

Stylized Adversarial Defense

Muzammal Naseer, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan and Fatih Porikli

Abstract—Deep Convolution Neural Networks (CNNs) can easily be fooled by subtle, imperceptible changes to the input images. To address this vulnerability, adversarial training creates perturbation patterns and includes them in the training set to robustify the model. In contrast to existing adversarial training methods that only use class-boundary information (e.g., using a cross-entropy loss), we propose to exploit additional information from the feature space to craft stronger adversaries that are in turn used to learn a robust model. Specifically, we use the *style* and *content* information of the target sample from another class, alongside its class-boundary information to create adversarial perturbations. We apply our proposed *multi-task* objective in a deeply supervised manner, extracting multi-scale feature knowledge to create maximally separating adversaries. Subsequently, we propose a max-margin adversarial training approach that minimizes the distance between source image and its adversary and maximizes the distance between the adversary and the target image. Our adversarial training approach demonstrates strong robustness compared to state-of-the-art defenses, generalizes well to naturally occurring corruptions and data distributional shifts, and retains the models accuracy on clean examples.

Index Terms—Adversarial Training, Style Transfer, Max-Margin Learning, Adversarial Attacks, Multi-task Objective.



1 INTRODUCTION

Although deep networks excel on a variety of learning tasks, they remain vulnerable to adversarial perturbations. These perturbations are imperceptible to humans, but significantly degrade the prediction accuracy of a trained model. Adversarial training [1] has emerged as a simple and successful mechanism to achieve robustness against adversarial perturbations. In this process, blind-spots of the model are first found by crafting malicious perturbations and subsequently included in the training set to learn a robust model.

In this paper, we interpret adversarial training from a margin maximization perspective. We consider margin as the shortest distance from a data-point to the classifiers boundary in the decision and perceptual (feature) spaces. Intuitively, the highest robustness can be achieved by learning a margin maximizing model that first crafts a maximally separated adversarial example and then readjusts the boundary to correctly classify such perturbed samples. However, in practice, this task turns out to be a nested max-min optimization problem, whose solution is non-trivial [2]. Therefore, we propose an alternate way to maximize classifier margins. Our approach is motivated by the fact that adversarial training maximizes a lower bound on the classifier’s margin. Towards this goal, our main idea is to identify a target image from a different class for guidance, and create perturbations that can push the source image towards the target in both feature and output spaces using a multi-task objective function.

In the pursuit of creating highly deceptive adversaries, we propose an attack based on multi-modal information including classifiers boundary information, image style and visual content. In this manner, the perturbations cause

significant changes to the intermediate feature as well as output decision space using a diverse multi-task attack objective based on multiple supervisory cues. The existing targeted attacks in the literature create adversaries by moving towards the least likely (or the most confusing) class [3], [4], [5]. In this manner, only the class boundary information is used to craft adversarial perturbations. Different to those, we create targeted adversaries by pushing the sample towards a randomly picked sample from a different class such that its style and content representations are also reshaped besides the output prediction. This is done by incorporating multi-scale information from the feature hierarchy in a deeply supervised manner [6].

Based on the proposed carefully crafted perturbations, we develop our Stylized Adversarial Training (SAT) approach to achieve robustness. Specifically, we enforce a margin-maximizing objective during adversarial training, which minimizes the distance between clean and perturbed images while maximizing the distance between clean image and target sample (used to create adversaries). The model thus learns corrective measures with respect to a reference sample from a different class, thereby enhancing the model’s robustness. Since our attack objective uses supervision from the style and content of a target image from different class, it forces the perturbations to lie close to the natural image manifold. As a result, our adversarially trained model performs significantly better than other adversarial training approaches [1], [7], [8] on the clean images. Simultaneously, our proposed defense shows strong robustness against naturally occurring image degradations such as contrast changes, blurring and rain, that cause distributional shifts. We further demonstrate that the model trained with our proposed scheme performs much better on the style transfer task despite having less parametric complexity (see Fig. 1).

The major contributions of our work are:

- We propose to set-up priors in the form of fooling target samples during adversarial training and propose a multi-task objective for adversary creation

- M. Naseer, S. Khan and F. Porikli are with the College of Engineering and Computer Science, Australian National University, Canberra, ACT 2614. E-mail: muzammal.naseer@anu.edu.au
- M. Naseer is also with Data61-CSIRO, Canberra, Australia.
- S. Khan, M. Hayat and F. Khan are with the Inception Institute of Artificial Intelligence, Abu Dhabi, UAE.

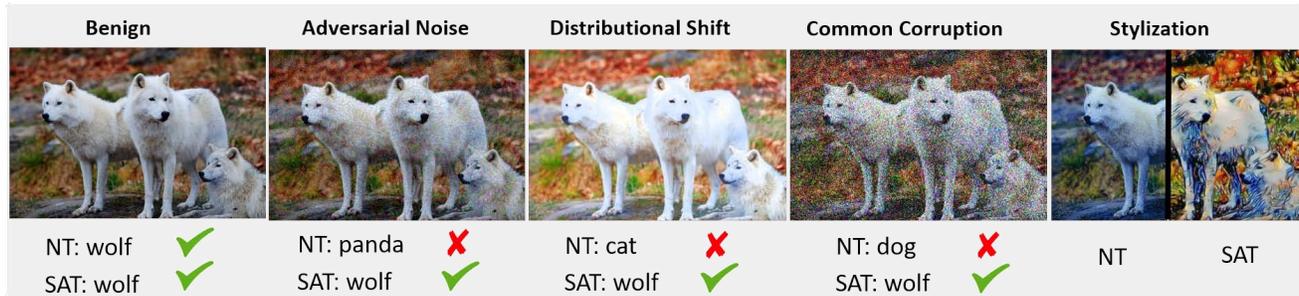


Fig. 1. A robust model trained with our proposed Stylized Adversarial Training (SAT) framework generalizes not only to adversarial noise but also handles naturally occurring distributional shifts (e.g., contrast change in the above example), common corruptions (e.g., sensor noise) and performs better stylization compared to a Naturally Trained (NT) model.

that seeks to fool the model in terms of image style, visual content as well as the decision boundary for the true class.

- Based on a high-strength perturbation, we develop a margin-maximizing (*contrastive*) adversarial training procedure that maps perturbed image close to clean one and maximally separates it from the target image used to craft the adversary.
- With extensive evaluation, we demonstrate that transferring information from multi-task objectives helps us perform favorably well against the strongest adversarial training methods such as PGD based Adversarial Training (PAT) [1], Trades [7] and Feature scattering [8].
- Compared to conventional adversarial training, our approach does not cause a drop in clean accuracy, and performs well against the real-world common image corruptions [9]. We further demonstrate robustness and generalization capabilities of the proposed training regime when the underlying data distribution shifts (Sec. 4.1).

2 RELATED WORKS

Adversarial Training: Training a model on adversarial examples can regularize it and increase its adversarial robustness. Goodfellow *et al.* [10] proposed a computationally fast adversary generation algorithm known as ‘Fast Gradient Sign Method’ (FGSM). FGSM suffers from label leakage [11] that allows the model to overfit on FGSM’s generated adversaries, hence hampering its adversarial generalization. Tamer *et al.* [12] proposed to mitigate this issue by taking a small random step before running FGSM. Their attack is known as ‘Random Fast Gradient Sign Method’ (RFGSM). Their method performs relatively better, but still suffers under iterative attacks [1], [13]. Madry *et al.* [1] solved the overfitting problem by adversarially training models on iterative attack known as ‘Projected Gradient Descent’ (PGD). PGD is an untargetted, label-dependent attack and models trained on PGD adversaries show significant robustness to the strongest white-box attacks [11], [13]. However, PGD gains robustness at the cost of a significant drop in clean accuracy and lacks a clear mechanism to control the accuracy-robustness trade-off. This is where Zhang *et al.* [7] contributed and proposed a method to control the trade-off with an untargetted, label-independent (unsupervised)

attack to create adversaries along with a surrogate clustering loss to minimize the model’s empirical risk.

