

# Adversarial Robustness Comparison of Vision Transformer and MLP-Mixer to CNNs

Philipp Benz\*  
phibenz@gmail.com

Soomin Ham\*  
smham@kaist.ac.kr

Chaoning Zhang\*  
chaoningzhang1990@gmail.com

Adil Karjauv  
mikolez@gmail.com

In So Kweon  
iskweon77@kaist.ac.kr

Robotics and Computer Vision Lab  
Korea Advanced Institute of Science  
and Technology (KAIST)  
Daejeon, South Korea

\*Equal contribution

---

## Abstract

Convolutional Neural Networks (CNNs) have become the *de facto gold standard* in computer vision applications in the past years. Recently, however, new model architectures have been proposed challenging the *status quo*. The Vision Transformer (ViT) relies solely on attention modules, while the MLP-Mixer architecture substitutes the self-attention modules with Multi-Layer Perceptrons (MLPs). Despite their great success, CNNs have been widely known to be vulnerable to adversarial attacks, causing serious concerns for security-sensitive applications. Thus, it is critical for the community to know whether the newly proposed ViT and MLP-Mixer are also vulnerable to adversarial attacks. To this end, we empirically evaluate their adversarial robustness under several adversarial attack setups and benchmark them against the widely used CNNs. Overall, we find that the two architectures, especially ViT, are more robust than their CNN models. Using a toy example, we also provide empirical evidence that the lower adversarial robustness of CNNs can be partially attributed to their shift-invariant property. Our frequency analysis suggests that the most robust ViT architectures tend to rely more on low-frequency features compared with CNNs. Additionally, we have an intriguing finding that MLP-Mixer is extremely vulnerable to universal adversarial perturbations. Code: [https://github.com/phibenz/robustness\\_comparison\\_vit\\_mlp-mixer\\_cnn](https://github.com/phibenz/robustness_comparison_vit_mlp-mixer_cnn).

## 1 Introduction

Convolutional Neural Networks (CNNs) [37] have been the *gold standard* architecture in computer vision. In Natural Language Processing (NLP), however, attention-based transformers are the dominant go-to model architecture [13, 53, 56]. Various attempts have been made to apply such transformer architectures to computer vision tasks [8, 10, 53, 58]. A

breakthrough moment was achieved with the advent of the Vision Transformer (ViT) [15], presenting a transformer architecture achieving comparable performance to state-of-the-art CNN architectures. Recently, another MLP-Mixer model architecture [66] which does not rely on convolutions or self-attention, has been presented competing with CNN and ViT. Following [66], MLP-Mixer is termed Mixer in the remainder of this work for simplicity.

Despite the success of CNNs, they are widely known to be vulnerable to adversarial examples [18, 65] whose small additive perturbations of the input cause the CNN to misclassify a sample. This vulnerability causes serious concerns in security-sensitive applications, and thus it is also important to know whether the recently proposed ViT and Mixer are also vulnerable to adversarial attacks. This work sets out to evaluate the adversarial vulnerability of ViT and Mixer architectures and compare their robustness against the CNN models. Therefore, a wide range of adversarial attack methods has been adopted for a comprehensive study. Specifically, first, the performance of the different architectures is compared under the white-box attack, where an adversary has full knowledge of the model parameters to attack. Overall, the two newly proposed architectures, especially ViT, exhibit significantly higher robustness than CNNs against adversarial examples. We further compare their robustness under both query-based and transfer-based black-box attacks. In both cases, we observe a similar trend that ViT is the most robust architecture while CNN is the least robust.

To facilitate the understanding of why CNN is more vulnerable, we design a toy task of binary classification where each class is only represented by a single image. The image from each class has either a vertical or horizontal black stripe in the center. We find that the adversarial example for a CNN exhibits repetitive stripes over the image, while that of an FC network mainly exhibits a single stripe in the center. This observation indicates that the vulnerability of a CNN might be partially attributed to the fact that a CNN, which exploits local connections and shared weights by convolving kernels, is shift-invariant [88, 83]. We also attempt to provide an analysis from the perspective of frequency, investigating whether the different model architectures are biased toward learning more high-frequency or low-frequency features. We find that the ViT seems to learn more low-frequency features, while the CNN is biased towards high-frequency features. Finally, we also investigate their robustness against common corruptions [26] and universal adversarial perturbations [47].

## 2 Related Work

**Beyond CNNs for vision applications.** In Natural Language Processing (NLP), transformers [70], solely based on the attention mechanisms, are the predominant model architecture [13, 65, 66]. In contrast, CNNs have been the *de facto* standard in deep learning for vision applications, while the application of transformers to vision tasks is an emerging trend [8, 10, 63, 68]. The Vision Transformer (ViT) [15] was recently introduced, demonstrating that transformers can achieve state-of-the-art performance, by sequencing the images into patches and pre-training the model on large amounts of data. To address the data issue, DeiT [67] introduced a teacher-student strategy specific to transformers and trained a transformer architecture only on the ImageNet-1K dataset. Concurrently, the T2T-ViT had been proposed [78] introducing an advanced Tokens-to-Tokens strategy. Further works are attempting to extend the ViT architecture to increase the efficiency and performance of transformer architectures [10, 22, 81, 74]. ViTs have further been explored beyond the task of image classification [0, 25, 64, 61, 71]. Tolstikhin *et al.* [66] challenge the *status-quo* of convolutions and attention in current computer vision models and proposes MLP-Mixer, a

pure Multi-Layer Perceptron (MLP)-based architecture.

