of the yellow adversarial regions defined by  $\delta \in S$ , the obtained classifier 2 has the identical standard performance with improved adversarial robustness.

We depict the detailed systematic diagram of the proposed adversarial training strategy in Fig. 2. Specifically, the proposed training strategy divides the training into two phases: vanilla standard training and adversarial robust training. In standard training, we exploit clean data to train an accurate

model,  $\theta_{std} = \underset{\theta}{\operatorname{argmin}} \mathbb{E}_{(x,y) \sim D} [L(x, y; \theta)]$ . The standard training has two benefits. First, it provides the margin value  $\mathbb{E}_{(x,y) \sim D} [L(x, y; \theta_{std})]$  in (9). Second, due to inherent transfer learning property, the downstream adversarial robust training is more cost-efficient than SOTA adversarial training strategies that often require large training loads and long training time to handle the complexity introduced by adversarial features.

With the clean-data pre-trained model, adversarial samples are incorporated in the model refinement phase. To prevent model catastrophic forgetting in model fine-tuning, we follow the replay-based strategy and let the network iteratively update upon clean and adversarial samples. Note, unlike conventional fine-tuning tends to unfreeze several outer layers to preserve the knowledge learned from the source task, we argue that it is important to update all parameters in robust training stage. Briefly, adversarial noises propagate through each layer in the model and aggregate into large prediction distortions. All layers of a robust model must contribute to the defend against adversarial attacks.

### B. Dynamically Regulated Adversary in Adversarial Training

Adversarial training is usually referred to as the "MinMax" optimization game, where the adversarial samples are significant contributors. In fact, aggressive adversarial attacks are preferred in adversarial training, because models trained on aggressive adversarial attacks are more resistant to weaker adversaries. Therefore previous studies have usually adopted the PGD attack in adversarial training.

A good adversarial attack approach for adversarial training should be aggressive but without greatly increasing training complexity. Although PGD is an aggressive solution for introducing noise, we argue that PGD is not the best candidate for adversarial training in the sense that it introduces excessive noisy information in model robustness training. More speci-

cally, PGD uses a  $\text{Sign}()$  clipping method to project adversarial noises onto the  $L_\infty$  ball. It treats all image pixels equally and applies the same noise injection strategy to all pixels, regardless of their contribution to the prediction results. We argue that the treatment of adding adversarial noise to all pixels in PGD significantly increases the training complexity. We will show in the experimentation that the proposed DRA attack is a better candidate than PGD in adversarial training and facilitates various adversarial training.

This work introduces a novel gradient-based attack method, namely DRA attack, in Algorithm 1. Unlike PGD, which treats image pixels equally, DRA distinguishes important pixels from others by aggressively modifying only those image features that are highly predictive, but non-robust. In this way, DRA adversaries enforce the robustness training by focusing more on these predictive features, thus helping to improve model’s robustness against adversarial attacks. Specifically, DRA quantifies pixel significance by the gradient of the loss function with respect to its pixel value. A large gradient value suggests that the pixel contributes more to the image prediction. In this regard, DRA abandons the  $\text{Sign}(\cdot)$  method so that it can smoothly search for the optimal adversarial samples along the gradient. In addition, when generating adversarial samples, DRA uses a more resilient distance metric  $L_1$  to bound the adversarial noise instead of using the  $L_\infty$  constraint.

Fig. 3 presents a visual comparison of the DRA and PGD adversaries. Compared to PGD, DRA introduces stronger noise in the "cat" region, which is the most predictive pattern in the image. We claim that the lower noise level in the image background makes the learning more complex. We show in the experimental section that DRA samples help various training strategies achieve higher adversarial robustness than PDG for the same adversarial aggressiveness.

