BEYOND IMAGENET ATTACK: TOWARDS CRAFTING ADVERSARIAL EXAMPLES FOR BLACK-BOX DOMAINS

Anonymous authors
Paper under double-blind review

ABSTRACT

Adversarial examples have posed a severe threat to deep neural networks due to their transferability nature. Currently, various works have paid great efforts to enhance the cross-model transferability, which mostly assume the substitute model is trained in the same domain as the target model. However, in real life, the relevant information of the deployed model is unlikely to leak. Hence, it is vital to build a more practical black-box threat model to overcome this limitation and evaluate the vulnerability of deployed models. In this paper, with the only knowledge of the ImageNet domain, we propose a Beyond ImageNet Attack (BIA) to investigate the transferability towards black-box domains (unknown classification tasks). Specifically, we leverage a generative model to learn the adversarial function for disrupting low-level features of input images. Based on this framework, we further propose two variants to narrow the gap between the source and target domains from the data and model perspectives, respectively. Extensive experiments on coarse-grained and fine-grained domains demonstrate the effectiveness of our proposed methods. Notably, our methods outperform state-of-the-art approaches by up to 7.71% (towards the coarse-grained domain) and 25.91% (towards the fine-grained domain) on average. Our code is available in the supplementary materials.

1 INTRODUCTION

Deep neural networks (DNNs) have achieved remarkable success in the image classification task in recent years. Nonetheless, advances in the field of adversarial machine learning (Szegedy et al., 2014; Goodfellow et al., 2015) make DNNs no longer reliable. By adding a well-designed perturbation on a benign image (*a.k.a* adversarial attack), the resulting adversarial examples can easily fool state-of-the-art DNNs. To make the matter worse, the adversarial attack technique can even be applied in the physical world (Sharif et al., 2016; Kurakin et al., 2017a; Xu et al., 2020; Duan et al., 2021), which inevitably raises concerns about the stability of deployed models. Therefore, exposing as many "blind spots" of DNNs as possible to the spotlight is a top priority.

Generally, deployed models are mainly challenged with two threat models: white-box and black-box. For white-box threat model (Kurakin et al., 2017b; Moosavi-Dezfooli et al., 2016; Carlini & Wagner, 2017; Shi et al., 2019), the attacker can obtain complete knowledge of the target model, such as the gradient for any input. However, deployed models are usually transparent to unauthorized users. In this scenario, prior black-box works (Poursaeed et al., 2018; Dong et al., 2018; Xie et al., 2019; Inkawhich et al., 2020; Wang et al., 2021) mostly assume that the source data for training the target model is available and mainly explore the cross-model transferability among models trained in the same data distribution. Specifically, perturbations are crafted via accessible white-box model (*a.k.a* substitute model), and resulting adversarial examples sometimes can fool other black-box models as well. Yet, these works still ignore a pivotal issue: *A model owner is unlikely to leak the relevant information of the deployed model.* To overcome this limitation, query-based black-box attacks (Papernot et al., 2016; Brendel et al., 2018; Chen et al., 2020; Li et al., 2021) are proposed, which adjust adversarial examples just according to the output of the target model. However, the resource-intensive query budget is extremely costly and inevitably alerts the model owner.

Therefore, we need a more "practical" black-box threat model to address this concern, *i.e.*, without any clue about the training data distribution as well as the pre-trained model based on it, and even

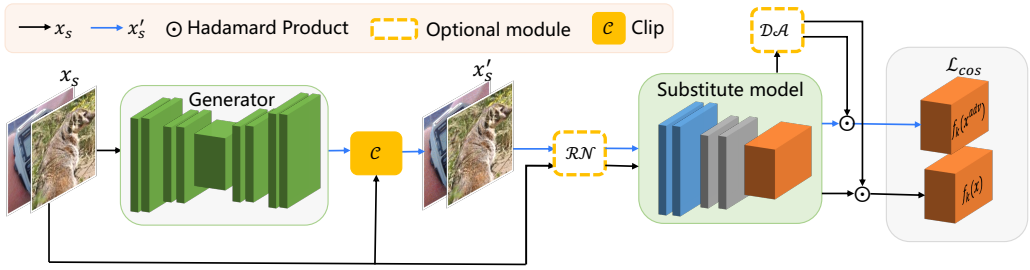

Figure 1: Our proposed generator framework aims to enlarge the cosine similarity of feature between benign image $\boldsymbol{x_s}$ and adversarial example $\boldsymbol{x'_s}$ during the training phase. Training data and substitute model are all from the ImageNet domain. $\mathcal{C}$ module is applied to constrain $\boldsymbol{x'_s}$ in the ℓ_∞-ball of $\boldsymbol{x_s}$. $\mathcal{RN}$ and $\mathcal{DA}$ are optional, which can further improve the transferability.

the query is forbidden. Intuitively, this threat model is more challenging to build for attackers but more threatening to model owners. To the best of our knowledge, a recent work called CDA (Naseer et al., 2019) is the first to attempt such an attack. Specifically, it learns a transferable adversarial function via the generator network against a different domain (training data and pre-trained model are all from ChestX-ray (Wang et al., 2017) domain). During inference, it directly crafts adversarial examples for benign ImageNet images to fool target ImageNet pre-trained models. However, its cross-domain transfer strength is still moderate. Besides, relying on small-scale datasets to train a generator may limit the generalization of the threat model.

Considering that ImageNet is a large-scale dataset containing most of common categories in real life and there are various off-the-shelf pre-trained models, one can easily dig out much useful information to build a strong threat model. Therefore, in this paper, solely relying on the knowledge of the ImageNet domain, we introduce an effective **Beyond ImageNet Attack (BIA)** framework to enhance the cross-domain transferability of adversarial examples. To reflect the applicability of our approach, we consider eight different image classification tasks (listed in Table 1). Figure 1 illustrates an overview of our method. Particularly, we learn a flexible generator network $\mathcal{G}_\theta$ against ImageNet domain. Instead of optimizing the domain-specific loss function like CDA, our method focuses on disrupting low-level features. Furthermore, we propose two variants based on the vanilla BIA to narrow the gap between source and target domains. Specifically, from the data perspective, we propose a random normalization ($\mathcal{RN}$) module to simulate different data distributions; from the model perspective, we propose a domain-agnostic attention ($\mathcal{DA}$) module to capture essential features for perturbing. In the inference phase, our $\mathcal{G}_\theta$ accepts images of any domain as the input and crafts adversarial examples with one forward propagation. Extensive experiments demonstrate the effectiveness of our proposed methods. Towards the coarse-grained and fine-grained domains, we can outperform state-of-the-art approaches by up to 7.71% and 25.91% on average, respectively. Besides, our methods can also enhance the cross-model transferability in the source domain.

Table 1: A comparison of datasets from different domains.

Dataset	Resolution	Type	Test size	Classes
ImageNet (Russakovsky et al., 2015)	224×224	-	50,000	1,000
CIFAR-10 (Krizhevsky, 2009)	32×32	coarse-grained	10,000	10
CIFAR-100 (Krizhevsky, 2009)	32×32	coarse-grained	10,000	100
STL-10 (Coates et al., 2011)	96×96	coarse-grained	8,000	10
SVHN (Netzer et al., 2011)	32×32	coarse-grained	26,032	10
CUB-200-2011 (Wah et al., 2011)	448×448	fine-grained	5,740	200
Stanford Cars (Krause et al., 2013)	448×448	fine-grained	8,041	196
FGVC Aircraft (Maji et al., 2013)	448×448	fine-grained	3,333	100

2 Related Works

Iterative Optimization Approaches. Under the black-box threat model, iterative attack methods are a popular branch, which usually adopt domain-specific loss or intermediate feature loss to craft

Figure 2: Benign images sampled from each domain. From the top to the bottom rows are the first category, the middle category and the last category of their label space, respectively.

adversarial examples. For the former, Madry et al. (2018) extend Goodfellow et al. (2015) to perform projected gradient descent from randomly chosen starting points inside ϵ-ball. Dong et al. (2018) introduce momentum term to stable the update direction. Xie et al. (2019) apply random transformations of the input at each iteration, thus mitigating overfitting. Gao et al. (2020) propose patch-wise perturbation to better cover the discriminative region. Wu et al. (2020) explore the security weakness of skip connections (He et al., 2016; Huang et al., 2017) to boost adversarial attacks.