**Adversarial attacks and robustness.** CNNs are known to be vulnerable to adversarial examples [18, 36, 65]. Depending on the accessibility to the target model, adversarial attacks can be divided into white box ones [8, 18, 42, 46] that require full access to the target model, query-based black-box attacks [9, 19, 30, 31, 52, 57, 62, 69], and transfer-based black-box attacks [14, 21, 33, 40, 68, 73, 75]. Adversarial attacks can be divided into image-dependent ones [8, 18, 42, 46, 50] and universal ones [9, 47, 48, 79, 80, 82]. Specifically, contrary to image-dependent attacks, a single perturbation, *i.e.* universal adversarial perturbation (UAP) exists to fool the model for most images [82]. The vulnerability of transformers in the context of NLP has also been investigated [16, 23, 27, 29, 35, 39, 63]. This work focuses on the empirical robustness evaluation of CNN, ViT, and Mixer in image classification.

**Concurrent works.** Recently, also other works [2, 9, 20, 44, 45, 49, 50, 54, 51, 72] investigated transformers from the perspective of adversarial robustness. Specifically, [2, 9, 44, 49, 54, 51] concurrently compare the robustness of transformers to CNNs and independently derive conclusions resembling each other. Their main conclusions, ignoring their nuance difference, can be summarized as *vision transformers are more robust than CNNs*. Not surprisingly, our work also comes to the same main take-way message but differs in multiple aspects, such as co-analysis of MLP-Mixer, perturbation-minimization (C&W and Deep-Fool) results in white-box setting, evaluation under universal attack. Our works also leads to some additional insight, such that MLP-Mixer shows an increased vulnerability to universal attacks. In another parallel line, [45, 50] have investigated how to improve the adversarial robustness of vision transformers.

### 3 Research Goal and Experimental Setup

**Research goal and scope.** CNNs have achieved dominant success in numerous vision applications in the last few years, however, they are also vulnerable to adversarial attacks. Such vulnerability causes serious concern in security-sensitive applications, such as autonomous driving. This concern has motivated an extensive study on model robustness against various attack methods. With the recent popularity of ViT and Mixer as alternatives to CNNs, it is vital for the community to understand their adversarial robustness and to benchmark them against the widely used CNNs. To this end, this work *empirically* investigates the adversarial robustness of the three architectures. In other words, this work has no intention to understand the reason behind why a certain architecture is more or less robust. Note that there is still no consensus [1] on the explanation of CNN being sensitive to adversarial examples despite a large body of works in this field. As an early attempt to investigate the adversarial robustness of ViT and Mixer, our work focuses on the empirical evaluation and it is out of the scope of this work to theoretically understand why they might be vulnerable. Nonetheless, our work attempts to provide better understandings of the robustness gap between models from a shift-invariance perspective and a frequency perspective. Admittedly, our attempt for explanation is limited and future work is needed for better understanding.

**Models and dataset.** In our experiments, we compare the ViT [15] models, MLP-Mixer [64] and CNN architectures [24]. Note that they all adopt shortcut [24] in their architecture design. For the ViT models, we consider ViT-B/16 and ViT-L/16, where B and L stand for “base” and “large”, respectively, while 16 indicates the patch size. The considered ViT models were pre-trained on ImageNet-21K and fine-tuned on ImageNet-1K [12]. We also evaluate ViT models that are directly trained on ImageNet-1K from [64] (indicated by

“ImageNet-1K”). Corresponding to the ViT models, we also investigated Mixer-B/16 and Mixer-L/16 [66] which are trained on ImageNet-1K. We further consider CNN architectures, ResNet-18 and ResNet-50 [24] trained on ImageNet-1K as well as the semi-weakly supervised (SWSL) variant [46], which is pre-trained on IG-1B-Targeted [43] with associated hashtags from 1,000 ImageNet-1K classes followed by fine-tuning on ImageNet-1K dataset. To evaluate adversarial attacks, if not otherwise mentioned we evaluate different adversarial attacks in the untargeted setting on an ImageNet-compatible dataset. This dataset was originally introduced in the NeurIPS 2017 adversarial challenge<sup>1</sup>. We compare different architectures for the corresponding most widely used models, such as ResNet-18/50 for CNN, and ViT and Mixer for B/16 and-L/16. However, we also note that many other factors other than the architecture itself might also have a role in influencing the robustness. Given the publicly available models, it is nearly impossible to rule out all other factors.

## 4 Experiment Results

### 4.1 Robustness Against White-Box Attacks

We first investigate the robustness under white-box attacks. Particularly, we deploy PGD [42] and FGSM [48]. For both attacks we consider  $\epsilon = \{d/255 \mid d \in \{0.1, 0.3, 0.5, 1, 3\}\}$  for images in range  $[0, 1]$ . For the PGD attack, we set the number of iterations to 20 and keep the other parameters as the default settings of Foolbox [59]. For these two attacks, we report the attack success rate (ASR), *i.e.* the percentage of samples which were classified differently from the ground-truth class. Additionally, we evaluate the models on the  $\ell_2$ -variants of the C&W attack [6] and DeepFool [46]. These two attacks have the objective to minimize the perturbation magnitude given the ASR of 100%. Hence, we report the  $\ell_2$ -norm of the adversarial perturbation and the results are available in Table 1. Overall a trend can be observed that compared with CNN architectures, the ViT and Mixer models have a lower attack success rate, suggesting they are more robust than CNN architectures. The increased robustness of ViT and Mixer models is further supported by a higher  $\ell_2$ -norm for the C&W and DeepFool attacks. However, when the perturbation magnitude is very small, the opposite phenomenon can be observed. For example for Mixer-L/16 for the PGD or FGSM with  $\epsilon = 0.1$  the model can exhibit decreased robustness compared to the CNN models.

