

# Bilateral Adversarial Training: Towards Fast Training of More Robust Models Against Adversarial Attacks

Jianyu Wang  
Baidu Research USA  
wjyouch@gmail.com

Haichao Zhang  
Baidu Research USA  
hczhang1@gmail.com

## Abstract

*In this paper, we study fast training of adversarially robust models. From the analyses of the state-of-the-art defense method, i.e., the multi-step adversarial training [34], we hypothesize that the gradient magnitude links to the model robustness. Motivated by this, we propose to perturb both the image and the label during training, which we call Bilateral Adversarial Training (BAT). To generate the adversarial label, we derive an closed-form heuristic solution. To generate the adversarial image, we use one-step targeted attack with the target label being the most confusing class. In the experiment, we first show that random start and the most confusing target attack effectively prevent the label leaking and gradient masking problem. Then coupled with the adversarial label part, our model significantly improves the state-of-the-art results. For example, against PGD100 white-box attack with cross-entropy loss, on CIFAR10, we achieve 63.7% versus 47.2%; on SVHN, we achieve 59.1% versus 42.1%. At last, the experiment on the very (computationally) challenging ImageNet dataset further demonstrates the effectiveness of our fast method.*

## 1. Introduction

Deep learning has achieved great success in many visual recognition tasks in computer vision. However, deep neural networks are extremely vulnerable to adversarial attacks [50]. Specifically, the network is easily fooled to make wrong predictions in the face of adversarial examples, which are adversarially manipulated images by adding small and imperceptible perturbations. This poses a great danger to deploying real-world machine learning systems. Therefore, training an adversarially robust model is of great value towards commercialized AI technology.

In the recent years, many approaches have been proposed to defend against adversarial examples. As demonstrated by [1], multi-step adversarial training [34] is currently the best defense method. Specifically, adversarial

training solves a minimax (saddle point) problem. The inner maximization generates adversarial examples by multi-step projected gradient descent (PGD), which are then used in the outer minimization to optimize the network parameters.

To understand the working mechanism of the multi-step adversarial training, we first perform two diagnostic experiments on CIFAR10. The first experiment is to test a seemingly correct assumption: using stronger adversarial attacks during training will lead to more robust models. To this end, we compare two adversarially trained models, which only differ in the hyper-parameters of the inner maximization: 1) the default setting in [34], denoted by PGD7-2, where the number of iterations is 7 and the step size is 2 pixels; 2) a model trained using a weaker attack, denoted by PGD2-8, where the number of iteration is 2 and the step size is 8 pixels. We observe that PGD2-8 is largely as robust as PGD7-2 under different white-box attacks, even though PGD2-8 attack is weaker than PGD7-2 attack. This result leads us to hypothesize that robustness may not be achieved by simply fitting sufficient adversarial examples during training, and to re-consider if there is more essential ingredients that directly relate to network robustness. With this in mind, we conduct the second experiment where we compare the gradient magnitude of both undefended models and adversarially trained models. We observe that the gradient magnitude of adversarially trained models is much smaller than that of the undefended models. Intuitively speaking, if the gradient (with respect to input images) becomes extremely small, the gradient-based adversarial attacks are likely to fail no matter how many iterations are used. This inspires us that gradient magnitude may directly link to model robustness.

Based on the above observations, in order to achieve adversarial robustness, we would like a model to satisfy the following two conditions: 1) low loss (zero-order condition); 2) small gradient magnitude (first-order condition). To this end, in this paper, we propose a formulation to achieve these two conditions by perturbing both input images and labels during training, which we call Bilateral Adversarial Training (BAT). As for generating the adversarial image, we adopt one-step PGD which speeds up the train-

ing by multiple times compared to [34]. In order to avoid the troublesome label leaking and gradient masking problem often induced by one-step PGD [29, 51], we do the following: 1) using targeted attack with the target label being the most confusing class; 2) adding random uniform noise to the original image as initialization, i.e., random start as in [34]. As for generating the adversarial label, we derive a formula to perturb the groundtruth label (in the form of one-hot vector) based on the gradient with respect to input label (i.e., the negative logarithm probability). As a special case, this solution reduces to label smoothing when the gradient of non-groundtruth classes are equal.

In the experiment, we first empirically demonstrate that random start and the MC targeted attack are very effective at avoiding label leaking and gradient masking problem. The model trained by using these two techniques alone can achieve similar robustness as multi-step adversarially trained models in [34]. Next, after adding adversarial label part, our model significantly improves the state-of-the-art results in [34]. In order for rigorous robustness evaluation, we use the strong white-box attacks such as PGD100 and PGD1000 with both cross-entropy loss and margin-based loss. For example, against PGD100 under cross-entropy loss, on CIFAR10, we achieve 63.7% versus 47.2%; on SVHN, we achieve 59.1% versus 42.1%. At last, we apply our fast method to the very challenging ImageNet dataset. Our model is successfully trained using only 8 GPUs, compared with 53 GPUs [26] and 128 GPUs [55]. Compared with the recent state-of-the-art [55], our model is better on clean images and against non-targeted attacks, but worse against randomly targeted attacks, using an order-of-magnitude less computational resources.

In summary, our contribution is threefold. First, we empirically show that small gradient magnitude may improve the adversarial robustness. Second, we propose a fast adversarial training method called BAT, which perturbs both the image and the label. Third, our method significantly improves the state-of-the-art results on several datasets.

## 2. Related Work

### 2.1. Adversarial Attacks

Adversarial examples have long been studied in machine learning [12, 23, 3, 4]. In the time of modern deep learning, [50] first pointed out that CNNs are vulnerable to adversarial examples, and proposed a box-constrained L-BFGS method to compute them. Later on, [19] proposed the fast gradient sign method (FGSM) to efficiently generate adversarial examples. FGSM was then extended to an iterative version in [28], which showed that adversarial examples can exist in the physical world. In [38], the authors proposed DeepFool to compute the adversarial perturbations, and define and quantify the robustness of classifiers. In [7], the

famous CW attack was proposed, which used the margin-based loss, and applied change-of-variables to remove the constraint. In spite of being very slow, CW attack is currently one of the strongest attacks. Later [9] modified the loss function in [7] by applying elastic net regularization.

