Boosting Transferability of Targeted Adversarial Examples via Hierarchical Generative Networks

Xiao Yang¹ Yinpeng Dong^{1,2} Tianyu Pang³ Hang Su^{1,4} Jun Zhu^{1,2,4}*

Dept. of Comp. Sci. and Tech., Institute for AI, Tsinghua-Bosch Joint ML Center, THBI Lab, BNRist Center, Tsinghua University¹

RealAI² Sea AI Lab, Singapore³

Peng Cheng Laboratory; Pazhou Laboratory (Huangpu), Guangzhou, China⁴

Abstract. Transfer-based adversarial attacks can evaluate model robustness in the black-box setting. Several methods have demonstrated impressive untargeted transferability, however, it is still challenging to efficiently produce *targeted* transferability. To this end, we develop a simple yet effective framework to craft targeted transfer-based adversarial examples, applying a hierarchical generative network. In particular, we contribute to amortized designs that well adapt to multi-class targeted attacks. Extensive experiments on ImageNet show that our method improves the success rates of targeted black-box attacks by a significant margin over the existing methods — it reaches an average success rate of 29.1% against six diverse models based only on one substitute whitebox model, which significantly outperforms the state-of-the-art gradientbased attack methods. Moreover, the proposed method is also more efficient beyond an order of magnitude than gradient-based methods.

1 Introduction

Recent progress in adversarial machine learning demonstrates that deep neural networks (DNNs) are highly vulnerable to adversarial examples [13, 47], which are maliciously generated to mislead a model to produce incorrect predictions. It has been demonstrated that adversarial examples possess an intriguing property of transferability [50, 19, 4] — the adversarial examples crafted for a white-box model can also mislead other unknown models, making *black-box attacks* feasible. The threats of adversarial examples have raised severe concerns in numerous security-sensitive applications, such as autonomous driving [10] and face recognition [54, 55, 53].

Tremendous efforts have been made to develop more effective black-box attacking methods based on transferability, since they can serve as an important surrogate to evaluate the model robustness in real-world scenarios [32, 8]. The current methods have achieved impressive performance of untargeted black-box

^{*} corresponding author.

2 X.Yang et al.

attacks, intending to cause misclassification of the black-box models. However, *targeted* black-box attacks, aiming at misleading the black-box models by outputting the adversary-desired target class, perform unsatisfactorily or require computation scaling with number of classes [7, 57]. Technically, the inefficiency of targeted adversarial attacks could result in an over-estimation of model robustness under the challenging black-box attack setting [16].

Existing efforts on targeted black-box attacks can be categorized as *instance-specific* and *instance-agnostic* attacks. Specifically, the instance-specific attack methods [12, 34, 25, 8, 27] craft adversarial examples by performing gradient updates iteratively, which obtain unsatisfactory performance for targeted black-box attacks due to easy overfitting to a white-box model [8, 51]. Recently, [58] propose several improvements for instance-specific targeted attacks, thus we treat the method in [58] as one of the strong instance-specific baselines compared in our experiments.

On the other hand, the instance-agnostic attack methods learn a universal perturbation [57] or a universal function [43, 35] on the data distribution independent of specific instances. They can promote more general and transferable adversarial examples since the universal perturbation or function can alleviate the data-specific overfitting problem by training on an unlabeled dataset. CD-AP [35], as one of the effective instance-agnostic methods, adopts a generative model as a universal function to obtain an acceptable performance when facing one specified target class. However, CD-AP needs to learn a generative model for each target class while performing multi-target attack [14], i.e., crafting adversarial examples targeted at different classes. Thus it is not scalable to the increasing number of targets such as hundreds of classes, limiting practical efficiency.

To address the aforementioned issues and develop a targeted black-box attack in the practical scenario, in this paper we propose a conditional generative model as the universal adversarial function to craft adversarial perturbations. Thus we can craft adversarial perturbations targeted at different classes, using a single model backbone with different class embeddings. The proposed generative method is simple yet practical to obtain superior performance of targeted blackbox attacks, meanwhile with two technical improvements including (i) smooth projection mechanism that better helps the generator probe targeted semantic knowledge from the classifier; (ii) adaptive Gaussian smoothing with the focus of making generated results obtain adaptive ability against adversarially trained models. Therefore, our approach have several advantages over existing generative attacks [35, 38, 39], as described in the followings.

One model for multiple target classes. The previous generative methods [35, 38] require costly training N models while performing a multi-target attack with N classes. However, ours only trains one model and reaches an average success rate of 51.1% against six naturally trained models and 36.4% against three adversarially trained models based only on one substitute white-box model in ImageNet dataset, which outperforms CD-AP by a large margin of 6.0% and 31.3%, respectively. Hierarchical partition of classes. While handling plenty of classes (e.g., 1,000 classes in ImageNet), the effectiveness of generating targeted adversarial examples will be affected by a single generative model due to the difficulty of loss convergence in adversarial learning [52, 1]. Thus we train a feasible number of models (e.g., $10\sim20$ models on ImageNet) to further promote the effectiveness beyond the single model backbone. Specifically, each model is learned from a subset of classes specified by a designed hierarchical partition mechanism by considering the diversity property among subsets, for seeking a balance between effectiveness and scalability. It reaches an average success rate of 29.1% against six different models, outperforming the state-of-the-art gradient-based methods by a large margin, based only on one substitute white-box model. Moreover, the proposed method achieves substantial speedup over the mainstream gradient-based methods.

Strong semantic patterns. We experimentally find that these adversarial perturbations generated by the proposed Conditional Generative models can arise as a result of strong Semantic Pattern (C-GSP) as shown in Fig. 1(a). Furthermore, we present more valuable analyses in Sec. 4.6, illustrating that the generated semantic pattern itself achieves well-generalizing performance among the different models and is robust to the influence of data. These analyses are very instructive for the understanding and design of adversarial examples.