---

#### Algorithm 1: DRA attack

---


**Input :** Clean-data pre-trained model  $f$ , loss function  $L$ , clean training pairs  $(x, y)$ , regulation term  $\epsilon$ , significant pixel percentage  $p$ , iteration  $T$

**Output:** Adversarial sample  $x^*$

- 1:  $\alpha = \epsilon/T$ ;
  - 2:  $x_0^* = x, t = 0$ ;
  - 3: **if**  $t < T$  **then**
    - Input  $x_t^*$  to  $f$  and obtain gradient  $\nabla_x L(x_t^*, y)$ ;
    - update  $x_{t+1}^* = x_t^* + \frac{\alpha \cdot \nabla_x L(x_t^*, y)}{\|\nabla_x L(x_t^*, y)\|_1}$
  - end**
  - 4: Rank current pixel values based on  $\nabla_x L(x_t^*, y)$  and set threshold  $T_h$  for the  $p$  most significant values;
  - 5: **if**  $x_{ij}^* < T_h$  **then**
    - $x_{ij}^* = x_{ij}$ ;
  - end**
- 

## V. EXPERIMENTS

In this section, we present extensive experiments to evaluate the accuracy-robustness performance of the proposed adversar-



**Fig. 3:** A visual comparison of PGD (left) and DRA (right) adversaries. They are both able to fool DNNs with imperceptible noises, however, the overall noise budget of DRA is smaller. Furthermore, DRA focuses on highly discriminate pixels (patterns that contribute most to final prediction outcomes), where PGD equally distributes adversarial noises over the whole image. As we can see from the example, DRA noises align well with the salient map of the cat.

ial training strategy on MNIST and CIFAR-10 datasets.

### A. Experimental Setup

In the proposed adversarial training strategy, we first train a model on clean data only so that the model has high standard performance. The model is then fine-tuned with clean and DRA samples. In both training processes, we utilize image augmentation including random cropping and random flipping. In addition, the image batch size is set to 128, SGD with momentum to 0.9 and weight decay to  $2e - 4$  for model optimization.

On the MNIST dataset, we follow the TRADES study and use a simple CNN model with 2 convolutional layers and 2 fully connected layers. We set  $\epsilon = 0.1$ ,  $p = 2/3$  and iterate 20 times to generate DRA samples for 50 epochs of robust fine-tuning.

On the CIFAR-10 dataset, we use ResNet as the backbone model and perform 60 epochs of adversarial fine-tune in the robust training phase. Adversarial samples used for our robust training are generated in 5 iterations. The noise constraint is linearly decayed from 2 to 0.5 in the first 50 epochs, and only clean images are fine-tuned in the last 10 epochs.

For evaluation purposes, adversarial samples are generated by PGD under various noise constraints (e.g.  $\epsilon = 0.1, 0.3$  out of 1 for MNIST images and  $\epsilon = 2, 5, 8$  out of 255 for CIFAR-10) in 20 iterations with step size  $\epsilon/20$ , unless otherwise specified. We compare our DRA fine-tuned model with TRADES [11] on both datasets using relatively small regularization parameter ( $1/\lambda$ ). We find that this setting is consistent with our objective, as the small regularization

**TABLE II:** Accuracy-Robustness performance against PGD attacks on MNIST. Note, the target is to improve adversarial robustness without sacrificing standard performance.

Model	natural	PGD $_{\epsilon=0.1}$	PGD $_{\epsilon=0.3}$
Vanilla trained	99.3%	88.3%	18.6%
Trades( $1/\lambda = 1$ )	99.3%	98.4%	93.5%
Trades( $1/\lambda = 0.5$ )	99.3%	97.9%	92.1%
Trades( $1/\lambda = 0.1$ )	99.4%	97.2%	90.8%
Ours	99.3%	98.0%	95.9%

**TABLE III:** Accuracy-Robustness performance against PGD attacks on CIFAR-10, with ResNet-50 as the backbone model.

Model	natural	PGD $_{\epsilon=2}$	PGD $_{\epsilon=5}$	PGD $_{\epsilon=8}$
Vanilla trained	93.7%	45.0%	15.9%	5.6%
Trades( $1/\lambda = 0.05$ )	91.4%	53.5%	25.8%	8.6%
Trades( $1/\lambda = 0.01$ )	92.6%	49.4%	19.2%	6.3%
Ours	93.8%	64.9%	31.0%	10.9%

parameter intuitively enforces TRADES to primarily optimize the clean data accuracy while "trading off" a small amount of adversarial robustness.

### B. Accuracy-Robustness Performance

The numerical results on the MNIST are shown in Table II, where "natural" indicates the standard performance of the model on clean data and "vanilla trained" indicates that the model is trained with clean data only. On the MNIST dataset, both TRADES and our method maintain high standard accuracy while improving model's adversarial robustness; the proposed method obtains 2%-5% higher robustness than TRADES with various settings.

