

Adversarial Defense by Restricting the Hidden Space of Deep Neural Networks

Aamir Mustafa^{1,3} Salman Khan^{1,2} Munawar Hayat^{1,3}

Roland Goecke³ Jianbing Shen^{1,4} Ling Shao¹

¹Inception Institute of Artificial Intelligence, ²Australian National University,

³University of Canberra, ⁴Beijing Institute of Technology

Abstract

Deep neural networks are vulnerable to adversarial attacks, which can fool them by adding minuscule perturbations to the input images. The robustness of existing defenses suffers greatly under white-box attack settings, where an adversary has full knowledge about the network and can iterate several times to find strong perturbations. We observe that the main reason for the existence of such perturbations is the close proximity of different class samples in the learned feature space. This allows model decisions to be totally changed by adding an imperceptible perturbation in the inputs. To counter this, we propose to class-wise disentangle the intermediate feature representations of deep networks. Specifically, we force the features for each class to lie inside a convex polytope that is maximally separated from the polytopes of other classes. In this manner, the network is forced to learn distinct and distant decision regions for each class. We observe that this simple constraint on the features greatly enhances the robustness of learned models, even against the strongest white-box attacks, without degrading the classification performance on clean images. We report extensive evaluations in both black-box and white-box attack scenarios and show significant gains in comparison to state-of-the-art defenses¹.

1. Introduction

Adversarial examples contain small, human-imperceptible perturbations specifically designed by an adversary to fool a learned model [37, 10]. These examples pose a serious threat for security critical applications, e.g. autonomous cars [1], bio-metric identification [34] and surveillance systems [28]. Furthermore, if a slight perturbation added to a benign input drastically changes the deep network’s output with a high-confidence, it reflects that our current models are not distinctively learning the fundamental visual concepts. Therefore, the design of robust deep networks goes a long way towards developing reliable and trustworthy artificial intelligence systems.

To mitigate adversarial attacks, various defense methods have recently been proposed. These can be broadly

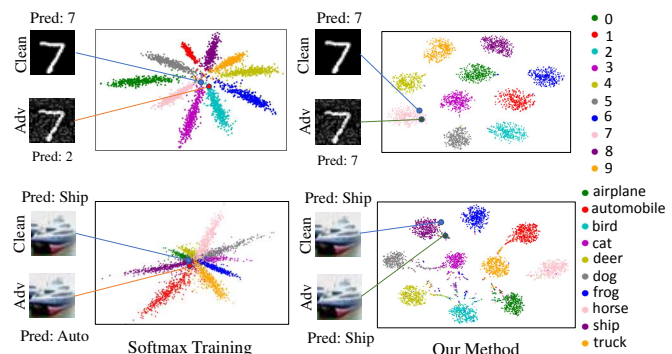


Figure 1: 2D penultimate layer activations of a clean image and its adversarial counterpart (PGD attack) for standard softmax trained model and our method on MNIST (top row) and CIFAR-10 (bottom row) datasets. Note that our method correctly maps the attacked image to its true-class feature space.

classified into two categories: (a) *Reactive defenses* that modify the inputs during testing time, using image transformations to counter the effect of adversarial perturbation [21, 5, 41, 26], and (b) *Proactive defenses* that alter the underlying architecture or learning procedure e.g. by adding more layers, ensemble/adversarial training or changing the loss/activation functions [38, 16, 31, 4, 19, 29, 22, 14]. Proactive defenses are generally more valued, as they provide relatively better robustness against *white-box* attacks. Nevertheless, both proactive and reactive defenses are easily circumvented by the iterative *white-box* adversaries [2].

This paper introduces a new proactive defense based on a novel training procedure, which maximally separates the learned feature representations at multiple depth levels of the deep model. We note that the addition of perturbations in the input domain leads to a corresponding polytope in the high-dimensional manifold of the intermediate features and the output classification space. Based upon this observation, we propose to maximally separate the polytopes for different class samples, such that there is a minimal overlap between any two classes in the decision and intermediate feature space. This ensures that an adversary can no longer fool the network within a restricted perturbation budget. In other words, we build on the intuition that two different class samples, which are visually dissimilar in the input domain, must be mapped to different regions in the output space. There-

¹Code and models are available at: <https://github.com/aamir-mustafa/pcl-adversarial-defense>

fore, we must also enforce that their feature representations are well separated along the hierarchy of network layers. This is achieved by improving within-class proximities and enhancing between-class differences of the activation maps, along multiple levels of the deep model. As illustrated in Fig. 1, the penultimate layer features learnt by the proposed scheme are well separated and hard to penetrate compared with the easily attacked features learnt using standard loss without any deep supervision. As evidenced with empirical evaluations (Sec. 5), the proposed method provides an effective and robust defense by significantly outperforming current state-of-the-art defenses under both *white-box* and *black-box* settings. Also, we experimentally show that our method does not suffer from the obfuscated gradient problem, which is otherwise the case for most existing defenses.

Our approach provides strong evidence towards the notion that the adversarial perturbations exist not only due to the properties of data (*e.g.* high-dimensionality) and network architecture (*e.g.* non-linearity functions) but also are greatly influenced by the choice of objective functions used for optimization. The deeply supervised multi-layered loss based defense provides a significant boost in robustness under strictest attack conditions where the balance is shifted heavily towards the adversary. These include *white-box* attacks and iterative adversaries including the strongest first-order attacks (Projected Gradient Descent). We demonstrate the robustness of the proposed defense through extensive evaluations on five publicly available datasets and achieve a robustness of 46.7% and 36.1% against the strongest PGD attack ($\epsilon = 0.03$) for the CIFAR-10 and CIFAR-100 datasets, respectively. To the best of our knowledge, these are significantly higher levels of robustness against a broad range of strong adversarial attacks.

2. Related Work

Generating adversarial examples to fool a deep network and developing defenses against such examples have gained significant research attention recently. Adversarial perturbations were first proposed by Szegedy *et al.* [37] using an L-BFGS based optimization scheme, followed by Fast Gradient Sign Method (FGSM) [10] and its iterative variant [16]. Moosavi-Dezfooli *et al.* [25] then proposed DeepFool, which iteratively projects an image across the decision boundary (form of a polyhedron) until it crosses the boundary and is mis-classified. One of the strongest attacks proposed recently is the Projected Gradient Descent (PGD) [22], which takes maximum loss increments allowed within a specified l_∞ norm-ball. Other popular attacks include the Carlini and Wagner Attack [3], Jacobian-based Saliency Map Approach [30], Momentum Iterative Attack [8] and Diverse Input Iterative Attack [42].

