

LFAA: Crafting Transferable Targeted Adversarial Examples with Low-Frequency Perturbations

Kunyu Wang^{a,*}, Juluan Shi^a and Wenxuan Wang^a

^aThe Chinese University of Hong Kong

Abstract. Deep neural networks are susceptible to adversarial attacks, which pose a significant threat to their security and reliability in real-world applications. The most notable adversarial attacks are transfer-based attacks, where an adversary crafts an adversarial example to fool one model, which can also fool other models. While previous research has made progress in improving the transferability of untargted adversarial examples, the generation of targeted adversarial examples that can transfer between models remains a challenging task. In this work, we present a novel approach to generate transferable targeted adversarial examples by exploiting the vulnerability of deep neural networks to perturbations on high-frequency components of images. We observe that replacing the high-frequency component of an image with that of another image can mislead deep models, motivating us to craft perturbations containing high-frequency information to achieve targeted attacks. To this end, we propose a method called Low-Frequency Adversarial Attack (LFAA), which trains a conditional generator to generate targeted adversarial perturbations that are then added to the low-frequency component of the image. Extensive experiments on ImageNet demonstrate that our proposed approach significantly outperforms state-of-the-art methods, improving targeted attack success rates by a margin from 3.2% to 15.5%.

1 Introduction

Deep neural networks (DNNs) have achieved remarkable progress in numerous domains [9]. However, the security and dependability of DNNs remain difficult to ensure when confronted with adversarial examples [6, 28]. Adversarial examples are created with malicious intent and include slight alterations that are imperceptible to the human eye but are sufficient to deceive deep models into making incorrect predictions. This weakness poses a threat to security-sensitive applications such as face recognition.

The study of adversarial attacks can be divided into two categories: white-box attacks and black-box attacks. White-box attacks have access to the victim model’s architecture and parameters to create adversarial perturbations. In contrast, black-box attacks have limited or no knowledge of the victim model. Existing white-box attacks use gradient information to generate adversarial examples [6, 19]. Adversarial examples created on white-box models can deceive unknown neural models [23], making black-box attacks feasible. Several methods have been proposed to enhance the transferability of adversarial examples [30, 2, 16], and they have achieved remarkable performance in the untargted black-box scenario, where they only aim to

produce different predictions. However, these methods are not satisfactory in the targeted attack scenario, where adversaries intend to trick the models into making specific predictions. Several approaches have been proposed to improve the transferability of adversarial examples in the targeted setting [39, 14, 32, 18], with the generator-based method being the most effective. This approach trains a generator based on the source model to produce adversarial examples for the target class [36, 24].




Figure 1: The frequency component of original image and adversarial images crafted by LFAA

Prior research on targeted attacks [6, 39, 14] has focused on generating adversarial perturbations, but these perturbations typically do not contain any information from the target class image, leading to overfitting of the models. To overcome this issue, it is important to modify the texture and shape of the input image to match the characteristics of the target label image. By incorporating such information, adversaries can generate adversarial examples with improved transferability compared to directly adding noise.

The high-frequency component of an image, as depicted in Fig. 1, typically captures the detailed textures and noise, while the low-frequency component represents the object’s shape. Since the background of an image is unnecessary for classification and can easily cause overfitting, it is essential to directly manipulate the image frequency domain information to embed the target class image information into the input image, which can improve transferability. With this insight in mind, we proposed a novel generation-based attack called Low-Frequency Adversarial Attack (LFAA). The method trains a conditional generator to produce adversarial frequency do-

* Corresponding Author. Email: kunyuwang@link.cuhk.edu.hk

main information corresponding to the target label, which is then added to the low-frequency component of the input image to generate adversarial examples (See Fig. 2). As depicted in Fig. 1, the perturbation modifies the high-frequency gesture and blurs the low-frequency component to emphasize high-frequency component. As a result, the classifier misclassifies the high-frequency of the adversarial examples as the target label image, indicating that we have embedded the target label information in the frequency component. Specifically, we optimize the generator by minimizing the classification loss between the predictions of the adversarial examples and the target label.

In summary, we highlight our contributions as follows:

- We demonstrate that modifying the frequency domain information can effectively mislead deep models.
- We propose a novel attack called Low-Frequency Adversarial Attack (LFAA), which is the first generation-based targeted attack that generates adversarial frequency information to embed target label information into the source image.
- Empirical evaluations on the ImageNet dataset demonstrate LFAA achieves much better transferability than the state-of-the-art target attacks.

2 Related Work

In this section, we review related adversarial attacks in the black-box setting, frequency-based attacks, and defenses related to our proposed attack.

Adversarial Attack Szegedy *et al.* were the first to demonstrate the vulnerability of deep neural networks to adversarial examples [28]. Subsequently, several works have investigated the susceptibility of deep models to black-box attacks, where access to the target model is restricted. Such attacks can be categorized into three types: (a) score-based attacks, which can access the predicted probability [11]; (b) decision-based attacks, which can only obtain the predicted label [1]; and (c) transfer-based attacks, which are effective in real-world settings where the attacker cannot query the target model. In this approach, the adversary crafts the adversarial examples on a white-box model and transfers them to fool the target model [16, 30]. Therefore, the transferability of adversarial examples is crucial for deceiving unknown models.

