

# Feature Separation and Recalibration for Adversarial Robustness

Woo Jae Kim Yoonki Cho Junsik Jung Sung-Eui Yoon

Korea Advanced Institute of Science and Technology (KAIST)

## Abstract

*Deep neural networks are susceptible to adversarial attacks due to the accumulation of perturbations in the feature level, and numerous works have boosted model robustness by deactivating the non-robust feature activations that cause model mispredictions. However, we claim that these malicious activations still contain discriminative cues and that with recalibration, they can capture additional useful information for correct model predictions. To this end, we propose a novel, easy-to-plugin approach named Feature Separation and Recalibration (FSR) that recalibrates the malicious, non-robust activations for more robust feature maps through Separation and Recalibration. The Separation part disentangles the input feature map into the robust feature with activations that help the model make correct predictions and the non-robust feature with activations that are responsible for model mispredictions upon adversarial attack. The Recalibration part then adjusts the non-robust activations to restore the potentially useful cues for model predictions. Extensive experiments verify the superiority of FSR compared to traditional deactivation techniques and demonstrate that it improves the robustness of existing adversarial training methods by up to 8.57% with small computational overhead. Codes are available at <https://github.com/wkim97/FSR>.*

## 1. Introduction

Despite the advancements of deep neural networks (DNNs) in computer vision tasks [2, 10, 20, 39], they are vulnerable to adversarial examples [19, 43] that are maliciously crafted to subvert the decisions of these models by adding imperceptible noise to natural images. Adversarial examples are also known to be successful in real-world cases, including autonomous driving [17] and biometrics [27, 41], and to be effective even when target models are unknown to the attacker [26, 31, 43]. Thus, it has become crucial to devise effective defense strategies against this insecurity.

To this end, numerous defense techniques have been proposed, including defensive distillation [38], input denoising [30], and attack detection [36, 54]. Among these meth-

ods, adversarial training [19, 33], which robustifies a model by training it on a set of worst-case adversarial examples, has been considered to be the most successful and popular.

Even with adversarial training, however, small adversarial perturbations on the pixel-level accumulate to a much larger degree in the intermediate feature space and ruin the final output of the model [50]. To solve this problem, recent advanced methods disentangled and deactivated the non-robust feature activations that cause model mispredictions. Xie *et al.* [50] applied classical denoising techniques to deactivate disrupted activations, and Bai *et al.* [4] and Yan *et al.* [55] deactivated channels that are irrelevant to correct model decisions. These approaches, however, inevitably neglect discriminative cues that potentially lie in these non-robust activations. Ilyas *et al.* [22] have shown that a model can learn discriminative information from non-robust features in the input space. Based on this finding, we argue that there exist potential discriminative cues in the non-robust activations, and deactivating them could lead to loss of these useful information that can provide the model with better guidance for making correct predictions.

For the first time, we argue that with appropriate adjustment, the non-robust activations that lead to model mispredictions could recapture discriminative cues for correct model decisions. To this end, we propose a novel Feature Separation and Recalibration (FSR) module that aims to improve the *feature robustness*. We first separate the intermediate feature map of a model into the malicious *non-robust* activations that are responsible for model mispredictions and the *robust* activations that still provide useful cues for correct model predictions even under adversarial attacks. Exploiting only the robust feature just like the existing methods [4, 50, 55], however, could lead to loss of potentially useful cues in the non-robust feature. Thus, we recalibrate the non-robust activations to capture cues that provide additional useful guidance for correct model decisions. These additional cues can better guide the model to make correct predictions and thus boost its robustness.

Fig. 1 visualizes the attention maps [40] on the features of natural images by a naturally trained model ( $f_{nat}$ ) and the robust ( $f^+$ ), non-robust ( $f^-$ ), and recalibrated features ( $\tilde{f}^-$ ) on adversarial examples ( $x'$ ) obtained from an adver-




Figure 1. Visualization of attention maps [40] on the features of natural images on a naturally trained model ( $f_{nat}$ ) and the robust ( $f^+$ ), non-robust ( $f^-$ ), and recalibrated features ( $\tilde{f}^-$ ) of adversarial examples ( $x'$ ) obtained from an adversarial training [33] model equipped with our FSR module on CIFAR-10 dataset. The robust feature captures discriminative cues regarding the ground truth class, while the non-robust feature captures irrelevant cues. To further boost feature robustness, we recalibrate the non-robust feature ( $f^- \rightarrow \tilde{f}^-$ ) and capture additional useful cues for model predictions. Adversarial images  $x'$  are upscaled for clearer visual representations.

arial training [33] model equipped with our FSR module. Given an adversarial example, while the non-robust feature ( $f^-$ ) captures cues irrelevant to the ground truth class, the robust feature ( $f^+$ ) captures discriminative cues (e.g., horse’s leg). Our FSR module recalibrates the non-robust activations ( $f^- \rightarrow \tilde{f}^-$ ), which are otherwise neglected by the existing methods, and restores additional useful cues not captured by the robust activations (e.g., horse’s body). With these additional cues, FSR further boosts the model’s ability to make correct decisions on adversarial examples.

Thanks to its simplicity, our FSR module can be easily plugged into any layer of a CNN model and is trained with the entire model in an end-to-end manner. We extensively evaluate the robustness of our FSR module on benchmark datasets against various white-box and black-box attacks and demonstrate that our approach improves the robustness of different variants of adversarial training (Sec. 4.2) with small computational overhead (Sec. 4.4). We also show that our approach of recalibrating non-robust activations is superior to existing techniques [4, 50, 55] that simply deactivate them (Sec. 4.2). Finally, through ablation studies, we demonstrate that our Separation stage can effectively disentangle feature activations based on their effects on model decision and that our Recalibration stage successfully recaptures useful cues for model predictions (Sec. 4.3).

In summary, our contributions are as follow:

- In contrast to recent methods that deactivate distorted feature activations, we present a novel point of view

that these activations can be recalibrated to capture useful cues for correct model decisions.

- We introduce an easy-to-plugin *Feature Separation and Recalibration (FSR)* module, which separates non-robust activations from feature maps and recalibrates these feature units for additional useful cues.
- Experimental results demonstrate the effectiveness of our FSR module on various white- and black-box attacks with small computational overhead and verify our motivation that recalibration restores discriminative cues in non-robust activations.

