

Exploring the Relationship between Architectural Design and Adversarially Robust Generalization

Aishan Liu¹, Shiyu Tang¹, Siyuan Liang², Ruihao Gong^{1,6}, Boxi Wu³,
Xianglong Liu^{1,4,5}, Dacheng Tao⁷

¹Beihang University, ²Chinese Academy of Sciences, ³Zhejiang University,
⁴Zhongguancun Laboratory, ⁵Hefei Comprehensive National Science Center,
⁶SenseTime, ⁷JD Explore Academy

Abstract

Adversarial training has been demonstrated to be one of the most effective remedies for defending adversarial examples, yet it often suffers from the huge robustness generalization gap on unseen testing adversaries, deemed as the adversarially robust generalization problem. Despite the preliminary understandings devoted to adversarially robust generalization, little is known from the architectural perspective. To bridge the gap, this paper for the first time systematically investigated the relationship between adversarially robust generalization and architectural design. In particular, we comprehensively evaluated 20 most representative adversarially trained architectures on ImageNet and CIFAR-10 datasets towards multiple ℓ_p -norm adversarial attacks. Based on the extensive experiments, we found that, under aligned settings, Vision Transformers (e.g., PVT, CoAtNet) often yield better adversarially robust generalization while CNNs tend to overfit on specific attacks and fail to generalize on multiple adversaries. To better understand the nature behind it, we conduct theoretical analysis via the lens of Rademacher complexity. We revealed the fact that the higher weight sparsity contributes significantly towards the better adversarially robust generalization of Transformers, which can be often achieved by the specially-designed attention blocks. We hope our paper could help to better understand the mechanism for designing robust DNNs. Our model weights can be found at <http://robust.art>.

1. Introduction

DNNs have achieved remarkable performance across a wide variety of applications [15, 28, 41, 79], yet they are vulnerable to adversarial examples (malicious inputs that could fool DNNs by adding human-imperceptible perturbations

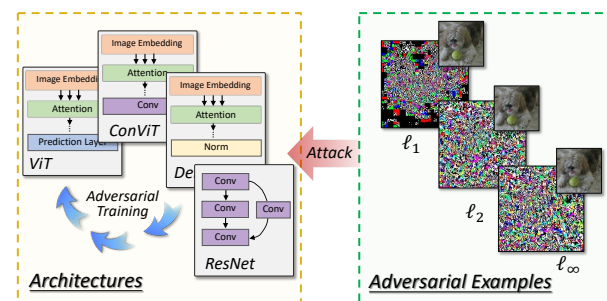


Figure 1. This paper aims to better understand adversarially robust generalizations from the architectural design perspective. We comprehensively evaluated 20 adversarially trained architectures on several datasets toward multiple ℓ_p adversaries and found that Transformers often yield better adversarially robust generalization. We also undertake theoretical analysis through the Rademacher complexity lens to further comprehend this.

[12, 33, 37, 54]). Numerous defensive strategies have been offered to combat this issue. The most effective and promising of these approaches is adversarial training [42, 58], which improves adversarial robustness by adding adversarial examples into training data.

However, adversarial training tends to exhibit a huge generalization gap between training and testing data [53, 59, 76]. Unlike models trained on clean data that usually generalize well on clean testing data, differences between training and testing robustness for adversarially-trained models tend to be large. Moreover, the robustness of adversarial training fails to generalize to unseen adversaries during testing [29, 31]. The aforementioned *adversarially robust generalization problem* puts a barrier in the way of the potential applications of adversarial training in practice, where attacks are often drawn from multiple unknown attacks, therefore attracting intensive research interests. Prior works have studied the characteristic of adversarially robust generalization and proposed many approaches for mitigation (e.g., unlabeled data [4, 47], robust local features [50], genera-

The first three authors contribute equally. Corresponding author: Xianglong Liu, xlliu@buaa.edu.cn.

tive adversarial training [44]), while little is known about adversarially robust generalization in terms of architectural design. Model architecture reflects the inherent nature of model robustness [55]. To better comprehend and create robust DNNs, it is important to carefully explore the links between architectural design and robust generalization.

In this paper, we present the first comprehensive investigation on adversarially robust generalization *w.r.t.* architectural design. In particular, we first provide an extensive evaluation on 20 adversarially trained architectures, including the prevailing Vision Transformers (*e.g.*, ViT, PVT, Swin Transformer) as well as some representative CNNs (*e.g.*, ResNet, VGG) for ImageNet and CIFAR-10 datasets towards multiple adversaries (ℓ_1 , ℓ_2 , ℓ_∞ adversarial attacks). Based on our experiments, we found that, under aligned setting, Vision Transformers (*e.g.*, PVT, CoAtNet) often yield better adversarially robust generalization compared to CNNs; yet CNNs often tend to overfit on specific attacks and fail to generalize on multiple adversaries. To better understand the mechanism, we then conduct theoretical analysis via the lens of Rademacher complexity. We theoretically revealed that higher weight sparsity contributes significantly to the better adversarially robust generalization of Transformers, which can often be achieved by the specially-designed architectural ingredient attention blocks. At last, we also provided more detailed analyses to better understand generalization (*e.g.*, generalization on common corruptions, larger dataset ImageNet) and discussed the potential directions for designing robust architectures. Our **main contributions** can be summarized as:

- We, for the first time, systematically studied 20 adversarially-trained architectures against multiple attacks and revealed the close relationship between architectural design and robust generalization.
- We theoretically revealed that higher weight sparsity contributes to the better adversarially robust generalization of Transformers, which can often be achieved by attention blocks.
- We provide more detailed analyses of the generalizability from several viewpoints and discuss potential pathways that may improve architecture robustness.

We open-sourced the weights of our adversarially-trained model zoo. Together with our analysis, we hope this paper could help researchers build more robust DNN architectures in the future.

2. Related Work

Adversarial Attacks. Adversarial examples are intentionally designed inputs that are imperceptible to human vision but intended to mislead deep learning models [12, 54],

which is denoted as

$$f_\theta(\mathbf{x}_i + \delta_i) \neq \mathbf{y}_i, \text{ s. t. } \delta_i \in \mathbb{B}(\epsilon), \quad (1)$$

where δ_i is the small perturbation bounded by budgets ϵ under a certain distance metric \mathbb{B} (usually ℓ_p -norm) and \mathbf{y}_i is the ground-truth label for the image. A long line of work has been proposed to attack deep learning models [12, 30, 34, 35, 64, 65], which could be roughly split into white-box and black-box attacks based on how they got access to the victim model.

