

# Improving the Transferability of Adversarial Samples by Path-Augmented Method

Anonymous CVPR submission

Paper ID 10640

## Abstract

Deep neural networks have achieved unprecedented success on diverse vision tasks. However, they are vulnerable to adversarial noise that is imperceptible to humans. This phenomenon negatively affects their deployment in real-world scenarios, especially security-related ones. To evaluate the robustness of a target model in practice, transfer-based attacks craft adversarial samples with a local model and have attracted increasing attention from researchers due to their high efficiency. The state-of-the-art transfer-based attacks are generally based on data augmentation, which typically augments multiple training images from a linear path when learning adversarial samples. However, such methods selected the image augmentation path heuristically and may augment images that are semantics-inconsistent with the target images, which harms the transferability of the generated adversarial samples. To overcome the pitfall, we propose the Path-Augmented Method (PAM). Specifically, PAM first constructs a candidate augmentation path pool. It then settles the employed augmentation paths during adversarial sample generation with greedy search. Furthermore, to avoid augmenting semantics-inconsistent images, we train a Semantics Predictor (SP) to constrain the length of the augmentation path. Extensive experiments confirm that PAM can achieve an improvement of over 3.7% on average compared with the state-of-the-art baselines in terms of the attack success rates.

## 1. Introduction

Deep neural networks (DNNs) appear to be the state-of-the-art solutions for a wide variety of vision tasks [18, 23]. However, DNNs are vulnerable to adversarial samples [9], which are elaborately designed by adding human-imperceptible noise to the clean image to mislead DNNs into wrong predictions. The existence of adversarial samples causes negative effects on security-sensitive DNN-based applications, such as self-driving and face recogni-

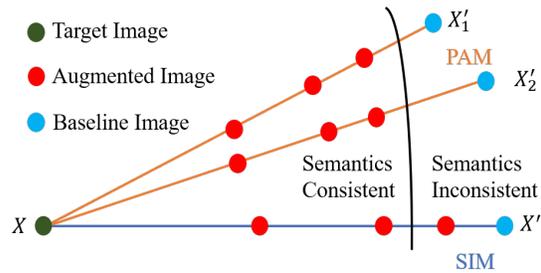


Figure 1. Illustration of how SIM and our PAM augment images (red dots) during the generation of adversarial samples. SIM only considers one linear path from the target image  $X$  to a baseline image  $X'$ . Besides, SIM may augment images that are semantics-inconsistent with the target image. In contrast, our PAM augments images along multiple augmentation paths. We also constrain the length of the path to avoid augmenting images that are semantics-inconsistent with the target one.

tion [22]. Therefore, it is necessary to enhance attack algorithms to better identify the DNN model's vulnerability, which is the first step to improve their robustness against adversarial samples [9].

There are generally two kinds of attacks in the literature [7]. One is the white-box attacks, which consider the white-box setting where attackers can access the architectures and parameters of the victim models. The other is the black-box attacks, which focus on the black-box situation where attackers fail to get access to the specifics of the victim models [8, 30]. Black-box attacks are more applicable than the white-box counterparts to real-world systems. There are two basic black-box attack methodologies: the query-based [1, 2] and the transfer-based attacks [29, 31]. Query-based attacks interact with the victim model to generate adversarial samples, but they may incur excessive queries. In contrast, transfer-based attacks craft adversarial samples with a local source model and do not need to query the victim model. Therefore, transfer-based attacks have attracted more attention recently because of their high efficiency [8, 30].

108 However, transfer-based attacks generally craft adversarial  
109 samples by employing white-box strategies like the Fast  
110 Gradient Sign Method (FGSM) [9] to attack a local model,  
111 which often leads to limited transferability due to overfit-  
112 ting to the employed local model. Most existing solutions  
113 address the overfitting issue from the perspective of opti-  
114 mization and generalization, which regards the local model  
115 and the target image as the training data of the adversarial  
116 sample. Therefore, the transferability of the learned  
117 adversarial sample corresponds to its generalization abil-  
118 ity across attacking different models [17]. Such method-  
119 ologies to improve adversarial transferability can be cate-  
120 gorized into two groups. One is the optimizer-based ap-  
121 proaches [7, 17, 27], which adopt more advanced optimizers  
122 to escape from poor local optima during the generation of  
123 adversarial samples. The other is the augmentation-based  
124 methods [8, 17, 28, 33], which resort to data augmentation  
125 and exploit multiple training images to learn a more trans-  
126 ferable adversarial sample.

127 Current state-of-the-art augmentation-based attacks gen-  
128 erally apply a heuristics-based augmentation method. For  
129 example, the Scale-Invariant attack Method (SIM) [17] aug-  
130 ments multiple scale copies of the target image, while Ad-  
131 mix [28] augments multiple scale copies of the mixtures of  
132 the target image and the images from other categories. SIM  
133 exponentially augments images along a linear path from the  
134 target image to a baseline image, which is the origin. Ad-  
135 mix, in contrast, first augments the target image with the  
136 mixture of the target image and the images from other cate-  
137 gories. Then it also exponentially augments images along a  
138 linear path from the mixture image to the origin. Therefore,  
139 such methods only consider the image augmentation path to  
140 one baseline image, i.e., the origin. Besides, although they  
141 attempt to augment images that are semantics-consistent to  
142 the target image [17, 28], they fail to constrain the length of  
143 the image augmentation path, which may result in augment-  
144 ing semantics-inconsistent images.

145 To overcome the pitfalls of existing augmentation-based  
146 attacks, we propose a transfer-based attack called Path-  
147 Augmented Method (PAM). PAM proposes to augment im-  
148 ages from multiple image augmentation paths to improve  
149 the transferability of the learned adversarial sample. How-  
150 ever, due to the continuous space of images, the possible  
151 image augmentation paths starting from the target image  
152 are countless. In order to cope with the efficiency prob-  
153 lem, we first select representative path directions to con-  
154 struct a candidate augmentation path pool. Then we settle  
155 the employed augmentation paths during adversarial sam-  
156 ple generation with greedy search. Furthermore, to avoid  
157 augmenting semantics-inconsistent images, we train a Se-  
158 mantics Predictor, which is a lightweight neural network, to  
159 constrain the length of each augmentation path.

