



# Robust Single-Step Adversarial Training with Regularizer

Lehui Xie<sup>1,2</sup>, Yaopeng Wang<sup>1,2</sup>, Jia-Li Yin<sup>1,2</sup>, and Ximeng Liu<sup>1,2</sup>(✉)

<sup>1</sup> College of Mathematics and Computer Science, Fuzhou University, Fuzhou 350108, China

<sup>2</sup> Fujian Provincial Key Laboratory of Information Security of Network Systems, Fuzhou University, Fuzhou 350108, China

**Abstract.** High cost of training time caused by multi-step adversarial example generation is a major challenge in adversarial training. Previous methods try to reduce the computational burden of adversarial training using single-step adversarial example generation schemes, which can effectively improve the efficiency but also introduce the problem of “catastrophic overfitting”, where the robust accuracy against Fast Gradient Sign Method (FGSM) can achieve nearby 100% whereas the robust accuracy against Projected Gradient Descent (PGD) suddenly drops to 0% over a single epoch. To address this issue, we focus on single-step adversarial training scheme in this paper and propose a novel Fast Gradient Sign Method with PGD Regularization (FGSM<sub>PR</sub>) to boost the efficiency of adversarial training without catastrophic overfitting. Our core observation is that single-step adversarial training can not simultaneously learn robust internal representations of FGSM and PGD adversarial examples. Therefore, we design a PGD regularization term to encourage similar embeddings of FGSM and PGD adversarial examples. The experiments demonstrate that our proposed method can train a robust deep network for  $L_\infty$ -perturbations with FGSM adversarial training and reduce the gap to multi-step adversarial training.

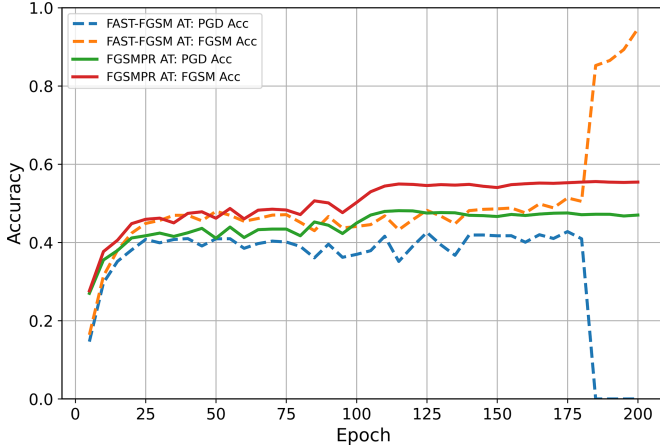
**Keywords:** Deep learning · Adversarial training · Adversarial defense

## 1 Introduction

Deep learning has shown outstanding success in near all machine learning fields. However, it has been proved that deep neural networks are vulnerable to adversarial examples, i.e., small disturbances to the input signal, which are usually invisible to the human eyes, are enough to induce large changes in model output [17]. This phenomenon has aroused people’s concerns about the safety of deep learning in the adversarial environment, where malicious attackers may significantly degrade the robustness of deep learning based applications. To mitigate the harm caused by the adversarial examples, numerous defensive methods

---

The first author is a student.



**Fig. 1.** (CIFAR-10) Visualization of the FGSM and PGD robustness of the model trained with FAST-FGSM AT (dashed), FGSMMPR AT (solid). All statistics are evaluated against FGSM attacks and 50 steps PGD attacks with 10 random restarts on the test dataset. FAST-FGSM AT occurs catastrophic overfitting at 180 epochs, characterized by a sudden drop of PGD robustness and a rapid increase of FGSM robustness. FGSMMPR AT (ours) does not suffer from catastrophic overfitting and maintains stable robustness during the whole training process.

including pre-processing based [6], modified networks based [3] and detection based [13] have been proposed. Among these methods, adversarial training [12] is one of the most powerful approaches for robust defense against adversarial attacks since [2] a set of purportedly robust defenses by the adaptive attack.