Two main lines of defense mechanisms have been proposed in the literature to counter adversarial attacks. First,

by applying different pre-processing steps and transformations on the input image at inference time [41, 11]. The second category of defenses improve network’s training regime to counter adversarial attacks. An effective scheme in this regards is *adversarial training*, where the model is jointly trained with clean images and their adversarial counterparts [17, 10]. Ensemble adversarial training is used in [38] to soften the classifier’s decision boundaries. Virtual Adversarial Training [24] smoothes the model distribution using a regularization term. Papernot *et al.* [31] used distillation to improve the model’s robustness by retraining with soft labels. Parseval Networks [4] restrict the Lipschitz constant of each layer of the model. Input Gradient Regularizer [33] penalizes the change in model’s prediction w.r.t input perturbations by regularizing the gradient of cross-entropy loss. The Frobenius norm of the Jacobian of the network has been shown to improve model’s stability in [13]. [20] proposed defensive quantization method to control the Lipschitz constant of the network to mitigate the adversarial noise during inference. [7] proposed Stochastic Activation Pruning as a defense against adversarial attacks. Currently the strongest defense method is Min-Max optimization [22] which augments the training data with a first order attacked samples. Despite significant research activity in devising defenses against adversarial attacks, it was recently shown in [2] that the currently existing state-of-the-art defenses [15, 32, 35] are successfully circumvented under *white-box* settings. Only Min-Max optimization [22] and Cascade adversarial machine learning [27] retained 47% and 15% accuracy respectively, and withstood the attacks under *white-box* settings. In our experiments (see Sec. 5), we extensively compare our results with [22] and make a compelling case by achieving significant improvements.

At the core of our defense are the proposed objective function and multi-level deep supervision, which ensure feature space discrimination between classes. Our training objective is inspired from center loss [40], which clusters penultimate layer features. We propose multiple novel constraints (Sec. 3) to enhance between-class distances, and ensure maximal separation of a sample from its non-true classes. Our method is therefore fundamentally different from [40], since the proposed multi-layered hierarchical loss formulation and the notion of maximal separation has not been previously explored for adversarial robustness.

3. Prototype Conformity Loss

Below, we first introduce the notations used, then provide a brief overview of the conventional cross entropy loss followed by a detailed description of our proposed method.

Notations: Let $\mathbf{x} \in \mathbb{R}^m$ and \mathbf{y} denote an input-label pair and $\mathbf{1}_y$ be the one-hot encoding of \mathbf{y} . We denote a deep neural network (DNN) as a function $\mathcal{F}_\theta(\mathbf{x})$, where θ are the trainable parameters. The DNN outputs a feature represen-

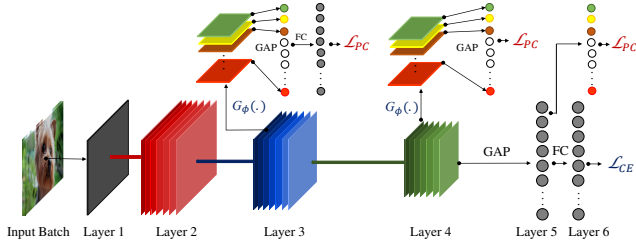


Figure 2: An illustration of our training with joint supervision of \mathcal{L}_{PC} and \mathcal{L}_{CE} . $G_\phi(\cdot)$ is an auxiliary branch to map features to a low dimensional output, which is then used for loss in Eq. 8

tation $\mathbf{f} \in \mathbb{R}^d$, which is then used by a classification layer to perform multi-class classification. Let k be the number of classes; the parameters of the classifier can then be represented as $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_k] \in \mathbb{R}^{d \times k}$. To train the model, we find the optimal θ and \mathbf{W} that minimize a given objective function. Next, we introduce a popular loss function for deep CNNs.

Cross-entropy Objective: The cross-entropy objective function maximizes the dot product between an input feature \mathbf{f}_i and its true class representative vector \mathbf{w}_y , such that $\mathbf{w}_y \in \mathbf{W}$. In other words, cross-entropy loss forces the classifier to learn a mapping from feature to output space such that the projection on to the correct class vector is maximized:

$$\mathcal{L}_{CE}(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^m -\log \frac{\exp(\mathbf{w}_{y_i}^T \mathbf{f}_i + \mathbf{b}_{y_i})}{\sum_{j=1}^k \exp(\mathbf{w}_j^T \mathbf{f}_i + \mathbf{b}_j)}, \quad (1)$$

where, m is the number of images, and \mathbf{f}_i is the feature of i^{th} image \mathbf{x}_i with the class \mathbf{y}_i . \mathbf{W} and \mathbf{b} are, respectively, the weights and the bias terms for the classification layer.

Adversarial Perspective: The main goal of an attack algorithm is to force a trained DNN \mathcal{F}_θ to make wrong predictions. Attack algorithms seek to achieve this goal within a minimal perturbation budget. The attacker’s objective can be represented by:

$$\arg \max_{\delta} \mathcal{L}(\mathbf{x} + \delta, \mathbf{y}), \quad s.t., \|\delta\|_p < \epsilon, \quad (2)$$

where \mathbf{y} is the ground-truth label for an input sample \mathbf{x} , δ denotes the adversarial perturbation, $\mathcal{L}(\cdot)$ denotes the error function, $\|\cdot\|_p$ denotes the p-norm, which is generally considered to be an ℓ_∞ -ball centered at \mathbf{x} , and ϵ is the available perturbation budget.

In order to create a robust model, the learning algorithm must consider the allowed perturbations in the input domain and learn a function that maps the perturbed images to the correct class. This can be achieved through the following min-max (saddle point) objective that minimizes the empirical risk in the presence of perturbations:

$$\min_{\theta} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} \left[\max_{\delta} \mathcal{L}(\mathbf{x} + \delta, \mathbf{y}; \theta) \right], \quad s.t., \|\delta\|_p < \epsilon, \quad (3)$$

where \mathcal{D} is the data distribution.

CE Loss in Adversarial Setting: The CE loss is the default choice for conventional classification tasks. However, it simply assigns an input sample to one of the pre-defined classes. It therefore does not allow one to distinguish between normal and abnormal inputs (adversarial perturbations in our case). Further, it does not explicitly enforce any margin constraints amongst the learned classification regions. It can be seen from Eq. 3 that an adversary’s job is to maximize $\mathcal{L}(\cdot)$ within a small perturbation budget ϵ . Suppose, the adversarial polytope in the output space² with respect to an input sample \mathbf{x} is given by:

$$\mathcal{P}_\epsilon(\mathbf{x}; \theta) = \{\mathcal{F}_\theta(\mathbf{x} + \delta) \text{ s.t.}, \|\delta\|_p \leq \epsilon\}. \quad (4)$$

An adversary’s task is easier if there is an overlap between the adversarial polytopes for different input samples belonging to different classes.

Definition 1: The overlap $\mathcal{O}_\epsilon^{i,j}$ between polytopes for each data sample pair (i, j) can be defined as the volume of intersection between the respective polytopes:

$$\mathcal{O}_\epsilon^{i,j} = \mathcal{P}_\epsilon(\mathbf{x}_{y_i}^i; \theta) \cap \mathcal{P}_\epsilon(\mathbf{x}_{y_j}^j; \theta).$$

Note that the considered polytopes can be non-convex as well. However, the overlap computation can be simplified for convex polytopes [6].