The Fast Gradient Sign Method (FGSM) was the first gradient-based attack proposed, which generates perturbations in the direction of the gradient [6]. The iterative FGSM is an extension of FGSM that produces better white-box performance but worse black-box performance. However, to avoid overfitting to the white-box model, several methods have been proposed to escape local maxima. For instance, MI-FGSM [2] adds momentum to I-FGSM to stabilize the optimization process. Lin *et al.* further improve the method by incorporating the Nesterov accelerated gradient, which provides an effective look-ahead strategy [16].

In addition to gradient-based attacks, data augmentation techniques have also been found to be effective in improving the transferability of adversarial examples. For instance, the Diverse Input Method (DIM) [35] rescales the input image to random sizes and adds padding to a fixed size before calculating the gradient. The Transition Invariant Attack [3] approximates the gradient calculation for a set of translated images by convolving the gradient with a Gaussian kernel. Admix [30] mixes the input image with a small portion of images from other categories. Some methods also focus on modifying surrogate models to enhance the transferability of iterative methods. For example, skip connections [33] have been used to

improve the transferability of adversarial examples. Finally, other attacks aim to increase the feature difference between the source image and adversarial examples [32].

Although the methods mentioned above have shown good performance in untargeted attacks, their performance in targeted attacks is usually poor. To address this, several researchers have proposed advanced loss functions for targeted attacks. Li *et al.* [14] adopt Poincare distance and Triplet loss to replace cross-entropy, which has a vanishing gradient in iterative targeted attacks. Zhao *et al.* [39] use the logits output of the target class as the loss function and perform a large number of iterations to achieve state-of-the-art performance. Gao *et al.* [4] minimize the distance between the features of the target sample and adversarial examples in the reproducing kernel Hilbert space, which is transition-invariant. Another branch of attack is generative-based attacks, where attackers train a generator to generate perturbations given an input image. Generative-based attacks are more efficient because they can learn the adversarial pattern of a target label using a large dataset. The adversarial pattern depends on the entire data distribution and is not limited to a single image, which may overfit the source model. UAP [20] was the first proposed method to fool models by learning universal noise, GAP [24] learns a generator that can produce image-agnostic perturbations for targeted attacks, CDA [21] uses relativistic training objectives to boost cross-domain transferability, and C-GSP [36] uses a conditional generator to generate targeted adversarial perturbations. Our method belongs to the generative-based method, but instead of generating perturbations directly, we generate frequency content that can easily fool deep models. We seek to learn the frequency components of a target class image based on a data distribution.

Frequency-based Attack According to recent studies, researchers have explored the generalization and adversarial vulnerability of deep neural networks (DNNs) from a frequency perspective [37, 31, 25]. These studies indicate that DNNs can capture high-frequency components that are imperceptible to humans. Yin *et al.* [37] demonstrate that naturally trained models are vulnerable to perturbations on high-frequency components, while adversarially trained models are less sensitive to such perturbations. As a result, several approaches have been proposed to craft adversarial examples from the perspective of frequency. For example, Long *et al.* [18] perturb input images with Gaussian noise in the frequency domain as a data augmentation technique, then convert them back to the spatial domain for gradient calculation to enhance transferability. Guo *et al.* [7] restrict the search space to the low-frequency domain to craft adversarial examples, showing that low-frequency components are important in model decision making for query-based attacks. Sharma *et al.* [25] craft adversarial examples by randomly masking low-frequency components and demonstrate that even adversarially trained models are still vulnerable to low-frequency perturbations. Finally, Zhang *et al.* [38] propose a method of crafting adversarial examples by replacing the high-frequency portion of an image with a handcrafted adversarial patch. However, this method cannot be used for targeted attacks, and the selection of the adversarial patch remains a concern.

Adversarial Defense As the threat of adversarial attacks on deep neural networks (DNNs) continues to increase, various defense methods have been proposed to mitigate this problem. Adversarial training, which involves injecting adversarial examples into the training process, is a promising method that has shown success in improving model robustness [6, 19]. Tramèr *et al.* [29] introduced ensemble adversarial training, which uses adversarial examples generated on multiple models to improve the robustness of the resulting model. Another approach to adversarial defense involves using denoising fil-




Figure 2: The procedure of our proposed LFAA

ters, which remove strange patterns from adversarial examples before feeding them to the classifier. For example, Liao *et al.* [15] proposed a High-level representation guided denoiser (HGD) to suppress perturbations, while Naseer *et al.* [22] trained a neural representation purifier (NRP) that learns to purify perturbed input images. Other defense methods utilize input transformations to mitigate the effects of adversarial perturbations, such as random resizing and padding (R&P) [35], and feature distillation [17].

3 Method

In this section, we will present the adversarial setup and provide a comprehensive description of our proposed LFAA.

