

# Defending Adversarial Examples by Negative Correlation Ensemble

Wenjian Luo, *Senior Member, IEEE*, Hongwei Zhang, Linghao Kong, Zhijian Chen, Ke Tang, *Senior Member, IEEE*

**Abstract**—The security issues in DNNs, such as adversarial examples, have attracted much attention. Adversarial examples refer to the examples which are capable to induce the DNNs return completely predictions by introducing carefully designed perturbations. Obviously, adversarial examples bring great security risks to the development of deep learning. Recently, Some defense approaches against adversarial examples have been proposed, however, in our opinion, the performance of these approaches are still limited. In this paper, we propose a new ensemble defense approach named the Negative Correlation Ensemble (NCEn), which achieves compelling results by introducing gradient directions and gradient magnitudes of each member in the ensemble negatively correlated and at the same time, reducing the transferability of adversarial examples among them. Extensive experiments have been conducted, and the results demonstrate that NCEn can improve the adversarial robustness of ensembles effectively.

**Index Terms**—Deep learning, adversarial examples, ensemble, negative correlation

## I. INTRODUCTION

Deep Neural Networks (DNN) have achieved significant improvements in various domains, typically like image classification, face recognition and autonomous driving [1]–[6]. However, there are serious security issues in DNNs which have attracted much attention in recent years [7]. Specifically, studies have shown that DNN are vulnerable to the attacks from adversarial examples, which are generated by adding delicate and imperceptible perturbation to the benign samples, and aims to prompt the DNNs to make incorrect predictions [8].

Many algorithms have been proposed for generating adversarial examples [8]–[16]. These algorithms can be divided into white-box attacks and black-box attacks according to the accessibility. White-box attack algorithms, such as Fast Gradient Sign Method (FGSM) [8], Iterative Fast Gradient Sign Method (I-FGSM) [9], DeepFool [10] and C&W [11], exploit both

the network structures, hyperparameters and other available information of the target model, to calculate the gradient and such that generate adversarial examples. On the contrary, black-box attack algorithms, such as Zero Order Optimization (ZOO) [17], Autoencoder based Zero Order Optimization (AutoZoom) [18], Boundary Attack [12], HopJumpSkipAttack [13], Momentum Iterative Fast Gradient Sign Method (MI-FGSM) [14], Diverse Input Iterative Fast Gradient Sign Method (DIM) [15] and Skip Gradient Method (SGM) [16], are capable to generate adversarial examples without accessing the all the information of the target model (e.g. not accessing the target model but leveraging a substitution model). Notably, the black-box algorithms such as MI-FGSM can also be used in white-box setting scenarios, where they can directly access the target model.

Effective defenses against adversarial examples are usually achieved by detecting adversarial perturbations or improving the robustness of the model. The approaches of detecting adversarial perturbations, such as the key-based network in [19], the MagNet framework in [20] and the feature squeezing in [21], are implemented mainly by detecting or cleaning the input data through technical means, so that can discover the adversarial examples in advance or destroy some key structures constituting the adversarial examples. Differently, the approaches of improving the robustness of the model, such as the gradient regularization [22], adversarial training [23] and the defensive distillation [11], are implemented by changing some specific properties of the target model, such that enhance the robustness of the model to small perturbations.

Recently, some studies have found that, for a same learning task, different models will learn different decision boundaries due to the differences in their model structures, initial weights, and training methods, and by which we can infer that an adversarial example which can fool one model may not be capable to fool another ones [24]. Intuitively, the vulnerability of a single model can be avoided through model ensemble. Kariyappa et al. [25] have proposed the Gradient Alignment Loss (GAL) to improve the diversity of ensemble by reducing the dimension of shared adversarial subspace (ADV-SS). Pang et al. [26] have proposed a training method named Adaptive Diversity Promoting (ADP) to make the non-maximum predictions of different members in the ensemble orthogonal to each other, so that can improve the diversity of the ensemble without compromising the prediction accuracy. What's more, considering the influence of the gradient magnitude and the gradient angle in an ensemble, Dabouei et al. [27] have proposed joint Gradient Phase and Magnitude Regularization

This study is supported by the Major Key Project of PCL (Grant No. PCL2022A03, PCL2021A02, PCL2021A09), and the Guangdong Provincial Key Laboratory (Grant No. 2020B121201001). (Corresponding author: Wenjian Luo.)

Wenjian Luo, Hongwei Zhang, Linghao Kong and Zhijian Chen are with the School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen 518055, Guangdong, China. Wenjian Luo is also with Peng Cheng Laboratory, Shenzhen 518055, China.

Ke Tang is with Guangdong Provincial Key Laboratory of Brain-inspired Intelligent Computation, the School of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen 518055, Guangdong, China.

Email: luowenjian@hit.edu.cn, 20S151127@stu.hit.edu.cn, (20S151073, 21B951010)@stu.hit.edu.cn, tangk3@sustech.edu.cn.

(GPMR) based on GAL to obtain the effective defense interaction by considering both the optimal geometric boundary and the gradient magnitude of the members in the ensemble. However, although the ensemble methods can achieve adversarial robustness by increasing the diversity, the complementarity among members is neglected during the training which may lead to low Performance when facing various adversarial examples attack. For example, an adversarial example which can fool one member of the ensemble may also be capable to fool the other members due to its transferability.

In this paper, in order to leverage the interactions between members in the ensemble and reduce the transferability of adversarial examples among members, we propose to train classifiers in the ensemble based on the negative correlation principle and accordingly, design a new ensemble defense strategy (NCEn) to improve the adversarial robustness of ensembles. In NCEn, we make the gradient angle and the gradient magnitude of each member with respect to the input  $x$  negatively correlated, thus prompting the gradient directions to have the greatest diversity, and the gradient magnitude to be balanced. Specifically, for realizing the greatest diversity of gradient directions based on negative correlation, the cosine similarity between the gradients of any member and the ensemble is trained negatively correlated with the cosine similarity between the ensemble and all other members. So that we can reduce the number of members that are vulnerable to adversarial perturbation  $\epsilon$ . After this, we make the gradient magnitude of any member in the ensemble to be negatively correlated with the gradient magnitude of other members in the ensemble, which indicates that the prediction results of the ensemble will not be greatly compromised by the maximum gradient magnitude of individual members in the ensemble.

The contributions of this paper can be summarized as follows.

- We propose a novel negative correlation ensemble (NCEn). Based on the principle of negative correlation, NCEn maximizes the diversity of members in the ensemble and reduces the transferability of adversarial examples among members in the ensemble by constraining the gradient direction and the gradient magnitude of each member.
- Extensive experiments have been conducted, and the results show that the performance of the negative correlation ensemble (NCEn) exceeds the state-of-the-art ensemble-based defense strategies. What's more, the experimental results also show that NCEn can effectively reduce the transferability of adversarial examples as well as improving the diversity and robustness of the ensemble.

The rest of this paper is organized as follows. Section II introduces related work. We then introduce our approach in Section III. Section IV presents our experimental details and results. A brief conclusion of this paper is given in Section V.

## II. RELATED WORK

In this section, we first introduce the concept of adversarial examples as well as the generation and defense algorithms.

After this, we will introduce the adversarial robustness of ensembles and negative correlation learning, respectively.

### A. Adversarial Examples

In 2014, Szegedy et al. [7] have proposed the concept of adversarial examples. Given a trained classifier  $f(\cdot)$  and a clean input sample  $x$  with the corresponding label  $y$ , the adversarial example  $x' = x + \epsilon$  satisfies  $f(x') \neq y$ , where  $\epsilon$  is a subtle perturbation that is not easily detected by human eyes.

Numerous adversarial examples generation algorithms have been proposed in past years [7]–[11], [14], [28], [29]. Goodfellow et al. [8] have proposed the FGSM algorithm based on the linear hypothesis of the high-dimensional space of deep neural networks. The advantage of FGSM is that it can generate adversarial examples quickly, and can be applied to a variety of deep neural networks. The momentum iterative FGSM (MI-FGSM) [14] is an iterative attack method based on FGSM. By adding a momentum term to the iterative attack method, the transferability of adversarial examples is improved. The project Gradient Descent (PGD) [28] is another iterative attack method based on FGSM. In each iteration process, the perturbation is clipped to the specified range to generate efficient adversarial examples. In 2017, Kurakin et al. [9] have proposed the basic iterative method (BIM). By limiting each pixel of the adversarial example to be within the  $l_p$  field of the original example  $x$ , they constructed the adversarial examples in the real scenarios. Compared to FGSM, BIM makes fewer changes to the original sample and has better attack Performance.

In order to defend against the adversarial examples, a lot of methods have been proposed, such as adversarial training, networks modification, feature squeezing as well as using additional networks [11], [19]–[23], [30]–[35]. Specifically, the method of adversarial training improves the robustness of the model by generating adversarial examples and updating model parameters alternatively [7]–[10], [28]. In order to keep the effectiveness, this method needs to use high-intensity adversarial examples, and at the same time, the network should have sufficient expressive power. Therefore, such method requires a large amount of adversarial example training data, so it is also called brute force adversarial training. In addition, existing works show that adversarial training can help regularize the models and mitigate overfitting [1], [36]. However, these adversarial trained models still lack robustness to unseen adversarial examples.

### B. Adversarial Robustness of Ensembles

An efficient way to defend against adversarial examples is to use an ensemble of deep neural networks [24]–[26], [37]. By using different model structures, initial weights, and training methods for different models can significantly improve the diversity. Strauss et al. [38] have shown that neural networks the ensemble can not only improve the prediction accuracy, but also improve the robustness to adversarial examples. What's more, Tramèr et al. [39] have proposed to use the adversarial training to enhance the adversarial robustness of an ensemble.