Technically, our main contributions can be summarized as:

- We propose a simple yet practical conditional generative targeted attack with a scalable hierarchical partition mechanism, which can generate targeted adversarial examples without tuning the parameters.
- Extensive experiments demonstrate that our method significantly improves the success rates of targeted black-box attacks over the existing methods.
- As a by-product, our baseline experiments provide a systematical evaluation on previous targeted black-box attacks, either instance-specific or instanceagnostic, on the ImageNet dataset with plenty of classes and face recognition.

2 Related Work

In this section, we review related work on adversarial attacks belonging to different types.

Instance-specific attacks. Some recent works [34, 58] adopt gradient-based optimization methods to generate the data-dependent perturbations. MIM [8] introduces the momentum term into the iterative attack process to improve the black-box transferability. DIM [51] and TI [9] aim to achieve the better transferability by input or gradient diversity. Recent works [20, 21] also attempt to costly train the multiple auxiliary classifiers to improve the black-box performance of iterative methods. In contrast, we improve the transferability performance over instance-specific methods simultaneously with the inference-time efficiency.

Instance-agnostic attacks. Different from instance-specific attacks, instanceagnostic attacks belong to image-independent (universal) methods. The first pipeline is to learn a universal perturbation. UAP [33] proposes to fool a model



Fig. 1: (a) shows the targeted adversarial examples crafted by MIM [8] and C-GSP given the target class *Viaduct* with the maximum perturbation $\epsilon = 16$. The predicted labels and probabilities are shown by another black-box model. (b) presents an overview of our proposed generative method for crafting C-GSP, including modules of conditional generator and classifier. The generator integrates the image and conditional class vector from Map network into a hidden incorporation. The generator is only trained in the whole pipeline to probe the target boundaries of the classifier.

by adding a learned universal noise vector. Another pipeline of attacks introduces learned generative models to craft adversarial examples. GAP [38] and AAA [39] craft adversarial perturbations based on target data directly and compress impressions, respectively. Previous methods, including universal perturbation and function, require costly training the same number of models for multiple target classes. Our method is capable of simultaneously generating adversarial samples for specifying multiple targets with better attack performance.

Multi-target attacks. Instance-specific attacks have the ability for specifying any target in the optimization phase. As elaborated in the introduction, these methods have degraded transferability and time-consuming iterative procedures. MAN [14] trains a generative model in the ImageNet under the constraint of ℓ_2 norm to explore the targeted attacks, which specifies all 1,000 categories from ImageNet for seeking extreme speed and storage. However, MAN does not fully compare multi-target black-box performance with previous instance-specific or instance-agnostic attacks, and the authors also claim that too many categories make it hard to transfer to another model. Recent approaches [57, 36] reveal better single-target transferability by learning universal perturbation or function, whereas they require to train multiple times while specifying multiple targets. As a comparison, our method can generate adversarial samples for specifying multiple targets, meanwhile generated strong semantic patterns can outperform existing attacks by a significant margin.

3 Method

In this section, we introduce a conditional generative model to learn a universal adversarial function, which can achieve effective multi-target black-box attacks. While handing plenty of classes, we design a hierarchical partition mechanism to make the generative model capable of specifying any target class under a feasible number of models, regarding both the effectiveness and scalability.

3.1 Problem Formulation

We use \boldsymbol{x}_s to denote an input image belonging to an unlabeled training set $\mathcal{X}_s \subset \mathbb{R}^d$, and use $c \in \mathcal{C}$ to denote a specific target class. Let $\mathcal{F}_{\phi} : \mathcal{X}_s \to \mathbb{R}^K$ denote a classification network that outputs a class probability vector with K classes. To craft a targeted adversarial example \boldsymbol{x}_s^* from a real example \boldsymbol{x}_s , the targeted attack aims to fool the classifier \mathcal{F}_{ϕ} by outputting a specific label c as arg $\max_{i \in \mathcal{C}} \mathcal{F}_{\phi}(\boldsymbol{x}_s^*)_i = c$, meanwhile the ℓ_{∞} norm of the adversarial perturbation is required to be no more than a threshold ϵ as $\|\boldsymbol{x}_s^* - \boldsymbol{x}_s\|_{\infty} \leq \epsilon$.

Although some generative methods [38, 35] can learn targeted adversarial perturbation, they do not take into account the effectiveness of multi-target generation, thus leading to inconvenience. To make the generative model learn how to specify multiple targets, we propose a conditional generative network \mathcal{G}_{θ} that effectively crafts multi-target adversarial perturbations by modeling class-conditional distribution. Different from previous single-target methods [35, 38], the target label c is regarded as a discrete variable rather than a constant. As illustrated in Fig. 1(b), our model contains a conditional generator \mathcal{G}_{θ} and a classification network \mathcal{F}_{ϕ} parameterized by θ and ϕ , respectively. The conditional generative model $\mathcal{G}_{\theta} : (\mathcal{X}_s, \mathcal{C}) \to \mathcal{P}$ learns a perturbation $\boldsymbol{\delta} = \mathcal{G}_{\theta}(\boldsymbol{x}_s, c) \in \mathcal{P} \subset \mathbb{R}^d$ on the training data. The output $\boldsymbol{\delta}$ of \mathcal{G}_{θ} is projected within the fixed ℓ_{∞} norm, thus generating the perturbed image $\boldsymbol{x}_s^* = \boldsymbol{x}_s + \boldsymbol{\delta}$.