3.1 Preliminaries

We begin by considering an image $\mathbf{x} \in \mathbf{R}^{H \times W \times C}$ from a dataset \mathcal{X} . An adversarial example \mathbf{x}_{adv} is generated such that it is nearly identical to the original image, i.e., $\|\mathbf{x} - \mathbf{x}_{adv}\|_p \leq \epsilon$, where $\|\cdot\|_p$ is the L_p norm distance and ϵ is the perturbation budget. We adopt the L_∞ distance metric in this work. We consider a classification neural network \mathcal{F} with parameters ϕ and loss function \mathbf{J} that is trained to classify images into a set of classes $\mathcal{C} = c_1, c_2, \dots, c_d$. The function $\mathcal{F}_\theta : \mathbf{R}^{H \times W \times C} \rightarrow \mathbf{R}^d$ maps an image to a class probability vector with d classes. The predicted class for a given sample image \mathbf{x} is $\text{argmax}_{i \in \mathcal{C}} \mathcal{F}_\theta(\mathbf{x})_i$. A targeted adversarial example is generated to mislead the classifier \mathcal{F} into predicting a target label $c \in \mathcal{C}$.

The optimization problem for generating a targeted adversarial example can be formulated as follows:

$$\mathbf{x}_{adv} = \underset{\|\mathbf{x}_{adv} - \mathbf{x}\|_\infty \leq \epsilon}{\text{argmin}} J(\mathbf{x}, c; \phi) \quad (1)$$

Some methods like FGSM[6], solve this problem by adding perturbation along the direction of the gradient:

$$\mathbf{x}_{adv} = \text{clip}_{\mathbf{x}, \epsilon} [\mathbf{x} - \text{sign}(\nabla_{\mathbf{x}} J(\mathbf{x}, c; \phi))]. \quad (2)$$

where clip is a function, that clip $\mathbf{x} - \text{sign}(\nabla_{\mathbf{x}} J(\mathbf{x}, c; \phi))$ into ϵ -ball of \mathbf{x} .

3.2 Frequency filter

Prior research has revealed that CNNs trained on the ImageNet dataset have a strong bias towards the texture and shape of an object [5]. As illustrated in Fig. 1, the high-frequency component of an image represents its texture, while the low-frequency component represents its shape. To separate the different frequency components of

Table 1: Untargeted and targeted attack success rate of several pre-trained models predicting adversarial images with substituted high-frequency component

Metric	ResNet101	DenseNet121	Vgg19 _{BN}	MobileNet _{v2}	Inc-v3
UASR (%)	21.4	40.1	35.4	41.5	32.5
TASR (%)	6.10	0.30	1.60	0.40	0.9

an input image, various methods such as DCT and Fourier transform can be employed. In this study, we use convolution with a low-pass Gaussian filter W as an approximation to obtain the low-frequency part of the image, i.e., $W * \mathbf{x}$. W is a $(4k + 1) \times (4k + 1)$ kernel matrix, and its formulation is as follows:

$$W_{i,j} = \frac{1}{2\pi\sigma^2} \exp - \frac{i^2 + j^2}{2\sigma^2} \quad (3)$$

where $\sigma = k$ denotes the radius of W . By using a larger σ , more high-frequency parts will be filtered. The high-frequency part of an image can be obtained by subtracting its low-frequency part from the original image, i.e., $\mathbf{x} - W * \mathbf{x}$.

Inspired by these findings, several works have been proposed to analyze and attack the vulnerability of deep models from the perspective of frequency, such as Fourier analysis [37], high-frequency perturbations [31], and practical frequency-based attacks [38]. These works show that deep models are vulnerable to perturbations in both high-frequency and low-frequency components. However, there is still a lack of discussion on how to use frequency components to craft targeted class perturbations.

In Fig. 1, it is evident that image models are highly sensitive to high-frequency components, which makes them capable of predicting these components accurately. To leverage this capability, we attempted to replace the high-frequency components of an image with the high-frequency components of a target class image. Specifically, we used the equation $W * \mathbf{x} + (\mathbf{x}_{target} - W * \mathbf{x}_{target})$ to replace the high-frequency components, as described in Tab. 1. We sampled 1000 images from ImageNet [13] and measured the performance of targeted and untargeted attacks after replacing the high-frequency components with those from a 'bull mastiff' image. Our results showed a promising attack success rate for untargeted attacks, but the performance for targeted attacks was unsatisfactory. However, it was still better than several existing methods, such as MIFGSM. The results described above demonstrate the feasibility of embedding critical frequency information to achieve a targeted attack.

Yin *et al.* [37] finds that crafting only high-frequency content in a perturbation results in a small and easily denoised perturbation,

which can be classified by adversarially trained models. However, Long *et al.* [18] suggest addressing the vulnerability of deep models by considering different frequency components. Thus, perturbations should also affect the low-frequency content of an image. Merely replacing the low-frequency content of an image with that of a target class image may cause the resulting image to fall outside the ϵ -ball of the original image \mathbf{x} . Conversely, high-frequency components of an image can be considered as small perturbations that can be entirely replaced by a targeted texture. Moreover, to incorporate the targeted information, the low-frequency components must be altered. A generator can be trained to create perturbations based on these two principles.

3.3 LFAA

Taking inspiration from the previous discussion, we suggest a generative method that directly perturbs the low-frequency component of images to create adversarial examples. This perturbation will be added directly to the low-frequency portion of the image to distort low-frequency content and include targeted class information while eliminating the high-frequency information from the original image. Besides, Yang *et al.* [36] have pointed out that many existing generative methods train multiple generators for multiple target labels, which is computationally inefficient. They propose a conditional generative approach that enables multi-target class attacks by training only one generator and generating perturbations based on the target class and input images. We hence adopt their generator architecture.