160 The difference between our PAM and SIM is illustrated  
161

162 in Figure 1. During the generation of adversarial samples,  
163 PAM augments images along multiple image augmentation  
164 paths from the target image to different baseline images,  
165 while SIM only augments images along a single image aug-  
166 mentation path from the target image to the origin. Be-  
167 sides, PAM constrains the length of the image augmen-  
168 tation path to avoid augmenting images that are far away from  
169 the target image and preserve the semantic meaning of the  
170 target image. In contrast, SIM may augment images that  
171 are semantics-inconsistent with the target image due to the  
172 overlong image augmentation path.

173 To confirm the superiority of our PAM, we conduct ex-  
174 tensive experiments against both undefended and defended  
175 models on the ImageNet dataset. Experimental results show  
176 that our PAM can achieve an improvement of over 3.7%  
177 on average compared with the state-of-the-art baselines in  
178 terms of the attack success rates. Since our method can be  
179 combined with other attack strategies, we also evaluate the  
180 performance of the combination of PAM with other compat-  
181 ible attack methods. Again, experimental results confirm  
182 that our method can significantly outperform the state-of-  
183 the-art baselines by about 7.2% on average.

184 In summary, our contributions in this paper are threefold:  
185

- 186 • We discover that the state-of-the-art augmentation-  
187 based attacks (SIM and Admix) actually augment  
188 training images from a linear path for learning adver-  
189 sarial samples. We argue that they suffer from limited  
190 and overlong augmentation paths.
- 191 • To address their pitfalls, We propose the Path-  
192 Augmented Method (PAM). PAM augments images  
193 from multiple augmentation paths during the genera-  
194 tion of adversarial samples. Besides, to make the aug-  
195 mented images preserve the semantic meaning of the  
196 target image, we train a Semantics Predictor (SP) to  
197 constrain the length of each augmentation path.
- 198 • We conduct extensive experiments to validate the ef-  
199 fectiveness of our methodologies. Experimental re-  
200 sults confirm that our approaches can outperform the  
201 state-of-the-art baselines by a margin of over 3.7% on  
202 average. Besides, when combined with other compati-  
203 ble strategies, our method can significantly surpass the  
204 state-of-the-art baselines by 7.2% on average.

## 2. Related Work

### 2.1. Adversarial Attack Method

205 According to the knowledge of the attacker, there are two  
206 categories of attacks in general: white-box and black-box  
207 attacks [4]. White-box attacks assume the white-box set-  
208 ting, where attackers have full access to the victim model,  
209 including the model structures and parameters. Fast Gradi-  
210 ent Sign Method (FGSM) [9] is the first white-box attack  
211  
212  
213  
214  
215

216 that utilizes the sign of the input gradient to maximize the  
217 classification loss to generate adversarial samples in one  
218 step. Basic Iterative Method (BIM) [15] deploys FGSM  
219 to iteratively perturb images to improve the attack perfor-  
220 mance. Project Gradient Descent (PGD) [20] extends BIM  
221 with random start to generate diverse adversarial samples.  
222 Current white-box attacks can achieve nearly 100% attack  
223 success rates in white-box settings. However, they cannot  
224 handle black-box situations, where the model structures and  
225 parameters are unseen.

226 As a result, black-box attacks have attracted increas-  
227 ing attention from researchers recently, which can work  
228 in the black-box setting. There are generally two cate-  
229 gories of black-box attacks. One is the query-based at-  
230 tacks [2, 3, 10, 21], and the other is the transfer-based at-  
231 tacks [7, 8, 17, 33]. Query-based attacks generally determine  
232 the susceptible direction of the victim model through query-  
233 ing it with deliberately designed inputs [2, 3, 10]. However,  
234 query-based attacks may incur prohibitive query costs, hin-  
235 dering their practical application. Transfer-based attacks  
236 exploit the transferability of adversarial samples, which  
237 means that the adversarial samples generated by a local  
238 source model can also mislead a different target model. Due  
239 to their high efficiency, transfer-based attacks are a research  
240 hot spot. However, adversarial samples crafted by white-  
241 box attacks generally possess limited transferability.

242 There are mainly two methodologies to improve the  
243 transferability of white-box attacks. The first one is the  
244 optimizer-based approach, which aims to escape from poor  
245 local optima by adjusting the employment of vanilla gra-  
246 dients during the generation of adversarial samples. For  
247 example, Momentum Iterative Fast Gradient Sign Method  
248 (MI-FGSM) [7] integrates the momentum term into BIM to  
249 improve its adversarial transferability.

250 The other one is the augmentation-based method, which  
251 can be further categorized into two lines. The first one ac-  
252 tually augments images from a linear path. For example,  
253 Scale Invariant Method (SIM) [17] exponentially augments  
254 images along the linear path from the target image to the  
255 origin. Admix [28] follows a similar image augmentation  
256 path while modifying the starting points as the mixture of  
257 the target image and the images from other classes. The  
258 other line banks on affine transformations to augment im-  
259 ages. For example, the Diverse Input Method (DIM) [33]  
260 applies random resizing and padding, while Translation In-  
261 variant Method (TIM) [8] employs shifting. Since affine  
262 transformations focus on changing the pixel positions of an  
263 image, the augmented images are less diverse than those  
264 from a linear path, leading to inferior transferability [28].

265 Unfortunately, state-of-the-art augmentation-based at-  
266 tacks, like SIM and Admix, only consider the image aug-  
267 mentation path to one baseline image, i.e., the origin. Be-  
268 sides, they fail to constrain the length of the image aug-

270 mentation path, which may be overlong and result in aug-  
271 menting images that are far away from and semantics-  
272 inconsistent with the target image. To overcome the defi-  
273 ciencies of such augmentation-based attacks, we propose  
274 the Path-Augmented Method (PAM). To make the aug-  
275 mented images more diverse, we propose to augment im-  
276 ages from multiple augmentation paths during the gener-  
277 ation of adversarial samples. Besides, to make the aug-  
278 mented images preserve the semantic meaning of the tar-  
279 get image, we train a Semantics Predictor (SP) to constrain  
280 the length of each augmentation path. As a result, our  
281 scheme can achieve superior performance over state-of-the-  
282 art transfer-based attacks.