Proposition 1: For an i^{th} input sample $\mathbf{x}_{y_i}^i$ with class label \mathbf{y}_i , reducing the overlap $\mathcal{O}_\epsilon^{i,j}$ between its polytope $\mathcal{P}_\epsilon(\mathbf{x}_{y_i}^i; \theta)$ and the polytopes of other class samples $\mathcal{P}_\epsilon(\mathbf{x}_{y_j}^j; \theta)$, s.t., $\mathbf{y}_j \neq \mathbf{y}_i$ will result in lower adversary success for a bounded perturbation $\|\delta\|_p \leq \epsilon$.

Proposition 2: For a given adversarial strength ϵ , assume λ is the maximum distance from the center of the polytope to the convex outer bounded polytope. Then, a classifier maintaining a margin $m > 2\lambda$ between two closest samples belonging to different classes will result in a decision boundary with guaranteed robustness against perturbation within the budget ϵ .

In other words, if the adversarial polytopes for samples belonging to different classes are non-overlapping, the adversary cannot find a viable perturbation within the allowed budget. We propose that an adversary’s task can be made difficult by including a simple maximal separation constraint in the objective of deep networks. The conventional CE loss does not impose any such constraint, which makes the resulting models weaker against adversaries. A more principled approach is to define convex category-specific classification regions for each class, where any sample outside all of such regions is considered an adversarial perturbation. Consequently, we propose the prototype

²Note that the output space in our case is not the final prediction space, but the intermediate feature space.

conformity loss function, described below.

Proposed Objective: We represent each class with its prototype vector, which represents the training examples of that class. Each class is assigned a fixed and non-overlapping p-norm ball and the training samples belonging to a class i are encouraged to be mapped close to its hyper-ball center:

$$\mathcal{L}_{PC}(\mathbf{x}, \mathbf{y}) = \sum_i \left\{ \|\mathbf{f}_i - \mathbf{w}_{y_i}^c\|_2 - \frac{1}{k-1} \sum_{j \neq y_i} \left(\|\mathbf{f}_i - \mathbf{w}_j^c\|_2 + \|\mathbf{w}_{y_i}^c - \mathbf{w}_j^c\|_2 \right) \right\}. \quad (5)$$

During model inference, a feature’s similarity is computed with all the class prototypes and it is assigned the closest class label if and only if the sample lies within its decision region:

$$\hat{y}_i = \underset{j}{\operatorname{argmin}} \|\mathbf{f}_i - \mathbf{w}_j^c\|. \quad (6)$$

Here, \mathbf{w}^c denotes the trainable class centroids. Note that the classification rule is similar to the Nearest Class Mean (NCM) classifier [23], but we differ in some important aspects: (a) the centroids for each class are not fixed as the mean of training samples, rather learned automatically during representation learning, (b) class samples are explicitly forced to lie within respective class norm-balls, (c) feature representations are appropriately tuned to learn discriminant mappings in an end-to-end manner, and (d) to avoid inter-class confusions, disjoint classification regions are considered by maintaining a large distance between each pair of prototypes. We also experiment with the standard softmax classifier and get equivalent performance compared to nearest prototype rule mentioned above.

Deeply Supervised Learning: The overall loss function used for training our model is given by:

$$\mathcal{L}(\mathbf{x}, \mathbf{y}) = \mathcal{L}_{CE}(\mathbf{x}, \mathbf{y}) + \mathcal{L}_{PC}(\mathbf{x}, \mathbf{y}). \quad (7)$$

The above loss enforces the intra-class compactness and an inter-class separation using learned prototypes in the output space. In order to achieve a similar effect in the intermediate feature representations, we include other auxiliary loss functions $\{\mathcal{L}^n\}$ along the depth of our deep networks, which act as companion objective functions for the final loss. This is achieved by adding an auxiliary branch $\mathcal{G}_\phi(\cdot)$ after the defined network depth, which maps the features to a lower dimension output, and is then used in the loss definition. For illustration, see Fig. 2.

$$\mathcal{L}^n(\mathbf{x}, \mathbf{y}) = \mathcal{L}_{CE}(\mathbf{f}^l, \mathbf{y}) + \mathcal{L}_{PC}(\mathbf{f}^l, \mathbf{y}) \quad (8)$$

$$s.t., \mathbf{f}^l = \mathcal{G}_\phi^l(\mathcal{F}_\theta^l(\mathbf{x})). \quad (9)$$

These functions avoid the vanishing gradients problem and act as regularizers that encourage features belonging to the same class to come together and the ones belonging to different classes to be pushed apart.

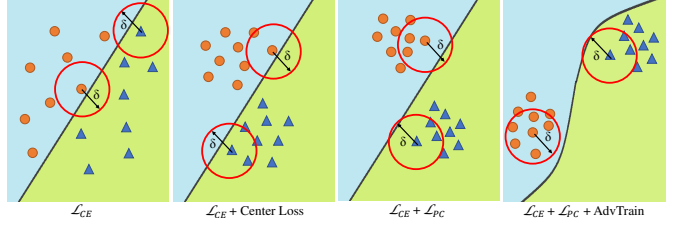


Figure 3: Comparison between different training methods. The red circle encompasses the adversarial sample space within a perturbation budget $\|\delta\|_p < \epsilon$.

4. Adversarial Attacks

We evaluate our defense model against five recently proposed state-of-the-art attacks, which are summarized below, for completeness.

Fast Gradient Sign Method (FGSM) [10] generates an adversarial sample \mathbf{x}_{adv} from a clean sample \mathbf{x} by maximizing the loss in Eq. 2. It finds \mathbf{x}_{adv} by moving a single step in the opposite direction to the gradient of the loss function, as:

$$\mathbf{x}_{adv} = \mathbf{x} + \epsilon \cdot \operatorname{sign}(\nabla_x \mathcal{L}(\mathbf{x}, \mathbf{y})). \quad (10)$$

Here, ϵ is the allowed perturbation budget.

Basic Iterative Method (BIM) [16] is an iterative variant of FGSM and generates an adversarial sample as:

$$\mathbf{x}_m = \operatorname{clip}_\epsilon(\mathbf{x}_{m-1} + \frac{\epsilon}{i} \cdot \operatorname{sign}(\nabla_{\mathbf{x}_{m-1}}(\mathcal{L}(\mathbf{x}_{m-1}, \mathbf{y}))), \quad (11)$$

where \mathbf{x}_0 is clean image \mathbf{x} and i is the iteration number.

Momentum Iterative Method (MIM) [8] introduces an additional momentum term to BIM to stabilize the direction of gradient. Eq. 11 is modified as:

$$g_m = \mu \cdot g_{m-1} + \frac{\nabla_{\mathbf{x}_{m-1}} \mathcal{L}(\mathbf{x}_{m-1}, \mathbf{y})}{\|\nabla_{\mathbf{x}_{m-1}}(\mathcal{L}(\mathbf{x}_{m-1}, \mathbf{y}))\|_1} \quad (12)$$

$$\mathbf{x}_m = \operatorname{clip}_\epsilon(\mathbf{x}_{m-1} + \frac{\epsilon}{i} \cdot \operatorname{sign}(g_m)), \quad (13)$$

where μ is the decay factor.