## 2. Related Works

### 2.1. Adversarial Training as Adversarial Defense

Adversarial training guides a model to be robust against adversarial attacks by training it with adversarially generated data and has been widely considered as one of the most effective defense strategies. It solves the following minimax optimization problem:

$$\min_{\theta} \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[ \max_{\delta} \mathcal{L}_{cls}(F_{\theta}(x + \delta), y) \right], \quad (1)$$

where  $F_{\theta}$  is a model parameterized by  $\theta$ ,  $x$  is a natural image with label  $y$  from dataset  $\mathcal{D}$ ,  $\delta$  is a perturbation bounded within the  $\ell_p$ -norm of magnitude  $\epsilon$  such that  $\|\delta\|_p \leq \epsilon$ , and  $\mathcal{L}_{cls}(\cdot, \cdot)$  represents the classification loss. The inner maximization aims to find the strongest possible perturbation  $\delta$  that maximizes the classification loss, and the outer minimization learns the model to minimize the loss with respect to the worst-case adversarial examples. To optimize the inner maximization, Goodfellow *et al.* [19] used the Fast Gradient Sign Method (FGSM), and Madry *et al.* [33] used the Projected Gradient Descent (PGD) attack.

Many variants of adversarial training have also been studied in recent years. ALP [25] reduced the distance between the logits from a natural image and its adversarial counterpart. TRADES [57] decomposed the prediction error on adversarial examples into the natural error and the boundary error to improve both robustness and accuracy. MART [47] additionally considered misclassified examples during training. Inspired by curriculum learning [5], CAT [7] and FAT [58] trained models with increasingly stronger adversarial examples to improve generalization. SEAT [45] proposed a self-ensemble method that combines the weights of different models through the training process, and S<sup>2</sup>O [24] applied the second-order statistics to the model weights to improve adversarial training robustness. Thanks to its simplicity, our method can be easily plugged into any of these adversarial training methods to further improve their robustness.

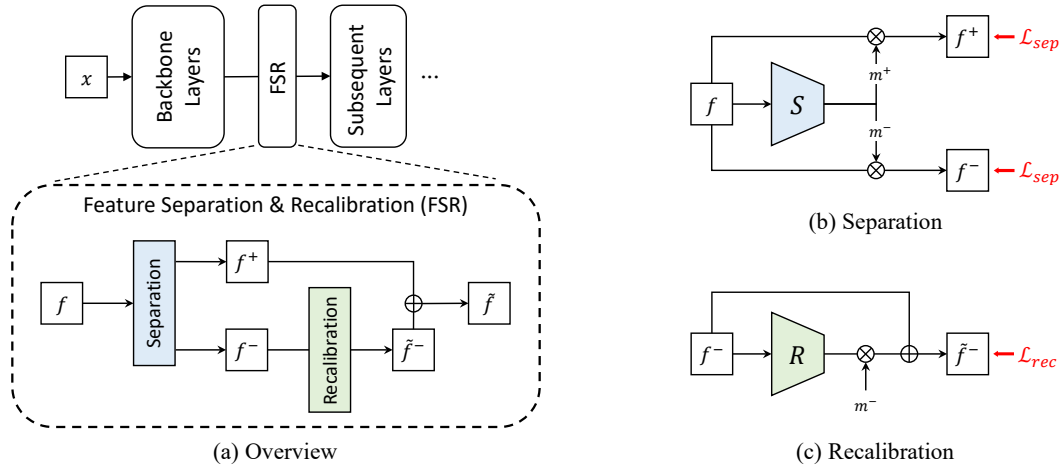


Figure 2. (a) Overview of the Feature Separation and Recalibration module. During the (b) Separation stage, we disentangle the input feature  $f$  into the robust feature  $f^+$  and the non-robust feature  $f^-$  by applying the learnable positive mask  $m^+$  and the negative mask  $m^-$ , respectively. Then, during the (c) Recalibration stage, we recalibrate the activations of  $f^-$  into  $\tilde{f}^-$  to restore the useful cues for correct model predictions. Finally, we combine the recalibrated feature and the robust feature to obtain the output feature  $\tilde{f}$  and pass it down to subsequent layers of the model.

## 2.2. Adversarial Defense on Feature Space

In a parallel line of research, it has been found that some prior models learn non-robust features from the dataset [22] and that input perturbations of adversarial examples are often accumulated through intermediate layers to misguide the final prediction [50]. To solve these problems, several works tried to learn robust feature representations by modifying the network structure or applying regularizations. Galloway *et al.* [18], Benz *et al.* [6], and Wang *et al.* [46] studied the adversarial vulnerability from the perspective of batch normalization. Dhillon *et al.* [14] applied pruning to a random set of activations, especially those with small magnitudes, and Madaan *et al.* [32] pruned out activations that are vulnerable to adversarial attacks. Mustafa *et al.* [34] proposed a class-wise feature disentanglement and pushed the centers of each class from each other to learn more discriminative feature representations.

There also have been attempts to reduce abnormalities in the feature maps by explicitly manipulating the feature activations. Xiao *et al.* [49] proposed  $k$ -Winner-Takes-All activation to deactivate all feature units except for  $k$  units with the largest magnitudes. Xu *et al.* [53] interpreted the effects of adversarial perturbations on the pixel, image, and network levels and masked out feature units sensitive to perturbations. Zoran *et al.* [59] applied attention mechanism to emphasize important regions on the feature map. Xie *et al.* [50] proposed Feature Denoising (FD) that applies classical denoising techniques to deactivate abnormal activations. Bai *et al.* [4] and Yan *et al.* [55] studied the effects of perturbations on the feature activation from the channel perspective and proposed Channel Activation Suppression (CAS) and Channel-wise Importance-based Feature Selec-

tion (CIFS), respectively, to deactivate the activation of non-robust channels.

In contrast to these methods, we propose a recalibration strategy. Existing deactivation strategies simply discard non-robust feature activations responsible for model mistakes. Taking a step further, we adjust such activations to instead *recapture* potentially discriminative cues and thus boost model robustness.