Different from the methods mentioned above, intermediate feature-based methods focus on disrupting low-level features. For example, Inkawhich et al. (2019) make the source image close to the target image in feature space. Gao et al. (2021) attempt to measure similarities between feature maps by high-order statistics. Lu et al. (2020) propose a dispersion reduction attack to make the low-level features featureless. Naseer et al. (2020) design a self-supervised perturbation mechanism for enabling a transferable defense approach.

Generator-oriented Approaches. Compared with iterative optimization approaches, generator-oriented attacks are more efficient (*i.e.*, only need one inference) to generate adversarial examples. In this branch, Baluja & Fischer (2017) propose an adversarial transformation network to modify the output of the classifier given the original input. Poursaeed et al. (2018) present trainable deep neural networks for producing both image-agnostic and image-dependent perturbations. Naseer et al. (2019) leverage datasets from other domain instead of ImageNet to train generator networks against pre-trained ImageNet models, and inference is performed on ImageNet domain with the aim of fooling black-box ImageNet model. They also attempt a practical black-box threat model (from ChestX-ray to ImageNet), and the attack success rate can outperform the result of Gaussian noise.

3 Transferable Adversarial Examples beyond ImageNet

3.1 Problem Formulation

Given a target deep learning classifier $f_t(\cdot)$ trained in a specific data distribution χ_t, we aim to craft a human-imperceptible perturbation for the benign image $\boldsymbol{x_t} \sim \chi_t$ from the target domain with the only available knowledge of source ImageNet domain (including pre-trained model $f_s(\cdot)$ and data distribution χ_s). Formally, suppose we have a threat model $\mathcal{M}_{\theta^*}$ whose parameter θ^* is solely derived from the source domain, our goal is to craft adversarial examples for $\boldsymbol{x_t}$ from target domain so that they can fool the $f_t(\cdot)$ successfully:

$$f_t(\mathcal{M}_{\theta^*}(\boldsymbol{x_t})) \neq f_t(\boldsymbol{x_t}) \quad s.t. \ ||\mathcal{M}_{\theta^*}(\boldsymbol{x_t}) - \boldsymbol{x_t}||_\infty \leq \epsilon, \tag{1}$$

where ϵ is the maximum perturbation to ensure $\boldsymbol{x_t}$ is minimally changed. Intuitively, crafting adversarial examples for the black-box domain is very challenging. As shown in Table 1 and Figure 2, images from different domain vary greatly.

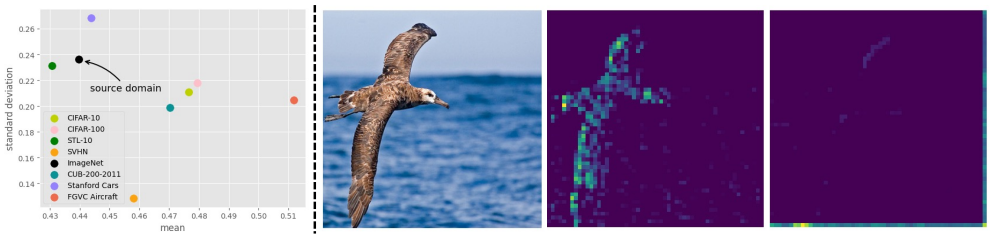

Figure 3: **Left:** The data distribution (*i.e.*, mean and standard deviation) for datasets from different domains. The result is the average over the three channels. **Right:** Two intermediate feature maps (*Maxpool.3*) of VGG-16 (Simonyan & Zisserman, 2015) (trained in ImageNet domain) for the input image from CUB-200-2011 (Wah et al., 2011).

3.2 Preliminary

Iterative/Single-step optimization methods (Goodfellow et al., 2015; Madry et al., 2018; Mao et al., 2021; Li et al., 2021) and generator-oriented methods (Baluja & Fischer, 2017; Poursaeed et al., 2018; Naseer et al., 2019; Li et al., 2020) are two popular branches for building the threat model. Since the attacker has the large-scale ImageNet training set at hand, there is no reason not to take full advantage of them. Therefore, in this paper, we adopt the generator-oriented framework which learns a transferable adversarial function via a generative model $\mathcal{G}_\theta$.

Given that the threat model aims at crafting transferable adversarial examples for unseen black-box domains, relying on the last layer with domain-specific loss functions (*e.g.*, relativistic cross-entropy loss adopted by Naseer et al. (2019)) is less effective because this might lead to overfitting in the source domain. In contrast, the intermediate layers of the DNN presumably extract general features (Yosinski et al., 2014) which may share across different models. Hence, our proposed Beyond ImageNet Attack (BIA) framework turns to destroy the low-level features of the substitute model at a specific layer L as much as possible. As illustrated in Figure 1, $\mathcal{G}_\theta$ is learned to decrease the cosine similarity between adversarial example $\boldsymbol{x_s'}$ and benign image $\boldsymbol{x_s} \in \mathbb{R}^{N \times H_s \times W_s}$ (sampled from χ_s) to make the feature featureless:

$$\theta^* = \arg\min_\theta \mathcal{L}_{cos}(f_s^L(\boldsymbol{x_s'}), f_s^L(\boldsymbol{x_s})). \tag{2}$$

In the inference phase, our generator $\mathcal{G}_{\theta^*}$ can directly craft adversarial examples for input images $\boldsymbol{x_t} \in \mathbb{R}^{N \times H_t \times W_t}$ from the target domain:

$$\boldsymbol{x_t'} = \min(\boldsymbol{x_t} + \boldsymbol{\epsilon}, \max(\mathcal{G}_{\theta^*}(\boldsymbol{x_t}), \boldsymbol{x_t} - \boldsymbol{\epsilon}). \tag{3}$$

The resulting adversarial examples $\boldsymbol{x_t'}$ are depicted in Figure 12 of Appendix. Compared with CDA, our BIA is more effective in both source (white-box) and target (black-box) domains. Yet, as shown in Figure 3, crafting more transferable adversarial examples still has some challenges:

- **Data Perspective:** The distribution (*i.e.*, mean and standard deviation) of source domain is largely different from the target domain. For example, the standard deviation of ImageNet is about twice that of SVHN.
- **Model Perspective:** Although some feature map of $f_s^L(\cdot)$ can capture the object ($\in \chi_t$) for feature representation (*e.g.*, the first feature map in Figure 3), there are also some feature maps that are significantly biased (*e.g.*, the second feature map in Figure 3).

To alleviate the concern of generating poor transferable adversarial examples that may arise from the above limitations, we propose two variants, equipped with random normalization ($\mathcal{RN}$) module or domain-agnostic attention ($\mathcal{DA}$) module, respectively.