Since the gradients of the models in the ensemble share similar directions, an adversarial example which can fool one model may also be capable to fool the others in the ensemble. For this, Kariyappa et al. [25] have proposed Gradient Alignment Loss (GAL) to improve the adversarial robustness of the ensemble by considering the diversity of gradient directions. Specifically, it focuses on training the ensemble whose members have irrelevant loss functions by using diversified training. Intuitively, GAL reduces the dimension of adversarial subspace of the models through diversified training. However, GAL does not consider the optimal geometric boundary for the diversification of gradient directions in the ensemble, and does not balance the gradient magnitude of each member in the ensemble. Based on this, Dabouei et al. [27] have proposed gradient phase and magnitude regularization (GPMR) [27]. The basic principle of GPMR is to increase the lower bound of the adversarial perturbation that changes the score of the classifier by considering the optimal geometric boundary to diversify the gradient direction in the ensemble, and to balance the gradient magnitude of the members, thereby constructing the first-order defense interaction of the members in the ensemble. However, GPMR does not fully consider the interaction between members in the ensemble. In contrast, our work builds a good ensemble defense system by considering the interaction between members in the ensemble, thereby improving the robustness of the ensemble.

### C. Negative Correlation Learning

Negative correlation ensemble has been proposed by Liu and Yao in 1999 [40], [41]. Liu and Yao [40] have proposed a cooperative ensemble learning system (CELS), and the purpose is to interactively and simultaneously train all individual models in the ensemble in a single learning process through negative correlation learning (NCL). Through negative correlation learning, different models in the ensemble can learn different features of the training set, so that the ensemble can work more holistically. Theoretical and experimental results show that NCL can promote the diversity of the models as well as keeping a high prediction accuracy.

Liu et al. [42] also have tried to solve the problem of the optimal number of neural networks in the ensemble on the premise of maintaining the good interaction of individual members in the ensemble through negative correlation learning and evolutionary learning. Chan et al. [43] proposed NCCD, which implemented negative correlation learning via correlation corrected data. NCCD does not add penalty items in the training process, but adds error related information to the training data for negative correlation learning. NCCD can reduce the communication bandwidth between individual networks, and can be applied to the ensemble of any type of network structures. In addition, because NCCD does not modify the error function of the network for negative correlation learning, but modifies the training data for negative correlation learning, it can accelerate the learning speed through parallel computing.

Wang et al. [44] proposed a new negative correlation learning algorithm AdaBoost.NC. The flexibility of AdaBoost is used to overcome the shortcomings of NCL, such as sensitivity

to parameters setting and long training time, and the overfitting problem of AdaBoost is solved by introducing diversity [45], [46]. The experimental results show that AdaBoost.NC has better generalization performance than NCCD, and the time cost is significantly less than CELS and NCCD.

In this paper, for the first time, we propose a defense approach against adversarial examples by leveraging the negative correlation method. We first make the gradient direction of each member in the ensemble negatively correlated. Secondly, we make the gradient magnitude of each member in the ensemble negatively correlated with the average gradient magnitude, such that help improve the diversity of the ensemble and balance the gradient magnitude of each member in the ensemble. The details will be given in the next section.

## III. THE PROPOSED METHOD

The purpose of defense is to enable the classifier  $f$  give a correct prediction on examples  $x'$  which is with the adversarial perturbation. The ensemble strategy exploits multiple trained models to make decisions together. Due to the differences in structures, initial weights and training methods, different models could learn different decision boundaries. This means that an adversarial example which can fool one model may not be capable to fool the others. So the ensemble could have better adversarial robustness compared to a single model.

In this section, we will show how negative correlation theory can be utilized in ensemble training and how it can be exploited to improve the adversarial robustness.

### A. On Gradient Directions of Members

The gradient of the loss to the input  $x$  refers to a direction where the directional derivative of the loss achieves a maximum along this direction, i.e., the loss changes fastest when a perturbation  $\epsilon$  is introduced along this direction. When the members in the ensemble have different gradient directions, we put the adversarial example  $x'$  with perturbation  $\epsilon$  along a certain direction into the ensemble for prediction, and  $x'$  could significantly change only the loss function of classifiers with gradient direction similar to  $\epsilon$ , but not affect others.

We treat the mean value of the gradients of all members as the gradient of the ensemble, which can be represented by  $\nabla J_{ensemble}$ . The adversarial perturbation added along the direction of  $\nabla J_{ensemble}$  can affect most members of the ensemble. We use  $\nabla J_i$  to represent the gradient of the  $i^{th}$  member with respect to input  $x$ . Firstly, we prompt the angle between  $\nabla J_i$  and  $\nabla J_{ensemble}$  negatively correlated with the angle between  $\nabla J_j$  ( $j \neq i$ ) of other members and  $\nabla J_{ensemble}$ . After negative correlation training, the gradient direction of all members in the ensemble with respect to whole dataset will be maximally different, which means that an adversarial perturbation increasing the loss function of one member will not absolutely increase the loss of other members.

In detail, we use cosine similarity (CS) to measure the gradient angle between members, and it can be calculated as Formula (1).

$$CS(\nabla J_i, \nabla J_j) = \frac{\langle \nabla J_i, \nabla J_j \rangle}{\|\nabla J_i\|_2 \cdot \|\nabla J_j\|_2} \quad (1)$$

The smaller the CS, the larger the angle between the gradients, and the greater the diversity among the members. This means that the loss of each member will not grow in a positive correlation manner for a same adversarial perturbation. Therefore, it is hard for an adversarial perturbation to fool all members simultaneously. Meanwhile, the ensemble with increased diversity is still capable to keep a high prediction accuracy. The relevant regularization term can be expressed as Formula (2).

$$\text{Loss}_{\text{cos}} = \text{CS}(\nabla J_i, \nabla J_{\text{ensemble}}) \sum_{j \neq i}^k \text{CS}(\nabla J_j, \nabla J_{\text{ensemble}}) \quad (2)$$

$$\nabla J_{\text{ensemble}} = \frac{1}{k} \sum_{i=1}^k \nabla J_i \quad (3)$$

Where  $\nabla J_{\text{ensemble}}$  is the averaged gradients of all members,  $k$  is the number of members, and  $\nabla J_i$  is the gradient of the  $i^{\text{th}}$  member.

In the training process, the regularization term aims to prompt the gradient direction of each member in the ensemble negatively correlated with the others. From Fig. 1, where  $\nabla J_i$  indicates the model gradient of the  $i^{\text{th}}$  member ( $i = 0, 1, 2$ ), and  $\nabla J_{\text{ensemble}}$  indicates averaged gradients of all members, the follows can be observed:

- When  $\sum_{j \neq i}^k \text{CS}(\nabla J_j, \nabla J_{\text{ensemble}})$  is greater than zero, the gradient directions of  $\nabla J_j$  are closer to  $\nabla J_{\text{ensemble}}$ . In such a situation, minimizing the negative correlation formula (2) will result in a larger gradient angle between  $\nabla J_i$  and  $\nabla J_{\text{ensemble}}$ .
- When  $\sum_{j \neq i}^k \text{CS}(\nabla J_j, \nabla J_{\text{ensemble}})$  is negative, the gradient directions of  $\nabla J_j$  is inconsistent with that of  $\nabla J_{\text{ensemble}}$ . In such a situation, minimizing the negative correlation formula (2) will result in the gradient direction of  $\nabla J_i$  to approach the gradient direction of  $\nabla J_{\text{ensemble}}$ .

In both cases, the regularization term will help all members train interactively, such that facilitates the maximization of the gradient direction diversity.

We use CS as the regularization term for enhancing the adversarial robustness, and the loss function can be expressed as Formula (4), where  $CE$  is the mean of the cross-entropy loss and  $\lambda_{\text{cos}}$  is the weight coefficient.

$$\text{Loss} = CE + \lambda_{\text{cos}} \text{Loss}_{\text{cos}} \quad (4)$$

$$CE = \frac{1}{k} \sum_{i=1}^k ce_{\text{loss}} \quad (5)$$

The algorithm for calculating the regularization term of the  $i^{\text{th}}$  member on the gradient direction using negative correlation is shown in Algorithm 1, where  $k$  is the number of members in the ensemble. Firstly, it calculates the cosine similarity  $\text{cos1}$  between the gradient of the  $i^{\text{th}}$  member and the ensemble as well as the sum of cosine similarity  $\text{cos2}$  between the gradients of the other members and the ensemble. After this, the negative correlation coefficient of  $\text{cos1}$  and  $\text{cos2}$  are calculated as the

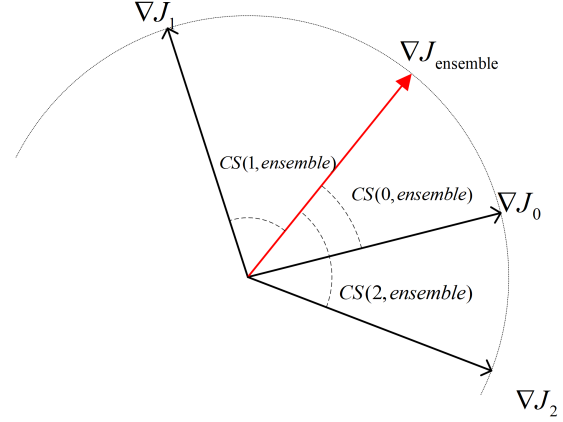


Figure 1: Improve the ensemble diversity from the view of gradient direction;  $\nabla J_i$  indicate the model gradient of the  $i^{\text{th}}$  member ( $i = 0, 1, 2$ ), and  $\nabla J_{\text{ensemble}}$  indicate the ensemble gradient