Given a pretrained network \mathcal{F}_{ϕ} parameterized by ϕ , we propose to generate the targeted adversarial perturbations by solving

$$\min_{\theta} \mathbb{E}_{(\boldsymbol{x}_s \sim \mathcal{X}_s, c \sim \mathcal{C})} [\mathbb{CE} (\mathcal{F}_{\phi}(\mathcal{G}_{\theta}(\boldsymbol{x}_s, c) + \boldsymbol{x}_s), c)], \text{ s.t. } \|\mathcal{G}_{\theta}(\boldsymbol{x}_s, c)\|_{\infty} \leq \epsilon, \quad (1)$$

where \mathbb{CE} is the cross-entropy loss. By solving problem (1), we can obtain a targeted conditional generator by minimizing the loss of specific target class in the unlabeled training dataset. Note that we only optimize the parameter θ of the generator \mathcal{G}_{θ} using the training data \mathcal{X}_s , then the targeted adversarial example \boldsymbol{x}_t^* can be crafted by $\boldsymbol{x}_t^* = \boldsymbol{x}_t + \mathcal{G}_{\theta}(\boldsymbol{x}_t, c)$ for any given image \boldsymbol{x}_t in the test data \mathcal{X}_t , which only requires an inference for this targeted image \boldsymbol{x}_t .

We experimentally find that the objective (1) can enforce the transferability for the generated perturbation δ . A reasonable explanation is that δ can arise as a result of **strong** and **well-generalizing** semantic pattern inherent to the target class, which is robust to the influence of any training data. In Sec. 4.5, we illustrate and corroborate our claim by directly feeding scaled adversarial perturbations¹ from different methods into the classifier. Indeed, we find that our semantic pattern can be classified as the target class with a high degree of confidence while the perturbation from MIM [8] performs like the noise, meanwhile the scaled semantic pattern performs well transferability in different black-box models.

¹ The perturbation is linearly scaled from $[-\epsilon, \epsilon]$ to [0, 255].

3.2 Network Architecture

We now present the details of the conditional generative model for targeted attack, as illustrated in Fig. 1(b). Specifically, we design a mapping network to generate a target-specific vector in the implicit space of each target and train conditional generator \mathcal{G}_{θ} to reflect this vector by constantly misleading the classifier \mathcal{F}_{ϕ} .

Mapping network. Given the one-hot class encoding $\mathbb{1}_c \in \mathbb{R}^K$ from target class c, the mapping network aims to generate the targeted latent vector $\boldsymbol{w} = \mathcal{W}(\mathbb{1}_c)$, where $\boldsymbol{w} \in \mathbb{R}^M$ and $\mathcal{W}(\cdot)$ consists of a multi-layer perceptron (MLP) and a normalization layer, which can construct diverse targeted vectors \boldsymbol{w} for a given target class c. Thus \mathcal{W} is capable of learning effective targeted latent vectors by randomly sampling different classes $c \in \mathcal{C}$ in training phase.

Generator. Given an input image \boldsymbol{x}_s , the encoder first calculates the feature map $\boldsymbol{F} \in \mathbb{R}^{N \times H \times W}$, where N, H and W refer to the number of channels, height and width of the feature map, respectively. The target latent vector \boldsymbol{w} , derived from the mapping network \mathcal{W} by introducing a specific target class c, is expanded along height and width directions to obtain the label feature map $\boldsymbol{w}_s \in \mathbb{R}^{M \times H \times W}$. Then the above two feature maps are concatenated along the channels to obtain $\boldsymbol{F}' \in \mathbb{R}^{(N+M) \times H \times W}$. The obtained mixed feature map is then fed to the subsequent network. Therefore, our generator \mathcal{G}_{θ} translates an input image \boldsymbol{x}_s and latent target vector \boldsymbol{w} into an output image $\mathcal{G}_{\theta}(\boldsymbol{x}_s, \boldsymbol{w})$, which enables \mathcal{G}_{θ} to synthesize adversarial images of a series of targets. For the output of feature map $\boldsymbol{f} \in \mathbb{R}^d$ in the decoder, we adopt a **smooth projection** $P(\cdot)$ to perform a change of variables over \boldsymbol{f} rather than directly minimizing its ℓ_2 norm as [14] or clipping values outside the fixed norm [35], which can be denoted as

$$\boldsymbol{\delta} = P(\boldsymbol{f}) = \boldsymbol{\epsilon} \cdot \tanh(\boldsymbol{f}), \tag{2}$$

where ϵ is the strength of perturbation. Since $-1 \leq \tanh(\mathbf{f}) \leq 1$, δ can automatically satisfy the ℓ_{∞} -ball bound with perturbation budget ϵ . This transformation can be regarded as a better smoothing of gradient than directly clipping values outside the fixed norm, which is also instrumental for \mathcal{G}_{θ} to probe and learn the targeted semantic knowledge from \mathcal{F}_{ϕ} .

Training objectives. The training objectives seek to minimize the classification error on the perturbed image of the generator as

$$\theta^* \leftarrow \arg\min_{\theta} \mathbb{CE}\Big(F_{\phi}\big(\boldsymbol{x}_s + \mathcal{G}_{\theta}(\boldsymbol{x}_s, \mathcal{W}(\mathbb{1}_c))\big), c\Big),$$
(3)

which adopts an end-to-end training paradigm with the goal of generating adversarial images to mislead the classifier the target label, and \mathbb{CE} is the cross entropy loss. Previous studies attempt different classification losses in their works [57, 35], and we found that cross-entropy loss works well in our settings. The detailed optimization procedure is summarized in Algorithm 1.

Algorithm 1	Training	Algorithm	for the	Conditional	Generative	Attack
-------------	----------	-----------	---------	-------------	------------	--------

Require: Training Data \mathcal{D}_s ; a generative network \mathcal{G}_{θ} ; a classification network \mathcal{F}_{ϕ} ; a mapping network \mathcal{W} .

Ensure: Adversarial perturbations $\boldsymbol{\theta}$.