For the more complex dataset, CIFAR-10, Table III shows the results with ResNet-50 as the backbone. Unlike the MNIST dataset, ResNet-50 with the proposed adversarial training strategy significantly outperforms TRADES in terms of adversarial robustness. Note that in our experiment, the maximal  $\epsilon$  of DRA in adversarial training is 2, which corresponds to  $\epsilon = 4 - 5$  in PDG attacks. We believe that training the model with a large DRA  $\epsilon$  would further improve the robustness.

In addition, we vary the backbone models for the CIFAR-10 experiments and report the accuracy-robustness performance in Table IV. First, for clean CIFAR-10 image classification, vanilla-trained models and our adversarial fine-tuned models achieve comparable performance in all examined backbone models. In particular, ResNet-50 and ResNet-101 can even outperform their vanilla training counterparts in terms of clean data accuracy. Second, to defend against PGD attacks, we observe that models with larger capacity are able to boost their adversarial robustness to a larger margin, which supports our hypothesis discussed in section 3.

### C. Effect of DRA on Adversarial Training

Generation of strong adversarial samples (with higher attack success rates) is a critical factor in adversarial training. In this experiment, we investigate the impact of different types of adversarial samples in three different adversarial training strategies: Madry's [7], TRADES [11], and ours. Concretely, we replace PGD samples with DRA data in Madry's, and

**TABLE IV:** Accuracy-Robustness performance against PGD attacks on CIFAR-10, with various backbone models.

Models	natural	PGD $_{\epsilon=2}$	PGD $_{\epsilon=5}$	PGD $_{\epsilon=8}$
ResNet-18 (vanilla)	93.1%	44.1%	14.6%	5.8%
	(-0.7%)	(+16.3%)	(+10.9%)	(+2.7%)
ResNet-18 (ours)	92.4%	60.4%	25.5%	8.5%
ResNet-34 (vanilla)	93.3%	46.6%	15.3%	4.5%
	(-0.4%)	(+14.8%)	(+12.6%)	(+4.4%)
ResNet-34 (ours)	92.9%	61.4%	27.9%	8.9%
ResNet-50 (vanilla)	93.7%	45.0%	15.9%	5.6%
	(+0.1%)	(+19.9%)	(+15.1%)	(+5.3%)
ResNet-50 (ours)	93.8%	64.9%	31.0%	+10.9%
ResNet-101 (vanilla)	93.7%	46.9%	15.9%	6.8%
	(+0.1%)	(+22.9%)	(+16.3%)	(+7.4%)
ResNet-101 (ours)	93.8%	69.8%	32.3%	+14.2%

**TABLE V:** ResNet-50 trained with PGD Vs. DRA on CIFAR-10


Model	Clean data	PGD $_{\epsilon=2}$	PGD $_{\epsilon=5}$	PGD $_{\epsilon=8}$
Madry's + PGD	88.8%	66.1%	30.9%	11.0%
	(+0.6%)	(+0.6%)	(+0.7%)	(+1.5%)
Madry's + DRA	89.4%	66.7%	31.6%	12.5%
TRADES + PGD	87.9%	65.0%	34.3%	12.1%
	(-0.7%)	(+3.3%)	(+4.8%)	(+1.8%)
TRADES + DRA	87.2%	68.3%	39.1%	13.9%
ours + PGD	90.6%	59.2%	27.0%	10.1%
	(+0.7%)	(+0.8%)	(+1.9%)	(-0.3%)
ours + DRA	91.3%	60.0%	28.9%	9.8%

TRADES. Similarly, we use PDG in the proposed method and compare its performance with DRA. We note that DRA is a stronger adversary than PGD for the same  $\epsilon$  conditions due to the presence of soft-bounded constraint. To make a fair comparison, we choose different  $\epsilon$  for them so that DRA-generated samples and PGD-generated samples can have similar attack strengths. Specifically, we set  $\epsilon = 1$  for DRA and  $\epsilon = 2.55$  for PGD, as both settings resulted in a robust accuracy of 44.5% on vanilla-trained ResNet-50.