---

**Algorithm 1** Calculating  $\text{Loss}_{\text{cos}}$  for the  $i^{\text{th}}$  member

---

**Require:** Gradient of the  $i^{\text{th}}$  member:  $\nabla J_i$ ; Gradient of other members in the ensemble:  $\nabla J_j$  ( $j \neq i$ ); Ensemble gradient:  $\nabla J_{\text{ensemble}}$ ; Number of models in the ensemble:  $k$

**Ensure:** Regularization term of the  $i^{\text{th}}$  member:  $\text{Loss}_{\text{cos}}$

```

cos1  $\leftarrow$  CS( $\nabla J_i$ ,  $\nabla J_{\text{ensemble}}$ )
cos2  $\leftarrow$  0;
for  $j \leftarrow 1$  to  $k$  do
  if  $j \neq i$  then
    cos2  $\leftarrow$  cos2 + CS( $\nabla J_j$ ,  $\nabla J_{\text{ensemble}}$ )
  end if
end for
Losscos  $\leftarrow$  cos1 * cos2
Return the regularization term: Losscos

```

---

regularization term  $\text{Loss}_{\text{cos}}$  of the  $i^{\text{th}}$  member and will be minimized in the training process.

### B. On the Gradient Magnitudes of Members

The gradient magnitude represents the magnitude of the change in the loss caused by the adversarial perturbation  $\epsilon$ . Adding an adversarial perturbation  $\epsilon$  along the gradient direction will influence the loss more if there is a larger gradient magnitude. In the white-box attack scenario, the attacker can easily attack the classifiers with the largest gradient magnitude on the original input. We use  $g$  to represent the mean value of the gradient magnitudes of all members, and use  $\nabla J_i$  as the gradient of the  $i^{\text{th}}$  member in the ensemble. According to the negative correlation principle, we make the gradient magnitude of the  $i^{\text{th}}$  member and the ensemble gradient magnitude negatively correlated with the gradient magnitude of other members in the ensemble. After the training, the gradient magnitude of all members in the ensemble will be negatively correlated. This means that when the gradient magnitude of one member in the ensemble decrease, the gradient magnitude of other members will increase or decrease. However, no matter how the gradient magnitude of a single member changes, after all

members in the ensemble conduct simultaneous interactive training through negative correlation, each member will get the best gradient magnitude over dataset.

Therefore, we propose the second regularization term, which is shown in Formula (6).

$$\text{Loss}_{\text{norm}} = \frac{1}{g^2} (\|\nabla J_i\|_2 - g) \sum_{j \neq i}^k (\|\nabla J_j\|_2 - g) \quad (6)$$

$$g = \frac{1}{k} \sum_{i=1}^k \|\nabla J_i\|_2 \quad (7)$$

Where  $g$  is the mean value of the gradient magnitudes of all members,  $k$  is the number of members in the ensemble and  $\nabla J_i$  represents the gradient of the  $i^{\text{th}}$  member. We use the  $L_2$  norm of the gradient to calculate the gradient magnitude of each member in the ensemble. By minimizing the regularization term, the gradient magnitudes of all members in the ensemble is negatively correlated, which means different members can learn better over different features.

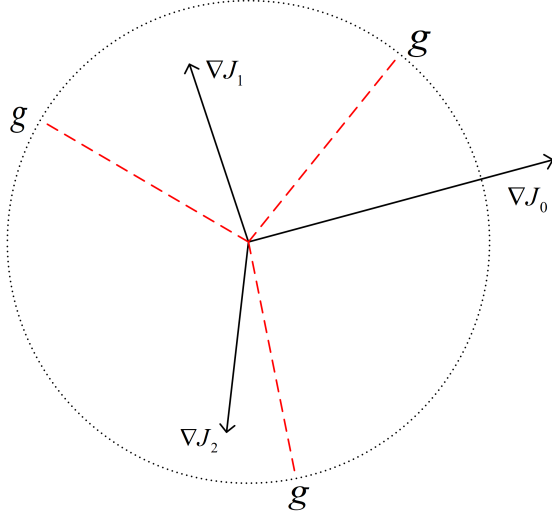


Figure 2: Improving the ensemble diversity from the view of the gradient magnitude;  $\nabla J_i$  indicates model gradients in the ensemble ( $i = 0, 1, 2$ ), and  $g$  indicates the mean value of the gradient magnitudes of all members in the ensemble

As shown in Fig. 2, where  $g$  is the mean value of the gradient magnitudes of all members in the ensemble, and there must be some gradient magnitudes less than  $g$ , so it can be inferred that the value of the regularization term will always be negative. Based on this, we can further observe that:

- When  $\sum_{j \neq i}^k (\|\nabla J_j\|_2 - g)$  is greater than zero,  $\|\nabla J_i\|_2$  is less than  $g$ . In such a situation, minimizing the negative correlation formula (6) will result in the decrease of the gradient magnitude  $\|\nabla J_i\|_2$ .
- When  $\sum_{j \neq i}^k (\|\nabla J_j\|_2 - g)$  is negative,  $\|\nabla J_i\|_2$  is greater than  $g$ . In such a situation, minimizing the negative correlation formula (6) will result in the increase of the gradient magnitude  $\|\nabla J_i\|_2$ .

Considering all members are trained simultaneously, different members can learn different features over the training

dataset and have different gradient magnitudes on the same input since it is not easy to always own very small gradient magnitudes over the whole input space. Based on this, it can be avoided that the adversaries attack the ensemble successfully by attacking only a few models with larger gradient magnitude.

We use this to regularize the training for improving the ensemble diversity and the loss function can be expressed as Formula (8). Where  $CE$  is the cross-entropy loss and  $\lambda_{\text{norm}}$  is the weight coefficient.

$$\text{Loss} = CE + \lambda_{\text{norm}} \text{Loss}_{\text{norm}} \quad (8)$$

The algorithm for calculating the  $\text{Loss}_{\text{norm}}$  with negative correlation is shown in Algorithm 2, where  $k$  is the number of members in ensemble. Firstly, we calculate the difference  $\text{norm1}$  between the gradient magnitude of the  $i^{\text{th}}$  member and the mean value  $g$  as well as the difference  $\text{norm2}$  between the gradient magnitude of other members in the ensemble and the mean value  $g$ . After this, according to the negative correlation method, the  $\text{Loss}_{\text{norm}}$  of  $\text{norm1}$  and  $\text{norm2}$  is calculated and treated as a regularization term for the training.

---

**Algorithm 2** Calculating  $\text{Loss}_{\text{norm}}$  for the  $i^{\text{th}}$  member

---

**Require:** Gradient of the  $i^{\text{th}}$  member:  $\nabla J_i$ ; Gradient of other members in the ensemble:  $\nabla J_j (j \neq i)$ ; the mean value of the gradient magnitude:  $g$

**Ensure:** Regularization term:  $\text{Loss}_{\text{norm}}$

$\text{norm1} \leftarrow (\|\nabla J_i\|_2 - g) / g$

$\text{norm2} \leftarrow 0$

**for**  $j \leftarrow 1$  **to**  $k$  **do**

**if**  $j \neq i$  **then**

$\text{norm2} \leftarrow \text{norm2} + (\|\nabla J_j\|_2 - g) / g$

**end if**

**end for**

$\text{Loss}_{\text{norm}} \leftarrow \text{norm1} * \text{norm2}$

**Return** the regularization term:  $\text{Loss}_{\text{norm}}$

---

### C. The Proposed NCEn

We have proposed two methods to increase the adversarial robustness, however, we find that using each method alone is still limited in performance.

Specifically, if only considering the influence of the gradient direction of each member in the ensemble, an adversary can attack a few members with large gradient magnitudes to make the ensemble predict incorrectly. While if only considering the influence of the gradient magnitude of each member in the ensemble, there could be a phenomenon that the gradient directions of all members are similar. At this time, the loss of each member in the ensemble will grow positively corrected, i.e., adversarial examples generated along the gradient direction can make most members of the ensemble predict incorrectly. So that the defense performance of the ensemble is similar to that of a single model, which is not robust enough.

Therefore, we will consider both the influence of member gradient directions and the influence of member gradient magnitudes, and use two regularization terms simultaneously

to improve the adversarial robustness. The loss function is shown in Formula (9), where  $CE$  is cross-entropy loss and the combined regularization term  $NCE$  can be calculated by Formula (10).

$$\text{Loss} = CE + NCE \quad (9)$$

$$NCE = \frac{1}{k} \sum_{i=1}^k (\lambda_{\text{norm}} \text{Loss}_{\text{norm}} + \lambda_{\text{cos}} \text{Loss}_{\text{cos}}) \quad (10)$$

The specific implementation of the ensemble training process is shown in Algorithm 3. In each epoch, we first get the predicted value  $pred$  of each member in the ensemble. Next, we use the predicted value  $pred$  and the real label  $label$  to calculate the cross entropy loss  $ce_{loss}$  of each member in the ensemble. Then, we calculate the mean value  $CE$  of the cross-entropy loss of all members in the ensemble. Finally, we use Formula (10) to calculate the regularization term  $NCE$ , and by which we can update the parameters of all members.