1: for iter in MaxIterations T do

- 2: Randomly sample B images $\{\boldsymbol{x}_{s_i}\}_{i=1}^B$;
- 3: Randomly sample B target classes $\{c_i\}_{i=1}^B$;
- 4: Forward pass c_i into W to compute the targeted latent vectors w_i ;
- 5: Obtain the perturbed images by $\boldsymbol{x}_{s_i}^* = \epsilon \cdot \tanh(\mathcal{G}(\boldsymbol{x}_{s_i}, \boldsymbol{w}_i)) + \boldsymbol{x}_{s_i};$
- 6: Forward pass $\boldsymbol{x}_{s_i}^*$ to \mathcal{F}_{ϕ} and compute loss in Eq. (3);
- 7: Backward pass and update the \mathcal{G}_{θ} ;

```
8: end for
```

3.3 Hierarchical Partition for Classes

While handling plenty of classes, the effectiveness of a conditional generative model will decrease as illustrated in Fig. 4, because the representative capacity is limited with a single generator. Therefore, we propose to divide all classes into a feasible number of subsets to train models when the class number K is large, e.g., 1,000 classes in ImageNet, with the aim of seeking the effectiveness of targeted black-box attack. To obtain a good partition, we introduce a representative target class space, which is nearly equivalent to the original class space C. Specifically, we utilize the weights $\phi_{cls} \in \mathbb{R}^{D \times K}$ in the classifier layer for the classification network \mathcal{F}_{ϕ} . Therefore, ϕ_{cls} can be regarded as the alternative class space since the weight vector $d_c \in \mathbb{R}^D$ from ϕ_{cls} can represent a class center of the feature embeddings of input images with same class c.

Note that once those subsets with closer metric distance (e.g., larger cosine similarity) in the target class space ϕ_{cls} are regarded as conditional inputs of generative network, they obtain worse loss convergence and transferability than diverse them due to mutual influence among these input conditions, as illustrated in Fig. 5. Thus we focus on selecting target classes that do not tend to overlap or be close to each other as accessible subsets. To capture more diverse examples in a given sampling space, we adopt K-determinantal point processes (DPP) [24, 23] to achieve a hierarchical partition, which can take advantage of the diversity property among subsets by assigning subset probabilities proportional to determinants of a kernel matrix.

First, we compute the RBF kernel matrix L of ϕ_{cls} and eigendecomposition of L, and a random subset V of the eigenvectors is chosen by regarding the eigenvalues as sampling probability. Second, we select a new class c_i to add to the set and update V in a manner that de-emphaseizes items similar to the one selected. Each successive point is selected and V is updated by Gram-Schmidt orthogonalization, and the distribution shifts to avoid points near those already chosen. By performing the above procedure, we can obtain a subset with k size. Thus while handling the conditional classes with K, we can hierarchically adopt this algorithm to get the final K/k subsets, which are regarded as conditional 8 X.Yang et al.

variables of generative models to craft adversarial examples. The details are presented in Appendix A.

4 Experiments

In this section, we present extensive experiments to demonstrate the effectiveness of proposed method for targeted black-box attacks².

4.1 Experimental Settings

Datasets. We consider the following datasets for training, including a widely used object detection dataset MS-COCO [30] and ImageNet training set [5]. We focus on standard and comprehensive testing settings, thus the inference is performed on ImageNet validation set (50k samples), a subset (5k) of ImageNet proposed by [28] and ImageNet-NeurIPS (1k) proposed by [37].

Networks. We consider some naturally trained networks, i.e., Inception-v3 (Inv3) [46], Inception-v4 (Inv4) [44], Resnet-v2-152 (R152) [15] and Inception-Resnet-v2 (IR-v2) [44], which are widely used for evaluating transferability. Besides, we supplement DenseNet-201 (DN) [17], GoogleNet (GN) [45] and VGG-16 (VGG) [42] to fully evaluate the transferability. Some adversarially trained networks [48] are also selected to evaluate the performance, i.e., ens3-adv-Inception-v3 (Inv3_{ens3}), ens4-adv-Inception-v3 (Inv3_{ens4}) and ens-adv-Inception-ResNet-v2 (IR-v2_{ens}).

Implementation details. As for instance-specific attacks, we compare our method with several attacks, including MIM [8], DIM [51], TI [9], SI [29] and the state-of-the-art targeted attack named Logit [58]. All instance-specific attacks adopt optimal hyperparameters provided in their original work. Specifically, the attack iterations M of MIM, DIM and Logit are set as 10, 20, 300, respectively. And $||W||_1 = 5$ is used for TI [9] as suggested by [11, 58]. We choose the same ResNet autoencoder architecture in [22, 35] as the basic generator networks, which consists of downsampling, residual and upsampling layers. We initialize the learning rate as 2e-5 and set the mini-batch size as 32. Smoothing mechanism is proposed to improve the transferability against adversarially trained models [9]. Instead of adopting smoothing for generated perturbation while the training is completed as CD-AP [35], we introduce adaptive Gaussian smoothing kernel to compute δ from Eq. (2) in the training phase, named adaptive Gaussian smoothing, with the focus of making generated results obtain adaptive ability. More implementation details and discussion with other networks (e.g., BigGAN [2]) are illustrated in Appendix B.

4.2 Transferability Evaluation

We consider 8 different target classes from [57] to form the multi-target black-box attack testing protocol with 8k times in 1k ImageNet NeurIPS set.

² Code at https://github.com/ShawnXYang/C-GSP.

Table 1: Transferability comparison for multi-target attacks on ImageNet NeurIPS validation set (1k images) with the perturbation budget of $\ell_{\infty} \leq 16$. The results are averaged on 8 different target classes. Note that CD-AP[†] indicates that training 8 **models** can obtain results, while our method only train **one** conditional generative model. * indicates white-box attacks.