There are some works devoted to improving the transferability of adversarial examples, which leads to stronger black-box attacks. [33] proposed to compute the adversarial perturbation by attacking an ensemble of networks simultaneously, and demonstrated improved transferability. In [40], the authors assumed a scenario where the attackers have access to the prediction results of a few examples. They then trained a substitute/surrogate model based on the limited number of examples, and generate adversarial examples using the substitute model. [15] demonstrated that momentum-based iterative attacks achieve better transferability. There are some works proposing zeroth-order attacks, i.e., using the logit to generate the attacks [52, 10]. Besides, [5] proposed the boundary attack, which is based on the final model decision, instead of the gradient or logit.

In addition to image classification, adversarial examples were also studied in many other tasks, including object detection [54], semantic segmentation [54, 37], speech recognition [11], image captioning [8], deep reinforcement learning [24, 42]. Besides the additive perturbation model, [17] studied how to generate adversarial examples under rotation and translation. [18] studied physically adversarial example in the context of detecting stop sign in real world. Another interesting topic is given by [2], where the authors synthesized robust adversarial examples in 3D.

### 2.2. Adversarial Defenses

In recent years, many methods have been proposed to defend against adversarial examples. One line of research is on detecting adversarial examples, such as [36, 35]. But later [6] showed that their CW attack is able to bypass most detection methods. Another line of research tries to break the special structure in adversarial perturbation by random or non-differentiable operations [53, 20, 48, 46, 31, 43, 32]. Recently, [1] showed that many existing defense methods relied on gradient masking, which leads to a false sense of robustness against adversarial attacks. Besides, gradient-based regularization [25, 45] and nearest neighbor [16] have been demonstrated to improve robustness.

Adversarial training [19, 29, 51, 34, 47, 44, 55, 57] is currently the best defense method against adversarial attacks. [29] first scaled up adversarial training to ImageNet dataset, where the authors used one-step least likely targeted attack to generate adversarial examples during training. Later in [51], the authors pointed out that such adversarially trained models suffer from gradient masking, and proposed ensemble adversarial training, which augmented the training data with perturbations computed from a set of

held-out models. [34] demonstrated that multi-step adversarial training is very effective at achieving robustness, and also managed to avoid the gradient masking problem. According to [1], this is currently the best defense method.

### 3. Motivation

In this section, we empirically analyze two aspects of the multi-step adversarial training method in [34]: 1) if more iterations in the inner maximization improves the robustness, and 2) the gradient magnitude of both undefended models and adversarially trained models. The experiments are conducted on CIFAR10. Based on the analyses, we hypothesize that making the loss surface locally flat (i.e., small local gradient magnitude) helps achieve better robustness. The proposed algorithm will then be given in the next section.

#### 3.1. Background

We first give a short description of the adversarial training method in [34]. This method achieves the currently best adversarial robustness according to [1]. Specifically, it solves the following saddle point (minimax) problem

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

Here  $(x, y)$  denotes the original data point,  $x'$  denotes the adversarially perturbed image,  $L(\cdot)$  denotes the loss function, and  $\epsilon_x$  denotes the perturbation budget. The feasible region  $\mathcal{S}_x$  is defined as <sup>1</sup>

$$\mathcal{S}_x = \{z \mid z \in B(x, \epsilon_x) \cap [-1.0, 1.0]^n\}, \quad (2)$$

where  $B(x, \epsilon_x) \triangleq \{z \mid \|z - x\|_{\infty} \leq \epsilon_x\}$  denotes the  $\ell_{\infty}$ -ball with center  $x$  and radius  $\epsilon_x$ . In the following, for the sake of notational simplicity, without loss of generality, we present the formulation based on a single sample. The outer minimization is minimizing the cross-entropy loss as in the standard classification. The inner maximization corresponds to the adversarial attack. In order to better explore the solution in  $B(x, \epsilon_x)$ , [34] uses random start before taking a number of PGD steps, i.e.,

$$x^0 \sim B(x, \epsilon_x), \quad (3)$$

$$x^{t+1} = \Pi_{\mathcal{S}_x}(x^t + \epsilon_x \cdot \text{sign}(\nabla_x L(x^t, y; \theta))). \quad (4)$$

The original image  $x$  is at first randomly (uniform) perturbed to some point  $x^0$  in  $B(x, \epsilon_x)$  as in (3), and then goes through several PGD steps as in (4). The  $\Pi_{\mathcal{S}_x}(\cdot)$  operator projects the input into the feasible region  $\mathcal{S}_x$ .

#### 3.2. Analyses

**Do more iterations help?** We first examine if more iterations in inner maximization help improve the robustness.

<sup>1</sup>In implementation, we rescale all images with pixel values in  $[-1, 1]$ .

Acc.(%)	clean	FGSM	PGD2-8	PGD7-2	PGD20-2	PGD100-2
PGD7-2	88.0	57.0	53.0	51.2	47.6	47.2
PGD2-8	88.2	56.9	53.2	50.5	46.7	46.2

Table 1: Comparison between the model in [34] (top) and an adversarially trained model using a weaker attack during training (bottom). They achieve similar robustness.

	undefended			adversarially trained		
	all	correct	wrong	all	correct	wrong
min	3.0e-32	3.0e-32	264.1	2.6e-26	2.6e-26	0.2
mean	395.0	23.6	7.4e3	3.8	0.4	28.9
max	4.5e4	7.0e3	4.5e4	236.2	85.9	236.2

Table 2: The minimal, average, and maximal value of gradient magnitude of the test images on CIFAR10. Overall Adversarially trained models have much smaller gradient magnitude than undefended models. All models are trained using the same regularization, epoch, learning rate, etc.

To this end, we compare two adversarially trained models with different hyper-parameters for generating the adversarial examples during training. The first one is the default in [34], denoted by PGD7-2, which uses 7 steps of PGD, and step size is 2.0. The second one is a seemingly weaker variant, denoted by PGD2-8, meaning only 2 steps of PGD are used and step size is 8.0. As in [34], the perturbation budget is 8.0 in training and evaluation, and random start is used. From Table 1, we see that PGD2-8 performs roughly the same as PGD7-2, against PGD attacks with different steps (strength). This result leads us to hypothesize that using stronger attacks during training may not necessarily lead to more robust models.