Carlini & Wagner Attack [3] defines an auxiliary variable ζ and minimizes the objective function:

$$\min_\zeta \left\| \frac{1}{2}(\tanh(\zeta) + 1) - \mathbf{x} \right\| + c \cdot f\left(\frac{1}{2}(\tanh(\zeta) + 1)\right), \quad (14)$$

where $\frac{1}{2}(\tanh(\zeta) + 1) - \mathbf{x}$ is the perturbation δ , c is the constant chosen and $f(\cdot)$ is defined as:

$$f(\mathbf{x}_{adv}) = \max(\mathcal{Z}(\mathbf{x}_{adv})_{\mathbf{y}} - \max\{\mathcal{Z}(\mathbf{x}_{adv})_k : k \neq \mathbf{y}\}, -\kappa). \quad (15)$$

Here, κ controls the adversarial sample’s confidence and $\mathcal{Z}(\mathbf{x}_{adv})_k$ are the logits values corresponding to a class k .

Projected Gradient Descent (PGD) [22] is similar to BIM, and starts from a random position in the clean image neighborhood $\mathcal{U}(\mathbf{x}, \epsilon)$. This method applies FGSM for m iterations with a step size of γ as:

$$\mathbf{x}_m = \mathbf{x}_{m-1} + \gamma \cdot \text{sign}(\nabla_{\mathbf{x}_{m-1}} \mathcal{L}(\mathbf{x}_{m-1}, \mathbf{y})). \quad (16)$$

$$\mathbf{x}_m = \text{clip}(\mathbf{x}_m, \mathbf{x}_m - \epsilon, \mathbf{x}_m + \epsilon). \quad (17)$$

It proves to be a strong iterative attack, relying on the first order information of the target model.

5. Experiments

Datasets and Models: We extensively evaluate the proposed method on five datasets: MNIST, Fashion-MNIST (F-MNIST), CIFAR-10, CIFAR-100 and Street-View House Numbers (SVHN). For the MNIST and F-MNIST datasets, the CNN model chosen has six layers, as in [40]. For the CIFAR-10, CIFAR-100 and SVHN datasets, we use a ResNet-110 model [12] (see Table 1). The deep features for the prototype conformity loss are extracted from different intermediate layers using an auxiliary branch, which maps the features to a lower dimension output (see Fig. 2). We first train for T' epochs ($T' = 50$ for F/MNIST, $T' = 200$ for CIFAR-10/100 and SVHN) with \mathcal{L}_{CE} and then use the loss in Eq. 8 for 300 epochs. A batch size of 256 and a learning rate of 0.1 ($\times 0.1$ at $T=200, 250$) are used. Further training details are summarized in Algorithm 1.

Algorithm 1: Model training with Prototype Conformity Loss.

Input: Classifier $\mathcal{F}_\theta(\mathbf{x})$, training data $\{\mathbf{x}\}$, ground truth labels $\{\mathbf{y}\}$, trainable parameters θ , trainable class centroids $\{\mathbf{w}_j^c : j \in [1, k]\}$, perturbation budget ϵ , epochs T , number of auxiliary branches L .

Output: Updated parameters θ

```

1 Initialize  $\theta$  in convolution layers.
2 for  $t = 0$  to  $T$ :
3   if  $t < T'$ :
4     Converge softmax objective,  $\theta := \arg \min_{\theta} \mathcal{L}_{\text{CE}}$ .
5   else:
6     Compute joint loss  $\mathcal{L} = \mathcal{L}_{\text{CE}} + \sum_l^L \mathcal{L}_{\text{PC}}$ 
7     Compute gradients w.r.t  $\theta$  and  $\mathbf{x}$ ,  $\nabla_{\theta} \mathcal{L}(\mathbf{x}, \mathbf{y})$  and
       $\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \mathbf{y})$  respectively.
8     Update model weights,  $\theta := \arg \min_{\theta} \mathcal{L}$ .
9     Update class centroids  $\mathbf{w}_j^c \forall j$ 
10    Generate adversarial examples as:
11      if FGSM: then  $\mathbf{x}_{adv} = \mathbf{x} + \epsilon \cdot \text{sign}(\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \mathbf{y}))$ 
12      elif PGD: then  $\mathbf{x}_{adv} = \text{clip}(\mathbf{x}, \mathbf{x} - \epsilon, \mathbf{x} + \epsilon)$ 
13    Augment  $\mathbf{x}$  with  $\mathbf{x}_{adv}$ 
14 return  $\theta$ 

```

5.1. Results and Analysis

White-Box vs Black-Box Settings: In an adversarial setting, there are two main threat models: *white-box* attacks where the adversary possesses complete knowledge of the target model, including its parameters, architecture and the training method, and *black-box attacks* where the adversary feeds perturbed images at test time (which are generated without any knowledge of the target model). We evaluate the robustness of our proposed defense against both *white-box* and *black-box* settings. Table 2 shows our results for

Table 1: Two network architectures: CNN-6 (MNIST, FMNIST) and ResNet-110 (CIFAR-10,100 and SVHN). Features are extracted in CNN-6 (after Layer 3 and two FC layers) and ResNet-110 (after Layer 3, 4 and FC layer) to impose the proposed \mathcal{L}_{PC} . Auxiliary branches are shown in green color.

Layer #	6-Conv Model	ResNet-110
1	$\left[\begin{array}{c} \text{Conv}(32, 5 \times 5) \\ \text{PReLU}(2 \times 2) \end{array} \right] \times 2$	$\begin{array}{c} \text{Conv}(16, 3 \times 3) + \text{BN} \\ \text{ReLU}(2 \times 2) \end{array}$
2	$\left[\begin{array}{c} \text{Conv}(64, 5 \times 5) \\ \text{PReLU}(2 \times 2) \end{array} \right] \times 2$	$\left[\begin{array}{c} \text{Conv}(16, 1 \times 1) + \text{BN} \\ \text{Conv}(16, 3 \times 3) + \text{BN} \\ \text{Conv}(64, 1 \times 1) + \text{BN} \end{array} \right] \times 12$
3	$\left[\begin{array}{c} \text{Conv}(128, 5 \times 5) \\ \text{PReLU}(2 \times 2) \end{array} \right] \times 2$	$\left[\begin{array}{c} \text{Conv}(32, 1 \times 1) + \text{BN} \\ \text{Conv}(32, 3 \times 3) + \text{BN} \\ \text{Conv}(128, 1 \times 1) + \text{BN} \end{array} \right] \times 12$
	$\text{GAP} \rightarrow \mathcal{L}_{\text{PC}}$	$(\text{GAP} \rightarrow \text{FC}(512) \rightarrow \mathcal{L}_{\text{PC}})$
4	$\text{FC}(512) \rightarrow \mathcal{L}_{\text{PC}}$	$\left[\begin{array}{c} \text{Conv}(64, 1 \times 1) + \text{BN} \\ \text{Conv}(64, 3 \times 3) + \text{BN} \\ \text{Conv}(256, 1 \times 1) + \text{BN} \end{array} \right] \times 12$
		$\text{GAP} \rightarrow \mathcal{L}_{\text{PC}}$
5	$\text{FC}(64) \rightarrow \mathcal{L}_{\text{PC}}$	$\text{FC}(1024) \rightarrow \mathcal{L}_{\text{PC}}$
6	$\text{FC}(10) \rightarrow \mathcal{L}_{\text{CE}}$	$\text{FC}(100/10) \rightarrow \mathcal{L}_{\text{CE}}$

the different attacks described in Sec. 4. The number of iterations for BIM, MIM and PGD are set to 10 with a step size of $\epsilon/10$. The iteration steps for C&W are 1,000, with a learning rate of 0.01. We report our model’s robustness with and without adversarial training for standard perturbation size *i.e.* $\epsilon = 0.3$ for F/MNIST and $\epsilon = 0.03$ for the CIFAR-10/100 and SVHN datasets.