---

**Algorithm 3** Negative Correlation Ensemble

---

**Require:** Dataset:  $X$ ; Correct label for dataset  $X$ :  $label$

**Ensure:** Trained ensemble

Get a list of all the members in ensemble:  $f_i (i = 1, 2, \dots, k)$

**for**  $epoch \leftarrow begin\_epoch$  **to**  $end\_epoch$  **do**

**for**  $model \leftarrow f_1$  **to**  $f_k$  **do**

$pred \leftarrow model(X)$

$ce_{loss} \leftarrow CE(pred, label)$

**end for**

    Calculate  $CE \leftarrow \frac{1}{k} \sum_{i=1}^k ce_{loss}$

    Get the value of the regularization term:  $NCE$

    Back propagation using the sum of  $CE$  and  $NCE$

    Update the model parameters of each model

**end for**

---

#### IV. EXPERIMENTS

In this section, we conduct experiments on the FashionMNIST and CIFAR-10 datasets. We first give our experimental settings. Secondly, we compare the adversarial robustness of NCEn with different ensemble defense approaches. Furthermore, we find that NCEn is capable to reduce the transferability of adversarial examples, which proves that NCEn leverages the interactions of different members to improve the adversarial robustness. The source codes of this work are available at <https://github.com/MiLabHITSZ/2022ZhangNCEn>

##### A. Experimental Setup

In our experiments, we use three ensemble defense strategies as baselines to evaluate the performance of NCEn. The first one is the ensemble training without any regularization term, which  $\lambda_{cos} = \lambda_{norm} = 0$ . The second one is to use GAL for diversified training to improve the adversarial robustness of the ensemble [25]. The third one is the GPMR proposed by Dabouei et al. [27], which constructs the first-order defense interaction of the members in the ensemble to improve the adversarial robustness of the ensemble. Both GAL and GPMR are state-of-the-art ensemble defense strategies.

Table I: Ensemble model structures.  $k = (3, 4, 5)$

Name	Ensemble Structures
Ensemble1	$k * ResNet20$
Ensemble2	$k * ResNet26$
Ensemble3	$k * ResNet32$
Ensemble4	$(k - \lfloor \frac{k-1}{2} \rfloor) * ResNet20 + \lfloor \frac{k-1}{2} \rfloor * ResNet26 + 1 * ResNet32$

In our experiments, we use  $k = 3, 4, 5$  models in the ensemble respectively for analyzing the impact of the number of members, as shown in Table I. We evaluate the adversarial robustness of the ensemble on the FashionMNIST and CIFAR-10 datasets. For GAL, the coefficient of the regularization term is set to 0.5. For GPMR,  $\lambda_{div}$  of FashionMNIST is set to 0.1,  $\lambda_{div}$  of CIFAR-10 is set to 0.04, and  $\lambda_{eq}$  is set to 10 for all datasets.

We train the ensemble with Adam optimizer [47]. The initial learning rate is set to  $10^{-3}$ , and decays with the factor of 0.1 every 15 epochs until reaching the final learning rate  $10^{-5}$ . We train 40 epochs on FashionMNIST, and 60 epochs on CIFAR-10. In the experiments, the batch size of all training processes is set to 64. We set  $\lambda_{cos} = 0.02$  and  $\lambda_{norm} = 0.02$  for FashionMNIST,  $\lambda_{cos} = 0.06$  and  $\lambda_{norm} = 0.04$  for CIFAR-10.

In FashionMNIST, we use random cropping, and in CIFAR-10 we use random horizontal flipping and random cropping. In addition, we add a new dataset  $D_{noise}$  to the dataset  $D$ , which is generated by the perturbation from the truncated normal distribution:  $N(\mu = 0, \sigma = \epsilon/2)$ ,  $\epsilon = 0.3$  for FashionMNIST and  $\epsilon = 0.09$  for CIFAR-10. We use the synthetic data set  $D + D_{noise}$  to train the ensemble. Such an operation can effectively distort the useless high-frequency features and prompt the ensemble to avoid overfitting during the training.

We use test accuracy of clean examples (ACE) and test accuracy of adversarial examples (AAE) as the evaluation metrics. AAE refers to the ratio that a new example set, which is constructed by adding adversarial perturbations to the clean examples, can still be correctly classified by the ensemble. We use the averaged  $CE$  loss of all models in the ensemble as the objective function of the attack.

##### B. Defense Performance

We compare the defense performance of NCEn with different defense methods, as shown in Tables II~IV, and the best defense results are bolded. We use several powerful white box attack methods to test the defense performance of ensembles, including the fast gradient sign method (FGSM) [8], the momentum iterative fast gradient sign method (MI-FGSM) [14], projected gradient descent (PGD) [28], the basic iterative method (BIM) [9]. The specific settings of attack methods are as follows. FGSM adds a perturbation with a step length of  $\epsilon$  in the gradient direction. For MI-FGSM, each step size is set to 0.01, while keeping the maximum distortion always within  $\epsilon$  of the initial point. For PGD, the number of iterations is set to 40, and the maximum perturbation and single-step attack steps are both  $\epsilon$ . For BIM, the number of iterations is set to 10,

and the maximum perturbation and single-step attack steps are both  $\epsilon$ . All attack methods are implemented by AdvTorch.

When the number of models in the ensemble is 3, the results of different approaches on dataset FashionMNIST/CIFAR-10 are shown in Table II. Compared with BL, GAL and GPMR, the adversarial robustness of NCEn is significantly improved. When we change the ensemble model structure to Ensemble2, Ensemble3 and Ensemble4 respectively, the ACE of GAL drops below 0.8. At this time, the AAE of GAL is unworthy of consideration. Such AAE are marked with gray background as well as the corresponding ACE. Unlike GAL, NCEn can not only maintain high ACE in all ensemble model structures, but also achieve better adversarial robustness. Specifically, on FashionMNIST, except that FGSM and MI-FGSM use a large attack step, NCEn has better adversarial robustness in all ensemble structures. AAE has the highest improvement of 0.14 on Ensemble4 and MI-FGSM ( $\epsilon = 0.1$ ). On CIFAR-10, the adversarial robustness of NCEn is better or very close to the existing best baselines, and AAE can be improved by 0.11 at most.

When the number of models in the ensemble is 4 or 5, the results of different defense strategies on FashionMNIST/CIFAR-10 are shown in Tables III~IV. The experimental results are consistent with the previous ones, and NCEn still has better adversarial robustness. Specifically, when the number of models in the ensemble is 4, the AAE score of NCEn can be increased by 0.12 on FashionMNIST. On CIFAR-10, the AAE of NCEn is higher than all baselines, and the highest AAE can be improved by 0.12. When the number of models in the ensemble is 5, the AAE of NCEn can be increased by 0.12 on FashionMNIST. On CIFAR-10, the highest AAE of NCEn can be improved by 0.11.

### C. Transferability Between Members

Transferability refers to the success rate of being able to attack other models at the same time when an adversarial example is designed to attack a particular model. The less transferable the adversarial example is, the better the diversity of the models

Therefore, we can use the transferability of adversarial examples between different models to evaluate the similarity between members in the ensemble. We perform transferability experiments using PGD [28] and MI-FGSM [14], which MI-FGSM performs well in black-box attacks and PGD is the most powerful first-order attack algorithm. We generate adversarial examples for each member, and then evaluate their transferability on other members by calculating AAE. The perturbation magnitude of all attacks is set to  $\epsilon = 0.05$ , we use Ensemble4 in Table I to conduct the transferability experiments, and the results are shown in Fig. 3.

As shown in Fig. 3, we use the heat map of the confusion matrix to show transferability. The  $i$ th row and  $j$ th column in the heat map represents the test accuracy of the adversarial examples on the  $j$ th member, while the adversarial examples are generated over the  $i$ th member. When the values of other positions in the confusion matrix are close to those of the diagonal, it means that the adversarial examples generated

by the  $i$ th model can successfully attack other models in the ensemble. The closer the values in the confusion matrix are, the higher the transferability of the adversarial examples and the higher the similarity between members in the ensemble. On the contrary, it means that the diversity among members in the ensemble is higher, and the ensemble has better adversarial robustness. It can be seen from the confusion matrix that the transferability of adversarial examples in NCEn ensemble is poor, which indicates that NCEn has better diversity, such that it can provide better defense interactions for members.

### D. Parameter Analysis

From the experiment, we can find that the weights of the regularization terms will affect the defense performance. In order to explore the relationship between the weights and the defense performance, we summarize the experimental results and plot the contour heat map on  $\lambda_{cos}$  and  $\lambda_{norm}$  (see Figure 4). We use the product of ACE and AAE (which is averaged over four attack methods) in Table II to represent the intensities of the heat map. The larger the intensity, the higher the attack success rate. We use Ensemble1 in Table I and test on CIFAR-10. From Figure 4, we can get that the best results are obtained when we set  $\lambda_{cos} = 0.06$  and  $\lambda_{norm} = 0.04$ .

## V. CONCLUSION

In this paper, we propose a practical and feasible adversarial examples defense scheme based on negative correlation ensemble, which is named NCEn. We use the negative correlation principle to make the gradient direction and gradient magnitude of each member in the ensemble negatively correlated, and to train all members in the ensemble interactively and simultaneously. The purpose of negative correlation training in NCEn is to produce the best defense performance for the whole ensemble. Experiments have shown that NCEn can reach a better adversarial robustness than other ensemble defense schemes. Then, we demonstrate that NCEn can reduce the transferability of adversarial examples between members in the ensemble through the confusion matrix. Finally, we discuss the influence of hyperparameters  $\nabla\lambda_{cos}$  and  $\nabla\lambda_{norm}$  on defense performance, and give the best settings of hyperparameters. In general, we can conclude that NCEn is capable to improve the diversity and the robustness of the ensemble by making the gradient direction and gradient magnitude negatively correlated. However, for a given task, how to find the optimal number of the models in the ensemble is still a question. Many models could not be better. In the future, we will study how to automatically set the optimal number of the models in the ensemble.

## REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. E. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Advances in neural information processing systems*, vol. 25, pp. 1097–1105, 2012.
- [3] S. Yan, Y. Xiong, and D. Lin, "Spatial temporal graph convolutional networks for skeleton-based action recognition," in *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*. AAAI, 2018, pp. 7444–7452.



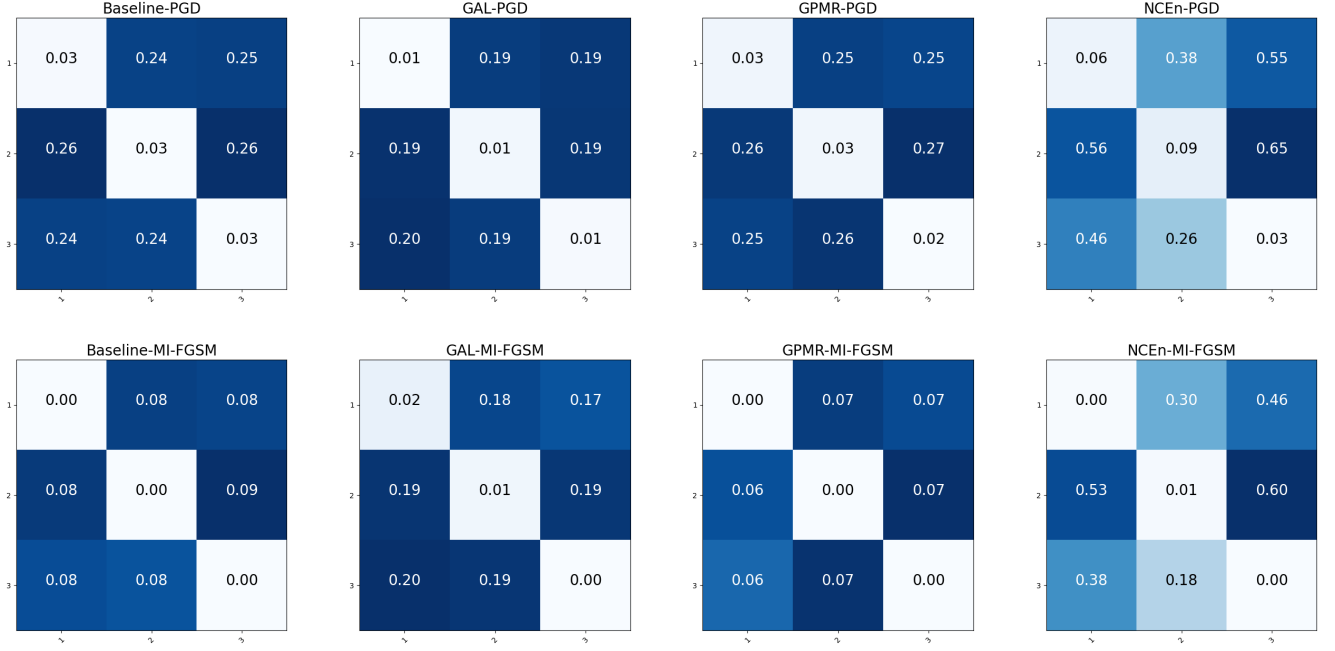


Figure 3: The confusion matrix of the accuracy of the adversarial examples test on CIFAR-10 by each member in the ensemble. Rows and columns respectively represent the model for generating adversarial examples and the model for testing the accuracy of adversarial examples.

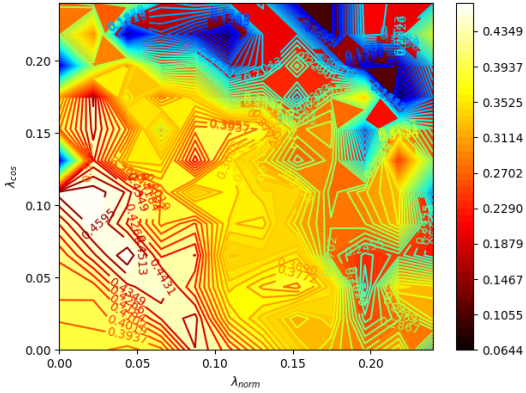


Figure 4: Contour Heat Map of  $\lambda_{cos}$  and  $\lambda_{norm}$

- [4] Y. Sun, Y. Chen, X. Wang, and X. Tang, "Deep learning face representation by joint identification-verification," in *Proceedings of the 27th International Conference on Neural Information Processing Systems*, vol. 2, 2014, pp. 1988–1996.
- [5] Z. Akhtar, M. R. Mouree, and D. Dasgupta, "Utility of deep learning features for facial attributes manipulation detection," in *Proceedings of the 2020 IEEE International Conference on Humanized Computing and Communication with Artificial Intelligence (HCCAI)*. IEEE, 2020, pp. 55–60.
- [6] M. Al-Qizwini, I. Barjasteh, H. Al-Qassab, and H. Radha, "Deep learning algorithm for autonomous driving using GoogLeNet," in *Proceedings of the IEEE Intelligent Vehicles Symposium*. IEEE, 2017, pp. 89–96.
- [7] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. J. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," in *Proceedings of the 2nd International Conference on Learning Representations*, 2014.
- [8] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *Proceedings of the 3rd International Conference on Learning Representations*, 2015.
- [9] A. Kurakin, I. J. Goodfellow, and S. Bengio, "Adversarial examples in the physical world," in *Proceedings of the 5th International Conference on Learning Representations*, 2017.
- [10] S. Moosavi-Dezfooli, A. Fawzi, and P. Frossard, "DeepFool: A simple and accurate method to fool deep neural networks," in *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2016, pp. 2574–2582.
- [11] N. Carlini and D. A. Wagner, "Towards evaluating the robustness of neural networks," in *Proceedings of the 2017 IEEE Symposium on Security and Privacy*. IEEE, 2017, pp. 39–57.
- [12] W. Brendel, J. Rauber, and M. Bethge, "Decision-based adversarial attacks: Reliable attacks against black-box machine learning models," in *Proceedings of the 6th International Conference on Learning Representations*, 2018.
- [13] J. Chen, M. I. Jordan, and M. J. Wainwright, "HopSkipJumpAttack: A query-efficient decision-based attack," in *Proceedings of the 2020 IEEE Symposium on Security and Privacy*. IEEE, 2020, pp. 1277–1294.
- [14] Y. Dong, F. Liao, T. Pang, H. Su, J. Zhu, X. Hu, and J. Li, "Boosting adversarial attacks with momentum," in *Proceedings of the 2018 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2018, pp. 9185–9193.
- [15] C. Xie, Z. Zhang, Y. Zhou, S. Bai, J. Wang, Z. Ren, and A. L. Yuille, "Improving transferability of adversarial examples with input diversity," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 2019, pp. 2730–2739.
- [16] D. Wu, Y. Wang, S. Xia, J. Bailey, and X. Ma, "Skip connections matter: On the transferability of adversarial examples generated with resnets," in *Proceedings of the 8th International Conference on Learning Representations*, 2020.
- [17] P. Chen, H. Zhang, Y. Sharma, J. Yi, and C. Hsieh, "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*. ACM, 2017, pp. 15–26.
- [18] C. Tu, P. Ting, P. Chen, S. Liu, H. Zhang, J. Yi, C. Hsieh, and S. Cheng, "Autozoom: Autoencoder-based zeroth order optimization method for attacking black-box neural networks," in *Proceedings of the AAAI Conference on Artificial Intelligence*. AAAI, 2019, pp. 742–749.
- [19] P. Zhao, Z. Fu, O. Wu, Q. Hu, and J. Wang, "Detecting adversarial examples via key-based network," *arXiv preprint arXiv:1806.00580*, 2018.



Table II: the classification accuracy of four ensemble defense strategies adversarial examples on FashionMNIST and CIFAR-10, when model number is 3. ACE means accuracy of clean examples.

