

# Diversifying the High-level Features for better Adversarial Transferability

Zhiyuan Wang\*<sup>1</sup>  
vinlee624@gmail.com  
Zeliang Zhang\*<sup>1</sup>  
hust0426@gmail.com  
Siyuan Liang<sup>2</sup>  
liangsiyuan@iie.ac.cn  
Xiaosen Wang<sup>†3</sup>  
xiaosen@hust.edu.cn

<sup>1</sup> School of Computer Science and Technology, Huazhong University of Science and Technology  
<sup>2</sup> Institute of Information Engineering Chinese Academy of Sciences  
<sup>3</sup> Huawei Singularity Security Lab  
\*Equal contribution  
<sup>†</sup>Corresponding author

---

## Abstract

Given the great threat of adversarial attacks against Deep Neural Networks (DNNs), numerous works have been proposed to boost transferability to attack real-world applications. However, existing attacks often utilize advanced gradient calculation or input transformation but ignore the white-box model. Inspired by the fact that DNNs are over-parameterized for superior performance, we propose diversifying the high-level features (DHF) for more transferable adversarial examples. In particular, DHF perturbs the high-level features by randomly transforming the high-level features and mixing them with the feature of benign samples when calculating the gradient at each iteration. Due to the redundancy of parameters, such transformation does not affect the classification performance but helps identify the invariant features across different models, leading to much better transferability. Empirical evaluations on ImageNet dataset show that DHF could effectively improve the transferability of existing momentum-based attacks. Incorporated into the input transformation-based attacks, DHF generates more transferable adversarial examples and outperforms the baselines with a clear margin when attacking several defense models, showing its generalization to various attacks and high effectiveness for boosting transferability. Code is available at <https://github.com/Trustworthy-AI-Group/DHF>.

## 1 Introduction

Recent studies have shown that Deep Neural Networks (DNNs) are vulnerable to adversarial examples [47], *i.e.*, the inputs with indistinguishable perturbation can mislead the DNNs. Such vulnerability brings a significant threat to the security of widely deployed DNNs in the physical world, *e.g.*, image classification [20, 22, 27, 45, 48], face recognition [42, 50, 62, 70], autopilot [4, 7, 15, 26], *etc.* However, existing adversarial attacks exhibit excellent attack performance when the attacker can access the full knowledge of the target model, but show poor transferability across different models, making it inefficient in the real world.

To improve the transferability of adversarial examples, many methods have been proposed, such as input transformations [66], ensemble-model attacks [34], and model-specific methods [63]. Among them, the model-specific methods, which modify or utilize the inner structure of DNNs, are one of the most effective methods to improve the transferability of adversarial examples and are compatible with the other two types of methods. For instance, Wu *et al.* [63] employ the structure of skip connection in ResNet-like DNNs to generate highly transferable adversarial examples. Li *et al.* [31] propose the ghost network, which virtually ensembles a vast set of diverse models by randomly perturbing the existing model to boost adversarial transferability. However, both of them lack the full utilization of model structure and image features, remaining potential for further improvements.

Many researchers have pointed out that DNNs are over-parameterization [1, 3, 6, 9, 13]. The increasing number of layers in DNNs brings better feature extraction and classification performance, but correspondingly induces a big redundancy of parameters, especially for the deeper layers. As shown in Fig. 1, when we randomly mask the high-level features of DNNs, the classification accuracy maintains at a high level. Thus, it naturally inspires us to exploit the over-parameterization of high-level features for better adversarial transferability. Unlike Wu *et al.* [63] and Li *et al.* [31], we focus on perturbing only the high-level features and propose a unified perturbing operation in any architectures, not restricted to dropout or scale operation and ResNet-like DNNs. We summarize our contributions as follows:

- To the best of our knowledge, it is the first work that establishes a relationship between over-parameterization and adversarial transferability.
- We propose a novel approach called Diversifying the High-level Features (DHF) that linearly transforms the high-level features and mixes up them with that of benign samples. Such transformation helps identify the invariant features across different models.
- Extensive experiments on ImageNet dataset demonstrate that DHF can achieve better adversarial transferability than the existing approaches and is general to other attacks.

## 2 Related Work

**Adversarial Attack:** After Szegedy *et al.* [47] identified the vulnerability of DNNs against adversarial examples, various adversarial attacks have been proposed, *e.g.*, gradient-based attack [17, 28, 38], transfer-based attack [10, 37, 61, 68], score-based attack [5, 24, 30], decision-based attack [2, 29, 58], generation-based attack [54, 65], *etc.* Among these, transfer-based attacks do not access the target model, making it popular to attack the deep models in the real world. To improve adversarial transferability, various momentum-based attacks have been proposed, such as MI-FGSM [10], NI-FGSM [33], VMI-FGSM [53], EMI-FGSM [56], *etc.* Researchers also propose several input transformation methods (*e.g.*, DIM [68], TIM [11], SIM [33], Admix [55], SIA [59], STM [16], BSR [52], *etc.*), which transform the input image before gradient calculation for better transferability. On the contrary, few works focus on the white-box model itself to craft more transferable adversarial examples. TAP [74] maximizes the difference of the feature maps for all layers between the benign sample and adversarial example to enhance the adversarial transferability, while ILA [23] finetunes an adversarial example by enlarging the similarity of the feature difference at a given layer. Ghost network [31] densely adds a dropout layer [46] after each convolutional layer to create a huge set of diverse models for better transferability. Wu *et al.* [63] argues that adopting more gradient from the skip connections could boost the transferability of adversarial examples across various ResNet-based models [20]. In this work, we