Recent literature has shown transferability amongst deep models [39, 17, 10], where adversarial images are effective even for the models they were never generated on. An adversary can therefore exploit this characteristic of deep models and generate generic adversarial samples to attack unseen models. Defense against *black-box* attacks is therefore highly desirable for secure deployment of machine learning models [30]. To demonstrate the effectiveness of our proposed defense under *black-box* settings, we generate adversarial samples using a VGG-19 model, and feed them to the model trained using our proposed strategy. Results in Table 2 show that *black-box* settings have negligible attack potential against our model. For example, on the CIFAR-10 dataset, where our model’s accuracy on clean images is 91.89%, even the strongest iterative attack (PGD-0.03) fails, and our defense retains an accuracy of 88.8%.

Adversarial Training has been shown to enhance many recently proposed defense methods [18]. We also evaluate the impact of adversarial training (AdvTrain) in conjunction with our proposed defense. For this, we jointly train our model on clean and attacked samples, which are generated using FGSM [10] and PGD [22] by uniformly sampling ϵ from an interval of [0.1, 0.5] for MNIST and F-MNIST and [0.01, 0.05] for CIFAR and SVHN. Results in Table 2 indicate that AdvTrain further complements our method and

Table 2: Robustness of our model in *white-box* and *black-box* settings. Adversarial samples generated in the *black-box* settings show negligible attack potential against our models. Here ϵ is the perturbation size and c is the initial constant for C&W attack. It can be seen that AdvTrain further complements the robustness of our models.

Training	No Attack	White-Box Setting					Black-Box Setting				
		FGSM	BIM	C&W	MIM	PGD	FGSM	BIM	C&W	MIM	PGD
MNIST ($\epsilon = 0.3, c = 10$)											
Softmax	98.71	4.9	0.0	0.2	0.01	0.0	23.0	17.8	20.9	14.8	11.9
Ours	99.53	31.1	23.3	29.1	24.7	19.9	78.3	72.7	77.2	74.5	69.5
Ours + AdvTrain _{FGSM}	99.44	53.1	36.6	40.9	37.0	34.5	85.6	81.0	82.3	81.4	78.2
Ours + AdvTrain _{PGD}	99.28	49.8	40.3	46.0	41.4	39.8	85.2	81.9	83.5	82.8	80.8
CIFAR-10 ($\epsilon = 0.03, c = 0.1$)											
Softmax	90.80	21.4	0.0	0.6	0.0	0.01	39.0	30.1	31.8	30.9	29.1
Ours	90.45	67.7	32.6	37.3	33.2	27.2	85.5	83.7	83.3	81.9	76.4
Ours + AdvTrain _{FGSM}	91.28	75.8	45.9	45.7	44.7	42.5	88.9	87.6	87.4	88.2	84.5
Ours + AdvTrain _{PGD}	91.89	74.9	46.0	51.8	49.3	46.7	88.5	88.3	88.2	88.5	88.8
CIFAR-100 ($\epsilon = 0.03, c = 0.1$)											
Softmax	72.65	20.0	4.2	1.1	3.52	0.17	40.9	34.3	37.1	35.5	30.7
Ours	71.90	56.9	28.0	31.1	28.7	25.9	65.3	64.5	64.1	64.8	62.8
Ours + AdvTrain _{FGSM}	69.11	61.3	32.3	35.2	33.3	31.4	66.1	65.2	65.7	65.5	63.4
Ours + AdvTrain _{PGD}	68.32	60.9	34.1	36.7	33.7	36.1	65.9	66.1	66.7	66.1	66.7
F-MNIST ($\epsilon = 0.3, c = 10$)											
Softmax	91.51	8.7	0.1	0.2	0.0	0.0	46.7	29.3	30.8	29.5	26.0
Ours	91.32	29.0	22.0	23.9	21.8	20.3	84.8	79.0	79.2	78.4	76.3
Ours + AdvTrain _{FGSM}	91.03	55.1	37.5	41.7	40.6	35.3	89.1	87.0	87.7	87.9	85.2
Ours + AdvTrain _{PGD}	91.30	47.2	40.1	44.6	41.3	40.7	88.2	88.0	88.2	88.3	89.7
SVHN ($\epsilon = 0.03, c = 0.1$)											
Softmax	93.45	30.6	6.2	7.1	7.3	9.6	48.1	30.3	31.4	33.5	21.5
Ours	94.36	69.3	37.1	39.2	41.0	33.7	77.4	73.1	76.4	74.0	70.1
Ours + AdvTrain _{FGSM}	94.18	80.1	47.4	51.9	45.6	40.5	90.1	87.4	88.0	87.6	84.4
Ours + AdvTrain _{PGD}	94.36	76.5	48.8	54.8	47.1	47.7	88.7	88.2	89.2	88.6	89.3

provides an enhanced robustness under both *black-box* and *white-box* attack settings.

Adaptive White-box Settings: Since at inference time, our model performs conventional softmax prediction, we evaluated the robustness against standard *white-box* attack settings, to be consistent with existing defenses. Now, we also experiment in an *adaptive white-box* setting where the attack is performed on the joint PC+CE loss (with access to learned prototypes). Negligible performance drop is observed in adaptive settings (see Table 3).

Table 3: Robustness in *adaptive white-box* attack settings. The performance for conventional attacks (where CE is the adversarial loss) is shown in blue. * indicates adversarially trained models.

Training	No Attack	FGSM	BIM	MIM	PGD
CIFAR-10 ($\epsilon = 8/255$)					
Ours	90.45	66.90 (67.7)	31.29 (32.6)	32.84 (33.2)	27.09 (27.2)
Ours* _{FGSM}	91.28	74.24 (75.8)	44.05 (45.9)	43.77 (44.7)	41.32 (42.5)
Ours* _{PGD}	91.89	74.31 (74.9)	44.85 (46.0)	47.31 (49.3)	44.75 (46.7)
F-MNIST ($\epsilon = 0.3/1$)					
Ours	91.32	28.1 (29.0)	21.7 (22.0)	20.3 (21.8)	19.5 (20.3)
Ours* _{FGSM}	91.03	53.3 (55.1)	36.0 (37.5)	39.3 (40.6)	34.7 (35.3)
Ours* _{PGD}	91.30	46.0 (47.2)	40.1 (40.1)	40.7 (41.3)	39.7 (40.7)