Model	Clean		PGD ( $\ell_\infty$ ) ↓					FGSM ( $\ell_\infty$ ) ↓					C&W ( $\ell_2$ ) ↑	DeepFool ( $\ell_2$ ) ↑
	ImageNet	NeurIPS	0.1	0.3	0.5	1	3	0.1	0.3	0.5	1	3		
ViT-B/16	81.4	90.7	<b>22.6</b>	63.6	86.5	97.5	99.9	<b>19.1</b>	38.7	52.8	66.3	79.7	<b>0.468</b>	0.425
ViT-L/16	82.9	89.3	<b>22.8</b>	<b>60.1</b>	80.9	95.8	100	19.5	<b>35.9</b>	<b>44.9</b>	<b>57.9</b>	<b>67.3</b>	0.459	<b>0.548</b>
ViT-B/16 (ImageNet-1K)	76.7	86.6	29.2	68.3	87.1	96.8	99.5	26.7	52.2	67.1	83.4	91.0	0.408	0.308
ViT-L/16 (ImageNet-1K)	72.8	79.2	44.9	77.0	90.6	97.1	<b>99.2</b>	36.8	51.4	59.8	70.3	78.8	0.335	0.261
Mixer-B/16	76.5	86.2	29.5	63.4	82.0	96.2	100	27.7	49.3	59.5	69.3	78.0	0.375	0.339
Mixer-L/16	71.8	80.0	41.1	67.3	<b>80.4</b>	<b>92.1</b>	99.4	36.7	51.8	56.9	61.6	67.4	0.297	0.377
ResNet-18 (SWSL)	73.3	90.4	47.9	93.7	98.7	99.5	99.6	38.0	76.3	89.9	96.2	97.6	0.295	0.132
ResNet-50 (SWSL)	81.2	96.3	39.4	90.2	97.0	98.4	99.4	26.3	60.9	73.0	83.8	87.5	0.380	0.149
ResNet-18	69.8	83.7	46.1	90.0	97.8	99.9	100	42.0	75.2	88.5	95.7	98.2	0.302	0.237
ResNet-50	76.1	93.0	35.8	86.3	97.9	99.5	100	27.5	63.1	77.6	89.4	93.9	0.371	0.287

Table 1: **White-box attacks on benchmark models with different epsilons.** We report the clean accuracy on the ImageNet and NeurIPS dataset. The attack success rate (%) of PGD and FGSM under  $\ell_\infty$  distortion, and the  $\ell_2$ -norm of C&W and DeepFool, respectively are reported for the NeurIPS dataset. All models are trained with an image size of 224. A model with a **lower ASR** ↓ or **higher  $\ell_2$ -norm** ↑ is considered to be more robust.

<sup>1</sup><https://github.com/rwightman/pytorch-nips2017-adversarial>

**Class-wise Robustness.** To provide a more detailed robustness evaluation, we perform a class-wise robustness study. We perform the class-wise robustness study on the ImageNet validation dataset, where for each class 50 validation images are present and attack the models with the  $\ell_\infty$ -PGD attack ( $\epsilon = 0.3$ ). Figure 1 (left) shows every 50th class from the most robust classes to the least robust classes for ViT-L/16, where the class “screen” exhibits no robustness, while the class “yellow lady’s slipper” exhibits robustness of 100%. This indicates an imbalance in the class-wise robustness between different classes. We further calculate the similarity of the class-wise accuracies between different models. We treat the different class-wise accuracies as a vector and calculate the cosine similarity between the class-wise accuracies of different models. From Figure 1 (center), it can be observed that the ViT and Mixer models exhibit relatively high similarity values, but also that ResNet18 and ResNet50 are similar in their class-wise robust accuracies. Finally, we examine the relative class-wise robustness of the models, by calculating how many classes of one model are more robust than those of another model. These results are presented in Figure 1 (right). Here it can be observed that the ViT and Mixer models exhibit higher class-wise accuracies than the CNNs, with consistently more than 945 classes being more robust than the examined CNNs.




Figure 1: **(Left):** Most to least robust classes in increments of 50 for ViT-L/16; **(Center):** Cosine similarity of class-wise robust accuracies between different models; **(Right):** Count of all class-wise accuracies of one model (row) which are more robust than those of another (column). The robustness of all models is evaluated under the  $\ell_\infty$ -PGD attack with  $\epsilon = 0.3$  on the ImageNet validation dataset.

## 4.2 Robustness Against Black-Box Attacks

We evaluate and compare the robustness of different model architectures for the black-box attacks in two setups: query-based black-box attack and transfer-based black-box attack.

	ViT-B	ViT-L	Mixer-B	Mixer-L	RN18 (SWSL)	RN50 (SWSL)	RN18	RN50
Boundary ( $\ell_2$ )	3.980	7.408	1.968	1.951	1.403	1.846	1.468	1.740

Table 2: **Boundary attack on benchmark models.** We use 100 randomly selected images from NeurIPS dataset against decision-based attack, and the  $\ell_2$ -norm of adversarial perturbation is presented.