Adversarial Training (AT) aims to augment each small batch of training data with adversarial examples for learning a robust model. It is generally considered to be more expensive than traditional training because it is necessary to construct adversarial examples through first-order methods such as projection gradient descent (PGD). To combat the increased computational overhead of PGD AT, a recent line of work focused on improving the efficiency of AT. [20] proposed to perform multi-step PGD adversarial attacks by chopping off redundant computations during backpropagation when computing adversarial examples to obtain additional speedup. [15] proposed a variant of  $K$  steps PGD AT with a single-step Fast Gradient Sign Method (FGSM) AT overhead, called “FREE AT”, which can update model weights as well as input perturbations simultaneously by using a single backpropagation in a way that is less expensive than PGD AT overheads. Inspired by [15, 19] found that previously non-robustness FGSM AT, with a random initialization, could reach similar robustness to PGD AT, called “FAST-FGSM AT”. However, FGSM-based AT suffers from catastrophic overfitting where the robustness against PGD attacks increases in the early stage of training, but suddenly drops to 0 over a single epoch, as shown in Fig. 1. Several methods have been proposed to prevent the overfitting of AT

[1, 11, 18, 19]. However, these methods are either computationally inefficient or decrease the robustness accuracy.

In this paper, we first analyze the reason why FGSM AT suffers catastrophic overfitting in the training process. We observe that FGSM AT is prone to learn spurious functions that excessively fit the FGSM adversarial data distribution but have undefined behavior off the FGSM adversarial data manifold. Then we discuss the difference behind the logits output between the FGSM and PGD adversarial examples in the model trained with FGSM AT and PGD AT, where we show that the logits become significantly different when the FGSM AT trained model suffers from overfitting, while the robust model trained with PGD AT remains stable. We additionally provide for this case an experimental analysis that helps to explain why the FGSM AT trained model generates vastly different logit outputs for single-step and multi-step adversarial examples when catastrophic overfitting occurs. Finally, we propose a novel Fast Gradient Sign Method with PGD regularization (FGSM<sub>PR</sub>), in which a PGD regularization item is utilized to prompt the model to learn logits that are a function of the truly robust features in the image and ignore the spurious features, thus preventing catastrophic overfitting, as shown in Fig. 1.

The contribution of this paper is summarized as follows:

- We analyze the reason why FGSM AT suffers from catastrophic overfitting and demonstrate that the logit distribution of the FGSM AT trained model evaluated against FGSM and PGD adversarial examples have significant difference when suffering from catastrophic overfitting.
- We propose a Fast Gradient Sign Method with PGD regularization (FGSM<sub>PR</sub>), which can effectively prevent FGSM AT from catastrophic overfitting by explicitly minimizing the difference in the logit of the model against FGSM and PGD adversarial examples.
- The extensive experiments show that the FGSM<sub>PR</sub> can learn a robust model comparable to PGD AT with low computational overhead while in relief of catastrophic overfitting. Specially, the FGSM<sub>PR</sub> takes only 30 min to train a CIFAR-10 model with 46% robustness against 50 steps PGD attacks.


## 2 Related Work

### 2.1 Adversarial Training

Previous work [12] formalized the training of adversarial robust model into the following non-convex non-concave min-max robust optimization problem:

$$\min_{\theta} \mathbb{E}_{(x,y) \sim \mathcal{D}} [\max_{\delta \in \mathcal{S}} \mathcal{L}(\theta, x + \delta, y)]. \quad (1)$$

The parameter  $\theta$  of the network is learned by Eq. 1 on the example  $(x, y) \sim \mathcal{D}$ , where  $\mathcal{D}$  is the data generating distribution.  $\mathcal{S}$  denotes the region within the



**Fig. 2.** (CIFAR-10) Visualization of the FGSM and PGD accuracy/loss of the model trained with FGSM AT, FAST-FGSM AT, PGD-7 AT and tested against FGSM adversarial attacks and 50 steps PGD attack with 10 random restarts during the training process. All results are averaged over three independent runs. FGSM AT and FAST-FGSM AT occurs catastrophic overfitting around 30 and 180 epochs, respectively, characterized by a sudden drop in PGD accuracy and FGSM loss and a rapid increase in PGD loss and FGSM accuracy.

$\epsilon$  perturbation range under the  $\ell_\infty$  threat model for each example, i.e.,  $\mathcal{S} = \{\delta : \|\delta\|_\infty \leq \epsilon\}$ , which is usually chosen so that it contains only visually imperceptible perturbations. The procedure for AT is to use adversarial attacks to approximate the internal maximization over  $\mathcal{S}$ .