However, [1], [7] deploy iterative attacks which are computationally expensive and less scalable to high-dimensional datasets. Further, adversarial training done using untargetted attacks whether computationally expensive [1], [7] or efficient [14], [15] results in only a limited robustness. To improve it, [8] proposed a faster attack that operates in logit space in an unsupervised way to maximize optimal transport distance. Combined with label smoothing, their method produced state-of-the-art results on SVHN, CIFAR10 and CIFAR100. In this work, we propose a conceptually simple and efficient adversarial training process, exploiting a multi-task loss that helps us perform favorably well against previous state-of-the-art methods.

Augmentation based Adversarial Training: Since traditional adversarial training results in a significant drop in clean accuracy, augmentation based methods have been proposed to overcome this limitation. [16] increased clean accuracy of adversarially trained models by using augmentation methods [17], [18]. Similarly, Zhang *et al.* [5] also proposed to create adversaries using augmentation while updating the model on these adversaries using perturbed labels. However, these methods [5], [16], [17] generate adversaries by mixing a sample with another which may or may not come from the same class, therefore the inter-class margins might not be enforced during training. In comparison, we carefully craft adversaries by style transfer followed by max-margin adversarial training, that results in enhanced robustness.

Metric Learning Defenses: More recently, some adversarial training efforts maximize the margin between clean and adversarial examples of different class samples. Mustafa *et al.* [19] proposed a contrastive loss function to maximize the inter-class distances along the feature hierarchy of a deep network. [2] dynamically selects the right perturbation budget for each data point to better enforce margin constraints during adversarial training. Triplet loss has also been explored to enforce margin constraints during training [20], [21], [22]. Similar to these approaches, our proposed method is model-agnostic and incorporates distance-based learning scheme. However, different from the previous works, we first transfer the style, content and boundary information from a target sample to the input and then maximize the distance between the target and the perturbed samples. In this manner, our triplet creation is automatic and does not

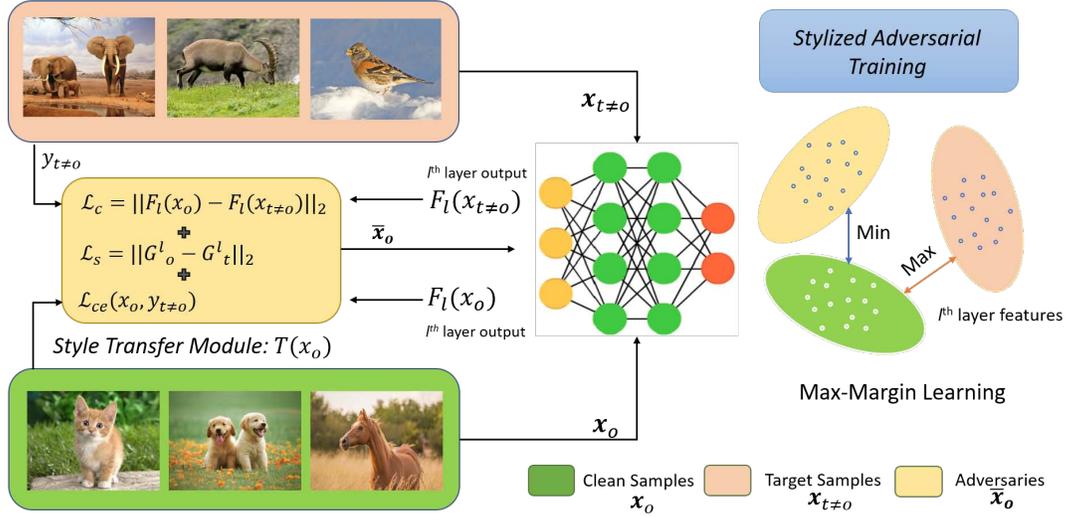


Fig. 2. *Stylized Adversarial Training (SAT)*. Our style transfer module (*left*) crafts perturbations based on three complimentary cues, that include content (\mathcal{L}_c) and style (\mathcal{L}_s) of the target image as well as the classifier boundary information (\mathcal{L}_{ce}). Based on the generated perturbations, our adversarial training approach seeks to minimize the distance between clean and adversarial examples of the same class and maximize the inter-class distances (*right*).

need careful sample selection as in [22].

3 METHODOLOGY

Consider a classifier, $\mathcal{F}(\cdot)$ that maps input samples, $\mathbf{x} \in \mathbb{R}^d$, drawn from a dataset, \mathbf{X} , to a discriminative space $\mathcal{F}(\mathbf{x}) \in \mathbb{R}^n$, where n represents the number of categories. Classifier can learn this mapping by minimizing an empirical risk defined on \mathbf{X} . Further, suppose that $\mathcal{F}_l(\cdot)$ represents a feature map at the l^{th} layer, and \mathcal{T} denotes a transformation operation that keeps the output close to input i.e., $\|\mathbf{x}_o - \mathcal{T}(\mathbf{x}_o)\| < \epsilon$, where ϵ is the perturbation budget. We present a generic training mechanism focused on robustifying neural networks by *minimizing* feature difference between the original examples \mathbf{x}_o and the transformed positive samples $\mathcal{T}(\mathbf{x}_o)$ and *maximizing* feature difference between \mathbf{x}_o and targeted class samples $\mathbf{x}_{t \neq 0}$. The contrastive constraints can be achieved by minimizing the following loss function:

$$\mathcal{L}_m(\mathbf{x}_o, \mathbf{x}_{t \neq 0}; \mathcal{T}, \mathcal{F}) = \max\{\|\mathcal{F}_l(\mathbf{x}_o) - \mathcal{F}_l(\mathcal{T}(\mathbf{x}_o))\|_p - \|\mathcal{F}_l(\mathbf{x}_o) - \mathcal{F}_l(\mathbf{x}_{t \neq 0})\|_p + m, 0\}, \quad (1)$$

where m represents the margin and $\|\cdot\|_p$ denotes p -norm. The transformation, \mathcal{T} , plays a significant role in training and should satisfy the following two properties:

- The transformation maps output close to the input i.e., $\mathcal{T}(\mathbf{x}_o) \approx \mathbf{x}_o$.
- \mathcal{T} should correlate with adversarial noise that fools the network.

Next in Sec. 3.1, we elaborate our proposed stylized perturbation generation mechanism (\mathcal{T}) that is central to our proposed defense described in Sec. 3.2.