**Query-based black-box attacks.** Query-based black-box attacks work by evaluating a sequence of perturbed images through the model. We adopt a popular decision-based attack, the Boundary Attack [5], which only requires a model’s final decision (*i.e.*, class label) and aims to minimize the perturbation while remaining adversarial. As with the white-box attack, a trend can be observed in the black-box attack that the ViT and Mixer models are more robust, indicated by the relatively higher  $\ell_2$ -norm of the adversarial perturbation (see Table 2).

Methods	ViT-B/16		ViT-L/16		Mixer-B/16		Mixer-L/16		RN18 (SWSL)		RN50 (SWSL)		RN18		RN50	
	ASR	#Q	ASR	#Q	ASR	#Q	ASR	#Q	ASR	#Q	ASR	#Q	ASR	#Q	ASR	#Q
<b>Bandits<sub>TD</sub></b> ( $\ell_\infty$ )	93%	1032	86%	<b>1588</b>	93%	1079	91%	1308	99%	428	91%	974	99%	311	98%	440

Table 3: **Results of Bandits<sub>TD</sub> (with time and data-dependent priors) attack on benchmark models.** ASR is the attack success rate and #Q denotes the average number of queries of the successful attacks. We test on the total NeurIPS dataset with a query limit of 10,000.

We further test and compare the models using a more recently proposed approach [54] that reduces the query cost. To achieve a higher success rate with fewer queries, Bandits<sub>TD</sub> integrates the prior information about the gradient (*i.e.* time and data) using bandit algorithm to reduce the query number. Following the setting in [54], the maximum  $\ell_\infty$ -norm of perturbations is set to 0.05 and the others are set to default as well. As shown in Table 3, overall we observe that ViT and Mixer require a larger average number of queries with a lower average ASR, suggesting ViT and Mixer are more robust than their CNN counterparts.

**Transfer-based black-box attacks.** Transfer-based black-box attacks exploit the transferable property of adversarial examples, *i.e.*, the adversarial examples generated on a source model transfer to another unseen target model. For the source model, we deploy the I-FGSM [56] attack with 7 steps and evaluate the transferability on the target model. From the result in Table 4, we have two major observations. First, adversarial examples from the same family (or similar structure) exhibit higher transferability, suggesting models from the same family learn similar features. Second, when a different model architecture is used as the source model, there is also a trend that CNNs are relatively more vulnerable (*i.e.*, transfer poorly toward foreign architectures). For example, the transferability from CNN to ViT is often lower than 20%, while the opposite scenario is higher.

Source model	Variant	Target model							
		ViT-B/16	ViT-L/16	Mixer-B/16	Mixer-L/16	ResNet-18 (SWSL)	ResNet-50 (SWSL)	ResNet-18	ResNet-50
ViT-B/16	I-FGSM	100	84.7	48.8	50.5	32.0	20.5	40.9	31.7
ViT-L/16	I-FGSM	90.9	99.9	45.7	48.0	30.4	22.2	40.8	30.9
Mixer-B/16	I-FGSM	33.9	25.3	100	89.1	30.6	20.5	40.8	32.0
Mixer-L/16	I-FGSM	27.7	20.1	80.3	99.7	27.7	17.0	38.2	28.4
ResNet-18 (SWSL)	I-FGSM	16.2	13.6	24.8	29.5	99.6	57.1	73.5	63.4
ResNet-50 (SWSL)	I-FGSM	15.3	13.5	23.6	29.9	56.5	99.5	49.4	51.0
ResNet-18	I-FGSM	18.2	14.7	28.9	35.6	84.6	49.9	100	81.6
ResNet-50	I-FGSM	17.7	13.6	28.4	34.5	73.9	63.9	80.6	100

Table 4: **Transfer-based black-box attacks on benchmark models.** We report the attack success rate (%) and a model with a lower ASR is considered to be more robust. All models are trained with an image size of 224, and attacked with a maximum  $\ell_\infty$  perturbation of  $\varepsilon = 16$ .

### 4.3 Toy Example

ViT and Mixer are more robust to adversarial attacks than conventional CNNs. In other words, CNN tends to be the least robust in most setups. To facilitate the understanding of the mechanisms, we design a toy example of binary classification where each class is represented by a single image with a size of 224. The two images shown in Figure 2 consist of a single black stripe on a grey background, differing in the stripe orientation, namely a vertical and a horizontal stripe.



Figure 2: Images for our binary classification toy example.

We train a Fully Connected network (FC), Convolution Neural Network (CNN), and Vision Transformer (ViT) on the images. Note that we designed the networks to be of relatively similar small capacity ( $< 5M$ ). We evaluate the adversarial robustness of these models with the commonly used  $\ell_2$  attacks C&W [6] and DDN [50]. We report the  $\ell_2$ -norm of the adversarial perturbation in Table 5. We observe that the CNN is also less robust than the FC and the ViT in this toy example setup.

	C&W ( $\ell_2$ )	DDN ( $\ell_2$ )	#params
CNN	12.55	17.36	4.59M
FC	25.06	25.39	4.82M
ViT	27.82	59.99	4.88M

Table 5:  $\ell_2$ -norm of the adversarial perturbation on our toy example.




Figure 3: Adversarial examples and perturbations generated with C&W (top) and PGD (bottom) using different architectures trained on our toy example.