**FGSM AT.** One of the earliest versions of AT used the FGSM attack to find adversarial examples  $x'$  to approximate the internal maximization, formalized as follows [5]:

$$x' = x + \epsilon \cdot \text{sign}(\nabla_x \mathcal{L}(\theta, x, y)). \quad (2)$$

FGSM AT is cheap since it only relies on computing the gradient once. However, the FGSM AT is easily defeated by multi-step adversarial attacks.

**PGD AT.** PGD attacks [12] used multi-step gradient projection descent to approximate the inner maximization, which is more accurate than FGSM but computationally expensive, formalized as follows:

$$x^{t+1} = \Pi_{x+\mathcal{S}}(x^t + \alpha \text{sign}(\nabla_x \mathcal{L}(\theta, x, y))), \quad (3)$$

where  $x^0$  initialized as the clean input  $x$ ,  $\Pi$  refers to the projection operator, which ensures projecting the adversarial examples back to the ball within the radius  $\epsilon$  of the clean data point. The number of iterations  $K$  in the PGD attacks (PGD- $K$ ) determines the strength of the attack and the computational cost. Further,  $N$  random restarts are usually employed to verify robustness under strong attacks (PGD- $K$ - $N$ ).

## 2.2 Single-Step Adversarial Training

FREE AT [15], a single-step training method that generates adversarial examples while updating network weights, is quite similar to FGSM AT. By deeply analyzing the differences between FREE AT and FGSM AT, [19] found that an important property of FREE AT is that the perturbation of the previous sign of gradient is used as the initial perturbation of the next iteration. Based on this observation, [19] proposed a FAST-FGSM AT with almost the same robustness as the PGD AT, but the spent time close to the normal training by adding non-zero initialization perturbations to FGSM AT and further combining some standard techniques [14, 16] to accelerate the training. However, FAST-FGSM AT suffers from catastrophic overfitting, where the robustness for PGD adversarial examples suddenly drop to 0% over a single epoch.

A recent line of work focus on addressing the catastrophic overfitting problem in single-step AT. [19] used the early stopping method to stop training the model when the model robustness decreases beyond a threshold. [18] introduced dropout layers after each non-linear layer of the model and further decay its dropout probability as the training progresses. [11] monitored the FGSM AT process and performed PGD AT with a few batches to help the FGSM model recover its robustness when the robustness decreases beyond a threshold. [1] proposed the Gradient Alignment (GradAlign) regularization item that maximizes the gradient alignment based on the connection between FAST-FGSM AT overfitting and local linearization of the model as a way to prevent the occurrence of catastrophic overfitting. Although these methods provide a better understanding of catastrophic overfitting prevention, but still cannot essentially explain the problem of catastrophic overfitting. Moreover, these methods can improve the robustness of single-step AT models to some extent, but sacrifice a large amount of computational overhead and lose the efficient advantage of single-step AT, even up to the training time of multi-step AT.

## 3 Proposed Approach

### 3.1 Observation

To investigate catastrophic overfitting, we begin by recording the robust accuracy of FGSM AT on CIFAR-10 [9]. We evaluate the robust accuracy of the model against 50 steps PGD attacks with 10 random restarts (PGD-50-10) for step size  $\alpha = 2/255$  and maximum perturbation  $\epsilon = 8/255$ . Figure 2 visualizes the accuracy and loss of the FGSM AT trained, FAST-FGSM AT trained, and PGD-7 AT trained model and evaluated against FGSM and PGD-50-10 attack during the training phase. As we can see, when FGSM AT and FAST-FGSM AT occur catastrophic overfitting around 30 and 180 epochs respectively, the robustness against PGD-50-10 attack of the model trained with FGSM AT and FAST-FGSM AT begin to drop suddenly, whereas the accuracy against FGSM increases rapidly. However, for the robust PGD-7 AT, the accuracy and loss of the model tend to stabilize after a certain number of epochs.