3.1 Transformation: Stylized Adversary Generation

The choice of transformation \mathcal{T} is critical to the strength of robustness achieved with adversarial training. Here, we

present our transformation mechanism, achieved with a style transfer module (Fig. 2), that jointly utilizes the style, content and class-boundary information to craft deceptive perturbations. The overall objective for learning the adversarial transformation is,

$$\begin{aligned} \operatorname{argmax}_{\mathcal{T}} \quad & \mathcal{L}(\mathcal{F}(\mathcal{T}(\mathbf{x}_o)), \mathcal{F}(\mathbf{x}_o)), \\ \text{s.t.} \quad & \|\mathcal{T}(\mathbf{x}_o) - \mathbf{x}_o\|_{\infty} \leq \epsilon, \end{aligned} \quad (2)$$

where, \mathcal{L} denotes any loss function and ϵ is the allowed perturbation budget. The aim is to remain in the vicinity of input sample \mathbf{x}_o , but maximally alter the predicted output by the model \mathcal{F} . The above objective is pursued in previous adversary generation methods as well, however, our main difference is the way we incorporate target samples ($\mathbf{x}_{t \neq 0}$) while crafting adversaries. Specifically, we extract three types of information about the target sample including class-boundary information, image style and visual content. For example, in an effort to robustify the model, the adversarial transformation should create adversarial examples that contain style and texture of the samples from target classes $\mathbf{x}_{t \neq 0}$, within a given perturbation budget, ϵ . The following adversarial loss is minimized to learn the transformation \mathcal{T} :

$$\mathcal{L}_{adv} = \underbrace{\alpha \cdot \mathcal{L}_s}_{\text{Style loss}} + \underbrace{\gamma \cdot \mathcal{L}_c}_{\text{Content loss}} + \underbrace{\beta \cdot \mathcal{L}_{ce}}_{\text{Cross-entropy loss}}, \quad (3)$$

where α, γ, β denote the hyper-parameters used for loss re-weighting which are set via validation. Notably, the style and content loss components are computed within feature space while boundary information comes from the logit space. We explain the individual losses below.

Style loss: The objective of the style loss \mathcal{L}_s is to transfer the texture of the target image to fool the classifier, \mathcal{F} . Style transfer [23] can be achieved by minimizing the mean-squared distance between the Gram matrices obtained from

the feature maps at layer l of the original and targeted images, as follows:

$$\mathcal{L}_s = \|\mathbf{G}_o^l - \mathbf{G}_{t \neq o}^l\|_2^2, \quad \text{s.t., } \mathbf{G}^l = \mathbf{f}\mathbf{f}^T, \quad (4)$$

where $\mathbf{G} \in \mathbb{R}^{c \times c}$ represents the Gram matrix, $\mathbf{f} \in \mathbb{R}^{c \times (h \cdot w)}$ denotes the matrix formed by stacking together the channel-wise features from $\mathcal{F}_l(\mathbf{x}) \in \mathbb{R}^{c \times h \times w}$. Here, h , w and c denote the height, width and channel dimensions of the feature tensor from layer l , respectively.

Content loss: The objective of the content loss \mathcal{L}_c is to mix the content of the target class image with that of the original image. This is achieved by minimizing the mean-squared distance between the feature representations of \mathbf{x}_o and $\mathbf{x}_{t \neq o}$, as follows:

$$\mathcal{L}_c = \|\mathcal{F}_l(\mathbf{x}_o) - \mathcal{F}_l(\mathbf{x}_{t \neq o})\|_2^2. \quad (5)$$

Boundary loss: The objective of this loss is to push the transformed sample into the boundary of targeted class. For this purpose, we use the regular cross-entropy loss. If \mathbf{y}_t represents target label then boundary-based targeted attack can be achieved by minimizing $\mathcal{L}_{ce}(\mathbf{x}_o, \mathbf{y}_t)$.

3.2 Stylized Adversarial Training

We describe our robust training framework in Algorithm 1 that applies adversarial transformation to push clean samples \mathbf{x}_o toward the targeted samples $\mathbf{x}_{t \neq o}$ within a predefined budget ϵ , and then robustify the network with adversarial training. The adversarial training procedure employs cross-entropy loss alongside the contrastive loss \mathcal{L}_m that seeks to maximize inter-class margins (Eq. 1).

Algorithm 1 SAT: A Robust Training Framework

Require: A classifier \mathcal{F} , clean sample \mathbf{x}_o and their corresponding labels \mathbf{y} , targeted sample $\mathbf{x}_{t \neq o}$ and their corresponding labels \mathbf{y}_t , margin loss \mathcal{L}_m , cross-entropy loss \mathcal{L}_{ce} , w_1 , w_2 and no. of iterations T .

- 1: **for** $t = 1$ to T **do**
- 2: Forward pass \mathbf{x}_o and $\mathbf{x}_{t \neq o}$ to \mathcal{F} and compute adversarial loss \mathcal{L}_{adv} (Eq. 3);
- 3: Compute gradient noise, $\mathbf{g}_t = \nabla_{\mathbf{x}} \mathcal{L}_{adv}$;
- 4: Generate adversaries using;

$$\bar{\mathbf{x}}_o = \mathbf{x}_o - \epsilon \cdot \text{sign}(\mathbf{g}_t); \quad (6)$$

- 5: Forward pass $\bar{\mathbf{x}}_o$ through \mathcal{F} ;
- 6: Backpass and update the parameters of \mathcal{F} to minimize the combined loss:

$$\mathcal{L} = w_1 \cdot \mathcal{L}_m(\mathbf{x}, \bar{\mathbf{x}}_o, \mathbf{x}_{t \neq o}) + w_2 \cdot \mathcal{L}_{ce}(\bar{\mathbf{x}}_o, \mathbf{y}) \quad (7)$$

- 7: **end for**
 - 8: **return** Robust classifier, \mathcal{F} .
-

Non-Adversarial \mathcal{T} : In order to emphasize on the significance of transformation \mathcal{T} in the SAT framework, here we consider a non-adversarial transformation function. In this case, \mathcal{T} can be as simple as adding Gaussian noise to the clean samples. Such a transformation is computationally less expensive and has been studied before in [24], [25], [26]. In this case, the perturbation generation process in lines 2-4 of Algorithm 1 is simply replaced with adding randomly sampled Gaussian noise in the image. We explore the effect of non-adversarial transformation on adversarial

robustness and compare our training Algorithm 1 with a recently proposed ‘Guided Cross-Entropy’ (GCE) [27] method in Sec. 4.2. We note that in the non-adversarial scenario, targeted samples are not playing any significant role other than providing a reference to contrastive margin loss which leads to a sub-optimal solution (see Sec. 4.2). This shows the efficacy of proposed stylized perturbation generation approach (Sec. 3.1).

4 EXPERIMENTAL PROTOCOL

We experiment on the widely used SVHN [29], CIFAR10, and CIFAR100 datasets [30]. We show comparative studies on the ResNet18 and WideResNet models. The models are trained using SAT (Algorithm 1) using SGD optimizer. The pixel values are normalized within $[-1, +1]$ and, and label smoothing [8] is used during training. Unless otherwise mentioned, the perturbation budget ϵ is set to 8 (out of 255). NT and AT respectively denote naturally and adversarially trained models. Our code and pretrained models are available at <https://github.com/Muzammal-Naseer/SAT>.

4.1 SAT: Defense Results and Insights

We thoroughly investigate the effect of our proposed adversarial transformation (Sec. 3) to maximize adversarial robustness without compromising clean accuracy. Our analysis is divided into the four categories: (a) Robustness against constrained adversarial attacks ($l_\infty \leq 8$), (b) Robustness against unconstrained adversarial attacks, namely Rectangular Occlusion Attacks (ROA) [31], which completely destroy the image content within a given window size, (c) Robustness against natural distributional shifts in data, and (d) Robustness against common corruptions.

4.1.1 Robustness against constrained adversarial attacks

We compare our approach with the state-of-the-art methods including Trades [7], Feature Scattering (FS) [8] and the metric learning based Prototype Conformity Loss (PCL) [19]. For a fair and direct comparison, we follow the same threat models (attack settings) and network architectures as recommended in the papers of the respective methods.