**Explanation from a shift-invariance perspective.** Recently, [47] has shown that the shift-invariance property of CNNs can be one cause for its vulnerability to adversarial attacks. Their conclusion is derived with the quantitative analysis and theoretical proof, while we focus on an intuitive qualitative analysis. The qualitative results of adversarial perturbations generated by the attacks are shown in Figure 3. For the ViT, one phenomenon can be observed that the adversarial perturbation consists of square patches. This is likely due to the division of the input image into patches in the ViT architecture. Without this split process on the image, we observe clear stripes but with different patterns for CNN and FC. While the CNN model generates perturbations with repeated stripes, the FC model generates perturbations with only a single stripe centered in the image. It should be noted that perturbations are generated toward the adversary, *i.e.*, in the direction of the opposite class' stripe. The observation that the CNN model yields stripes all over the image can be attributed to the shift-invariant property of the CNN model. From the perspective of shift-invariance, the CNN model recognizes features, *i.e.* horizontal or vertical stripe in this setup, regardless of the position of the features on the image. Thus, it is somewhat expected that the perturbation has stripes in a different direction all over the image. For the FC model without the shift-invariant property, it only recognizes the stripes in the center; thus, the resulting perturbation mainly has the stripe in the center. Since our toy example only consists of two samples, the ASR has only limited informative value about the model's robustness. However, the qualitative results can still be observed. For the  $\ell_2$ -PGD attack, we choose a sufficiently high  $\epsilon$  of 40, such that both samples are misclassified. The qualitative results of the PGD attack in

Figure 3 (bottom) resemble those of the C&W attack. These qualitative results provide an interesting insight into the possible link between shift-invariant property and CNN vulnerability. Admittedly, this link is vague and future work is needed to establish a more concrete link between them.

## 4.4 Frequency Analysis

We further attempt to explain the lower robustness of CNN from a frequency perspective [27, 80]. We apply a low-pass filter and a high-pass filter to the input images before feeding them to the model. We then evaluate the Top-1 accuracy of images from the NeurIPS dataset by applying low-pass or high-pass filtering, and the results are shown in Figure 4. For the low-pass filtering, a sharper decline of the CNN architectures can be observed than for the ViT and Mixer, indicating that the CNN architectures are more reliant on the high-frequency features compared to other models. For example, ResNet-50 (SWSL) has higher accuracy than ViT-L/16 when filtered with large bandwidths; however, with smaller bandwidths, the accuracy of ViT-L/16 becomes higher than the ResNet. Additionally, the Mixers have a sharper decline than the ViTs at low-pass bandwidths from 60 to 180. For the high-pass filtering, the ViT models show the steepest decline among the models, indicating that the ViT models rely more on low-frequency features. Note that non-robust features tend to have high-frequency properties [52, 27, 80] and explain the decreased model robustness. This demonstrates why ViT models are more robust than CNN architectures from the frequency perspective. Comparing the results from low-pass and high-pass filtering, we observe that Mixers, regardless of their absolute value of accuracy, exhibit a similar trend to CNNs rather than ViTs.




Figure 4: Top-1 accuracy across a range of frequency bandwidths from low/high-pass filtering (The right direction of x-axis is more oriented to the low and high frequency, respectively). **Left:** Low-pass filtering. **Right:** High-pass filtering.

## 5 Additional Investigations

### 5.1 Robustness against common corruptions

Additional to our investigation of the adversarial vulnerability of ViT, Mixer, and CNN models, we examine the robustness of these models to common, *i.e.* naturally occurring, corruptions [26]. ImageNet-C was proposed by [26] benchmarking neural network robustness




against these common corruptions. In essence, ImageNet-C is a perturbed version of the original ImageNet validation dataset that has 1000 classes and 50 images per class. ImageNet-C has 15 test corruptions, and four hold-out corruptions, each consisting of 5 severities.

Following [28], we evaluate on 15 test corruptions and the results are shown in Table 6. First, there is a clear trend that models that have higher accuracy on the original (clean) ImageNet tend to also have higher accuracy on ImageNet-C, which is somewhat expected. Second, with comparable accuracy on the original (clean) ImageNet, ViT, and Mixer architectures tend to have higher robustness against corruptions. For example, ViT-B/16 has similar accuracy as RN50(SWSL), *i.e.* 81.43% vs. 81.18%, but the robustness of ViT-B/16 is notably higher than that of RN50(SWSL), *i.e.* 58.85% vs. 52.03%. A similar phenomenon can be observed for Mixer-B/16 and RN50. More detailed corruption-wise results are shown in Figure 5. We find that under some corruptions, such as zoom blur and snow, the superiority of ViT is more significant than other corruptions, such as Gaussian noise.

	ImageNet	ImageNet-C	mCE
ViT-B/16	81.43	58.85	51.98
ViT-L/16	82.89	64.11	45.46
Mixer-B/16	76.47	47.00	67.35
Mixer-L/16	71.77	40.47	75.84
RN18 (SWSL)	73.29	37.84	78.30
RN50 (SWSL)	81.18	52.03	60.49
RN18	69.76	32.92	84.67
RN50	76.13	39.17	76.70

**Table 6: Evaluation of benchmark models on the ImageNet-C dataset.** Accuracy the higher the better, mCE the lower the better.