As shown in Fig. 2, we train a conditional generator \mathcal{G}_θ with parameters θ on the entire dataset and class labels. The generator learns a mapping: $(\mathcal{X}, \mathbf{R}^d) \rightarrow \mathbf{R}^{H \times W \times C}$. Given sample input image \mathbf{x} , and one-hot encoding $\mathcal{I}_c \in \mathbf{R}^d$ of target class label c , \mathcal{G}_θ output adversarial perturbation containing high-frequency texture of given target class image and noise to perturb low-frequency shape. To create an adversarial example, the generated perturbation is directly added to the low-frequency component of the image using Eq. 3, and then projected onto the ϵ -ball of the original image \mathbf{x} :

$$\mathbf{x}_{adv} = \text{clip}_{\mathbf{x}, \epsilon}(\mathcal{G}_\theta(\mathbf{x}, \mathcal{I}_c) + W * \mathbf{x}) \quad (4)$$

Given a pretrained network \mathcal{F}_ϕ , and dataset \mathcal{X} , the training objective is

$$\mathbf{E}_{\mathbf{x} \sim \mathcal{X}} [\text{CE}(\mathcal{F}_\phi(\text{clip}_{\mathbf{x}, \epsilon}(\mathcal{G}_\theta(\mathbf{x}, \mathcal{I}_c) + W * \mathbf{x})), c)] \quad (5)$$

where CE is cross-entropy loss. By minimizing this objective, the generator can learn the patterns of target class images that force the classifier to make targeted predictions based on a data distribution. As a result, the generated frequency information is independent of the input image and more generalizable to different models. The optimization procedure is outlined in Algorithm 1.

Our proposed method, LFAA, and C-GSP [36] generate adversarial patterns using cross-entropy loss as the loss function. The patterns generated by both methods are structural, repeated, and semantic, with the noise containing information about the target class image. In contrast, methods such as MI-FGSM [6] generate random noise that lacks semantic information and can easily overfit a specific model. This is why generative models, which learn the pattern of the target class image with respect to the data distribution, often outperform gradient-based methods.

Furthermore, our method focuses more on high-frequency texture and perturbs the low-frequency part of the image, which is more gen-

Algorithm 1 LFAA

Input: A classifier \mathcal{F}_ϕ with parameters ϕ , randomly initialized generative network \mathcal{G}_θ , training dataset \mathbf{X} , target label c and perturbation budget ϵ .

Output: adversarial generator \mathcal{G}_θ

```

1: repeat
2:   Randomly sample batch of images  $\mathbf{x} \sim \mathcal{X}$ ;
3:   Randomly sample batch of class  $c \sim \mathcal{C}$ ;
4:   Forward images  $\mathbf{x}$  and one-hot encoding  $\mathcal{I}_c$  to Generator  $\mathcal{G}_\theta$ 
   to obtain the perturbed images  $\mathbf{x}_{adv}$  by Eq 4;
5:   Forward perturbed images to classifier  $\mathcal{F}_\phi$  to calculate the
   loss by Eq5;
6:   Backward pass and update parameters of  $\mathcal{G}_\theta$ ;
7: until  $\mathcal{G}_\theta$  converge;
8: return  $\mathcal{G}_\theta$ ;

```




Figure 3: The Adversarial Perturbation (first row) and Adversarial examples (second row) crafted by different methods

eral and less biased. The noise generated by our method is more generalizable to other models and performs better than other generative models, such as CDA.

4 Experiment

This section presents our experimental results using the ImageNet dataset [13] to assess the efficacy of our proposed LFAA method for targeted black-box attacks. We describe our experimental setup and implementation details in Section 4.1, followed by our evaluations of transferability in Section 4.2, real-world vision systems in Section 4.3, and adversarial defenses in Section 4.4. We also present an ablation study in Section 4.5.

4.1 Experimental Setup

Dataset This paper uses the ImageNet dataset for both training and testing purposes. We trained the generator on 10,000 randomly selected images from the ImageNet training set and evaluated its performance on 1000 images belonging to 1000 categories from the ImageNet validation dataset [13].

Models We adopt 6 popular models that pre-trained on ImageNet *i.e.*, ResNet50 [9], Vgg-19_{BN} [12] and DenseNet-121 [10], Inception-v3 (Inc-v3) [27], Inception-v4 (Inc-v4) [26], Inception-ResNet-v2(IncRes-v2) [26]. To evaluate the robustness of our attacks against various defense mechanisms, we consider several adversarial defense

Table 2: Targeted attack success rates (%) on six models with various adversarial attack. The adversaries are crafted on ResNet-50, DenseNet-121 and Vgg-19_{BN} respectively. * indicates white-box attacks.