#### Gradient magnitude of adversarially trained models.

Next we examine the gradient magnitude of the undefended models and adversarially trained models. We consider three collections of all test images in CIFAR10, and for each collection we compute the minimal, average, and maximal of gradient magnitude, i.e.,  $\|\nabla_x L(x, y; \theta)\|_2^2$ . The three collections are: 1) entire images, denoted by “all”, 2) correctly predicted images, denoted by “correct”, 3) wrongly predicted images, denoted by “wrong”. The numbers are shown in Table 2. First we see that for any collection, the gradient magnitude of undefended model is much larger than that of adversarially trained model. Also, for each model, the gradient magnitude of correctly predicted images is much smaller than that of wrongly predicted images.

#### 3.3. Hypothesis

From the above analyses, we hypothesize that small gradient magnitude directly links to the adversarial robustness. Intuitively speaking, if the loss surface is locally “flat” around the data points, the model is hard to attack, no mat-

ter how many steps are used. This hypothesis is aligned with [32] where they call it local Lipschitz. Note that there are papers [22, 14] studying the possible relationship between the flatness of loss surface and the generalization of the model. In this paper, we simply use gradient magnitude to measure the “flatness”. A rigorous treatment is beyond the scope of the current paper and left as future work.

A straightforward idea to reduce the gradient magnitude is to augment the loss function with some form of gradient regularization during training, e.g., [45, 25]. However, the key problem of this idea is that training requires the computation of second-order derivatives, which becomes extremely slow and expensive for large-scale networks.

## 4. Formulation

In this section, in order to improve the adversarial robustness, we propose Bilateral Adversarial Training (BAT), which simultaneously perturbs both the image and the label during adversarial training.

We first approximately relate the general adversarial training framework to small gradient magnitude. Let  $x, x'$  denote the original and the slightly perturbed image, and  $y, y'$  denote the original and the slightly perturbed groundtruth (in the form of a probability distribution lying in the probability simplex). Let  $L(\cdot)$  denote the loss function. The first-order Taylor expansion of the loss is

$$\begin{aligned} L(x', y'; \theta) &\approx L(x, y; \theta) \\ &+ \nabla_x L(x, y; \theta) \cdot (x' - x) \\ &+ \nabla_y L(x, y; \theta) \cdot (y' - y). \end{aligned} \quad (5)$$

We use the perturbation budget constraint in  $\ell_\infty$ -norm, i.e.,

$$\|x' - x\|_\infty \leq \epsilon_x, \quad \|y' - y\|_\infty \leq \epsilon_y. \quad (6)$$

By Holder’s inequality, from (5) we can approximately have the upper bound

$$\begin{aligned} L(x', y'; \theta) &\leq L(x, y; \theta) \\ &+ \epsilon_x \|\nabla_x L(x, y; \theta)\|_1 \\ &+ \epsilon_y \|\nabla_y L(x, y; \theta)\|_1. \end{aligned} \quad (7)$$

Intuitively speaking, adversarial training, by minimizing  $L(x', y'; \theta)$ , translates to 1) minimizing  $L(x, y; \theta)$  and 2) minimizing the gradient magnitude  $\|\nabla_x L(x, y; \theta)\|_1$  and  $\|\nabla_y L(x, y; \theta)\|_1$ . The second point explains the results in Table 2. Note that the first point makes the network predict the correct class, and the second point makes it difficult to generate adversarial examples for gradient-based attacks, because the gradient magnitude becomes very small.

The above formulation does not specify how to generate  $x', y'$ . Mathematically, the optimization problem can be written as

$$\max_{x' \in \mathcal{S}_x, y' \in \mathcal{S}_y} L(x', y'; \theta). \quad (8)$$

Here  $\mathcal{S}_x$  is defined by (2), and  $\mathcal{S}_y$  is defined as

$$\mathcal{S}_y = \{z \mid z \in B(y, \epsilon_y), z \geq 0, \sum_i z_i = 1\}. \quad (9)$$

The final formulation for adversarial training is as follows

$$\min_{\theta} \left[ \max_{x' \in \mathcal{S}_x, y' \in \mathcal{S}_y} L(x', y'; \theta) \right], \quad (10)$$

where  $(x, y) \sim \mathcal{D}$ . Our simple strategy to solve (10) is to decompose it into two separate sub-problems, and optimize over  $x'$  or  $y'$  conditioned on the other respectively. After obtaining  $x', y'$ , we use them in place of the original  $x, y$  as the training data points and optimize over  $\theta$ . In other words, the training batch only contains adversarially perturbed images. In the next two subsections, we will describe the solution to each sub-problem respectively.

### 4.1. Generating Adversarial Labels

We first study how to compute the adversarial label  $y'$  given the original data point  $x, y$ . We need to solve

$$\max_{y' \in \mathcal{S}_y} L(x, y'; \theta). \quad (11)$$

Here the original groundtruth  $y$  is a one-hot vector, i.e.,  $y_c = 1$  and  $y_k = 0, k \neq c$ . We use  $k$  to denote the class index and  $c$  to denote the groundtruth class. The most straightforward idea is to use the one-step PGD

$$y' = \Pi_{\mathcal{S}_y}(y + \alpha \nabla_y L(x, y; \theta)), \quad (12)$$

$$\nabla_{y_k} L(x, y; \theta) = -\log(p_k). \quad (13)$$

Here the  $\Pi_{\mathcal{S}_y}(\cdot)$  operator projects the input into the feasible region  $\mathcal{S}_y$ . Basically it ensures that the adversarial label  $y'$  is 1) in  $B(y, \epsilon_y)$  and 2) in the probability simplex. Next we propose a heuristic solution to achieve both. In the following, we will use some short notations.

$$v_k = \nabla_{y_k} L(x, y; \theta), \quad v_{MC} = \min_{k \neq c} v_k, \quad v_{LL} = \max_{k \neq c} v_k.$$