5.2. Comparison with Existing Defenses

We compare our method with recently proposed state-of-the-art proactive defense mechanisms, which alter the network or use modified training loss functions. To this end, we compare with [17], which injects adversarial examples into the training set and generates new samples at each iteration. We also compare with [29], which introduces an Adaptive Diversity Promoting (ADP) regularizer to improve adversarial robustness. Further, we compare with an

input gradient regularizer mechanism [33] that penalizes the degree to which input perturbations can change a model’s predictions by regularizing the gradient of the cross-entropy loss. Finally, we compare with the current state-of-the-art Min-Max optimization based defense [22], which augments the training data with adversarial examples, causing the maximum gradient increments to the loss within a specified l_∞ norm. The results in Tables 8, 9 and 4 in terms of retained classification accuracy on different datasets show that our method significantly outperforms all existing defense schemes by a large margin. The performance gain is more pronounced for the strongest iterative attacks (e.g. C&W and PGD) with large perturbation budget ϵ . For example, our method achieves a relative gain of 20.6% (AdvTrain models) and 41.4% (without AdvTrain) compared to the 2nd best methods on the CIFAR-10 and MNIST datasets respectively for the PGD attack. On CIFAR-100 dataset, for the strongest PGD attack with $\epsilon = 0.01$, the proposed method achieves 38.9% compared with 18.3% by ADP [29], which, to the best of our knowledge, is the only method in the literature evaluated on the CIFAR-100 dataset. Our results further indicate that adversarial training consistently complements our method and augments its performance across all evaluated datasets.

Additionally we compare our model’s performance with a close method proposed by Song *et al.* [36] in Table 5, where our approach outperforms them by a significant margin. Besides a clear improvement, we discuss our main distinguishing features below: **(a)** Our approach is based on a “deeply-supervised” loss that prevents changes to the out-

Table 4: Comparison on **CIFAR-100** dataset for *white-box* adversarial attacks (numbers shows robustness, higher is better). * sign denotes adversarially trained models. For our model, we report results without adversarial training (Ours) and with adversarially generated images from FGSM (Ours_f^{*}) and PGD (Ours_p^{*}) attacks.

Attacks	Params.	Baseline	ADP [29]	Ours	Ours _f [*]	Ours _p [*]
No Attack	-	72.6	70.2	71.9	69.1	68.3
BIM	$\epsilon = 0.005$	21.6	26.2	44.8	55.1	55.7
	$\epsilon = 0.01$	10.1	14.8	39.8	46.2	46.9
MIM	$\epsilon = 0.005$	24.2	29.4	46.1	56.7	57.1
	$\epsilon = 0.01$	11.2	17.2	40.6	43.8	45.9
PGD	$\epsilon = 0.005$	26.6	32.1	42.2	53.6	55.0
	$\epsilon = 0.01$	11.7	18.3	38.9	40.1	44.0

puts within the limited perturbation budget. Such supervision paradigm is the main contributing factor towards our improved results (See Table 7). (b) [36] focuses on domain adaption between adversarial and natural samples without any constraint on the intermediate feature representations. In contrast, we explicitly enforce the hidden layer activations to be maximally separated in our network design. (c) [36] only considers adversarially trained models, while we demonstrate clear improvements both with and without adversarial training (a more challenging setting). In Table 5 we have followed the exact model settings used in [36].

Dataset	Clean	FGSM	MIM	PGD
F-MNIST ($\epsilon = 0.1$)	85.5	78.2	68.8	68.6
	91.3	86.6	80.1	79.4
SVHN ($\epsilon = 0.02$)	82.9	57.2	53.9	53.2
	94.4	87.1	82.2	80.7
CIFAR-10 ($\epsilon = 4/255$)	84.8	60.7	59.0	58.1
	91.9	85.3	70.1	69.4
CIFAR-100 ($\epsilon = 4/255$)	61.6	29.3	27.3	26.2
	71.9	49.1	40.7	38.6

Table 5: Comparison of our approach with [36] on 4 datasets. Green rows show results for [36] and blue for our models.

5.3. Transferability Test

We investigate the transferability of attacks on CIFAR-10 dataset between a standard VGG-19 model, adversarially trained VGG-19 [17], Madry *et al.*'s [22] and our model. We report the accuracy of target models (columns) on adversarial samples generated from source models (rows) in Table 6. Our results yield the following findings:

Improved *black-box* robustness: As noted in [2], a model that gives a false sense of security due to obfuscated gradients can be identified if the *black-box* attacks are stronger than *white-box*. In other words, robustness of such a model under *white-box* settings is higher than under *black-box* settings. It was shown in [2] that most of the existing defenses obfuscate gradients. Madry *et al.*'s approach [22] was endorsed by [2] to not cause gradient masking. The comparison in Table 6 shows that our method outperforms [22].

Similar architectures increase transferability: Changing the source and target network architectures decreases the transferability of an attack. The same architectures (*e.g.* VGG-19 and its AdvTrain counterpart, as in Table 6) show increased robustness against *black-box* attacks generated from each other.

Table 6: Transferability Test on CIFAR-10: PGD adversaries are generated with $\epsilon = 0.03$, using the source network, and then evaluated on target model. Underline denotes robustness against *white-box* attack. Note that adversarial samples generated on our model are highly transferable to other models as *black-box* attacked images.

Source \ Target	VGG-19	AdvTrain [17]	Madry <i>et al.</i> [22]	Ours
VGG-19	<u>0.0</u>	16.20	52.71	88.80
AdvTrain [17]	12.43	<u>0.0</u>	49.80	72.53
Madry <i>et al.</i> [22]	58.91	67.32	<u>43.70</u>	71.72
Ours	50.31	61.02	66.70	<u>49.10</u>

5.4. Ablation Analysis

\mathcal{L}_{PC} at Different Layers: We investigate the impact of our proposed prototype conformity loss (\mathcal{L}_{PC}) at different depths of the network. Specifically, as shown in Table 7, we apply \mathcal{L}_{PC} individually after each layer (see Table 1 for architectures) and in different combinations. We report the achieved results on the CIFAR-10 dataset for clean and perturbed samples (using FGSM and PGD attacks) in Table 7. The network without any \mathcal{L}_{PC} loss is equivalent to a standard softmax trained model. It achieves good performance on clean images, but fails under both *white-box* and *black-box* attacks (see Table 2). The models with \mathcal{L}_{PC} loss at initial layers are unable to separate deep features class-wise, thereby resulting in inferior performance. Our proposed \mathcal{L}_{PC} loss has maximum impact in the deeper layers of the network. This justifies our choice of different layers for \mathcal{L}_{PC} loss, indicated in Table 1.

Table 7: **Ablation Analysis** with \mathcal{L}_{PC} applied at different layers of ResNet-110 (Table 1) for CIFAR-10 dataset.

Layer #	No Attack $\epsilon = 0$	FGSM $\epsilon = 0.03$	PGD $\epsilon = 0.03$
None	90.80	21.40	0.01
Layer 1	74.30	23.71	0.01
Layer 2	81.92	30.96	8.04
Layer 3	88.75	33.74	10.47
Layer 4	90.51	39.90	11.90
Layer 5	91.11	47.02	13.56
Layer 4+5	90.63	55.36	20.70
Layer 3+4+5	90.45	67.71	27.23

5.5. Identifying Obfuscated Gradients

Recently, Athalye *et al.* [2] were successful in breaking several defense mechanisms in the *white-box* settings by identifying that they exhibit a false sense of security. They call this phenomenon *gradient masking*. Below, we discuss how our defense mechanism does not cause gradient masking on the basis of characteristics defined in [2, 9].