Model	Attack	ResNet-50	DenseNet-121	Vgg19 _{BN}	Inc-v3	Inc-v4	IncRes-v2
ResNet-50	MI-FGSM	98.4*	0.4	0.2	0.1	0.1	0.1
	DIM	84.5*	2.4	1.1	0.8	0.2	0.1
	TIM	99.9*	1.3	0.8	0.5	0.4	0.5
	Po-Trip	100.0*	3.8	2.5	1.0	0.2	0.6
	Logits	98.1*	9.8	3.9	1.5	0.9	2.5
	CDA	77.3*	29.2	35.1	4.2	7.0	2.8
	C-GSP	92.4*	55.2	36.2	29.9	24.3	11.5
	LFAA	95.8*	80.1	66.6	38.9	33.8	15.2
DenseNet121	MI-FGSM	0.8	100.0*	0.5	0.5	0.2	0.3
	DIM	2.6	89.4*	2.6	1.8	0.5	0.1
	TIM	1.0	100.0*	0.4	0.3	0.5	0.3
	Po-Trip	2.6	100.0*	1.1	1.0	0.4	0.4
	Logits	4.3	100.0*	2.6	2.5	1.2	1.1
	CDA	55.6	89.5*	27.3	22.3	8.7	2.4
	C-GSP	51.7	92.7*	33.6	29.6	17.9	12.1
	LFAA	72.9	94.5*	48.7	33.2	39.2	26.0
Vgg19 _{BN}	MI-FGSM	0.4	0.3	99.9*	0.2	0.2	0.2
	DIM	0.7	0.7	83.4*	0.4	0.3	0.2
	TIM	0.4	0.4	100.0*	0.2	0.1	0.6
	Po-Trip	0.7	0.7	100.0*	0.6	0.3	0.3
	Logits	1.8	2.2	100.0*	0.7	1.0	0.9
	CDA	10.3	12.3	96.3*	0.7	1.1	0.1
	C-GSP	19.7	22.6	92.0*	8.2	11.7	1.3
	LFAA	29.6	29.3	93.7*	11.2	7.7	1.8

models including adversarially trained models, denoising defense, and input transformation defense methods. Specifically, we consider the adversarially trained models Inc-v3_{adv} and ensemble adversarially trained network IncRes-v2_{ens2}[29]. For denoising defense, we consider HGD [15] and NRP [22]. For input transformation defense, we consider R&P [34], NIPS-r3¹, and FD [17].

Baselines To evaluate the effectiveness of our proposed LFAA, we choose several attacks method including MI-FGSM[2], DIM [35], TIM [3] and several competitive methods on improving the transferability of target attack including Po-Trip [14] and Logits attack [39]. We also consider several generative approaches *i.e.*, CDA [21], and C-GSP[36].

Implementation Details In our experiments, we set the maximum perturbation budget ϵ to be 16. For all the baseline methods, we follow the implementation details specified in their respective papers. Our proposed LFAA is based on ResNet architecture[9] and we adopt the same architecture as a previous generative method[36]. Specifically, our generator \mathcal{G}_θ generates perturbations on the low-frequency part of the image with the same input size as the original image. The size of the Gaussian kernel k is set to be 17×17 *i.e.* $k = 4$. The classifier $\mathcal{F}\phi$ used in our experiments is a standard pre-trained model on ImageNet, and we fixed the parameters ϕ in the classifier \mathcal{F} while training the generator \mathcal{G}_θ . We use the Adam optimizer with a learning rate of 5×10^{-4} , $\beta_1 = 0.5$, and $\beta_2 = 0.999$. All experiments were conducted on a GeForce RTX 3090 GPU using a PyTorch implementation.

4.2 Evaluation on Targeted Transferability

We compare the target transferability of our proposed LFAA method with several other attack methods, including MI-FGSM[2], DIM [35], TIM [3], Logits[39], Po+Trip[14], CDA[21] and C-GSP[36]. For training-free approaches, we craft adversarial adversaries on three standard trained models and test them on six models. For generative approaches, we train the generator using three standard trained models *i.e.* ResNet50, DenseNet-121, and VGG-19_{BN}. For each target label, we generated adversarial examples for each image and evaluated their targeted attack success rate on other models. Fig. 6 displays some crafted adversarial examples by LFAA. We randomly sample four target labels from the ImageNet class set, and train generators based on each label. We evaluate the targeted attack success rate, that is the success rate of the victim models to make targeted predictions. The results are summarized in Tab. 2, each column of this table represents the model to be attacked, while each row indicates the attacker generates the adversarial examples based on the corresponding methods.

Our experiments show that instance-specific methods *i.e.* MI-FGSM, DIM, TIM, Po-Trip, and Logits, perform well in white-box targeted attacks, and they outperform generation-based attacks in white-box setting. As gradient-based approaches directly update adversarial examples by following the gradient direction from the original image. In contrast, the generative approach learns the semantic patterns from the entire dataset, which may not cover all possible cases encountered during the attack, and hence results in better performance in white-box setting. However, in black-box settings, they tend to overfit specific models, which makes it difficult to transfer the adversarial examples to other models. Generative approaches, including CDA, G-GSP, and our proposed LFAA, have higher black-box targeted attack success rates than the most powerful instance-

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




Figure 4: Targeted attack success rates (%) on eight models by G-GSP and LFAA, the generators are trained on ResNet-50, and Vgg19_{BN}.

Table 3: Target attack success rates (%) on six models under single model setting with various attack methods. The adversaries are crafted on DenseNet-121.