try to diversify the high-level feature to craft more transferable adversarial examples, which are general to any DNNs.

**Adversarial Defense:** To mitigate the threat of adversarial attacks, various defenses have been proposed, such as adversarial training [38, 51, 57], input preprocessing [41, 66], feature denoising [32, 67, 71], certified defense [8, 18, 43], *etc.* JPEG [19] eliminates the adversarial perturbation by applying the JPEG compression to the input image. High-level representation guided denoiser (HGD) [32] trains a denoising autoencoder based on U-Net [44] to purify the image. Randomized resizing and padding (R&P) [66] randomly resizes the image and adds padding to mitigate the adversarial effect. Bit depth reduction (Bit-Red) [69] reduces the number of bits for each pixel to squeeze the perturbation. FD [36] employs a JPEG-based compression framework to defend against adversarial attacks. Cohen *et al.* [8] adopt randomized smoothing (RS) to train a certifiably robust ImageNet classifier. Neural Representation Purifier (NRP) [41] adopts a self-supervised adversarial training mechanism to eliminate perturbation effectively.

**Over-parametering for DNNs:** Since Krizhevsky *et al.* [27] achieved superior performance using Convolutional Neural Networks (CNNs) on ImageNet, DNNs are becoming deeper and wider with millions of parameters [20, 22, 40, 48, 64]. Over-parameterization, characterized by deeper and wider neuron layers, has been widely recognized to significantly boost performance [20, 22, 40, 48, 64]. On the other hand, model quantization has been employed to enable the integration of deep neural networks into mobile phones and embedded devices. This technique reduces the memory requirements of over-parameterized models by replacing the floating point weights with low-precision weights [1, 3, 6, 9, 13]. The success of model quantization suggests that DNNs contain redundant information, as the decrease in precision does not significantly impact the final results. In this work, we aim to utilize such redundancy to generate more transferable adversarial examples.

## 3 Methodology

In this section, we first provide our motivation and detail our proposed approach called Diversifying the High-level Features (DHF). Then we analyze why high-level features are better than low-level features when diversifying the features and summarize the difference between DHF and Ghost.

### 3.1 Motivation

Lin *et al.* [33] analogize the generation of adversarial examples with the training process of models. In this case, the transferability of adversarial examples is equivalent to model generalization. Recent studies have demonstrated that input transformation-based methods, akin to data augmentation in model training, can significantly enhance adversarial transferability [11, 55, 68]. On the other hand, Li *et al.* [31] boosts transferability by densely applying dropout and random scaling operations on the models (Ghost networks) to diversify the inner features. It sheds light on the potential to enhance transferability by diversifying the features. Motivated by the Ghost networks, we further study how to improve transferability through feature diversification.

It is widely known that DNNs are over-parameterized for superior performance. Such redundant parameters enable small perturbations among the features to have little impact on the performance, especially in *deeper layers*. To validate this argument, we randomly mask




Figure 1: Classification accuracy of Res-101 when randomly masking  $\rho$  features of the last  $l$  layers among the residual blocks.



Figure 2: Visualization of the benign images and their corresponding adversarial examples generated by DHF.

$\rho$  ( $0 \leq \rho \leq 1$ ) elements of features at the last  $l$  ( $0 \leq l \leq 1$ ) layers (we only focus on the features of convolutional layers in the residual block) on Res-101 [20] and report the accuracy in Fig. 1. We can observe that randomly masking 10% elements of the last 16.7% layers does not significantly decay the accuracy. In particular, the accuracy is still 60.0% when we mask half elements of the last 16.7% layers. When we mask more layers or elements at each layer, the accuracy will decrease but can maintain more than 50% most time, highlighting the possibility of diversifying the features using a single model. Moreover, we argue that despite the models’ diversity in parameters and structures, they share a similar concentration of features that impacts the performance. Thus, attackers can craft adversarial perturbations on diverse features by perturbing the features, which are excellent for transferring across models. This naturally prompts us to diversify the high-level features to identify the invariant features for better transferability. As shown in Fig. 2, DHF generates high-quality adversarial examples with high transferability. More visualization results can be referred to Sec. 2 in the supplementary.