Iterative attacks perform better than one-step attacks: Our evaluations in Fig. 4 indicate that stronger iterative attacks (*e.g.* BIM, MIM, PGD) in the *white-box* settings are more successful at attacking the defense models than single-step attacks (FGSM in our case).

Table 8: Comparison on **CIFAR-10** dataset for *white-box* adversarial attacks (numbers shows robustness, higher is better). * sign denotes adversarially trained models. For our model, we report results without adversarial training (Ours) and with adversarially generated images from FGSM (Ours_f^{*}) and PGD (Ours_p^{*}) attacks.

Attacks	Params.	Baseline	AdvTrain [17]*	Yu et al. [43]*	Ross et al. [33]*	Pang et al. [29]*	Madry et al. [22]*	Ours	Ours _f [*]	Ours _p [*]
No Attack	-	90.8	84.5	83.1	86.2	90.6	87.3	90.5	91.3	91.9
FGSM	$\epsilon = 0.02$	36.5	44.3	48.5	39.5	61.7	71.6	72.5	80.8	78.5
	$\epsilon = 0.04$	19.4	31.0	38.2	20.8	46.2	47.4	56.3	70.5	69.9
BIM	$\epsilon = 0.01$	18.5	22.6	62.7	19.0	46.6	64.3	62.9	67.9	74.5
	$\epsilon = 0.02$	6.1	7.8	39.3	6.9	31.0	49.3	40.1	51.2	57.3
MIM	$\epsilon = 0.01$	23.8	23.9	-	24.6	52.1	61.5	64.3	68.8	74.9
	$\epsilon = 0.02$	7.4	9.3	-	9.5	35.9	46.7	42.3	53.8	60.0
C&W	$c = 0.001$	61.3	67.7	82.5	72.2	80.6	84.5	84.3	91.0	91.3
	$c = 0.01$	35.2	40.9	62.9	47.8	54.9	65.7	63.5	72.9	73.7
	$c = 0.1$	0.6	25.4	40.7	19.9	25.6	47.9	41.1	55.7	60.5
PGD	$\epsilon = 0.01$	23.4	24.3	-	24.5	48.4	67.7	60.1	68.3	75.7
	$\epsilon = 0.02$	6.6	7.8	-	8.5	30.4	48.5	39.3	50.6	58.5

Table 9: Comparison on **MNIST** dataset for *white-box* adversarial attacks (numbers shows robustness, higher is better). * sign denotes adversarially trained models. For our model, we report results without adversarial training (Ours) and with adversarially generated images from FGSM (Ours_f^{*}) and PGD (Ours_p^{*}) attacks.

Attacks	Params.	Baseline	AdvTrain [17]*	Yu et al. [43]	Ross et al. [33]	Pang et al. [29]	Ours	Ours _f [*]	Ours _p [*]
No Attack	-	98.7	99.1	98.4	99.2	99.5	99.5	99.4	99.3
FGSM	$\epsilon = 0.1$	58.3	73.0	91.6	91.6	96.3	97.1	97.2	96.5
	$\epsilon = 0.2$	12.9	52.7	70.3	60.4	52.8	70.6	80.0	77.9
BIM	$\epsilon = 0.1$	22.5	62.0	88.1	87.9	88.5	90.2	92.0	92.1
	$\epsilon = 0.15$	12.2	18.7	77.1	32.1	73.6	76.3	76.5	77.3
MIM	$\epsilon = 0.1$	58.3	64.5	-	83.7	92.0	92.1	92.7	93.0
	$\epsilon = 0.15$	16.1	28.8	-	29.3	77.5	77.7	80.2	82.0
C&W	$c = 0.1$	61.6	71.1	89.2	88.1	97.3	97.7	97.1	97.6
	$c = 1.0$	30.6	39.2	79.1	75.3	78.1	80.4	87.3	91.2
	$c = 10.0$	0.2	17.0	37.6	20.0	23.8	29.1	39.7	46.0
PGD	$\epsilon = 0.1$	50.7	62.7	-	77.0	82.8	83.6	93.7	93.9
	$\epsilon = 0.15$	6.3	31.9	-	44.2	41.0	62.5	78.8	80.2

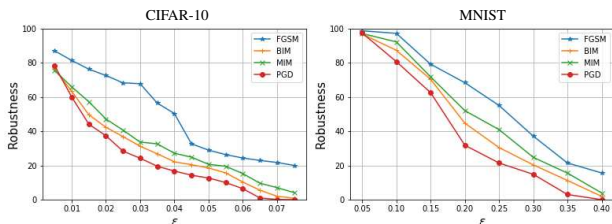


Figure 4: Robustness of our model (without adversarial training) against *white-box* attacks for various perturbation budgets.

Robustness against black-box settings is higher than white-box settings: In *white-box* settings, the adversary has complete knowledge of the model, so attacks should be more successful. In other words, if a defense does not suffer from obfuscated gradients, robustness of the model against *white-box* settings should be inferior to that in the *black-box* settings. Our extensive evaluations in Table 2 show that the proposed defense follows this trend and therefore does not obfuscate gradients.

Increasing the distortion bound (ϵ) decreases the robustness of defense: On increasing the perturbation size, the success rate of the attack method should significantly increase monotonically. For an unbounded distortion, the classifier should exhibit 0% robustness to the attack, which

again is true in our case (see Fig. 4).

6. Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement n° 725253–EyeCode)

7. Conclusion

Our findings provide evidence that the adversary’s task is made difficult by incorporating a maximal separation constraint in the objective function of DNNs, which conventional cross-entropy loss fails to impose. Our theory and experiments indicate, if the adversarial polytopes for samples belonging to different classes are non-overlapping, the adversary cannot find a viable perturbation within the allowed budget. We extensively evaluate the proposed model against a diverse set of attacks (both single-step and iterative) in *black-box* and *white-box* settings and show that the proposed model maintains its high robustness in all cases. Through empirical evaluations, we further demonstrate that the achieved performance is not due to obfuscated gradients, thus the proposed model can provide significant security against adversarial vulnerabilities in deep networks.