We maintain that the reason the models trained using FGSM AT suffer from catastrophic overfitting is that it is prone to learn spurious functions that fit the FGSM data distribution but have undefined behavior off the FGSM data manifold. Therefore, the FGSM AT is highly susceptible to overfitting due to a single-step adversarial perturbation, resulting in a sudden drop in the PGD robustness of the model, while the FGSM accuracy increases instantaneously. To study the differences in the performance of the models trained with FGSM AT and PGD-7 AT for evaluating at the FGSM and PGD adversarial examples, we utilize a distance function  $\mathcal{L}$  to measure the difference between the output of the model evaluated at single-step and multi-step adversarial attacks. For a model that take inputs  $x$  and output logits  $f(x)$ , we have:

$$\mathcal{L}(f(x^{fgsm}), f(x^{pgd})), \quad (4)$$


where  $x^{fgsm}$  and  $x^{pgd}$  are adversarial examples crafted by FGSM and PGD-7, respectively. Here, we choose  $L_2$  for  $\mathcal{L}$ . For a well-generalized and robustness model, we assume that the logit  $f(x^{fgsm})$  and  $f(x^{pgd})$  of the model evaluated at FGSM and PGD adversarial examples should be as similar as possible, i.e.,  $\|f(x^{fgsm}) - f(x^{pgd})\|_2$  should be very small.

To demonstrate our intuition, we firstly train several CIFAR-10 models using FGSM AT and PGD-7 for 200 epochs. For each model, we compute the difference between the output of the model evaluated at FGSM and PGD-7 adversarial examples by using Eq. 4, and performed data processing using a logarithmic function to visualize the differences more clearly, as shown in Fig. 3. In Fig. 3 (b), it can be observed that there is no significant difference in the logits from FGSM and PGD adversarial examples during the early phase of training, which matches our intuition. Once catastrophic overfitting occurs, the gap between the logit of the model evaluated at single-step and multi-step adversarial attacks are increasing rapidly around 30 and 180 epochs respectively, which is consistent with PGD loss. In contrast, the PGD-7 AT does not suffer catastrophic overfitting and the difference of the logit of the model is keeping stable. This phenomenon will also appear on the simple MNIST dataset [10], but it is not as clear as CIFAR-10, as shown in Fig. 3(a).

### 3.2 PGD Regularization

Based on the analysis in Sect. 3.1, the only FGSM adversarial loss is not enough for the model to learn the robust features of both single-step and multi-step adversarial examples. To solve this problem, inspired by [8], we use the logit pairing to encourage the model to learn robust internal representation of FGSM and PGD adversarial examples so that the logit outputs  $f(x^{fgsm})$  and  $f(x^{pgd})$  of the model for FGSM and PGD adversarial examples to be as similar as possible:

$$\lambda \frac{1}{m} \sum_{i=1}^m \mathcal{L}(f(x_i^{fgsm}; \theta), f(x_i^{pgd}; \theta)), \quad (5)$$



**Fig. 3.** Visualization of the  $\mathcal{L}_2$  distance of logit of the FGSM AT trained, FAST-FGSM AT trained, PGD-7 AT trained model and evaluated against FGSM and PGD adversarial attack. (a) In MNIST, when the model is not robust, the difference in  $\mathcal{L}_2$  distance starts to fluctuate, while PGD AT is relatively smooth. (b) In CIFAR10, FGSM AT and FAST-FGSM AT occurs catastrophic overfitting around 30 and 180 epochs, respectively, and is characterized by a rapid increase of  $\mathcal{L}_2$  distance.

where  $\mathcal{L}$  is  $L_2$  norm;  $x_i^{fgsm}$  and  $x_i^{pgd}$  are adversarial examples crafted by FGSM and PGD attacks, respectively;  $\lambda$  is a hyperparameter to balance FGSM loss and PGD regularization item. Combining with the proposed regularization, the FGSM AT can learn a robustness model comparable with PGD-7 AT, as validated in Sect. 4.

We hold that PGD regularization works well because it provides an additional prior that regularizes the model toward a more accurate understanding of adversarial examples. If we train the model with only the single-step FGSM adversarial loss, it is prone to learn spurious functions that excessively fit the FGSM adversarial data distribution but have undefined behavior off the FGSM data manifold (e.g., multi-step adversarial examples). PGD regularization forces the explanations of the FGSM adversarial example and multi-step adversarial example to be similar. This is essentially a prior encouraging the model to learn logits that are a function of the truly significant features in the image and ignore the spurious features.