3.3 Random Normalization Module

Generally, DNNs (Simonyan & Zisserman, 2015; He et al., 2016; Huang et al., 2017) are usually equipped with normalization for input images[1] so that they can subject to the standard normal dis-

[1]For Pytorch pre-trained ImageNet model, input images should be normalized using mean = [0.485, 0.456, 0.406] and standard deviation = [0.229, 0.224, 0.225]

tribution. However, as illustrated in Figure 3 (left), the distribution of dataset from the different domain can vary dramatically. Thus, training against a specific domain may limit the generalization of the resulting $\mathcal{G}_{\theta^*}$. Besides, the commonly used strategy of label-preserving data augmentation (Krizhevsky et al., 2012; Simonyan & Zisserman, 2015) is less effective because it has little effect on changing the distribution of the inputs (the experiment result is shown in Appendix A.4).

To overcome this limitation, fusing knowledge from different domains might be helpful. As shown in Table 3 of Naseer et al. (2019), using training data from other domains against the pre-trained model in the source domain usually enhances the transferability of adversarial examples towards the target domain. However, this setup is not feasible because the available knowledge for a more practical threat model may be limited, *i.e.*, restricted to one domain. Therefore, we instead propose a random normalization ($\mathcal{RN}$) module to simulate different data distribution in the training phase:

$$\mathcal{RN}(\boldsymbol{x_s}) = \boldsymbol{\sigma} \cdot \frac{\boldsymbol{x_s} - \mu'}{\sigma'} + \boldsymbol{\mu}, \tag{4}$$

where $\boldsymbol{\sigma}$ and $\boldsymbol{\mu}$ are default standard deviation and mean vectors for ImageNet, and $\mu' \sim \mathcal{N}(\mu'_{mean}, \mu'_{std})$ and $\sigma' \sim \mathcal{N}(\sigma'_{mean}, \sigma'_{std})$ are two random scales[2] sampled from Gaussian distribution. Combined with $\mathcal{RN}$, and the update rule can be expressed as:

$$\theta^* = \arg\min_{\theta} \mathcal{L}_{cos}(f_s^L(\mathcal{RN}(\boldsymbol{x'_s})), f_s^L(\mathcal{RN}(\boldsymbol{x_s}))). \tag{5}$$

3.4 Domain-agnostic Attention Module

Unlike the random normalization module, the domain-agnostic attention module aims to narrow the domain gap from the model perspective. Our inspiration is from prior works (Hansen & Salamon, 1990; Caruana et al., 2004; Dong et al., 2018), which demonstrate that the ensemble strategy can avoid getting trapped in the local optimum and improve performance. Since there are many feature maps at layer L and each of them can model the input, *i.e.,* extracts the features, we can also integrate them to produce a more robust feature representation $\mathcal{A}^L$, thus mitigating the impact of several biased feature maps, *e.g.*, the second feature map of Figure 3. Specifically, we apply a channel-wise average pooling to the feature maps at layer L:

$$\mathcal{A}^L = \frac{|\sum_{i=0}^{C} f_s^L(\boldsymbol{x_s})|}{C}, \tag{6}$$

where C is the channel number of $f_s^L(\boldsymbol{x_s})$.

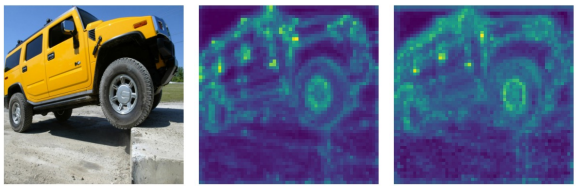

Figure 4: **Left:** A benign image from Stanford Cars (Krause et al., 2013). **Middle & Right:** We apply the channel-wise average pooling to the intermediate feature maps ($Maxpool.3$) of VGG-16 and ($Conv3_8$) of DCL (Chen et al., 2019) with backbone Res-50.

As depicted in Figure 4, even if our source model is not trained in the target domain, the robust feature representation is still able to capture the essential feature of object very well. Surprisingly, it is even similar to the one that derived from a completely different target model. Therefore, this robust feature representation can serve as a domain-agnostic attention ($\mathcal{DA}$) to enhance the cross-domain transferability of adversarial examples. Specifically, in the training phase, we leverage $\mathcal{A}^L$ to assign weights for each pixel of feature maps at the same layer. The resulting object function can be written as:

$$\theta^* = \arg\min_{\theta} \mathcal{L}_{cos}(\mathcal{A}^L \odot f_s^L(\boldsymbol{x'_s}), \mathcal{A}^L \odot f_s^L(\boldsymbol{x_s})), \tag{7}$$

where $\odot$ is Hadamard product. After finishing the training, even the inputs are not from the source domain, our generator $\mathcal{G}_{\theta^*}$ is supposed to own an ability to capture the essential feature to disrupt.

4 Experiments

Source (white-box) Domain. Our training data is the large-scale ImageNet (Russakovsky et al., 2015) training set which including about 1.2 million $3 \times 224 \times 224$ images. The generator is

[2] μ'_{mean}, μ'_{std}, σ'_{mean} and σ'_{std} are all hyper-parameters.

trained against four ImageNet pre-trained models including VGG-16, VGG-19 (Simonyan & Zisserman, 2015), ResNet152 (Res-152) (He et al., 2016) and DenseNet169 (Dense-169) (Huang et al., 2017). In this domain, we also consider another three pre-trained models, *i.e.*, DenseNet121 (Dense-121) (Huang et al., 2017), ResNet50 (Res-50) (He et al., 2016) and Inception-v3 (Inc-v3[3]) (Szegedy et al., 2016), to analyze the cross-model transferability. All models are available in the Torchvision[4].

Target (black-box) Domain. In addition to ImageNet (white-box domain), we also consider another seven black-box domains (shown in Table 1) including four coarse-grained (CIFAR-10, CIFAR-100 (Krizhevsky, 2009), STL-10 (Coates et al., 2011) and SVHN (Netzer et al., 2011)) and three fine-grained (CUB-200-2011 (Wah et al., 2011), Standford Cars (Krause et al., 2013) and FGVC Aircraft (Maji et al., 2013)) classification tasks. For the fine-grained classification, we use DCL framework (Chen et al., 2019) with three different backbones: ResNet50 (Res-50) (He et al., 2016), SENet154 and SE-ResNet101 (SE-Res101) (Hu et al., 2018). The pre-trained models for the coarse-grained classification are from Github[5].

Implementation Details. Our generator $\mathcal{G}_\theta$ adopts the same architecture as (Naseer et al., 2019), which is a composite of downsampling, residual (He et al., 2016) and upsampling blocks (more details can be found in Figure 11 of Appendix). The output size of $\mathcal{G}_\theta$ is equal to the input size. For CDA and our methods, we use Adam optimizer (Kingma & Ba, 2015) with a learning rate of 2e-4 and the exponential decay rate for first and second moments is set to 0.5 and 0.999, respectively. All generators are trained for one epoch with the batch size 16. For the layer L, we attack the output of $Maxpool.3$ for VGG-16 and VGG-19, the output of $Conv3_8$ for Res-152 and the output of $DenseBlock.2$ for Dense-169 (the ablation study can be found in Appendix A.2). For brevity, we will only refer the output of a specific layer/block using its layer/block name. For our $\mathcal{RN}$ module, we set $\mu' \sim \mathcal{N}(0.50, 0.08)$ and $\sigma' \sim \mathcal{N}(0.75, 0.08)$, and the ablation study is shown in Appendix A.3.

Competitors. We compare our proposed methods with projected gradient descent (PGD) (Madry et al., 2018), diverse inputs method (DIM) (Dong et al., 2018; Xie et al., 2019), dispersion reduction (DR) (Lu et al., 2020), self-supervised perturbation (SSP) (Naseer et al., 2020) and cross-domain attack (CDA) (Naseer et al., 2019). The maximum perturbation ε is set to 10. Follow Lu et al. (2020), we set the step size $\alpha = 4$ and the number of iterations $T = 100$ for all iterative methods. For DIM, we set the default decay factor $\mu = 1.0$ and the transformation probability $p = 0.7$. The Gaussian smoothing (gs) for CDA is applied by 3×3 Gaussian kernel.