Comparison with Trades [7] is presented in Table 1. We note that [7] offers a trade-off parameter (λ) to increase robustness on the cost of losing clean accuracy. Our defense not only achieves 24.1% higher robustness when compared with best adversarial results from [7] ($\lambda = 6$) against PGD attack with 20 iterations but also improves clean accuracy by 4.7% when compared with best clean results from [7] ($\lambda = 1$). Furthermore, our defense can withstand large number of PGD attack iterations e.g. the drop in robustness of our defense is only 2.5% when attack iterations increase from 20 to 1000. As a result of our proposed max-margin learning, class-wise latent features of our defense are well clustered and separated as compared to [7] (see Figure 3). This leads to significantly better robustness than Trades on major attacks including PGD, CW, MIFGSM and DeepFool. **Comparison with FS [8]** is presented in Tables 2 and 3. In terms of worst-case robustness measure (CW attack with 100 iterations), our defense offers 11.2%, 10.0% and 8.8% (Tables 2 and 3) robustness gain on CIFAR10, CIFAR100

Model	Defense	Clean	PGD		CW	MIFGSM	DeepFool
			20	1000	100	100	100
ResNet18	NT	94.5	0.0	0.0	0.0	0.0	1.2
	Trades [7] ($\lambda = 1$)	91.3	26.5	-	-	-	-
	Trades [7] ($\lambda = 5$)	81.7	50.6	50.1	49.3	52.1	63.0
	SAT (ours)	92.1	69.1	66.8	62.5	72.8	66.0
WideResNet	NT	95.3	0.0	0.0	0.0	0.0	3.2
	Trades [7] ($\lambda = 1$)	88.6	49.1	48.9	48.3	53.5	69.0
	Trades [7] ($\lambda = 6$)	84.9	56.6	56.4	53.7	58.5	72.0
	SAT (ours)	93.3	80.7	78.2	71.8	82.8	88.3

TABLE 1

White-box attack scenario. Comparisons of our defense with Trades [7] on CIFAR10 test set under perturbation budget $\epsilon \leq 8$. Models trained via our proposed approach withstand PGD attack with 1000 iteration while providing high accuracy on clean samples. We used DeepFool [28] with default settings and project the adversarial noise found by the attack on l_∞ ball to respect the perturbation budget.

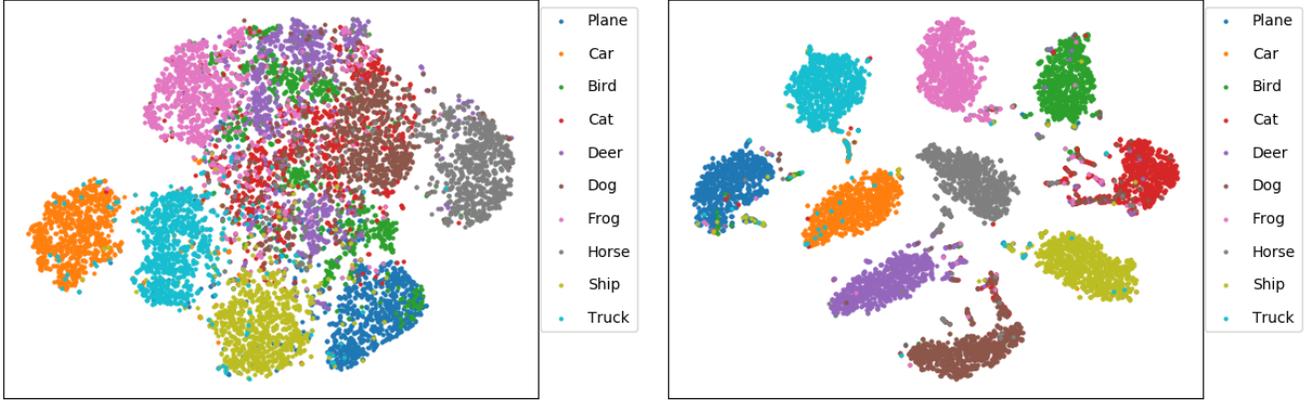


Fig. 3. Latent space t-SNE visualization of intermediate features extracted from Trades [7] (left) and our SAT model (right) on CIFAR10 test set. Compared to Trades [7], our SAT model forms distinct class-wise clusters.

Defense	Clean	FGSM	PGD				CW			
			10	20	40	100	10	20	40	100
NT	95.6	36.9	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
PAT [1]	85.7	54.9	45.1	44.9	44.8	44.8	45.9	45.7	45.6	45.4
BL [3]	91.2	70.7	-	57.5	-	55.2	-	56.2	-	53.8
FS [8]	90.0	78.4	70.9	70.5	70.3	68.6	62.6	62.4	62.1	60.6
Ours	93.3	85.0	81.1	80.7	79.8	78.5	75.0	74.9	73.2	71.8

TABLE 2

White-box attack scenario. Comparison (%) of our approach with naturally trained (NT), Madry (PAT) [1], bilateral (BL) [3] and feature scattering (FS) [8] methods on CIFAR10 test set under different threat models. Attacks ran for maximum of 100 iterations. Models trained using our approach show significant robustness without losing noticeable clean accuracy.

Defense	Clean	FGSM	PGD		CW		Defense	Clean	FGSM	PGD		CW	
			20	100	20	100				20	100	20	100
NT	97.2	53.0	0.3	0.1	0.3	0.1	NT	79.0	10.0	0.0	0.0	0.0	0.0
PAT [1]	93.9	68.4	47.9	46.0	48.7	47.3	PAT [1]	59.9	28.5	22.6	22.3	23.2	23.0
BL [3]	94.1	69.8	53.9	50.3	-	48.9	BL [3]	68.2	60.8	26.7	25.3	-	22.1
FS [8]	96.2	83.5	62.9	52.0	61.3	50.8	FS [8]	73.9	61.0	47.2	46.2	34.6	30.6
Ours	96.2	86.0	73.2	71.5	70.0	68.0	Ours	74.1	64.9	49.7	49.1	44.2	40.6

TABLE 3

White-box attack scenario. Comparison (%) is shown on SVHN (left) and CIFAR100 (right) test sets.

and SVHN datasets, respectively. We further observe that our defense simultaneously improves clean accuracy while achieving significant robustness gains. In contrast to SAT

(Algorithm 1), FS [8] is dependant on optimal transport distance to increase model loss towards the unknown class. It does not leverage target image information and neither

Defense	Clean	ROA (Gradient Search)				ROA (Exhaustive Search)			
		5x5	7x7	9x9	11x11	5x5	7x7	9x9	11x11
NT	96.0	56.0	42.3	27.3	13.0	38.5	21.0	8.6	2.4
PAT [1]	86.8	47.2	29.7	15.4	7.0	33.8	23.7	7.9	3.0
Trades [7]	84.9	49.8	30.5	16.1	7.1	39.2	28.0	8.8	3.2
FS [8]	90.0	66.5	56.2	44.3	32.6	49.1	42.2	22.3	12.5
Ours	93.3	74.7	65.4	54.2	42.6	56.4	51.0	31.1	23.1

TABLE 4

Adversarial robustness against unconstrained adversarial attack, ROA [31] at different window sizes. Our defense perform significantly better than other training approaches.

(a) CIFAR10. Perturbation budget is 8/255 in ℓ_∞ norm.

Defense	Clean	FGSM	IFGSM	CW	MIFGSM	PGD
PCL [19]	91.9	74.9	46.0	51.8	49.3	46.7
SAT (ours)	92.3	84.7	83.5	81.2	83.8	83.5

(b) CIFAR100. Perturbation budget is 8/255 in ℓ_∞ norm.