## 283 2.2. Adversarial Defense 284

285 Many adversarial defense methods have been proposed  
286 to alleviate the threat of adversarial samples, which can be  
287 generally grouped into two categories. The first category  
288 is adversarial training, which keeps the state-of-the-art de-  
289 fense methods [15, 26]. Adversarial training retrains the  
290 model by injecting the adversarial samples into the train-  
291 ing data to improve its robustness [9]. Ensemble adversar-  
292 ial training augments the training data with perturbations  
293 transferred from several other models to defend against  
294 transfer-based attacks [15]. The other category is to pu-  
295 rify the adversarial samples. They rectify adversarial per-  
296 turbations by pre-processing inputs without losing classifi-  
297 cation performance on benign images. The state-of-the-art  
298 defense methods in this category include utilizing a high-  
299 level representation guided denoiser [16], random resiz-  
300 ing and padding [32], a JPEG-based defensive compression  
301 framework [19], a compression module [13], and random-  
302 ized smoothing [6]. In this paper, we exploit these state-of-  
303 the-art defenses to evaluate the effectiveness of our attack  
304 against defended models.

## 305 3. Method 306

307 In this section, we first describe the state-of-the-art  
308 augmentation-based attacks (SIM and Admix). Then we an-  
309alyze the limitation of such approaches. We finally present  
310 our Path-Augmented Method (PAM) to overcome the pit-  
311 falls of such attacks.

### 312 3.1. Augmentation-based Attacks 313

314 We first set up some notations. We denote the benign in-  
315 put image as  $x$  and the corresponding true label as  $y$ . We  
316 represent the output of a DNN classifier by  $f(x)$ .  $J(x, y)$   
317 stands for the classification loss function of the classifier,  
318 which is usually the cross-entropy loss. Given the target  
319 image  $x$ , adversarial attacks aim to find an adversarial sam-  
320 ple  $x^{adv}$ , which can mislead the classifier, i.e.,  $f(x^{adv}) \neq$   
321  $f(x)$ , while it is human-imperceptible, i.e., satisfying the  
322 constraint  $\|x - x^{adv}\|_p < \epsilon$ .  $\|\cdot\|_p$  represents the  $L_p$  norm,  
323

and we focus on the  $L_\infty$  norm here to align with previous papers [7, 17].

Prevailing white-box attacks like FGSM [9] usually craft adversarial samples by solving the following constrained maximization problem:

$$\max_{x^{adv}} J(x^{adv}, y) \quad s.t. \|x - x^{adv}\|_\infty < \epsilon.$$

**Scale Invariant Method (SIM)** first computes the average gradient  $\bar{g}$  of the classification loss with respect to  $m$  scaled copies of the target image. Then it updates the target image with the sign of  $\bar{g}$  by a small step size  $\epsilon' = \frac{\epsilon}{T}$  in each iteration, where  $T$  is the iteration number. The update rule is formulated below:

$$\begin{aligned} \bar{g}_{t+1} &= \frac{1}{m} \sum_{i=0}^{m-1} \nabla_{x_t^{adv}} J\left(\frac{1}{2^i} \cdot x_t^{adv}, y\right), \\ x_{t+1}^{adv} &= x_t^{adv} + \epsilon' \cdot \text{sgn}\{\bar{g}_{t+1}\}. \end{aligned} \quad (1)$$

**Admix** first replaces the target image with  $m_2$  mixtures of the target image and the images from other categories ( $x' \in X'$ ). Then it follows SIM by using  $m_1$  scale copies of the mixed images. Therefore, Admix computes the update gradient as follows:

$$\begin{aligned} \bar{g}_{t+1} &= \\ \frac{1}{m_1 \cdot m_2} \sum_{x' \in X'} \sum_{i=0}^{m_1-1} \nabla_{x_t^{adv}} J\left(\frac{1}{2^i} \cdot (x_t^{adv} + \eta \cdot (x')), y\right), \end{aligned}$$

where  $\eta$  is the strength of  $x'$  in the mixture image.

## 3.2. Analysis

After pre-processing, the pixel value of an image will be normalized. We denote the image with pixel values all equal to 0 as the origin  $\mathbf{0}$  in the normalized space. We note that the origin is a pure color image, since all its pixels have constant RGB values when we transform the origin in the normalized space back to the original color space.

We find that when generating adversarial samples, SIM and Admix actually augment images from a linear path. Specifically, SIM augments multiple scaled copies of the target image:  $\frac{1}{2^i} \cdot x_t^{adv} = \frac{1}{2^i} \cdot x_t^{adv} + (1 - \frac{1}{2^i}) \cdot \mathbf{0}$ , which is a linear combination of the target image and the origin. Therefore, SIM exponentially augments images along a linear path from the target image to the origin. Admix first replaces the target image with the mixture of the target image and the image from other categories ( $x' \in X'$ ):  $x_t^{adv} + \eta \cdot x'$ . Then it follows SIM to augment multiple scaled copies of the mixture image:  $\frac{1}{2^i} \cdot (x_t^{adv} + \eta \cdot x') = \frac{1}{2^i} \cdot (x_t^{adv} + \eta \cdot x') + (1 - \frac{1}{2^i}) \cdot \mathbf{0}$ , which is also a linear combination of the mixed target image and the origin. Therefore,

Admix exponentially augments images along a linear path from the mixed target image to the origin.