### 3.2 Diversifying the High-level Feature

Given a  $L$ -layer DNN  $f$ ,  $y_l = f^{(l-1)}(x)$  is the output w.r.t. the input  $x$  at the  $l$ -th layer ( $1 \leq l \leq L$ ). To generate highly transferable adversarial examples on  $f$ , we propose a new approach called Diversifying the High-level Feature (DHF) with two operations detailed as follows:

**Mixing up the feature.** To diversify the features of adversarial examples without changing the recognition performance, we mix up the features with the features of benign samples:

$$y_l^* = (1 - \eta) \cdot y_l^{adv} + \eta \cdot y_l, \quad \eta \sim \mathcal{U}(0, \eta_{max}), \quad (1)$$

where  $y_l = f^{(l-1)}(x)$ ,  $y_l^{adv} = f^{(l-1)}(x^{adv})$  and  $\mathcal{U}(0, \eta_{max})$  indicates the uniform distribution from 0 to the hyper-parameter  $\eta_{max}$ . Different from *Admix* [55], which mixes up the image from other categories, we only mix up the features with that of the benign samples. This diversifies the intermediate layers’ features but does not significantly affect the recognition results, making the forward and backward propagation more stable than mixing up some unrelated features. Furthermore, as the input samples become adversarial in the attack

progress, the gradients of such samples might be unstable since the neighborhood of adversarial examples mainly consists of correctly classified samples. Leveraging benign images helps the optimization process to escape such a region, making the gradient more reliable and stabilizing the optimization.

**Randomly adjusting the feature.** As shown in Fig. 1, randomly masking the features does not significantly affect the classification accuracy. This validates the redundancy of parameters, which would result in multiple local minima when optimizing the cross-entropy loss for training. Such redundancy would introduce variance into the intermediate features, which differs in the same models trained twice and is more severe across different models, making it difficult for adversarial examples to transfer across models. To eliminate such variance, we randomly replace  $\rho$  elements of features at each layer with the mean of features, which stabilizes the propagation and helps boost the transferability. Note that replacing the feature with its mean is still differentiable, allowing a more precise gradient w.r.t. the adversarial examples. Compared with adjusting parameters, which will affect the entire feature map, replacing features locally will only change specific elements in the feature map. This distinction makes feature replacement different from parameter adjustment.

Our proposed DHF is a general attack applicable to many existing attacks. It applies the above two operations on the high-level features to other attack methods, *e.g.* MI-FGSM, NI-FGSM, during the forward propagation of features to prevent the adversarial examples from over-fitting the decision surface of the surrogate model. These two operations help the attacker eliminate the variance in the features and focus on the invariant features across different models, which stabilizes the gradient calculation and results in highly transferable adversarial examples.

### 3.3 Features under Quantitative Approaches

In this section, we analyze why high-level features are better than low-level features when diversifying the features. Without considering the memory footprint of models, the learnable parameters have a specific distribution of float-point numbers. The quantized neural networks reduce the size of the models by limiting the float point numbers to finite values:

$$Q(\theta_i) = q_j, \theta_i \in (r_j, r_{j+1}], \quad (2)$$

where  $(r_j, r_{j+1}]$  represents the range of real values depending on the quantization bits  $k$ , *i.e.*,  $j = 0, \dots, 2^k - 1$ . The quantization process map this range  $(r_j, r_{j+1}]$  to the value  $q_j$ .

However, directly using a low-bit precision to quantize a model will decrease the model accuracy. A promising approach to achieve a trade-off between model size and accuracy is to mix different precisions to quantize models. The HAWQ algorithm [12] utilizes the Hessian eigenvalues as the sensitivity measurement matrix for hierarchical features in the network and adopts mixed-precision quantization for network layers with divergent sensitivities. This motivates us to consider sensitivity when modifying network features. In other words, changing parts of the model at different levels will have different effects on how the model works. To measure network layers, we adopt the average Hessian trace  $H$  proposed by the HAWQ-V2 algorithm as the sensitivity matrix. We have the following Lemma:

**Lemma 1 [13].** *Suppose that the model is twice differentiable and converges to a local minimum when we quantize two layers with the same perturbation, we have:*

$$\mathcal{L}(\theta_1^* + \Delta\theta_1^*, \theta_2^*, \dots, \theta_N^*) \leq \mathcal{L}(\theta_1^*, \theta_2^* + \Delta\theta_2^*, \dots, \theta_N^*), \text{ if } \frac{1}{d_1} \text{Tr}(\nabla_{\theta_1}^2 \mathcal{L}(\theta_1^*)) \leq \frac{1}{d_2} \text{Tr}(\nabla_{\theta_2}^2 \mathcal{L}(\theta_2^*)), \quad (3)$$

where  $d_i$  represents the dimension of the  $i$ -th parameter  $\theta_i$  and  $\theta_i^*$  represents the converging result of the  $i$ -th layer in deep networks. The average Hessian trace of the  $i$ -th layer can be represented by  $\frac{1}{d_i} \text{Tr}(\nabla_{\theta_i}^2 \mathcal{L}(\theta_i^*))$ .