Evaluation Metrics. We use the top-1 accuracy on the whole test set (test size is shown in Table 1) after attacking to evaluate the performance of different methods.

Table 2: Transferability comparisons on four coarse-grained classification tasks. Here we report the top-1 accuracy after attacking (the lower, the better). The generator $\mathcal{G}_\theta$ is trained in the ImageNet domain, and adversarial examples are within the perturbation budget of $\ell_\infty \leq 10$.

Model	Attacks	CIFAR-10	CIFAR-100	STL-10	SVHN	Model	Attacks	CIFAR-10	CIFAR-100	STL-10	SVHN
	Clean	93.78	74.27	77.59	96.03		Clean	93.78	74.27	77.59	96.03
VGG-16	PGD	79.63	48.02	74.32	94.66	Res-152	PGD	86.17	56.38	74.51	93.94
	DIM	77.16	44.75	72.74	91.53		DIM	80.50	48.03	71.2	90.87
	DR	72.49	39.04	72.56	93.27		DR	78.86	48.62	71.66	93.26
	SSP	68.54	33.63	72.77	93.98		SSP	75.54	42.38	72.66	92.63
	CDA	66.41	32.37	72.91	92.17		CDA	66.47	39.30	69.81	88.09
	CDA+gs	86.70	59.43	73.38	91.61		CDA+gs	85.61	57.27	73.06	90.34
	BIA (Ours)	57.38	22.47	69.45	90.44		BIA (Ours)	65.49	33.48	69.91	89.46
	BIA+$\mathcal{DA}$ (Ours)	55.16	21.71	70.00	91.76		BIA+$\mathcal{DA}$ (Ours)	65.34	**32.68**	69.65	91.38
	BIA+$\mathcal{RN}$ (Ours)	**52.81**	**20.82**	**67.55**	**88.03**		BIA+$\mathcal{RN}$ (Ours)	**61.23**	32.84	**68.04**	**85.79**
VGG-19	PGD	79.15	47.73	74.71	94.86	Dense-169	PGD	84.55	54.29	74.55	93.83
	DIM	77.54	43.81	72.96	91.68		DIM	80.89	49.06	72.64	89.47
	DR	70.72	37.59	71.98	93.73		DR	78.24	48.67	70.75	93.2
	SSP	70.46	35.28	73.21	93.67		SSP	77.13	42.18	72.53	91.64
	CDA	81.60	51.53	71.43	92.64		CDA	67.75	35.03	69.00	88.76
	CDA+gs	88.55	61.90	73.64	92.18		CDA+gs	85.01	54.71	72.61	88.69
	BIA (Ours)	57.88	23.12	69.84	88.89		BIA (Ours)	72.02	38.99	69.80	86.12
	BIA+$\mathcal{DA}$ (Ours)	57.26	23.04	70.16	90.08		BIA+$\mathcal{DA}$ (Ours)	71.69	38.95	70.60	88.02
	BIA+$\mathcal{RN}$ (Ours)	**54.47**	**22.61**	**68.23**	**88.08**		BIA+$\mathcal{RN}$ (Ours)	**66.67**	**34.41**	**68.79**	**81.54**

[3]Inc-v3 expects inputs with a size of $3 \times 299 \times 299$. Therefore, the adversarial examples are also craft on the same size inputs. Note the size of training data for our generator is still $3 \times 224 \times 224$.

[4]https://pytorch.org/vision/stable/models.html

[5]https://github.com/aaron-xichen/pytorch-playground

4.1 TRANSFERABILITY COMPARISONS

In this section, we first conduct experiments for the black-box domain in Section 4.1.1 (coarse-grain) and Section 4.1.2 (fine-grain), then we report the results for the white-box (source) domain in Section 4.1.3. For the discussion of $\mathcal{RN}$ + $\mathcal{DA}$ variant, we leave them in Appendix A.5.

4.1.1 RESULTS ON COARSE-GRAINED DOMAIN

In this section, we craft transferable adversarial examples for coarse-grained classification tasks. The results are shown in Table 2, where we leverage four ImageNet pre-trained models, including VGG-16, VGG-19, Res-152 and Dense-169, to train generators, respectively.

From Table 2, a first glance shows that our proposed methods consistently surpass state-of-the-art approaches. For example, if the substitute model is VGG-16, the most effective CDA remains a top-1 accuracy of 66.41% for CIFAR-10 after attacking, while our vanilla BIA can effectively bring down it to **57.38%** and $\mathcal{DA}$ variant can further drop the top-1 accuracy to **55.16%**. Among all tasks, the STL-10 and SVHN domains are the most difficult to attack, and the performance gap among existing attacks is moderate. Nonetheless, we can significantly enhance the transferability with the help of our $\mathcal{RN}$ variant in these domains. Notably, if the substitute model is Res-152, $\mathcal{RN}$ variant can further decrease the top-1 accuracy from 89.46% (vanilla BIA) to **85.79%** on SVHN. Compared with state-of-the-art CDA on all domains, $\mathcal{RN}$ variant significantly outperforms it by **7.71%** on average. This demonstrates that our proposed $\mathcal{RN}$ module is effective in coping with different distributions of inputs, thus improving the generalization of the resulting generator.

Table 3: Transferability comparisons on three fine-grained classification tasks. Here we report the top-1 accuracy after attacking (the lower, the better). The generator $\mathcal{G}_\theta$ is trained in ImageNet domain and adversarial examples are within the perturbation budget of $\ell_\infty \leq 10$.