Adversarial Training. Among the proposed defense methods, adversarial training [10, 12, 36, 42, 48, 62, 66, 78] has been demonstrated to be the most effective one. This technique improves the robustness of DNNs by injecting adversarial examples during training, which can be viewed as a min-max optimization. Formally, given a DNN f_θ , input image \mathbf{x} and ground truth label \mathbf{y} sampled from data distribution \mathbb{D} , the standard adversarial training (SAT) is to minimize the adversarial expected risk as

$$\hat{R}(f_\theta) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathbb{D}} \left[\max_{\delta \in \mathbb{B}(\epsilon)} \mathcal{L}(f_\theta, \mathbf{x} + \delta, \mathbf{y}) \right], \quad (2)$$

where $\mathcal{L}(\cdot)$ is the loss function.

Based on SAT, many different types of adversarial training have been suggested to make adversarial robustness even better. For example, Adversarial Logit Pairing (ALP) [23] was proposed to minimize the distance of model logits between clean examples and corresponding adversarial examples; [77] stated the trade-off between adversarial robustness and standard accuracy and proposed TRADES, which introduces a regularization term to induce a shift of the decision boundaries away from the training data points; [63] considered misclassified examples and incorporated a differentiation regularizer into the adversarial loss.

Adversarially Robust Generalization. For adversarial training, extensive studies have shown that there is a large gap between training and testing robustness for both seen and unseen adversaries. Several studies have been conducted to better understand this adversarially robust generalization gap. For example, [73] aimed to analyze adversarially robust generalizations via bias-variance decomposition, and found the bias-variance characteristic of adversarial training. [52] disentangled robustness and generalization by introducing on-manifold adversarial examples. [57] demonstrated that some types of adversarial perturbations (different ℓ_p -norms) are mutually exclusive, and training on one type of perturbation would not generalize well to other types. [20] explored the influence of depth and width on the adversarially robust generalization for WideResNet. Meanwhile, some research aims to improve adversarially robust generalization to adversarial examples or even other unanticipated perturbations. [4, 47] confirmed that more unla-

beled training data could close the adversarially robust generalization gap. Besides the huge generalization gap on adversaries, adversarial training is also demonstrated to not generalize well towards clean data, and extensive evidence has been demonstrated that adversarial robustness may be inherently at odds with natural accuracy [59, 77].

By contrast, this paper takes the first step toward a systematic study and understanding of adversarially robust generalization from an architectural design point of view.

3. Empirical Evaluation

In this section, we first systematically evaluate 20 adversarially-trained architectures on multiple adversaries.

3.1. Evaluation protocols

We first briefly introduce our evaluation protocols. *More details can be found in Supplementary Material.*

Datasets. We conduct extensive experiments on both CIFAR-10 [27] and ImageNette [18]. CIFAR-10 contains 10 classes with input size $32 \times 32 \times 3$; ImageNette is a subset of 10 classes from ImageNet with input size $224 \times 224 \times 3$. We also conduct experiments on large-scale dataset ImageNet [27] with 1,000 classes as shown in Section 5.1.

Architectures. We consider 20 representative model architectures with different categories and designs (*e.g.*, CNN with convolutions, ViT with attentions, hybrids with both attention/convolutions, and newly designed attentions), aiming to find the influential parts. Specifically, we use ViT [11], MLP-Mixer [56], PoolFormer [72], Swin Transformer [39], ViTAE [69], CCT [14], MobileViT [43], CPVT [6], BoTNet [51], CeiT [74], CoAtNet [9], CvT [14], LeViT [13], PVTv2 [61]), ResNet [15], ResNeXt [68], WideResNet [75], PreActResNet [16], VGG [49], and DenseNet [19]. *Note that model sizes differ greatly among architectures, and simply aligning model sizes would cause training collapse. Our studies in Section 3.2 also reveal that model sizes have a comparatively small influence on adversarially robust generalizations. Thus, we primarily focus on discussing the influence of architectural design.*

Training settings. For each model architecture, we conduct standard training (vanilla training) and PGD- ℓ_∞ adversarial training ($\epsilon = 8/255$). To avoid the influence of training techniques on adversarially robust generalization, for all models we use the aligned training settings following [1, 55]. Specifically, for CIFAR-10 we use AutoAugment [8] data augmentation, AdamW [40] optimizer and cosine learning rate with warmup; we set batch size=128 and training epoch=100. For ImageNette we also use AdamW and cosine learning rate with warmup, and set batch size=128 and training epoch=100. Due to the bad convergence of adversarially trained vision transformers using autoaugment on ImageNet [1], we use the standard data augmentation for

ImageNette dataset. To demonstrate that our adversarial-trained models do not suffer from adversarial overfitting [46], we report the robust training/testing accuracy at every ten epochs in Supplementary Material. *We also conduct comparisons for different architectures using optimal settings (see Supplementary Material).*

Evaluation strategy. To evaluate the adversarially robust generalization, we adopt commonly-used adversarial attacks, including PGD- ℓ_1 [42], PGD- ℓ_2 , PGD- ℓ_∞ and AutoAttack- ℓ_∞ attacks [7]. For CIFAR-10 we set $\epsilon = 8/255, 1.5, 40.0$ for ℓ_∞, ℓ_2 and ℓ_1 attacks respectively, while for ImageNette we set $\epsilon = 8/255, 8.0, 1600.0$ for the three attack norms. As for evaluation metrics, following [3], we use the classification accuracy for clean examples and measure the sample-wise worst-case robust accuracy for adversaries, which is the lower bound for robustness under multiple adversarial attacks. The metric is formulated as

$$W(f_\theta, \mathbb{A}) = \frac{1}{n} \sum_{i=1}^n \left\{ \min_{\mathcal{A} \in \mathbb{A}} \mathbf{1} [f_\theta(\mathcal{A}(x_i)) = y_i] \right\}, \quad (3)$$

where (x_i, y_i) is the clean sample pair, \mathcal{A} is attack in the set of multiple adversaries \mathbb{A} , and $\mathbf{1}(\cdot)$ is the indicator function. For the above metrics, the higher the better.

3.2. Empirical results

Here, we show the evaluation results of these architectures on CIFAR-10 and ImageNette. Due to space limitation, our main paper only reports the evaluation results of standard trained models and PGD- ℓ_∞ adversarially trained models for all the architectures on CIFAR-10; the results on ImageNette are shown in Supplementary Material.

As shown in Table 1, for the commonly-adopted ℓ_∞ -norm adversarial attacks, CNNs achieve better robustness than Transformers. However, for ℓ_1 and ℓ_2 -norm adversarial attacks, CNNs (especially ResNet and its variants) have much lower robust accuracy than most of the Transformers (*e.g.*, on CIFAR-10, accuracy under PGD- ℓ_1 attack for ResNet and PVTv2 are 15.90 and 46.14). Considering the much more comprehensive worst-case robust accuracy, some representative Transformers (*e.g.*, PVTv2, CoAtNet) achieve the highest robustness and outperform CNNs (*e.g.*, ResNet) by large margins. We conjecture that CNNs might overfit the ℓ_∞ adversarial attacks during training thus achieving better test accuracy under PGD- ℓ_∞ while worse accuracy under PGD- ℓ_1 and ℓ_2 adversaries. In general, we could reach several **observations** as follows.