Table. V reports CIFAR-10 classification performance with ResNet-50 as the backbone model. We observe that in all three settings, DRA is more beneficial for improving model robustness. Furthermore, the clean data accuracy of DRA trained models is higher than that of the PGD trained models using Madry's and our method. Since the loss function of TRADES (Eqn.6) is essentially designed to improve robustness through trading accuracy, we believe it is reasonable to assume that the standard accuracy of the DRA trained model is slightly lower than that of the model trained with PGD. In short, Table. V with numerical results validate our hypothesis that DRA is a better adversary that benefits model's adversarial training.

### D. Ablation on DRA Hyperparameter Setting

Our DRA attack algorithm filters out unimportant pixels by a pre-fixed percentage and applies adversarial noises only to important image features that contribute to the final prediction. In this ablation study, we investigate the effect of hyperparameter settings (e.g. the value of significant feature percentage  $p$  and noise budget  $\epsilon$  in Algorithm 1) on CIFAR-10 dataset. Specifically, CIFAR-10 images are in size of  $3 \times 32 \times 32$ . We



**Fig. 4:** Ablation on DRA hyperparameter settings: DRA’s attack success rates versus the significant feature percentage  $p$  on CIFAR-10 vanilla-trained ResNet models. A marginal improvement on attack success rate is observed when  $p > 1/3$ .

**TABLE VI:** Ablation on DRA hyperparameter settings: adversarial training performance versus different threshold value  $p$ .

$p$	natural	PGD $_{\epsilon=2}$	PGD $_{\epsilon=5}$	PGD $_{\epsilon=8}$
0	93.7%	45.0%	15.9%	5.6%
1/6	91.8%	58.5%	28.0%	9.2%
1/3	91.3%	60.0%	28.9%	9.8%
1/2	90.4%	61.4%	29.1%	10.0%
1	89.6%	61.9%	29.3%	10.1%

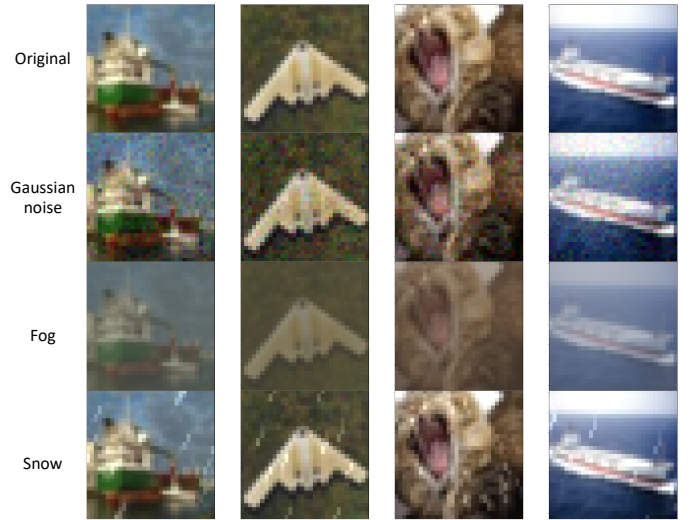
treat the values of red, green and blue independently and thus obtain  $N = 3072$  pixel values per image. To comprehensively study this problem, we vary the values of  $p$  and  $\epsilon$  and report DRA’s attack success rates on CIFAR-10 vanilla-trained ResNet-18, ResNet-34 and ResNet-50.

Fig. 4 reports adversarial attack success rate versus the percentage  $p$  out of the 3072 pixels. With the settings  $\epsilon = 1$ , we note that the successful attack rate of DRA grows linearly in the most significant 500 (i.e.  $p = 1/6$ ) pixels and saturates in about 1000 ( $p = 1/3$ ) pixels. For  $\epsilon = 2$ , the successful attack rate also almost saturates at approximately 1000 ( $p = 1/3$ ) pixels.

To further investigate the impact of DRA thresholds on the overall adversarial training method, we train ResNet-50 with different values of  $p$  and report the performance in Table. VI. We notice that a higher threshold value  $p$  does contribute to better model robustness, however, the improvement becomes marginal as  $p > 1/3$ . In particular, increasing  $p$  from 1/2 to 1 only leads to less than 0.5% robustness increment for PGD attacks, but causes the downgrade of natural accuracy by 0.8%. These results also support our claim that introducing too much unnecessary noise in the images complicates model optimization and tends to lead to a standard performance loss. Similar to the results in Fig. 4, Table VI indicates that  $p = 1/3 - 1/2$  is a good setting for generating adversarial samples in the proposed method.

