

Structure Invariant Transformation for better Adversarial Transferability

Xiaosen Wang
Huawei Singularity Security Lab
xiaosen@hust.edu.cn

Zeliang Zhang
HUST
hust0426@gmail.com

Jianping Zhang
Chinese University of Hong Kong
jpzhang@cse.cuhk.edu.hk

Abstract

Given the severe vulnerability of Deep Neural Networks (DNNs) against adversarial examples, there is an urgent need for an effective adversarial attack to identify the deficiencies of DNNs in security-sensitive applications. As one of the prevalent black-box adversarial attacks, the existing transfer-based attacks still cannot achieve comparable performance with the white-box attacks. Among these, input transformation based attacks have shown remarkable effectiveness in boosting transferability. In this work, we find that the existing input transformation based attacks transform the input image globally, resulting in limited diversity of the transformed images. We postulate that the more diverse transformed images result in better transferability. Thus, we investigate how to locally apply various transformations onto the input image to improve such diversity while preserving the structure of image. To this end, we propose a novel input transformation based attack, called Structure Invariant Transformation (SIA), which applies a random image transformation onto each image block to craft a set of diverse images for gradient calculation. Extensive experiments on the standard ImageNet dataset demonstrate that SIA exhibits much better transferability than the existing SOTA input transformation based attacks on CNN-based and transformer-based models, showing its generality and superiority in boosting transferability. Code is available at <https://github.com/xiaosen-wang/SIT>.

1. Introduction

With the unprecedented progress of Deep Neural Networks (DNNs) [25, 20, 23, 50, 12], they have been deployed in many security-sensitive applications, such as face recognition [41, 51, 43, 44], autonomous driving [14, 31], etc. On the other hand, recent works have found that DNNs are vulnerable to adversarial examples [47, 15], which mislead the deep models with imperceptible perturbations. This brings a huge threat to the real-world applications [42, 13, 45, 57, 71, 78] and makes it imperative for an

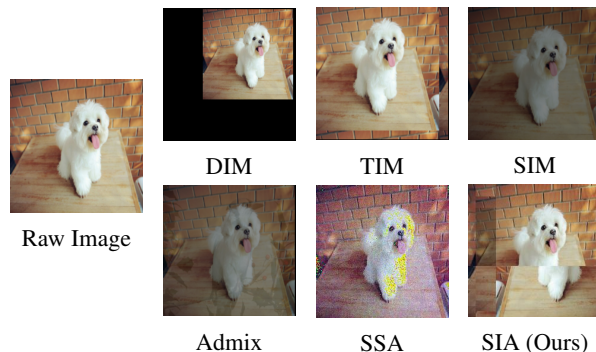


Figure 1: The raw image and its transformed images by DIM (resize factor of 0.8) [64], TIM (translated by 15 pixels) [11], SIM (scale factor of 0.5), Admix (admix strength of 0.2 and scale factor of 0.5) [55], SSA (turning factor of 0.5) [36] and our proposed SIA (3×3 blocks).

effective attack to identify the deficiencies of DNNs when we robustify the deep models or deploy them for commercial applications.

Existing adversarial attacks usually fall into two categories: *white-box attacks* [15, 38, 37, 26, 54] can fetch any information of the target model, including (hyper-)parameters, gradient, architecture, while *black-box attacks* [24, 7, 4, 27, 33, 77] are only allowed limited access to the target model. One of the significant properties of adversarial examples is their transferability [64, 10, 53, 76], in which the adversarial examples generated on one model can still mislead other models, making it possible to attack the real-world applications in the black-box setting. However, existing adversarial attacks [26, 37] often exhibit superior white-box attack performance but poor transferability.

To craft more transferable adversarial examples, various techniques have been proposed, such as momentum-based methods [10, 32, 53, 56], input transformations [64, 11, 32, 68, 72], ensemble attacks [33, 66, 36], advanced objective functions [61, 80] and model-specific approaches [28, 60]. Among which, input transformations (*e.g.*, random resizing and padding [64], translation [11], scale [32] admix [55], etc.) that transform the image before gradient calculation,

have achieved superior transferability and attracted broad attention. Nevertheless, we find that the existing input transformation based attacks, even the SOTA attack *Admix* [55] and SSA [36], transform the image globally without changing the local relationship among objects in the input image. We argue that the diversity of such transformed images is still not enough, leading to limited transferability.

In this work, we postulate and empirically validate that the more diverse transformed images lead to better transferability. Based on this observation, instead of applying a single transformation on the input image, we apply different transformations locally on different parts of the image to enhance the diversity of transformed images. As shown in Fig. 1, such transformation can bring much more difference to the generated image compared with the raw image but still preserve the global structure of the object. Based on this transformation, we propose a novel input transformation based attack called structure invariant attack (SIA), which utilizes the gradient of these transformed images to update the adversarial examples for better transferability.

Our contributions are summarized as follows:

- We empirically validate that the high diversity of transformed images is beneficial to improve transferability, which sheds new light on how to design new input transformations for more transferable adversarial examples.
- We design a new image transformation method, which locally applies different transformations on different parts of the image to generate more diverse images but preserve its global structure.
- Based on the proposed image transformation, we devise a new input transformation based attack, called structure invariant attack (SIA), to generate more transferable adversarial examples.
- Extensive experiments on ImageNet dataset demonstrate that SIA outperforms the baselines with a clear margin on CNN-based as well as transformer-based models, showing its superiority and generality.

2. Related Work

Adversarial examples [47] have brought an impressive threat to the DNN-enabled applications, such as computer vision [15, 58, 3], natural language processing [1, 52, 69], speech recognition [5, 70], *etc.* To identify the vulnerability of DNNs, various attack methods have been proposed recently, such as gradient-based attacks [15, 38, 26, 37], score-based attacks [24, 49, 17, 29, 2], decision-based attacks [4, 6, 27, 57] and transfer-based attacks [33, 10, 64, 32, 55]. Among these, transfer-based attacks do not access any information of the target model, making it applicable to attack any model in the physical world. In this work, we focus on generating more transferable adversarial examples and briefly introduce the existing transfer-based attacks, the corresponding defense methods, and data augmentations.