Defense	Clean	FGSM	IFGSM	CW	MIFGSM	PGD
PCL [19]	68.3	60.9	34.1	36.7	33.7	36.1
SAT (ours)	72.5	65.2	48.3	47.5	49.2	48.0

(c) SVHN. Perturbation budget is 8/255 in ℓ_∞ norm.

Defense	Clean	FGSM	IFGSM	CW	MIFGSM	PGD
PCL [19]	94.4	76.5	48.8	54.8	47.1	47.7
SAT (ours)	95.3	85.7	67.8	65.3	68.0	66.8

TABLE 5

White-box: Comparison between PCL and our proposed defense (SAT) under the threat model of PCL. SAT shows significantly better robustness as compared to PCL.

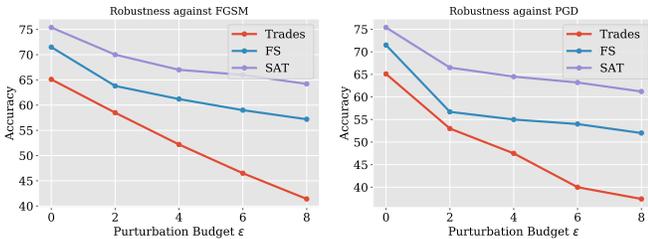


Fig. 4. **White-Box analysis (%)**: Our proposed SAT handles the data distributional shifts significantly better than FS [8] and Trades [7]. Models are evaluated on CINIC-10 [32] test set. PGD attack ran for 10 iterations. SAT’s state-of-the-art robustness to such distributional shifts complements the strength of our proposed approach to enhance the generalizability of deep neural networks.

it offers any mechanism to increase inter-class margins. FS [8] does perform better than Madry [1] and Trades [7] but its performance is sub-optimal when compared with our defense approach.

Comparison with PCL [19]: The results in Table 5 indicate that our defense demonstrates significant robustness gains as compared to metric-learning based prototype conformity loss (PCL) [19], boosting adversarial accuracy by 36.8%, 11.9% and 19.1% on CIFAR10, CIFAR100 and SVHN, respectively. We further observe that alongside significant gains for adversarial robustness, our approach also increases the clean accuracy on all the evaluated datasets. We note that PCL [19] enhances separation between class centers and is dependent on PGD untargetted attack, while SAT increases distance between the features of original and target samples,

which are responsible for adversarial perturbations.

4.1.2 Robustness against unconstrained attack

Here, we consider the unconstrained occlusion attack called ROA [31] against our proposed adversarial training approach. We run the gradient and exhaustive search versions of ROA on CIFAR10 dataset with four different window sizes (ranging from 5x5 to 11x11). We observe that as the window size is increased for ROA, the robustness of PAT and Trades is matched with a model trained without adversarial training (NT). In comparison, our defense shows a higher robustness e.g., an relative increment of 84% and 620% more over FS and Trades respectively, at the window size of 11x11 (see Table 4). This is attributed to the fact that our training approach constructs a smooth loss surface and requires large input distortions to deceive the model. We empirically demonstrate such smoothness by analyzing the intermediate features of PAT [1] and SAT in terms of correlation loss (CL) between features of adversarial image with respect to the clean image. Lower the correlation, better the feature space as it indicates that model feature space does not change significantly in response to the attack. We ran 100 iterations of PGD attack and extract features from the last layer before the logit layer. Correlation loss is measured in terms Frobenius norm between covariance of adversarial and clean features. The averaged correlation loss of our method is 0.74 as compared to 8.8 from PAT [1] on the CIFAR10 dataset. This further supports the robustness of our proposed adversarial training (SAT) approach.

4.1.3 Robustness against natural distributional shifts

Generalization of deep networks goes beyond adversarial robustness e.g. robustness to non-adversarial distributional shifts. For example, [33] showed that models trained on CIFAR10 suffer from accuracy drop when there are small natural distributional shifts in the data. We evaluated robustness of adversarially trained models on CINIC-10 test set (90k images). CINIC-10 is down-sampled from ImageNet [34] and contains the same classes as in CIFAR10 but with a significantly different distribution. Fig. 4 shows that SAT outperforms both Trades [7] and feature scattering [8], and generalizes well to the shift in underlying data distribution. This is potentially because the SAT simulates distribution shifts in style during its training procedure.

4.1.4 Robustness to Common Corruptions

Distributional shifts in the data can also come in the form of natural corruptions [9]. Hendrycks *et al.* [9] simulated

Corruption	NT	Trades [7]	FS [8]	SAT (Ours)
Brightness	93.7	80.6	88.3	92.1
Contrast	92.1	43.1	82.5	88.9
Defocus Blur	92.3	80.0	86.1	90.4
Elastic Transform	86.4	78.9	83.6	87.6
Fog	91.8	60.3	78.6	86.9
Gaussian Blur	91.4	78.0	85.4	89.7
Gaussian Noise	78.6	79.1	85.9	90.4
Glass Blur	71.7	77.9	80.9	82.6
Implus Noise	76.1	73.8	81.9	86.0
JPEG Compression	78.8	82.8	85.8	89.9
Motion Blur	89.6	76.5	84.1	88.3
Pixelate	88.3	82.7	86.1	90.1
Saturate	93.3	81.5	87.3	91.4
Shot Noise	81.9	80.4	86.2	90.8
Snow	86.3	80.4	84.0	89.0
Spatter	88.3	80.7	84.1	87.8
Speckle Noise	82.1	80.2	86.0	90.6
Zoom Blur	91.1	78.9	86.0	90.2
Mean	86.3	76.6	84.6	89.0
Variance	41.2	90.2	5.5	4.9

TABLE 6

Comparative analysis of robustness (%) to common corruptions is shown. SAT showed significant improvement over majority of the corruptions and did specially well against those that are most difficult for naturally trained (NT) models such as glass blur, Gaussian noise and impulse noise. Mean accuracy (higher is better) and variance (lower is better) are reported.

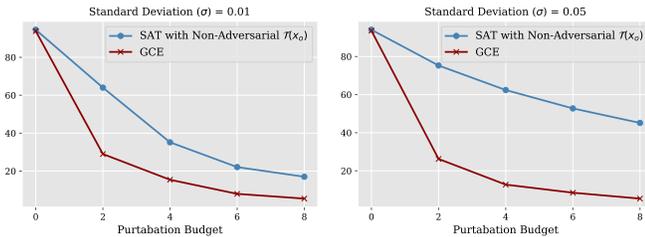


Fig. 5. **White-box analysis:** Our approach with non-adversarial transformation is compared against GCE [27]. Robustness is measured on CIFAR-10 test set against PGD with 20 iterations and random restarts. Models trained with our approach are significantly robust compared to GCE method [27].

multiple such corruptions including snow, fog and glass blur. We study 18 of such corruptions. Depending upon the severity, each corruption is sub-divided into 5 levels resulting in a total of 50k images for every corruption type. Analysis of robust models on such distributional shifts is presented in Table 6. Interestingly, theoretically robust model, Trades [7] loses significant accuracy on such corruptions as compared to naturally trained (NT) models. Feature scattering performed better than Trades. We observe that, compared with Trades [7] and FS [8], a naturally trained model shows better robustness to these image corruptions. Our proposed SAT, however, demonstrates improved generalization to common corruptions while simultaneously providing adversarial robustness. Contrastive nature of our adversarial training helps the model to adapt to such distribution shifts which aligns with findings that discriminative learning boosts domain adaptation and generalization [35]. Further, it is consistent with the experiments where random sterilizations also help in boosting domain adaptation [36].