Model	Attacks	CUB-200-2011			Standford Cars			FGVC Aircraft		
		Res-50	SENet154	SE-Res101	Res-50	SENet154	SE-Res101	Res-50	SENet154	SE-Res101
	Clean	87.35	86.81	86.56	94.35	93.36	92.97	92.23	92.08	91.90
VGG-16	PGD	80.65	79.58	80.69	87.45	89.04	90.30	84.88	83.92	82.15
	DIM	70.02	62.86	70.57	74.72	78.10	84.33	73.54	66.88	62.38
	DR	81.08	82.05	82.52	90.82	90.59	91.12	84.97	87.55	85.54
	SSP	62.27	60.44	71.52	58.02	75.71	83.02	54.91	68.74	63.79
	CDA	69.69	62.51	71.00	75.94	72.45	84.64	71.53	58.33	63.39
	CDA+gs	70.19	63.19	68.92	85.03	79.52	83.52	78.55	65.62	68.38
	BIA (Ours)	32.74	52.99	58.04	39.61	69.90	70.17	28.92	60.31	51.43
	BIA+$\mathcal{DA}$ (Ours)	**25.00**	**40.27**	**53.24**	22.24	**59.48**	**61.52**	**15.36**	**47.91**	**40.83**
	BIA+$\mathcal{RN}$ (Ours)	26.13	46.15	55.07	**20.61**	62.64	63.38	16.50	52.54	45.48
VGG-19	PGD	80.98	79.00	80.60	87.54	88.87	90.56	84.70	84.01	83.35
	DIM	69.93	61.60	70.90	75.02	78.55	84.63	74.59	67.69	65.26
	DR	80.83	81.57	81.95	91.00	90.23	91.15	84.43	85.96	84.57
	SSP	62.94	58.34	70.45	61.90	76.27	83.77	58.78	69.52	66.88
	CDA	59.48	61.08	68.50	58.53	70.70	80.70	59.26	52.24	62.26
	CDA+gs	67.88	59.42	67.57	82.83	78.61	82.64	79.36	65.89	68.35
	BIA (Ours)	48.90	52.33	56.47	66.34	72.45	75.08	50.95	54.04	51.79
	BIA+$\mathcal{DA}$ (Ours)	**27.46**	**37.61**	**50.14**	**35.27**	**61.40**	**64.41**	**17.97**	**45.81**	**44.01**
	BIA+$\mathcal{RN}$ (Ours)	31.77	43.41	51.09	42.81	68.90	66.27	27.75	52.48	45.57
Res-152	PGD	73.21	75.89	76.11	83.99	86.89	88.24	79.00	79.30	75.64
	DIM	56.30	59.35	63.36	67.88	76.37	79.77	64.21	62.95	54.49
	DR	77.58	82.07	80.60	87.96	90.11	90.67	77.89	82.27	80.08
	SSP	47.64	66.17	67.90	53.29	79.09	83.45	58.35	73.36	69.34
	CDA	45.15	53.69	52.86	57.72	63.09	73.05	64.87	46.74	59.11
	CDA+gs	59.32	64.24	60.08	81.15	82.91	82.56	75.52	73.72	66.58
	BIA (Ours)	43.55	49.50	55.54	31.65	54.12	67.21	38.49	32.46	51.19
	BIA+$\mathcal{DA}$ (Ours)	25.80	**39.71**	52.83	33.43	52.41	68.03	**27.51**	**27.42**	47.28
	BIA+$\mathcal{RN}$ (Ours)	**23.54**	40.13	**51.36**	**12.39**	**47.92**	**60.59**	32.34	32.85	**46.89**
Dense-169	PGD	79.29	81.14	79.74	87.66	90.13	89.80	83.23	84.31	81.24
	DIM	63.17	62.01	65.96	72.88	78.29	81.25	68.89	65.26	54.79
	DR	74.84	78.89	77.93	86.54	88.82	89.52	77.80	78.73	74.47
	SSP	41.80	49.95	59.72	26.65	68.71	74.36	16.80	55.78	44.07
	CDA	52.92	60.96	57.04	53.64	73.66	75.51	62.23	61.42	59.83
	CDA+gs	60.86	61.34	60.10	74.95	76.35	78.86	72.94	68.68	64.66
	BIA (Ours)	21.79	**29.29**	39.13	9.58	44.46	49.06	8.04	27.84	33.87
	BIA+$\mathcal{DA}$ (Ours)	12.36	29.31	**35.78**	**4.56**	**40.85**	**31.82**	3.90	**13.59**	**14.55**
	BIA+$\mathcal{RN}$ (Ours)	**10.67**	32.50	46.62	7.76	44.14	42.76	**3.81**	34.20	27.00

Table 4: Transferability comparisons on ImageNet (source domain). Here we report the top-1 accuracy after attacking (the lower, the better). The generator $\mathcal{G}_\theta$ is trained in ImageNet domain ("*" denotes white-box model) and adversarial examples are within the perturbation budget of $\ell_\infty \leq 10$.

Model	Attack	VGG-16	Dense-169	VGG-19	Res-50	Res-152	Dense-121	Inc-v3
	Clean	70.14	75.75	70.95	74.61	77.34	74.22	76.19
VGG-16	PGD	2.49*	53.22	4.27	45.52	58.69	48.28	61.08
	DIM	3.32*	**22.72**	3.39	19.13	33.03	18.50	30.05
	DR	20.95*	67.91	43.59	64.90	70.00	64.73	69.22
	SSP	0.95*	39.56	3.42	26.22	41.68	34.45	47.46
	CDA	**0.40***	42.67	**0.77**	36.27	51.05	38.89	54.02
	CDA+gs	12.10*	57.09	20.48	51.87	60.83	52.21	55.12
	BIA (Ours)	1.55*	32.35	3.61	25.36	42.98	26.97	41.20
	BIA+$\mathcal{DA}$ (Ours)	1.04*	24.52	2.07	18.63	36.43	19.97	34.54
	BIA+$\mathcal{RN}$ (Ours)	1.44*	25.96	2.58	**16.52**	**31.80**	**18.25**	**28.54**
Dense-169	PGD	38.51	5.03*	40.53	33.91	44.97	21.18	58.30
	DIM	12.31	5.25*	13.20	8.98	12.93	5.91	21.44
	DR	38.45	23.99*	41.59	50.19	58.70	49.95	63.70
	SSP	11.53	1.32*	12.54	12.97	25.66	9.74	25.58
	CDA	7.26	**0.63***	7.91	6.46	15.56	5.13	43.78
	CDA+gs	27.98	22.95*	28.56	31.52	43.67	31.66	49.18
	BIA (Ours)	4.76	6.45*	7.15	6.97	13.83	6.60	38.58
	BIA+$\mathcal{DA}$ (Ours)	**3.17**	3.32*	**4.09**	**4.44**	**5.85**	**3.98**	26.51
	BIA+$\mathcal{RN}$ (Ours)	3.66	4.05*	5.23	6.91	13.25	4.21	**14.24**

4.1.2 Results on Fine-grained Domain

We also analyze the transferability of adversarial examples towards fine-grained classification tasks. For each domain, three black-box models with different backbones trained via the DCL framework are the target. The results are summarized in Table 3, where the leftmost column is the substitute model and the top row shows the target model.

In this scenario, the performance gap between the existing state-of-the-art algorithms and our proposed methods is further enlarged. Remarkably, when attacking against Dense-169, even the vanilla BIA can drop the average top-1 accuracy to **29.23%**, while CDA+gs, CDA, SSP, DR, DIM and PGD are still with the high average top-1 accuracy of 68.75%, 61.91%, 48.65%, 80.84%, 68.06% and 84.06% after attacking, respectively. Furthermore, by adding $\mathcal{DA}$ or $\mathcal{RN}$ modules in the training phase, the generator $\mathcal{G}_{\theta^*}$ is capable of crafting more transferable adversarial examples. On average, our $\mathcal{RN}$ variant can reduce the top-1 accuracy from 46.65% (vanilla BIA) to **39.33%**, and $\mathcal{DA}$ variant can further drop it to **36.42%**, which remarkably outperforms SSP by **25.91%**. This demonstrates our proposed $\mathcal{DA}$ module can effectively alleviate the bias caused by several feature maps, thus focusing on disrupting essential features.

4.1.3 Results on Source Domain

Although our proposed methods are mainly designed to improve the threat of adversarial examples towards black-box domains, they are also effective for enhancing the cross-model black-box transferability in the white-box domain. For example, by training against Dense-169, CDA still remains a top-1 accuracy of 43.78% on Inc-v3, while our BIA can achieve a relatively low top-1 accuracy of 38.58% on it. Besides, $\mathcal{RN}$ variant can further decrease the top-1 accuracy on Dense-169 (white-box model) and Inc-v3 (black-box model) by **2.4%** and **24.34%**, respectively. This demonstrates that our proposed $\mathcal{RN}$ module is also able to avoid getting stuck in the local optimum of a specific model when training on a large-scale dataset. For $\mathcal{DA}$ variant, since the target domain and source domain are identical, it is naturally able to improve the transferability as well. As shown in Table 4, compared with vanilla BIA, $\mathcal{DA}$ variant can further degrade the top-1 accuracy from 18.45% (vanilla BIA) to **13.47%** on average.