2.1. Adversarial Attack

As the first gradient-based attack, Fast Gradient Sign Method (FGSM) [15] adds the perturbation in the gradient direction to the benign sample, leading to high attack efficiency but limited performance. Later, I-FGSM [26] extends FGSM into an iterative version, which achieves much better white-box attack performance but poor transferability. Given the high efficiency and effectiveness of I-FGSM, numerous transfer-based attacks are proposed to boost transferability based on I-FGSM to attack the deep model in the black-box setting.

Momentum-based methods. MI-FGSM [10] introduces momentum into I-FGSM to stabilize the optimization direction and escape local maxima. NI-FGSM [32] adopts Nesterov Accelerated Gradient to accumulate the momentum, which achieves better transferability. Variance tuning [53] adopts the gradient variance of the previous iteration to tune the current gradient in MI-FGSM and NI-FGSM, significantly improving transferability. EMI-FGSM [56] enhances the momentum by accumulating several samples' gradients in the previous gradient's direction to further stabilize the optimization direction.

Input transformation based attacks. DIM [64] is the first input transformation based attack, which adds padding to a randomly resized image for fixed size before gradient calculation. TIM [11] optimizes the perturbation over an ensemble of translated images, which is further approximated by convolving the gradient at the untranslated image with a pre-defined kernel. SIM [32] scales the images with different scale factors for gradient calculation. DEM [81] averages the gradient on several diverse transformed images similar to DIM with various resizing factors. *Admix* [55] calculates the gradient on the input image admixed with a small portion of each add-in image from other categories while using the original label of the input. Wu *et al.* [62] train an adversarial transformation network to destroy the adversarial perturbation and require the synthesized adversarial examples resistant to such transformations. SSA [36] adds Gaussian noise and randomly masks the image in the frequency domain to transform the input image.

Ensemble attacks. Liu *et al.* [33] first found that the adversarial examples generated on multiple models denoted as ensemble attack, are more transferable. Recently, Xiong *et al.* [66] reduce the gradient variance between various models to boost the ensemble attack.

Advanced objective functions. The above attacks often take the cross-entropy loss as the objective function. Researchers also find that some regularizers are beneficial to boost transferability. For instance, Zhou *et al.* [80] additionally maximize the difference of the intermediate feature maps between the benign sample and adversarial example. Wu *et al.* [61] adopt a regularizer about the distance of the attention maps between these samples.

Model-specific approaches. Some works utilize the surrogate model’s architecture to improve transferability. Li *et al.* [28] densely add dropout [46] after each layer to create several ghost networks for better transferability. SGM [60] adopts more gradient from the skip connection in ResNets [20] to boost adversarial transferability. LinBP [19] replaces the zeros with ones in the derivative of ReLU to make the model more linear, leading to improved transferability.

2.2. Adversarial Defense

Numerous adversarial defenses have been proposed to mitigate the threat of adversarial attacks. Adversarial training [15, 48], which adopts the adversarial examples during the training process, has been shown as one of the most effective methods [3] but taking huge computation cost. Recently, Wong *et al.* [59] find that single-step adversarial examples can bring satisfiable robustness, making it possible to be applied on large-scale datasets (*e.g.*, ImageNet [25]). Pre-processing the input samples before the model is shown to be another effective way. Liao *et al.* [30] design a high-level representation guided denoiser (HGD) to eliminate the adversarial perturbation. Xie *et al.* [63] find that random resizing and padding on the input image can mitigate the adversarial threat. Naseer *et al.* [39] train a neural representation purifier (NRP) by a self-supervised adversarial training mechanism to purify the input sample, which exhibits superior effectiveness against transfer-based adversarial examples. On the other hand, certified defense methods aim to provide provable defense in a given radius [16, 74, 8]. For instance, randomized smoothing (RS) trains a robust ImageNet classifier with a tight robustness guarantee [8].

2.3. Data Augmentations

Data augmentations often transform (*e.g.*, flipping, rotation, cropping, *etc.*) the image during the training process for better generalization. Mixup [75] interpolates two images and their labels to generate virtual samples for training, which also inspires *Admix* to enhance transferability. Cutmix [73] pastes an image patch to the original patch and mixes the labels accordingly. AutoAugment [9] automatically searches for improved data augmentation policies (operations and parameters) on the dataset for better generalization, which has been widely adopted in deep learning. Unlike these data augmentation strategies, we aim to construct a set of diverse images by transforming the image block using various transformations, which can be used for gradient calculation to achieve better transferability.

3. Methodology

In this section, we introduce our motivation, provide a detailed description of the proposed SIA, and highlight the difference between SIA and AutoAugment.

	TIM	DIM	SIM	SSA	<i>Admix</i>
Transferability	57.4	77.6	79.3	80.6	83.6
LPIPS	0.25	0.43	0.48	0.54	0.73

Table 1: The transferability of TIM, DIM, SIM, *Admix*, SSA, and similarity between 1,000 images and the transformed images evaluated by LPIPS. The transferability is evaluated by the attack success rate of Inception-v3 on the adversarial examples generated on ResNet-18.

3.1. Motivation

Since Xie *et al.* [64] found that transforming the image by random resizing and padding before gradient calculation can generate more transferable adversarial examples, various input transformation based attacks are proposed to further improve transferability. As shown in Fig. 1, the input transformation (*e.g.*, DIM [64] *vs.* *Admix* [55]) with better transferability tends to bring more obvious visual change to the input image. This inspires us with a new assumption:

Assumption 1 *Without harming the semantic information, the more diverse the transformed image is, the better transferability the adversarial examples have.*

To validate this hypothesis, we utilize the learned perceptual image patch similarity (LPIPS) [79] to evaluate the semantic similarity between the benign samples and the transformed images:

$$\text{LPIPS}(\mathbf{x}, \hat{\mathbf{x}}) = \frac{1}{H \times W} \sum_l \sum_{h,w} \|z_{h,w}^l - \hat{z}_{h,w}^l\|_2 \quad (1)$$

where z^l and \hat{z}^l are the extracted feature from l -th layer of SqueezeNet [22] with \mathbf{x} and $\hat{\mathbf{x}}$ as input, respectively. A smaller LPIPS value indicates better similarity between the two images. The semantic similarity between the raw images and transformed images by DIM, TIM, SIM and *Admix* are summarized in Tab. 1. As we can see, the similarity decreases when the transferability of attack increases because the transformed images are more diverse. Moreover, as shown in Fig. 1, such diverse images are significantly different from the raw images, which introduces instability when calculating the gradient. Hence, the more powerful input transformation based attack needs to calculate the gradient on multiple transformed images to eliminate the instability. For instance, DIM utilizes single image, SIM adopts 5 images, while the SOTA *Admix*/SSA takes 15/20 images.

