

# Preventing Catastrophic Overfitting in Fast Adversarial Training: A Bi-level Optimization Perspective

Zhaoxin Wang<sup>1</sup>, Handing Wang<sup>1</sup><sup>✉</sup>, Cong Tian<sup>2</sup>, and Yaochu Jin<sup>3</sup>

<sup>1</sup> School of Artificial Intelligence, Xidian University, Xi'an, China  
zxwang74@163.com

<sup>2</sup> School of Computer Science and Technology, Xidian University, Xi'an, China  
ctian@mail.xidian.edu.cn

<sup>3</sup> School of Engineering, Westlake University, Zhejiang Hangzhou, China  
jinyaochu@westlake.edu.cn

**Abstract.** Adversarial training (AT) has become an effective defense method against adversarial examples (AEs) and it is typically framed as a bi-level optimization problem. Among various AT methods, fast AT (FAT), which employs a single-step attack strategy to guide the training process, can achieve good robustness against adversarial attacks at a low cost. However, FAT methods suffer from the catastrophic overfitting problem, especially on complex tasks or with large-parameter models. In this work, we propose a FAT method termed FGSM-PCO, which mitigates catastrophic overfitting by averting the collapse of the inner optimization problem in the bi-level optimization process. FGSM-PCO generates current-stage AEs from the historical AEs and incorporates them into the training process using an adaptive mechanism. This mechanism determines an appropriate fusion ratio according to the performance of the AEs on the training model. Coupled with a loss function tailored to the training framework, FGSM-PCO can alleviate catastrophic overfitting and help the recovery of an overfitted model to effective training. We evaluate our algorithm across three models and three datasets to validate its effectiveness. Comparative empirical studies against other FAT algorithms demonstrate that our proposed method effectively addresses unresolved overfitting issues in existing algorithms.

**Keywords:** Fast adversarial training · catastrophic overfitting · fusion adversarial examples

## 1 Introduction

The security of deep neural networks [5, 21, 22] has raised increasing concerns as they always face malicious threats [13, 14, 32, 34, 44]. Among these threats, adversarial examples (AEs) pose a significant risk by causing neural networks

---

<sup>✉</sup> Corresponding Author: hdwang@xidian.edu.cn

to make an incorrect prediction or classification without participating in the training phase [3, 6, 12, 16, 25]. Adversarial training (AT) [23, 31, 35–37, 40] is one of the most effective ways to resist this adversarial attack, which can significantly enhance the models’ classification ability for AEs rather than merely identifying them. AT can be framed as a bi-level optimization problem [23] as follows:

$$\min_{\theta} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[ \max_{\|\delta\|_p \leq \epsilon} \mathcal{L}(\mathbf{f}_{\theta}(\mathbf{x} + \delta), \mathbf{y}) \right], \quad (1)$$


where  $\mathbf{f}_{\theta}$  represents a deep model with parameters  $\theta$ ,  $\mathcal{D}$  represents the data distribution of clean examples  $\mathbf{x}$  and their labels  $\mathbf{y}$ ,  $\delta$  denotes the perturbation,  $\epsilon$  denotes the perturbation threshold, and  $\|\cdot\|_p$  is defined as  $l_p$  norm.

Typically, the univariate search technique is commonly employed for this bi-level optimization problem [7], where attack strategies generate AEs with fixed model parameters in the inner maximum process, and these AEs are used to update the model parameters in the outer minimum optimization process. However, this maximum process requires multiple iterative processes to find an effective  $\delta$ , leading to significant computation demands.

To accelerate the training process of AT, fast adversarial training (FAT) methods have been developed [2, 13, 17, 19, 38, 42], which utilize the fast gradient sign attack method (FGSM) [13], a single-step attack method, as the inner attack strategy. However, coupling the univariate search technique with FGSM is highly susceptible to the collapse of FAT, leading to the catastrophic overfitting problem where the classification accuracy of AEs suddenly drops to 0% under multi-step attacks like Projected Gradient Descent (PGD) attack. Fig. 1a shows an example of catastrophic overfitting observed in FGSM-AT [13].

To address the catastrophic overfitting problem, FGSM-RS [38] initializes the perturbation with a random uniform distribution  $U[-\epsilon, \epsilon]$  and adopts a large attack step size in FGSM strategy. Zhang *et al.* [2, 42] propose integrating regularization into the loss function to mitigate the issue. FGSM-MEP [17] proposes that the previous perturbation can guide the perturbation initialization in FAT, leveraging the accumulation of the perturbation momentum to overcome the catastrophic overfitting problem. These methods significantly improve the performance of FAT, allowing FAT methods to achieve comparable robustness to the multi-step attack-guided AT at a low cost.

Nonetheless, these methods are not immune to catastrophic overfitting on more complex tasks [11] or with large-parameter models. Fig. 1b shows catastrophic overfitting phenomenon under different FAT methods. Compared with FGSM-AT, these algorithms primarily delay the onset of overfitting but cannot completely prevent it. The reason for this catastrophic overfitting phenomenon is the collapse of the inner optimization problem, and the collapse derives from coupling FGSM this single, large-step attack strategy with the alternate optimization method between the inner and outer problem. As shown in Fig. 1a, the classification accuracy under FGSM attack steep rise. Once this inner optimization fails, the entire optimization process becomes ineffective. Besides, existing FAT methods lack mechanisms to rectify this flaw.