From the above analysis, we argue that SIM and Admix suffer from two pitfalls. The first one is the limited augmentation path. SIM and Admix only consider the augmentation path to one baseline image, which is the origin. However, there are other possible augmentation paths that can increase the diversity of the augmented images. Therefore, the limited diversity of the augmented images can incur limited transferability of the resultant adversarial sample. Besides, the augmentation path of SIM and Admix may be overlong. They may augment images that are too far away from the target image. As a result, the augmented images are close to the origin, which contains no information about the target image. Augmenting such images can distract the learning of adversarial samples against the target image, thus harming adversarial transferability.

## 3.3. Path-Augmented Method

To overcome the pitfalls of state-of-the-art augmentation-based attacks, we propose the Path-Augmented Method (PAM). We first describe how we explore more augmentation paths to increase the diversity of augmented images. Then we introduce our method to constrain the length of the augmentation path to make the augmented images preserve the semantic meanings of the target image.

### 3.3.1 Augmentation Path Exploration

In order to diversify the augmented images, we propose to explore more augmentation paths. In fact, the augmentation paths starting from the target image are numerous, considering the continuous image space. In order to deal with the efficiency problem, we first construct a candidate augmentation path pool by selecting representative augmentation paths. Then, we employ the augmentation path candidate in a greedy manner when crafting adversarial samples.

We first demonstrate the construction of the candidate augmentation path pool. To reduce the numerous searching space and align with SIM, we only consider the pure color images as the baseline image for the augmentation path. Moreover, we select distinct baseline images to guarantee the augmented images on the paths are diverse. The close augmented images have similar augmented gradients having similar effects on transferability. Therefore, we divide the whole image space into multiple regions and select one baseline from each region as the representative augmentation path to form a candidate augmentation path pool. In general, we regard the image space is normalized to  $[-1, 1]$  for the RGB channel. We divide each channel by three points  $(-1, 0, 1)$  to largely diversify the path, so we have  $3^3 = 27$  representative augmentation paths for the image

space. Although we can divide each channel more precisely, the number of augmentation paths increase in cubic degree. Therefore, our way of constructing the augmentation path pool is efficient in improving the transferability.

Afterward, we discuss how to utilize the constructed augmentation path pool for generating adversarial samples. Intuitively, we combine more augmentation paths to compute the gradient, the higher transferability we can obtain, but the computation complexity will increase. Thus, we should balance the transferability and the computation complexity. In consequence, the number of augmentation paths  $n$  we select is a hyperparameter to tune. After the determination of the augmentation paths number for computing the gradient, we should also figure out the augmentation paths we select from the candidate augmentation path pool. We first rank the augmentation paths in the candidate path pool by deploying the following adversarial attack and measuring the average transferability on a development dataset to rank each augmentation path. For simplicity, we denote the baseline image from the path pool as  $x'$ . Therefore the  $i$ -th scaled augmented image along the path from the target image  $x$  to the baseline image  $x'$  is represented by  $\frac{1}{2^i} \cdot x + (1 - \frac{1}{2^i}) \cdot x'$ .

$$\bar{g}_{t+1} = \frac{1}{m} \sum_{i=0}^{m-1} \nabla_{x_t^{adv}} J\left(\frac{1}{2^i} \cdot x_t^{adv} + \left(1 - \frac{1}{2^i}\right) \cdot x', y\right)$$

We follow a greedy manner in that we choose the top- $n$  augmentation paths and directly combine the gradient of augmented images from those augmentation paths together for generating adversarial samples.

### 3.3.2 Semantics Preservation

In order to keep the semantics of the augmented images on the augmentation paths consistent with the target image, we can constrain the length of the augmentation path and augment the images in the semantics-consistent part to avoid the overlong path. However, it is hard to directly know the semantics-consistent part of the augmentation path. We can use the prediction of the classifier on the image along the augmentation path to identify the semantics-consistent length. If the augmented image is semantics-consistent, the augmented image should have the same prediction as the target image. Therefore, the semantics-consistent length is actually to find the decision boundary of the target image class along the augmentation path. Thus, we train a Semantics Predictor (SP) to constrain the length of each augmentation path. The SP takes the image as the input and predicts the semantic ratio on each augmentation path. The semantic ratio is represented by a scaling factor  $r \in [0, 1]$  on each augmentation path. Therefore, we can utilize the semantic ratio to constrain the length of the augmentation path. We

augment the gradient in the semantics-consistent length to obtain meaningful gradients. Therefore, the  $i$ -th scaled image along the augmentation path from the target image  $x$  to the baseline image  $x'$  with a semantic scaling factor  $r$  is represented by  $(1 - r(1 - \frac{1}{2^i})) \cdot x + (1 - \frac{1}{2^i})r \cdot x'$ .

The Semantics Predictor (SP) is a lightweight neural network consisting of five layers: two Convolutional layers, two Average Pooling layers, and one Fully Connected layer. The image is fed into one Convolutional layer with a kernel size of  $5 \times 5$  and one Average Pooling layer with a stride of 4, which can largely reduce the dimension. Then the feature map is sent into another Convolutional layer and Average Pooling layer with the same setting. After that, the feature map is fed into a Fully Connected layer with Sigmoid activation, and the output size is set to be the number of augmentation paths. The output of the lightweight neural network is exactly the semantic scaling factor of each augmentation path. The training objective is to minimize the difference between the confidence score of the true label and the highest confidence score from other classes, as shown below. We train the Semantics Predictor with Adam optimizer for ten epochs and set the learning rate to be  $1 \times 10^{-4}$ .

$$x_b = SP(x) \cdot x' + (1 - SP(x)) \cdot x$$

$$loss = \left\| F(x_b, y) - \max_{y' \neq y} F(x_b, y') \right\|_2$$

### 3.3.3 Attacking Equation and Comparison

The attacking equation of PAM is shown below, where  $x'_j$  is the baseline image of  $j$ -th augmentation path in the augmentation path pool, and  $r_j$  is the semantic ratio of  $j$ -th augmentation path from the Semantics Predictor.  $n$  is the number of augmentation paths, and  $m$  is the number of copies. The detailed PAM algorithm is shown in the Appendix.