❶ Architectural design is highly correlated with adversarially robust generalization, resulting in significant influences (*e.g.*, at most 20+%).

❷ Transformer-based architectures PVTv2 and CoAtNet achieve the best adversarially robust generalization among all architectures. Compared to other ViTs, they contain

Table 1. Evaluation results on CIFAR-10. We use worst-case robust accuracy to measure robust generalization towards adversarial data (higher the better). Vanilla Acc denotes the clean accuracy of standard training. Results are ranked by the worst-case robust accuracy.

Architecture	Params (M)	Vanilla Acc	Clean Acc	PGD- ℓ_∞ Adversarial Training				Worst-case Acc
				PGD- ℓ_∞	AA- ℓ_∞	PGD- ℓ_2	PGD- ℓ_1	
PVTv2	12.40	88.34	75.99	46.48	38.18	35.77	46.14	33.54
CoAtNet	16.99	90.73	77.73	48.27	39.85	33.80	42.30	32.17
ViT	9.78	86.73	78.76	46.02	38.00	30.86	39.27	29.24
CPVT	9.49	90.34	78.57	45.02	36.73	30.15	39.22	28.47
ViTAE	23.18	88.24	75.42	40.53	33.22	29.67	40.02	28.13
MLP-Mixer	0.68	83.43	62.86	38.93	31.81	29.27	36.50	27.42
PoolFormer	11.39	89.26	73.66	46.33	38.93	28.84	34.32	27.36
CCT	3.76	92.27	81.23	49.21	40.97	28.29	34.59	26.82
VGG	14.72	94.01	84.30	50.87	41.66	26.78	31.48	25.32
Swin Transformer	27.42	91.58	80.44	48.61	41.31	26.58	30.47	25.04
LeViT	6.67	89.01	77.10	47.16	39.87	26.28	29.58	25.04
MobileViT	5.00	91.47	77.52	49.51	41.50	26.96	29.35	24.41
BoTNet	18.82	94.16	80.76	51.29	42.95	25.84	27.38	23.15
WideResNet	55.85	96.47	89.54	55.17	44.13	22.55	23.68	20.88
DenseNet	1.12	94.42	83.23	53.06	44.02	22.55	21.87	19.48
PreActResNet	23.50	95.86	87.96	54.85	45.81	18.60	16.46	15.11
CeiT	5.56	85.24	71.55	36.20	28.02	15.31	16.77	14.35
ResNet	23.52	95.60	87.92	54.18	45.40	17.52	15.90	14.32
ResNeXt	9.12	95.64	87.12	51.51	42.66	15.07	13.64	12.18
CvT	19.54	87.81	73.76	41.36	33.67	12.75	9.25	8.76

more attention blocks rather than convolutional layers (ViTAE, CeiT), more global attention rather than window attention (Swin Transformer), and have a hierarchical structure with diverse sizes, which are harder to overfit specific attacks during adversarial training.

④ CNNs (especially ResNet and its variations) tend to overfit trained ℓ_∞ adversarial attacks, making them less resistant to other attacks and less able to generalize. Besides the worst-case robustness, we also calculate the rank based on each attack and take the average, where we reach similar observations (*i.e.*, CoAtNet and PVTv2 rank first, and most ViTs outperform CNNs).

Other adversarial training methods. Aside from the conclusions reached on classic PGD- ℓ_∞ adversarial training, we investigate whether similar conclusions hold up on other variants of adversarial training strategies, where we conduct the same experiments theoretically on CIFAR-10 using TRADES [77]. Overall speaking, as shown in Supplementary Material, we could still reach similar observations, demonstrating architectures’ consistent influence on adversarially robust generalization.

Model size. Besides the architectural design, model capacity (parameter sizes) also influences robustness [1]. We here choose 3 model pairs that have similar model parameters, *i.e.*, ResNeXt (9.12M) and ViT (9.78M), ResNet (23.52M) and ViTAE (23.18M), DenseNet (1.12M) and MLP-Mixer (0.68M), and compare them in pairs. Table 1 shows a similar trend under aligned model sizes, namely that transformers frequently have better adversarially robust generalizations than CNNs. Besides, according to Table 1,

the relationship between model size and adversarially robust generalization is comparatively small, and larger models (*e.g.*, WideResNet) do not necessarily exhibit better generalization than smaller models (*e.g.*, ViT).

Based on the above results, we could observe that architectural design significantly influences an adversarially robust generalization and that Transformers show superiority over other models. In the next section, we do a theoretical analysis to learn more about how it works.

4. Theoretical Understanding

Recently, many studies [25, 60, 71] have theoretically revealed the relationship between adversarial risk and Rademacher complexity, which is a classical metric for measuring generalization error. In this part, we use Rademacher complexity [2] to better understand the above phenomenon and conclusions.

4.1. Overall understanding: weight sparsity

We first define the Rademacher complexity of the set of all possible evaluation functions $f \in \mathcal{F}$ concerning the training dataset S as follows:

$$R_S(\mathcal{F}) = \frac{1}{n} \mathbb{E}_\sigma \left[\sup_{f \in \mathcal{F}} \sum_{i=1}^n \sigma_i f(\mathbf{x}_i) \right], \quad (4)$$

where the random variable $\sigma_i \in [-1, +1]$ with probability 0.5 and \mathbb{E}_σ denotes the expectation. We discuss a linear classifier with binary classification, *i.e.*, the label $\mathbf{y} \pm 1$. We consider imposing different ℓ_p norm on the weight θ




Figure 2. Visualization of weight distribution for adversarially trained models. We calculate the sparsity ratio, which represents the proportion of weights whose absolute value is below a small margin in all weights (we use $0.01 * \text{maximum weight value}$). Note that the ratio does not indicate the number of zero weights. For example, PVTv2/CoAtNet are narrower and have more zero weights, which are sparser than ViT (flatter) and show better robustness.

of classifier f_{θ} and define the set of different weight norm classifiers as the function class $\mathcal{F} = \{f_{\theta} : \|\theta\|_p \leq \Theta\}$. For adversarial training, the function can be written as $\hat{\mathcal{F}} = \{\min_{\delta \in \mathbb{B}(\epsilon)} y f_{\theta}(x + \delta) : \|\theta\|_p \leq \Theta\}$.