**Fig. 1:** Catastrophic overfitting phenomenon in FAT. (a) is on the CIFAR10 dataset with a multi-step learning rate. (b) shows the overfitting on the Tiny-ImageNet dataset with a cyclic learning rate. Most FAT algorithms cannot prevent catastrophic overfitting.

Therefore, we propose a FAT algorithm termed FGSM-PCO, which can effectively address the catastrophic overfitting problem in FAT. Specifically, we fuse the historical and current perturbations for training according to an adaptive mechanism. This adaptive mechanism can adjust the ratio of the two types of AEs according to the performance on the model. It ensures that FGSM-PCO can correct the training course, avoiding the collapse of the inner optimization problem. In addition, we propose a regularization to assist this methodology, which advocates a consistent prediction for training samples before and after fusion. In summary, our main contributions are:

- We propose a FAT framework that incorporates both the historical and current perturbation into training through an adaptive mechanism. This framework significantly mitigates the issue of catastrophic overfitting and helps the recovery of an overfitted model to effective training.
- A tailored regularization strategy is proposed to prevent the collapse of the inner optimization problem within the FAT framework, which advocates consistent predictions for samples both before and after the fusion.
- Experimental results show that our algorithm improves the classification accuracy on both clean examples and AEs.

## 2 Related work

### 2.1 Fast Adversarial Training

Most AT methods [24, 26, 28, 37] can be formulated as a bi-level optimization problem as shown in Eq.(1). In the AT framework, a multi-step attack method PGD is frequently employed to generate a high-quality  $\delta$  in the inner optimization process. PGD can be described as:

$$\mathbf{x}_t^* = \mathbf{x} + \Pi_{[-\epsilon, \epsilon]} \left[ \delta_{t-1} + \alpha \text{sign} \left( \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{f}_{\theta}(\mathbf{x}_{t-1}^*), \mathbf{y}) \right) \right], \quad (2)$$

where  $\Pi_{[-\epsilon, \epsilon]}$  represents the projection that limits the input to the range  $[-\epsilon, \epsilon]$ ,  $\delta_0$  is the randomly initialized perturbation,  $\mathbf{x}$  are the clean examples,  $\alpha$  is the attack step size.


This multi-step attack strategy makes AT requires high computational cost, which is impractical in some tasks. Therefore, researchers propose FAT, which employs FGSM, a single-step attack method instead of the PGD method, to generate the training examples. FGSM can be described as:

$$\mathbf{x}^* = \mathbf{x} + \epsilon \text{sign}(\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{f}_{\theta}(\mathbf{x} + \delta), \mathbf{y})), \quad (3)$$

where  $\delta$  is the perturbation initialization. Compared with the PGD attack, FGSM applies a single large step, using the perturbation threshold  $\epsilon$  directly as the step size, which significantly accelerates the training process. Despite its efficiency, FGSM increases the risk of FAT encountering the collapse of the inner optimization problem, leading to catastrophic overfitting.

## 2.2 Dilemma in Bi-Level Optimization Problems

Bi-level optimization involves two interconnected optimization problems, where the solution to the outer-level problem depends on the outcome of the inner-level problem [10] and vice versa. This nested structure can lead to conflicts and potential collapse, particularly when the inner and outer levels are optimized alternately. In the FAT framework, the bi-level optimization process can be described as Fig. 2.



**Fig. 2:** Optimizing the inner and outer problems alternately is easy to cause the collapse of the bi-level optimization.

In the bi-level optimization framework, the objective function of the outer layer  $Robust(\theta)$ , solely depends on the outer variable  $\theta$ , while the inner-level objective function,  $Loss(\theta, \mathbf{x}^*)$ , is dependent on both the outer variable  $\theta$  and the inner variable  $\mathbf{x}^*$ . Given an outer variable  $\theta^{(A)}$ , we obtain the corresponding loss surface on the right side in Fig. 2. Once the inner optimization collapse

occurs under  $\theta^{(B)}$ , the nested structure of the bi-level optimization exacerbates the collapse phenomenon in optimization. In the FAT, alternating optimization of inner and outer levels, coupled with FGSM’s large step size and constrained perturbation search space, increases the risk of collapse in bi-level optimization.

### 2.3 Prior-Guided Fast Adversarial Training

FGSM-MEP [17] proposes that prior knowledge of adversarial perturbations can effectively guide subsequent AT. This method employs the accumulation of the gradient momentum over all previous training epochs to initialize the FGSM perturbation  $\boldsymbol{\eta}$ . The process is detailed as follows:

$$\begin{aligned} \mathbf{g}_c &= \text{sign} \left( \nabla_{\mathbf{x}} \mathcal{L} \left( \mathbf{f}_{\theta} \left( \mathbf{x} + \boldsymbol{\eta}_{E_t} \right), \mathbf{y} \right) \right), \\ \mathbf{g}_{E_{t+1}} &= \mu \cdot \mathbf{g}_{E_t} + \mathbf{g}_c, \\ \boldsymbol{\delta}_{E_{t+1}} &= \Pi_{[-\epsilon, \epsilon]} \left[ \boldsymbol{\eta}_{E_t} + \alpha \cdot \mathbf{g}_c \right], \\ \boldsymbol{\eta}_{E_{t+1}} &= \Pi_{[-\epsilon, \epsilon]} \left[ \boldsymbol{\eta}_{E_t} + \alpha \cdot \text{sign} \left( \mathbf{g}_{E_{t+1}} \right) \right], \end{aligned} \quad (4)$$

where  $\mathbf{g}_{E_{t+1}}$  represents the gradient momentum in the  $(t+1)$ -th epoch,  $\boldsymbol{\delta}$  is the perturbation and  $\mu$  is the decay factor.