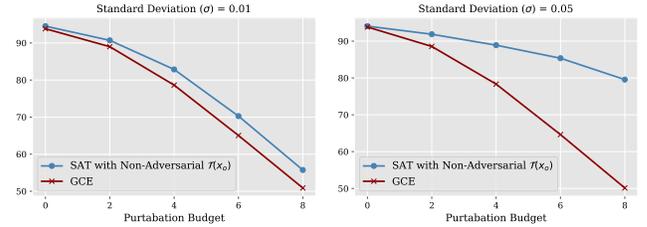


Fig. 6. **Black-Box analysis:** Our approach trained with non-adversarial perturbations is compared against GCE [27]. Adversaries are generated using MIFGSM with 10 iterations on CIFAR-10 test set. Our trained models show high resistance to transferable attack as compared to GCE [27].

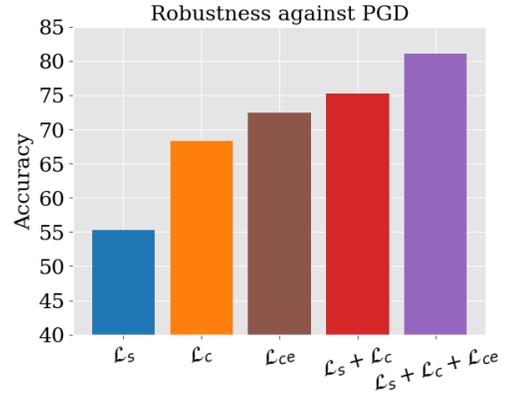


Fig. 7. White-box robustness analysis shows the effectiveness of different losses introduced in Eq. 3. Results are reported for WideResNet on CIFAR10 dataset.

4.2 Non-adversarial T: Defense Results and Insights

In this section, we analyse the performance of training SAT with non-adversarial transformation to establish empirical evidence of adversarial robustness with non-adversarial transformation. For this case, Gaussian noise is considered and ResNet18 [37] is trained using our approach and GCE [27] on CIFAR-10 training set. White-box robustness is measured against PGD [1] with 20 iterations. Black-box robustness is evaluated by transferring adversaries from VGG19 [38] using momentum iterative fast sign gradient (MIFGSM) [39] attack on the CIFAR-10 test set. From our experimental results shown in Figures 5 and 6, we observe that, as the transformation, \mathcal{T} , becomes better *e.g.* by increasing the standard deviation (σ) of Gaussian noise, our robustness against PGD attack increases significantly. This suggests that better transformation can lead to more robust models. We further notice that the Gaussian transformation does not have noticeable effect on GCE performance in terms of robustness. It is interesting to note that our approach maintains accuracy on clean examples while its robustness improves with better transformation. This behavior complements our design approach that takes into account the relationship of original and transformed samples unlike GCE [21] which only relies on minimizing the probability of other classes with respect to the true class.

4.3 Ablation Study

We dissect our proposed adversarial training for WideResNet architecture on CIFAR10 dataset to develop further

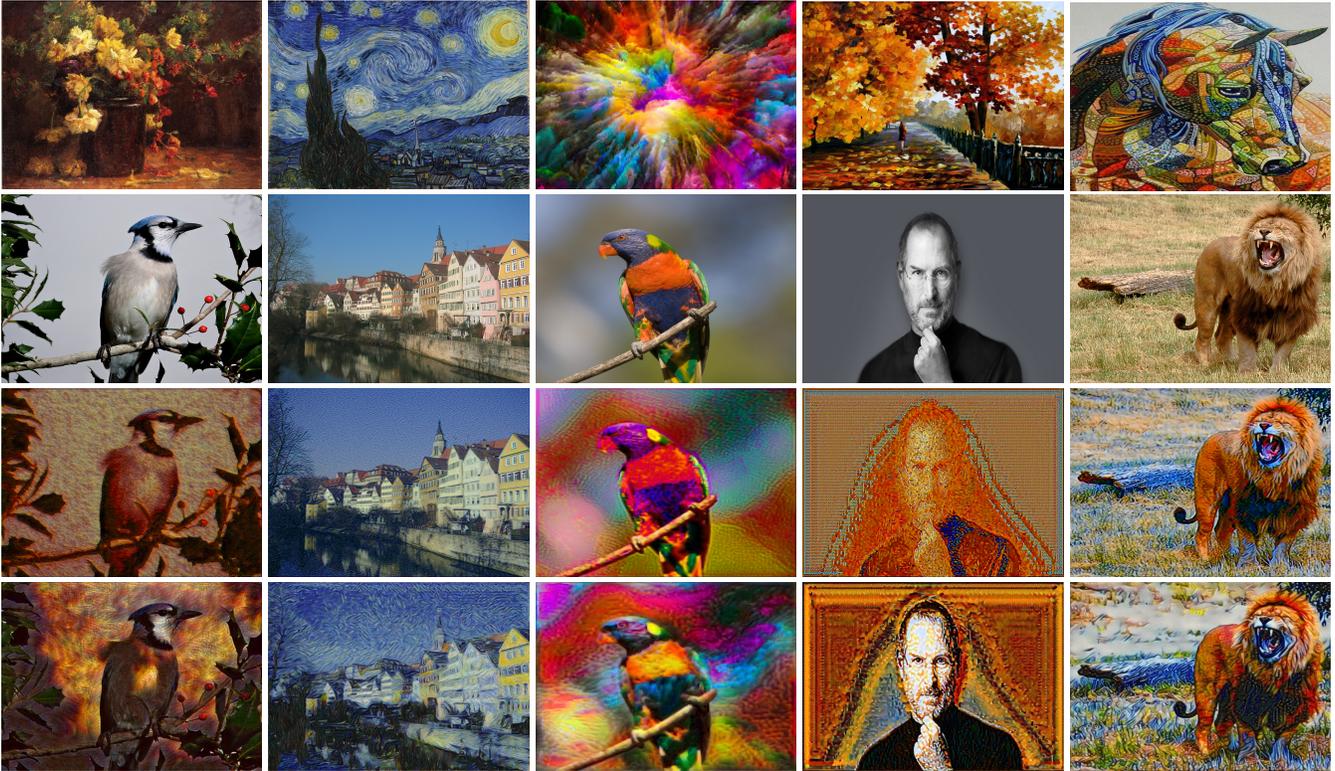


Fig. 8. **Style transfer** using features of ResNet18 trained on CIFAR10 dataset. Robust features obtained using SAT produced more perceptually appealing style transfer as compared to naturally trained (NT) feature space of ResNet18. Comparison is made under the same hyper-parameters and number of iterations (100). (Top to bottom) 1st and 2nd rows show style and target images while 3rd and 4th rows show style transfer using NT and SAT models, respectively.

insights, as follows:

Contribution of proposed losses: We show in Figure 7 that each loss proposed in Eq. 3 contributes towards SAT’s robustness. Individually, the content loss demonstrates better robustness compared to style loss while the classification loss (CE) provides better robustness than content loss. One potential reason for this behavior is that the style loss encodes more abstract information about the target sample compared to content and classification losses. Interestingly, the combination of style and content losses could beat the case when only CE loss is used. Overall, SAT performance increases significantly when adversaries are computed using style, content and boundary information of the target samples.

Convergence Analysis: Each SAT model is trained for 200 epochs. We set the initial learning rate to 0.1 and decrease it by a factor of 0.1 at epochs 50 and 95. We observe that robustness increases with the number of iterations as shown in Table 7. We did not observe a noticeable gain in training beyond 200 epochs.

Number of Epochs (→)	100	150	200	250
Robustness (→)	77.5	80.0	81.1	81.1

TABLE 7
Convergence analysis of SAT.