Here “MC” (most confusing) corresponds to the non-groundtruth class with the highest probability, and “LL” (least likely) corresponds to the non-groundtruth class with the lowest probability. The idea is that we decrease  $y_c$  from 1 to  $1 - \epsilon_y$ , and distribute  $\epsilon_y$  to other non-groundtruth classes. The share for each class is based on their respective gradient  $\nabla_{y_k} L(x, y; \theta)$ , while the share for the MC class (i.e.,  $y'_{MC}$ ) is set to be very small. This way, we can obtain

$$y'_k = \frac{\epsilon_y}{n-1} \cdot \frac{v_k - v_{MC} + \gamma}{\sum_{k \neq c} \frac{v_k}{n-1} - v_{MC} + \gamma}, \quad k \neq c. \quad (14)$$

Here  $\gamma$  is a very small value, e.g., 0.01. Please refer to the supplementary material for another heuristic solution and the comparison. It is easy to see that if the gradient of

non-groundtruth classes are equal, the second multiplicative term becomes 1 and we then obtain

$$y'_k = \frac{\epsilon_y}{n-1}, \quad k \neq c. \quad (15)$$

This is exactly the label smoothing [49]. In other words, label smoothing can be thought of as an adversarial perturbation of the groundtruth label.

Note that  $\epsilon_y$  controls the perturbation budget of  $y$ . We are interested in finding the largest  $\epsilon_y$  that leads to the most adversarially perturbed label. The idea is that we want to keep the probability of the groundtruth class (i.e.,  $y'_c$ ) at least  $\beta$  times larger than the maximal probability over non-groundtruth classes (i.e.,  $y'_{LL}$ ). Mathematically, we want

$$y'_c \geq \beta \cdot \max_{k \neq c} y'_k. \quad (16)$$

Solving for the following equation

$$1 - \epsilon_y \geq \frac{\beta \epsilon_y}{n-1} \cdot \frac{v_{LL} - v_{MC} + \gamma}{\frac{\sum_{k \neq c} v_k}{n-1} - v_{MC} + \gamma}, \quad (17)$$

we obtain

$$\epsilon_y \leq \frac{1}{1 + \frac{\beta}{n-1} \frac{v_{LL} - v_{MC} + \gamma}{\frac{\sum_{k \neq c} v_k}{n-1} - v_{MC} + \gamma}}. \quad (18)$$

Next we consider two extreme cases.

1) The probabilities of non-groundtruth classes are evenly distributed, i.e., label smoothing. In this case,  $v_{LL} = v_{MC}$ . Then we have

$$\epsilon_y = \frac{1}{1 + \frac{\beta}{n-1}}. \quad (19)$$

Take CIFAR10 for example ( $n = 10$ ). We have  $\epsilon_y = 0.1$ ,  $\beta = 81$ , or  $\epsilon_y = 0.5$ ,  $\beta = 9$ , or  $\epsilon_y = 0.9$ ,  $\beta = 1$ .

2) The probabilities of non-groundtruth classes are centered on one class. In this case,  $v_{LL} = \sum_{k \neq c} v_k$ ,  $v_{MC} = 0$ . Then we have

$$\epsilon_y = \frac{1}{1 + \beta \frac{v_{LL} + \gamma}{v_{LL} + (n-1)\gamma}} \approx \frac{1}{1 + \beta}. \quad (20)$$

Note that  $\gamma$  is usually very small, e.g.  $\gamma = 0.01$ .

We can see that given the multiplier  $\beta$ , the range of perturbation budget  $\epsilon_y$  is

$$\epsilon_y \in \left( \frac{1}{1 + \beta}, \frac{1}{1 + \frac{\beta}{n-1}} \right]. \quad (21)$$

Note that we only need to specify a proper  $\beta$ . As a special case,  $\beta = \infty$  corresponds to the original one-hot label.

## 4.2. Generating Adversarial Images

Next we study how to compute the adversarial image  $x'$  given the original data point  $x, y$ . Mathematically, we need to solve the problem

$$\max_{x' \in \mathcal{S}_x} L(x', y; \theta). \quad (22)$$

This is the adversarial attack problem. For non-targeted attacks, we directly maximize (22). The downside of using non-targeted attacks is label leaking [29]. This is because during training the model implicitly learns to infer the true label from the adversarial perturbation (gradient). In other words, the model smartly finds a shortcut (degenerate minimum [51]) towards the local optima. A more general and severe problem is gradient masking [41]. It refers to the fact that the loss surface of the model is very jagged, and so it becomes harder for the attackers to find good gradient during the iterative attack. As demonstrated in [1], gradient masking (a.k.a. gradient obfuscation) gives a false sense of robustness, and the model gets broken in the face of strong attacks with large number of iterations.

Recently, two techniques were proposed to reduce or avoid gradient masking problem: 1) using multi-step PGD [34]; 2) using an ensemble of models to generate adversarial examples [51]. However, the effectiveness comes with expensive time cost [34] or memory cost [51]. Since one of our design consideration is speed, in this paper, we focus on two simple techniques: 1) using targeted attack [29]; 2) adding random noise as in Eq. (3) [51, 34].

As for targeted attack, in [29], the authors used the Least Likely (LL) class as the targeted class, i.e.,

$$y' = \operatorname{argmax}_{\hat{y} \neq y} L(x, \hat{y}; \theta). \quad (23)$$

Differently, here in this paper, we use the Most Confusing (MC) class as the targeted class, i.e.,

$$y' = \operatorname{argmin}_{\hat{y} \neq y} L(x, \hat{y}; \theta). \quad (24)$$

In order for fast training, we use one-step PGD (in the experiment on the difficult ImageNet we use two-step). Note that the update equations (3) and (4) are for non-targeted attacks. For targeted attacks, we simply replace the groundtruth label  $y$  by the targeted label  $y'$ , and also replace the plus sign by the minus sign, in order to minimize the loss with respect to the targeted label.

## 5. Experiments

**Dataset and Network.** In the experiment, we use CIFAR10 [27], and SVHN [39] and the large-scale ImageNet [13]. We do not use MNIST [30] because it is not a good testbed due to the near-binary nature of the images [51]. For CIFAR10 and SVHN, we use Wide

ResNet [56] (WRN-28-10). For ImageNet, we use ResNet family [21]. Most of the diagnostic experiments are conducted on CIFAR10, because it is currently the most commonly used dataset for adversarial training.