## 3. Methods

While the distorted feature activations upon adversarial attack are known to be responsible for model mispredictions, we argue that with adjustment, we can recapture useful cues for model predictions. To fully utilize such potentially useful cues, we present a novel Feature Separation and Recalibration (FSR) module (Fig. 2) that can restore these cues through a separation-and-recalibration scheme. During the Separation stage, we disentangle the feature map into the *robust* and *non-robust* features by masking out the non-robust and robust activations, respectively. Then, during the Recalibration stage, we recalibrate the activations of the non-robust features such that they provide useful information for correct model prediction. Thanks to its simplicity, as shown in Fig. 2, FSR module can be inserted to any layer of a model to improve its robustness. We elaborate on the details of our FSR module in the following subsections.

### 3.1. Feature Separation

As compared to recent approaches [4, 50, 55] that robustify feature maps by deactivating non-robust activations, we aim to recalibrate them to restore potentially useful cues for model predictions. To this end, during the Separation stage,

we extract the non-robust activations to be recalibrated from the input intermediate feature map. To determine which activation is robust or non-robust, we introduce the Separation Net  $S$ , which learns the robustness of each feature unit. Given the intermediate feature map  $f \in \mathbb{R}^{C \times H \times W}$  as input, where  $C$ ,  $H$ , and  $W$  represent the channel, height, and width dimensions of  $f$ , respectively, the Separation Net outputs a robustness map  $r \in \mathbb{R}^{C \times H \times W}$  that represents a robustness scores of the corresponding units of  $f$ , where a higher score means a more robust feature activation.

In order to extract the non-robust feature, we disentangle the feature map into the robust feature and the non-robust feature in an element-wise manner based on the robustness score. One way to achieve this goal is to apply to the feature map a binary mask  $b \in \{0, 1\}^{C \times H \times W}$  generated based on the robustness score. However, such discrete sampling is non-differentiable [23] and discontinuous [52], which could cause gradient masking that would give a false sense of robustness [3, 37].

To avoid such problem, we approximate a binary mask  $b$  with a differentiable soft mask  $m \in [0, 1]^{C \times H \times W}$  using Gumbel softmax [23] such that:

$$m = \frac{e^{((\log(\sigma(r)) + g_1)/\tau)}}{e^{((\log(\sigma(r)) + g_1)/\tau)} + e^{((\log(1 - \sigma(r)) + g_2)/\tau)}}, \quad (2)$$

where  $r$  is the robustness map, and  $\sigma$  is a sigmoid function used to normalize the robustness map.  $g_1$  and  $g_2$  represent the samples drawn from the Gumbel distribution such that  $g = -\log(-\log(u))$ , where  $u \sim \text{Uniform}(0, 1)$ , and  $\tau$  is a temperature used to control the effects of  $g_1$  and  $g_2$ . Note that during inference, to avoid stochasticity from sampling  $g_1$  and  $g_2$ , we fix them as  $-\log(-\log(u_c))$ , where  $u_c \in \mathbb{R}^{C \times H \times W}$  is the expected value of Uniform distribution. By computing two-class Gumbel softmax between the normalized robustness map and its inverted version, we obtain a mask  $m$  with values close to 1 for high robustness scores and values close to 0 for low robustness scores. Then, to mask out the non-robust activations from the input feature  $f$  and obtain the robust feature  $f^+$ , we compute element-wise product between the feature  $f$  and the positive mask  $m^+ = m$  such that  $f^+ = m^+ \otimes f$ . Similarly, we obtain non-robust feature  $f^- = m^- \otimes f$  by masking out the robust activations with the negative mask  $m^- = 1 - m$ .

Without any guidance, however, the Separation Net may not learn the correct robustness score for each activation. To this end, we design an objective that guides the Separation Net to learn robustness scores specifically based on the influence of feature activations on model making a correct or incorrect prediction. We attach an MLP-based auxiliary layer  $h$  that takes each of the two feature  $f^+$  and  $f^-$  as inputs and outputs prediction scores  $p^+$  and  $p^-$ , respectively.

Then, we compute the separation loss  $\mathcal{L}_{sep}$  as follows:

$$\mathcal{L}_{sep} = - \sum_{i=1}^N (y_i \cdot \log(p_i^+) + y'_i \cdot \log(p_i^-)), \quad (3)$$

where  $N$  is the number of classes,  $y$  is the ground truth label, and  $y'$  is the label corresponding to the wrong class with the highest prediction score from the final model output. By training the auxiliary layer to make correct predictions based on the activations that are preserved from the positive mask  $m^+$ ,  $\mathcal{L}_{sep}$  guides the Separation Net to assign high robustness scores to units that help the auxiliary layer make correct predictions. At the same time, we aim to disentangle the highly disrupted activations that cause the model to make mispredictions upon adversarial attack.  $\mathcal{L}_{sep}$  also guides the Separation Net to assign low robustness scores to non-robust activations that are specifically responsible for the most probable misprediction.

With  $\mathcal{L}_{sep}$ , the Separation Net can effectively separate the robust and the non-robust feature activations. Discarding the non-robust activations is one way to improve feature robustness; however, this approach would ignore potentially useful cues that can be recaptured through recalibration (see Sec. 4.3). In the following subsection, we discuss how we recalibrate the disentangled non-robust activations to capture additional useful cues for improved feature robustness.

### 3.2. Feature Recalibration

Exploiting only the robust feature obtained through the Separation stage just like recent techniques [4, 50, 55] could lead to loss of potentially useful cues for model predictions that could further boost model robustness. Therefore, for the first time, we adjust the non-robust feature activations to capture the additional useful cues during the Recalibration stage. We first introduce the Recalibration Net  $R$  that takes the non-robust feature  $f^-$  as input and outputs recalibrating units that are designed to adjust the activations of  $f^-$  accordingly. To recalibrate the non-robust activations designated by the robustness map, we apply the negative mask  $m^-$  to the recalibrating units. Finally, we compute the recalibrated feature  $\tilde{f}^-$  by adding the result to  $f^-$ , i.e.,  $\tilde{f}^- = f^- + m^- \otimes R(f^-)$ .