Attack	ResNet50	Vgg19 _{BN}	Inc-v3	Inc-v4	IncRes-v2
MI-FGSM	2.5	1.5	0.9	1.3	2.0
DIM	0.8	0.6	0.6	0.3	0.6
TIM	1.0	0.4	0.3	0.5	0.3
Po-Trip	2.5	1.5	0.9	1.3	2.0
Logits	5.8	3.6	1.8	1.6	2.7
C-GSP	18.0	16.3	5.2	3.6	2.8
LFAA	39.8	31.0	18.0	10.9	9.2

specific method Logits, with a range from 3.2% to 76.3%. Compared to instance-specific methods, the average black-box targeted attack success rate of LFAA is significantly better, with a margin ranging from 14.9% to 45.5%. LFAA is also competitive with generative models, with the average targeted success rate of LFAA higher than that of the most powerful generative approach, C-GSP, with a margin ranging from 3.2% to 15.5%. These results show that LFAA is superior in generating transferable adversarial examples and support the claim that our approach can significantly improve transferability.

To further validate the performance of LFAA, we measured the average targeted attack success rate within a family of models. We trained the generator on ResNet50 and Vgg19_{BN} and transferred it to models within the same family (e.g. ResNet50 → ResNet101, Vgg19_{BN} → Vgg16_{BN}). The results are shown in Fig. 4, our method and C-GSP have similar targeted attack success rates for white-box performance, but for the black-box setting, the performance of our method is better than C-GSP in ResNet family, with a clear margin ranging from 18.7% to 28.1%, and for Vgg family, the average targeted attack success rate is higher with a margin ranging from 5.5% to 34.4%. This means that for different architectures or within the same architecture family, our method is superior to the state-of-the-art method C-GSP. These results further validate the effectiveness of our proposed LFAA.

To evaluate LFAA's ability to generate latent representations for targeted labels in the presence of a large label set by using a randomly generated target label dataset. Each image was randomly assigned a target label from a pool of 100 randomly selected labels. To address the challenges posed by large label sets, we follow the approach proposed by Yang *et al.* [36] to divide the label set into smaller and more diverse groups. We train a conditional generator for each group and evaluate the transferability of LFAA using two models, each trained on 50 diverse categories. We perform the targeted attack using the corresponding model for the corresponding group when the target label falls into that group. To ensure a fair comparison, we also eval-




Figure 5: Successful targeted adversarial images on Google Cloud Vision generated by LFAA, the given target class is rock snake.

uated our approach against instance-based methods, and due to computing resources, most generative approaches require training multiple generators for multiple labels, which was not feasible in our case, so we only evaluated LFAA against C-GSP. We present the results of our evaluation in Tab. 3, where we summarize the targeted attack success rates of targeted attacks. In the case of instance-based methods, Logits still has the best transferability, but it is weaker than the generative-approach C-GSP. Our LFAA exhibited the strongest transferability, outperforming the most powerful C-GSP with a range of 6.4% to 21.8%. Our approach extends the capabilities of C-GSP. Our evaluation demonstrated that LFAA can effectively handle large label sets and further confirmed the superiority of LFAA.

4.3 Evaluation on Real-world Recognition System

The majority of previous works [8, 11] have used score-based attacks to fool image recognition systems in the real world, which require thousands of queries to the victim system. In contrast, we conduct an assessment of the effectiveness of targeted transfer attacks by LFAA on the widely-used Google Cloud Vision API. In particular, we generate the adversaries by our conditional generator trained on DenseNet-121 using 100 target classes and transfer the adversarial examples to fool the vision system.

The API provides a list of labels with corresponding scores indicating the confidence of the model for each label. It only returns labels with a score of 50% or higher and at most 10 labels are shown. Fig. 5 displays one of the examples where we generated the adversarial examples using the target label "rock snake." The recognition system classified the adversarial images as "snake," which validates that our generator can learn the semantic information of the target class to fool the real-world system.

4.4 Evaluation on Defense Method

To thoroughly evaluate the effectiveness of our proposed method, we assess the attack performance of LFAA against several defense mechanisms, including adversarial training where we consider adversarially trained models Inception-v3 (Inc-v3_{adv}), and ensemble adversarially trained network Inception-Resnet-v2 (IncRes-v2_{ens2}), as well as input-transformation based defenses (R&P, NIPS-r3, and FD), denoising methods (HGD and NRP). We compared our approach with C-GSP. The results are presented in Tab. 4, where we report the average target attack success rates against each defense method.

The LFAA method performs slightly worse against adversarially trained models, particularly IncRes-v2_{ens2}, but has stronger transferability against Inc-v3_{adv}. For input-transformation based defense methods, LFAA outperforms C-GSP and achieves a higher average target attack success rate with a margin of 15% against three defense methods. For denoising methods, LFAA has slightly weaker transferability against NRP, but it can still bypass HGD with a higher targeted attack success rate with a margin of 6.6%. Overall, LFAA