Lemma 1. [71] Suppose that $\frac{1}{p} + \frac{1}{q} = 1$, there exists a constant $c \in (0, 1)$ such that

$$\frac{c}{2}(R_S(\mathcal{F}) + \epsilon\Theta \frac{d^{1-\frac{1}{p}}}{\sqrt{n}}) \leq R_S(\hat{\mathcal{F}}) \leq R_S(\mathcal{F}) + \epsilon\Theta \frac{d^{1-\frac{1}{p}}}{\sqrt{n}}, \quad (5)$$

where d denotes the dimension of the input sample, p is the matrix norm of model weights, and q is the norm associated with the input image in Rademacher complexity.

From Lemma 1, we can observe that the Rademacher complexity of adversarial training depends on the classifier f , which implies that adversarially robust generalization is influenced by the architecture of the classifier.

Rademacher complexity relates to the input sample dimension d and the matrix norm p of weight. When $p = 1$, $d^{1-\frac{1}{p}} \rightarrow 1$, the upper and lower bounds of Rademacher complexity are reduced (compared to $p > 1$). More specifically, we can observe that $R_S(\hat{\mathcal{F}})$ has an unavoidable polynomial dimension dependence on d unless the $p = 1$. In other words, imposing the ℓ_1 norm constraints on the weight of the classifier would lead to a sparser weight matrix and reduce the Rademacher complexity of adversarial training. Thus, we visualize the weight distribution by plotting the histogram for multiple adversarially trained architectures on CIFAR-10 in Figure 2. From the weight distribution visu-

alization above, we can observe that for PVTv2 and CoAtNet, most weight values are very close to zero, indicating that the two networks have the **highest weight sparsity**. As for the CNNs, ResNet and its variants (e.g., WideResNet, PreActResNet) exhibit bad adversarially robust generalization due to the overfitting on the trained adversarial attacks, and their weight distributions are also **much denser** than PVTv2 and CoAtNet according to the plots.

4.2. Attention contributes to sparseness

Supposing a fundamental block in Transformers $f_w(x) = \mathbf{W}_1 \phi(\mathbf{A}_1 x)$, where $\mathbf{A}_1 = \text{SoftMax}(\mathbf{Q}\mathbf{K}^T/\sqrt{d} + \mathbf{B})\mathbf{V}$ denotes a Multi-head Self-Attention layer module, \mathbf{W}_1 represents a Multi-Layer Perceptron layer module, and ϕ denotes a non-linear activation function, such as ReLU. \mathbf{B} is the bias term of models. Thus, we extend the derivation from [71] that imposing ℓ_1 constraint in transformer will improve the generalization. In the theory of [45], the model's generalization cannot be directly optimized, so the upper bound of the model's generalization is found step by step through SDP relaxation, and the generalization is improved by minimizing the upper bound of the model. We can find a generalization upper bound for the transformer model by SDP relaxation and optimize this upper bound to get Theorem 1 as follows:

Theorem 1. Suppose the Transformer network function $\mathcal{F} = \{f_w(x) : \mathbf{W} = (\mathbf{A}_1, \mathbf{W}_1), \|\mathbf{A}_1\|_p \leq s_1, \|\mathbf{W}_1\|_p \leq s_2, \|\mathbf{A}_1\|_1 \leq b_1, \|\mathbf{W}_{1,1}\|_1 \leq b_2\}$. $\forall \gamma > 0$, with probability at least $1 - \delta$, we have $\forall f_w \in \mathcal{F}$,




Figure 3. Weight distribution for different architectural ingredients in adversarially trained CoAtNet. (a): whole weights in CoAtNet, (b): weights in convolution blocks, (c): QKV weights in attention layers, and (d): weights in feedforward blocks. The horizontal axis represents the value of model weights, while the vertical axis represents the frequency (density) of the corresponding value.

5. Analyses and Discussions

5.1. More Detailed Analyses

$$\begin{aligned}
& \mathbb{P}_{(x,y) \sim \mathcal{S}} \{ \exists \delta \in \mathbb{B}(\epsilon), \text{ s.t. } \mathbf{y} \neq \arg \max [f_{\mathbf{w}}(\mathbf{x} + \delta)]_{\mathbf{y}'} \} \\
& \leq \frac{1}{n} \sum_{i=1}^n E_i + \frac{1}{\gamma} \left(\frac{4}{n^{3/2}} + \frac{60 \log(n) \log(2d_{\max})}{n} s_1 s_2 C \right) \\
& \quad + \frac{2\epsilon b_1 b_2}{\gamma \sqrt{n}} + 3 \sqrt{\frac{\log(2/p)}{2n}}, \quad (6)
\end{aligned}$$

where $\mathbf{w}_{1,k}$ denotes the k -th column of \mathbf{W}_1 , $C = ((\frac{b_1}{s_1})_{2/3} + (\frac{b_2}{s_2})_{2/3})_{3/2} \|\mathbf{X}\|_F$, $E_i = \mathbb{1}([f_{\mathbf{w}}(\mathbf{x}_i)]_{\mathbf{y}_i} + \frac{\epsilon}{2} \max_{k \in [K], z = \pm 1} \mathbf{P}_{z \geq 0, \text{diag} \mathbf{P} \leq 1} \langle z \mathbf{Q}(\mathbf{w}_{1,k}, \mathbf{A}_1), \mathbf{P} \rangle)$, and $\mathbf{Q}(\mathbf{w}_{1,k}, \mathbf{A}_1) = \begin{bmatrix} 0 & 0 & \mathbf{1}^\top \mathbf{A}_1^\top \mathbf{b} \\ 0 & 0 & \mathbf{A}_1^\top \mathbf{b} \\ \mathbf{b}^\top \mathbf{A}_1 \mathbf{1} & \mathbf{b}^\top \mathbf{A}_1 & 0 \end{bmatrix}$. At this time, there is $\mathbf{b} = \text{diag}(\mathbf{w}_{1,k})$.

Similar to the linear classifier in Lemma 1, if we have an ℓ_1 norm constraint on the matrix of \mathbf{A}_1 , then the generalization bound of the surrogate adversarial loss does not have an explicit dependence on the dimension d and will be reduced. The proof is inspired by [71] and can be found in the supplementary material. Thus, we can draw the following remark.

Remark 1. *The purpose of imposing the ℓ_1 constraint is to increase the sparsity of the weight matrix \mathbf{A}_1 , and we can conclude that the better generalization of Transformers compared with CNNs may come from the sparseness of transformer’s weights caused by attention module \mathbf{A}_1 .*

To further verify the reason for the high sparsity of Transformers, we visualize the weight distribution for different components in CoAtNet (a hybrid Transformer network consisting of convolution and attention), which is shown in Figure 3. We can observe that QKV weights (main weights in the attention layer) have the highest sparsity while the weights of convolution have the lowest sparsity. Hence it is demonstrated that the attention layer could contribute to the sparsity of Transformers, which in turn improves its adversarially robust generalization. These results are consistent with Theorem 1 and the Remark 1.