The goal of the Recalibration stage is to make the non-robust activations recapture cues that can help the model make correct decisions. To guide the Recalibration Net to achieve this goal, we again attach the auxiliary layer  $h$  after the recalibrated feature  $\tilde{f}^-$  and compute the recalibration loss  $\mathcal{L}_{rec}$  as follows:

$$\mathcal{L}_{rec} = - \sum_{i=1}^N y_i \cdot \log(\tilde{p}_i^-), \quad (4)$$

where  $\tilde{p}^-$  is the output prediction score of the auxiliary layer given  $\tilde{f}^-$  as input. By training the same auxiliary layer

to make correct decisions based on the recalibrated feature, we guide the Recalibration Net to adjust the non-robust activations such that they provide cues relevant to the ground truth class. After the Recalibration stage, we add the robust feature  $f^+$  and the recalibrated non-robust feature  $\tilde{f}^-$  in an element-wise manner to obtain the output feature map  $\tilde{f} = f^+ + \tilde{f}^-$ , which is passed to subsequent layers of the model. Through the Recalibration stage, we can capture additional useful cues from the non-robust activations, which are neglected in previous approaches.

### 3.3. Model Training

The proposed FSR module can be easily inserted to any layer of a model and is trained with the entire model in an end-to-end manner thanks to its simplicity. We can also apply the proposed method with any classification loss  $\mathcal{L}_{cls}$  for different types of adversarial training [33, 47, 57], and the overall objective function is as follows:

$$\mathcal{L} = \mathcal{L}_{cls} + \frac{1}{|L|} \sum_{l \in L} (\lambda_{sep} \cdot \mathcal{L}_{sep}^l + \lambda_{rec} \cdot \mathcal{L}_{rec}^l), \quad (5)$$

where  $L$  represents the set of positions in which the FSR module is inserted, and  $\mathcal{L}_{sep}^l$  and  $\mathcal{L}_{rec}^l$  each represents the separation loss and the recalibration loss applied on the FSR module at  $l$ -th layer. The hyperparameters  $\lambda_{sep}$  and  $\lambda_{rec}$  are used to control the weights of  $\mathcal{L}_{sep}$  and  $\mathcal{L}_{rec}$ , respectively. Addition of this simple FSR module can improve the robustness of adversarial training methods against both white-box and black-box attacks with small computational overhead, as described in the following section.

## 4. Experiments

### 4.1. Experimental Setups

**Evaluation Protocols.** We evaluate our method on CIFAR-10/100 [28], SVHN [35], and Tiny ImageNet [13] datasets using ResNet-18 [20], VGG16 [42], and WideResNet-34-10 [56] as the baseline models. We apply our method to PGD adversarial training (AT) [33] and other variants of adversarial training, *i.e.*, TRADES [57] and MART [47], to verify its wide applicability. For training, we use PGD-10 [33] with perturbation bound  $\epsilon = 8/255$  (step size  $\epsilon/4$  for CIFAR-10/100 and Tiny ImageNet,  $\epsilon/8$  for SVHN) under  $\ell_\infty$ -norm to craft adversarial examples. For evaluation, we use FGSM [19], PGD-20 [33] (step size  $\epsilon/10$ ), PGD-100 [33] (step size  $\epsilon/10$ ), and C&W [9] (PGD optimization for 30 steps with step size  $\epsilon/10$ ) bounded within  $\epsilon = 8/255$  under  $\ell_\infty$ -norm. As suggested by Carlini *et al.* [8], to better compare the robustness of different defense techniques, we also report the average per-example Ensemble robustness of

the model as formulated below [15]:

$$\text{Ensemble} = \frac{1}{N_{test}} \sum_{i=1}^{N_{test}} \min_{a \in \mathcal{A}} \mathbb{1}(F_\theta(a(x)) = y), \quad (6)$$

where  $N_{test}$  is the number of images in the test dataset,  $\mathbb{1}(\cdot)$  is an indicating function,  $F_\theta$  is a target model with parameter  $\theta$ ,  $y$  is the ground truth label, and  $\mathcal{A}$  is a set of adversarial attacks (FGSM, PGD-20, PGD-100, and C&W).

**Implementation Details.** We train all models for 100 epochs using an SGD optimizer (momentum 0.9, weight decay  $5 \times 10^{-4}$ ). We set the initial learning rate to 0.1 for CIFAR-10/100 and Tiny ImageNet and to 0.01 for SVHN, and reduce it by a factor of 10 after 75-th and 90-th epochs. We empirically set  $\lambda_{sep} = 1$ ,  $\lambda_{rec} = 1$ , and  $\tau$  for Gumbel softmax as 0.1. We implement the Separation Net as a series of three blocks each comprised of a convolutional layer, a batch normalization layer, and a ReLU activation, except for the last block which consists of a single convolutional layer. The Recalibration Net also consists of a series of three blocks each comprised of a convolutional layer, a batch normalization layer, and a ReLU activation. We insert our FSR module after `block4` for ResNet-18, `block4` for VGG16, and `block3` for WideResNet-34-10.

### 4.2. Robustness Evaluation

**Defense against White-box Attacks.** To evaluate the ability of our FSR module to improve the model robustness of various adversarial training techniques, we report in Table 1 the effectiveness of applying FSR to three different methods (AT, TRADES, and MART) on ResNet-18. Applying our FSR module consistently improves the robustness of all defense techniques under all of the individual attacks and the Ensemble attack. Similar trends are observed in the SVHN dataset (Table 1) and on VGG16 (Table 2); for example, FSR improves the robustness of TRADES on VGG16 under PGD-20 by 8.57% on the SVHN dataset. By recalibrating the malicious non-robust activations, our method provides the model with additional useful cues for improved robustness. With only a slight increase in the number of computations (see Sec. 4.4), we can improve the robustness of various adversarial training methods regardless of the dataset and the model. Results on WideResNet-34-10, CIFAR-100, Tiny ImageNet are provided in the supplementary materials.