While this approach significantly enhances the performance of FAT, it still has some issues. Although the accumulation of momentum provides prior knowledge for AT, as training progresses, continuous accumulation of gradient momentum can inadvertently increase the divergence between adversarial and clean examples, potentially degrading the model’s performance on clean data, even though the momentum is reset periodically during training. Furthermore, while leveraging momentum to initialize subsequent perturbations decreases the risk of ineffective optimization and mitigates catastrophic overfitting, it fails to rectify the algorithm once the inner-level optimization collapses. In addition, although the accumulation of momentum does not increase the algorithm’s runtime, it significantly raises memory requirements.

## 3 Proposed Method

To prevent the collapse of the inner optimization problem in FAT, we propose the FAT method FGSM-PCO. It generates AEs based on those from the previous training epoch and then fuses these two types of AEs based on their performance on the model. The fusion AEs are then utilized for training within a framework guided by a novel loss function, ensuring FGSM-PCO can rectify the trend toward inner-layer optimization collapse. Specifically, when AEs from the current stage prove ineffective for FAT, the adaptive fusion mechanism prioritizes the prior-stage AEs, thus maintaining the training efficacy.

### 3.1 Procedure of Proposed Method

As illustrated in Fig. 1a, catastrophic overfitting is marked by a sudden spike in the model’s accuracy on FGSM-generated AEs, signaling the collapse of the

inner optimization within FAT. This collapse makes the AEs generated by FGSM ineffective for current training. To prevent this and correct the trend when it emerges, the fusion process is skewed towards retaining a higher proportion of previous-stage AEs, thereby averting the entrapment of bi-level optimization in local minima and the subsequent collapse of inner optimization. The training process can be specifically described by Eq.(5).

$$\begin{aligned}
\mathbf{x}_{\text{train}} &= \lambda_t \mathbf{x}_{t-1}^* + (1 - \lambda_t) \mathbf{x}_{\text{am}}^*, \\
\mathbf{x}_{\text{am}}^* &= \mathbf{x}_{t-1}^* + \gamma \boldsymbol{\delta}_t, \\
\boldsymbol{\delta}_t &= \epsilon \mathbf{g}_t, \\
\mathbf{g}_t &= \text{sign} \left( \nabla_{\mathbf{x}} \mathcal{L} \left( \mathbf{f}_{\theta} \left( \mathbf{x}_{t-1}^* \right), \mathbf{y} \right) \right), \\
\mathbf{x}_0^* &= \mathbf{x} + U[-\epsilon, \epsilon],
\end{aligned} \tag{5}$$

where  $\mathbf{x}$  and  $\mathbf{y}$  represent the clean examples and their corresponding labels,  $\mathbf{x}_{t-1}^*$  are the AEs in the  $(t-1)$ -th epoch,  $\mathbf{x}_{\text{am}}^*$  are the current stage AEs with a amplified delta  $\gamma \boldsymbol{\delta}_t$ ,  $\gamma$  is the amplification factor,  $\lambda_t$  represents the fusion factor,  $U[-\epsilon, \epsilon]$  represents a uniform distribution. The detailed description is presented in Algorithm 1.

### 3.2 Adaptive Fusion Ratio

Different from other FAT methods, which utilize FGSM to generate perturbations and attach them to the clean examples, we scale the perturbations with an amplification factor  $\gamma$  as mentioned in Eq.(5). This amplification factor ensures the effectiveness of perturbations after fusion, *i.e.*, it balances the perturbation attenuation brought by the fusion factor  $\lambda_t$ .

To determine a good fusion factor  $\lambda_t$ , which can make the algorithm escape the collapse of the inner optimization problem. The fusion factor value is based on the performance of the AEs on the model. Specifically, through the model's classification confidence for the ground-truth category, which directly influences  $\lambda_t$ . The formula for  $\lambda_t$  is:

$$\lambda_t = \mathbf{f}_{\theta}^i(\mathbf{x}_t^*) \tag{6}$$

where  $\mathbf{f}_{\theta}^i(\mathbf{x}_t^*)$  is the classification confidence,  $i$  is the index of the ground-truth category. As shown in Eq.(5), when the catastrophic overfitting occurs, the training examples  $\mathbf{x}_{\text{train}}$  skewed towards retaining a higher proportion of previous-stage AEs  $\mathbf{x}_{t-1}^*$ . This adaptive fusion mechanism ensures that FGSM-PCO can correct the catastrophic overfitting trend before the collapse of the inner optimization problem.

### 3.3 Regularization Loss

To better guide our proposed optimization process, we introduce a new loss function. For fusion AEs generated by Eq.(6), we hope to achieve prediction

results consistent with both the previous and the current AEs. Thus, apart from utilizing cross-entropy loss to direct AT, we propose a regularization strategy as defined in Eq.(7).

$$\begin{aligned} \mathcal{L}_{PCO}(\mathbf{x}_{train}, \mathbf{y}) &= \mathcal{L}_{CE}(\mathbf{f}_{\theta}(\mathbf{x}_{train}), \mathbf{y}) \\ &+ \beta [\mathcal{L}_1(\mathbf{f}_{\theta}(\mathbf{x}_t^*), \mathbf{f}_{\theta}(\mathbf{x}_{t-1}^*)) - \mathcal{L}_1(\mathbf{f}_{\theta}(\mathbf{x}_{train}), \mathbf{f}_{\theta}(\mathbf{x}_t^*))], \end{aligned} \quad (7)$$

where  $\beta$  is the trade-off parameter, the first item represents the cross-entropy loss on the fusion AEs and the second item  $\mathcal{L}_1$  represents the mean square loss with  $l_2$  norm.