Weight loss landscape analysis. Previous works [21, 24, 26] have demonstrated the strong relation between model generalization ability and model weight loss landscape, and flatter loss landscape usually leads to better generalization ability. Some recent works [5, 22, 38, 67, 70] further studied this relationship under robust generalization. For example, [38, 67, 70] have shown that adversarial training would sharpen the weight loss landscape; [5] demonstrated that ViT and MLP-Mixer converge to sharper regions than ResNet under standard training. By contrast, we investigate the weight loss landscape for different model architectures under standard and adversarial training to help better understand adversarially robust generalization. Specifically, we visualize the weight loss landscape for 2 representative architectures (ResNet and PVTv2) under standard training, and PGD- ℓ_∞ adversarial training ($\epsilon = 8/255$) leveraging the filter normalization method [32].

As shown in Figure 5, we could draw several observations: ❶ the loss landscapes for PVTv2 under standard training are sharper than ResNet, which is consistent with [5]; ❷ adversarial training make the loss landscape for ResNet sharper, which is consistent to [38, 67, 70]; ❸ however, different to ResNet, the loss landscapes for PVTv2 become flatter after adversarial training. To a certain extent, these visualization results could explain our adversarially robust generalization results in Section 3.2, where PVTv2 shows comparatively better adversarially robust generalization than ResNet.

Observations on ImageNet. We further explore whether the tendency between Transformer and CNNs on adversarially robust generalization still exists on large-scale dataset ImageNet. In particular, we choose the representative Transformer ViT and typical CNN ResNet for experiments on ImageNet. For fair comparisons, we trained both models using similar settings as Section 3.1.

On ImageNet, the Worst-Case Acc of ViT is 24.77 still better than 20.08 in the ResNet, which indicates a better

adversarially robust generalization of Transformers. However, the gap in the generalization ability between them shrinks as the training dataset size increases. We further analyze the weight sparsity of these adversarially-trained ViT and ResNet models, where we found that from CIFAR-10 to ImageNet the sparsity of ViT parameters improves from 0.65 to 0.95, and the sparsity of ResNet increases from 0.15 to 0.92. Recall our analysis in Equation 1 where n denotes the number of samples, and increasing sample numbers would decrease the minimum upper bound of the Rademacher complexity. Thus, we can conclude that the adversarially robust generalization of the model is also positively correlated with the dataset size. The generalization ability of both ViT and ResNet is generally improved by increasing the size of datasets in the same experimental setup. We speculate that the diversity of the dataset increases the sparsity of the model, and the data volume and richness of CIFAR-10 are far less than that of ImageNet. *More results can be found in the Supplementary Material.*

5.2. Potential Pathways

In this part, we further provide some potential pathways for robust architectures in the future. Unless otherwise specified, the experiments in this part adopt the same experimental settings from Section 3.1.

Imposing ℓ_1 regularization for sparsity. In Section 4, we revealed that weight sparsity has a close relation with adversarially robust generalization. We then explore: *could we improve adversarially robust generalization by explicitly improving the sparsity via weight regularization (e.g., ℓ_1 regularization) during training?*

To answer this question, we choose ResNet on CIFAR-10 and impose ℓ_1 regularization during adversarial training with different regularization strengths. From Figure 4, we can observe that imposing ℓ_1 regularization can indeed improve the adversarially robust generalization. However, the improvement on adversarially robust generalization is comparatively small (*i.e.*, +2%) considering that ResNet may fall behind some Transformers for 10-20% as shown in Table 1. Moreover, we could also notice the following obstacles: ❶ with the increasing of ℓ_1 regularization strength, model clean accuracy drops, which indicates that there may exist a trade-off between clean example accuracy and adversarially robust generalization; and ❷ robust accuracy is highly sensitive to the ℓ_1 regularization hyper-parameters, which may entail the careful adoption in practice considering the different model architectures and sizes.

Moreover, we also conduct experiments on directly changing the weight sparsity of Transformers during training, which shows a similar tendency (*Supplementary Material*). The above analyses imply that *to improve adversarially robust generalization, it would be better to introduce weight sparsity by designing architectural ingredients (e.g.,*




Figure 4. Adversarial training on ResNet with different ℓ_1 regularization strengths to improve sparsity.

attention layers) instead of adding training regularization.

Hybrid architecture with increased sparsity. Since attention layers contribute sparsity to ViTs, we investigate the influence of different attention layer numbers on adversarially robust generalization. In particular, we adopt CoAtNet architecture, which is a hybrid architecture containing four network stages and each stage could be set as either the attention layer (denote T) or convolution layer (denote C). We choose 4 settings and evaluate their adversarially robust generalization on CIFAR-10. According to Table 2, we highlight an interesting observation that the increasing attention layer improves weight sparsity, which might be beneficial to improving adversarially robust generalization.

Patch size as receptive fields. Patch splitting and embedding is an important procedure for ViTs, we, therefore, investigate the patch size on adversarially robust generalization. Specifically, we select 2 representative Transformers (ViT and PVTv2) and conduct evaluations on different patch sizes after adversarial training on CIFAR-10. As shown in Table 3, we can observe that with the decreasing of patch size, ViT and PVTv2 both achieve better accuracy on clean examples and PGD- ℓ_∞ adversarial examples (the perturbations that are trained on), while tend to show worse results towards PGD- ℓ_1 and PGD- ℓ_2 adversarial examples. Similar results can be found in Section 3.2, where Swin Transformer (which has window attention instead of global attention) performs better on ℓ_∞ noises while worse on ℓ_1 and ℓ_2 compared to ViT. In other words, *Transformers tend to overfit the trained adversarial attacks with the decreasing of patch size.* Intuitively, decreasing patch size would make each pixel interact with the closer pixels thus making the whole model show the locality characteristic like CNNs, which are easier to fit the texture of trained adversarial examples yet harder to generalize on other attacks.

Considering generalization on common corruptions. In our main experiments, we discuss the adversarially robust generalization towards different adversarial data distributions, while the generalization of adversarial training

Similar to [9], we do not consider the first stage to be the attention layer due to the huge memory consumption.