Ensemble	Attack	Setting	FashionMNIST				Setting	CIFAR-10			
			BL	GAL	GPMR	NCEn		BL	GAL	GPMR	NCEn
Ensemble1		ACE	0.9337	0.8925	0.9205	0.9184	ACE	0.8639	0.8070	0.8433	0.8436
	FGSM	$\epsilon=0.1$	0.8459	0.8168	0.8178	<b>0.8656</b>	$\epsilon=0.03$	0.6744	0.6102	0.6118	<b>0.7473</b>
		$\epsilon=0.3$	0.6131	<b>0.6885</b>	0.6847	0.6414	$\epsilon=0.09$	0.5462	0.4421	0.4986	<b>0.6217</b>
	MI-FGSM	$\epsilon=0.1$	0.6310	0.5253	0.5720	<b>0.6882</b>	$\epsilon=0.03$	0.2067	0.0131	0.1530	<b>0.3000</b>
		$\epsilon=0.3$	<b>0.4759</b>	0.3396	0.3867	0.4412	$\epsilon=0.09$	0.0625	0.0004	0.0500	<b>0.1157</b>
	PGD	$\epsilon=0.1$	0.7002	0.6406	0.5982	<b>0.7694</b>	$\epsilon=0.01$	0.5661	0.5708	0.4963	<b>0.6473</b>
		$\epsilon=0.15$	0.6163	0.5434	0.4983	<b>0.7362</b>	$\epsilon=0.02$	0.4307	0.2117	0.3459	<b>0.5305</b>
	BIM	$\epsilon=0.1$	0.7079	0.6674	0.6291	<b>0.7783</b>	$\epsilon=0.01$	0.6119	0.5382	0.5241	<b>0.6740</b>
$\epsilon=0.15$		0.6292	0.5620	0.5264	<b>0.7383</b>	$\epsilon=0.2$	0.4722	0.2469	0.3916	<b>0.5716</b>	
Ensemble2		ACE	0.9326	0.6724	0.9203	0.9208	ACE	0.8688	0.7842	0.8385	0.8571
	FGSM	$\epsilon=0.1$	0.8396	0.9284	0.7995	<b>0.8525</b>	$\epsilon=0.03$	<b>0.6862</b>	0.6884	0.5933	0.6844
		$\epsilon=0.3$	0.6025	0.7754	<b>0.6975</b>	0.6564	$\epsilon=0.09$	<b>0.5606</b>	0.5673	0.4697	0.5570
	MI-FGSM	$\epsilon=0.1$	0.6155	0.7751	0.5605	<b>0.6725</b>	$\epsilon=0.03$	0.2107	0.0175	0.1404	<b>0.2405</b>
		$\epsilon=0.3$	0.4461	0.4531	0.4018	<b>0.4549</b>	$\epsilon=0.09$	0.0709	0.0051	0.0374	<b>0.0851</b>
	PGD	$\epsilon=0.1$	0.6904	0.8334	0.6014	<b>0.7516</b>	$\epsilon=0.01$	0.5719	0.5728	0.4819	<b>0.6013</b>
		$\epsilon=0.15$	0.6210	0.7360	0.4883	<b>0.6917</b>	$\epsilon=0.02$	0.4391	0.3555	0.3378	<b>0.4735</b>
	BIM	$\epsilon=0.1$	0.7060	0.8511	0.6220	<b>0.7671</b>	$\epsilon=0.01$	0.6109	0.5618	0.5197	<b>0.6472</b>
$\epsilon=0.15$		0.6107	0.7473	0.5182	<b>0.6933</b>	$\epsilon=0.2$	0.4762	0.3816	0.3831	<b>0.5061</b>	
Ensemble3		ACE	0.9316	0.5259	0.9198	0.9212	ACE	0.8665	0.7589	0.8387	0.8408
	FGSM	$\epsilon=0.1$	0.8535	0.8836	0.8259	<b>0.8782</b>	$\epsilon=0.03$	0.6873	0.5577	0.6076	<b>0.7614</b>
		$\epsilon=0.3$	0.6327	0.7649	<b>0.6797</b>	0.6651	$\epsilon=0.09$	0.5644	0.4507	0.4864	<b>0.6321</b>
	MI-FGSM	$\epsilon=0.1$	0.6456	0.6280	0.5842	<b>0.7028</b>	$\epsilon=0.03$	0.2148	0.0052	0.1423	<b>0.3159</b>
		$\epsilon=0.3$	0.4913	0.4772	0.4016	<b>0.5593</b>	$\epsilon=0.09$	0.0679	0.0006	0.0337	<b>0.1202</b>
	PGD	$\epsilon=0.1$	0.7075	0.7837	0.6133	<b>0.7871</b>	$\epsilon=0.01$	0.5764	0.3464	0.4914	<b>0.6678</b>
		$\epsilon=0.15$	0.6412	0.6773	0.5310	<b>0.7147</b>	$\epsilon=0.02$	0.4482	0.2582	0.3369	<b>0.5448</b>
	BIM	$\epsilon=0.1$	0.7183	0.7961	0.6428	<b>0.7864</b>	$\epsilon=0.01$	0.6168	0.3390	0.5218	<b>0.6912</b>
$\epsilon=0.15$		0.6384	0.7003	0.5472	<b>0.7206</b>	$\epsilon=0.2$	0.4859	0.2205	0.3851	<b>0.5936</b>	
Ensemble4		ACE	0.9309	0.7503	0.9205	0.8428	ACE	0.8695	0.6851	0.8426	0.8492
	FGSM	$\epsilon=0.1$	0.8505	0.9338	0.8116	<b>0.8992</b>	$\epsilon=0.03$	0.6813	0.6293	0.6011	<b>0.7081</b>
		$\epsilon=0.3$	0.6319	0.8501	0.6793	<b>0.7257</b>	$\epsilon=0.09$	0.5587	0.4936	0.4831	<b>0.5827</b>
	MI-FGSM	$\epsilon=0.1$	0.6360	0.7722	0.6015	<b>0.7734</b>	$\epsilon=0.03$	0.2118	0.1224	0.1482	<b>0.2731</b>
		$\epsilon=0.3$	0.4820	0.5316	0.4285	<b>0.5202</b>	$\epsilon=0.09$	0.0672	0.0389	0.0431	<b>0.0959</b>
	PGD	$\epsilon=0.1$	0.6954	0.8527	0.6160	<b>0.8054</b>	$\epsilon=0.01$	0.5738	0.5831	0.4897	<b>0.6295</b>
		$\epsilon=0.15$	0.6235	0.7651	0.5177	<b>0.7469</b>	$\epsilon=0.02$	0.4426	0.3595	0.3332	<b>0.4981</b>
	BIM	$\epsilon=0.1$	0.7073	0.8571	0.6337	<b>0.8156</b>	$\epsilon=0.01$	0.6165	0.5739	0.5185	<b>0.6637</b>
$\epsilon=0.15$		0.6231	0.7654	0.5434	<b>0.7543</b>	$\epsilon=0.2$	0.4833	0.3908	0.3876	<b>0.5449</b>	

- [20] D. Meng and H. Chen, "MagNet: A two-pronged defense against adversarial examples," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2017, pp. 135–147.
- [21] W. Xu, D. Evans, and Y. Qi, "Feature squeezing: Detecting adversarial examples in deep neural networks," in *Proceedings of the 25th Annual Network and Distributed System Security Symposium*. The Internet Society, 2018.
- [22] A. Sinha, Z. Chen, V. Badrinarayanan, and A. Rabinovich, "Gradient adversarial training of neural networks," *arXiv preprint arXiv:1806.08028*, 2018.
- [23] A. Kurakin, I. J. Goodfellow, and S. Bengio, "Adversarial machine learning at scale," in *Proceedings of the 5th International Conference on Learning Representations*, 2017.
- [24] A. Bagnall, R. Bunescu, and G. Stewart, "Training ensembles to detect adversarial examples," *arXiv preprint arXiv:1712.04006*, 2017.
- [25] S. Kariyappa and M. K. Qureshi, "Improving adversarial robustness of ensembles with diversity training," *arXiv preprint arXiv:1901.09981*, 2019.
- [26] T. Pang, K. Xu, C. Du, N. Chen, and J. Zhu, "Improving adversarial robustness via promoting ensemble diversity," in *Proceedings of the 36th International Conference on Machine Learning*, vol. 97. PMLR, 2019, pp. 4970–4979.
- [27] A. Dabouei, S. Soleymani, F. Taherkhani, J. M. Dawson, and N. M. Nasrabadi, "Exploiting joint robustness to adversarial perturbations," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 2020, pp. 1119–1128.
- [28] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," in *Proceedings of the 6th International Conference on Learning Representations*, 2018.
- [29] S. Moosavi-Dezfooli, A. Fawzi, O. Fawzi, and P. Frossard, "Universal adversarial perturbations," in *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2017, pp. 86–94.
- [30] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. C. Courville, and Y. Bengio, "Generative adversarial nets," in *Proceedings of the Advances in Neural Information Processing Systems*, 2014, pp. 2672–2680.
- [31] N. Papernot, P. D. McDaniel, X. Wu, S. Jha, and A. Swami, "Distillation as a defense to adversarial perturbations against deep neural networks," in *Proceedings of the IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2016, pp. 582–597.
- [32] P. Samangouei, M. Kabkab, and R. Chellappa, "Defense-GAN: Protecting classifiers against adversarial attacks using generative models," in *Proceedings of the 6th International Conference on Learning Representations*, 2018.
- [33] G. K. Dziugaite, Z. Ghahramani, and D. M. Roy, "A study of the effect of jpg compression on adversarial images," *arXiv preprint arXiv:1608.00853*, 2016.
- [34] X. Yuan, P. He, Q. Zhu, and X. Li, "Adversarial examples: Attacks and defenses for deep learning," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, no. 9, pp. 2805–2824, 2019.

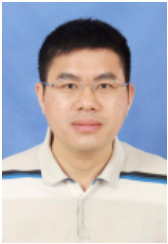
Table III: the classification accuracy of four ensemble defense strategies adversarial examples on FashionMNIST and CIFAR-10, when model number is 4. ACE means accuracy of clean examples.