References

- [1] Evan Ackerman. How drive. ai is mastering autonomous driving with deep learning. *IEEE Spectrum Magazine*, 2017. [1](#)
- [2] Anish Athalye, Nicholas Carlini, and David Wagner. Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. *arXiv preprint arXiv:1802.00420*, 2018. [1](#), [2](#), [7](#)
- [3] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017. [2](#), [4](#)
- [4] Moustapha Cisse, Piotr Bojanowski, Edouard Grave, Yann Dauphin, and Nicolas Usunier. Parseval networks: Improving robustness to adversarial examples. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 854–863. JMLR. org, 2017. [1](#), [2](#)
- [5] Nilaksh Das, Madhuri Shanbhogue, Shang-Tse Chen, Fred Hohman, Li Chen, Michael E Kounavis, and Duen Horng Chau. Keeping the bad guys out: Protecting and vaccinating deep learning with jpeg compression. *arXiv preprint arXiv:1705.02900*, 2017. [1](#)
- [6] Mark De Berg, Otfried Cheong, Olivier Devillers, Marc Van Kreveld, and Monique Teillaud. Computing the maximum overlap of two convex polygons under translations. *Theory of computing systems*, 31(5):613–628, 1998. [3](#)
- [7] Guneet S Dhillon, Kamyar Azizzadenesheli, Zachary C Lipton, Jeremy Bernstein, Jean Kossaifi, Aran Khanna, and Anima Anandkumar. Stochastic activation pruning for robust adversarial defense. *arXiv preprint arXiv:1803.01442*, 2018. [2](#)
- [8] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9185–9193, 2018. [2](#), [4](#)
- [9] Justin Gilmer, Ryan P Adams, Ian Goodfellow, David Andersen, and George E Dahl. Motivating the rules of the game for adversarial example research. *arXiv preprint arXiv:1807.06732*, 2018. [7](#)
- [10] Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015. [1](#), [2](#), [4](#), [5](#)
- [11] Chuan Guo, Mayank Rana, Moustapha Cisse, and Laurens van der Maaten. Countering adversarial images using input transformations. *arXiv preprint arXiv:1711.00117*, 2017. [2](#)
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [5](#)
- [13] Daniel Jakubovitz and Raja Giryes. Improving dnn robustness to adversarial attacks using jacobian regularization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 514–529, 2018. [2](#)
- [14] Harini Kannan, Alexey Kurakin, and Ian Goodfellow. Adversarial logit pairing. *arXiv preprint arXiv:1803.06373*, 2018. [1](#)
- [15] J Zico Kolter and Eric Wong. Provable defenses against adversarial examples via the convex outer adversarial polytope. *arXiv preprint arXiv:1711.00851*, 1(2):3, 2017. [2](#)
- [16] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial examples in the physical world. *arXiv preprint arXiv:1607.02533*, 2016. [1](#), [2](#), [4](#)
- [17] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, 2016. [2](#), [5](#), [6](#), [7](#), [8](#)
- [18] Alexey Kurakin, Ian Goodfellow, Samy Bengio, Yinpeng Dong, Fangzhou Liao, Ming Liang, Tianyu Pang, Jun Zhu, Xiaolin Hu, Cihang Xie, et al. Adversarial attacks and defenses competition. In *The NIPS'17 Competition: Building Intelligent Systems*, pages 195–231. Springer, 2018. [5](#)
- [19] Fangzhou Liao, Ming Liang, Yinpeng Dong, Tianyu Pang, Jun Zhu, and Xiaolin Hu. Defense against adversarial attacks using high-level representation guided denoiser. *arXiv preprint arXiv:1712.02976*, 2017. [1](#)
- [20] Ji Lin, Chuang Gan, and Song Han. Defensive quantization: When efficiency meets robustness. 2018. [2](#)
- [21] Yan Luo, Xavier Boix, Gemma Roig, Tomaso Poggio, and Qi Zhao. Foveation-based mechanisms alleviate adversarial examples. *arXiv preprint arXiv:1511.06292*, 2015. [1](#)
- [22] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [23] Thomas Mensink, Jakob Verbeek, Florent Perronnin, and Gabriela Csurka. Distance-based image classification: Generalizing to new classes at near-zero cost. *IEEE transactions on pattern analysis and machine intelligence*, 35(11):2624–2637, 2013. [4](#)
- [24] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, Ken Nakae, and Shin Ishii. Distributional smoothing with virtual adversarial training. *arXiv preprint arXiv:1507.00677*, 2015. [2](#)
- [25] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2574–2582, 2016. [2](#)
- [26] Aamir Mustafa, Salman H Khan, Munawar Hayat, Jianbing Shen, and Ling Shao. Image super-resolution as a defense against adversarial attacks. *arXiv preprint arXiv:1901.01677*, 2019. [1](#)
- [27] Taesik Na, Jong Hwan Ko, and Saibal Mukhopadhyay. Cascade adversarial machine learning regularized with a unified embedding. *arXiv preprint arXiv:1708.02582*, 2017. [2](#)
- [28] Maryam M Najafabadi, Flavio Villanustre, Taghi M Khoshgoftaar, Naeem Seliya, Randall Wald, and Edin Muharemagic. Deep learning applications and challenges in big data analytics. *Journal of Big Data*, 2(1):1, 2015. [1](#)
- [29] Tianyu Pang, Kun Xu, Chao Du, Ning Chen, and Jun Zhu. Improving adversarial robustness via promoting ensemble diversity. *arXiv preprint arXiv:1901.08846*, 2019. [1](#), [6](#), [7](#), [8](#)

- [30] Nicolas Papernot, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Z Berkay Celik, and Ananthram Swami. The limitations of deep learning in adversarial settings. In *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 372–387. IEEE, 2016. [2](#), [5](#)
- [31] Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. Distillation as a defense to adversarial perturbations against deep neural networks. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 582–597. IEEE, 2016. [1](#), [2](#)
- [32] Aditi Raghunathan, Jacob Steinhardt, and Percy Liang. Certified defenses against adversarial examples. *arXiv preprint arXiv:1801.09344*, 2018. [2](#)
- [33] Andrew Slavin Ross and Finale Doshi-Velez. Improving the adversarial robustness and interpretability of deep neural networks by regularizing their input gradients. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. [2](#), [6](#), [8](#)
- [34] Conrad Sanderson. *Biometric person recognition: Face, speech and fusion*, volume 4. VDM Publishing, 2008. [1](#)
- [35] Aman Sinha, Hongseok Namkoong, and John Duchi. Certifiable distributional robustness with principled adversarial training. *stat*, 1050:29, 2017. [2](#)
- [36] Chuanbiao Song, Kun He, Liwei Wang, and John E Hopcroft. Improving the generalization of adversarial training with domain adaptation. *arXiv preprint arXiv:1810.00740*, 2018. [6](#), [7](#)
- [37] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013. [1](#), [2](#)
- [38] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. Ensemble adversarial training: Attacks and defenses. *arXiv preprint arXiv:1705.07204*, 2017. [1](#), [2](#)
- [39] Florian Tramèr, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. The space of transferable adversarial examples. *arXiv preprint arXiv:1704.03453*, 2017. [5](#)
- [40] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *European conference on computer vision*, pages 499–515. Springer, 2016. [2](#), [5](#)
- [41] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, and Alan Yuille. Mitigating adversarial effects through randomization. In *International Conference on Learning Representations*, 2018. [1](#), [2](#)
- [42] Cihang Xie, Zhishuai Zhang, Jianyu Wang, Yuyin Zhou, Zhou Ren, and Alan Yuille. Improving transferability of adversarial examples with input diversity. *arXiv preprint arXiv:1803.06978*, 2018. [2](#)
- [43] Fuxun Yu, Chenchen Liu, Yanzhi Wang, and Xiang Chen. Interpreting adversarial robustness: A view from decision surface in input space. *arXiv preprint arXiv:1810.00144*, 2018. [8](#)