**Evaluation.** Based on the amount of knowledge that attackers have, there are several types of attacks: i) Gradient-based (white-box), where the attackers have full knowledge of the model (structure, parameters); ii) Score(logit)-based, where the attackers know the score/logit vector of the model (e.g. SPSA [52], ZOO [10]); iii) Decision-based, where the attackers only know the predicted class (e.g. boundary attack [5]). Note that the more information the attackers have, the stronger the adversarial attacks will be. In the experiment we use the strongest gradient-based white-box attacks.

For CIFAR10 and SVHN, we follow the evaluation setup in [34]. Specifically, the test perturbation budget is 8 pixels. In order to use strong attacks to evaluate the model robustness, we 1) always use non-targeted attack, and 2) drop random start for one-step PGD attack (i.e. FGSM), and use random start for multi-step PGD attack. Also, for one-step attack, the step size is the perturbation budget (i.e., 8 pixels), and for multi-step attack, the step size is 2 pixels.

For ImageNet, we follow the evaluation setup in [26, 55]. Specifically, the test perturbation budget is 16 pixels. We use the non-targeted attack, and the targeted attack where the label is randomly selected. The step size is 1 pixel, except for PGD10 attack where the step size is 1.6 pixel.

**Implementation Details.** For CIFAR10 and SVHN, we largely follow the released code in [34]. The learning rate schedule is [0.1, 0.01, 0.001] for CIFAR and [0.01, 0.001, 0.0001] for SVHN. For the short training, the decay epoch schedule is [60, 90, 100]. And for the long training, the epoch schedule [100, 150, 200]. In all the tables, the models of long training are postfixed by “+”. For ImageNet, we use the Tensorpack package and perform distributed adversarial training with 8 GPUs. We largely follow the code for distributed training ResNet models on ImageNet. The learning rate schedule is [0.1, 0.01, 0.001, 0.0001], and the decay epoch schedule is [30, 60, 90, 100]. For ResNet50, the training takes about 2 days on a DGX machine. We plan to release the code and models after this work is published.

### 5.1. Random Start and MC Targeted Attack

In this subsection, we demonstrate that for one-step adversarial training, random start and MC targeted attack are effective at preventing label leaking and gradient masking problem. This diagnostic experiment is conducted on CIFAR10. To this end, we evaluate three different ways of generating adversarial examples during training: FGSM attack, LL targeted attack, and MC targeted attack. For each option, we perform one-step adversarial training with or without random start. This leads to six adversarially trained

Acc.(%)	clean	FGSM		CE7		black
		w.o. RS	RS	w.o. RS	RS	
FGSM	55.2	99.1	68.6	0.0	0.0	56.2
R-FGSM	89.8	55.8	63.6	46.4	48.0	88.0
LL	92.6	97.9	86.2	0.0	0.0	80.9
R-LL	91.4	46.6	56.6	34.1	36.0	88.2
MC	86.4	70.7	73.0	37.6	40.3	84.3
R-MC	89.9	62.6	70.2	46.8	48.4	87.1

Table 3: The classification accuracy of one-step adversarially trained models, using different attacks, and, with or without random start. The models trained using random start are prefixed with “R-”. We see that random start and MC targeted attack are effective at preventing label leaking and gradient masking problem.

Acc.(%)	FGSM		MC		LL	
	w.o. RS	RS	w.o. RS	RS	w.o. RS	RS
R-FGSM	55.8	63.6	55.4	63.6	75.5	79.8
R-LL	46.6	56.6	44.0	55.6	70.7	76.4
R-MC	62.6	70.2	63.9	71.3	80.1	83.8

Table 4: The classification accuracy of three attacks, i.e., FGSM attack, LL targeted attack and MC targeted attack, with or without random start. The rows correspond to different adversarially trained models. We see that MC targeted attack has similar strength as FGSM attack, and both are much stronger than LL targeted attack.

models in total. The perturbation budget is 8 pixels in training. The results are shown in Table 3. The rows correspond to different models, where the prefix “R” means that random start is used. The columns correspond to non-targeted attacks using one-step (FGSM) or 7-step (CE7, CE is short for Cross-Entropy), with or without random start (denoted by “RS”). The last column corresponds to the black-box attack using the undefended model and FGSM attack (w.o. RS). Firstly we see that the target models trained by FGSM and LL suffer badly from the label leaking problem because the accuracy against FGSM attack is even higher than the clean accuracy. But this is just false robustness and the accuracy drops to zero under CE7 attacks. Next, after applying random start, R-FGSM and R-LL become quite robust, demonstrating random start helps the model avoid the label leaking problem during adversarial training. Lastly, we see that the model R-MC performs the best, while R-LL performs the worst, against FGSM and CE7 attacks.

We hypothesize that the adversarially trained mode by one-step LL targeted attack is weak, because the LL targeted attack is weak by itself. Table 4 shows the strength of these three attacks using one-step PGD. The rows correspond to different models trained using random start. The columns correspond to three attacks with or without random start. We see that MC targeted attack is roughly as strong as FGSM attack, and both of them are much stronger than LL

Acc.(%)	clean	FGSM	CE7	CE20
R-FGSM	89.8	55.8	48.0	42.9
R-FGSM-LS ( $\epsilon_y = 0.5$ )	89.1	62.0	54.6	49.0
R-MC	89.9	62.6	48.4	43.4
R-MC-LS ( $\epsilon_y = 0.5$ )	91.1	70.6	59.2	53.3
R-MC-LS+ ( $\epsilon_y = 0.5$ )	<b>91.8</b>	<b>71.4</b>	62.7	55.9
R-MC-LA ( $\beta = 9$ )	90.7	69.6	59.9	55.3
R-MC-LA+ ( $\beta = 9$ )	91.2	70.7	<b>63.0</b>	<b>57.8</b>
Madry [34]	87.3	56.1	50.0	45.8
Madry*	88.0	57.0	51.2	47.6
Madry-LA	86.8	63.4	57.8	53.2
Madry-LA+	87.5	65.9	61.3	57.5

Table 5: The classification accuracy of R-MC-LA models and variants under various white-box attacks on CIFAR10.

targeted attack. This is probably because it is usually hard to slightly manipulate the original image so that it becomes a visually very different class. Therefore, we argue for using MC targeted attack during adversarial training because 1) it is much stronger than LL targeted attack; 2) FGSM attack risks label leaking and gradient masking problem.