---

**Algorithm 1: FGSM-PCO**


---

**Input:** Target model  $\mathbf{f}_{\theta}$ , maximum perturbation  $\epsilon$ , training epoch  $T$ ,  
mini-batch data  $\mathbb{B}$ , learning rate  $\alpha$ , amplification parameter  $\gamma$ .  
**Output:** Model weights  $\theta$

```

1 for  $t = 1, \dots, T$  do
2   for  $\{\mathbf{x}, \mathbf{y}\} \sim \mathbb{B}$  do
3     if  $t == 1$  then
4        $\mathbf{x}_{t-1}^* \leftarrow \mathbf{x} + U[-\epsilon, \epsilon]$ ;
5     end
6      $\mathbf{g}_t \leftarrow \text{sign}(\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{f}_{\theta}(\mathbf{x}_{t-1}^*), \mathbf{y}))$ ;
7      $\delta_t = \epsilon \mathbf{g}_t$ ;
8      $\mathbf{x}_{am}^* = \mathbf{x}_{t-1}^* + \gamma \delta_t$ ;
9      $\lambda_t \leftarrow f_{\theta}^k(\mathbf{x}_{t-1}^* + \delta_t)$  where  $\mathbf{y}[k] = 1$ ;
10     $\mathbf{x}_{train} \leftarrow \lambda_t \mathbf{x}_{t-1}^* + (1 - \lambda_t) \mathbf{x}_{am}^*$ ;
11     $\mathbf{x}_t^* \leftarrow \mathbf{x}_{t-1}^* + \epsilon \mathbf{g}_t$ ;
12     $\theta_{t+1} \leftarrow \theta_t + \alpha \frac{1}{|\mathbb{B}|} \sum_{\mathbb{B}} \nabla \mathcal{L}_{PCO}(\mathbf{f}_{\theta}(\mathbf{x}_{train}), \mathbf{y})$ ;
13  end
14 end
15 return  $\theta$ 
```

---


### 3.4 Mitigating Catastrophic Overfitting in FAT

Current FAT methods fall short of reversing the trend of catastrophic overfitting. Once the inner-level optimization process collapses, the algorithm can only progress towards overfitting. However, in FGSM-PCO, training examples are generated by a fusion method. When catastrophic overfitting occurs, this adaptive fusion mechanism ensures retain more of the previous AEs, which avoids the reliance on current ineffective AEs as shown in Fig.2. Moreover, the inner optimization problem is closely related to the model parameters. After updating these parameters, the inner layer optimization problem can also escape from being trapped in the current overfitting solution. In contrast, other FAT methods

solely focus on averting models from falling into catastrophic overfitting without addressing how to rectify this trend once it occurs.

To validate this capability, we use FGSM-AT and FGSM-MEP to train ResNet18 [15] on CIFAR10 [20] and PreActResNet18 [15] on Tiny-ImageNet [11], respectively. We switch to other FAT strategies when the model occurs catastrophic overfitting. Fig. 3 shows FGSM-AT occurs overfitting at 16 – *th* training epoch and FGSM-MEP occurs at 50 – *th* training epoch. When the inner optimization collapses, in both cases, our method FGSM-PCO can effectively correct the catastrophic overfitting issue.

In summary, the proposed framework with tailored training loss, can prevent the collapse of the inner optimization problem, thereby preventing the catastrophic overfitting phenomenon in FAT. Meanwhile, the adaptive fusion mechanism corrects the potential overfitting trend and ensures the efficacy of training.



**Fig. 3:** The performance of different FAT methods when the FGSM-AT and FGSM-MEP occur catastrophic overfitting.

## 4 Experimental Results

In this section, we present experimental results to evaluate the effectiveness of FGSM-PCO in comparison to other AT methods. After that, we conduct an ablation analysis to highlight the contributions of the proposed components. Experiments of CIFAR datasets are conducted on an NVIDIA RTX 4090 GPU, and experiments of Tiny-ImageNet are on an NVIDIA V100 GPU. The code will be released at <https://github.com/HandingWangXDGroup/FGSM-PCO>.

### 4.1 Experimental Settings

**Datasets and Models** We evaluate the performance of our method on CIFAR10, CIFAR100 [20] and Tiny-ImageNet [11] datasets. The CIFAR10 dataset consists of 50000 images across 10 classes, with 5000 images per class, while



CIFAR100 contains 100 classes with 500 images each. The size of each image in these two datasets is  $32 \times 32 \times 3$ . Tiny-ImageNet contains 100000 images of 200 classes (500 for each class), and each image is  $64 \times 64 \times 3$ . Our base models are ResNet18 [15], PreActResNet18 [15], and WideResNet34-10 [39].

**Compared Methods** We choose the PGD-AT [23] and FGSM-AT [13] as baselines for multi-step and single-step guided AT methods. In addition, we compare our FGSM-PCO method with six state-of-the-art AT methods, including FGSM-RS [38], FGSM-GA [2], FreeAT [38], TRADES [41], and FGSM-PGI [17], which includes FGSM-EP, and FGSM-MEP.