### 3.3 Training Route

The overall training procedure of the FGSMPR AT is summarized in Algorithm 1. We first perform FGSM adversarial attack to generate FGSM adversarial examples  $x_i^{fgsm}$  and compute FGSM AT loss  $fgsm\_loss$  using cross-entropy. Then, we perform PGD adversarial attack for  $m$  examples, from a batch of natural examples, to generate  $m$  PGD adversarial examples. After generating FGSM and PGD adversarial examples, the regularization loss  $reg\_loss$  of  $m$  FGSM and PGD adversarial examples are calculated using Eq. 5 and used as part of the total loss  $total\_loss$ . Finally, the parameter  $\theta$  of the model is updated using a proper optimizer (e.g., stochastic gradient descent). The hyperparameter

---

**Algorithm 1: FGSM-PR AT**

---

**Input** : Training data  $(X, Y)$ , perturbation bound  $\epsilon$ , learning rate  $\gamma$ ,  
 hyperparameter  $\alpha, \lambda$ .

**Output**: Trained model  $f(\cdot)$  with parameter  $\theta$

```

1 for  $epoch = 1 \dots N_{epoch}$  do
2   for  $i = 1 \dots B$  do
3     // Perform FGSM adversarial attack
4      $x_i^{fgsm} = x_i + \alpha \cdot \text{sign}(\nabla_{\delta} \ell(f_{\theta}(x_i), y_i))$ 
5      $fgsm\_loss = J(f_{\theta}, x_i^{fgsm}, y_i)$ 
6     // Perform PGD adversarial attack
7     for  $k = 1 \dots K$  do
8        $\delta = \delta + \alpha \cdot \text{sign}(\nabla_{\delta} \ell(f_{\theta}(x_i + \delta), y_i))$ 
9        $\delta = \max(\min(\delta, \epsilon), -\epsilon)$ 
10    end for
11     $x_i^{pgd} = x_i + \delta$ 
12     $reg\_loss = \lambda \frac{1}{m} \sum_{j=1}^m \mathcal{L}(f(x_{i,j}^{fgsm}; \theta), f(x_{i,j}^{pgd}; \theta))$ 
13     $total\_loss = fgsm\_loss + reg\_loss$ 
14    Update model parameter  $\theta$  based on  $total\_loss$ 
15  end for
16 end for
17 return  $f(\cdot)$ .
```

---

$\lambda$  shall be properly chosen to balance FGSM loss  $fgsm\_loss$  and PGD regularization  $reg\_loss$  item. In practice, we take  $\alpha = \epsilon/K$ ,  $K = 3$  and  $m = 1$ . In other words, we only pick a single example from a batch for generating a PGD-3 adversarial example, which is then used for regularization to encourage the model to learn similar logit output. The experiments show that a single PGD adversarial example for regularization is enough to learn a robustness model.

## 4 Experiments

In this section, we demonstrate that the proposed FGSM-PR is robust against strong PGD attacks. All experiments are run on a single RTX 2070, in which we use half-precision computation recommended in [19] to speed up the training of CIFAR-10 model, which was incorporated with the Apes amp package at the O1 optimization level for all CIFAR-10 experiments.

**Attacks:** We attack all models using PGD attacks with  $K$  iterations and 10 random restarts on both cross-entropy loss (PGD-K-10) and the Carlini-Wagner loss (CW-K-10) [4]. All PGD attacks used at evaluation for MNIST [10] are run with 10 random restarts for 20/40 iterations. All PGD attacks used at evaluation for CIFAR-10 [9] are run with 10 random restarts for 20/50 steps.

**Perturbation:** For MNIST, we set the maximum perturbation  $\epsilon$  to 0.3 and the PGD step size  $\alpha$  to 0.1. For CIFAR-10, we set the maximum perturbation  $\epsilon$  to  $8/255$  and the PGD step size  $\alpha$  to  $2/255$ .



**Table 1.** Validation accuracy (%) and robustness of MNIST models trained with FGSM AT, FAST-FGSM AT, GradAlign AT, FREE AT, PGD-40 AT, FGSMPR AT without early stopping and the corresponding training time. All statistics are evaluated against PGD/CW attacks with 20/40 iterations and 10 random restarts for  $\alpha = 0.1$ ,  $\epsilon = 0.3$  over three independent runs. The bold indicates the best performance except for PGD-40 AT.