	Mathad Time Mod		Model	Naturally Trained						Adversarially Trained			
	Method	(ms)	Number	Inv3	Inv4	IR-v2	R152	DN	$_{\rm GN}$	VGG-16	$Inv3_{ens3}$	$Inv3_{ens4}$	IR-v2 _{ens}
	MIM	~ 130	-	99.9^{*}	0.8	1.0	0.4	0.2	0.2	0.3	< 0.1	0.1	< 0.1
	TI-MIM	~ 130	-	99.9^{*}	0.9	1.1	0.4	0.4	0.3	0.5	0.1	0.2	0.1
	SI-MIM	~ 130	-	99.8^{*}	1.5	2.0	0.8	0.7	0.7	0.5	0.3	0.3	0.1
3	DIM	~ 260	-	95.6^{*}	4.0	4.8	1.3	1.9	0.8	1.3	0.1	0.2	0.1
Ň	TI-DIM	~ 260	-	96.0^{*}	4.4	5.1	1.4	2.4	1.1	1.8	0.3	0.4	0.2
д	SI-DIM	~ 260	-	98.4^{*}	5.6	5.9	2.8	3.0	2.3	1.6	0.9	0.9	0.3
	Logit	~ 3900	-	99.6^{*}	5.6	6.5	1.7	3.0	0.8	1.5	0.2	0.3	0.1
	$CD-AP^{\dagger}$	~ 15	8	94.2^{*}	57.6	60.1	37.1	41.6	32.3	41.7	1.5	2.2	1.2
	CD-AP-gs [†]	~ 15	8	69.7^{*}	31.3	30.8	18.6	20.1	14.8	20.2	5.0	5.8	4.5
	Ours	~ 15	1	93.4^{*}	66.9	66.6	41.6	46.4	40.0	45.0	39.7	37.2	32.2
	MIM	$\sim \! 185$	-	0.5	0.4	0.6	99.7^{*}	0.3	0.3	0.2	0.1	0.1	< 0.1
	TI-MIM	~ 185	-	0.3	0.3	1.0	96.5^{*}	0.5	0.3	0.3	0.3	0.2	0.3
	SI-MIM	~ 185	-	1.3	1.2	1.6	99.5^{*}	1.0	1.4	0.7	0.3	0.4	0.2
2	DIM	~ 370	-	2.8	3.1	5.0	93.6^{*}	3.5	1.7	1.3	0.4	0.4	0.3
15	TI-DIM	~ 370	-	4.3	4.1	5.8	92.9^{*}	4.3	2.1	1.4	0.8	0.7	0.4
н	SI-DIM	~ 370	-	7.2	8.4	10.4	97.4^{*}	7.6	6.4	2.6	0.8	0.7	1.3
	Logit	~ 5550	-	10.1	10.7	12.8	95.7^{*}	12.4	3.7	3.5	1.1	0.9	0.4
	$CD-AP^{\dagger}$	~ 10	8	33.3	43.7	42.7	96.6^{*}	53.8	36.6	34.1	15.7	15.2	12.0
	CD-AP-gs [†]	~ 10	8	7.8	11.3	10.0	53.6^{*}	20.4	8.7	12.5	4.9	6.4	6.2
	Ours	~ 10	1	37.7	47.6	45.1	93.2^{*}	64.2	41.7	45.9	31.6	32.0	29.9

Efficiency of multi-target black-box attack. Among comparable methods, instance-specific methods, i.e., MIM, DIM, and Logit, require iterative mechanism with M steps by computing gradients to obtain adversarial examples. Given the cost t_C^{FP} and t_C^{BP} of forward and backward passing the classifier, computing cost T^{IS} of single data can be defined as $T^{IS} = t_C^{FP} * M + t_C^{BP} * M$ in Table 1. Instance-agnostic methods only require the inference cost from the trained generator as $T^{IA} = t_G^{FP}$, thus possessing the priority for those attack scenarios within limited time. However, instance-agnostic methods require to train 8 models to obtain all predictions from 8 different classes. Due to timeconsuming training and more storage, we only reproduce an excellent generative method CD-AP [35] as a baseline, which already fully demonstrate the superior performance than other generative methods such as GAP [38] in their work. As a comparison, our conditional generative method only trains one model to inference the results and outperforms other methods w.r.t efficiency.

Effectiveness of multi-target black-box attack. Table 1 shows the transferability comparison of different methods on both naturally and adversarially trained models. The success rate of instance-specific attacks are very unsatisfactory, possibly explained by the data-point overfitting that makes it hard to transfer another model. The instance-agnostic attack CD-AP obtains acceptable performance, yet inferior to proposed method w.r.t black-box transferability. The **primary reason** for such a trend lies in some distinctions as 1) direct clip projection in CD-AP and our smooth projection in Eq. (2) and 2) their Gaussian





Fig. 2: Comparison of different projection functions and modes of Gaussian Smoothing. Results are reported with Inv3 network on ImageNet NeurIPS validation set.

Table 2: The untargeted fooling ratio (UT-FR) and targeted fooling ratio (T-FR) for adversarial attacks on ImageNet validation set (50k images) with the perturbation budget of $\ell_{\infty} \leq 10$. The attack is performed in same setting [57] with the target class 'sea lion' and the training dataset MS-COCO. * indicates white-box attacks.

	Mothod	VGC	G-16	VGC	G-19	R152		
	Method	UT-FR	T-FR	UT-FR	T-FR	UT-FR	T-FR	
VGG-16	UAE [57]	93.62^{*}	82.90^{*}	82.99	13.69	36.03	0.01	
	Ours	95.30^*	83.54^*	90.13	38.59	35.15	0.14	
VGG-19	UAE [57]	83.40	44.53	92.53^{*}	75.61^{*}	35.36	0.01	
	Ours	88.20	48.96	92.69^{*}	73.96^{*}	35.96	0.14	
R152	UAE [57]	55.05	1.63	55.12	1.05	82.58^{*}	70.20^{*}	
	Ours	83.90	29.81	83.24	24.81	91.14^{*}	80.47^*	

Smoothing and our adaptive Gaussian Smoothing, as described in Sec. 4.1 and Appendix B. Fig. 2 empirically shows the comparison results of single-target black-box attacks based on the CD-AP framework. Thus proposed conditional generative method can be a reliable baseline w.r.t targeted black-box attacks, regarding both *effectiveness* and *efficiency*.