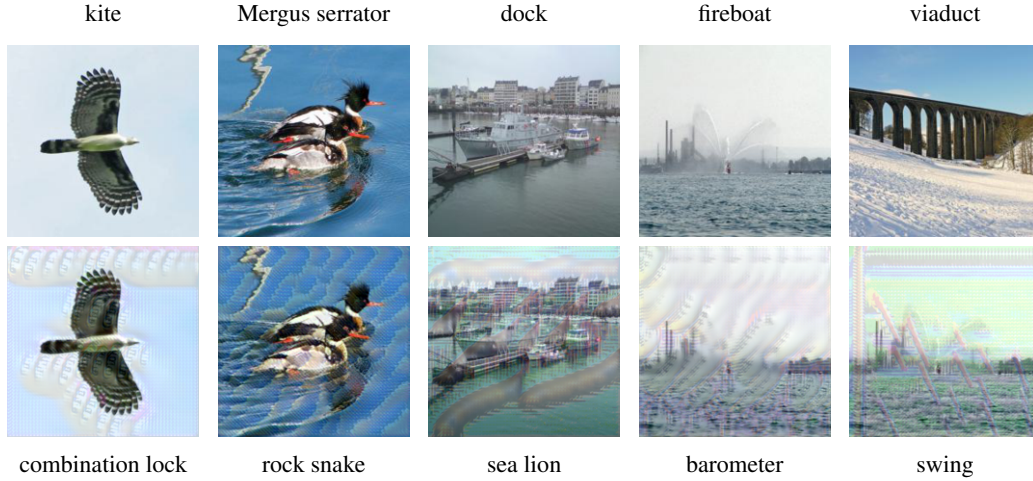


Figure 6: Targeted adversaries produced by generator trained against ResNet50. 1st row shows original images, while 2nd row shows adversaries examples. The caption in the top represents the original category. The caption below represents the target category.

Table 4: Target attack success rates (%) of seven defense methods by C-GSP and LFAA. The generators are trained on ResNet-50

Method	Inc-v3 _{adv}	IncRes-v2 _{ens}	HGD	R&P	NIPS-r3	FD	NRP	Average
C-GSP	15.9	23.5	62.0	38.3	37.1	16.0	13.2	29.4
LFAA	24.2	18.2	68.6	57.5	57.1	21.8	10.1	36.8

has a higher average targeted attack success rate against all seven defense methods than C-GSP, with a margin of 7.4%. As LFAA contains more semantics regarding the target label, and the underlying pattern still persists under different defense methods. Hence, LFAA can effectively pass some adversarial defense methods. These findings further confirm the effectiveness of our proposed LFAA method.




Figure 7: Targeted attack success rates (%) on six models by LFAA with different k .

4.5 Ablation Study

To further gain insight into the performance improvement of LFAA, we conduct ablation and hyper-parameter studies, which train our generator on DenseNet-121 and generate adversarial examples to validate them on the six models.

On the effectiveness of Gaussian kernel size $4k + 1 \times 4k + 1$ To capture the low-frequency components of the image, a Gaussian low-pass filter is utilized for approximation. As shown in Fig. 7, the

performance of transferring adversarial examples into Vgg19_{BN} and Inception-v3 models significantly drops as k increases above 5. On the other hand, when k is less than or equal to 3, the performance of transferring adversarial examples from DenseNet-121 to ResNet50 is not satisfactory. Although $k = 4$ shows weak performance in some models, it achieves the highest average targeted attack success rate. Therefore, $k = 4$ is chosen as the hyperparameter for the Gaussian low-pass filter.

5 Conclusion

In this paper, we introduce a novel method called LFAA for generating transferable targeted adversarial examples that exploit the vulnerability of deep neural models from a frequency perspective. The proposed approach is capable of generating perturbations that can cause misclassification on multiple black-box target models and real-world vision systems, regardless of the image's source class. LFAA trains a conditional generator to generate targeted adversarial perturbations, which are then added to the low-frequency components of the image. Experimental results on ImageNet demonstrate that LFAA significantly outperforms state-of-the-art methods. This work suggests that different frequency components play a crucial role in deep learning models, and targeted attacks based on perturbing these components can be an effective and efficient approach for generating transferable attacks.

References

- [1] Wieland Brendel, Jonas Rauber, and Matthias Bethge, 'Decision-Based Adversarial Attacks: Reliable Attacks Against Black-Box Machine Learning Models', *International Conference on Learning Representations*, (2018).