$$x_t^{i,j} = \left(1 - r_j \left(1 - \frac{1}{2^i}\right)\right) \cdot x_t^{adv} + \left(1 - \frac{1}{2^i}\right) r_j \cdot x'_j$$

$$\bar{g}_{t+1} = \frac{1}{m \cdot n} \sum_{j=0}^{n-1} \sum_{i=0}^{m-1} \nabla_{x_t^{adv}} J(x_t^{i,j}, y)$$

Finally, we regard the current state-of-the-art methods SIM [17], and Admix [28] are special cases of the PAM because both SIM and Admix treat the origin as the baseline and augment the gradient along a linear path. SIM utilizes the target image as the starting point, but Admixs select mixtures of the target image with images from other classes as starting points. Our PAM tries to solve two problems of the previous methods: the limited and overlong augmentation path. We first augment images from multiple augmentation paths to explore other augmentation directions. Besides, we train a lightweight neural network Semantic Pre-

dictor to constrain the length of each augmentation path for providing a semantics-consistent gradient.

## 4. Experiments

In this section, we conduct experiments to validate the effectiveness of our proposed approach. We first specify the setup of the experiments. Then, we present the attacking results of our approach against both state-of-the-art undefended and defended models. Finally, we present the ablation study on the number of augmentation paths and the Semantic Predictor.

### 4.1. Experimental Setup

We focus on attacking image classification models trained on ImageNet [23], which is the most widely recognized benchmark task for transfer-based attacks [5, 14, 31] and is a more challenging dataset compared to MNIST and CIFAR-10. We follow the protocol of the baseline method [17] to set up the experiments, whose details are shown as follows.

**Dataset.** We randomly sample 1000 images of different categories from the ILSVRC 2012 validation set [23]. We ensure that nearly all selected test images can be correctly classified by all of the models deployed in this paper. We also randomly sample another 1000 images as the development set to train Semantic Predictor and rank representative augmentation paths.

**Target Model.** We consider both undefended (normally trained) models and defended models as the target models. For undefended models, we choose four top-performance models with different architectures, containing Inception-v3 (Inc-v3) [25], Inception-v4 (Inc-v4) [24], Inception-Resnet-v2 (IncRes-v2) [24], and Resnet-v2-101 (Res-v2) [11, 12]. For defended models, we consider three adversarially trained models, because adversarial training is the most simple but effective way to defend attacks [20]. The selected defended models include Inception v3 trained with adversarial samples from an ensemble of three models (Inc-v3<sub>ens3</sub>), and four models (Inc-v3<sub>ens4</sub>), and adversarially trained Inception-Resnet-v2 (IncRes-v2<sub>adv</sub>). Furthermore, we include six advanced defense models that are robust against black-box attacks on the ImageNet dataset. These defenses cover high-level representation guided denoiser (HGD) [16], random resizing and padding (R&P) [32], NIPS-r3<sup>1</sup>, feature distillation (FD) [19], compression defense (ComDefend) [13], and randomized smoothing (RS) [6].

**Baseline.** We take an advanced optimizer-based attack: MI-FGSM [7] as our baseline because it exhibits better transferability than white-box attacks [9, 15]. Furthermore, SIM [17] and Admix [28] can be viewed as special cases of

our proposed PAM, so we select them as baselines. In order to show that our approaches achieve state-of-the-art performance, we select Variance Tuning Method [27] (VMI) because Admix and VMI are the current state-of-the-art transfer-based attack methods. In addition, we integrate all the methods with other augmentation-based methods: DIM [33] and TIM [8] for further comparison. We denote the approaches with DT extension as the method combined with DIM and TIM.

**Metric.** We evaluate the performance of attack methods via the attack success rate against the target model. The attack success rate is the percentage of adversarial samples that successfully mislead the target model over the total number of the generated adversarial sample.

**Parameter.** Following [7], we set the maximum perturbation budget  $\epsilon = 16$ , the number of attack iterations  $T = 10$ , and the step length  $\epsilon' = 1.6$ . We set the decay factor  $\mu = 1.0$  for all the methods. We follow the source code of SIM [17] and Admix [28] to change the number of scale copies to 32 and 8 for a fair comparison with the same computation complexity as PAM. For DIM, we set the transformation probability to 0.5. We deploy the  $7 \times 7$  Gaussian kernels for TIM. We take  $n = 8$  and  $m = 4$  for PAM.

### 4.2. Attack Transferability

First, we study the performance of our attack method PAM against both undefended and defended models. We fix a source model and produce adversarial samples with different attack methods. The generated samples are then fed into the target models to compute the attack success rates. Our attack achieves nearly 100% success rates under the white-box scenarios in Table 1. More importantly, on the evaluation of transferability, our technique can drastically outperform VMI over 10% and Admix about 3.7% under the black-box setting on average. In addition, PAM improves the transferability to adversarially trained models, largely showing a high threat to adversarial training. Besides, our attack consistently outperforms other baselines by a significant margin under the black-box setting, which confirms the superiority of our strategies on transferable adversarial sample generation.

Then, we combine all the baselines with augmentation-based methods: DIM and TIM to further enhance the transferability. As shown in Table 2, the attack success rates against black-box models are promoted by a large margin with our approaches. In general, our attacks consistently outperform the state-of-the-art baselines by about 7.2%, which further corroborates the effectiveness of our method.