The relation between the diversity of transformed images and transferability inspires us to generate more diverse images for gradient calculation so that we can craft more transferable adversarial examples. In this work, we apply various input transformations on different blocks of a single input image to obtain more diverse images, detailed in Sec. 3.2.




Raw	VShift	HShift	VFlip	HFlip	Rotate	Scale	Add Noise	Resize	DCT	Dropout
										

Table 2: The raw and transformed images using various transformations adopted by SIA (See details in Appendix A).




Figure 2: The randomly sampled raw image and its transformed images by scaling on the full image and 3×3 blocks.

3.2. Structure Invariant Attack

Considering the limited number of the existing image transformations, we first investigate that *how to generate a more diverse image using a single transformation?* Without loss of generality, we take scale as an example, which is the base operation of SIM [32] (also a particular case of *Admix* [55]). As shown in Fig. 2, scaling the image changes its intensity uniformly on all the pixels, resulting in limited diverse images, which is also validated in Tab. 1. To improve the diversity of scaled images, we scale the image with different factors for different parts. Specifically, we randomly split the image into several image blocks and independently scale each image block with different scaled factors. As illustrated in Fig. 2, scaling the image blocks can bring a much more diverse transformed image while humans can still catch the visual information from the image. To explain why such image does not confuse humans, we define the structure of image as follows:

Definition 1 (Structure of Image). *Given an image x , which is randomly split into $s \times s$ blocks, the relative relation between each anchor point is the structure of image, where the anchor point is the center of the image block.*

We argue that the structure of image depicts important semantic information for human recognition. For instance, the dog’s body should be between its head and tail, while its legs should be under its body. Scaling the image blocks with various factors does not change the structure of image so that the generated image can be correctly recognized by humans as well as deep models.

In summary, transforming the image blocks does not harm the recognition but crafts more diverse images, which is of great benefit to improve transferability. To further

Algorithm 1: Structure Invariant Attack

Input: Classifier $f(\cdot)$ with the loss function J ; The benign sample x with ground-truth label y ; The maximum perturbation ϵ , number of iterations T and decay factor μ ; Splitting number s ; Number of transformed images N

Output: An adversarial example.

- 1 $\alpha = \epsilon/T, \mathbf{g}_0 = 0, \mathbf{x}_0^{adv} = \mathbf{x}$
 - 2 **for** $t = 0 \rightarrow T - 1$ **do**
 - 3 Constructing a set \mathcal{X} of N transformed images using SIT
 - 4 Calculating the average gradient on \mathcal{X} :

$$\bar{\mathbf{g}}_{t+1} = \frac{1}{N} \sum_{\mathbf{x}_i \in \mathcal{X}} \nabla_{\mathbf{x}} J(\mathbf{x}_i, y) \quad (2)$$
 - 5 Updating the momentum:

$$\mathbf{g}_{t+1} = \mu \mathbf{g}_t + \frac{\bar{\mathbf{g}}_{t+1}}{\|\bar{\mathbf{g}}_{t+1}\|_1} \quad (3)$$
 - 6 Updating the adversarial example:

$$\mathbf{x}_{t+1}^{adv} = \text{Clip}(\mathbf{x}_t^{adv} + \alpha \cdot \text{sign}(\mathbf{g}_{t+1}), 0, 1) \quad (4)$$
 - 7 **return** \mathbf{x}_T^{adv}
-

boost the diversity of transformed images, we apply various image transformations to different image blocks, where the adopted transformations are summarized in Tab. 2. To avoid information loss, we add the constraints on some transformations. For instance, the rotation can only rotate the image 180° to avoid dropping some pixels. We denote such transformation as structure invariant transformation (SIT).

Since SIT can generate more diverse images while maintaining the critical semantic information, we adopt SIT as the input transformation to conduct the adversarial attack, denoted as structure invariant attack (SIA). Instead of directly calculating the gradient on the input image, SIA calculates the gradient on several images transformed by SIT. Note that SIA is generally applicable to any gradient-based attacks. Due to the limited space, we integrate SIA into MI-FGSM [10] and summarize the algorithm in Algorithm 1.

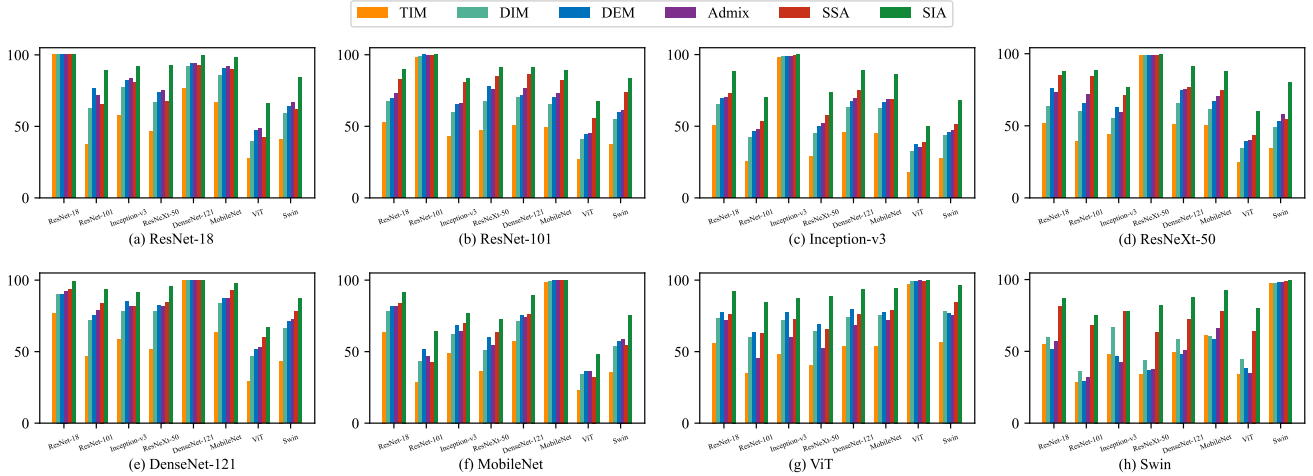


Figure 3: Attack success rates (%) of eight deep models on the adversarial examples crafted on each model by TIM, DIM, DEM, *Admix*, SSA, and SIA.