It can be concluded from Lemma 1 that the average Hessian trace can measure the sensitivity of the layers of the deep model. When the same changes are made, the effect on the loss function is bigger if the average Hessian trace is bigger. As the network layer increases, the average Hessian trace descends, the model becomes less sensitive, and the decision boundary is harder to change, as shown in [13].

A well-designed quantization scheme does not affect the decision boundary of the model but reduces the computational and memory overhead of the model, indicating that the over-parameterized model has redundancy. Lemma 1 tells that appropriately changing the point numbers in the weight matrix will not change the loss function. Hence, we have that:

**Corollary 1.** *Since the over-parameterized network has redundancy, appropriately changing the high-level features will not change the model outputs.*

With Corollary 1, we can obtain diverse features by perturbing the high-level features without changing model outputs. While enhancing the diversity of features, we should choose the less sensitive ones so that the deep model’s predictions do not change. We can use these diverse features to craft more transferable adversarial examples. Hence, we tend not to change the features of the latter layers.

### 3.4 DHF vs. Ghost

DHF pays attention to the white-box model by utilizing the redundancy of parameters to perturb the high-level features for better transferability. Li *et al.* [31] propose to generate a set of Ghost networks by densely applying dropout to each convolutional layer or random scaling on the skip connection for ResNet-based models to enhance the transferability. We summarize the difference between DHF and Ghost as follows:

- **Motivation.** DHF diversifies the features based on the redundancy of parameters to obtain invariant features while Ghost aims to reduce the training cost of ensemble model attack.
- **Strategy.** DHF perturbs the high-level features since the parameters of lower layers tend to be less redundant. However, Ghost argues that perturbation on the latter layer cannot provide transferability and densely perturbs the features.
- **Transformation.** DHF mixups the feature of current examples and benign samples and randomly replaces the features with their means while Ghost densely adopts a dropout layer for convolutional layer or random scaling on the skip connection.
- **Generalization.** As shown in Sec. 4.2, DHF could consistently boost the transferability of various attacks while Ghost sometimes degrades the performance.

## 4 Experiments

In this section, we conduct empirical evaluations to validate the effectiveness of DHF. We specify the experimental setting, evaluate DHF using momentum-based as well as input transformation-based attacks, and provide parameter studies.

Attack	Method	Res-18	Res-50	Res-101	Res-152	IncRes-v2	DenseNet-121	MobileNet	ViT	Swin
MI-FGSM	Org.	52.0	54.1	45.7	51.3	40.5	52.6	52.3	23.0	35.1
	TAP	62.8	61.6	46.5	50.8	51.9	63.9	59.0	16.7	31.7
	ILA	63.9	57.7	45.4	52.1	45.1	58.7	53.2	25.6	33.5
	SGM	64.8	73.0	47.7	53.7	51.9	70.5	63.3	29.1	45.5
	Ghost	67.3	74.5	47.5	67.0	52.9	71.4	<b>65.3</b>	28.0	45.5
	DHF	<b>71.9</b>	<b>76.7</b>	<b>47.9</b>	<b>70.2</b>	<b>57.5</b>	<b>74.7</b>	62.9	<b>35.2</b>	<b>53.2</b>
NI-FGSM	Org.	54.8	58.9	46.8	53.0	44.0	55.9	54.2	23.9	37.6
	TAP	63.4	63.9	45.7	58.2	57.6	57.3	59.7	20.1	34.5
	ILA	64.7	62.0	47.6	63.9	48.1	57.1	62.4	28.3	48.1
	SGM	65.3	75.9	47.8	55.3	50.6	73.9	66.1	30.2	48.7
	Ghost	69.3	75.8	47.9	69.5	54.6	72.3	68.9	30.9	48.1
	DHF	<b>73.0</b>	<b>77.3</b>	<b>48.5</b>	<b>74.8</b>	<b>60.3</b>	<b>77.1</b>	<b>71.5</b>	<b>35.8</b>	<b>55.9</b>

Table 1: Average black-box attack success rates (%) on nine models by two momentum-based attacks. The adversarial examples are generated on Res-101, Res-152 and IncRes-v2, respectively.

## 4.1 Experimental Setting

**Dataset.** We adopt 1,000 images [53] in 1,000 categories, which are randomly sampled from ILSVRC 2012 validation set [27]. All the chosen models can correctly classify the images.