4.2 Insight into the Generator

Although many prior works (Poursaeed et al., 2018; Naseer et al., 2019) have leveraged the generator to craft adversarial examples, they hardly analyze the role of each block of the generator. This section will give an insight into the generator and understand how it processes an input. In Figure 5, we feed an image into the generator trained with the vanilla BIA (against Res-152) and visualize the

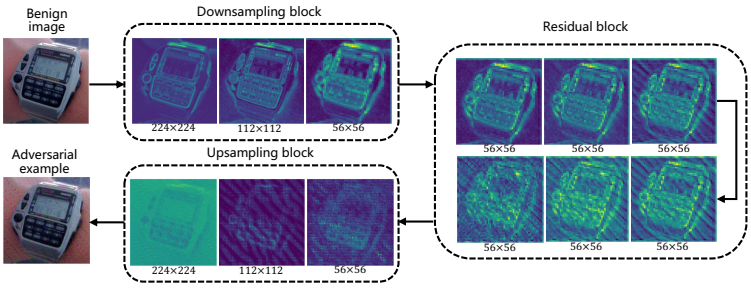

Figure 5: We apply channel-wise average pooling to visualize each block of our generator $\mathcal{G}_{\theta^*}$.

output of each block. As we can observe, these blocks play different roles in crafting adversarial examples:

- Downsampling block: it mainly extracts the features of the input image. As the size of the down-sampled images gets smaller, there is no significant noise overall.
- Residual block: Unlike the downsampling block, this block is responsible for adding noise and its behavior looks very similar to the iterative algorithm.
- Upsampling block: This block gradually reconstructs the adversarial example from abstract features.

In sum, the downsampling module has an essential impact on the resulting adversarial examples (the residual block mainly works as suppressing or reversing the features extracted by the downsampling module).

To better understand the effectiveness of our proposed modules, here we investigate them from the perspective of the generator mechanism. Specifically, we apply channel-wise average pooling to the output of the downsampling block and calculate the difference map between vanilla BIA and our proposed variants (The formula definitions is shown in Appendix A.1).

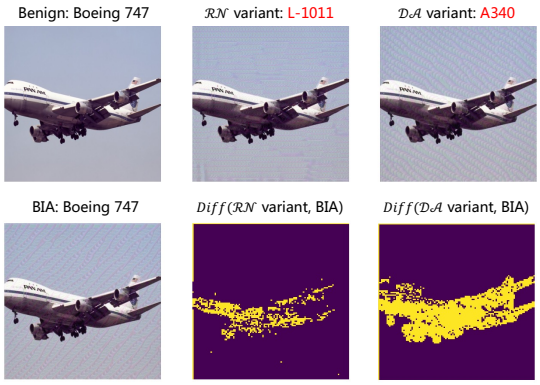

Figure 6: A benign image (label is "Boeing 747") from FGVC Aircraft (Maji et al., 2013) and its corresponding adversarial examples and difference map.

From the result of Figure 6, we can observe that the generators derived from our proposed variants concentrate more on the body of the object (especially for $\mathcal{DA}$ variant) than that of vanilla BIA (which pays more attention to the background, *i.e.*, the black region in the difference map). This demonstrates that our proposed $\mathcal{RN}$ and $\mathcal{DA}$ variants do narrow the domain gap, and thus be capable of yielding more transferable adversarial examples.

5 Conclusion

In this paper, we present a practical black-box threat model for the cross-domain attack. Specifically, we train a generator network in the large-scale ImageNet domain to disrupt low-level features better, thus generating transferable adversarial examples for the black-box domain. Based on this framework, we further propose two variants to narrow the gap between the source and target domains from the data and model perspectives, respectively. Extensive experiments demonstrate the effectiveness of our proposed methods. This also reminds the model owner that "Your deployed model is not safe even you do not leak any information to the public". We hope our proposed approaches can serve as a benchmark for evaluating the stability of various deployed models.

REFERENCES

Shumeet Baluja and Ian Fischer. Adversarial transformation networks: Learning to generate adversarial examples. *CoRR*, abs/1703.09387, 2017.

Wieland Brendel, Jonas Rauber, and Matthias Bethge. Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. In *ICLR*, 2018.

Nicholas Carlini and David A. Wagner. Towards evaluating the robustness of neural networks. In *Symposium on Security and Privacy*, 2017.

Rich Caruana, Alexandru Niculescu-Mizil, Geoff Crew, and Alex Ksikes. Ensemble selection from libraries of models. In *ICML*, 2004.

Jianbo Chen, Michael I. Jordan, and Martin J. Wainwright. Hopskipjumpattack: A query-efficient decision-based attack. In *SP*, 2020.

Yue Chen, Yalong Bai, Wei Zhang, and Tao Mei. Destruction and construction learning for fine-grained image recognition. In *CVPR*, 2019.

Adam Coates, Andrew Y. Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In Geoffrey J. Gordon, David B. Dunson, and Miroslav Dudík (eds.), *AISTATS*, 2011.

Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *CVPR*, 2018.

Ranjie Duan, Xiaofeng Mao, A. Kai Qin, Yuefeng Chen, Shaokai Ye, Yuan He, and Yun Yang. Adversarial laser beam: Effective physical-world attack to dnns in a blink. In *CVPR*, 2021.

Lianli Gao, Qilong Zhang, Jingkuan Song, Xianglong Liu, and Hengtao Shen. Patch-wise attack for fooling deep neural network. In *ECCV*, 2020.

Lianli Gao, Yaya Cheng, Qilong Zhang, Xing Xu, and Jingkuan Song. Feature space targeted attacks by statistic alignment. In *IJCAI*, 2021.

Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In Yoshua Bengio and Yann LeCun (eds.), *ICLR*, 2015.

Lars Kai Hansen and Peter Salamon. Neural network ensembles. *IEEE Trans. Pattern Anal. Mach. Intell.*, 1990.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018.

Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *CVPR*, 2017.

Nathan Inkawhich, Wei Wen, Hai (Helen) Li, and Yiran Chen. Feature space perturbations yield more transferable adversarial examples. In *CVPR*, 2019.

Nathan Inkawhich, Kevin J. Liang, Binghui Wang, Matthew Inkawhich, Lawrence Carin, and Yiran Chen. Perturbing across the feature hierarchy to improve standard and strict blackbox attack transferability. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *NeurIPS*, 2020.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.

Jonathan Krause, Jia Deng, Michael Stark, and Li Fei-Fei. Collecting a large-scale dataset of fine-grained cars. 2013.

Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *NeurPIS*, 2012.

Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. Adversarial examples in the physical world. In *ICLR*, 2017a.

Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. Adversarial machine learning at scale. In *ICLR*, 2017b.

Qizhang Li, Yiwen Guo, and Hao Chen. Practical no-box adversarial attacks against dnns. In *NeurIPS*, 2020.

Xiaodan Li, Jinfeng Li, Yuefeng Chen, Shaokai Ye, Yuan He, Shuhui Wang, Hang Su, and Hui Xue. QAIR: practical query-efficient black-box attacks for image retrieval. In *CVPR*, 2021.

Yantao Lu, Yunhan Jia, Jianyu Wang, Bai Li, Weiheng Chai, Lawrence Carin, and Senem Velipasalar. Enhancing cross-task black-box transferability of adversarial examples with dispersion reduction. In *CVPR*, 2020.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018.

Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew B. Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. volume abs/1306.5151, 2013.

Xiaofeng Mao, Yuefeng Chen, Shuhui Wang, Hang Su, Yuan He, and Hui Xue. Composite adversarial attacks. In *AAAI*, 2021.

Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: A simple and accurate method to fool deep neural networks. In *CVPR*, 2016.

Muzammal Naseer, Salman H. Khan, Muhammad Haris Khan, Fahad Shahbaz Khan, and Fatih Porikli. Cross-domain transferability of adversarial perturbations. In *NeurPIS*, 2019.