We briefly summarize the role of random start. In training, it effectively prevents the label leaking and gradient masking problem, but in attack, it weakens the strength for one-step PGD attack (shown in Table 4). As we will show later, random start has very little effect for multi-step PGD attack, especially when the number of steps becomes large.

## 5.2. CIFAR10 Dataset

In this subsection, we report results against white-box attacks on CIFAR10 [27] dataset. It has 10 classes, 50K training images (5K per class) and 10K test images. As summarized above, to generate adversarial images, we use random start and MC targeted attack (the perturbation budget is 8 pixels). To generate adversarial labels, we use (18) to compute the budget  $\epsilon_y$  and (14) to compute the adversarially perturbed labels  $y'$ . The resulting model is denoted as R-MC-LA where LA stands for label adversary. We also experiment with label smoothing (LS for short), a special case of adversarial label, and denote this model by R-MC-LS. Our baseline is the multi-step adversarial training method by [34]. We report the original numbers in their paper, denoted by Madry, and also report the numbers by our reproduced model, denoted by Madry\*. The accuracy against various steps of PGD attacks is given in Table 5. We see that perturbing labels during training significantly improves the robustness over the baseline. Label smoothing (R-MC-LS) works pretty well, and label adversary (R-MC-LA) is even better. Besides, we also apply label adversary to the multi-step adversarial training. We see that the resulting models, denoted by Madry-LA and Madry-LA+, significantly improves the original version, further verifying the effec-

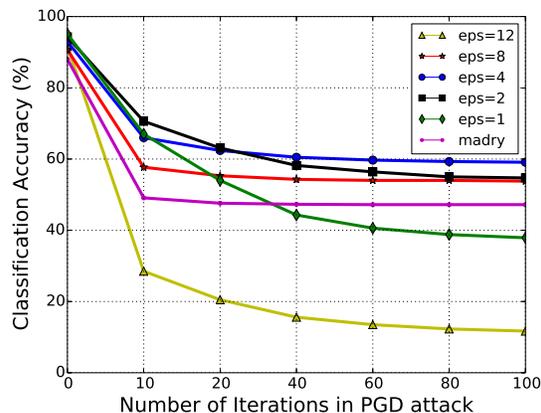


Figure 1: The classification accuracy of the proposed R-MC-LA models against white-box PGD attacks of different number of iterations on CIFAR10. The models are trained by using different perturbation budget. We use  $\beta = 9$ .

Acc.(%)	clean	CE20	CW20	CE100	CW100	CW200
R-MC-LA ( $\epsilon_x = 8$ )	90.8	54.6	53.7	52.9	51.9	51.7
R-MC-LA+ ( $\epsilon_x = 8$ )	91.0	57.5	56.2	55.2	53.8	53.6
R-MC-LA ( $\epsilon_x = 4$ )	<b>93.0</b>	63.1	61.5	60.1	58.0	57.6
R-MC-LA+ ( $\epsilon_x = 4$ )	92.9	<b>66.9</b>	<b>64.2</b>	<b>63.7</b>	<b>60.7</b>	<b>60.3</b>
Madry*	88.0	47.6	48.6	47.2	48.1	48.1

Table 6: The classification accuracy of the proposed R-MC-LA models under various white-box attacks on CIFAR10. To rule out randomness, the numbers are averaged over 3 independently trained models. We use  $\beta = 9$ .

tiveness of label adversary. Interestingly, R-MC-LA(+) performs even better than Madry-LA(+). Lastly, we observe that longer training is helpful to all models.

### 5.2.1 Different Perturbation Budgets during Training

Next we study whether using larger perturbation budget during training leads to more robust models. We train models using the budget  $\epsilon_x \in \{1, 2, 4, 8, 12\}$  pixels during training, and we use 8 pixels for evaluation. Figure 1 shows the classification accuracy with respect to the number of steps in PGD attack. Firstly we observe the general trend that as the number of steps increases, the accuracy drops quickly and then plateaus. Secondly, we find that big budget (i.e.,  $\epsilon_x = 12$ ) or small budget (i.e.,  $\epsilon_x = 1$ ) lead to less robust models. Interestingly, we see that the model trained using  $\epsilon_x = 4$  achieves the best robustness. The exact numbers are given in Table 6. Note that to rule out randomness, the numbers are averaged over 3 independently trained models. We also test the attacks using the margin-based CW loss [7]. For example, CW200 attack means using CW loss and 200 steps PGD. We notice that 1) the baseline model [34] performs similarly against either (cross-entropy-based) CE at-

Acc.(%)	clean	Undefended		another R-MC-LA	
		FGSM	CE20	FGSM	CE20
R-MC-LA	90.7	87.8	88.8	74.4	71.0
R-MC-LA+	91.2	88.5	89.9	74.6	74.4

Table 7: The classification accuracy of R-MC-LA models against black-box attacks on CIFAR10. We use  $\beta = 9$ .

Acc.(%)	clean	FGSM	CE20	CE100	CW100
R-MC-LA ( $\epsilon_x = 8$ )	94.1	66.9	46.7	42.0	40.9
R-MC-LA+ ( $\epsilon_x = 8$ )	94.1	69.8	53.9	50.3	48.9
R-MC-LA ( $\epsilon_x = 4$ )	<b>95.7</b>	72.6	54.4	47.2	45.7
R-MC-LA+ ( $\epsilon_x = 4$ )	95.5	<b>74.2</b>	<b>63.0</b>	<b>59.1</b>	<b>58.5</b>
Madry*	91.8	61.0	43.2	42.1	43.4

Table 8: The classification accuracy of R-MC-LA models under various white-box attacks on SVHN. The models are trained using different perturbation budget. We use  $\beta = 9$ .

tack or (margin-based) CW attack; 2) CW attack is more effective than CE attack when attacking our models (R-MC-LA). Furthermore, we evaluate our best model, R-MC-LA+ ( $\epsilon_x = 4$ ), against 1000-step PGD attacks using CE loss and CW loss. The accuracy is **61.4%** for CE and **59.3%** for CW, which is very close to that against 200-step PGD attack.