**Models.** We first evaluate the attack performance on nine models of two popular architectures, namely Convolutional Network Works (CNNs), *i.e.*, ResNet-18 (Res-18) [20], ResNet-50 (Res-50) [20], ResNet-101 (Res-101) [20], ResNet-152 (Res-152) [20], Inception-ResNet-v2 (IncRes-v2) [49], DenseNet-121 [22], MobileNet [21], and Transformers, *i.e.*, Vision Transformer (ViT) [14] and Swin [35]. We also consider several models with defense mechanisms, including the top-3 submissions in the NIPS 2017 defense competition: HGD [32], R&P [66], NIPS-r3<sup>1</sup>, input preprocessing: Bit-Red [69], FD [36], JPEG [19], a certified defense RS [8], and an adversarial perturbation denoiser NRP [41].

**Evaluation Settings.** To validate the effectiveness of DHF, we consider two momentum-based attacks, *i.e.*, MI-FGSM [10] and NI-FGSM [33]. We also adopt two input transformations, namely DIM [68] and TIM [11]. DIM adopts the transformation probability 0.5, and TIM utilizes the Gaussian kernel with size  $7 \times 7$ . Our DHF perturbs the last  $\frac{5}{6}$ s layers with the upper bound of mixup weight  $\eta_{max} = 0.2$  and the portion of perturbed elements  $\rho = 10\%$ . We set the the perturbation budget  $\epsilon = 16$ , which is a typical black-box setting [10, 23, 31, 63, 74], and the number of iteration  $T = 10$ , step size  $\alpha = 1.6$  and decay factor  $\mu = 1.0$ .

## 4.2 Numeric Results

We evaluate our DHF using two momentum-based attacks, namely MI-FGSM [10] and NI-FGSM [33], and two input transformation-based attacks, *i.e.*, DIM [68] and TIM [11]. We adopt the original model and several architecture-related approaches as our baselines, namely TAP [74], ILA [23], SGM [63], ghost network [31]. Since some baselines (*e.g.*, SGM) are limited to ResNets, we take Res-101, Res-152 and IncRes-v2 as the white-box models but DHF is general to any architectures. The attack performance is quantified by the attack success rates, which are the misclassification rates of the corresponding models on the generated adversarial examples. In Sec. 1.1 of our supplementary, we provide the results of the

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

Attack	Method	Res-18	Res-50	Res-101	Res-152	IncRes-v2	DenseNet-121	MobileNet	ViT	Swin
DIM	Org.	61.2	58.7	51.5	69.5	47.7	58.9	62.2	30.6	40.5
	TAP	68.4	66.9	53.9	57.1	56.6	69.0	63.9	30.1	41.1
	ILA	72.5	63.9	52.8	60.2	59.3	69.2	65.4	38.5	46.2
	SGM	74.8	75.9	56.6	64.3	62.1	75.6	71.9	40.6	54.9
	Ghost	78.1	80.2	61.9	70.8	70.3	85.2	80.9	47.5	59.6
	DHF	<b>85.3</b>	<b>86.9</b>	<b>64.2</b>	<b>75.8</b>	<b>72.9</b>	<b>85.7</b>	<b>83.6</b>	<b>52.9</b>	<b>62.7</b>
TIM	Org.	58.3	60.2	49.9	57.1	50.2	60.8	59.5	27.1	43.2
	TAP	65.1	68.0	50.2	58.9	60.2	63.8	65.9	28.8	38.4
	ILA	69.4	65.7	51.6	65.2	62.6	64.9	65.9	40.3	52.3
	SGM	71.9	78.3	55.6	67.5	64.3	70.0	67.2	44.5	54.2
	Ghost	76.2	81.9	63.5	68.2	69.3	79.1	75.7	48.3	58.7
	DHF	<b>82.6</b>	<b>88.6</b>	<b>67.2</b>	<b>74.0</b>	<b>74.3</b>	<b>86.9</b>	<b>80.7</b>	<b>53.9</b>	<b>61.3</b>

Table 2: Average black-box attack success rates (%) on nine models by two input transformation-based attacks. The adversarial examples are generated on Res-101, Res-152 and IncRes-v2, respectively.

comparison with more advanced attacks.

**Evaluations on Momentum-based Attacks.** We first evaluate the performance improvement of MI-FGSM and NI-FGSM.

The average attack success rates on the black-box models are summarized in Table 1. We can observe that all the approaches can boost the adversarial transferability than generating on the original models, in which our DHF achieves the best attack performance on either CNNs or transformers. In particular, DHF outperforms the best baseline (Ghost) with a clear margin of 3.4% and 4.1% on average for MI-FGSM and NI-FGSM, respectively. With the same computation cost, such superior attack performance and its generality to CNNs and ViTs support the high effectiveness of DHF.

**Evaluations on Input Transformation-based Attacks.** Input transformation-based attack is another mainstream approach to boost adversarial transferability. To further validate the effectiveness of DHF, we adopt two input transformation-based methods under the same setting, namely DIM, and TIM. These two approaches are integrated into MI-FGSM and the results are presented in Tab. 2. Either DIM or TIM achieves better transferability than MI-FGSM (in Tab. 1) on the original model and the baselines can consistently boost the transferability. Among these attacks, DHF outperforms the runner-up method Ghost on the adversarial transferability with a notable gap of 4.0% and 5.4% on DIM and TIM, respectively. Such superior performance further supports the high effectiveness of DHF and validates our hypothesis that utilizing the redundancy of parameters to perturb the high-level features can significantly improve the transferability.