**Figure 5: Detailed corruption performance comparison of the different models.**

## 5.2 Robustness against Universal Adversarial Perturbations

We further compare the robustness of the different model architectures against Universal Adversarial Perturbations (UAPs). UAPs have been proposed in [47], however, their algorithm has been identified to be relatively less effective yet slow. In this work, we adopt the state-of-the-art UAP algorithm in [80]. Following [47, 80], we set the image size to 224 and constrain the UAP with an  $\ell_\infty$ -norm of  $\epsilon = 10/255$ . The white-box and the corresponding transferability results are shown in Table 7. Several observations can be made. First, MLP-Mixer models are very vulnerable to UAPs in both white-box and black-box scenarios. Especially in the black-box scenario, the attack success rate is always higher than 95% regardless of the surrogate model architecture. On the contrary, ViT models and CNN models are more robust against UAPs, particularly in the black-box scenarios. Second, comparing ViT and CNN models, there is no obvious robustness gap with an ASR of around 90% for all models. However, in the more challenging black-box scenarios, the ViT models are more robust than their CNN counterparts. This trend aligns well with our previous finding that ViT models are more robust than CNN models. We visualize amplified versions of the resulting UAPs in Figure 6. It is noticeable that for the ViT and Mixer architectures, a tile pattern is observable, which is caused by the operation of dividing the images into patches as tokens. Another interesting observation is that the UAPs of Mixer tend to be less locally smooth

than those generated on ViT and CNN. The ViT-L/16 with the highest robustness against UAP seems to have the most locally smooth patterns. Overall, it is interesting that Mixer is extremely vulnerable to UAPs and the finding from a qualitative result that they have locally non-smooth patterns. However, we have no clear explanation for the observed phenomenon since the Mixer is still a very recent architecture. An important message to the community is that the adversarial threat to Mixer cannot be ignored since it is vulnerable in the *practical* transferable universal attack scenario.

Source model	Target model							
	ViT-B/16	ViT-L/16	Mixer-B/16	Mixer-L/16	RN18 (SWSL)	RN50 (SWSL)	RN18	RN50
ViT-B/16	<b>93.9</b>	32.7	95.7	96.0	43.3	29.2	46.4	37.2
ViT-L/16	42.1	<b>84.8</b>	95.6	95.6	39.0	21.2	41.6	31.7
Mixer-B/16	16.6	14.9	<b>98.6</b>	96.7	35.4	19.6	43.3	30.6
Mixer-L/16	15.2	13.7	96.3	<b>98.9</b>	30.3	15.5	38.7	26.0
RN18 (SWSL)	12.6	12.1	95.4	95.6	<b>91.2</b>	49.2	56.6	56.5
RN50 (SWSL)	11.8	11.9	95.3	95.6	50.4	<b>87.1</b>	46.8	46.6
RN18	13.5	11.6	95.7	95.7	66.2	46.1	<b>93.7</b>	63.3
RN50	12.2	12.0	95.4	95.6	63.3	57.0	63.4	<b>91.9</b>

Table 7: **Universal adversarial attacks on benchmark models.** We report the attack success rate (%) and a model with a lower ASR is considered to be more robust. All models are trained with an image size of 224, and evaluated on the NeurIPS dataset. The **bold** font indicates white-box attacks.




Figure 6: Visualization of the amplified UAPs generated on different model architectures.

## 6 Conclusion

Our work performs an empirical study on the adversarial robustness comparison of ViT and MLP-Mixer to the widely used CNN on image classification. Our results show that ViT is significantly more robust than CNN in a wide range of white-box attacks. A similar trend is also observed in the query-based and transfer-based black-box attacks. Our toy task of classifying two simple images with a vertical or horizontal black stripe in the image center provides interesting insight on the possible link between shift-invariant property and CNN vulnerability and future work is necessary for further investigating this link. Our analysis from the feature perspective further suggests that ViTs are more reliant on low-frequency (robust) features while CNNs are more sensitive to high-frequency features. We also investigate the robustness of the newly proposed MLP-Mixer and find that its robustness generally locates in the middle of ViT and CNN. We have also performed additional investigations on robustness against common corruptions and UAPs. One very intriguing finding is that Mixer is extremely vulnerable to UAPs, even in transfer-based black-box attacks. Future work is needed for a better understanding of the reported empirical results.

## References

- [1] Naveed Akhtar and Ajmal Mian. Threat of adversarial attacks on deep learning in computer vision: A survey. *IEEE Access*, 2018.
- [2] Ahmed Aldahdooh, Wassim Hamidouche, and Olivier Deforges. Reveal of vision transformers robustness against adversarial attacks. *arXiv preprint arXiv:2106.03734*, 2021.
- [3] Philipp Benz, Chaoning Zhang, Tooba Imtiaz, and In So Kweon. Double targeted universal adversarial perturbations. In *ACCV*, 2020.
- [4] Srinadh Bhojanapalli, Ayan Chakrabarti, Daniel Glasner, Daliang Li, Thomas Unterthiner, and Andreas Veit. Understanding robustness of transformers for image classification. *arXiv preprint arXiv:2103.14586*, 2021.
- [5] Wieland Brendel, Jonas Rauber, and Matthias Bethge. Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. *arXiv preprint arXiv:1712.04248*, 2017.
- [6] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *SP*, 2017.
- [7] Jieneng Chen, Yongyi Lu, Qihang Yu, Xiangde Luo, Ehsan Adeli, Yan Wang, Le Lu, Alan L Yuille, and Yuyin Zhou. Transunet: Transformers make strong encoders for medical image segmentation. *arXiv preprint arXiv:2102.04306*, 2021.
- [8] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pretraining from pixels. In *ICML*, 2020.
- [9] Pin-Yu Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui Hsieh. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *ACM workshop on artificial intelligence and security*, 2017.
- [10] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- [11] Xiangxiang Chu, Bo Zhang, Zhi Tian, Xiaolin Wei, and Huaxia Xia. Conditional positional encodings for vision transformers. *arXiv preprint arXiv:2102.10882*, 2021.
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [14] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *CVPR*, 2018.
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.