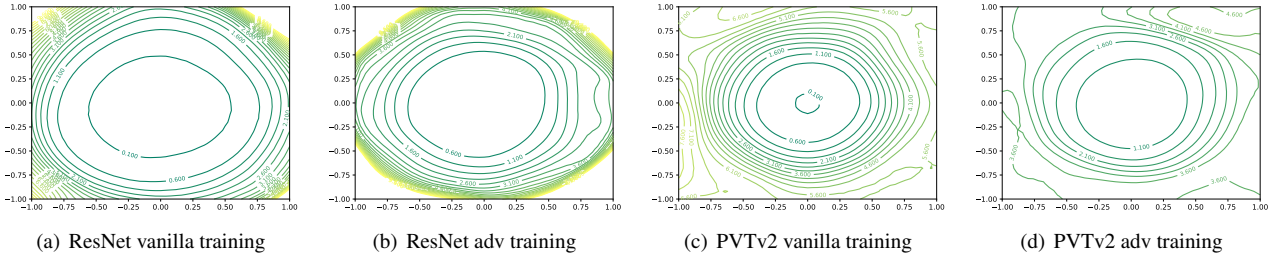


Figure 5. The weight loss landscape visualization for ResNet and PVTv2. The horizontal and vertical axes represent the neighborhood around the trained weights, and the contour means all model weights that have the same corresponding cross-entropy loss.

Table 2. The influence of convolution-attention combinations on robust generalization. C means convolution layer and T means attention layer, e.g., CTTT means the first stage is the convolution layer and the last three stages are attention layer.

Architecture	Params (M)	Vanilla	PGD- ℓ_∞ Adversarial Training					
			Clean	PGD- ℓ_∞	AA- ℓ_∞	PGD- ℓ_2	PGD- ℓ_1	Worst-case
CoAtNet-CTTT	17.36	90.83	74.25	45.70	37.27	31.95	40.15	30.30
CoAtNet-CTTC	17.02	91.13	77.84	46.61	37.89	31.61	39.88	29.96
CoAtNet-CTCC	16.38	90.69	78.71	42.36	34.56	27.84	37.19	26.69
CoAtNet-CCCC	16.02	91.41	79.14	43.71	35.59	29.03	38.68	27.64

Table 3. The influence of Transformer patch size on robust generalization for ViT and PVTv2. For each architecture we choose 3 different patch size (8, 4, 2 for ViT and 4, 2, 1 for PVT).

Architecture	Patch Size	Vanilla	PGD- ℓ_∞ Adversarial Training					
			Clean	PGD- ℓ_∞	AA- ℓ_∞	PGD- ℓ_2	PGD- ℓ_1	Worst-case
PVTv2	$p = 4$	88.34	75.99	46.48	38.18	35.77	46.14	33.54
	$p = 2$	93.03	83.80	52.34	44.04	32.49	39.63	31.16
	$p = 1$	94.60	87.50	54.59	46.58	23.47	24.76	21.10
ViT	$p = 8$	82.30	72.39	42.77	35.04	32.74	42.61	30.72
	$p = 4$	86.73	78.76	46.02	38.00	30.86	39.27	29.24
	$p = 2$	85.99	77.37	45.45	37.95	25.36	30.15	23.78

towards common corruptions is also of great significance. Thus, for all the 20 adversarially trained architectures, we also evaluate their performance on clean examples and common corruptions using CIFAR-10-C and ImageNette-C dataset [17] that consists of 15 diverse types of corruptions (e.g., fog, gaussian noise).

As shown in Table 4, we can observe that as for generalization of adversarial training towards clean examples and corruptions on CIFAR-10, CNNs behave better than Transformers, showing a much smaller clean accuracy drop (e.g., 95.60 \rightarrow 87.92 for ResNet, while 91.58 \rightarrow 80.44 for Swin Transformer) and better CIFAR-10-C accuracy; while on ImageNette some Transformers (e.g., CoAtNet, ViTAE) surpass CNNs and achieve the best robust generalization on clean examples and corruptions. It is important to note that our studies differ from recent studies that primarily benchmark Transformers and CNNs robustness on ImageNet under standard training [80], where we focus on adversarial training. Thus, the generalization of different adversarially-trained architectures towards common corruptions and clean examples is still ambiguous. We conjecture that ❶ common corruptions and adversarial examples might be drawn from different distributions with a comparatively

Table 4. Generalization of adversarially-trained models towards clean data and common corruptions on CIFAR-10. For common corruptions, we leverage mean Corruption Error (mCE) [17] and use 1-mCE as our CIFAR-10-C Acc (the higher the better). Results on ImageNette are in Supplementary Material.

Architecture	CIFAR-10 Dataset			
	Vanilla Training		PGD- ℓ_∞ Training	
	Vanilla Acc	CIFAR-C Acc	Clean Acc	CIFAR-C Acc
WideResNet	96.47	83.91	89.54	81.48
ResNet	95.60	81.20	87.92	79.24
PreActResNet	95.86	82.18	87.96	78.99
ResNeXt	95.64	80.43	87.12	77.76
VGG	94.01	81.22	84.30	75.85
DenseNet	94.42	79.73	83.23	74.60
BoTNet	94.16	81.04	80.76	72.72
CCT	92.27	78.99	81.23	72.56
ViT	86.73	77.06	78.76	71.87
Swin Transformer	91.58	77.61	80.44	71.36
CPVT	90.34	79.66	78.57	70.74
LeViT	89.01	78.31	77.10	70.48
CoAtNet	90.73	79.91	77.73	70.27
MobileViT	91.47	80.48	77.52	70.15
PVTv2	88.34	79.84	75.99	69.12
ViTAE	88.24	75.86	75.42	67.58
PoolFormer	89.26	77.57	73.66	66.45
CVT	87.81	75.10	73.76	66.28
CeIT	85.24	73.99	71.55	65.07
MLP-Mixer	83.43	70.70	62.86	57.09

huge gap, and it would be highly non-trivial to achieve a universally better architecture after adversarial training; and ❷ there exists several factors that may influence the robustness on common corruptions (e.g., datasets, input size) [80].

6. Conclusion

We systematically examined the relationship between adversarially robust generalization and architectural design. Based on the extensive experiments, we found that Vision Transformers (e.g., PVT, CoAtNet) often yield better adversarially robust generalization while CNNs tend to overfit on specific attacks and fail to generalize on multiple adversaries. We then conduct theoretical analysis using Rademacher complexity, where we found that the higher weight sparsity contributes significantly towards the better adversarially robust generalization of Transformers, which can be often achieved by the attention layers.

Acknowledgement. This work was supported by the National Key Research and Development Plan of China (2021ZD0110601), the National Natural Science Foundation of China (62022009 and 62206009), and the State Key Laboratory of Software Development Environment.