**Evaluations on Defense Models.** Recently, various defense methods have been proposed to mitigate the threat of adversarial examples. To fully evaluate the effectiveness of DHF, we further consider several models with defense mechanisms, including HGD [32], R&P [66], NIPS-r3<sup>2</sup>, Bit-Red [69], FD [36], JPEG [19], RS [8], and NRP [41]. We generate the adversarial examples using DIM on Res-101, Res-152, and IncRes-v2 models, respectively and test them on the aforementioned eight defense models. We select Res-18 model as the target model for the input processing-based defense, *i.e.*, Bit-Red, FD, JPEG, and NRP. The other four defense approaches adopt the official models provided in the corresponding papers. The average success attack rates are presented in Tab. 3. Surprisingly, Ghost, which

<sup>2</sup><https://github.com/anlhms/nips-2017/tree/master/mmd>



Method	HGD	R&P	NIPS-r3	Bit-Red	FD	JPEG	RS	NRP	Average
Org.	52.9	45.5	50.9	34.0	48.1	56.2	20.7	36.0	43.1
TAP	54.1	47.2	51.2	45.1	50.6	62.4	25.5	46.4	47.8
ILA	56.3	48.0	53.7	47.6	52.3	64.8	22.9	46.5	49.0
SGM	54.9	48.1	54.6	47.4	52.9	70.2	23.8	48.3	50.9
Ghost	35.1	29.8	30.6	31.4	38.9	49.3	19.0	33.9	33.5
DHF	<b>60.1</b>	<b>54.7</b>	<b>55.6</b>	<b>52.0</b>	<b>61.7</b>	<b>75.3</b>	<b>34.3</b>	<b>55.7</b>	<b>56.2</b>

Table 3: Average black-box attack success rates (%) by DIM on various defense models. The adversarial examples are crafted on Res-101, Res-152, and IncRes-v2, respectively.

achieves the best performance among the baselines in the above evaluations, achieves the poorest performance on these defense methods, even poorer than the adversarial examples generated on the original model. By contrast, DHF consistently performs better than the baselines and outperforms the runner-up method (SGM) with a clear margin of 5.3% on average, showing its high effectiveness and generality when attacking the black-box models with different defense mechanisms.

### 4.3 Parameter Studies

DHF randomly perturbs the high-level features by mixing up the high-level features with clean features and randomly adjusting the mixed features. In this section, we conduct parameter studies to investigate the impact of three hyper-parameters, namely the upper bound of mixup weight  $\eta_{max}$ , the ratio of elements to be adjusted  $\rho$  and the number of layers to be changed. We generate the adversarial examples on Res-101, Res-152, and IncRes-v2, and report the average attack performance on the other eight models.

**On the upper bound of mixup weight  $\eta_{max}$ .** As shown in Eq. (1),  $\eta_{max}$  balances the feature of clean sample and adversarial example among the mixed feature. To find a proper value for  $\eta_{max}$ , we conduct DHF attack to perturb the last five-sixths layers with  $\eta_{max}$  from 0 to 0.5 in a step of 0.05 using  $\rho = 0.1$ . As shown in Fig. 3 (a), when  $\eta_{max} = 0$ , the mixup operation cannot take effect and DHF achieves low transferability. When we increase  $\eta_{max}$ , the attack performance consistently increases and achieves the peak around  $\eta_{max} = 0.2$ . This further validates that adopting an appropriate portion of clean features can boost adversarial transferability. When  $\eta_{max}$  continually increases, the clean features take a more significant portion in the mixed feature, making it hard to calculate the accurate gradient. As a result, the attack performance starts to decay. Hence, we adopt  $\eta_{max} = 0.2$  in our experiments.

**On the ratio of elements to be adjusted  $\rho$ .**  $\rho$  reduces the variance among the features to identify the invariant features. To determine the value of  $\rho$ , we conduct DHF attack to perturb the last five-sixths layers with  $\rho$  from 0 to 0.5 in a step of 0.05 using  $\eta_{max} = 0.2$ . As shown in Fig. 3 (b), when increasing  $\rho$ , the overall performance slightly increases when  $\rho \leq 0.1$ , and arrives at the peak when  $\rho = 0.1$ . Larger  $\rho$  will replace more elements with the mean of feature, which decays the classification accuracy as shown in Fig. 1 and results in the inaccurate gradient. Hence, the attack performance significantly drops and we adopt  $\rho = 0.1$  in our experiments.

**On the ratio of layers to be adjusted  $r$ .** To determine a suitable ratio of layers for adjusting the features, we conduct DHF attack to perturb a different number of layers with  $\rho = 0.1$  and  $\eta_{max} = 0.2$ .