- [16] Siddhant Garg and Goutham Ramakrishnan. Bae: Bert-based adversarial examples for text classification. *arXiv preprint arXiv:2004.01970*, 2020.
- [17] Songwei Ge, Vasu Singla, Ronen Basri, and David Jacobs. Shift invariance can reduce adversarial robustness. *arXiv preprint arXiv:2103.02695*, 2021.
- [18] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015.
- [19] Chuan Guo, Jacob R Gardner, Yurong You, Andrew Gordon Wilson, and Kilian Q Weinberger. Simple black-box adversarial attacks. *ICML*, 2019.
- [20] Chuan Guo, Alexandre Sablayrolles, Hervé Jégou, and Douwe Kiela. Gradient-based adversarial attacks against text transformers. *arXiv preprint arXiv:2104.13733*, 2021.
- [21] Yiwen Guo, Qizhang Li, and Hao Chen. Backpropagating linearly improves transferability of adversarial examples. *arXiv preprint arXiv:2012.03528*, 2020.
- [22] Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. Transformer in transformer. *arXiv preprint arXiv:2103.00112*, 2021.
- [23] Yaru Hao, Li Dong, Furu Wei, and Ke Xu. Self-attention attribution: Interpreting information interactions inside transformer. *arXiv preprint arXiv:2004.11207*, 2020.
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [25] Shuting He, Hao Luo, Pichao Wang, Fan Wang, Hao Li, and Wei Jiang. Transreid: Transformer-based object re-identification. *arXiv preprint arXiv:2102.04378*, 2021.
- [26] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *ICLR*, 2019.
- [27] Dan Hendrycks, Xiaoyuan Liu, Eric Wallace, Adam Dziedzic, Rishabh Krishnan, and Dawn Song. Pretrained transformers improve out-of-distribution robustness. *arXiv preprint arXiv:2004.06100*, 2020.
- [28] Dan Hendrycks, Norman Mu, Ekin D Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. *ICLR*, 2020.
- [29] Yu-Lun Hsieh, Minhao Cheng, Da-Cheng Juan, Wei Wei, Wen-Lian Hsu, and Cho-Jui Hsieh. On the robustness of self-attentive models. In *Annual Meeting of the Association for Computational Linguistics*, 2019.
- [30] Andrew Ilyas, Logan Engstrom, Anish Athalye, and Jessy Lin. Black-box adversarial attacks with limited queries and information. In *ICML*, 2018.
- [31] Andrew Ilyas, Logan Engstrom, and Aleksander Madry. Prior convictions: Black-box adversarial attacks with bandits and priors. *ICLR*, 2019.
- [32] Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. In *NeurIPS*, 2019.

- [33] Nathan Inkawhich, Kevin J Liang, Binghui Wang, Matthew Inkawhich, Lawrence Carin, and Yiran Chen. Perturbing across the feature hierarchy to improve standard and strict blackbox attack transferability. *NeurIPS*, 2020.
- [34] Yifan Jiang, Shiyu Chang, and Zhangyang Wang. Transgan: Two transformers can make one strong gan. *arXiv preprint arXiv:2102.07074*, 2021.
- [35] Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. Is bert really robust? a strong baseline for natural language attack on text classification and entailment. In *AAAI*, 2020.
- [36] Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. In *ICLR*, 2017.
- [37] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 2015.
- [38] Karel Lenc and Andrea Vedaldi. Understanding image representations by measuring their equivariance and equivalence. In *CVPR*, 2015.
- [39] Linyang Li, Ruotian Ma, Qipeng Guo, Xiangyang Xue, and Xipeng Qiu. Bert-attack: Adversarial attack against bert using bert. *arXiv preprint arXiv:2004.09984*, 2020.
- [40] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into transferable adversarial examples and black-box attacks. *ICLR*, 2017.
- [41] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*, 2021.
- [42] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018.
- [43] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens Van Der Maaten. Exploring the limits of weakly supervised pretraining. In *Proceedings of the European conference on computer vision (ECCV)*, pages 181–196, 2018.
- [44] Kaleel Mahmood, Rigel Mahmood, and Marten Van Dijk. On the robustness of vision transformers to adversarial examples. *arXiv preprint arXiv:2104.02610*, 2021.
- [45] Xiaofeng Mao, Gege Qi, Yuefeng Chen, Xiaodan Li, Shaokai Ye, Yuan He, and Hui Xue. Rethinking the design principles of robust vision transformer. *arXiv preprint arXiv:2105.07926*, 2021.
- [46] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *CVPR*, 2016.
- [47] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. Universal adversarial perturbations. In *CVPR*, 2017.
- [48] Konda Reddy Mopuri, Utkarsh Ojha, Utsav Garg, and R. Venkatesh Babu. Nag: Network for adversary generation. In *CVPR*, 2018.