Muzammal Naseer, Salman H. Khan, Munawar Hayat, Fahad Shahbaz Khan, and Fatih Porikli. A self-supervised approach for adversarial robustness. In *CVPR*, 2020.

Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. 2011.

Nicolas Papernot, Patrick D. McDaniel, and Ian J. Goodfellow. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *CoRR*, abs/1605.07277, 2016.

Omid Poursaeed, Isay Katsman, Bicheng Gao, and Serge J. Belongie. Generative adversarial perturbations. In *CVPR*, 2018.

Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Fei-Fei Li. Imagenet large scale visual recognition challenge. *IJCV*, 2015.

Mahmood Sharif, Sruti Bhagavatula, Lujo Bauer, and Michael K. Reiter. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *SIGSAC*, 2016.

Yucheng Shi, Siyu Wang, and Yahong Han. Curls & whey: Boosting black-box adversarial attacks. In *CVPR*, 2019.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In Yoshua Bengio and Yann LeCun (eds.), *ICLR*, 2015.

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In Yoshua Bengio and Yann LeCun (eds.), *ICLR*, 2014.

Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016.

C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical report, California Institute of Technology, 2011.

Xiaosen Wang, Xuanran He, Jingdong Wang, and Kun He. Admix: Enhancing the transferability of adversarial attacks. In *ICCV*, 2021.

Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M. Summers. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *CVPR*, 2017.

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 *ICLR*, 2020.

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 *CVPR*, 2019.

Kaidi Xu, Gaoyuan Zhang, Sijia Liu, Quanfu Fan, Mengshu Sun, Hongge Chen, Pin-Yu Chen, Yanzhi Wang, and Xue Lin. Adversarial t-shirt! evading person detectors in a physical world. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (eds.), *ECCV*, 2020.

Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In Zoubin Ghahramani, Max Welling, Corinna Cortes, Neil D. Lawrence, and Kilian Q. Weinberger (eds.), *NeurIPS*, 2014.

A APPENDIX

A.1 THE DEFINITION FOR DIFFERENCE MAP

This section will describe the formula definition for difference map $Diff(\cdot,\cdot)$. Without loss of generality, here we only show the definition of $Diff(\mathcal{RN}\,variant$, BIA), and $Diff(\mathcal{DA}\,variant$, BIA) can be easily deduced. Specifically, we first apply channel-wise pooling to the output of the downsampling block:

$$\mathcal{A}^L_{RN} = \frac{|\sum_{i=0}^{C} \mathcal{G}^d_{\mathcal{RN}_{\theta^*}}(\boldsymbol{x})|}{C}, \tag{8}$$

$$\mathcal{A}^L_{BIA} = \frac{|\sum_{i=0}^{C} \mathcal{G}^d_{BIA_{\theta^*}}(\boldsymbol{x})|}{C}, \tag{9}$$

where $\mathcal{G}^d_{\mathcal{RN}_{\theta^*}}$ and $\mathcal{G}^d_{\mathcal{BIA}_{\theta^*}}$ denote the output of the downsampling block for $\mathcal{RN}$ variant and vanilla BIA, respectively. Then our difference map can be expressed by:

$$Diff(\mathcal{RN}\,variant, BIA) = \begin{cases} 1, & \mathcal{A}^L_{RN} - \mathcal{A}^L_{BIA} > 0, \\ 0, & else. \end{cases} \tag{10}$$

A.2 SELECT LAYER FOR ATTACKING

In this section, we analyze the impact of different intermediate layers of the substitute model on the transferability of resulting adversarial examples. The results are illustrated in Figure 8.

In general, training against the shallow and middle layers yields more cross-domain transferable but less cross-model transferable adversarial examples than against the deep layer. For example, if the substitute model is Res-152, disrupting shallow layer $Conv2_3$ is more effective for transferring towards coarse-grained domain, and perturbing middle layer $Conv3_8$ is more effective in reducing the average top-1 accuracy of fine-grained models. In contrast, attacking deep layers like $Conv5_3$ can yield more transferable adversarial examples in the source domain. This demonstrates that low-level features are more similar across domains and high-level features are more domain-specific.

A.3 SELECT GAUSSIAN DISTRIBUTION FOR RANDOM NORMALIZATION

For $\mathcal{DA}$ module, it is parameter-free. Therefore, we only conduct the experiment to select an optimal distribution for $\mathcal{RN}$ module, *i.e.*, the μ' and σ' in Equation 4. Here we tune the mean of μ' and σ' from 0.25 to 0.75 with a granularity of 0.25. For the standard deviation of them, we fix it to 0.08 so that the sampled μ' and σ' can basically take values ranging from 0.0 to 1.0 (according to the three-sigma rule).

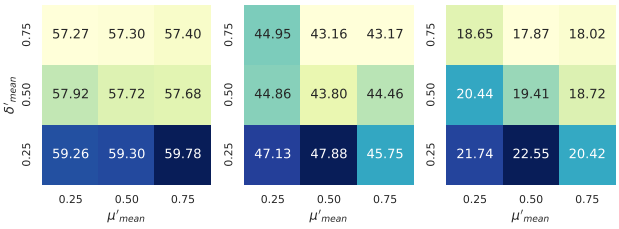

Figure 7: The average top-1 accuracy after attacking on coarse-grained (left), fine-grained (middle) and source (right) domains with different μ'_{mean} and σ'_{mean} for $\mathcal{RN}$.

The results are shown in Figure 7, where we craft adversarial examples via VGG-16 and report the average top-1 accuracy in both source and black-box domains. From the results, we observe that a bigger σ'_{mean} can effectively improve the transferability. For example, if we increase σ'_{mean} from 0.25 to 0.75, the top-1 accuracy on fine-grained models can be further decreased by 3.16% on average. Although the influence of μ'_{mean} is relatively moderate when $\sigma'_{mean} = 0.75$, setting μ'_{mean} to 0.5 is usually better. Therefore, we set $\mu' \sim \mathcal{N}(0.50, 0.08)$ and $\sigma' \sim \mathcal{N}(0.75, 0.08)$ in our paper.

A.4 DATA AUGMENTATION VS. RANDOM NORMALIZATION

Data augmentation (Krizhevsky et al., 2012; Simonyan & Zisserman, 2015) is a widely used strategy for improving the generalization of the model. Nonetheless, these label-preserving transformations are less effective for training a generator to craft more transferable adversarial examples.

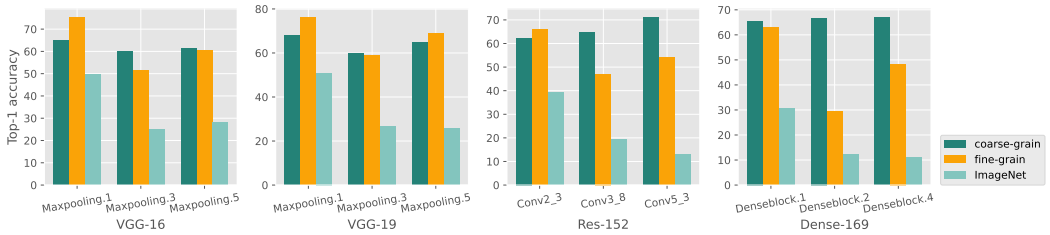

Figure 8: The average top-1 accuracy after attacking (the lower, the better) on coarse-grained, fine-grained and source domains. Our generator $\mathcal{G}_\theta$ is trained against different layers (from shallow to deep) of ImageNet pre-trained VGG-16, VGG-19, Res-152 and Dense-169, respectively.