3.3. SIT vs. AutoAugment

Both SIT and AutoAugment [9] leverage several image transformations to pre-process the images. We summarize the difference as follows:

- SIT improves the adversarial transferability when attacking the model while AutoAugment boosts the model generalization during the training process.
- SIT randomly samples the transformation while AutoAugment needs to search for a good policy for each transformation on each dataset.
- SIT applies various transformations on different image blocks locally but preserves the global structure. On the contrary, AutoAugment applies two transformations sequentially on the image.
- The fine-grained transformations by SIT generate more diverse images than AutoAugment.

4. Experiment

In this section, we conduct extensive evaluations on ImageNet dataset to validate the effectiveness of SIA.

4.1. Experimental Setting

Dataset. To align with previous works [10, 32, 53, 55], we randomly sample 1,000 images pertaining to 1,000 categories from the ILSVRC 2012 validation set [40], which are correctly classified by the adopted models.

Baselines. We compare our proposed SIA with five competitive input transformation based attacks, namely DIM [64], TIM [11], DEM [81], *Admix* [55], and SSA [36], which are integrated into MI-FGSM [10]. We also integrate the baseline methods with two model-specific approaches, namely LinBP [19] and SGM [60].

Victim Models. We evaluate the attack performance on two popular model architectures, namely Convolutional Network Works, *i.e.*, ResNet-18 [20], ResNet-101 [20], ResNeXt-50 [65], DenseNet-121 [23], MobileNet [21], and Transformers, *i.e.*, Vision Transformer (ViT) [12] and Swin Transformer (Swin) [34]. Furthermore, we study several SOTA defense methods, including one adversarial training method, *i.e.*, ensemble adversarially trained model (Inc-v3_{ens}) [48], the top-3 submissions in NIPS 2017 defense competition, *i.e.*, high-level representation guided denoiser (HGD) [30], random resizing and padding (R&P) [63] and NIPS-r3¹, three input pre-processing based defenses, namely FD [35], JPEG [18] and Bit-Red [67], a certified defense, *i.e.*, randomized smoothing (RS) [8] and a deep denoiser, *i.e.*, neural representation purifier (NRP) [39].

Evaluation Settings. We follow MI-FGSM [10] with the perturbation budget $\epsilon = 16$, number of iteration $T = 10$, step size $\alpha = \epsilon/T = 1.6$ and decay factor $\mu = 1$. DIM [64] adopts the transformation probability of 0.5 and TIM [11] utilizes the Gaussian kernel with the size of 7×7 . DEM [81] uses the resize ratios: [1.14, 1.27, 1.4, 1.53, 1.66]. *Admix* admixes 3 images from other categories with the strength of 0.2 and 5 scaled images for each admixed image. SSA [36] sets the turning factor as 0.5 and the standard deviation as ϵ . SIA sets the splitting number $s = 3$ and number of transformed images for gradient calculation $N = 20$.

4.2. Attacking a Single Model

To validate the effectiveness of the proposed SIA, we first compare SIA with five SOTA input transformation based attacks, namely DIM, TIM, DEM, *Admix* and SSA.

¹<https://github.com/anthms/nips-2017/tree/master/mmd>




Figure 4: Attack success rates (%) of eight models on the adversarial examples generated on ResNet-18 when integrating TIM, DIM, DEM, *Admix*, SSA, and SIA into SGM, respectively.

We generate the adversarial examples on a single model and test them on the other models. The attack success rates, *i.e.*, the misclassification rates of the victim model on the crafted adversarial examples, are summarized in Fig. 3.

From the figure, we can observe that all the attackers can achieve the attack success rate of 100.0% or near 100.0%, showing that the input transformation based attacks do not degrade the white-box attack performance. As for the black-box performance, TIM exhibits the poorest transferability on these normally trained models. *Admix* consistently achieves better transferability than DIM and DEM on CNN-based models. Surprisingly, DIM achieves even better transferability than *Admix* when generating adversarial examples on transformer-based models. SSA exhibits the best transferability among the baselines in most cases. Also, even for the transformer-based models, the adversarial examples generated on Swin show the poorest transferability on ViT than other CNN-based models. Hence, we argue that it is necessary to evaluate the effectiveness of transfer-based attacks on both CNN-based and transformer-based models. We need an in-depth analysis of the transferability among the emerging transformer-based models. Compared with the baselines, SIA consistently performs much better than the best baselines on all eight models with different architectures. In particular, SIA outperforms the winner-up method with a clear margin of 14.3% on average and achieves an attack success rate higher than the best baseline of at least 2.8% on all the models. Such consistent and superior performance demonstrates that the proposed SIA is general to various model architectures (either CNN or transformer) to boost transferability effectively.




Figure 5: Attack success rates (%) of eight models on the adversarial examples generated on ResNet-18 when integrating TIM, DIM, DEM, *Admix*, SSA, and SIA into LinBP, respectively.

4.3. Integration to model-specific approaches

To further verify the scalability of the proposed SIA, we integrate existing input transformation-based attacks into two model-specific approaches, namely SGM and LinBP. The adversarial examples are generated on ResNet-18 and test them on the other models.

As depicted in Fig. 4 and Fig. 5, SGM and LinBP can significantly boost the adversarial transferability of val-lina input transformation-based attacks. Compared with the baselines, SIA still achieves much better transferability, which has been boosted by an average margin of 3.3% and 3.6% when integrated into SGM and LinBP, respectively. These superior results validate the generality of SIA to various transfer-based attacks and underscore the potential of SIA in augmenting adversarial transferability through the fusion of different strategies.

4.4. Attacking Ensemble Models

Liu *et al.* [33] have shown that attacking ensemble models can effectively improve transferability. As shown in Sec. 4.2, the adversarial examples often exhibit poorer transferability across the CNN-based models and transformer-based models than being transferred among the CNN-based models. Hence, we generate adversarial examples on ensemble CNN-based and ensemble transformer-based models and test them on the remaining models to evaluate SIA when attacking ensemble models.