Results of single-target black-box attack. Recent related works, e.g., UAE [57] and TTP [36] report excellent single-target black-box performances based on universal perturbations or functions. We obtain single-target degraded version of our model by specifying an input target label during the training process. The performance of black-box targeted attack between different methods is presented in Table 2. Besides, we also make some analyses about TTP and present compared results in Appendix C. Furthermore, some other instanceagnostic adversarial methods, e.g., UAP [33], GAP [38] and RHP [28], have tendency towards the untargeted black-box problem. Despite this, we also follow the corresponding untargeted setting and compare these methods in Appendix C. Our method is steadily improved under black-box targeted and untargeted black-box manner.

Targeted Black-box Attack in NeurIPS 2017 Competition (1,000 target classes) VGG-16 Method Inv4 IRv2 R152 DN GN MIM 0.1< 0.1< 0.10.30.1< 0.1TI-MIM 0.30.3< 0.10.4< 0.10.1SI-MIM 0.60.60.10.40.30.1DIM 0.22.92.50.61.20.6TI-DIM 2.92.50.51.70.31.0SI-DIM 4.34.11.71.91.81.1Logit 2.41.22.40.40.84.7

Table 3: Transferability comparison with the perturbation budget of $\ell_{\infty} \leq 16$. Whitebox substitute model is Inv3 for all attacks, following the standard protocol [9] with **1,000 stochastic target classes**.

Table 4: The success rate of black-box *impersonation* attacks on face verification with the perturbation budget of $\ell_{\infty} \leq 16$. ArcFace is chosen as white-box model.

25.0

22.9

26.8

26.6

Black-box Impersonation Attack in Face Recognition									
Protocol	Method	FaceNet	CosFace	SphereFace	MobileFace				
	MIM	34.4	16.6	22.4	35.0				
Ι	DIM	38.8	21.2	27.4	44.3				
	Ours	65.2	56.2	52.2	83.5				
	MIM	31.3	13.6	21.1	22.3				
II	DIM	36.1	16.4	24.4	31.9				
	Ours	66.8	49.1	47.9	67.8				

4.3 Effectiveness on NeurIPS 2017 Competition

37.4

35.9

Ours

To illustrate the effectiveness of our proposed attack methods in practical 1.000 classification, we here follow the official setting from NeurIPS 2017 adversarial competition [26] for testing targeted black-box transferability. Considering limited resource, previous instance-agnostic attacks are not required as comparable methods due to training 1,000 models, thus we focus on various excellent instance-specific attacks for comparison, including official top attack methods in NeurIPS 2017 adversarial competition. Compared with other instance-agnostic attacks, our hierarchical partition mechanism can make conditional generative networks be capable of specifying any target class via a feasible number of models for the scalability. Specifically, we consider 20 models, with each specifying 50 diverse classes from k-DPP hierarchical partition in this setting, to implement targeted attack by only once inference for each target image. As shown in Table 3, our method can obviously outperform all other methods. In addition, these trained generative models can directly be applied to craft adversarial examples, which is more convenient and efficient when required to handle large-scale (e.g., millions of images) datasets than instance-specific attacks.



Fig. 3: Generative examples of adversarial images with perturbation budget of $\ell_{\infty} \leq 16$. We separately adopt the ImageNet and MS-COCO dataset as the training dataset to implement the generation of targeted perturbations. Our method can generate semantic pattern independent of training dataset.



Fig. 4: Asr vs. *numbers of conditional tar*- Fig. 5: Comparison of loss convergence *gets* curve against Inv3 and VGG-16 mod- and transferability between diverse and els. close conditional subset.

4.4 Effectiveness on Realistic Face Recognition

Adversarial perturbations added to original face images have ability to evade being recognized or impersonate another individual [41, 56]. In this section, we consider the transferability of impersonation attack to further illustrate the generalization of our method, which is also corresponding to targeted attack in the image classification task.

Dataset and models. We conduct the experiments on Labeled Faces in the Wild (LFW) [18] and introduce two test protocols. For *Protocol I* defined as single-target impersonation attack, we choose 1 target identity and 1k source face images belonging with different identities from LFW as the attackers, thus forming 1k pairs. For *Protocol II* named multi-target impersonation attack, 5 target identities and 1k source face images are selected to form 1k attack pairs, meaning that we need to implement 5k attacks. We involve some excellent face recognition models for conducting black-box testing, including Sphereface [31], CosFace [49], FaceNet [40] and MobileFace [3]. These models lie in different



Fig. 6: We show the adversarial examples and Fig. 7: Plots of logit vectors from the adextracted perturbation scaled in image-pixel versarial image L_{img} and scaled crafted space in the second column and the third col- perturbation L_{adv} of MIM and proposed umn. The predictive confidence is presented generative method, with their respective by **directly feeding** extracted perturbation PCC values. into the classifier in the last column.

model architectures and training objectives. In all experiments, we only use one model ArcFace [6] as substitute model to craft adversarial samples, and test attack performance against other unknown models.

Evaluation metrics. We first compute the optimal threshold of every face recognition models from LFW dataset by following standard protocols. If the similarity of a pair of images exceeds the threshold, we regard them as same identity, otherwise different identities.

Black-box attack results. We adjust the optimization object function to adapt face recognition for chosen attack methods (detailed in Appendix C), and report the success rate of black-box *impersonation* attacks in Table 4, which illustrates that our method can achieve nearly two times of the success rates than DIM in *Protocol I* and *Protocol II*. The results indicate that our method is superior to other methods not only in image classification.

4.5 Comparison Study about Target Classes

We conduct an extensive study to investigate two key points about target classes.

Different numbers of target classes. We conduct effectiveness for different numbers of target classes in Fig. 4. It can be seen that the results perform well within a feasible number of targets, whereas to a certain extent effectiveness tend to decay. Therefore, the effectiveness of conditional generative networks is influenced by the number of conditional classes, due to the representative capacity of single generator. We aim to divide all classes into a feasible number of set while handling plenty of classes.