Ensemble	Attack	Setting	FashionMNIST				Setting	CIFAR-10			
			BL	GAL	GPMR	NCEn		BL	GAL	GPMR	NCEn
Ensemble1		ACE	0.9340	0.9049	0.9235	0.9209	ACE	0.8659	0.8055	0.8527	0.8494
	FGSM	$\epsilon=0.1$	0.8455	0.8636	0.8351	<b>0.8708</b>	$\epsilon=0.03$	0.6642	0.6788	0.6010	<b>0.7356</b>
		$\epsilon=0.3$	0.6267	<b>0.7506</b>	0.7090	0.6523	$\epsilon=0.09$	0.5445	0.5333	0.4768	<b>0.6074</b>
	MI-FGSM	$\epsilon=0.1$	0.6358	0.6118	0.6021	<b>0.6940</b>	$\epsilon=0.03$	0.1984	0.0977	0.1470	<b>0.2727</b>
		$\epsilon=0.3$	<b>0.4867</b>	0.4050	0.4442	0.4364	$\epsilon=0.09$	0.0605	0.0022	0.0418	<b>0.0996</b>
	PGD	$\epsilon=0.1$	0.6784	0.6670	0.5916	<b>0.7620</b>	$\epsilon=0.01$	0.5505	0.6216	0.5013	<b>0.6266</b>
$\epsilon=0.15$		0.5962	0.5414	0.5046	<b>0.6968</b>	$\epsilon=0.02$	0.4072	0.2813	0.3342	<b>0.4901</b>	
BIM	$\epsilon=0.1$	0.6928	0.6894	0.6321	<b>0.7690</b>	$\epsilon=0.01$	0.6003	0.6300	0.5323	<b>0.6612</b>	
	$\epsilon=0.15$	0.6080	0.5638	0.5352	<b>0.7117</b>	$\epsilon=0.2$	0.4434	0.3492	0.3810	<b>0.5376</b>	
Ensemble2		ACE	0.9338	0.8913	0.9228	0.9043	ACE	0.8678	0.7942	0.8575	0.8496
	FGSM	$\epsilon=0.1$	0.8421	<b>0.8521</b>	0.8333	0.8454	$\epsilon=0.03$	0.6665	0.7076	0.6020	<b>0.7162</b>
		$\epsilon=0.3$	0.6081	<b>0.7198</b>	0.7190	0.6403	$\epsilon=0.09$	0.5436	0.5732	0.4772	<b>0.5836</b>
	MI-FGSM	$\epsilon=0.1$	0.6243	0.5796	0.6111	<b>0.6806</b>	$\epsilon=0.03$	0.2014	0.1126	0.1498	<b>0.2704</b>
		$\epsilon=0.3$	0.4499	0.3944	<b>0.4642</b>	0.4490	$\epsilon=0.09$	0.0667	0.0045	0.0415	<b>0.0952</b>
	PGD	$\epsilon=0.1$	0.6835	0.6753	0.5988	<b>0.7421</b>	$\epsilon=0.01$	0.5706	0.5849	0.4997	<b>0.6394</b>
$\epsilon=0.15$		0.6019	0.5270	0.5123	<b>0.6836</b>	$\epsilon=0.02$	0.4136	0.3283	0.3369	<b>0.4871</b>	
BIM	$\epsilon=0.1$	0.6855	0.6977	0.6324	<b>0.7451</b>	$\epsilon=0.01$	0.6114	0.5931	0.5320	<b>0.6673</b>	
	$\epsilon=0.15$	0.5999	0.5580	0.5409	<b>0.6861</b>	$\epsilon=0.2$	0.4566	0.3857	0.3884	<b>0.5387</b>	
Ensemble3		ACE	0.9320	0.7638	0.9255	0.8980	ACE	0.8718	0.7911	0.8537	0.8504
	FGSM	$\epsilon=0.1$	0.8505	0.8634	0.8294	<b>0.8678</b>	$\epsilon=0.03$	0.6670	0.7181	0.6098	<b>0.7683</b>
		$\epsilon=0.3$	0.6493	0.7052	<b>0.6975</b>	0.6347	$\epsilon=0.09$	0.5462	0.6059	0.4870	<b>0.6498</b>
	MI-FGSM	$\epsilon=0.1$	0.6433	0.5943	0.6128	<b>0.7239</b>	$\epsilon=0.03$	0.2068	0.1237	0.1451	<b>0.3226</b>
		$\epsilon=0.3$	0.5037	0.4409	0.4534	<b>0.5465</b>	$\epsilon=0.09$	0.0571	0.0297	0.0366	<b>0.1089</b>
	PGD	$\epsilon=0.1$	0.6979	0.7526	0.6024	<b>0.7899</b>	$\epsilon=0.01$	0.5724	0.5496	0.4979	<b>0.6456</b>
$\epsilon=0.15$		0.6113	0.6671	0.5121	<b>0.7097</b>	$\epsilon=0.02$	0.4183	0.3541	0.3409	<b>0.4995</b>	
BIM	$\epsilon=0.1$	0.7146	0.7717	0.6339	<b>0.7884</b>	$\epsilon=0.01$	0.6097	0.5729	0.5314	<b>0.6880</b>	
	$\epsilon=0.15$	0.6305	0.6804	0.5479	<b>0.7119</b>	$\epsilon=0.2$	0.4652	0.3953	0.3863	<b>0.5569</b>	
Ensemble4		ACE	0.9327	0.7846	0.9224	0.9226	ACE	0.8706	0.7961	0.8532	0.8668
	FGSM	$\epsilon=0.1$	0.8478	0.8817	0.8299	<b>0.8715</b>	$\epsilon=0.03$	0.6577	0.7283	0.6037	<b>0.7530</b>
		$\epsilon=0.3$	0.6155	0.7931	<b>0.6823</b>	0.6622	$\epsilon=0.09$	0.5342	0.6036	0.4731	<b>0.6481</b>
	MI-FGSM	$\epsilon=0.1$	0.6264	0.6707	0.6013	<b>0.7145</b>	$\epsilon=0.03$	0.1952	0.0771	0.1468	<b>0.2812</b>
		$\epsilon=0.3$	<b>0.4726</b>	0.5034	0.4340	0.4640	$\epsilon=0.09$	0.0593	0.0113	0.0406	<b>0.1016</b>
	PGD	$\epsilon=0.1$	0.6812	0.7178	0.5944	<b>0.7834</b>	$\epsilon=0.01$	0.5607	0.5288	0.5004	<b>0.6435</b>
$\epsilon=0.15$		0.6043	0.6178	0.5079	<b>0.7241</b>	$\epsilon=0.02$	0.4121	0.2905	0.3363	<b>0.5034</b>	
BIM	$\epsilon=0.1$	0.6959	0.7378	0.6201	<b>0.7818</b>	$\epsilon=0.01$	0.6046	0.5415	0.5303	<b>0.6801</b>	
	$\epsilon=0.15$	0.6096	0.6440	0.5388	<b>0.7194</b>	$\epsilon=0.2$	0.4501	0.3220	0.3809	<b>0.5532</b>	

- [35] Z. Zheng and P. Hong, “Robust detection of adversarial attacks by modeling the intrinsic properties of deep neural networks,” in *Proceedings of the Advances in Neural Information Processing Systems*, 2018, pp. 7924–7933.
- [36] S. Sankaranarayanan, A. Jain, R. Chellappa, and S. N. Lim, “Regularizing deep networks using efficient layerwise adversarial training,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, AAAI, 2018, pp. 4008–4015.
- [37] M. Abbasi and C. Gagné, “Robustness to adversarial examples through an ensemble of specialists,” in *Proceedings of the 5th International Conference on Learning Representations*, 2017.
- [38] T. Strauss, M. Hanselmann, A. Junginger, and H. Ulmer, “Ensemble methods as a defense to adversarial perturbations against deep neural networks,” *arXiv preprint arXiv:1709.03423*, 2017.
- [39] F. Tramèr, A. Kurakin, N. Papernot, I. J. Goodfellow, D. Boneh, and P. D. McDaniel, “Ensemble adversarial training: Attacks and defenses,” in *Proceedings of the 6th International Conference on Learning Representations*, 2018.
- [40] Y. Liu and X. Yao, “Simultaneous training of negatively correlated neural networks in an ensemble,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 29, no. 6, pp. 716–725, 1999.
- [41] —, “Ensemble learning via negative correlation,” *Neural networks*, vol. 12, no. 10, pp. 1399–1404, 1999.
- [42] Y. Liu, X. Yao, and T. Higuchi, “Evolutionary ensembles with negative correlation learning,” *IEEE Transactions on Evolutionary Computation*, vol. 4, no. 4, pp. 380–387, 2000.
- [43] Z. S. H. Chan and N. K. Kasabov, “A preliminary study on negative correlation learning via correlation-corrected data (NCCD),” *Neural Processing Letters*, vol. 21, no. 3, pp. 207–214, 2005.
- [44] S. Wang, H. Chen, and X. Yao, “Negative correlation learning for classification ensembles,” in *Proceedings of the International Joint Conference on Neural Networks*. IEEE, 2010, pp. 1–8.
- [45] J. R. Quinlan, “Bagging, boosting, and C4.5,” in *Proceedings of the Thirteenth National Conference on Artificial Intelligence and Eighth Innovative Applications of Artificial Intelligence Conference*. AAAI, 1996, pp. 725–730.
- [46] D. W. Opitz and R. Maclin, “Popular ensemble methods: An empirical study,” *Journal of Artificial Intelligence Research*, vol. 11, pp. 169–198, 1999.
- [47] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proceedings of the 3rd International Conference on Learning Representations*, 2015.

Table IV: the classification accuracy of four ensemble defense strategies adversarial examples on FashionMNIST and CIFAR-10, when model number is 5. ACE means accuracy of clean examples.