- [49] Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. *arXiv preprint arXiv:2105.10497*, 2021.
- [50] Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Fahad Shahbaz Khan, and Fatih Porikli. On improving adversarial transferability of vision transformers. *arXiv preprint arXiv:2106.04169*, 2021.
- [51] Daniel Neimark, Omri Bar, Maya Zohar, and Dotan Asselmann. Video transformer network. *arXiv preprint arXiv:2102.00719*, 2021.
- [52] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *ACM on Asia conference on computer and communications security*, 2017.
- [53] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *ICML*, 2018.
- [54] Sayak Paul and Pin-Yu Chen. Vision transformers are robust learners. *arXiv preprint arXiv:2105.07581*, 2021.
- [55] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding with unsupervised learning. 2018.
- [56] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 2019.
- [57] Ali Rahmati, Seyed-Mohsen Moosavi-Dezfooli, Pascal Frossard, and Huaiyu Dai. Geoda: a geometric framework for black-box adversarial attacks. In *CVPR*, 2020.
- [58] Prajit Ramachandran, Niki Parmar, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jonathon Shlens. Stand-alone self-attention in vision models. In *NeurIPS*, 2019.
- [59] Jonas Rauber, Roland Zimmermann, Matthias Bethge, and Wieland Brendel. Foolbox native: Fast adversarial attacks to benchmark the robustness of machine learning models in pytorch, tensorflow, and jax. *Journal of Open Source Software*, 2020.
- [60] Jérôme Rony, Luiz G Hafemann, Luiz S Oliveira, Ismail Ben Ayed, Robert Sabourin, and Eric Granger. Decoupling direction and norm for efficient gradient-based l2 adversarial attacks and defenses. In *CVPR*, 2019.
- [61] Rulin Shao, Zhouxing Shi, Jinfeng Yi, Pin-Yu Chen, and Cho-Jui Hsieh. On the adversarial robustness of visual transformers. *arXiv preprint arXiv:2103.15670*, 2021.
- [62] Yucheng Shi, Siyu Wang, and Yahong Han. Curls & whey: Boosting black-box adversarial attacks. In *CVPR*, 2019.
- [63] Zhouxing Shi, Huan Zhang, Kai-Wei Chang, Minlie Huang, and Cho-Jui Hsieh. Robustness verification for transformers. *arXiv preprint arXiv:2002.06622*, 2020.
- [64] Andreas Steiner, Alexander Kolesnikov, Xiaohua Zhai, Ross Wightman, Jakob Uszkoreit, and Lucas Beyer. How to train your vit? data, augmentation, and regularization in vision transformers. *arXiv preprint arXiv:2106.10270*, 2021.

- [65] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [66] Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. Mlp-mixer: An all-mlp architecture for vision. *arXiv preprint arXiv:2105.01601*, 2021.
- [67] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. *arXiv preprint arXiv:2012.12877*, 2020.
- [68] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. Ensemble adversarial training: Attacks and defenses. *ICLR*, 2018.
- [69] Chun-Chen Tu, Paishun Ting, Pin-Yu Chen, Sijia Liu, Huan Zhang, Jinfeng Yi, Cho-Jui Hsieh, and Shin-Ming Cheng. Autozoom: Autoencoder-based zeroth order optimization method for attacking black-box neural networks. In *AAAI*, 2019.
- [70] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [71] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. *arXiv preprint arXiv:2102.12122*, 2021.
- [72] Zhipeng Wei, Jingjing Chen, Micah Goldblum, Zuxuan Wu, Tom Goldstein, and Yu-Gang Jiang. Towards transferable adversarial attacks on vision transformers. *arXiv preprint arXiv:2109.04176*, 2021.
- [73] Dongxian Wu, Yisen Wang, Shu-Tao Xia, James Bailey, and Xingjun Ma. Skip connections matter: On the transferability of adversarial examples generated with resnets. *arXiv preprint arXiv:2002.05990*, 2020.
- [74] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. Cvt: Introducing convolutions to vision transformers. *arXiv preprint arXiv:2103.15808*, 2021.
- [75] Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, and Alan L Yuille. Improving transferability of adversarial examples with input diversity. In *CVPR*, 2019.
- [76] I Zeki Yalniz, Hervé Jégou, Kan Chen, Manohar Paluri, and Dhruv Mahajan. Billion-scale semi-supervised learning for image classification. *arXiv preprint arXiv:1905.00546*, 2019.
- [77] Dong Yin, Raphael Gontijo Lopes, Jon Shlens, Ekin Dogus Cubuk, and Justin Gilmer. A fourier perspective on model robustness in computer vision. In *NeurIPS*, 2019.

- [78] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Francis EH Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. *arXiv preprint arXiv:2101.11986*, 2021.
- [79] Chaoning Zhang, Philipp Benz, Tooba Imtiaz, and In-So Kweon. Cd-uap: Class discriminative universal adversarial perturbation. In *AAAI*, 2020.
- [80] Chaoning Zhang, Philipp Benz, Tooba Imtiaz, and In-So Kweon. Understanding adversarial examples from the mutual influence of images and perturbations. In *CVPR*, 2020.
- [81] Chaoning Zhang, Philipp Benz, Adil Karjauv, and In So Kweon. Universal adversarial perturbations through the lens of deep steganography: Towards a fourier perspective. *AAAI*, 2021.
- [82] Chaoning Zhang, Philipp Benz, Chenguo Lin, Adil Karjauv, Jing Wu, and In So Kweon. A survey on universal adversarial attack. *IJCAI*, 2021.
- [83] Richard Zhang. Making convolutional networks shift-invariant again. In *International Conference on Machine Learning*, pages 7324–7334. PMLR, 2019.