One observation is that our FSR module occasionally drops the accuracy on natural images. This is because our method is designed to disentangle and recalibrate the ‘‘malicious’’ activations that are *intentionally* crafted to fool model predictions. In natural images, since there are no intentionally disrupted feature activations, trying to identify and recalibrate these malicious cues can lead to potential loss of discriminative information and occasional accuracy drop. Nevertheless, the natural accuracy drops only by a

<i>ResNet-18</i>	CIFAR-10						SVHN					
	Method	Natural	FGSM	PGD-20	PGD-100	C&W	Ensemble	Natural	FGSM	PGD-20	PGD-100	C&W
AT	<b>85.02</b>	56.21	48.22	46.37	47.38	45.51	91.21	55.55	40.85	37.54	40.61	37.41
AT + FSR	81.46	<b>58.07</b>	<b>52.47</b>	<b>51.02</b>	<b>49.44</b>	<b>48.34</b>	<b>91.28</b>	<b>60.46</b>	<b>43.94</b>	<b>39.01</b>	<b>43.22</b>	<b>38.81</b>
TRADES	<b>86.31</b>	57.21	50.74	49.44	48.66	47.89	90.99	61.31	47.12	43.55	45.48	42.99
TRADES + FSR	84.49	<b>58.29</b>	<b>52.27</b>	<b>51.28</b>	<b>49.92</b>	<b>49.28</b>	<b>91.39</b>	<b>68.85</b>	<b>51.49</b>	<b>47.50</b>	<b>46.70</b>	<b>46.17</b>
MART	82.73	56.65	50.88	49.15	47.21	45.98	<b>90.50</b>	58.21	43.61	40.43	42.20	40.07
MART + FSR	<b>83.28</b>	<b>59.55</b>	<b>54.80</b>	<b>53.69</b>	<b>48.98</b>	<b>48.36</b>	89.87	<b>61.06</b>	<b>46.51</b>	<b>42.94</b>	<b>43.89</b>	<b>42.40</b>

Table 1. Robustness (accuracy (%)) of adversarial training strategies (AT, TRADES, MART) with (+ FSR) and without our FSR module against diverse white-box attacks on ResNet-18. Better results are marked in **bold**.

<i>VGG16</i>	CIFAR-10						SVHN					
	Method	Natural	FGSM	PGD-20	PGD-100	C&W	Ensemble	Natural	FGSM	PGD-20	PGD-100	C&W
AT	<b>80.56</b>	53.47	47.17	45.58	45.82	43.71	89.59	54.88	40.27	36.90	39.46	36.62
AT + FSR	80.06	<b>54.40</b>	<b>49.82</b>	<b>48.82</b>	<b>47.28</b>	<b>46.24</b>	<b>91.44</b>	<b>65.01</b>	<b>45.99</b>	<b>39.07</b>	<b>43.08</b>	<b>38.15</b>
TRADES	<b>82.44</b>	53.92	47.39	46.20	44.80	44.20	90.48	61.50	45.99	40.00	42.82	39.27
TRADES + FSR	80.78	<b>55.48</b>	<b>49.95</b>	<b>49.03</b>	<b>46.28</b>	<b>45.90</b>	<b>91.89</b>	<b>69.25</b>	<b>54.56</b>	<b>47.81</b>	<b>46.66</b>	<b>44.10</b>
MART	76.11	54.86	51.06	50.16	43.53	43.01	89.95	59.03	42.89	38.73	39.12	37.64
MART + FSR	<b>79.18</b>	<b>56.41</b>	<b>52.69</b>	<b>52.13</b>	<b>44.49</b>	<b>44.20</b>	<b>90.60</b>	<b>62.28</b>	<b>47.17</b>	<b>42.50</b>	<b>43.44</b>	<b>40.73</b>

Table 2. Robustness (accuracy (%)) of adversarial training strategies (AT, TRADES, MART) with (+ FSR) and without our FSR module against diverse white-box attacks on VGG16. Better results are marked in **bold**.

Method	TI-FGSM	DI-FGSM	$\mathcal{N}$ Attack	AutoAttack
AT	59.03	46.56	39.55	44.11
AT + FSR	<b>62.53</b>	<b>50.50</b>	<b>52.65</b>	<b>46.41</b>
TRADES	59.95	49.23	43.45	46.81
TRADES + FSR	<b>61.45</b>	<b>50.54</b>	<b>50.43</b>	<b>48.45</b>
MART	59.73	48.38	44.68	44.27
MART + FSR	<b>62.68</b>	<b>51.42</b>	<b>53.76</b>	<b>46.55</b>

Table 3. Robustness (accuracy (%)) of adversarial training strategies (AT, TRADES, MART) with (+ FSR) and without our FSR module against diverse black-box attacks and AutoAttack on CIFAR-10 using ResNet-18. Better results are marked in **bold**.

small amount or even improves upon the addition of FSR module (*e.g.*, on SVHN), showing that this phenomenon is also dependent on the dataset.

**Defense against Black-box Attacks and AutoAttack.** To show that FSR improves adversarial training methods even under different types of attacks, we evaluate our method against a variety of black-box attacks and AutoAttack [12]. For black-box attacks, we use two transfer-based attacks – TI-FGSM [16] and DI-FGSM [51] – crafted on a naturally trained ResNet-50, and  $\mathcal{N}$ Attack [29], which is a strong query-based attack. Following CAS [4], to evaluate each method against  $\mathcal{N}$ Attack, we sample 1,000 images from CIFAR-10 test set and limit the number of queries to 40,000. AutoAttack, which is an ensemble of two Auto-PGD attacks [12], Fast Adaptive Boundary attack [11], Square attack [1], has been shown to evaluate the robustness of defense techniques more reliably. As shown in Table 3, our method improves the robustness against diverse black-box attacks, especially against the stronger  $\mathcal{N}$ Attack. Our method also improves the robustness against AutoAt-

tack, showing that the improvement upon the addition of FSR module is truly thanks to making the model robust instead of obfuscated gradients [3] or improper evaluation.

**Comparison with Existing Methods.** To verify the effectiveness of recalibrating non-robust activations, we report in Table 4 the comparison of our method to existing feature manipulation methods (FD, CAS, and CIFS). We leave out from evaluation kWTA [49] and SAP [14], which also manipulate feature activations for robustness, because they are known to cause gradient masking [3, 44]. We report the robustness against various white-box attacks, the Ensemble of these attacks, and AutoAttack. Our FSR outperforms the three methods under most white-box attacks. CAS and CIFS excel at defending against C&W because they exploit the weights of auxiliary classifiers or the gradients of their logit outputs to manipulate feature activations and thus enlarge the prediction margins on the feature space [4]. Our method does not adopt such technique and slightly lags behind CAS and CIFS under C&W attack. Nevertheless, our method achieves the highest Ensemble robustness, verifying that overall, our method is the most robust. FSR also outperforms all methods under AutoAttack, showing that it is more reliable even under an ensemble of various white-box and black-box attacks. As opposed to these approaches that deactivate the non-robust feature activations, our method instead recalibrates them to capture additional useful cues that help the model make correct predictions, thus improving the model robustness.