In addition, we also evaluate the performance of different attacks against advanced defenses. Table 3 shows the results when adopting Inc-v3 as the source model to attack other advanced defense models. Our attacks reduce the accuracy of defended models to 52.7% on average, defeating

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

| Model     | Attack  | Inc-v3       | Inc-v4       | IncRes-v2   | Res-v2       | Inc-v3 <sub>ens3</sub> | Inc-v3 <sub>ens4</sub> | IncRes-v2 <sub>adv</sub> |
|-----------|---------|--------------|--------------|-------------|--------------|------------------------|------------------------|--------------------------|
| Inc-v3    | MI-FGSM | <b>100.0</b> | 44.1         | 43.1        | 35.1         | 13.2                   | 13.2                   | 6.2                      |
|           | SIM     | <b>100.0</b> | 69.9         | 67.7        | 63.2         | 36.7                   | 31.4                   | 17.5                     |
|           | VMI     | <b>100.0</b> | 71.7         | 67.1        | 59.9         | 36.3                   | 31.0                   | 17.8                     |
|           | Admix   | <b>100.0</b> | 80.1         | 79.1        | 70.1         | 36.9                   | 34.8                   | 19.0                     |
|           | PAM     | <b>100.0</b> | <b>82.9</b>  | <b>82.2</b> | <b>77.5</b>  | <b>44.8</b>            | <b>43.9</b>            | <b>22.0</b>              |
| Inc-v4    | MI-FGSM | 55.1         | 99.6         | 46.7        | 41.6         | 16.1                   | 15.0                   | 7.8                      |
|           | SIM     | 81.2         | 99.5         | 73.8        | 68.7         | 47.2                   | 44.6                   | 29.1                     |
|           | VMI     | 77.9         | 99.7         | 71.1        | 61.8         | 38.4                   | 36.5                   | 24.0                     |
|           | Admix   | 87.0         | 99.7         | 82.9        | 78.2         | 50.6                   | 47.5                   | 31.3                     |
|           | PAM     | <b>90.5</b>  | <b>100.0</b> | <b>83.9</b> | <b>79.7</b>  | <b>57.9</b>            | <b>52.9</b>            | <b>34.0</b>              |
| IncRes-v2 | MI-FGSM | 60.1         | 51.2         | 97.9        | 46.7         | 21.0                   | 16.0                   | 10.9                     |
|           | SIM     | 84.4         | 80.7         | 99.0        | 76.0         | 56.1                   | 48.6                   | 41.9                     |
|           | VMI     | 78.6         | 73.4         | 98.2        | 67.6         | 48.4                   | 39.9                   | 33.5                     |
|           | Admix   | 87.7         | 85.3         | 99.1        | 80.4         | 61.4                   | 54.6                   | 47.3                     |
|           | PAM     | <b>90.8</b>  | <b>88.3</b>  | <b>99.6</b> | <b>84.9</b>  | <b>68.6</b>            | <b>62.0</b>            | <b>51.0</b>              |
| Res-v2    | MI-FGSM | 57.2         | 51.4         | 48.7        | 99.2         | 24.2                   | 22.4                   | 12.7                     |
|           | SIM     | 74.2         | 70.4         | 68.9        | 99.8         | 42.9                   | 38.6                   | 25.2                     |
|           | VMI     | 75.0         | 68.8         | 69.4        | 99.3         | 45.6                   | 41.0                   | 29.6                     |
|           | Admix   | 80.3         | 75.6         | 76.1        | 99.8         | 45.5                   | 40.8                   | 27.5                     |
|           | PAM     | <b>81.8</b>  | <b>77.4</b>  | <b>76.9</b> | <b>100.0</b> | <b>53.1</b>            | <b>45.9</b>            | <b>31.0</b>              |

Table 1. The attack success rates (%) against seven models by various transfer-based attacks. The best results are marked in bold.

| Model     | Attack   | Inc-v3      | Inc-v4      | IncRes-v2   | Res-v2      | Inc-v3 <sub>ens3</sub> | Inc-v3 <sub>ens4</sub> | IncRes-v2 <sub>adv</sub> |
|-----------|----------|-------------|-------------|-------------|-------------|------------------------|------------------------|--------------------------|
| Inc-v3    | SIM-DT   | 99.0        | 85.7        | 80.3        | 75.1        | 67.6                   | 63.1                   | 46.0                     |
|           | VMI-DT   | 99.2        | 78.4        | 75.2        | 67.9        | 58.1                   | 57.4                   | 44.5                     |
|           | Admix-DT | <b>99.6</b> | 88.1        | 85.6        | 79.1        | 69.2                   | 66.1                   | 48.9                     |
|           | PAM-DT   | 99.4        | <b>92.5</b> | <b>91.5</b> | <b>89.4</b> | <b>80.1</b>            | <b>77.9</b>            | <b>55.9</b>              |
| Inc-v4    | SIM-DT   | 86.4        | 98.4        | 84.2        | 77.9        | 69.9                   | 67.1                   | 56.1                     |
|           | VMI-DT   | 81.4        | 98.4        | 76.4        | 67.0        | 58.8                   | 56.7                   | 49.8                     |
|           | Admix-DT | 88.8        | 99.4        | 85.8        | 80.2        | 72.4                   | 69.0                   | 57.6                     |
|           | PAM-DT   | <b>93.9</b> | <b>99.7</b> | <b>91.9</b> | <b>88.1</b> | <b>83.1</b>            | <b>78.1</b>            | <b>67.2</b>              |
| IncRes-v2 | SIM-DT   | 88.2        | 85.6        | 97.4        | 82.2        | 77.6                   | 73.2                   | 72.7                     |
|           | VMI-DT   | 78.8        | 77.2        | 94.8        | 71.8        | 63.9                   | 59.9                   | 59.3                     |
|           | Admix-DT | 88.2        | 87.4        | 98.2        | 84.0        | 80.0                   | 75.4                   | 71.8                     |
|           | PAM-DT   | <b>95.3</b> | <b>93.2</b> | <b>99.3</b> | <b>90.8</b> | <b>88.8</b>            | <b>86.4</b>            | <b>82.8</b>              |
| Res-v2    | SIM-DT   | 85.8        | 80.9        | 84.8        | 98.5        | 76.2                   | 70.3                   | 62.0                     |
|           | VMI-DT   | 81.0        | 78.8        | 78.3        | 98.1        | 69.5                   | 65.7                   | 57.2                     |
|           | Admix-DT | 89.0        | 85.5        | 86.2        | <b>99.9</b> | 78.2                   | 73.1                   | 64.5                     |
|           | PAM-DT   | <b>90.0</b> | <b>86.8</b> | <b>88.0</b> | 99.5        | <b>84.4</b>            | <b>79.6</b>            | <b>71.6</b>              |

Table 2. The attack success rates (%) on eight models by various transfer-based attacks combined with augmentation-based strategies. The best results are marked in bold.

all baseline attacks. It validates the effectiveness of our attack against advanced defense models, raising security concerns for developing more robust defenses.