Ensemble	Attack	Setting	FashionMNIST				Setting	CIFAR-10			
			BL	GAL	GPMP	NCEn		BL	GAL	GPMP	NCEn
Ensemble1		ACE	0.9315	0.8998	0.9247	0.9235	ACE	0.8694	0.8355	0.8573	0.8488
	FGSM	$\epsilon=0.1$	0.8446	0.8546	0.8330	<b>0.8619</b>	$\epsilon=0.03$	0.6473	<b>0.7655</b>	0.6100	0.7248
		$\epsilon=0.3$	0.6265	<b>0.7482</b>	0.7073	0.6248	$\epsilon=0.09$	0.5258	<b>0.6266</b>	0.4875	0.6054
	MI-FGSM	$\epsilon=0.1$	0.6333	0.6572	0.6038	<b>0.6832</b>	$\epsilon=0.03$	0.1808	0.0418	0.1564	<b>0.2611</b>
		$\epsilon=0.3$	<b>0.4813</b>	0.3860	0.4543	0.4092	$\epsilon=0.09$	0.0512	0.0021	0.0422	<b>0.0919</b>
	PGD	$\epsilon=0.1$	0.6673	0.6694	0.6053	<b>0.7386</b>	$\epsilon=0.01$	0.5365	0.5368	0.4934	<b>0.6183</b>
		$\epsilon=0.15$	0.5834	0.5660	0.5211	<b>0.6794</b>	$\epsilon=0.02$	0.3818	0.2794	0.3378	<b>0.4687</b>
	BIM	$\epsilon=0.1$	0.6809	0.6984	0.6350	<b>0.7524</b>	$\epsilon=0.01$	0.5845	0.5770	0.5273	<b>0.6528</b>
		$\epsilon=0.15$	0.5978	0.5965	0.5578	<b>0.6868</b>	$\epsilon=0.2$	0.4324	0.3244	0.3909	<b>0.5192</b>
Ensemble2		ACE	0.9344	0.8858	0.9256	0.9216	ACE	0.8687	0.8293	0.8627	0.8599
	FGSM	$\epsilon=0.1$	0.8358	0.8418	0.8316	<b>0.8534</b>	$\epsilon=0.03$	0.6573	<b>0.7886</b>	0.6080	0.6582
		$\epsilon=0.3$	0.6210	<b>0.7448</b>	0.6970	0.6392	$\epsilon=0.09$	0.5272	<b>0.6598</b>	0.4838	0.5256
	MI-FGSM	$\epsilon=0.1$	0.6268	0.6510	0.6165	<b>0.6866</b>	$\epsilon=0.03$	0.1871	0.1110	0.1473	<b>0.2061</b>
		$\epsilon=0.3$	0.4834	<b>0.4940</b>	0.4571	0.4326	$\epsilon=0.09$	0.0582	0.0101	0.0426	<b>0.0695</b>
	PGD	$\epsilon=0.1$	0.6702	0.6605	0.6030	<b>0.7501</b>	$\epsilon=0.01$	0.5479	<b>0.6049</b>	0.4913	0.5759
		$\epsilon=0.15$	0.5810	0.5705	0.5138	<b>0.6801</b>	$\epsilon=0.02$	0.4027	0.3867	0.3383	<b>0.4224</b>
	BIM	$\epsilon=0.1$	0.6865	0.6816	0.6366	<b>0.7640</b>	$\epsilon=0.01$	0.5929	<b>0.6386</b>	0.5249	0.6186
		$\epsilon=0.15$	0.5893	0.5908	0.5505	<b>0.6902</b>	$\epsilon=0.2$	0.4393	0.4382	0.3841	<b>0.4603</b>
Ensemble3		ACE	0.9339	0.8082	0.923	0.9228	ACE	0.8744	0.8286	0.8604	0.8532
	FGSM	$\epsilon=0.1$	0.8435	0.8603	0.8319	<b>0.8809</b>	$\epsilon=0.03$	0.6531	<b>0.8204</b>	0.6090	0.7535
		$\epsilon=0.3$	0.6292	<b>0.7216</b>	0.6821	0.6395	$\epsilon=0.09$	0.5370	<b>0.6778</b>	0.4894	0.6480
	MI-FGSM	$\epsilon=0.1$	0.6417	0.6471	0.6159	<b>0.7225</b>	$\epsilon=0.03$	0.1901	0.1776	0.1471	<b>0.3088</b>
		$\epsilon=0.3$	<b>0.5023</b>	0.3957	0.4596	0.4531	$\epsilon=0.09$	0.0564	0.0359	0.0399	<b>0.1159</b>
	PGD	$\epsilon=0.1$	0.6789	0.6978	0.6087	<b>0.7838</b>	$\epsilon=0.01$	0.5489	<b>0.6438</b>	0.4995	0.6431
		$\epsilon=0.15$	0.6085	0.5673	0.5126	<b>0.7207</b>	$\epsilon=0.02$	0.4055	0.4440	0.3410	<b>0.5206</b>
	BIM	$\epsilon=0.1$	0.6941	0.7024	0.6343	<b>0.7929</b>	$\epsilon=0.01$	0.5933	0.6751	0.5253	<b>0.6751</b>
		$\epsilon=0.15$	0.6106	0.5857	0.5510	<b>0.7197</b>	$\epsilon=0.2$	0.4486	0.4930	0.3916	<b>0.5672</b>
Ensemble4		ACE	0.9332	0.9015	0.9244	0.9240	ACE	0.8705	0.8370	0.8616	0.8592
	FGSM	$\epsilon=0.1$	0.8423	0.8293	0.8370	<b>0.8631</b>	$\epsilon=0.03$	0.6464	<b>0.8016</b>	0.6001	0.6646
		$\epsilon=0.3$	0.6255	<b>0.7114</b>	0.6966	0.6582	$\epsilon=0.09$	0.5236	<b>0.7076</b>	0.4821	0.5471
	MI-FGSM	$\epsilon=0.1$	0.6279	0.5801	0.6147	<b>0.6889</b>	$\epsilon=0.03$	0.1754	0.1144	0.1555	<b>0.2174</b>
		$\epsilon=0.3$	0.4813	0.3207	0.4709	<b>0.5001</b>	$\epsilon=0.09$	0.0454	0.0186	0.0470	<b>0.0690</b>
	PGD	$\epsilon=0.1$	0.6677	0.6207	0.6049	<b>0.7436</b>	$\epsilon=0.01$	0.5408	<b>0.5887</b>	0.4932	0.5861
		$\epsilon=0.15$	0.5903	0.4892	0.5209	<b>0.6814</b>	$\epsilon=0.02$	0.3885	0.3778	0.3398	<b>0.4306</b>
	BIM	$\epsilon=0.1$	0.6800	0.6495	0.6380	<b>0.7486</b>	$\epsilon=0.01$	0.5808	<b>0.6271</b>	0.5218	0.6228
		$\epsilon=0.15$	0.5970	0.5196	0.5553	<b>0.6879</b>	$\epsilon=0.2$	0.4279	0.4344	0.3844	<b>0.4696</b>

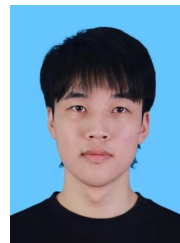


**Wenjian Luo** (SM'15, IEEE) received the BS and PhD degrees from Department of Computer Science and Technology, University of Science and Technology of China, Hefei, China, in 1998 and 2003. He is presently a professor of School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen, China. His current research interests include computational intelligence and applications, network security and data privacy, machine learning and data mining. He is a distinguished member of CCF and a senior member of IEEE, ACM and CAAI. He currently serves as an associate editor or editorial board member for several journals including *Information Sciences Journal*, *Swarm and Evolutionary Computation Journal*, *Journal of Information Security and Applications*, *Applied Soft Computing Journal* and *Complex & Intelligent Systems Journal*. Currently, he also serves as the chair of the IEEE CIS ECTC Task Force on Artificial Immune Systems. He has been a member of the organizational team of more than ten academic conferences, in various functions, such as program chair, symposium chair and publicity chair.

member of CCF and a senior member of IEEE, ACM and CAAI. He currently serves as an associate editor or editorial board member for several journals including *Information Sciences Journal*, *Swarm and Evolutionary Computation Journal*, *Journal of Information Security and Applications*, *Applied Soft Computing Journal* and *Complex & Intelligent Systems Journal*. Currently, he also serves as the chair of the IEEE CIS ECTC Task Force on Artificial Immune Systems. He has been a member of the organizational team of more than ten academic conferences, in various functions, such as program chair, symposium chair and publicity chair.



**Hongwei Zhang** received the B.S. degree from Harbin Institute of Technology, Harbin, China, in 2020. He is currently working toward a Master's degree at the School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen, China. His research interests include adversarial attack and defense.



**Linghao Kong** received the B.E. degree from Northwest University, Shaanxi, China, in 2020. He is currently working toward a Master's degree at the School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen, China. His research interests include adversarial attack and defense.



**Zhijian Chen** received the B.E. degree from Taiyuan University of Technology, Shanxi, China, in 2017. He is currently working towards a Ph.D. degree at the School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen, China. His research interests include machine learning and adversarial examples.



**Ke Tang** (S'05-M'07-SM'13) received the B.Eng. degree from the Huazhong University of Science and Technology, Wuhan, China, in 2002, and the Ph.D. degree from Nanyang Technological University, Singapore, in 2007. From 2007 to 2017, he was with the School of Computer Science and Technology, University of Science and Technology of China, Hefei, China, first as an Associate Professor from 2007 to 2011 and later a Professor from 2011 to 2017. He is currently a Professor with the Department of Computer Science and Engineering, Southern

University of Science and Technology, Shenzhen, China. He has published more than 70 journal papers and more than 80 conference papers. According to Google Scholar, his publications have received more than 8000 citations and the H-index is 42. His major research interests include evolutionary computation, machine learning, and their applications.