### 4.3. Ablation Studies

We first evaluate whether the individual components of FSR module work as desired by measuring the following:

Method	Natural	FGSM	PGD-20	PGD-100	C&W	Ensemble	AutoAttack
AT	85.02	56.21	48.22	46.37	47.38	45.51	44.11
FD	85.14	56.81	48.54	46.70	47.72	45.82	44.57
CAS	<b>85.78</b>	55.57	50.42	49.91	<b>53.47</b>	46.46	44.23
CIFS	79.87	56.53	49.80	48.17	49.89	47.26	43.94
FSR (Ours)	81.46	<b>58.07</b>	<b>52.47</b>	<b>51.02</b>	49.44	<b>48.34</b>	<b>46.41</b>

Table 4. Comparison of robustness (accuracy (%)) between existing methods and our method. All models are trained using AT with ResNet-18 on CIFAR-10. The best results are marked in **bold**, and more comprehensive Ensemble and AutoAttack are highlighted in grey.

Method	(a) Classification		(b) Weighted $k$ -NN	
	Ensemble	AutoAttack	5-NN	20-NN
$f^+$	47.89	45.82	66.21	61.58
$f^-$	33.11	28.39	54.69	53.89
$\tilde{f}^-$	46.93	44.52	66.34	65.64
$\tilde{f}$ (Ours)	48.34	46.41	70.91	65.88

Table 5. Ablation studies on the robustness of various feature maps obtained throughout our framework on CIFAR-10 using ResNet-18. (a) Robustness (%) of the model upon replacing the final feature map  $\tilde{f}$  with different feature maps. (b) Top-1 accuracy (%) of weighted  $k$ -NN on different feature maps.

(a) classification accuracy of model and (b) weighted  $k$ -NN accuracy. To evaluate (a) the classification accuracy of model, we pass down different features from FSR into the subsequent layers instead of the final feature  $\tilde{f}$ . By doing so, we compare the *model robustness* brought by these different features.

Different from the model robustness, we also explicitly measure the *feature robustness* based on how well each feature captures discriminative cues corresponding to the ground truth class. To do so, we embed each feature among the features of natural images and measure (b) the weighted  $k$ -NN accuracy [21, 48]. For each arbitrary feature  $f$ , we first compute the weight  $w_i$  for each neighbor  $v_i$  as its cosine similarity to  $f$  with the temperature parameter  $\gamma$  such that  $w_i = \exp(\cos(f, v_i)/\gamma)$ . Then, we compute the prediction score  $s_c$  for each class  $c$  by weighting the vote of each of the  $k$ -nearest neighbors  $\mathcal{N}_k$  as  $s_c = \sum_{i \in \mathcal{N}_k} w_i \cdot \mathbb{1}(c_i = c)$ . Through this way, we measure robustness of each feature by determining how close it lies to the unperturbed features.

**Evaluation on Separation.** We show that our Separation stage does appropriately disentangle the non-robust feature from the input feature, which is an essential step in our method before recalibrating non-robust activations. To do so, we evaluate how well the robust feature  $f^+$  and the non-robust feature  $f^-$  capture useful cues for correct model predictions. We measure the classification accuracy of the model against the Ensemble (of FGSM, PGD-20, PGD-100, and C&W) and AutoAttack upon using each feature (*i.e.*,  $\tilde{f} = f^+$  or  $\tilde{f} = f^-$ ), whose results are shown in the left side (a) of Table 5. We can observe that  $f^+$  leads to much

higher robustness against both Ensemble and AutoAttack than  $f^-$ , which drastically decreases the adversarial robustness. This implies that while  $f^+$  captures robust cues that lead to correct model decisions on adversarial examples,  $f^-$  captures non-robust cues that are responsible for mispredictions upon attack, showing that our Separation stage well disentangles the input feature activations as intended.

On the right side (b) of Table 5, we report the 5-NN and 20-NN accuracies on  $f^+$  and  $f^-$ . Both cases show that  $f^+$  results in higher accuracies than  $f^-$ . This is because we disentangle  $f^+$  as activations that help the model make correct predictions, and it thus captures more similar representations as the features of natural images than  $f^-$  does.

**Evaluation on Recalibration.** We also evaluate the ability of Recalibration to adjust the non-robust activations such that they capture cues that help the model make correct predictions. We again test the classification accuracy of our model upon replacing the final feature map  $\tilde{f}$  with the non-robust feature  $f^-$  or the recalibrated non-robust feature  $\tilde{f}^-$ . As shown in the left side (a) of Table 5,  $\tilde{f}^-$  leads to huge improvements in robustness compared to  $f^-$ , showing that our Recalibration stage appropriately adjusts the non-robust activations to capture cues that help the model make correct predictions.

We can also observe that  $f^+$ , which is equivalent to simply suppressing the non-robust activations with our mask  $m^+$ , lags behind using both the robust and recalibrated features, *i.e.*,  $\tilde{f} = f^+ + \tilde{f}^-$  (Ours). This shows that recalibrating the non-robust activations restores additional useful cues for model decisions that are not captured by the robust activations and further improves the adversarial robustness, advocating for the necessity of our Recalibration stage.

Similar trends can also be observed on the  $k$ -NN accuracy as shown in the right side (b) of Table 5. The recalibrated non-robust feature  $\tilde{f}^-$  leads to significantly higher  $k$ -NN accuracy than  $f^-$ , and the accuracy of  $f^+$  lags behind that of  $\tilde{f} = f^+ + \tilde{f}^-$  (Ours), again verifying the effectiveness of our Recalibration stage.

**Effectiveness of Separation and Recalibration.** We verify the necessity of the Separation and the Recalibration stages by comparing the model robustness upon removing each of them from FSR. As shown in Table 6, removing the Separation stage (*i.e.*, recalibrating the entire input feature map)

Method	FGSM	PGD-20	PGD-100	C&W	Ensemble	AutoAttack
AT	56.21	48.22	46.37	47.38	45.51	44.11
+ FSR	<b>58.07</b>	<b>52.47</b>	<b>51.02</b>	<b>49.44</b>	<b>48.34</b>	<b>46.41</b>
w/o Sep	57.51	50.71	48.98	49.32	47.60	45.47
w/o Rec	57.67	50.06	48.54	49.41	47.32	44.96

Table 6. Comparison of robustness (%) of FSR applied on AT upon removing the Separation or the Recalibration stage. Model and dataset used are ResNet-18 and CIFAR-10, respectively. Best results are marked in **bold**.

leads to drop in adversarial robustness. This is because recalibrating the entire feature map that contains both robust and non-robust activations leads to a suboptimal training of the Recalibration Net. More specifically, the Recalibration Net learns to minimize the recalibration loss through the discriminative information that is already captured by the robust activations instead of recalibrating the non-robust activations. Nevertheless, FSR without Separation still improves the robustness of AT, showing that Recalibration effectively captures useful cues for model predictions.