References

- [1] Y. Bai, J. Mei, A. L. Yuille, and C. Xie. Are transformers more robust than cnns? *Advances in Neural Information Processing Systems*, 34, 2021. 3, 4
- [2] P. L. Bartlett and S. Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002. 4
- [3] N. Carlini, A. Athalye, N. Papernot, W. Brendel, J. Rauber, D. Tsipras, I. Goodfellow, A. Madry, and A. Kurakin. On evaluating adversarial robustness. *arXiv preprint arXiv:1902.06705*, 2019. 3
- [4] Y. Carmon, A. Raghunathan, L. Schmidt, J. C. Duchi, and P. S. Liang. Unlabeled data improves adversarial robustness. *Advances in Neural Information Processing Systems*, 32, 2019. 1, 2
- [5] X. Chen, C.-J. Hsieh, and B. Gong. When vision transformers outperform resnets without pre-training or strong data augmentations. *arXiv preprint arXiv:2106.01548*, 2021. 6
- [6] X. Chu, Z. Tian, B. Zhang, X. Wang, X. Wei, H. Xia, and C. Shen. Conditional positional encodings for vision transformers. *arXiv preprint arXiv:2102.10882*, 2021. 3
- [7] F. Croce and M. Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *International conference on machine learning*, pages 2206–2216. PMLR, 2020. 3
- [8] E. D. Cubuk, B. Zoph, D. Mane, V. Vasudevan, and Q. V. Le. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*, 2018. 3
- [9] Z. Dai, H. Liu, Q. V. Le, and M. Tan. Coatnet: Marrying convolution and attention for all data sizes. *Advances in Neural Information Processing Systems*, 34:3965–3977, 2021. 3, 7
- [10] G. W. Ding, Y. Sharma, K. Y. C. Lui, and R. Huang. Mma training: Direct input space margin maximization through adversarial training. In *International Conference on Learning Representations*, 2019. 2
- [11] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. 3
- [12] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 1, 2
- [13] B. Graham, A. El-Nouby, H. Touvron, P. Stock, A. Joulin, H. Jégou, and M. Douze. Levit: a vision transformer in convnet’s clothing for faster inference. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12259–12269, 2021. 3
- [14] A. Hassani, S. Walton, N. Shah, A. Abuduweili, J. Li, and H. Shi. Escaping the big data paradigm with compact transformers. *arXiv preprint arXiv:2104.05704*, 2021. 3
- [15] 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*, pages 770–778, 2016. 1, 3
- [16] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. 3
- [17] D. Hendrycks and T. Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2018. 8
- [18] J. Howard. Imagenette. <https://github.com/fastai/imagenette>. 3
- [19] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 3
- [20] H. Huang, Y. Wang, S. Erfani, Q. Gu, J. Bailey, and X. Ma. Exploring architectural ingredients of adversarially robust deep neural networks. *Advances in Neural Information Processing Systems*, 34, 2021. 2
- [21] S. Jastrzębski, Z. Kenton, N. Ballas, A. Fischer, Y. Bengio, and A. Storkey. On the relation between the sharpest directions of dnn loss and the SGD step length. In *International Conference on Learning Representations*, 2018. 6
- [22] S. Kanai, M. Yamada, H. Takahashi, Y. Yamanaka, and Y. Ida. Smoothness analysis of adversarial training. *arXiv preprint arXiv:2103.01400*, 2021. 6
- [23] H. Kannan, A. Kurakin, and I. Goodfellow. Adversarial logit pairing. *arXiv preprint arXiv:1803.06373*, 2018. 2
- [24] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016. 6
- [25] J. Khim and P.-L. Loh. Adversarial risk bounds via function transformation. *arXiv preprint arXiv:1810.09519*, 2018. 4

- [26] B. Kleinberg, Y. Li, and Y. Yuan. An alternative view: When does sgd escape local minima? In *International Conference on Machine Learning*, pages 2698–2707. PMLR, 2018. 6
- [27] A. Krizhevsky, G. Hinton, et al. Learning multiple layers of features from tiny images. 2009. 3
- [28] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. 1
- [29] A. Kurakin, I. Goodfellow, and S. Bengio. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, 2016. 1
- [30] A. Kurakin, I. J. Goodfellow, and S. Bengio. Adversarial examples in the physical world. In *Artificial intelligence safety and security*, pages 99–112. Chapman and Hall/CRC, 2018. 2
- [31] C. Laidlaw, S. Singla, and S. Feizi. Perceptual adversarial robustness: Defense against unseen threat models. In *International Conference on Learning Representations*, 2020. 1
- [32] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein. Visualizing the loss landscape of neural nets. *Advances in neural information processing systems*, 31, 2018. 6
- [33] A. Liu, J. Guo, J. Wang, S. Liang, R. Tao, W. Zhou, C. Liu, X. Liu, and D. Tao. X-adv: Physical adversarial object attacks against x-ray prohibited item detection. In *USENIX Security Symposium*, 2023. 1
- [34] A. Liu, T. Huang, X. Liu, Y. Xu, Y. Ma, X. Chen, S. J. Maybank, and D. Tao. Spatiotemporal attacks for embodied agents. In *European Conference on Computer Vision*, 2020. 2
- [35] A. Liu, X. Liu, J. Fan, Y. Ma, A. Zhang, H. Xie, and D. Tao. Perceptual-sensitive gan for generating adversarial patches. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 1028–1035, 2019. 2
- [36] A. Liu, X. Liu, C. Zhang, H. Yu, Q. Liu, and J. He. Training robust deep neural networks via adversarial noise propagation. *IEEE Transactions on Image Processing*, 2021. 2
- [37] A. Liu, J. Wang, X. Liu, C. Zhang, B. Cao, and H. Yu. Bias-based universal adversarial patch attack for automatic check-out. In *European Conference on Computer Vision*, 2020. 1
- [38] C. Liu, M. Salzmann, T. Lin, R. Tomioka, and S. Ssstrunk. On the loss landscape of adversarial training: Identifying challenges and how to overcome them. *Advances in Neural Information Processing Systems*, 33:21476–21487, 2020. 6
- [39] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 3
- [40] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018. 3
- [41] Y. Ma, X. Liu, S. Bai, L. Wang, A. Liu, D. Tao, and E. R. Hancock. Regionwise generative adversarial image inpainting for large missing areas. *IEEE Transactions on Cybernetics*, 2022. 1
- [42] 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. 1, 2, 3
- [43] S. Mehta and M. Rastegari. Mobilevit: light-weight, general-purpose, and mobile-friendly vision transformer. *arXiv preprint arXiv:2110.02178*, 2021. 3
- [44] O. Poursaeed, T. Jiang, H. Yang, S. Belongie, and S.-N. Lim. Robustness and generalization via generative adversarial training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15711–15720, 2021. 2
- [45] A. Raghunathan, J. Steinhardt, and P. Liang. Certified defenses against adversarial examples. *arXiv preprint arXiv:1801.09344*, 2018. 5
- [46] L. Rice, E. Wong, and Z. Kolter. Overfitting in adversarially robust deep learning. In *International Conference on Machine Learning*. PMLR, 2020. 3
- [47] L. Schmidt, S. Santurkar, D. Tsipras, K. Talwar, and A. Madry. Adversarially robust generalization requires more data. *Advances in neural information processing systems*, 31, 2018. 1, 2
- [48] 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! *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [49] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3
- [50] C. Song, K. He, J. Lin, L. Wang, and J. E. Hopcroft. Robust local features for improving the generalization of adversarial training. In *International Conference on Learning Representations*, 2019. 1
- [51] A. Srinivas, T.-Y. Lin, N. Parmar, J. Shlens, P. Abbeel, and A. Vaswani. Bottleneck transformers for visual recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16519–16529, 2021. 3