Method	Standard accuracy	PGD-20-10	PGD-40-10	CW-20-10	CW-40-10	Training time (s)
FGSM AT	97.53 $\pm$ 0.39	39.31 $\pm$ 20.68	12.45 $\pm$ 12.15	40.14 $\pm$ 20.64	13.46 $\pm$ 12.89	481.03 $\pm$ 0.81
FAST-FGSM AT	98.52 $\pm$ 0.34	42.60 $\pm$ 12.47	11.19 $\pm$ 6.52	43.41 $\pm$ 12.55	11.85 $\pm$ 7.26	491.24 $\pm$ 1.71
GradAlign AT	<b>99.05 <math>\pm</math> 0.03</b>	91.42 $\pm$ 0.57	75.94 $\pm$ 3.23	91.23 $\pm$ 0.51	75.86 $\pm$ 3.16	633.96 $\pm$ 3.55
FREE AT	98.49 $\pm$ 0.05	92.90 $\pm$ 0.20	90.06 $\pm$ 0.36	92.70 $\pm$ 0.15	89.85 $\pm$ 0.32	<b>175.45 <math>\pm</math> 1.99</b>
PGD-40 AT	99.16 $\pm$ 0.03	94.72 $\pm$ 0.08	92.52 $\pm$ 0.14	94.75 $\pm$ 0.03	92.65 $\pm$ 0.10	3652.39 $\pm$ 1.00
FGSMPR AT (ours)	98.35 $\pm$ 0.09	<b>93.77 <math>\pm</math> 0.32</b>	<b>90.83 <math>\pm</math> 0.49</b>	<b>93.65 <math>\pm</math> 0.26</b>	<b>90.56 <math>\pm</math> 0.55</b>	626.57 $\pm$ 1.68

**Table 2.** Validation accuracy (%) and robustness of CIFAR-10 models trained with FGSM AT, FAST-FGSM AT, GradAlign AT, FREE AT, PGD-7 AT, FGSMPR AT without early stopping and the corresponding training time. All statistics are evaluated against PGD/CW attacks with 20/50 iterations and 10 random restarts for  $\alpha = 2/255$ ,  $\epsilon = 8/255$  over three independent runs. The bold indicates the best performance except for PGD-7 AT.

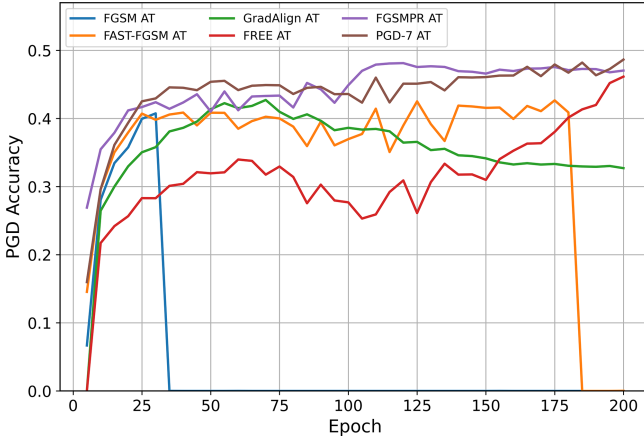
Method	Standard accuracy	PGD-20-10	PGD-50-10	CW-20-10	CW-50-10	Training time (m)
FGSM AT	88.51 $\pm$ 1.27	0.01 $\pm$ 0.17	0.00 $\pm$ 0.00	0.01 $\pm$ 0.11	0.00 $\pm$ 0.00	119.04 $\pm$ 0.41
FAST-FGSM AT	<b>90.33 <math>\pm</math> 0.42</b>	0.92 $\pm$ 0.49	0.32 $\pm$ 0.25	0.52 $\pm$ 0.30	00.17 $\pm$ 0.08	123.81 $\pm$ 0.18
GradAlign AT	82.82 $\pm$ 0.13	32.94 $\pm$ 0.83	32.50 $\pm$ 0.80	32.94 $\pm$ 0.83	32.52 $\pm$ 0.81	486.20 $\pm$ 0.67
FREE AT	82.32 $\pm$ 0.12	46.97 $\pm$ 0.05	46.07 $\pm$ 0.82	45.77 $\pm$ 0.23	45.64 $\pm$ 0.24	<b>61.94 <math>\pm</math> 0.15</b>
PGD-7 AT	84.75 $\pm$ 0.87	48.33 $\pm$ 0.62	47.99 $\pm$ 0.66	47.80 $\pm$ 0.31	47.59 $\pm$ 0.38	493.41 $\pm$ 0.04
FGSMPR AT (ours)	83.31 $\pm$ 0.40	<b>47.59 <math>\pm</math> 0.51</b>	<b>47.19 <math>\pm</math> 0.42</b>	<b>46.98 <math>\pm</math> 0.18</b>	<b>46.79 <math>\pm</math> 0.20</b>	211.65 $\pm$ 0.61