Figure 3: Average black-box attack success rates (%) on nine models by DHF using various hyper-parameters. The adversarial examples are generated on Res-101, Res-152 and IncRes-v2, respectively.

As shown in Fig. 3 (c), when DHF does not perturb the features (*i.e.*,  $r = 0\%$ ), DHF cannot take effect and achieves the lowest transferability. When we start to perturb the features (*e.g.*,  $r = 16\%$ ), the transferability can be significantly enhanced. The transferability can be consistently improved when DHF perturbs more features and achieves the peak when DHF perturbs the last  $\frac{5}{6}$  layers (*i.e.*,  $r = 83\%$ ). Perturbing all the layers results in lower performance than perturbing the last five-sixths layers, which is also consistent with our analysis in Sec. 3.3. In our experiments, we perturb the last five-sixths of layers for better performance.

## 5 Conclusion

DNNs are often over-parameterized for good generalization. In this work, we utilize such property to enhance adversarial transferability. Specifically, we find that small perturbations in high-level features have a negligible impact on overall performance. Motivated by this observation, we propose DHF, which perturbs high-level features by randomly transforming and mixing them with benign sample features during gradient calculation. We also provide a theoretical analysis of why high-level features are better than low-level features when diversifying the features. Our extensive evaluations demonstrate that the proposed method achieves significantly better adversarial transferability than existing state-of-the-art attacks.

## References

- [1] Ron Banner, Itay Hubara, Elad Hoffer, and Daniel Soudry. Scalable Methods for 8-bit Training of Neural Networks. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 5151–5159, 2018.
- [2] 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.

- [3] Yaohui Cai, Zhewei Yao, Zhen Dong, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. ZeroQ: A Novel Zero Shot Quantization Framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13166–13175, 2020.
- [4] Florian Chabot, Mohamed Chaouch, Jaonary Rabarisoa, Céline Teulière, and Thierry Chateau. Deep MANTA: A Coarse-to-Fine Many-Task Network for Joint 2D and 3D Vehicle Analysis from Monocular Image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1827–1836, 2017.
- [5] Pin-Yu Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui Hsieh. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *Proceedings of the 10th ACM workshop on artificial intelligence and security*, pages 15–26, 2017.
- [6] Brian Chmiel, Liad Ben-Uri, Moran Shkolnik, Elad Hoffer, Ron Banner, and Daniel Soudry. Neural Gradients are Near-lognormal: Improved Quantized and Sparse Training. In *Proceedings of the International Conference on Learning Representations*, 2021.
- [7] Felipe Codevilla, Antonio M. López, Vladlen Koltun, and Alexey Dosovitskiy. On Offline Evaluation of Vision-Based Driving Models. In *Proceedings of the European Conference on Computer Vision*, pages 246–262, 2018.
- [8] Jeremy 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.
- [9] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Training deep neural networks with low precision multiplications. *arXiv preprint arXiv:1412.7024*, 2014.
- [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.
- [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.
- [12] Zhen Dong, Zhewei Yao, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. HAWQ: Hessian AWARE Quantization of Neural Networks With Mixed-Precision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 293–302, 2019.
- [13] Zhen Dong, Zhewei Yao, Daiyaan Arfeen, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. HAWQ-V2: Hessian Aware Trace-Weighted Quantization of Neural Networks. In *Proceedings of the Advances in Neural Information Processing Systems*, 2020.

- [14] 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.
- [15] Hongbo Gao, Bo Cheng, Jianqiang Wang, Keqiang Li, Jianhui Zhao, and Deyi Li. Object classification using cnn-based fusion of vision and lidar in autonomous vehicle environment. *IEEE Transactions on Industrial Informatics*, 14(9):4224–4231, 2018.
- [16] Zhijin Ge, Fanhua Shang, Hongying Liu, Yuanyuan Liu, Liang Wan, Wei Feng, and Xiaosen Wang. Improving the Transferability of Adversarial Examples with Arbitrary Style Transfer. *arXiv preprint arXiv:2308.10601*, 2023.
- [17] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and Harnessing Adversarial Examples. In *Proceedings of the International Conference on Learning Representations*, 2015.
- [18] 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.
- [19] 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.
- [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.
- [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.
- [22] 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.
- [23] Qian Huang, Isay Katsman, Zeqi Gu, Horace He, Serge J. Belongie, and Ser-Nam Lim. Enhancing Adversarial Example Transferability With an Intermediate Level Attack. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4732–4741, 2019.
- [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.
- [25] Nathan Inkawhich, Wei Wen, Hai (Helen) Li, and Yiran Chen. Feature Space Perturbations Yield More Transferable Adversarial Examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7066–7074, 2019.