### E. Evaluation on Naturally Corrupted Images

Corrupted data with naturally occurring perturbations and distributed shifts pose challenges to model generalization. In this experiment, we explore the potential of DRA fine-tuned models on corrupted data. In this experiment, we train a ResNet-50 model with our adversarial fine-tuning strategy and evaluate its performance on the corrupted CIFAR-10



**Fig. 5:** Corrupted image samples in *Cifar-10-C* [31].

**TABLE VII:** The accuracy of DRA trained ResNet-50 vs. Vanilla trained ResNet-50 over different corrupted types.

Corrupted Type	ResNet-50 (DRA Trained)	ResNet-50 (Vanilla trained)
Snow	85.37%	84.10%
Frost	84.4%	81.41%
Zoom_blur	84.97%	82.13%
Motion_blur	76.89%	74.08%
JPEG compression	89.51%	84.51%
Gaussian noise	75.85%	50.92%

dataset (*CIFAR-10-C* [31]). In *CIFAR-10-C*, the clean *CIFAR-10* data are processed to mimic various image distortions under harsh conditions. Table VII presents classification accuracy on the *CIFAR-10-C* dataset with the DRA refined model and the vanilla trained model. Results show that the DRA fine-tuned model exhibits stronger robustness to different types of corruption.

## VI. CONCLUSION

In this paper, we aim to tackle a unique problem in adversarial training: improving the adversarial robustness of a model without sacrificing the standard performance. We explicitly formulate the problem and propose a cost-efficient adversarial training strategy. It decomposes adversarial training into two phases: standard training and robust training. In addition, we introduce a training-friendly adversary to further benefit adversarial training. Extensive experimentation on MNIST and *CIFAR-10* datasets suggests that the proposed adversarial training strategy serves better for the target objective.

## REFERENCES

- [1] A. Nguyen, J. Yosinski, and J. Clune, “Deep neural networks are easily fooled: High confidence predictions for unrecognizable images,” in *Conference on Computer Vision and Pattern Recognition*, 2015, pp. 427–436.