**Comparisons:** We compare the performance of our proposed method (FGSMPR) with FGSM: standard FGSM AT [5]; FAST-FGSM AT: FGSM AT with a random initialization [19]; FREE AT: recently proposed single-step AT method [15]; GradAlign AT: recently proposed method solving catastrophic overfitting [1]; PGD- $K$  AT: AT with a  $K$  iterations PGD attack [12].

**Evaluation:** We demonstrate that the performance of models against PGD- $K$ -10/CW- $K$ -10 adversarial attacks under white-box settings. For all experiments, the mean and standard deviation over three independent runs are reported.

#### 4.1 Results on MNIST

First, we conduct a study to demonstrate that our proposed approach is highly working in MNIST benchmark dataset [10]. We train models for MNIST dataset with the same architecture used by [19], using FGSM AT, FAST-FGSM AT,




**Fig. 4.** Visualization of the accuracy of the CIFAR-10 model trained for FGSM AT, FAST-FGSM AT, GradAlign AT, FREE AT, PGD-7 AT, and FGSMRP AT. All the statistics are tested against 50 steps PGD attacks with 10 random restarts for  $\alpha = 2/255$ ,  $\epsilon = 8/255$ . Catastrophic overfitting for the FGSM and FAST-FGSM AT occur around 30 and 180 epochs, respectively, and is characterized by a sudden drop in the PGD accuracy.

FREE AT, PGD-40 AT, FGSMRP AT. Except that the AT free replays each batch of  $m = 8$  for a total of 7 epochs, all other models are trained for 50 epochs. For the proposed method, we set the hyperparameter  $\lambda$ ,  $K$  and  $m$  to  $(0.1, 3, 1)$ . The experimental results are provided in Table 1. It can be observed that our proposed FGSMRP AT is more robust against both PGD and CW attacks on the MNIST dataset than the GradAlign AT and FREE AT, and is second only to the PGD AT model with a small difference. In the course of testing the robustness of FAST-FGSM AT on the MNIST dataset, we found an interesting problem where increasing the number of MNIST training epochs to 50 also resulted in catastrophic overfitting, although this phenomenon was previously found only in CIFAR-10. Besides, the GradAlign AT [1] can keep the model from suffering catastrophic overfitting to some extent, but it is far inferior to other comparison methods in defending against the higher iteration adversarial attacks.

## 4.2 Results on CIFAR-10

To verify whether AT scheme suffers from catastrophic overfitting, we train 200 epoch for all CIFAR-10 models using the Preact ResNet-18 [7] architecture without early stopping, especially the FREE AT replays each batch  $m = 8$  times for a total of 25 epochs as recommend in [15]. For the FGSMRP, we set the hyperparameter  $\lambda$ ,  $K$  and  $m$  to  $(0.5, 3, 1)$ . The experimental results are provided in Table 2. It can be observed that FGSMRP AT is quite similar to PGD-7 AT while our training time is half of PGD-7 AT. To demonstrate that the proposed FGSMRP does not suffer from catastrophic overfitting, we takes 211 min to



**Fig. 5.** Accuracy of the model trained for FGSM AT, FAST-FGSM AT, GradAlign AT, FREE AT, PGD-7 AT and FGSMR AT with early stopping. All the statistics are evaluated against 50 steps PGD attacks with 10 random restarts for different  $l_\infty$ -perturbation  $\epsilon$ .