Comparison of different multi-target conditions. We select closer conditional classes with larger cosine similarity in the target class space ϕ_{cls} and diverse conditional classes from k-DPP method. In Fig. 5, closer conditional classes have worse loss convergence and transferability than diverse them due to mutual influence among conditions.

14 X.Yang et al.

4.6 More Analyses

Targeted adversarial samples from proposed generative method can produce semantic pattern inherent to the target class in Fig. 3. Why does generative semantic pattern work?

First, generative methods can produce strong targeted semantic pattern that is robust to the influence of data, which is obtained by minimizing the loss of specific target class in the training phase. To corroborate our claim, we directly feed scaled crafted perturbations by instance-specific attack MIM and our generative method into the classifier. Indeed, we find that our generative perturbation is considered as target class with a high degree of confidence whereas the perturbation from MIM performs like the noise, as shown in Fig. 6. Furthermore, we plot the logit relationship by computing PCC (Pearson correlation coefficient) values from scaled crafted perturbation and adversarial image in Fig. 7. The numerical performance is also consistent with our mentioned claim.

Second, the generated adversarial semantic pattern achieves well-generalizing performance among the different models. We feed 1k images from ImageNet test set into the generator trained by Inv3 model to obtain 1k semantic patterns, which are scaled to image pixel space and then fed into different classifiers. We compute the mean confidence of **0.46** for DN, **0.44** for Inv4, and **0.35** for R152, whereas the perturbation from MIM is lower than 0.01. The results show that our scaled semantic pattern can directly achieve well-generalizing performance among models, possibly explained by utilizing similar feature knowledge from the same class on different classifiers trained on same training data distribution. Thus similar pattern can be instrumental for transferability among models.

5 Discussion and Conclusion

Transferability of targeted black-box attack is simultaneously affected by data and model. Therefore, instance-specific methods easily overfit the data point and white-box model, resulting in weak transferability. As a comparison, the proposed generative method with powerful learning capacity reduces the dependency for data point by adopting the unlabeled training data, thus enabling the model to learn semantic pattern and improve the transferability of targeted black-box attack. Extensive experiments demonstrate that proposed generative method can significantly improve the success rates of targeted black-box attacks against various models, meanwhile achieving faster speedup beyond an order of magnitude than gradient-based methods. Therefore, this method can be regarded as a new baseline method in terms of targeted black-box attacks, which provides a novel framework to explore the vulnerabilities of DNNs.

Acknowledgement. This work was supported by the National Key Research and Development Program of China (Nos. 2020AAA0104304, 2017YFA0700904), NSFC Projects (Nos. 62061136001, 61621136008, 62076147, U19B2034, U19A2081, U1811461), the major key project of PCL (No. PCL2021A12), Tsinghua-Alibaba Joint Research Program, Tsinghua-OPPO Joint Research Center, and the High Performance Computing Center, Tsinghua University.

References

- 1. Berthelot, D., Schumm, T., Metz, L.: Began: Boundary equilibrium generative adversarial networks. arXiv preprint arXiv:1703.10717 (2017)
- Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. arXiv preprint arXiv:1809.11096 (2018)
- Chen, S., Liu, Y., Gao, X., Han, Z.: Mobilefacenets: Efficient cnns for accurate real-time face verification on mobile devices. In: Chinese Conference on Biometric Recognition. pp. 428–438. Springer (2018)
- Demontis, A., Melis, M., Pintor, M., Jagielski, M., Biggio, B., Oprea, A., Nita-Rotaru, C., Roli, F.: Why do adversarial attacks transfer? explaining transferability of evasion and poisoning attacks. In: 28th {USENIX} Security Symposium ({USENIX} Security 19). pp. 321–338 (2019)
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A largescale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
- Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4690–4699 (2019)
- Dong, Y., Fu, Q.A., Yang, X., Pang, T., Su, H., Xiao, Z., Zhu, J.: Benchmarking adversarial robustness. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
- Dong, Y., Liao, F., Pang, T., Su, H., Zhu, J., Hu, X., Li, J.: Boosting adversarial attacks with momentum. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- Dong, Y., Pang, T., Su, H., Zhu, J.: Evading defenses to transferable adversarial examples by translation-invariant attacks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- Eykholt, K., Evtimov, I., Fernandes, E., Li, B., Rahmati, A., Xiao, C., Prakash, A., Kohno, T., Song, D.: Robust physical-world attacks on deep learning visual classification. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 1625–1634 (2018)
- Gao, L., Zhang, Q., Song, J., Liu, X., Shen, H.T.: Patch-wise attack for fooling deep neural network. In: European Conference on Computer Vision. pp. 307–322. Springer (2020)
- 12. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press (2016), http: //www.deeplearningbook.org
- 13. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: International Conference on Learning Representations (ICLR) (2015)
- Han, J., Dong, X., Zhang, R., Chen, D., Zhang, W., Yu, N., Luo, P., Wang, X.: Once a man: Towards multi-target attack via learning multi-target adversarial network once. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 5158–5167 (2019)
- 15. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: European Conference on Computer Vision (ECCV). pp. 630–645. Springer (2016)
- Hendrycks, D., Carlini, N., Schulman, J., Steinhardt, J.: Unsolved problems in ml safety. arXiv preprint arXiv:2109.13916 (2021)
- Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4700–4708 (2017)