Computational Training Cost: As mentioned above, SAT takes 200 epochs for convergence which takes around 11.6 hours of training time (see Table 8 for computational

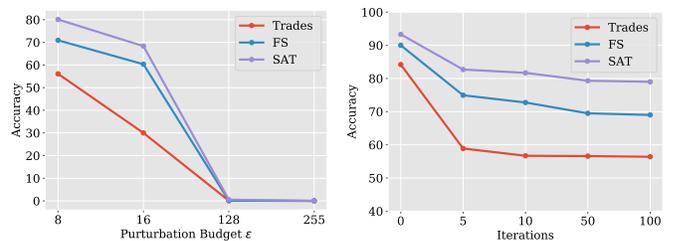


Fig. 9. Defenses are evaluated against large perturbation budget and number of iterations. Left: accuracy drop is shown as perturbation budget goes larger beyond the training regime. Right: performance of each defense against larger number of iteration under ($\epsilon \leq 8/255$).

analysis). This cost is less than the adversarial training approach of [1], Trades and FS. This is because SAT requires only one attack iteration to compute adversaries. However, it needs inference on clean, adversarial and target class samples to update the model.

SAT (ours)	Madry <i>et al.</i> [1]	Trades [7]	FS [8]
11.6	23.6	29.2	29.4

TABLE 8
Training time (hours) on a Tesla V100.

4.4 Sanity Checks for Gradient Obfuscation

Gradient obfuscation or gradient masking refers to the phenomenon where optimization based attacks fail thus leading

Source	MIFGSM		PGD		CW	
	FS [8]	SAT	FS [8]	SAT	FS [8]	SAT
NT	87.5	88.3	89.0	91.5	88.1	90.3
AT	80.8	81.6	80.0	83.5	79.5	82.4

TABLE 9

Black-box robustness evaluation. Adversaries are transferred from naturally and adversarially [8] trained models with the same architectures as of SAT and FS [8] (WideResNet). High accuracy on black-box adversaries indicate model convergence on non-degenerate solution that is without gradient masking.

to high but false adversarial robustness. Athalye *et al.* [40] devise certain tests to evaluate if the defense is relying on gradient masking. We perform the following sanity checks on our defense to show that it does not rely on gradient obfuscation:

- **Robustness to Black-Box Attacks:** If black-box attack (where adversaries are transferred from another model) are stronger than the white-box, this indicates that white-box attack is weak and gradients are being obfuscated. We evaluated SAT (Algorithm 1) under different black-box attacks (Table 9) and the accuracy of our defense remains higher than the white-box attacks (Table 2).
- **Iterative attack should be stronger than single-step:** Another test for gradient masking is that iterative attacks like PGD with small step-size should be more effective than single step attack like FGSM. In all of our evaluations (Tables 1, 2, 3 and 5), PGD with step size $2/255$ is always a stronger attack than FGSM.
- **Robustness should approach to zero for large enough perturbation:** Gradient masking occurs if defense accuracy does not approach to zero for large enough perturbation. Our defense also fulfills this sanity check as its accuracy decreases (see Figure 9) and follow the similar trends like feature scattering [8] on larger perturbations.

4.5 Improved Style Transfer with SAT

Image style transfer works well for VGG features compared with residual connection based models [41]. For the case of residual networks, it has been noted that compared with their naturally trained counterpart, features from adversarially robust models generate more visually appealing images for style transfer [41]. In our case, we observe (Figure 8) that ResNet18 trained on CIFAR10 using SAT performs a better style transfer, compared with a naturally trained ResNet18. This indicates that our adversarial training approach learns representations that can faithfully model the perceptual space.

5 CONCLUSION

We propose to maximize inter-class margins by setting target class samples as priors to adversarial perturbations for the original samples. Our framework pushes the clean images towards randomly selected targets by adding the style, content and boundary information of the target image from other classes in the form of adversarial perturbations within

an allowed perturbation budget. Our framework naturally fits within max-margin learning as it generates positive (adversaries) and negatives (target samples) for clean images. Adversarially trained models using our framework show significant robustness against adversarial attacks (both in white-box and black-box attack scenarios), naturally occurring distributional shifts as well as on common corruptions. Furthermore, robust features obtained via our proposed approach can also be used for style transfer.

REFERENCES

- [1] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=rjzIBfZAb1,2,5,6,7,8>
- [2] G. W. Ding, Y. Sharma, K. Y. C. Lui, and R. Huang, "Max-margin adversarial (mma) training: Direct input space margin maximization through adversarial training," *arXiv preprint arXiv:1812.02637*, 2018. 1, 2
- [3] J. Wang and H. Zhang, "Bilateral adversarial training: Towards fast training of more robust models against adversarial attacks," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 6629–6638. 1, 5
- [4] H. Zhang and J. Wang, "Joint adversarial training: Incorporating both spatial and pixel attacks," *arXiv preprint arXiv:1907.10737*, 2019. 1
- [5] H. Zhang and W. Xu, "Adversarial interpolation training: A simple approach for improving model robustness," 2020. [Online]. Available: <https://openreview.net/forum?id=Syeji0NYvr1,2>
- [6] C.-Y. Lee, S. Xie, P. Gallagher, Z. Zhang, and Z. Tu, "Deeply-supervised nets," 2014. 1
- [7] H. Zhang, Y. Yu, J. Jiao, E. P. Xing, L. E. Ghaoui, and M. I. Jordan, "Theoretically principled trade-off between robustness and accuracy," *arXiv preprint arXiv:1901.08573*, 2019. 1, 2, 4, 5, 6, 7, 8
- [8] H. Zhang and J. Wang, "Defense against adversarial attacks using feature scattering-based adversarial training," in *Advances in Neural Information Processing Systems*, 2019. 1, 2, 4, 5, 6, 7, 8, 9
- [9] D. Hendrycks and T. Dietterich, "Benchmarking neural network robustness to common corruptions and perturbations," *Proceedings of the International Conference on Learning Representations*, 2019. 2, 6
- [10] I. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *International Conference on Learning Representations*, 2015. [Online]. Available: <http://arxiv.org/abs/1412.6572> 2
- [11] A. Kurakin, I. Goodfellow, and S. Bengio, "Adversarial machine learning at scale," *arXiv preprint arXiv:1611.01236*, 2016. 2
- [12] F. Tramr, A. Kurakin, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel, "Ensemble adversarial training: Attacks and defenses," in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=rkZvSe-RZ2>
- [13] N. Carlini and D. Wagner, "Towards evaluating the robustness of neural networks," in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 39–57. 2
- [14] A. Shafahi, M. Najibi, M. A. Ghiasi, Z. Xu, J. Dickerson, C. Studer, L. S. Davis, G. Taylor, and T. Goldstein, "Adversarial training for free!" in *Advances in Neural Information Processing Systems*, 2019, pp. 3353–3364. 2
- [15] E. Wong, L. Rice, and J. Z. Kolter, "Fast is better than free: Revisiting adversarial training," in *International Conference on Learning Representations*, 2020. [Online]. Available: <https://openreview.net/forum?id=BJx040EFvH2>
- [16] A. Lamb, V. Verma, J. Kannala, and Y. Bengio, "Interpolated adversarial training: Achieving robust neural networks without sacrificing accuracy," *arXiv preprint arXiv:1906.06784*, 2019. 2
- [17] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," *arXiv preprint arXiv:1710.09412*, 2017. 2
- [18] V. Verma, A. Lamb, C. Beckham, A. Najafi, I. Mitliagkas, A. Courville, D. Lopez-Paz, and Y. Bengio, "Manifold mixup: Better representations by interpolating hidden states," *arXiv preprint arXiv:1806.05236*, 2018. 2