### 5.2.2 Black-box Attack Evaluation

We next evaluate our R-MC-LA models against black-box attacks. We use two models to generate the adversarial examples: the undefended model and another randomly initialized R-MC-LA model. All the models are trained using  $\epsilon_x = 8$ . The results are shown in Table 7. We see that attacks generated by R-MC-LA model are stronger than those by undefended model, because two independently trained R-MC-LA models share inherent structure. Besides, we see that all the black-box attacks are weaker than white-box attacks (by comparing the accuracy), demonstrating that our models do not suffer from the gradient masking problem.

### 5.3. SVHN Dataset

The SVHN [39] is a 10-way house number classification dataset. It contains 73257 training images, and 26032 test images. We don't use the additional training images. The results against white-box attacks are shown in Table 8. Similar to CIFAR10, we see that our models significantly outperform the state-of-the-art results on clean images and against PGD attacks of various strength.

### 5.4. ImageNet Dataset

The ImageNet dataset has 1.28 million training images with 1000 classes. We use the validation set with 50K images for evaluation. To the best of our knowledge, up to

Acc.(%)	clean	CE10-nt	CE100-nt	CE10-rd	CE100-rd
R-MC-LA-R50	58.9	14.9	4.0	45.8	24.5
R-MC-LA-R101	61.9	18.0	6.3	45.8	26.0
R-MC-LA-R152	63.9	<b>19.8</b>	<b>7.4</b>	46.5	26.6
[26]-IncepV3	<b>72.0</b>	NA	NA	27.9	NA
[55]-R152	62.3	17.1	7.3	<b>52.5</b>	<b>41.7</b>

Table 9: The classification accuracy of R-MC-LA models under various white-box attacks on ImageNet. We use  $\beta = 100$ . The budget is 16 pixels in training and evaluation.

now, there are only two papers that have applied multi-step adversarial training on ImageNet, because it is very computationally expensive. Specifically, the prior art [26] used 53 P100 GPUs and the recent paper [55] used 128 V100 GPUs. We train our models on a DGX machine with only 8 GPUs and it takes us about 2 days. For fair comparison, we use 16 pixels as test perturbation budget. In our experiment, we find that, using one-step attack during training (in this case step size is 16 pixels) suffers severely from label leaking and gradient masking. We also observe the similar problem on CIFAR10 and SVHN, when training with large budgets, e.g., 12 or 16. In order to make our method work for the 16-pixel evaluation setup, we use two-step MC targeted attack (in this case the step size become 8 pixels). In the experiment, we find that training without label adversary performs very bad, further demonstrating the effectiveness of label adversary. In evaluation, we use both the non-targeted attack and the targeted attack where the target label is uniformly randomly selected. Note that the non-targeted attack is much stronger than the randomly targeted attack, so we believe using both will lead to a more reliable robustness evaluation. Table 9 shows the top-1 accuracy of our method and two baseline methods, where the non-targeted attack is denoted by "nt", and the randomly targeted attack is denoted by "rd". We can see that our methods significantly outperform the prior art [26] against CE10-rd attack. Compared with the recent work [55], our models are better on clean accuracy and against non-targeted attacks, but are worse against randomly targeted attacks. We hypothesize that this may be because that the models in [55] are trained using randomly targeted attack (same as test), and, they are using an-order-of-magnitude more computational resources (30-step PGD during adversarial training).

## 6. Conclusion

In this paper, we proposed to use the adversarial image and the adversarial label during adversarial training. The adversarial image was generated by one-step or two-step MC targeted attack. The adversarial label was computed by an close-form formula. Comprehensive experiments on CIFAR10, SVHN and ImageNet, against various white-box attacks, demonstrate the effectiveness of our method.

## References

- [1] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *International Conference on Machine Learning*, 2018. 1, 2, 3, 5
- [2] Anish Athalye and Ilya Sutskever. Synthesizing robust adversarial examples. In *International Conference on Machine Learning*, 2018. 2
- [3] Battista Biggio, Iginio Corona, Davide Maiorca, Blaine Nelson, Nedim Šrđić, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. Evasion attacks against machine learning at test time. In *Joint European conference on machine learning and knowledge discovery in databases*, 2013. 2
- [4] Battista Biggio and Fabio Roli. Wild patterns: Ten years after the rise of adversarial machine learning. *Pattern Recognition*, 84, 2018. 2
- [5] Wieland Brendel, Jonas Rauber, and Matthias Bethge. Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. In *International Conference on Learning Representations*, 2018. 2, 6
- [6] Nicholas Carlini and David Wagner. Adversarial examples are not easily detected: Bypassing ten detection methods. In *ACM Workshop on Artificial Intelligence and Security*, 2017. 2
- [7] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy*, 2017. 2, 7
- [8] Hongge Chen, Huan Zhang, Pin-Yu Chen, Jinfeng Yi, and Cho-Jui Hsieh. Attacking visual language grounding with adversarial examples: A case study on neural image captioning. In *Annual Meeting of the Association for Computational Linguistics*, 2018. 2
- [9] Pin-Yu Chen, Yash Sharma, Huan Zhang, Jinfeng Yi, and Cho-Jui Hsieh. Ead: elastic-net attacks to deep neural networks via adversarial examples. In *AAAI Conference on Artificial Intelligence*, 2018. 2
- [10] Pin-Yu Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui Hsieh. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *ACM Workshop on Artificial Intelligence and Security*, 2017. 2, 6
- [11] Moustapha Cisse, Yossi Adi, Natalia Neverova, and Joseph Keshet. Houdini: Fooling deep structured prediction models. In *Advances in Neural Information Processing Systems*, 2017. 2
- [12] Nilesh Dalvi, Pedro Domingos, Sumit Sanghai, Deepak Verma, et al. Adversarial classification. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2004. 2
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition*, 2009. 5
- [14] Laurent Dinh, Razvan Pascanu, Samy Bengio, and Yoshua Bengio. Sharp minima can generalize for deep nets. In *Proceedings of the 34th International Conference on Machine Learning*, 2017. 4
- [15] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Xiaolin Hu, Jianguo Li, and Jun Zhu. Boosting adversarial attacks with momentum. In *Computer Vision and Pattern Recognition*, 2018. 2
- [16] Abhimanyu Dubey, Laurens van der Maaten, Zeki Yalniz, Yixuan Li, and Dhruv Mahajan. Defense against adversarial images using web-scale nearest-neighbor search. In *Computer Vision and Pattern Recognition*, 2019. 2
- [17] Logan Engstrom, Dimitris Tsipras, Ludwig Schmidt, and Aleksander Madry. A rotation and a translation suffice: Fooling CNNs with simple transformations. In *International Conference on Learning Representations*, 2018. 2
- [18] Ivan Evtimov, Kevin Eykholt, Earlene Fernandes, Tadayoshi Kohno, Bo Li, Atul Prakash, Amir Rahmati, and Dawn Song. Robust physical-world attacks on deep learning models. In *Computer Vision and Pattern Recognition*, 2018. 2
- [19] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015. 2
- [20] Chuan Guo, Mayank Rana, Moustapha Cissé, and Laurens van der Maaten. Countering adversarial images using input transformations. In *International Conference on Learning Representations*, 2018. 2
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, 2016. 6
- [22] Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural Computation*, 9(1), 1997. 4
- [23] Ling Huang, Anthony D Joseph, Blaine Nelson, Benjamin IP Rubinstein, and JD Tygar. Adversarial machine learning. In *ACM workshop on Security and artificial intelligence*, 2011. 2
- [24] Sandy Huang, Nicolas Papernot, Ian Goodfellow, Yan Duan, and Pieter Abbeel. Adversarial attacks on neural network policies, 2017. 2
- [25] Daniel Jakubovitz and Raja Giryes. Improving DNN robustness to adversarial attacks using Jacobian regularization. In *European Conference on Computer Vision*, 2018. 2, 4
- [26] Harini Kannan, Alexey Kurakin, and Ian Goodfellow. Adversarial logit pairing. *arXiv preprint arXiv:1803.06373*, 2018. 2, 6, 8
- [27] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009. 5, 7
- [28] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. In *International Conference on Learning Representations Workshop*, 2017. 2
- [29] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. In *International Conference on Learning Representations*, 2017. 2, 5
- [30] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 1998. 5
- [31] Fangzhou Liao, Ming Liang, Yinpeng Dong, and Tianyu Pang. Defense against adversarial attacks using high-level representation guided denoiser. In *Computer Vision and Pattern Recognition*, 2018. 2