train a CIFAR-10 model for 200 epochs, which is longer than time for FREE AT. However, our method was able to achieve 46% robustness by training 30 epochs in only 30 min, which is half less than FREE AT. Further, we visualize the robustness of the training process of different AT method and tested against a 10 random restart PGD-50 attack, as shown in Fig. 4. It can be observed that the robustness of the FGSMR against PGD has steadily increased, which is only 0.8% behind PGD-7 AT and does not suffer from catastrophic overfitting even when trained to 200 epochs. Instead, FAST-FGSM AT started to have a trend similar to PGD AT, but there is a sharp drop in robustness around 180 epochs when occurring catastrophic overfitting. The GradAlign AT was proposed to prevent the FGSM AT from catastrophic overfitting, but the accuracy still dropped by more than 10% and took more than two times longer compared to our FGSMR AT. Besides, we also test the model’s robustness under different  $l_\infty$  perturbation where all models are trained with early stopping. In the case of larger  $l_\infty$  perturbations, FGSMR AT is essentially indistinguishable from PGD-7 AT, and even slightly better than PGD-7 AT, as shown in Fig. 5.

## 5 Conclusion

In this paper, we analyze and address the catastrophic overfitting in FGSM AT. We empirically show that FGSM AT is prone to learn spurious functions that fit the FGSM adversarial data distribution but have undefined behavior off the FGSM data manifold. We therefore exploit the difference behind the logits between the FGSM and PGD adversarial examples in the model trained with FGSM AT and PGD AT, where the logit becomes significantly different

when FGSM AT suffers from overfitting, while PGD AT remains stable. Based on these observations, we propose a novel FGSM<sub>PR</sub> AT, where a PGD regularization term is used to encourage the model to learn similar embeddings of FGSM and PGD adversarial examples. The extensive experiments show that the FGSM<sub>PR</sub> can effectively keep FGSM AT from catastrophic overfitting with a low computational cost.

## References

1. Andriushchenko, M., Flammarion, N.: Understanding and improving fast adversarial training. *Virtual*, Online (2020)
2. Athalye, A., Carlini, N., Wagner, D.: Obfuscated gradients give a false sense of security: circumventing defenses to adversarial examples. In: *International conference on machine learning*, pp. 274–283. Stockholm, Sweden (2018)
3. Buckman, J., Roy, A., Raffel, C., Goodfellow, I.: Thermometer encoding: one hot way to resist adversarial examples. In: *International Conference on Learning Representations*. Vancouver, BC, Canada (2018)
4. Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. In: *2017 IEEE Symposium on Security and Privacy (sp)*, pp. 39–57. IEEE (2017)
5. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. San Diego, CA, United states (2015)
6. Guo, C., Rana, M., Cisse, M., Van Der Maaten, L.: Countering adversarial images using input transformations. Vancouver, BC, Canada (2018)
7. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016*. LNCS, vol. 9908, pp. 630–645. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46493-0\\_38](https://doi.org/10.1007/978-3-319-46493-0_38)
8. Kannan, H., Kurakin, A., Goodfellow, I.: Adversarial logit pairing (2018)
9. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf> (2009)
10. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proc. IEEE* **86**(11), 2278–2324 (1998)
11. Li, B., Wang, S., Jana, S., Carin, L.: Towards understanding fast adversarial training. *arXiv preprint arXiv:2006.03089* (2020)
12. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks (2018)
13. Metzen, J.H., Genewein, T., Fischer, V., Bischoff, B.: On detecting adversarial perturbations. Toulon, France (2017)
14. Narang, S., et al.: Mixed precision training. Vancouver, BC, Canada (2018)
15. Shafahi, A., et al.: Adversarial training for free! In: *Advances in Neural Information Processing Systems*, pp. 3358–3369 (2019)
16. Smith, L.N.: Cyclical learning rates for training neural networks. In: *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 464–472. IEEE (2017)
17. Szegedy, C., et al.: Intriguing properties of neural networks. Banff, AB, Canada (2014)
18. Vivek, B., Babu, R.V.: Single-step adversarial training with dropout scheduling. In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 947–956. IEEE (2020)

19. Wong, E., Rice, L., Kolter, J.Z.: Fast is better than free: Revisiting adversarial training. Addis Ababa, Ethiopia (2020)
20. Zhang, D., Zhang, T., Lu, Y., Zhu, Z., Dong, B.: You only propagate once: accelerating adversarial training via maximal principle. In: Advances in Neural Information Processing Systems, pp. 227–238 (2019)