As shown in Fig. 6, when attacking ensemble models, the adversarial examples generated by all attacks exhibit better transferability than that crafted on a single model, showing the excellent compatibility of all the input trans-




Figure 6: Attack success rates (%) of eight models on the adversarial examples crafted on ensemble models by TIM, DIM, DEM, *Admix*, SSA, and SIA, respectively. We generate adversarial examples on the CNN-based models and test them on the transformer-based models, and vice versa.

formation based attacks with such a setting. When generating the adversarial examples on ViT and Swin, DIM and DEM achieve better transferability than *Admix* and SSA, highlighting the difference between crafting adversarial examples on different architectures and the necessity to evaluate the transferability on these models. On all eight models, SIA achieves the attack success rate of at least 88.3%. It outperforms the winner-up approach with a clear margin of 6.5%, showing its superior effectiveness in generating transferable adversarial examples. In particular, SIA achieves 94.9% and 99.3% attack success rates on ViT and Swin, respectively, when the adversarial examples are generated on CNN-based models without access to the attention module in the transformer. This further supports our motivation that improving the diversity of transformed images can significantly boost transferability, even on models with completely different architectures.

4.5. Attacking Defense Methods

SIA has achieved superior attack performance on eight normally trained models with different architectures when attacking single model as well as ensemble models. Recently, several defenses have been proposed to mitigate the threat of adversarial examples on ImageNet dataset. To validate the effectiveness of these defenses, we adopt the adversarial examples generated on these eight models simultaneously to attack the defense methods, including Inc-v3_{ens}, HGD, R&P, NIPS-r3, FD, JPEG, Bit-Red, RS and NRP.

The attack results on these defense methods are summarized in Fig. 7. Overall, DIM is on par with TIM on




Figure 7: Attack success rates (%) of various defense methods on the adversarial examples generated by TIM, DIM, DEM, *Admix*, SSA, and SIA under ensemble model setting. The adversarial examples are generated on the eight models simultaneously.

these defense methods while *Admix* can consistently outperform the other four baselines. SIA can consistently achieve better transferability than the baselines. In particular, SIA achieves the attack success rate of 78.2% on the certified defense method (*i.e.*, RS), and SIA can achieve the attack success rate of at least 89.9% on the other eight defense methods, including the powerful denoising method NRP. Such high attack performance indicates the insufficiency of existing defense methods and raises a new security issue to designing more robust deep learning models.

4.6. Ablation Studies

To further gain insights on the superior attack performance achieved by SIA, we conduct a series of ablation studies to validate that all the transformations and splitting the image into blocks can help improve transferability. We generate the adversarial examples on ResNet-18 and test them on the other seven models for all the experiments.

Are all the transformations beneficial for boosting the transferability? There are ten different input transformations that SIA might apply to each image block. To investigate whether each transformation is of benefit to generate more transferable adversarial examples, we first adopt one or two transformations to implement SIA. The results are summarized in Fig. 8. The average attack success rate of MI-FGSM is 59.7%. On the diagonal line of Fig. 8, we only adopt a single transformation for each block and the lowest average attack success rate is 67.3% by adopting DCT, which outperforms MI-FGSM with a margin of 7.6%. This demonstrates the high effectiveness of SIA and

SIA	-VShift	-HShift	-VFlip	-HFlip	-Rotate	-Sclae	-Add Noise	-Resize	-DCT	-Dropout
92.1	89.7	90.1	90.1	88.4	88.3	90.1	90.6	90.2	90.1	90.7

Table 3: The average attack success rates of adversarial examples crafted by SIA and SIA without a single transformation. The adversaries are generated on ResNet-18 and tested on the other seven deep models. - indicates removing such transformation.

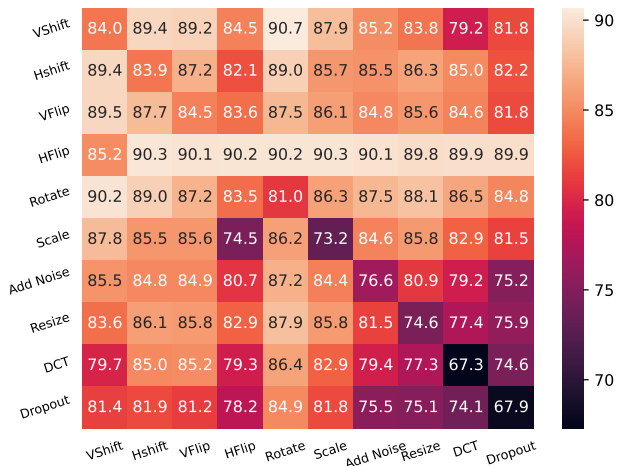


Figure 8: The average attack success rates (%) on the adversarial examples generated by SIA using one (diagonal) or two (other) transformations.

its generality to various transformations. We can also observe that the transferability can be further improved when we combine any two transformations, showing the advantage of combining these transformations. To further validate the necessity of each transformation, we generate adversarial examples by SIA without each transformation and summarize the results in Tab. 3. Since removing a single transformation does not significantly decrease the diversity of the transformation, these attacks achieve similar attack performance. However, no matter which transformation is removed, the transferability will be decreased, supporting that each transformation is of benefit to generate more transferable adversarial examples.

Are the image blocks beneficial for boosting transferability? We have shown that the abundant transformations can effectively improve transferability. Here we further explore whether it is useful to apply these transformations to the image block instead of the raw image. For comparison, we randomly apply the transformation to the raw image $s \times s$ times, denoted as SIA_r . The results are reported in Fig. 9. We can observe that SIA_r exhibits better transferability than MI-FGSM, which also validates our motivation that improving the diversity of transformed images can generate more transferable adversarial examples. Our SIA consistently performs better than SIA_r , indicating that applying




Figure 9: The average attack success rates (%) of adversarial examples generated by SIA and SIA_r (randomly transform the image without the block partition).

the transformation onto each image block is of great benefit to improve transferability.

4.7. Parameter Studies

In this subsection, we conduct parameter studies to explore the impact of two hyper-parameters, namely the number of blocks s and the number of transformed images for gradient calculation N . All the adversarial examples are generated on ResNet-18 and tested on the other seven models in the black-box setting.