- [2] 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*, pp. 9185–9193, (2018).
- [3] 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*, pp. 4312–4321, (2019).
- [4] Lianli Gao, Yaya Cheng, Qilong Zhang, Xing Xu, and Jingkuan Song, 'Feature Space Targeted Attacks by Statistic Alignment', in *Proceedings of the International Joint Conference on Artificial Intelligence*, pp. 671–677, (2021).
- [5] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel, 'Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness', *International Conference on Learning Representations*, (2019).
- [6] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy, 'Explaining and harnessing adversarial examples', in *International Conference on Learning Representations*, (2015).
- [7] Chuan Guo, Jared S. Frank, and Kilian Q. Weinberger, 'Low Frequency Adversarial Perturbation', in *Proceedings of the Thirty-Fifth Conference on Uncertainty in Artificial Intelligence*, volume 115, pp. 1127–1137, (2019).
- [8] Chuan Guo, Jacob Gardner, Yurong You, Andrew Gordon Wilson, and Kilian Weinberger, 'Simple Black-box Adversarial Attacks', in *International Conference on Machine Learning*, pp. 2484–2493, (2019).
- [9] 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*, pp. 770–778, (2016).
- [10] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger, 'Densely connected convolutional networks', in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4700–4708, (2017).
- [11] Andrew Ilyas, Logan Engstrom, Anish Athalye, and Jessy Lin, 'Black-box Adversarial Attacks with Limited Queries and Information', in *International Conference on Machine Learning*, pp. 2142–2151, (2018).
- [12] Simonyan Karen and Zisserman Andrew, 'Very Deep Convolutional Networks for Large-Scale Image Recognition', in *International Conference on Learning Representations*, (2015).
- [13] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton, 'ImageNet Classification with Deep Convolutional Neural Networks', in *Advances in Neural Information Processing Systems*, pp. 1106–1114, (2012).
- [14] Maosen Li, Cheng Deng, Tengjiao Li, Junchi Yan, Xinbo Gao, and Heng Huang, 'Towards Transferable Targeted Attack', in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 638–646, (2020).
- [15] 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*, pp. 1778–1787, (2018).
- [16] Jiadong Lin, Chuanbiao Song, Kun He, Liwei Wang, and John E. Hopcroft, 'Nesterov Accelerated Gradient and Scale Invariance for Adversarial Attacks', in *International Conference on Learning Representations*, (2020).
- [17] Zihao Liu, Qi Liu, Tao Liu, Nuo Xu, Xue Liu, 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*, pp. 860–868, (2019).
- [18] Yuyang Long, Qilong Zhang, Boheng Zeng, Lianli Gao, Xianglong Liu, Jian Zhang, and Jingkuan Song, 'Frequency Domain Model Augmentation for Adversarial Attack', *European Conference on Computer Vision*, (2022).
- [19] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu, 'Towards Deep Learning Models Resistant to Adversarial Attacks', in *International Conference on Learning Representations*, (2018).
- [20] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard, 'Universal Adversarial Perturbations', in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1765–1773, (2017).
- [21] Muhammad Muzammal Naseer, Salman H Khan, Muhammad Haris Khan, Fahad Shahbaz Khan, and Fatih Porikli, 'Cross-domain transferability of adversarial perturbations', volume 32, pp. 12905–12915, (2019).
- [22] Muzammal Naseer, Salman Khan, Munawar Hayat, Khan Fahad Shahbaz, and Fatih Porikli, 'A Self-supervised Approach for Adversarial Robustness', in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 262–271, (2020).
- [23] 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*, pp. 506–519, (2017).
- [24] Omid Poursaeed, Isay Katsman, Bicheng Gao, and Serge Belongie, 'Generative adversarial perturbations', in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4422–4431, (2018).
- [25] Yash Sharma, Gavin Weiguang Ding, and Marcus A. Brubaker, 'On the effectiveness of low frequency perturbations', in *Proceedings of the International Joint Conference on Artificial Intelligence*, p. 3389–3396, (2019).
- [26] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alex Alemi, 'Inception-v4, inception-resnet and the impact of residual connections on learning', in *AAAI Conference on Artificial Intelligence*, (2017).
- [27] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Wojna Zbigniew, 'Rethinking the inception architecture for computer vision', in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2818–2826, (2016).
- [28] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus, 'Intriguing properties of neural networks', in *International Conference on Learning Representations*, (2014).
- [29] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel, 'Ensemble Adversarial Training: Attacks and Defenses', *International Conference on Learning Representations*, (2018).
- [30] Xiaosen Wang, Xuanran He, Jingdong Wang, and Kun He, 'Admix: Enhancing the transferability of adversarial attacks', in *International Conference on Computer Vision*, pp. 16138–16147, (2021).
- [31] Xindi Wang, Haohan an Wu, Zeyi Huang, and Eric P. Xing, 'High-Frequency Component Helps Explain the Generalization of Convolutional Neural Networks', in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8684–8694, (2020).
- [32] Zhibo Wang, Hengchang Guo, Zhifei Zhang, Wenxin Liu, Qin Zhan, and Kui Ren, 'Feature importance-aware transferable adversarial attacks', in *International Conference on Computer Vision*, pp. 7639–7648, (2021).
- [33] 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*, (International Conference on Learning Representations).
- [34] Cihang Xie, Jianyu Wang, Zhishuai Zhang, Zhou Ren, and Alan Yuille, 'Mitigating adversarial effects through randomization', in *International Conference on Learning Representations*, (2018).
- [35] 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 Conference on Computer Vision and Pattern Recognition*, pp. 2730–2739, (2019).
- [36] Xiao Yang, Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu, 'Boosting Transferability of Targeted Adversarial Examples via Hierarchical Generative Network', in *Proceedings of the European Conference on Computer Vision (ECCV)*, p. 725–742, (2022).
- [37] Dong Yin, Raphael Gontijo Lopes, Jonathon Shlens, Ekin D. Cubuk, and Justin Gilmer, 'A fourier perspective on model robustness in computer vision', volume 1189, p. 13276–13286, (2019).
- [38] Qilong Zhang, Chaoning Zhang, Chaoqun Li, Jingkuan Song, and Lianli Gao, 'Practical No-box Adversarial Attacks with Training-free Hybrid Image Transformation', *arXiv preprint arXiv:2203.04607*, (2022).
- [39] Zhengyu Zhao, Zhuoran Liu, and Martha Larson, 'On Success and Simplicity: A Second Look at Transferable Targeted Attacks', in *Advances in Neural Information Processing Systems*, pp. 6115–6128, (2021).