**Training Details** For all compared methods, the maximum perturbation is set to  $8/255$ . For FGSM-based AT methods, the attack step size to  $8/255$ , while for PGD-based AT methods, the step size is  $2/255$  over 10 iterations. We utilize the SGD optimizer with a momentum of 0.9 and a weight decay of  $5 \times 10^{-4}$ . The initial learning rate is set to 0.1, following the settings in [17, 27, 29]. The training is conducted for 110 epochs, with the learning rate being divided by 10 at the 100<sup>th</sup> and 105<sup>th</sup> epochs. For our FGSM-PCO parameters, the regularization coefficient  $\beta$  is set to 10, and the amplification parameter  $\gamma$  is set to 2.

**Evaluation Metrics** We evaluate the classification accuracy of all the AT models under the FGSM [13], PGD [23], C&W [4], APGD [9], FAB [8], Square [1] and AutoAttack (AA) [9], where PGD attack includes PGD10, PGD20 and PGD50 three version, APGD uses the DLR loss version. We report the result of the last checkpoint and the best checkpoint under the PGD10 attack. In addition, we also evaluate the classification accuracy of the clean examples to show the generalization of our method.

## 4.2 Comparative Results

**Results on CIFAR10 and CIFAR100** To demonstrate the effectiveness of our proposed method, we compared our method with other AT methods on ResNet18 with the CIFAR datasets. The comparison results on CIFAR10 are presented in Table 1. For the catastrophic overfitting models, we do not report the classification accuracy under the FGSM attack because those models are overfitting to the FGSM attack, which provides a false sense of security. Our method FGSM-PCO achieves 56.32% classification accuracy under PGD10 attack, which is 3.1% higher than the multi-step guided AT baseline PGD-AT and 1.2% higher than the state-of-the-art FGSM-based AT method FGSM-MEP [17]. Compared to the FGSM-based methods baseline FGSM-RS, our method achieves a 14% improvement. It is worth mentioning that, our model achieves the best classification accuracy at the last training epoch, indicating our method can effectively solve the catastrophic overfitting problem. Our method improves a little training cost compared with the FGSM-MEP but has an improvement in

**Table 1:** Accuracy (%) and training time (min) of compared AT models on ResNet18 with the CIFAR10 dataset. The number in bold indicates the best.

Method		Clean	Acc	FGSM	PGD10	PGD20	PGD50	C&W	APGD	Square	AA	Time
PGD-AT	Best	82.57	63.93	53.19	52.42	52.21	48.01	51.22	55.70	48.77		
	Last	82.99	64.20	53.05	52.14	51.96	47.51	50.66	55.23	48.23		199
TRADES	Best	82.03	64.28	54.06	53.35	53.16	46.22	50.66	55.83	49.47		
	Last	81.80	63.57	53.89	53.23	53.05	46.31	50.39	55.54	49.56		241
FGSM-RS	Best	72.95	52.64	41.40	40.57	40.40	36.83	39.95	44.75	37.67		
	Last	84.24	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00		40
FGSM-GA	Best	81.27	59.52	45.49	43.84	43.52	42.84	43.15	40.96	39.29		
	Last	<b>86.40</b>	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00		132
Free-AT	Best	80.64	57.62	44.10	42.99	42.73	39.88	42.83	48.82	41.17		
	Last	79.88	55.40	42.00	40.49	40.18	38.83	40.37	46.03	40.34		140
FGSM-EP	Best	79.28	65.04	54.34	53.52	53.54	45.84	47.85	52.42	45.66		
	Last	80.91	<b>73.83</b>	40.89	37.76	36.05	38.46	28.41	45.29	25.16		57
FGSM-MEP	Best	81.72	64.71	55.13	54.45	54.29	47.05	<b>50.39</b>	55.47	<b>48.23</b>		
	Last	81.72	64.71	55.13	54.45	54.29	47.05	<b>50.39</b>	55.47	<b>48.23</b>		57
Ours	Best	<b>82.05</b>	<b>65.53</b>	<b>56.32</b>	<b>55.66</b>	<b>55.67</b>	<b>47.12</b>	50.05	<b>55.59</b>	48.04		
	Last	82.05	65.53	<b>56.32</b>	<b>55.66</b>	<b>55.67</b>	<b>47.12</b>	50.05	<b>55.59</b>	48.04		60

both adversarial and clean examples. Although our method incurs a higher computational cost than FGSM-MEP, it saves memory on computational devices, requiring only two-thirds of the memory compared to FGSM-MEP.

**Table 2:** Accuracy (%) and training time (min) of compared AT models on WideResNet34-10 with the CIFAR100 dataset. The number in bold indicates the best.