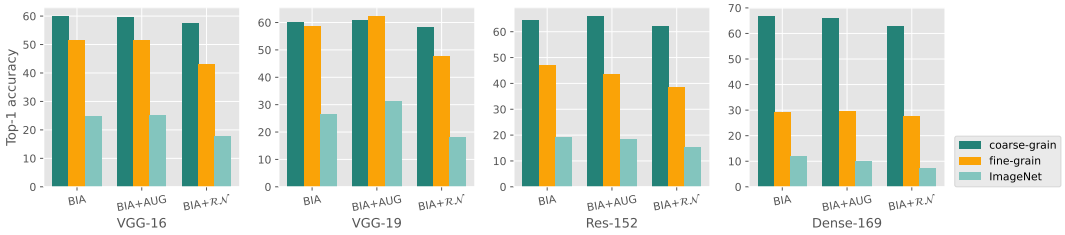

Figure 9: The average top-1 accuracy after attacking on coarse-grained, fine-grained and source domains. Here we compare the results of vanilla BIA, BIA with data augmentation (AUG) and BIA with $\mathcal{RN}$. Our generator $\mathcal{G}_\theta$ is trained against ImageNet pre-trained VGG-16, VGG-19, Res-152 and Dense-169, respectively.

As demonstrated in Figure 9, where we report the results for our vanilla BIA, BIA with data augmentation (BIA+AUG)[6] and BIA with $\mathcal{RN}$ (BIA+$\mathcal{RN}$), BIA+AUG is less effective than our proposed BIA+$\mathcal{RN}$ and might even degrade the performance of our vanilla BIA. For example, when training against VGG-19, BIA gets an average top-1 accuracy of 58.71% on the fine-grained domain, yet BIA+AUG degrades it to 62.29%. In contrast, our proposed BIA+$\mathcal{RN}$ can significantly decrease the result to **47.78%**. This is mainly because that the common data augmentation cannot effectively change the distribution of the training dataset, thus decreasing the generalization towards the black-box domain.

A.5 The Effect of Combining Domain-agnostic Attention and Random Normalization

In this section, we report the results for BIA equipped with both $\mathcal{RN}$ and $\mathcal{DA}$. The following is the update rule:

$$\theta^* = \arg\min_\theta \mathcal{L}_{cos}(\mathcal{A}^L \odot f_s^L(\mathcal{RN}(\boldsymbol{x'_s})), \mathcal{A}^L \odot f_s^L(\mathcal{RN}(\boldsymbol{x_s}))). \tag{11}$$

As illustrated in Figure 10, $\mathcal{RN}$ module and $\mathcal{DA}$ module are not always mutually reinforcing. For example, when training against Res-152 and transferring adversarial examples to fine-grained or source domains, combining $\mathcal{RN}$ module and $\mathcal{DA}$ module can further decrease the average top-1 accuracy to 35.75% and 13.11%, respectively. However, if the black-box domain is coarse-grain, only applying $\mathcal{RN}$ module is better than applying both $\mathcal{RN}$ and $\mathcal{DA}$ modules. We speculate that it may be because the $\mathcal{RN}$ module affects the low-level features extract by the substitute model, and thus be incompatible with the $\mathcal{DA}$ module sometimes.

[6]Specifically, we use three common data augmentation techniques including "RandomResizedCrop", "RandomHorizontalFlip" and "ColorJitter".

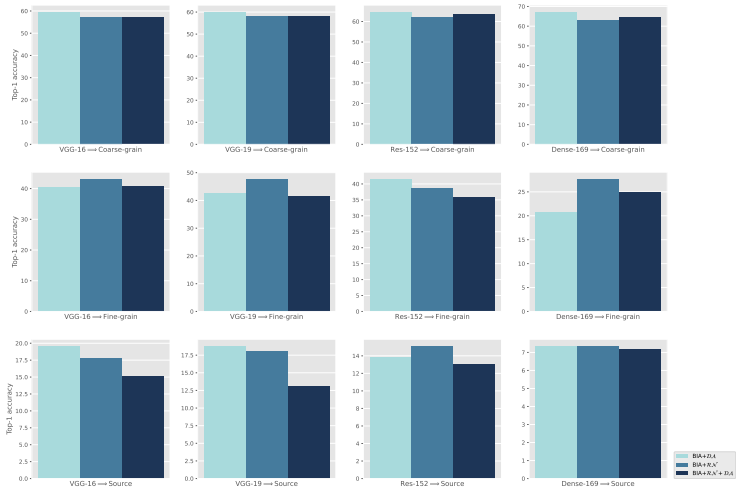

Figure 10: The average top-1 accuracy after attacking on coarse-grained, fine-grained and source domains. Here we compare the results of BIA with $\mathcal{DA}$, BIA with $\mathcal{RN}$ and BIA with both $\mathcal{DA}$ and $\mathcal{RN}$.

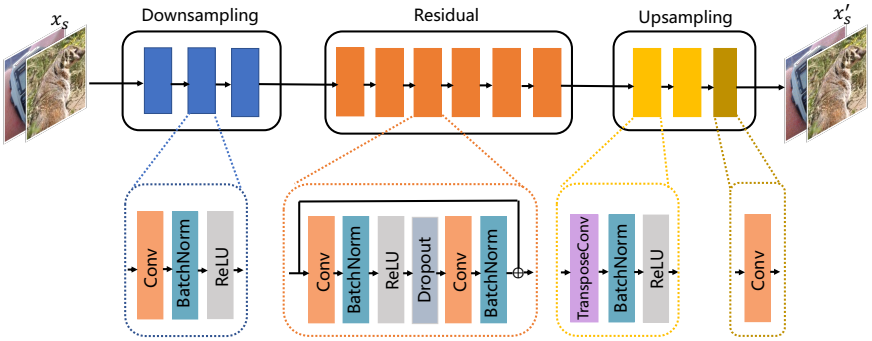

Figure 11: The structure of the generator.

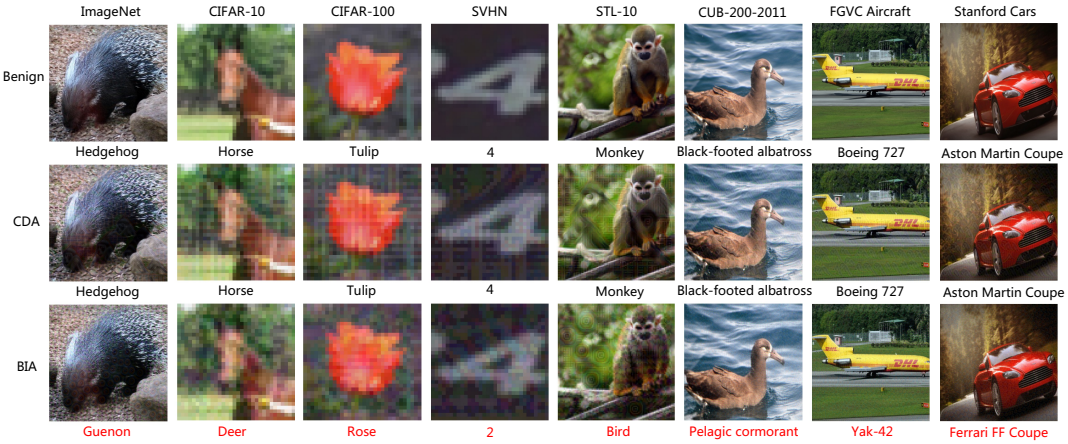

Figure 12: Adversarial examples crafted by CDA and our vanilla BIA ($\epsilon = 10$). Both generator networks are trained against ImageNet pre-trained VGG-16 (Simonyan & Zisserman, 2015). Red highlighted labels represent misclassification.