### 4.3. Ablation Study

We conduct ablation studies to examine two designs in our proposed PAM: the number of augmentation paths  $n$  and the Semantics Predictor. Adversarial samples are gener-

ated by attacking the Inc-v3 model without employing augmentation-based methods.

**Number of Augmentation Paths.** We investigate the effect of different augmentation path numbers on attack performance. We employ PAM with top- $n$  augmentation paths for generating adversarial samples based on the Inc-v3 model. The result is shown in Figure 2. With the increase of the number of augmentation paths, transferabil-

| Attack | HGD         | R&P         | NIPS-r3     | FD          | ComDefend   | RS          | Average     |
|--------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| SIM    | 15.1        | 28.1        | 36.6        | 59.5        | 55.1        | 22.3        | 36.1        |
| VMI    | 15.8        | 27.0        | 33.3        | 54.8        | 52.0        | 22.5        | 34.2        |
| Admix  | 32.4        | 30.5        | 41.3        | 64.4        | 60.8        | 23.7        | 42.2        |
| PAM    | <b>41.0</b> | <b>40.3</b> | <b>48.1</b> | <b>66.0</b> | <b>63.8</b> | <b>24.3</b> | <b>47.3</b> |

Table 3. The attack success rates (%) of six advanced defense mechanisms on adversarial samples. The adversarial samples are generated on the Inc-v3 model by various transfer-based attacks. The best results are marked in bold.

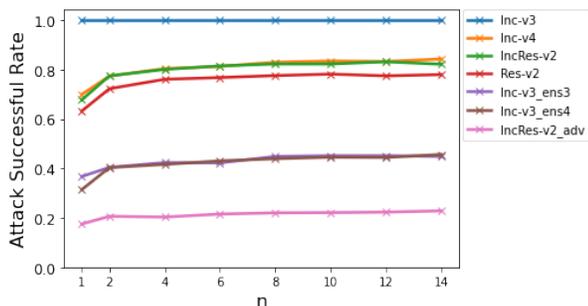


Figure 2. The attack success rates (%) of PAM with different number of augmentation paths  $n$ .

ity improves. However, the computation cost also rises as the number of augmentation paths increases. Therefore, we choose  $n = 8$  to balance the performance and computation cost. Besides, we find an intriguing observation that the selected augmentation path is not the same as SIM or Admix when  $n = 1$ . Our top-1 augmentation path improves the transferability of SIM with more than 1% on average without introducing additional computation complexity. This means the augmentation path of SIM and Admix is not optimal.

**Semantics Predictor.** We study the influence of Semantics Predictor on attack performance for PAM and the performance improvement for SIM. As shown in Table 4, the transferability of SIM can be improved by 1% on average by utilizing the Semantics Predictor because some of the augmented images are semantics-inconsistent with the target image as shown in Figure 3. We cannot recognize the object in the augmented image of SIM. However, the augmented image of SIM+SP demonstrates consistency with the target image, which shows the effectiveness of the Semantics Predictor. In addition, SIM+path+SP outperforms SIM+path by more than 4%, showing Semantics Predictor improves more transferability when we combine multiple augmentation paths together. Besides, SIM + path surpasses SIM by a large margin, which also demonstrates the effectiveness of exploring more augmentation paths.

| Model              | IncRes-v2 | Res-v2 | Inc-v3 <sub>ens4</sub> | IncRes-v2 <sub>adv</sub> |
|--------------------|-----------|--------|------------------------|--------------------------|
| SIM                | 67.7      | 63.2   | 31.4                   | 17.5                     |
| SIM+SP             | 68.3      | 64.3   | 32.5                   | 18.3                     |
| SIM+paths          | 79.3      | 74.2   | 38.5                   | 19.6                     |
| SIM+paths+SP (PAM) | 82.2      | 77.5   | 43.9                   | 22.0                     |

Table 4. The attack success rates (%) against selected four black-box models by various transfer-based attacks.



Figure 3. Visualization of original image and augmented images. We fail to identify the object in the augmented image of SIM. However, the object in the augmented image of SIM+SP is recognizable.

## 5. Conclusion

In this paper, we propose to investigate the problems of current state-of-the-art data augmentation-based attacks and improve their transferability. Specifically, we argue they suffer from the limited and overlong augmentation path. PAM proposes to augment images from multiple image augmentation paths to improve the transferability of the learned adversarial sample. However, due to the continuous space of images, the possible image augmentation paths starting from the target image are countless. In order to cope with the efficiency problem, we first select representative path directions to construct a candidate augmentation path pool. Then we settle the employed augmentation paths during adversarial sample generation with greedy search. Furthermore, to avoid augmenting semantics-inconsistent images, we train a Semantics Predictor, which is a lightweight neural network, to constrain the length of each augmentation path. Extensive experiments confirm the superiority of our approaches on generating transferable adversarial samples against both undefended and defended models over state-of-the-art baselines. In addition, our approaches can generally be combined with other transfer-based attacks to further boost their transferability.