- [52] D. Stutz, M. Hein, and B. Schiele. Disentangling adversarial robustness and generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6976–6987, 2019. [2](#)
- [53] D. Su, H. Zhang, H. Chen, J. Yi, P.-Y. Chen, and Y. Gao. Is robustness the cost of accuracy?—a comprehensive study on the robustness of 18 deep image classification models. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 631–648, 2018. [1](#)
- [54] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013. [1](#), [2](#)
- [55] S. Tang, R. Gong, Y. Wang, A. Liu, J. Wang, X. Chen, F. Yu, X. Liu, D. Song, A. Yuille, et al. Robustart: Benchmarking robustness on architecture design and training techniques. *arXiv preprint arXiv:2109.05211*, 2021. [2](#), [3](#)
- [56] I. O. Tolstikhin, N. Houlsby, A. Kolesnikov, L. Beyer, X. Zhai, T. Unterthiner, J. Yung, A. Steiner, D. Keysers, J. Uszkoreit, et al. Mlp-mixer: An all-mlp architecture for vision. *Advances in Neural Information Processing Systems*, 34, 2021. [3](#)
- [57] F. Tramer and D. Boneh. Adversarial training and robustness for multiple perturbations. *Advances in Neural Information Processing Systems*, 32, 2019. [2](#)
- [58] F. Tramèr, A. Kurakin, N. Papernot, I. Goodfellow, D. Boneh, and P. McDaniel. Ensemble adversarial training: Attacks and defenses. In *International Conference on Learning Representations*, 2018. [1](#)
- [59] D. Tsipras, S. Santurkar, L. Engstrom, A. Turner, and A. Madry. Robustness may be at odds with accuracy. In *International Conference on Learning Representations*, 2018. [1](#), [3](#)
- [60] Z. Tu, J. Zhang, and D. Tao. Theoretical analysis of adversarial learning: A minimax approach. *Advances in Neural Information Processing Systems*, 32, 2019. [4](#)
- [61] W. Wang, E. Xie, X. Li, D. Fan, K. Song, D. Liang, T. Lu, P. Luo, and L. Shao. Pvtv2: Improved baselines with pyramid vision transformer. *CoRR*, abs/2106.13797, 2021. [3](#)
- [62] Y. Wang, X. Ma, J. Bailey, J. Yi, B. Zhou, and Q. Gu. On the convergence and robustness of adversarial training. In *International Conference on Machine Learning*, pages 6586–6595. PMLR, 2019. [2](#)
- [63] Y. Wang, D. Zou, J. Yi, J. Bailey, X. Ma, and Q. Gu. Improving adversarial robustness requires revisiting misclassified examples. In *International Conference on Learning Representations*, 2019. [2](#)
- [64] Z. Wei, J. Chen, M. Goldblum, Z. Wu, T. Goldstein, and Y.-G. Jiang. Towards transferable adversarial attacks on vision transformers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2668–2676, 2022. [2](#)
- [65] Z. Wei, J. Chen, Z. Wu, and Y.-G. Jiang. Boosting the transferability of video adversarial examples via temporal translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2659–2667, 2022. [2](#)
- [66] E. Wong, L. Rice, and J. Z. Kolter. Fast is better than free: Revisiting adversarial training. In *International Conference on Learning Representations*, 2019. [2](#)
- [67] D. Wu, S.-T. Xia, and Y. Wang. Adversarial weight perturbation helps robust generalization. *Advances in Neural Information Processing Systems*, 33:2958–2969, 2020. [6](#)
- [68] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. [3](#)
- [69] Y. Xu, Q. Zhang, J. Zhang, and D. Tao. Vitae: Vision transformer advanced by exploring intrinsic inductive bias. *Advances in Neural Information Processing Systems*, 34, 2021. [3](#)
- [70] M. Yamada, S. Kanai, T. Iwata, T. Takahashi, Y. Yamanaka, H. Takahashi, and A. Kumagai. Adversarial training makes weight loss landscape sharper in logistic regression. *arXiv preprint arXiv:2102.02950*, 2021. [6](#)
- [71] D. Yin, R. Kannan, and P. Bartlett. Rademacher complexity for adversarially robust generalization. In *International conference on machine learning*, pages 7085–7094. PMLR, 2019. [4](#), [5](#), [6](#)
- [72] W. Yu, M. Luo, P. Zhou, C. Si, Y. Zhou, X. Wang, J. Feng, and S. Yan. Metaformer is actually what you need for vision. *arXiv preprint arXiv:2111.11418*, 2021. [3](#)
- [73] Y. Yu, Z. Yang, E. Dobriban, J. Steinhardt, and Y. Ma. Understanding generalization in adversarial training via the bias-variance decomposition. *arXiv preprint arXiv:2103.09947*, 2021. [2](#)
- [74] K. Yuan, S. Guo, Z. Liu, A. Zhou, F. Yu, and W. Wu. Incorporating convolution designs into visual transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 579–588, 2021. [3](#)
- [75] S. Zagoruyko and N. Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. [3](#)

- [76] H. Zhang, H. Chen, Z. Song, D. Boning, I. S. Dhillon, and C.-J. Hsieh. The limitations of adversarial training and the blind-spot attack. In *International Conference on Learning Representations*, 2018. [1](#)
- [77] H. Zhang, Y. Yu, J. Jiao, E. Xing, L. El Ghaoui, and M. Jordan. Theoretically principled trade-off between robustness and accuracy. In *International conference on machine learning*, pages 7472–7482. PMLR, 2019. [2](#), [3](#), [4](#)
- [78] J. Zhang, X. Xu, B. Han, G. Niu, L. Cui, M. Sugiyama, and M. Kankanhalli. Attacks which do not kill training make adversarial learning stronger. In *International conference on machine learning*, pages 11278–11287. PMLR, 2020. [2](#)
- [79] X. Zhao, X. Liu, Y. Ma, S. Bai, Y. Shen, Z. Hao, and A. Liu. Temporal speciation network for few-shot object detection. *IEEE Transactions on Multimedia*, 2023. [1](#)
- [80] D. Zhou, Z. Yu, E. Xie, C. Xiao, A. Anandkumar, J. Feng, and J. M. Alvarez. Understanding the robustness in vision transformers. *arXiv preprint arXiv:2204.12451*, 2022. [8](#)