We also compare robustness as we remove the Recalibration stage and pass down only the robust feature into subsequent layers. As shown in Table 6, removing the Recalibration stage also decreases robustness. This demonstrates the necessity of the Recalibration stage and also confirms that the non-robust activations contain additional useful cues which further boost model robustness. Still, FSR without Recalibration improves the robustness compared to vanilla AT. This shows that our Separation stage well disentangles the intermediate feature map based on feature robustness and outputs robust activations that provide useful cues for model predictions.

In the supplementary materials, we report additional ablation studies including the effects of different hyperparameters ( $\lambda_{sep}$ ,  $\lambda_{rec}$ ,  $\tau$ ), effectiveness of Gumbel softmax, addition of FSR on different layers, choice of  $y'$  in Eq. 3, and effects of replacing robustness map  $m$  with other strategies.

#### 4.4. Computational Efficiency

In Table 7, we provide the computational analysis of our method compared to a vanilla model in terms of the number of parameters (# params (M)) and the number of floating point operations (FLOPs (G)). With only a slightly more number of computations, we can improve the robustness of adversarial training and its variants with a significant margin. Additionally, on the CIFAR-10 dataset, one training epoch of PGD-10 adversarial training on the vanilla ResNet-18 takes 114 seconds, while it takes 120 seconds on ResNet-18 with FSR module. One epoch of evaluation on PGD-20 adversarial examples takes 19 seconds for the vanilla ResNet-18, while it takes 20 seconds for ResNet-18 equipped with FSR module. With slight computational overhead, our FSR module improves the robustness of traditional adversarial training models.

Method	<i>VGG16</i>		<i>ResNet-18</i>	
	# Params (M)	FLOPs (G)	# Params (M)	FLOPs (G)
Vanilla	15.25	0.6299	11.17	1.1133
+ FSR	16.52	0.6701	12.43	1.1535

Table 7. Comparison of computational costs (# params and FLOPs) on a vanilla model and a model with our FSR module.

## 5. Discussion

Since the goal of our work is to improve the robustness of adversarial training models against adversarial examples, it lies on the assumption that the input images contain malicious perturbations designed to fool the model. Thus, we deliberately design the Separation stage to disentangle activations that specifically lead the model to mispredictions. However, natural images may not contain such malicious cues, and thus, our FSR module occasionally decreases the natural accuracy by a small amount. A potential direction for future work could be applying curriculum learning [5] to make FSR be better aware of the relationship between feature robustness and attack strength.

Adversarial attack poses a huge threat to the deployment of deep neural networks (DNN) in real-world applications [17, 27, 41]. In this regard, our work contributes positively to the research field by designing an easy-to-plugin module to robustify the DNN models against adversarial attacks. Thanks to its simplicity, we expect that our method could also serve as a basis to design more robust DNN models in diverse and real-world applications.

## 6. Conclusion

In this paper, we have proposed the novel Feature Separation and Recalibration (FSR) module that recalibrates the non-robust activations to restore discriminative cues that help the model make correct predictions under adversarial attack. To achieve this goal, FSR first disentangles the intermediate feature map into the robust activations that capture useful cues for correct model decisions and the non-robust activations that are responsible for incorrect predictions. It then recalibrates the non-robust activations to capture potentially useful cues that could provide additional guidance for more robust predictions on adversarial examples. We have empirically demonstrated the ability of our method to improve the robustness of various models when applied to different adversarial training strategies across diverse datasets. We have also verified the superiority of our method to existing approaches that simply deactivate such non-robust activations.

**Acknowledgement** Prof. Sung-Eui Yoon is a corresponding author. This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No. RS-2023-00208506).