Method		Clean	Acc	FGSM	PGD10	PGD20	PGD50	C&W	APGD	Square	AA	Time
PGD-AT	Best	62.45	41.24	32.36	31.66	31.41	27.78	30.56	34.20	27.64		
	Last	62.46	40.66	30.97	30.31	30.02	27.39	30.39	33.54	27.40		1397
TRADES	Best	61.23	40.46	32.14	31.58	31.56	25.85	28.87	32.62	27.60		
	Last	61.09	40.26	31.89	31.56	31.57	25.73	28.72	32.40	27.56		1692
FGSM-RS	Best	51.27	30.92	22.95	22.41	21.55	23.74	25.91	27.11	16.78		
	Last	63.11	-	0.00	0.00	0.00	0.00	0.00	0.00	0.00		281
FGSM-MEP	Best	43.42	28.92	23.77	23.47	23.53	18.28	20.39	22.34	18.34		
	Last	<b>72.96</b>	<b>65.65</b>	18.99	13.86	9.63	12.82	3.62	12.19	1.49		407
FGSM-PCO	Best	<b>65.80</b>	<b>40.41</b>	<b>29.80</b>	<b>28.71</b>	<b>28.61</b>	<b>24.91</b>	<b>27.45</b>	<b>31.94</b>	<b>24.96</b>		
	Last	65.38	40.40	<b>29.11</b>	<b>28.26</b>	<b>28.15</b>	<b>25.51</b>	<b>27.37</b>	<b>31.37</b>	<b>24.20</b>		421

On the CIFAR100 dataset with WideResNet34-10, we achieve the best performance both on adversarial and clean examples as shown in Table 2. Our method achieves 29.80% under PGD10 attack and 65.80% classification accuracy for clean examples at the best checkpoints. Even on the last training checkpoint, our method achieves 29.11% under PGD10 attack which is higher than the

**Table 3:** Numbers of catastrophic overfitting. 10 independent repeated experiments on the WideResNet34-10 with the CIFAR100 dataset.


Method	FGSM-AT	FGSM-RS	FGSM-GA	FGSM-MEP	FGSM-PCO
Numbers of overfitting	10/10	10/10	9/10	6/10	<b>0/10</b>

state-of-the-art FGSM-based AT method FGSM-MEP. Our method significantly improves the classification accuracy on clean examples. On the last training checkpoint, we achieve 65.38% accuracy, which is 7% lower than the FGSM-MEP. That is because FGSM-MEP occurs overfitting phenomenon, resulting in an ineffective perturbation. To avoid the randomness of experimental outcomes, we conduct 10 independent repeated experiments on the WideResNet34-10 with the CIFAR100 dataset and record the occurrences of catastrophic overfitting as shown in Table 3. All the FAT algorithms occurs catastrophic overfitting problems except FGSM-PCO. More experimental results are presented in the Supplementary Material.

**Table 4:** Classification accuracy (%) on the PreActResNet with Tiny-ImageNet dataset. The number in bold indicates the best.

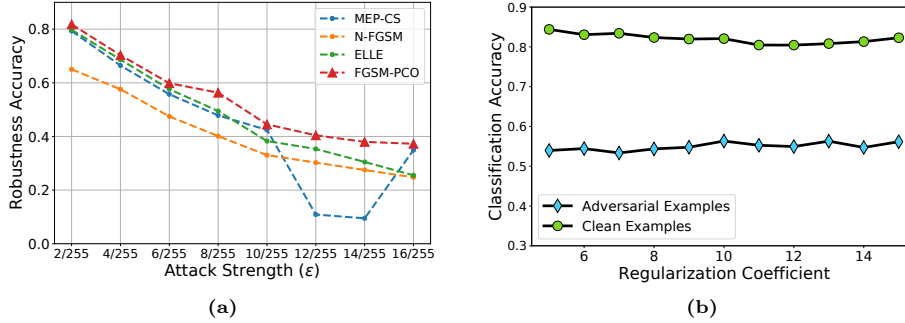
Method		Clean Acc	FGSM	PGD10	PGD20	PGD50	Training Time
PGD-AT	Best	33.99	19.56	15.35	15.29	15.16	1961.67
	Last	33.76	12.94	7.05	6.75	6.68	
FGSM-GA	Best	23.67	14.57	11.59	11.57	11.57	731.50
	Last	34.85	9.87	0.00	0.00	0.00	
FGSM-MEP	Best	31.70	20.51	16.81	16.74	16.69	523.23
	Last	<b>46.07</b>	-	3.09	1.98	1.36	
Ours	Best	34.96	<b>22.32</b>	<b>18.17</b>	<b>18.10</b>	<b>17.99</b>	686.67
	Last	37.67	<b>20.33</b>	<b>14.55</b>	<b>14.37</b>	<b>14.28</b>	

**Results on Tiny-ImageNet** To show the performance comprehensively and further invest in the catastrophic overfitting phenomenon, we test FGSM-PCO on PreActResNet18 with a more challenging dataset Tiny-ImageNet. We use the cyclic learning rate strategy [33] with a maximum learning rate of 0.2. The experimental results are shown in Table 4. We do not report the FGSM accuracy of the overfitting models. Although the classification accuracy of our method has a gap between the best and the last model under the PGD-10 attack, FGSM-PCO still achieves the best performance. FGSM-PCO obtains 18.17% accuracy under the PGD10 attack and 34.96% classification accuracy on clean examples. Al-



**Fig. 4:** The classification accuracy of PreActResNet18 with the Tiny-ImageNet dataset. The left figure shows the classification accuracy for AEs under PGD10 attack, and the right figure shows the accuracy for clean examples.

though FGSM-MEP utilizes the historical information to guide the initialization of the current perturbation, it occurs catastrophic overfitting shown in Fig. 4. Meanwhile, the classification on clean examples occurs in fluctuations. FGSM-MEP obtains 16.48% accuracy on AEs and 32.15% on clean examples, which is lower than our proposed method FGSM-PCO.



**Fig. 5:** 5a shows the classification accuracy under different attack strengths. 5b represents the sensitivity of  $\beta$  on classification for adversarial and clean examples.