- [26] Joel Janai, Fatma Güney, Aseem Behl, Andreas Geiger, et al. Computer vision for autonomous vehicles: Problems, datasets and state of the art. *Foundations and Trends® in Computer Graphics and Vision*, 12(1–3):1–308, 2020.
- [27] 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.
- [28] 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.
- [29] Huichen Li, Xiaojun Xu, Xiaolu Zhang, Shuang Yang, and Bo Li. QEBA: Query-Efficient Boundary-Based Blackbox Attack. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1218–1227, 2020.
- [30] 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.
- [31] 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.
- [32] 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.
- [33] 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.
- [34] 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.
- [35] 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.
- [36] 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.
- [37] Yuyang Long, Qilong Zhang, Boheng Zeng, Lianli Gao, Xianglong 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.

- [38] 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.
- [39] Krishna Kanth Nakka and Mathieu Salzmann. Learning Transferable Adversarial Perturbations. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 13950–13962, 2021.
- [40] Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, Boaz Barak, and Ilya Sutskever. Deep Double Descent: Where Bigger Models and More Data Hurt. In *Proceedings of the International Conference on Learning Representations*, 2020.
- [41] 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.
- [42] Haibo Qiu, Dihong Gong, Zhifeng Li, Wei Liu, and Dacheng Tao. End2End Occluded Face Recognition by Masking Corrupted Features. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 44(10):6939–6952, 2022.
- [43] Aditi Raghunathan, Jacob Steinhardt, and Percy Liang. Certified Defenses against Adversarial Examples. In *Proceedings of the International Conference on Learning Representations*, 2018.
- [44] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [45] Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *Proceedings of the International Conference on Learning Representations*, 2015.
- [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.
- [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.
- [48] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the Inception Architecture for Computer Vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, 2016.
- [49] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A. Alemi. Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4278–4284, 2017.

- [50] Xiaou Tang and Zhifeng Li. Video based face recognition using multiple classifiers. In *IEEE International Conference on Automatic Face and Gesture Recognition*. IEEE, 2004.
- [51] 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.
- [52] Kunyu Wang, Xuanran He, Wenxuan Wang, and Xiaosen Wang. Boosting Adversarial Transferability by Block Shuffle and Rotation. *arXiv preprint arXiv:2308.10299*, 2023.
- [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.
- [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.
- [55] Xiaosen Wang, Xuanran He, Jingdong Wang, and Kun He. Admix: Enhancing the Transferability of Adversarial Attacks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16138–16147, 2021.
- [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.
- [57] Xiaosen Wang, Chuanbiao Song, Liwei Wang, and Kun He. Multi-stage Optimization Based Adversarial Training. *arXiv preprint arXiv:2106.15357*, 2021.
- [58] 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.
- [59] Xiaosen Wang, Zeliang Zhang, and Jianping Zhang. Structure Invariant Transformation for better Adversarial Transferability. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023.
- [60] Zhibo Wang, Hengchang Guo, Zhifei Zhang, Wenxin Liu, Zhan Qin, and Kui Ren. Feature Importance-aware Transferable Adversarial Attacks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7619–7628, 2021.
- [61] 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.
- [62] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A Discriminative Feature Learning Approach for Deep Face Recognition. In *Proceedings of the European Conference on Computer Vision*, pages 499–515, 2016.

- [63] 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.
- [64] Zifeng Wu, Chunhua Shen, and Anton Van Den Hengel. Wider or deeper: Revisiting the resnet model for visual recognition. *Pattern Recognition*, 90:119–133, 2019.
- [65] Chaowei Xiao, Bo Li, Jun-Yan Zhu, Warren He, Mingyan Liu, and Dawn Song. Generating Adversarial Examples with Adversarial Networks. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 3905–3911, 2018.
- [66] 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.
- [67] Cihang Xie, Yuxin Wu, Laurens van der Maaten, Alan L. Yuille, and Kaiming He. Feature Denoising for Improving Adversarial Robustness. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 501–509, 2019.
- [68] 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.
- [69] 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.
- [70] Xiaolong Yang, Xiaohong Jia, Dihong Gong, Dong-Ming Yan, Zhifeng Li, and Wei Liu. LARNet: Lie Algebra Residual Network for Face Recognition. In *Proceedings of the International Conference on Machine Learning*, pages 11738–11750, 2021.
- [71] Yichen Yang, Xiaosen Wang, and Kun He. Robust Textual Embedding against Word-level Adversarial Attacks. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, page 2214–2224, 2022.
- [72] Jianping Zhang, Weibin Wu, Jen-tse Huang, Yizhan Huang, Wenxuan Wang, Yuxin Su, and Michael R. Lyu. Improving Adversarial Transferability via Neuron Attribution-based Attacks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14973–14982, 2022.
- [73] 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.
- [74] Wen Zhou, Xin Hou, Yongjun Chen, Mengyun Tang, Xiangqi Huang, Xiang Gan, and Yong Yang. Transferable Adversarial Perturbations. In *Proceedings of the European Conference on Computer Vision*, pages 471–486, 2018.