- [2] C. Szegedy *et al.*, “Intriguing properties of neural networks,” in *International Conference on Learning Representations*, 2014.
- [3] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” 2015.
- [4] A. Kurakin, I. Goodfellow, and S. Bengio, *Adversarial machine learning at scale*, 2017.
- [5] F. Tramèr, A. Kurakin, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel, “Ensemble adversarial training: Attacks and defenses,” 2018.
- [6] Y. Dong *et al.*, “Boosting adversarial attacks with momentum,” in *Conference on Computer Vision and Pattern Recognition*, 2018, pp. 9185–9193.
- [7] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards deep learning models resistant to adversarial attacks,” in *arXiv preprint arXiv:1706.06083*, 2017.
- [8] Y. Carmon, A. Raghunathan, L. Schmidt, P. Liang, and J. C. Duchi, “Unlabeled data improves adversarial robustness,” in *Conference on Neural Information Processing Systems*, 2019.
- [9] M. Cheng, Q. Lei, P.-Y. Chen, I. Dhillon, and C.-J. Hsieh, “Cat: Customized adversarial training for improved robustness,” in *arXiv preprint arXiv:2002.06789*, 2020.
- [10] Y. Balaji, T. Goldstein, and J. Hoffman, “Instance adaptive adversarial training: Improved accuracy tradeoffs in neural nets,” in *arXiv preprint arXiv:1910.08051*, 2019.
- [11] H. Zhang, Y. Yu, J. Jiao, E. Xing, L. El Ghaoui, and M. Jordan, “Theoretically principled trade-off between robustness and accuracy,” in *International Conference on Interactive Mobile Communication, Technologies and Learning*, Proceedings of Machine Learning Research, 2019, pp. 7472–7482.
- [12] D. Tsipras, S. Santurkar, L. Engstrom, A. Turner, and A. Madry, “Robustness may be at odds with accuracy,” in *arXiv preprint arXiv:1805.12152*, 2018.
- [13] A. Raghunathan, S. M. Xie, F. Yang, J. C. Duchi, and P. Liang, “Adversarial training can hurt generalization,” in *arXiv preprint arXiv:1906.06032*, 2019.
- [14] A. Raghunathan, S. M. Xie, F. Yang, J. Duchi, and P. Liang, “Understanding and mitigating the tradeoff between robustness and accuracy,” in *Proceedings of Machine Learning Research*, Proceedings of Machine Learning Research, 2020, pp. 7909–7919.
- [15] D. Stutz, M. Hein, and B. Schiele, “Disentangling adversarial robustness and generalization,” in *Conference on Computer Vision and Pattern Recognition*, Conference on Computer Vision and Pattern Recognition, 2019, pp. 6976–6987.
- [16] Y.-Y. Yang, C. Rashtchian, H. Zhang, R. Salakhutdinov, and K. Chaudhuri, “A closer look at accuracy vs. robustness,” in *Conference on Neural Information Processing Systems*, vol. 33, 2020.
- [17] C. Xie, M. Tan, B. Gong, J. Wang, A. L. Yuille, and Q. V. Le, “Adversarial examples improve image recognition,” in *Conference on Computer Vision and Pattern Recognition*, 2020, pp. 819–828.
- [18] Y. Liu, X. Chen, C. Liu, and D. Song, “Delving into transferable adversarial examples and black-box attacks,” in *International Conference on Learning Representations*, 2017.
- [19] S. Chen, Z. He, C. Sun, J. Yang, and X. Huang, “Universal adversarial attack on attention and the resulting dataset damagenet,” in *IEEE Trans. Pattern Analysis and Machine Intelligence*, 2020.
- [20] R. Huang, B. Xu, D. Schuurmans, and C. Szepesvári, “Learning with a strong adversary,” in *arXiv preprint arXiv:1511.03034*, 2015.
- [21] A. Kurakin, I. Goodfellow, S. Bengio, *et al.*, *Adversarial examples in the physical world*, 2017.
- [22] A. Jeddi, M. J. Shafiee, and A. Wong, “A simple fine-tuning is all you need: Towards robust deep learning via adversarial fine-tuning,” *arXiv preprint arXiv:2012.13628*, 2020.
- [23] L. Schmidt, S. Santurkar, D. Tsipras, K. Talwar, and A. Madry, “Adversarially robust generalization requires more data,” *Advances in neural information processing systems*, vol. 31, 2018.
- [24] D. Hendrycks, K. Lee, and M. Mazeika, “Using pre-training can improve model robustness and uncertainty,” in *International Conference on Machine Learning*, PMLR, 2019, pp. 2712–2721.
- [25] T. Chen, S. Liu, S. Chang, Y. Cheng, L. Amini, and Z. Wang, “Adversarial robustness: From self-supervised pre-training to fine-tuning,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 699–708.
- [26] T. H. Trinh, M.-T. Luong, and Q. V. Le, “Selfie: Self-supervised pretraining for image embedding,” *arXiv preprint arXiv:1906.02940*, 2019.
- [27] M. Noroozi and P. Favaro, “Unsupervised learning of visual representations by solving jigsaw puzzles,” in *European conference on computer vision*, Springer, 2016, pp. 69–84.
- [28] Y. Wang<sup>1</sup>, D. Zou<sup>2</sup>, J. Yi<sup>3</sup>, J. Bailey<sup>4</sup>, X. Ma<sup>4</sup>, and Q. Gu, “Improving adversarial robustness requires revisiting misclassified examples,” in *International Conference on Learning Representations* *International Conference on Learning Representations*, 2020.
- [29] A. Fawzi, H. Fawzi, and O. Fawzi, “Adversarial vulnerability for any classifier,” in *Conference on Neural Information Processing Systems*, Conference on Neural Information Processing Systems, 2018, 1186–1195.
- [30] A. Ilyas, S. Santurkar, D. Tsipras, L. Engstrom, B. Tran, and A. Madry, “Adversarial examples are not bugs, they are features,” *arXiv preprint arXiv:1905.02175*, 2019.
- [31] D. Hendrycks and T. Dietterich, “Benchmarking neural network robustness to common corruptions and perturbations,” 2019.