- 16 X.Yang et al.
- Huang, G.B., Mattar, M., Berg, T., Learned-Miller, E.: Labeled faces in the wild: A database forstudying face recognition in unconstrained environments. In: Technical report (2007)
- Huang, Q., Katsman, I., He, H., Gu, Z., Belongie, S., Lim, S.N.: Enhancing adversarial example transferability with an intermediate level attack. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4733–4742 (2019)
- Inkawhich, N., Liang, K., Wang, B., Inkawhich, M., Carin, L., Chen, Y.: Perturbing across the feature hierarchy to improve standard and strict blackbox attack transferability. Advances in Neural Information Processing Systems 33, 20791–20801 (2020)
- Inkawhich, N., Liang, K.J., Carin, L., Chen, Y.: Transferable perturbations of deep feature distributions. arXiv preprint arXiv:2004.12519 (2020)
- Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: European conference on computer vision. pp. 694–711. Springer (2016)
- Kulesza, A., Taskar, B.: k-dpps: Fixed-size determinantal point processes. In: ICML (2011)
- Kulesza, A., Taskar, B.: Determinantal point processes for machine learning. arXiv preprint arXiv:1207.6083 (2012)
- Kurakin, A., Goodfellow, I., Bengio, S.: Adversarial examples in the physical world. In: International Conference on Learning Representations (ICLR) Workshops (2017)
- Kurakin, A., Goodfellow, I., Bengio, S., Dong, Y., Liao, F., Liang, M., Pang, T., Zhu, J., Hu, X., Xie, C., et al.: Adversarial attacks and defences competition. In: The NIPS'17 Competition: Building Intelligent Systems, pp. 195–231. Springer (2018)
- Li, M., Deng, C., Li, T., Yan, J., Gao, X., Huang, H.: Towards transferable targeted attack. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 641–649 (2020)
- Li, Y., Bai, S., Xie, C., Liao, Z., Shen, X., Yuille, A.L.: Regional homogeneity: Towards learning transferable universal adversarial perturbations against defenses. arXiv preprint arXiv:1904.00979 (2019)
- Lin, J., Song, C., He, K., Wang, L., Hopcroft, J.E.: Nesterov accelerated gradient and scale invariance for adversarial attacks. In: International Conference on Learning Representations (2019)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
- Liu, W., Wen, Y., Yu, Z., Li, M., Raj, B., Song, L.: Sphereface: Deep hypersphere embedding for face recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 212–220 (2017)
- 32. Liu, Y., Chen, X., Liu, C., Song, D.: Delving into transferable adversarial examples and black-box attacks. In: ICLR (2017)
- Moosavi-Dezfooli, S.M., Fawzi, A., Fawzi, O., Frossard, P.: Universal adversarial perturbations. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1765–1773 (2017)
- Moosavi-Dezfooli, S.M., Fawzi, A., Frossard, P.: Deepfool: a simple and accurate method to fool deep neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)

- Naseer, M.M., Khan, S.H., Khan, M.H., Khan, F.S., Porikli, F.: Cross-domain transferability of adversarial perturbations. In: Advances in Neural Information Processing Systems. pp. 12905–12915 (2019)
- Naseer, M., Khan, S., Hayat, M., Khan, F.S., Porikli, F.: On generating transferable targeted perturbations. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7708–7717 (2021)
- 37. NeurIPS: https://www.kaggle.com/c/nips-2017-defense-against-\ \adversarial-attack/data. Kaggle, 2017
- Poursaeed, O., Katsman, I., Gao, B., Belongie, S.: Generative adversarial perturbations. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4422–4431 (2018)
- Reddy Mopuri, K., Krishna Uppala, P., Venkatesh Babu, R.: Ask, acquire, and attack: Data-free uap generation using class impressions. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 19–34 (2018)
- Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 815–823 (2015)
- 41. Sharif, M., Bhagavatula, S., Bauer, L., Reiter, M.K.: Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 1528–1540 (2016)
- 42. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
- Song, Y., Shu, R., Kushman, N., Ermon, S.: Constructing unrestricted adversarial examples with generative models. In: Advances in Neural Information Processing Systems (NeurIPS) (2018)
- 44. Szegedy, C., Ioffe, S., Vanhoucke, V., Alemi, A.: Inception-v4, inception-resnet and the impact of residual connections on learning. In: AAAI (2017)
- 45. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1–9 (2015)
- 46. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
- 47. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R.: Intriguing properties of neural networks. In: International Conference on Learning Representations (ICLR) (2014)
- Tramèr, F., Kurakin, A., Papernot, N., Boneh, D., McDaniel, P.: Ensemble adversarial training: Attacks and defenses. In: International Conference on Learning Representations (ICLR) (2018)
- 49. Wang, H., Wang, Y., Zhou, Z., Ji, X., Gong, D., Zhou, J., Li, Z., Liu, W.: Cosface: Large margin cosine loss for deep face recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5265–5274 (2018)
- Wu, D., Wang, Y., Xia, S.T., Bailey, J., Ma, X.: Skip connections matter: On the transferability of adversarial examples generated with resnets. arXiv preprint arXiv:2002.05990 (2020)
- 51. Xie, C., Zhang, Z., Zhou, Y., Bai, S., Wang, J., Ren, Z., Yuille, A.L.: Improving transferability of adversarial examples with input diversity. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
- 52. Xu, K., Li, C., Zhu, J., Zhang, B.: Understanding and stabilizing gans' training dynamics with control theory. arXiv preprint arXiv:1909.13188 (2019)

- 18 X.Yang et al.
- Yang, X., Dong, Y., Pang, T., Xiao, Z., Su, H., Zhu, J.: Controllable evaluation and generation of physical adversarial patch on face recognition. arXiv e-prints pp. arXiv-2203 (2022)
- 54. Yang, X., Dong, Y., Pang, T., Zhu, J., Su, H.: Towards privacy protection by generating adversarial identity masks. arXiv preprint arXiv:2003.06814 (2020)
- Yang, X., Wei, F., Zhang, H., Zhu, J.: Design and interpretation of universal adversarial patches in face detection. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII 16. pp. 174–191. Springer (2020)
- Yang, X., Yang, D., Dong, Y., Yu, W., Su, H., Zhu, J.: Delving into the adversarial robustness on face recognition. arXiv preprint arXiv:2007.04118 (2020)
- Zhang, C., Benz, P., Imtiaz, T., Kweon, I.S.: Understanding adversarial examples from the mutual influence of images and perturbations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14521– 14530 (2020)
- Zhao, Z., Liu, Z., Larson, M.: On success and simplicity: A second look at transferable targeted attacks. Advances in Neural Information Processing Systems 34 (2021)