- [19] A. Mustafa, S. Khan, M. Hayat, R. Goecke, J. Shen, and L. Shao, "Adversarial defense by restricting the hidden space of deep neural networks," in *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 2, 4, 6
- [20] Y. Zhong and W. Deng, "Adversarial learning with margin-based triplet embedding regularization," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 6549–6558. 2
- [21] P. Li, J. Yi, B. Zhou, and L. Zhang, "Improving the robustness of deep neural networks via adversarial training with triplet loss," *arXiv preprint arXiv:1905.11713*, 2019. 2, 7
- [22] C. Mao, Z. Zhong, J. Yang, C. Vondrick, and B. Ray, "Metric learning for adversarial robustness," in *Advances in Neural Information Processing Systems*, 2019, pp. 478–489. 2, 3
- [23] L. A. Gatys, A. S. Ecker, and M. Bethge, "Image style transfer using convolutional neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2414–2423. 3
- [24] J. M. Cohen, E. Rosenfeld, and J. Z. Kolter, "Certified adversarial robustness via randomized smoothing," *arXiv preprint arXiv:1902.02918*, 2019. 4
- [25] H. Kannan, A. Kurakin, and I. J. Goodfellow, "Adversarial logit pairing," *ArXiv*, vol. abs/1803.06373, 2018. 4
- [26] V. Zantedeschi, M.-I. Nicolae, and A. Rawat, "Efficient defenses against adversarial attacks," in *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, 2017, pp. 39–49. 4
- [27] H.-Y. Chen, J.-H. Liang, S.-C. Chang, J.-Y. Pan, Y.-T. Chen, W. Wei, and D.-C. Juan, "Improving adversarial robustness via guided complement entropy," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 4881–4889. 4, 7
- [28] S.-M. Moosavi-Dezfooli, A. Fawzi, and P. Frossard, "Deepfool: a simple and accurate method to fool deep neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2574–2582. 5
- [29] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng, "Reading digits in natural images with unsupervised feature learning," 2011. 4
- [30] A. Krizhevsky, "Learning multiple layers of features from tiny images," Citeseer, Tech. Rep., 2009. 4
- [31] T. Wu, L. Tong, and Y. Vorobeychik, "Defending against physically realizable attacks on image classification," in *International Conference on Learning Representations*, 2020. [Online]. Available: <https://openreview.net/forum?id=H1xscnEKDr> 4, 6
- [32] L. N. Darlow, E. Crowley, A. Antoniou, and A. J. Storkey, "Cinic-10 is not imagenet or cifar-10," *ArXiv*, vol. abs/1810.03505, 2018. 6
- [33] B. Recht, R. Roelofs, L. Schmidt, and V. Shankar, "Do cifar-10 classifiers generalize to cifar-10?" *arXiv preprint arXiv:1806.00451*, 2018. 6
- [34] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, pp. 211–252, 2015. 6
- [35] S. Motiian, M. Piccirilli, D. A. Adjeroh, and G. Doretto, "Unified deep supervised domain adaptation and generalization," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 5715–5725. 7
- [36] P. T. Jackson, A. Atapour-Abarghouei, S. Bonner, T. P. Breckon, and B. Obara, "Style augmentation: Data augmentation via style randomization," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2019, pp. 83–92. 7
- [37] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778. 7
- [38] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014. 7
- [39] Y. Dong, F. Liao, T. Pang, H. Su, J. Zhu, X. Hu, and J. Li, "Boosting adversarial attacks with momentum," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 9185–9193. 7
- [40] A. Athalye, N. Carlini, and D. Wagner, "Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples," in *Proceedings of the 35th International Conference on Machine Learning, ICML 2018*, Jul. 2018. [Online]. Available: <https://arxiv.org/abs/1802.00420> 9
- [41] R. Nakano, "A discussion of 'adversarial examples are not bugs, they are features': Adversarially robust neural

style transfer," *Distill*, 2019, <https://distill.pub/2019/advex-bugs-discussion/response-4>. 9

Muzammal Naseer received the B.Sc. degree (Hons.) in electrical engineering from the University of Punjab and the masters degree in electrical engineering from the King Fahd University of Petroleum and Minerals. He is currently pursuing the Ph.D. degree with Australian National University, where he is a recipient of a competitive postgraduate scholarship. He has published at top venues such as NeurIPS and CVPR. His current research interests are computer vision and machine learning. He received the M.Sc. degree Scholarship from the Ministry of Higher Education, Saudi Arabia, and also the Gold Medal for outstanding performance in the B.Sc. degree.

Salman Khan received the Ph.D. degree from The University of Western Australia, in 2016. His Ph.D. thesis received an honorable mention on the Deans List Award. From 2016 to 2018, he was a Research Scientist with Data61, CSIRO. He has been a Senior Scientist with Inception Institute of Artificial Intelligence, since 2018, and an Adjunct Lecturer with Australian National University, since 2016. He has served as a program committee member for several premier conferences, including CVPR, ICCV, and ECCV. In 2019, he was awarded the outstanding reviewer award at CVPR and the best paper award at ICPRAM 2020. His research interests include computer vision and machine learning.

Munawar Hayat received his PhD from The University of Western Australia (UWA). His PhD thesis received multiple awards, including the Deans List Honorable Mention Award and the Robert Street Prize. After his PhD, he joined IBM Research as a postdoc and then moved to the University of Canberra as an Assistant Professor. He is currently a Senior Scientist at Inception Institute of Artificial Intelligence, UAE. Munawar was granted two US patents, and has published over 30 papers at leading venues in his field, including TPAMI, IJCV, CVPR, ECCV and ICCV. His research interests are in computer vision and machine/deep learning.

Fahad Shahbaz Khan is currently a Lead Scientist at the Inception Institute of Artificial Intelligence (IIAI), Abu Dhabi, United Arab Emirates and an Associate Professor (Universitetslektor + Docent) at Computer Vision Laboratory, Linköping University, Sweden. He received the M.Sc. degree in Intelligent Systems Design from Chalmers University of Technology, Sweden and a Ph.D. degree in Computer Vision from Computer Vision Center Barcelona and Autonomous University of Barcelona, Spain. He has achieved top ranks on various international challenges (Visual Object Tracking VOT: 1st 2014 and 2018, 2nd 2015, 1st 2016; VOT-TIR: 1st 2015 and 2016; OpenCV Tracking: 1st 2015; 1st PASCAL VOC Segmentation and Action Recognition tasks 2010). He received the best paper award in the computer vision track at IEEE ICPR 2016. His research interests include a wide range of topics within computer vision and machine learning. He serves as a regular program committee member for leading computer vision and artificial intelligence conferences such as CVPR, ICCV, and ECCV.

Fatih Porikli (M96SM04F14) received the Ph.D. degree from New York University, in 2002. He was a Distinguished Research Scientist with Mitsubishi Electric Research Laboratories. He is currently a Professor with the Research School of Engineering, Australian National University, and a Chief Scientist with the Global Media Technologies Lab, Huawei, Santa Clara. He has authored over 300 publications, co-edited two books, and invented 66 patents. His research interests include computer vision, pattern recognition, manifold learning, image enhancement, robust and sparse optimization, and online learning with commercial applications in video surveillance, car navigation, robotics, satellite, and medical systems. He was a recipient of the Research and Development 100 Scientist of the Year Award, in 2006. He received five best paper awards at premier IEEE conferences and five other professional prizes. He is serving as an associate editor for several journals for the past 12 years. He has also served in the organizing committees of several flagship conferences, including ICCV, ECCV, and CVPR.