**Results under Different Attack Strength** To further show the effectiveness of our proposed method, we evaluate the classification accuracy under different levels of attack budgets. We compared FGSM-PCO with three methods, N-FGSM [18], MEP-CS [43] and GAT-ELLE [30], which perform well under various attack strengths. The results are shown in Fig. 5a.

**Sensitivity Analysis** We investigate the regularization parameter  $\beta$  in Eq.(7) on ResNet18 with the CIFAR10 dataset. We present the accuracy on adversarial and clean examples as shown in Fig. 5b. The regularization parameter is set to different values  $\beta \in [5, 15]$ . The sensitivity experiments show that our method achieves good stability and robustness across different choices of  $\beta$ .

### 4.3 Ablation Study

To better understand the effects of each component and highlight our contributions, we conduct the ablation experiments on ResNet18 with the CIFAR10 dataset.

**Table 5:** Ablation study of the proposed method.

Component	Loss <sup>†</sup>	✗	✗	✓	✓	✗	✓
	Fusion <sup>‡</sup>	✗	✓	✓	✗	✓	✓
	Adaptive <sup>*</sup>	✗	✗	✗	✓	✓	✓
Clean Acc	Best	87.44	76.40	82.01	80.01	88.88	82.48
	Last	87.90	89.81	81.97	80.30	88.65	82.48
PGD10	Best	48.74	39.91	54.27	49.04	50.67	<b>56.12</b>
	Last	48.42	10.18	53.56	48.81	50.57	<b>56.12</b>

<sup>†</sup> represents the proposed regularization.


<sup>‡</sup> represents the perturbation fusion strategy.

<sup>\*</sup> represents the adaptive mechanism.

**Effect of Each Component** For the proposed loss function, AEs fusion strategy and adaptive fusion factor, we conduct experiments on ResNet18 on the CIFAR10 dataset and the results are presented in Table 5. The result shows that our proposed regularization can effectively prevent catastrophic overfitting and well guide the proposed AT method, achieving the best performance when the AT algorithm contains the three components.

**Fusion strategies** To demonstrate that the enhanced classification accuracy does not result from the generalization introduced by the mixup technology. We conduct experiments on ResNet18 with the CIFAR10 dataset, and three fusion strategies are compared. The first is our fusion method using the previous and current perturbation with an amplified parameter. The second is to fuse the amplified current AEs with clean examples (Amplification), and the third one is to fuse the current AEs and clean examples without amplification (Without Amplification). The results shown in Fig. 6, our proposed method achieves the

best performance among the three strategies without decreasing the classification accuracy for clean examples.



**Fig. 6:** The classification accuracy of different fusion strategies on ResNet18 with the CIFAR10. The left figure shows the classification accuracy for AEs under PGD10 attack, and the right figure shows the accuracy for clean examples.

## 5 Conclusions

In this study, we demonstrate that avoiding the collapse of the inner optimization problem can effectively prevent catastrophic overfitting in FAT. Our proposed method FGSM-PCO, fuses the previous AEs and an amplification AEs with an adaptive mechanism, which can correct the catastrophic overfitting trend. Our method is compared with the seven state-of-the-art AT methods on three datasets and three models, and the experimental results demonstrate the effectiveness of our method.

Although our method achieves good performance on the chosen datasets, it also has a limitation. The fusion coefficient of our method is decided according to the classification confidence of the AEs, which may become less effective when the method faces a vast number of categories in an open-world scenario. In the future, we plan to explore a more efficient connection between previous and current training epochs.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (No. 62376202, 62192734) and the Innovation Fund of Xidian University (YJSJ24012).

## References

1. Andriushchenko, M., Croce, F., Flammarion, N., Hein, M.: Square attack: a query-efficient black-box adversarial attack via random search. In: European Conference on Computer Vision. pp. 484–501. Springer (2020)
2. Andriushchenko, M., Flammarion, N.: Understanding and improving fast adversarial training. *Advances in Neural Information Processing Systems* **33**, 16048–16059 (2020)
3. Brendel, W., Rauber, J., Bethge, M.: Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. In: International Conference on Learning Representations (2018)
4. Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. In: 2017 IEEE Symposium on Security and Privacy (SP). IEEE (2017)
5. Chakraborty, A., Alam, M., Dey, V., Chattopadhyay, A., Mukhopadhyay, D.: Adversarial attacks and defences: A survey. *arXiv preprint arXiv:1810.00069* (2018)
6. Chen, P.Y., Zhang, H., Sharma, Y., Yi, J., Hsieh, C.J.: Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In: Proceedings of the 10th ACM workshop on artificial intelligence and security. pp. 15–26 (2017)
7. Colson, B., Marcotte, P., Savard, G.: An overview of bilevel optimization. *Annals of operations research* **153**, 235–256 (2007)
8. Croce, F., Hein, M.: Minimally distorted adversarial examples with a fast adaptive boundary attack. In: International Conference on Machine Learning. pp. 2196–2205. PMLR (2020)
9. Croce, F., Hein, M.: Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In: International conference on machine learning. pp. 2206–2216. PMLR (2020)
10. Dempe, S., Kalashnikov, V., Pérez-Valdés, G.A., Kalashnykova, N.: Bilevel programming problems. *Energy Systems*. Springer, Berlin **10**, 978–3 (2015)
11. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
12. Duan, R., Mao, X., Qin, A.K., Chen, Y., Ye, S., He, Y., Yang, Y.: Adversarial laser beam: Effective physical-world attack to dnns in a blink. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16062–16071 (2021)
13. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572* (2014)
14. Gu, T., Dolan-Gavitt, B., Garg, S.: Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733* (2017)
15. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14. pp. 630–645. Springer (2016)
16. Ilyas, A., Engstrom, L., Athalye, A., Lin, J.: Black-box adversarial attacks with limited queries and information. In: International conference on machine learning. pp. 2137–2146. PMLR (2018)
17. Jia, X., Zhang, Y., Wei, X., Wu, B., Ma, K., Wang, J., Cao, X.: Prior-guided adversarial initialization for fast adversarial training. In: European Conference on Computer Vision. pp. 567–584. Springer (2022)