On the number of blocks s . The number of blocks determines how diverse the transformed images are, which can influence the attack performance. To find a good value for s , we conduct SIA with s from 1 to 5 and summarize the attack results in Fig. 10. When $s \leq 3$, increasing the value of s leads to better diversity, which can improve the attack performance. However, when we continually increase the value of s , the diversity is increased, but it also introduces more variance to the gradient, decaying the attack performance slightly. Hence, we adopt $s = 3$ to balance the diversity of images and variance of the gradient for better attack performance.

On the number of images for gradient calculation N . SIA calculates the average gradient on N images to eliminate the variance introduced by the transformation. To de-




Figure 10: Attack success rates (%) of adversarial examples generated by SIA with various number of blocks s .

termine a good value for N , we evaluate SIA with N from 1 to 30 and report the attack success rates in Fig. 11. As we can observe, when $N = 1$, SIA introduces massive variance into the gradient and achieves the lowest performance on all models. When we increase the value of N , the variance among the gradients can be eliminated, and SIA achieves better transferability before $N = 20$. When $N > 20$, increasing N can only introduce computation cost without performance improvement. To balance the attack performance and computation cost, we adopt $N = 20$ in our experiments.

5. Conclusion

In this work, we find that the existing input transformation based attack with better transferability often generates more diverse transformed images. Based on this finding, we design a new image transformation, called structure invariant transformation (SIT), which splits the image into several blocks and randomly transforms each image block independently. With such image transformation, we propose a novel input transformation based attack, dubbed structure invariant attack (SIA), which calculates the average gradient on several transformed images by SIT to update the perturbation. Extensive evaluations demonstrate that SIA can achieve remarkably better transferability than the existing SOTA attacks. In our opinion, SIA provides a new direction by applying the transformation onto the image block to effectively boost transferability, which sheds new light on generating more transferable adversarial examples with fine-grained transformations.




Figure 11: Attack success rates (%) of adversarial examples generated by SIA with various number of images N .

References

- [1] Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani B. Srivastava, and Kai-Wei Chang. Generating Natural Language Adversarial Examples. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 2890–2896, 2018. 2
- [2] Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion, and Matthias Hein. Square Attack: A Query-Efficient Black-Box Adversarial Attack via Random Search. In *Proceedings of the European Conference on Computer Vision*, pages 484–501, 2020. 2
- [3] Anish Athalye, Nicholas Carlini, and David A. Wagner. Obfuscated Gradients Give a False Sense of Security: Circumventing Defenses to Adversarial Examples. In *Proceedings of the International Conference on Machine Learning*, pages 274–283, 2018. 2, 3
- [4] Wieland Brendel, Jonas Rauber, and Matthias Bethge. Decision-Based Adversarial Attacks: Reliable Attacks Against Black-Box Machine Learning Models. In *Proceedings of the International Conference on Learning Representations*, 2018. 1, 2
- [5] Nicholas Carlini and David A. Wagner. Audio Adversarial Examples: Targeted Attacks on Speech-to-Text. In *Proceedings of the IEEE Symposium on Security and Privacy (Workshops)*, pages 1–7, 2018. 2
- [6] Minhao Cheng, Thong Le, Pin-Yu Chen, Huan Zhang, Jinfeng Yi, and Cho-Jui Hsieh. Query-Efficient Hard-label Black-box Attack: An Optimization-based Approach. In *Proceedings of the International Conference on Learning Representations*, 2019. 2
- [7] Shuyu Cheng, Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Improving Black-box Adversarial Attacks with a Transfer-based Prior. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 10932–10942, 2019. 1

- [8] Jeremy M. Cohen, Elan Rosenfeld, and J. Zico Kolter. Certified Adversarial Robustness via Randomized Smoothing. In *Proceedings of the International Conference on Machine Learning*, pages 1310–1320, 2019. 3, 5
- [9] Ekin D. Cubuk, Barret Zoph, Dandelion Mané, Vijay Vasudevan, and Quoc V. Le. AutoAugment: Learning Augmentation Strategies From Data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 113–123, 2019. 3, 5
- [10] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting Adversarial Attacks With Momentum. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9185–9193, 2018. 1, 2, 4, 5
- [11] Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Evading Defenses to Transferable Adversarial Examples by Translation-Invariant Attacks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4312–4321, 2019. 1, 2, 5
- [12] 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 *Proceedings of the International Conference on Learning Representations*, 2021. 1, 5
- [13] Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. Robust Physical-World Attacks on Deep Learning Visual Classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1625–1634, 2018. 1
- [14] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we Ready for Autonomous Driving? The KITTI Vision Benchmark Suite. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3354–3361, 2012. 1
- [15] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and Harnessing Adversarial Examples. In *Proceedings of the International Conference on Learning Representations*, 2015. 1, 2, 3
- [16] Sven Gowal, Krishnamurthy Dvijotham, Robert Stanforth, Rudy Bunel, Chongli Qin, Jonathan Uesato, Relja Arandjelovic, Timothy Arthur Mann, and Pushmeet Kohli. Scalable Verified Training for Provably Robust Image Classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4841–4850, 2019. 3
- [17] Chuan Guo, Jacob R. Gardner, Yurong You, Andrew Gordon Wilson, and Kilian Q. Weinberger. Simple Black-box Adversarial Attacks. In *Proceedings of the International Conference on Machine Learning*, pages 2484–2493, 2019. 2
- [18] Chuan Guo, Mayank Rana, Moustapha Cissé, and Laurens van der Maaten. Countering Adversarial Images Using Input Transformations. In *Proceedings of the International Conference on Learning Representations*, 2018. 5
- [19] Yiwen Guo, Qizhang Li, and Hao Chen. Backpropagating linearly improves transferability of adversarial examples. *Advances in neural information processing systems*, 33:85–95, 2020. 3, 5
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 1, 3, 5
- [21] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. *arXiv preprint arXiv:1704.04861*, 2017. 5
- [22] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-Excitation Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7132–7141, 2018. 3
- [23] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely Connected Convolutional Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2261–2269, 2017. 1, 5
- [24] Andrew Ilyas, Logan Engstrom, Anish Athalye, and Jessy Lin. Black-box Adversarial Attacks with Limited Queries and Information. In *Proceedings of the International Conference on Machine Learning*, pages 2142–2151, 2018. 1, 2
- [25] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 1106–1114, 2012. 1, 3
- [26] Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. Adversarial Examples in the Physical World. In *Proceedings of the International Conference on Learning Representations (Workshops)*, 2017. 1, 2
- [27] Huichen Li, Xiaojun Xu, Xiaolu Zhang, Shuang Yang, and Bo Li. QEBA: Query-Efficient Boundary-Based Black-box Attack. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1218–1227, 2020. 1, 2
- [28] Yingwei Li, Song Bai, Yuyin Zhou, Cihang Xie, Zhishuai Zhang, and Alan L. Yuille. Learning Transferable Adversarial Examples via Ghost Networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 11458–11465, 2020. 1, 3
- [29] Yandong Li, Lijun Li, Liqiang Wang, Tong Zhang, and Boqing Gong. NATTACK: Learning the Distributions of Adversarial Examples for an Improved Black-Box Attack on Deep Neural Networks. In *Proceedings of the International Conference on Machine Learning*, pages 3866–3876, 2019. 2
- [30] Fangzhou Liao, Ming Liang, Yinpeng Dong, Tianyu Pang, Xiaolin Hu, and Jun Zhu. Defense Against Adversarial Attacks Using High-Level Representation Guided Denoiser. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1778–1787, 2018. 3, 5
- [31] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and