## References

- [1] Maksym Andriushchenko, Francesco Croce, Nicolas Flammarion, and Matthias Hein. Square attack: a query-efficient black-box adversarial attack via random search. In *European Conference on Computer Vision*, pages 484–501. Springer, 2020. 1
- [2] Yang Bai, Yuyuan Zeng, Yong Jiang, Yisen Wang, Shu-Tao Xia, and Weiwei Guo. Improving query efficiency of black-box adversarial attack. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXV 16*, pages 101–116. Springer, 2020. 1, 3
- [3] Arjun Nitin Bhagoji, Warren He, Bo Li, and Dawn Song. Practical black-box attacks on deep neural networks using efficient query mechanisms. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 154–169, 2018. 3
- [4] Battista Biggio and Fabio Roli. Wild patterns: Ten years after the rise of adversarial machine learning. *Pattern Recognition*, 84:317–331, 2018. 2
- [5] Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian Goodfellow, Aleksander Madry, and Alexey Kurakin. On evaluating adversarial robustness. *arXiv preprint arXiv:1902.06705*, 2019. 6
- [6] Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *International Conference on Machine Learning*, pages 1310–1320. PMLR, 2019. 3, 6
- [7] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9185–9193, 2018. 1, 2, 3, 4, 6
- [8] 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, 3, 6
- [9] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 1, 2, 3, 4, 6
- [10] Chuan Guo, Jacob Gardner, Yurong You, Andrew Gordon Wilson, and Kilian Weinberger. Simple black-box adversarial attacks. In *International Conference on Machine Learning*, pages 2484–2493. PMLR, 2019. 3
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. 6
- [13] Xiaojun Jia, Xingxing Wei, Xiaochun Cao, and Hassan Foroosh. Comdefend: An efficient image compression model to defend adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6084–6092, 2019. 3, 6
- [14] Alexey Kurakin, Ian Goodfellow, Samy Bengio, Yinpeng Dong, Fangzhou Liao, Ming Liang, Tianyu Pang, Jun Zhu, Xiaolin Hu, Cihang Xie, et al. Adversarial attacks and defenses competition. In *The NIPS’17 Competition: Building Intelligent Systems*, pages 195–231. Springer, 2018. 6
- [15] Alexey Kurakin, Ian Goodfellow, Samy Bengio, et al. Adversarial examples in the physical world, 2016. 3, 6
- [16] 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 Conference on Computer Vision and Pattern Recognition*, pages 1778–1787, 2018. 3, 6
- [17] Jiadong Lin, Chuanbiao Song, Kun He, Liwei Wang, and John E Hopcroft. Nesterov accelerated gradient and scale invariance for adversarial attacks. *arXiv preprint arXiv:1908.06281*, 2019. 2, 3, 4, 5, 6
- [18] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 1
- [19] 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 *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 860–868. IEEE, 2019. 3, 6
- [20] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017. 3, 6
- [21] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia conference on computer and communications security*, pages 506–519, 2017. 3
- [22] Samira Pouyanfar, Saad Sadiq, Yilin Yan, Haiman Tian, Yudong Tao, Maria Presa Reyes, Mei-Ling Shyu, Shu-Ching Chen, and Sundaraja S Iyengar. A survey on deep learning: Algorithms, techniques, and applications. *ACM Computing Surveys (CSUR)*, 51(5):1–36, 2018. 1
- [23] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 1, 6
- [24] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-first AAAI conference on artificial intelligence*, 2017. 6
- [25] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016. 6
- [26] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. Ensemble

|      |   |      |
|------|---|------|
| 972  | adversarial training: Attacks and defenses. <i>arXiv preprint arXiv:1705.07204</i> , 2017. 3  | 1026 |
| 973  |   | 1027 |
| 974  | [27] Xiaosen Wang and Kun He. Enhancing the transferability of adversarial attacks through variance tuning. In <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> , pages 1924–1933, 2021. 2, 6   | 1028 |
| 975  |   | 1029 |
| 976  |   | 1030 |
| 977  |   | 1031 |
| 978  | [28] Xiaosen Wang, Xuanran He, Jingdong Wang, and Kun He. Admix: Enhancing the transferability of adversarial attacks. In <i>Proceedings of the IEEE/CVF International Conference on Computer Vision</i> , pages 16158–16167, 2021. 2, 3, 5, 6  | 1032 |
| 979  |   | 1033 |
| 980  |   | 1034 |
| 981  |   | 1035 |
| 982  | [29] Zhibo Wang, Hengchang Guo, Zhifei Zhang, Wenxin Liu, Zhan Qin, and Kui Ren. Feature importance-aware transferable adversarial attacks. <i>arXiv preprint arXiv:2107.14185</i> , 2021. 1  | 1036 |
| 983  |   | 1037 |
| 984  |   | 1038 |
| 985  |   | 1039 |
| 986  | [30] 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 <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> , pages 1161–1170, 2020. 1           | 1040 |
| 987  |   | 1041 |
| 988  |   | 1042 |
| 989  |   | 1043 |
| 990  | [31] Weibin Wu, Yuxin Su, Michael R Lyu, and Irwin King. Improving the transferability of adversarial samples with adversarial transformations. In <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> , pages 9024–9033, 2021. 1, 6                             | 1044 |
| 991  |   | 1045 |
| 992  |   | 1046 |
| 993  |   | 1047 |
| 994  |   | 1048 |
| 995  | [32] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, and Alan Yuille. Mitigating adversarial effects through randomization. <i>arXiv preprint arXiv:1711.01991</i> , 2017. 3, 6  | 1049 |
| 996  |   | 1050 |
| 997  |   | 1051 |
| 998  | [33] 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 <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> , pages 2730–2739, 2019. 2, 3, 6 | 1052 |
| 999  |   | 1053 |
| 1000 |   | 1054 |
| 1001 |   | 1055 |
| 1002 |   | 1056 |
| 1003 |   | 1057 |
| 1004 |   | 1058 |
| 1005 |   | 1059 |
| 1006 |   | 1060 |
| 1007 |   | 1061 |
| 1008 |   | 1062 |
| 1009 |   | 1063 |
| 1010 |   | 1064 |
| 1011 |   | 1065 |
| 1012 |   | 1066 |
| 1013 |   | 1067 |
| 1014 |   | 1068 |
| 1015 |   | 1069 |
| 1016 |   | 1070 |
| 1017 |   | 1071 |
| 1018 |   | 1072 |
| 1019 |   | 1073 |
| 1020 |   | 1074 |
| 1021 |   | 1075 |
| 1022 |   | 1076 |
| 1023 |   | 1077 |
| 1024 |   | 1078 |
| 1025 |   | 1079 |