- [32] Xuanqing Liu, Minhao Cheng, Huan Zhang, and Cho-Jui Hsieh. Towards robust neural networks via random self-ensemble. In *European Conference on Computer Vision*, 2018. 2, 4
- [33] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into transferable adversarial examples and black-box attacks. In *International Conference on Learning Representations*, 2017. 2
- [34] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018. 1, 2, 3, 5, 6, 7
- [35] Dongyu Meng and Hao Chen. Magnet: a two-pronged defense against adversarial examples. In *ACM SIGSAC Conference on Computer and Communications Security*, 2017. 2
- [36] Jan Hendrik Metzen, Tim Genewein, Volker Fischer, and Bastian Bischoff. On detecting adversarial perturbations. In *International Conference on Learning Representations*, 2017. 2
- [37] Jan Hendrik Metzen, Mummadi Chaithanya Kumar, Thomas Brox, and Volker Fischer. Universal adversarial perturbations against semantic image segmentation. In *International Conference on Computer Vision*, 2017. 2
- [38] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *Computer Vision and Pattern Recognition*, 2016. 2
- [39] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011. 5, 8
- [40] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *ACM on Asia Conference on Computer and Communications Security*, 2017. 2
- [41] 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*, 2016. 5
- [42] Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. Robust adversarial reinforcement learning. In *International Conference on Machine Learning*, 2017. 2
- [43] Aaditya Prakash, Nick Moran, Solomon Garber, Antonella DiLillo, and James Storer. Deflecting adversarial attacks with pixel deflection. In *Computer Vision and Pattern Recognition*, 2018. 2
- [44] Aditi Raghunathan, Jacob Steinhardt, and Percy Liang. Certified defenses against adversarial examples. In *International Conference on Learning Representations*, 2018. 2
- [45] Andrew Slavin Ross and Finale Doshi-Velez. Improving the adversarial robustness and interpretability of deep neural networks by regularizing their input gradients. In *AAAI Conference on Artificial Intelligence*, 2018. 2, 4
- [46] Pouya Samangouei, Maya Kabkab, and Rama Chellappa. Defense-GAN: Protecting classifiers against adversarial attacks using generative models. In *International Conference on Learning Representations*, 2018. 2
- [47] Aman Sinha, Hongseok Namkoong, and John Duchi. Certifying some distributional robustness with principled adversarial training. In *International Conference on Learning Representations*, 2018. 2
- [48] Yang Song, Taesup Kim, Sebastian Nowozin, Stefano Ermon, and Nate Kushman. Pixeldefend: Leveraging generative models to understand and defend against adversarial examples. In *International Conference on Learning Representations*, 2018. 2
- [49] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Computer Vision and Pattern Recognition*, 2016. 5
- [50] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations*, 2014. 1, 2
- [51] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Dan Boneh, and Patrick McDaniel. Ensemble adversarial training: Attacks and defenses. In *International Conference on Learning Representations*, 2018. 2, 5
- [52] Jonathan Uesato, Brendan O’Donoghue, Aaron van den Oord, and Pushmeet Kohli. Adversarial risk and the dangers of evaluating against weak attacks. In *International Conference on Machine Learning*, 2018. 2, 6
- [53] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, and Alan Yuille. Mitigating adversarial effects through randomization. In *International Conference on Learning Representations*, 2018. 2
- [54] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Yuyin Zhou, Lingxi Xie, and Alan Yuille. Adversarial Examples for Semantic Segmentation and Object Detection. In *International Conference on Computer Vision*, 2017. 2
- [55] Cihang Xie, Yuxin Wu, Laurens van der Maaten, Alan Yuille, and Kaiming He. Feature denoising for improving adversarial robustness. *Computer Vision and Pattern Recognition*, 2019. 2, 6, 8
- [56] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *British Machine Vision Conference*, 2016. 6
- [57] Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric P Xing, Laurent El Ghaoui, and Michael I Jordan. Theoretically principled trade-off between robustness and accuracy. In *International Conference on Machine Learning*, 2019. 2