- Daan Wierstra. Continuous Control with Deep Reinforcement Learning. In *Proceedings of the International Conference on Learning Representations*, 2016. 1
- [32] Jiadong Lin, Chuanbiao Song, Kun He, Liwei Wang, and John E. Hopcroft. Nesterov Accelerated Gradient and Scale Invariance for Adversarial Attacks. In *Proceedings of the International Conference on Learning Representations*, 2020. 1, 2, 4, 5
- [33] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into Transferable Adversarial Examples and Black-box Attacks. In *Proceedings of the International Conference on Learning Representations*, 2017. 1, 2, 6
- [34] 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. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9992–10002, 2021. 5
- [35] Zihao Liu, Qi Liu, Tao Liu, Nuo Xu, Xue Lin, Yanzhi Wang, and Wujie Wen. Feature Distillation: DNN-Oriented JPEG Compression Against Adversarial Examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 860–868, 2019. 5
- [36] Yuyang Long, Qilong Zhang, Boheng Zeng, Lianli Gao, Xi-anlong Liu, Jian Zhang, and Jingkuan Song. Frequency Domain Model Augmentation for Adversarial Attack. In *Proceedings of the European Conference on Computer Vision*, pages 549–566, 2022. 1, 2, 5
- [37] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards Deep Learning Models Resistant to Adversarial Attacks. In *Proceedings of the International Conference on Learning Representations*, 2018. 1, 2
- [38] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. DeepFool: A Simple and Accurate Method to Fool Deep Neural Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2574–2582, 2016. 1, 2
- [39] Muzammal Naseer, Salman H. Khan, Munawar Hayat, Fahad Shahbaz Khan, and Fatih Porikli. A Self-supervised Approach for Adversarial Robustness. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 259–268, 2020. 3, 5
- [40] Russakovsky Olga, Deng Jia, Su Hao, Krause Jonathan, Sathesh Sanjeev, Ma Sean, hUANG Zhiheng, Karpathy Andrej, Khosla Adiya, and Bernstein Michael et al. Imagenet large scale visual recognition challenge. In *International Journal of Computer Vision*, pages 211–252, 2015. 5
- [41] Florian Schroff, Dmitry Kalenichenko, and James Philbin. FaceNet: A Unified Embedding for Face Recognition and Clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 815–823, 2015. 1
- [42] Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K. Reiter. Accessorize to a Crime: Real and Stealthy Attacks on State-of-the-Art Face Recognition. In *Proceedings of the ACM Conference on Computer and Communications Security*, pages 1528–1540, 2016. 1
- [43] Luchuan Song, Zheng Fang, Xiaodan Li, Xiaoyi Dong, Zhenchao Jin, Yuefeng Chen, and Siwei Lyu. Adaptive face forgery detection in cross domain. In *European Conference on Computer Vision*, pages 467–484. Springer, 2022. 1
- [44] Luchuan Song, Xiaodan Li, Zheng Fang, Zhenchao Jin, Yuefeng Chen, and Chenliang Xu. Face forgery detection via symmetric transformer. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4102–4111, 2022. 1
- [45] Luchuan Song, Bin Liu, Guojun Yin, Xiaoyi Dong, Yufei Zhang, and Jia-Xuan Bai. Tacr-net: editing on deep video and voice portraits. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 478–486, 2021. 1
- [46] Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a Simple Way to prevent Neural Networks from Overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014. 3
- [47] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing Properties of Neural Networks. In *Proceedings of the International Conference on Learning Representations*, 2014. 1, 2
- [48] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian J. Goodfellow, Dan Boneh, and Patrick D. McDaniel. Ensemble Adversarial Training: Attacks and Defenses. In *Proceedings of the International Conference on Learning Representations*, 2018. 3, 5
- [49] Jonathan Uesato, Brendan O’Donoghue, Pushmeet Kohli, and Aäron van den Oord. Adversarial Risk and the Dangers of Evaluating Against Weak Attacks. In *Proceedings of the International Conference on Machine Learning*, pages 5032–5041, 2018. 2
- [50] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 5998–6008, 2017. 1
- [51] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. CosFace: Large Margin Cosine Loss for Deep Face Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5265–5274, 2018. 1
- [52] Xiaosen Wang, Jin Hao, Yichen Yang, and Kun He. Natural Language Adversarial Defense through Synonym Encoding. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 823–833, 2021. 2
- [53] Xiaosen Wang and Kun He. Enhancing the Transferability of Adversarial Attacks Through Variance Tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1924–1933, 2021. 1, 2, 5
- [54] Xiaosen Wang, Kun He, Chuanbiao Song, Liwei Wang, and John E Hopcroft. AT-GAN: An Adversarial Generator Model for Non-constrained Adversarial Examples. *arXiv preprint arXiv:1904.07793*, 2019. 1
- [55] Xiaosen Wang, Xuanran He, Jingdong Wang, and Kun He. Admix: Enhancing the Transferability of Adversarial At-