18. de Jorge Aranda, P., Bibi, A., Volpi, R., Sanyal, A., Torr, P., Rogez, G., Dokania, P.: Make some noise: Reliable and efficient single-step adversarial training. *Advances in Neural Information Processing Systems* **35**, 12881–12893 (2022)
19. Kim, H., Lee, W., Lee, J.: Understanding catastrophic overfitting in single-step adversarial training. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 35, pp. 8119–8127 (2021)
20. Krizhevsky, A.: Learning multiple layers of features from tiny images (2009)
21. Li, Y., Jiang, Y., Li, Z., Xia, S.T.: Backdoor learning: A survey. *IEEE Transactions on Neural Networks and Learning Systems* (2022)
22. Liu, X., Xie, L., Wang, Y., Zou, J., Xiong, J., Ying, Z., Vasilakos, A.V.: Privacy and security issues in deep learning: A survey. *IEEE Access* **9**, 4566–4593 (2020)
23. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. In: *International Conference on Learning Representations* (2018)
24. Mao, C., Zhong, Z., Yang, J., Vondrick, C., Ray, B.: Metric learning for adversarial robustness. *Advances in neural information processing systems* **32** (2019)
25. Moosavi-Dezfooli, S.M., Fawzi, A., Frossard, P.: Deepfool: a simple and accurate method to fool deep neural networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2574–2582 (2016)
26. Moosavi-Dezfooli, S.M., Fawzi, A., Uesato, J., Frossard, P.: Robustness via curvature regularization, and vice versa. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9078–9086 (2019)
27. Pang, T., Yang, X., Dong, Y., Su, H., Zhu, J.: Bag of tricks for adversarial training. *arXiv preprint arXiv:2010.00467* (2020)
28. Qin, C., Martens, J., Goyal, S., Krishnan, D., Dvijotham, K., Fawzi, A., De, S., Stanforth, R., Kohli, P.: Adversarial robustness through local linearization. *Advances in neural information processing systems* **32** (2019)
29. Rice, L., Wong, E., Kolter, Z.: Overfitting in adversarially robust deep learning. In: *International Conference on Machine Learning*. pp. 8093–8104. PMLR (2020)
30. Rocamora, E.A., Liu, F., Chrysos, G., Olmos, P.M., Cevher, V.: Efficient local linearity regularization to overcome catastrophic overfitting. In: *The Twelfth International Conference on Learning Representations* (2024)
31. Shafahi, A., Najibi, M., Ghiasi, M.A., Xu, Z., Dickerson, J., Studer, C., Davis, L.S., Taylor, G., Goldstein, T.: Adversarial training for free! *Advances in Neural Information Processing Systems* **32** (2019)
32. Shokri, R., Stronati, M., Song, C., Shmatikov, V.: Membership inference attacks against machine learning models. In: *2017 IEEE symposium on security and privacy (SP)*. pp. 3–18. IEEE (2017)
33. 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)
34. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., fellow, I., Fergus, R.: Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199* (2013)
35. Wang, Y., Zou, D., Yi, J., Bailey, J., Ma, X., Gu, Q.: Improving adversarial robustness requires revisiting misclassified examples. In: *International Conference on Learning Representations* (2019)
36. Wang, Z., Pang, T., Du, C., Lin, M., Liu, W., Yan, S.: Better diffusion models further improve adversarial training. *arXiv preprint arXiv:2302.04638* (2023)
37. Wang, Z., Wang, H., Tian, C., Jin, Y.: Adversarial training of deep neural networks guided by texture and structural information. In: *Proceedings of the 31st ACM International Conference on Multimedia*. pp. 4958–4967 (2023)



38. Wong, E., Rice, L., Kolter, J.Z.: Fast is better than free: Revisiting adversarial training. *Learning* (2020)
39. Zagoruyko, S., Komodakis, N.: Wide residual networks. *arXiv: Computer Vision and Pattern Recognition* (2016)
40. Zhang, D., Zhang, T., Lu, Y., Zhu, Z., Dong, B.: You only propagate once: Accelerating adversarial training via maximal principle. *Advances in neural information processing systems* **32** (2019)
41. Zhang, H., Yu, Y., Jiao, J., Xing, E., El Ghaoui, L., Jordan, M.: Theoretically principled trade-off between robustness and accuracy. In: *International conference on machine learning*. pp. 7472–7482. PMLR (2019)
42. Zhang, Y., Zhang, G., Khanduri, P., Hong, M., Chang, S., Liu, S.: Revisiting and advancing fast adversarial training through the lens of bi-level optimization (2022)
43. Zhao, M., Zhang, L., Kong, Y., Yin, B.: Fast adversarial training with smooth convergence. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4720–4729 (2023)
44. Zhu, L., Liu, Z., Han, S.: Deep leakage from gradients. *Advances in neural information processing systems* **32** (2019)