## References

- [1] Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion, and Matthias Hein. Square attack: a query-efficient black-box adversarial attack via random search. In *ECCV*, 2020. 6
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *ICCV*, 2015. 1
- [3] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *ICML*, 2018. 4, 6
- [4] Yang Bai, Yuyuan Zeng, Yong Jiang, Shu-Tao Xia, Xingjun Ma, and Yisen Wang. Improving adversarial robustness via channel-wise activation suppressing. In *ICLR*, 2021. 1, 2, 3, 4, 6
- [5] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *ICML*, 2009. 2, 8
- [6] Philipp Benz, Chaoning Zhang, and In So Kweon. Batch normalization increases adversarial vulnerability and decreases adversarial transferability: A non-robust feature perspective. In *ICCV*, 2021. 3
- [7] Qi-Zhi Cai, Min Du, Chang Liu, and Dawn Song. Curriculum adversarial training. In *IJCAI*, 2018. 2
- [8] Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian Goodfellow, Aleksander Madry, and Alexey Kurakin. On evaluating adversarial robustness. *arXiv preprint arXiv:1902.06705*, 2019. 5
- [9] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *IEEE symposium on security and privacy (SP)*, 2017. 5
- [10] Yoonki Cho, Woo Jae Kim, Seunghoon Hong, and Sung-Eui Yoon. Part-based pseudo label refinement for unsupervised person re-identification. In *CVPR*, 2022. 1
- [11] Francesco Croce and Matthias Hein. Minimally distorted adversarial examples with a fast adaptive boundary attack. In *ICML*, 2020. 6
- [12] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *ICML*, 2020. 6
- [13] Jiequan Cui, Shu Liu, Liwei Wang, and Jiaya Jia. Learnable boundary guided adversarial training. In *ICCV*, 2021. 5
- [14] Guneet S Dhillon, Kamyar Azizzadenesheli, Zachary C Lipton, Jeremy Bernstein, Jean Kossaifi, Aran Khanna, and Anima Anandkumar. Stochastic activation pruning for robust adversarial defense. In *ICLR*, 2018. 3, 6
- [15] Yinpeng Dong, Zhijie Deng, Tianyu Pang, Jun Zhu, and Hang Su. Adversarial distributional training for robust deep learning. In *NeurIPS*, 2020. 5
- [16] Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Evading defenses to transferable adversarial examples by translation-invariant attacks. In *CVPR*, 2019. 6
- [17] Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. Robust physical-world attacks on deep learning visual classification. In *CVPR*, 2018. 1, 8
- [18] Angus Galloway, Anna Golubeva, Thomas Tanay, Medhat Moussa, and Graham W Taylor. Batch normalization is a cause of adversarial vulnerability. In *ICML Workshop*, 2019. 3
- [19] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015. 1, 2, 5
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1, 5
- [21] Jiabo Huang, Qi Dong, Shaogang Gong, and Xiatian Zhu. Unsupervised deep learning by neighbourhood discovery. In *ICML*, 2019. 7
- [22] Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. In *NeurIPS*, 2019. 1, 3
- [23] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *ICLR*, 2017. 4
- [24] Gaojie Jin, Xiping Yi, Wei Huang, Sven Schewe, and Xiawei Huang. Enhancing adversarial training with second-order statistics of weights. In *CVPR*, 2022. 2
- [25] Harini Kannan, Alexey Kurakin, and Ian Goodfellow. Adversarial logit pairing. *arXiv preprint arXiv:1803.06373*, 2018. 2
- [26] Woo Jae Kim, Seunghoon Hong, and Sung-Eui Yoon. Diverse generative perturbations on attention space for transferable adversarial attacks. In *ICIP*, 2022. 1
- [27] Stepan Komkov and Aleksandr Petiushko. Advhat: Real-world adversarial attack on arcfac face id system. In *ICPR*, 2021. 1, 8
- [28] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 5
- [29] Yandong Li, Lijun Li, Liqiang Wang, Tong Zhang, and Boqing Gong. Nattack: Learning the distributions of adversarial examples for an improved black-box attack on deep neural networks. In *ICML*, 2019. 6
- [30] Fangzhou Liao, Ming Liang, Yinpeng Dong, Tianyu Pang, Xiaolin Hu, and Jun Zhu. Defense against adversarial attacks using high-level representation guided denoiser. In *CVPR*, 2018. 1
- [31] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into transferable adversarial examples and black-box attacks. In *ICLR*, 2017. 1
- [32] Divyam Madaan, Jinwoo Shin, and Sung Ju Hwang. Adversarial neural pruning with latent vulnerability suppression. In *ICML*, 2020. 3
- [33] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018. 1, 2, 5
- [34] Aamir Mustafa, Salman Khan, Munawar Hayat, Roland Goecke, Jianbing Shen, and Ling Shao. Adversarial defense by restricting the hidden space of deep neural networks. In *ICCV*, 2019. 3
- [35] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011. 5

- [36] Tianyu Pang, Chao Du, Yinpeng Dong, and Jun Zhu. Towards robust detection of adversarial examples. In *NeurIPS*, 2018. 1
- [37] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia conference on computer and communications security*, 2017. 4
- [38] Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. Distillation as a defense to adversarial perturbations against deep neural networks. In *IEEE symposium on security and privacy (SP)*, 2016. 1
- [39] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016. 1
- [40] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, 2017. 1, 2
- [41] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K Reiter. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proceedings of the 2016 ACM sigsac conference on computer and communications security*, 2016. 1, 8
- [42] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 5
- [43] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *ICLR*, 2014. 1
- [44] Florian Tramer, Nicholas Carlini, Wieland Brendel, and Aleksander Madry. On adaptive attacks to adversarial example defenses. In *NeurIPS*, 2020. 6
- [45] Hongjun Wang and Yisen Wang. Self-ensemble adversarial training for improved robustness. In *ICLR*, 2022. 2
- [46] Haotao Wang, Aston Zhang, Shuai Zheng, Xingjian Shi, Mu Li, and Zhangyang Wang. Removing batch normalization boosts adversarial training. In *ICML*, 2022. 3
- [47] Yisen Wang, Difan Zou, Jinfeng Yi, James Bailey, Xingjun Ma, and Quanquan Gu. Improving adversarial robustness requires revisiting misclassified examples. In *ICLR*, 2019. 2, 5
- [48] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. 7
- [49] Chang Xiao, Peilin Zhong, and Changxi Zheng. Enhancing adversarial defense by k-winners-take-all. In *ICLR*, 2020. 3, 6
- [50] Cihang Xie, Yuxin Wu, Laurens van der Maaten, Alan L Yuille, and Kaiming He. Feature denoising for improving adversarial robustness. In *CVPR*, 2019. 1, 2, 3, 4
- [51] Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, and Alan L Yuille. Improving transferability of adversarial examples with input diversity. In *CVPR*, 2019. 6
- [52] Yujia Xie, Hanjun Dai, Minshuo Chen, Bo Dai, Tuo Zhao, Hongyuan Zha, Wei Wei, and Tomas Pfister. Differentiable top-k with optimal transport. In *NeurIPS*, 2020. 4
- [53] Kaidi Xu, Sijia Liu, Gaoyuan Zhang, Mengshu Sun, Pu Zhao, Quanfu Fan, Chuang Gan, and Xue Lin. Interpreting adversarial examples by activation promotion and suppression. *arXiv preprint arXiv:1904.02057*, 2019. 3
- [54] Weilin Xu, David Evans, and Yanjun Qi. Feature squeezing: Detecting adversarial examples in deep neural networks. In *NDSS*, 2018. 1
- [55] Hanshu Yan, Jingfeng Zhang, Gang Niu, Jiashi Feng, Vincent Tan, and Masashi Sugiyama. Cifs: Improving adversarial robustness of cnns via channel-wise importance-based feature selection. In *ICML*, 2021. 1, 2, 3, 4
- [56] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016. 5
- [57] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In *ICML*, 2019. 2, 5
- [58] Jingfeng Zhang, Xilie Xu, Bo Han, Gang Niu, Lizhen Cui, Masashi Sugiyama, and Mohan Kankanhalli. Attacks which do not kill training make adversarial learning stronger. In *ICML*, 2020. 2
- [59] Daniel Zoran, Mike Chrzanowski, Po-Sen Huang, Sven Gowal, Alex Mott, and Pushmeet Kohli. Towards robust image classification using sequential attention models. In *CVPR*, 2020. 3