- tacks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16138–16147, 2021. 1, 2, 3, 4, 5
- [56] Xiaosen Wang, Jiadong Lin, Han Hu, Jingdong Wang, and Kun He. Boosting Adversarial Transferability through Enhanced Momentum. In *Proceedings of the British Machine Vision Conference*, page 272, 2021. 1, 2
- [57] Xiaosen Wang, Zeliang Zhang, Kangheng Tong, Dihong Gong, Kun He, Zhifeng Li, and Wei Liu. Triangle Attack: A Query-Efficient Decision-Based Adversarial Attack. In *Proceedings of the European Conference on Computer Vision*, pages 156–174, 2022. 1, 2
- [58] Xingxing Wei, Siyuan Liang, Ning Chen, and Xiaochun Cao. Transferable Adversarial Attacks for Image and Video Object Detection. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 954–960, 2019. 2
- [59] Eric Wong, Leslie Rice, and J. Zico Kolter. Fast is Better than Free: Revisiting Adversarial Training. In *Proceedings of the International Conference on Learning Representations*, 2020. 3
- [60] Dongxian Wu, Yisen Wang, Shu-Tao Xia, James Bailey, and Xingjun Ma. Skip Connections Matter: On the Transferability of Adversarial Examples Generated with ResNets. In *Proceedings of the International Conference on Learning Representations*, 2020. 1, 3, 5
- [61] Weibin Wu, Yuxin Su, Xixian Chen, Shenglin Zhao, Irwin King, Michael R. Lyu, and Yu-Wing Tai. Boosting the Transferability of Adversarial Samples via Attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1158–1167, 2020. 1, 2
- [62] Weibin Wu, Yuxin Su, Michael R Lyu, and Irwin King. Improving the Transferability of Adversarial Samples with Adversarial Transformations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9024–9033, 2021. 2
- [63] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, and Alan L. Yuille. Mitigating Adversarial Effects Through Randomization. In *Proceedings of the International Conference on Learning Representations*, 2018. 3, 5
- [64] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2730–2739, 2019. 1, 2, 3, 5
- [65] Saining Xie, Ross B. Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated Residual Transformations for Deep Neural Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5987–5995, 2017. 5
- [66] Yifeng Xiong, Jiadong Lin, Min Zhang, John E. Hopcroft, and Kun He. Stochastic Variance Reduced Ensemble Adversarial Attack for Boosting the Adversarial Transferability. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14963–14972, 2022. 1, 2
- [67] Weilin Xu, David Evans, and Yanjun Qi. Feature Squeezing: Detecting Adversarial Examples in Deep Neural Networks. In *Proceedings of the Network and Distributed System Security Symposium*, 2018. 5
- [68] Bo Yang, Hengwei Zhang, Zheming Li, Yuchen Zhang, Kaiyong Xu, and Jindong Wang. Adversarial example generation with adabelief optimizer and crop invariance. *Applied Intelligence*, 53(2):2332–2347, 2023. 1
- [69] Yichen Yang, Xiaosen Wang, and Kun He. Robust Textual Embedding against Word-level Adversarial Attacks. pages 2214–2224, 2022. 2
- [70] Zhuolin Yang, Bo Li, Pin-Yu Chen, and Dawn Song. Characterizing Audio Adversarial Examples Using Temporal Dependency. In *Proceedings of the International Conference on Learning Representations*, 2019. 2
- [71] Shengming Yuan, Qilong Zhang, Lianli Gao, Yaya Cheng, and Jingkuan Song. Natural Color Fool: Towards Boosting Black-box Unrestricted Attacks. In *Proceedings of the Advances in Neural Information Processing Systems*, 2022. 1
- [72] Zheng Yuan, Jie Zhang, and Shiguang Shan. Adaptive image transformations for transfer-based adversarial attack. In *European Conference on Computer Vision*, pages 1–17. Springer, 2022. 1
- [73] Sangdoon Yun, Dongyoon Han, Sanghyuk Chun, Seong Joon Oh, Youngjoon Yoo, and Junsuk Choe. CutMix: Regularization Strategy to Train Strong Classifiers With Localizable Features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6022–6031, 2019. 3
- [74] Huan Zhang, Hongge Chen, Chaowei Xiao, Sven Gowal, Robert Stanforth, Bo Li, Duane S. Boning, and Cho-Jui Hsieh. Towards Stable and Efficient Training of Verifiably Robust Neural Networks. In *Proceedings of the International Conference on Learning Representations*, 2020. 3
- [75] Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz. Mixup: Beyond Empirical Risk Minimization. In *Proceedings of the International Conference on Learning Representations*, 2018. 3
- [76] Jianping Zhang, Jen-tse Huang, Wenxuan Wang, Yichen Li, Weibin Wu, Xiaosen Wang, Yuxin Su, and Michael R. Lyu. Improving the Transferability of Adversarial Samples by Path-Augmented Method. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8173–8182, 2023. 1
- [77] Qilong Zhang, Xiaodan Li, Yuefeng Chen, Jingkuan Song, Lianli Gao, Yuan He, and Hui Xue. Beyond ImageNet Attack: Towards Crafting Adversarial Examples for Black-box Domains. In *Proceedings of the International Conference on Learning Representations*, 2022. 1
- [78] Qilong Zhang, Chaoning Zhang, Chaoqun Li, Jingkuan Song, Lianli Gao, and Heng Tao Shen. Practical No-box Adversarial Attacks with Training-free Hybrid Image Transformation. *arXiv preprint arXiv:2203.04607*, 2022. 1
- [79] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. 3
- [80] Wen Zhou, Xin Hou, Yongjun Chen, Mengyun Tang, Xi-anqi Huang, Xiang Gan, and Yong Yang. Transferable Ad-

versarial Perturbations. In *Proceedings of the European Conference on Computer Vision*, pages 471–486, 2018. [1](#), [2](#)

- [81] Junhua Zou, Zhisong Pan, Junyang Qiu, Xin Liu, Ting Rui, and Wei Li. Improving the Transferability of Adversarial Examples with Resized-Diverse-Inputs, Diversity-Ensemble and Region Fitting. In *Proceedings of the European Conference on Computer Vision*, pages 563–579, 2020. [2](#), [5](#)