

Feature Importance-aware Transferable Adversarial Attacks

Zhibo Wang^{†,‡}, Hengchang Guo^{‡,*}, Zhifei Zhang^{*}, Wenxin Liu[‡], Zhan Qin[†], Kui Ren^{†,‡}

[†]School of Cyber Science and Technology, Zhejiang University, P. R. China

[‡]School of Cyber Science and Engineering, Wuhan University, P. R. China, ^{*}Adobe Research

[#]Key Laboratory of Blockchain and Cyberspace Governance of Zhejiang Province, P. R. China

{zhibowang, qinzhn, kuiren}@zju.edu.cn, {hc_guo, wxliu111}@whu.edu.cn, zzhang@adobe.com

Abstract

Transferability of adversarial examples is of central importance for attacking an unknown model, which facilitates adversarial attacks in more practical scenarios, e.g., black-box attacks. Existing transferable attacks tend to craft adversarial examples by indiscriminately distorting features to degrade prediction accuracy in a source model without aware of intrinsic features of objects in the images. We argue that such brute-force degradation would introduce model-specific local optimum into adversarial examples, thus limiting the transferability. By contrast, we propose the Feature Importance-aware Attack (FIA), which disrupts important object-aware features that dominate model decisions consistently. More specifically, we obtain feature importance by introducing the aggregate gradient, which averages the gradients with respect to feature maps of the source model, computed on a batch of random transforms of the original clean image. The gradients will be highly correlated to objects of interest, and such correlation presents invariance across different models. Besides, the random transforms will preserve intrinsic features of objects and suppress model-specific information. Finally, the feature importance guides to search for adversarial examples towards disrupting critical features, achieving stronger transferability. Extensive experimental evaluation demonstrates the effectiveness and superior performance of the proposed FIA, i.e., improving the success rate by 9.5% against normally trained models and 12.8% against defense models as compared to the state-of-the-art transferable attacks. Code is available at: <https://github.com/hcguo00/FIA>

1. Introduction

Deep neural networks (DNNs) have achieved superior performance in many vision tasks, e.g., image classification [17, 12], object detection [9, 25], semantic segmenta-

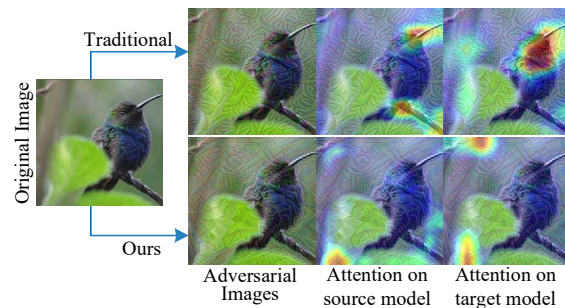


Figure 1. Comparison of traditionally indiscriminate attacks (top row) and our feature importance-aware attacks (bottom row). Adversarial images are generated on the source model (VGG16) and used to attack the target model (Inception-V3). Our attacks will suppress important features and promote trivial features, thus significantly defocusing/misleading the models as visualized by the attention maps, i.e., our adversarial example causes the source and target models not only failed to capture the important object but also focus on trivial regions.

tion [4, 22], face recognition [31, 27], etc. However, despite the impressive progress, recent studies showed that DNNs are vulnerable to adversarial examples [30] which are crafted by adding carefully designed perturbations to fool DNNs. Adversarial attacks have raised great concern for DNN-based applications, especially in safety- and security-sensitive areas like autonomous driving. Meanwhile, adversarial examples also play an important role in investigating the internal drawbacks of neural networks and improving their robustness.

Many works [30, 10, 18, 3, 34, 33, 35] have been proposed to generate adversarial examples, which can be divided into two categories, i.e., white-box attacks vs. black-box attacks, according to the knowledge owned by attackers. With the progress of adversarial attacks, the more challenging black-box attacks have attracted more attention. A common type of black-box attacks [33, 14, 2] is to craft adversarial examples by estimating gradients based on queried information (e.g., probability vectors and hard labels), which is referred to as the query-based attack. Those

*Hengchang Guo is the corresponding author.

query-based attacks may be impractical in the real world since excessive queries would not be allowed. By contrast, another typical black-box attack, called transfer-based attack, relies on the cross-model transferability of adversarial examples [21] (*i.e.*, adversarial examples crafted on one model could successfully attack other models for the same task), which is more practical and flexible.

However, adversarial examples crafted by traditional attacking methods (*e.g.*, FGSM [10], BIM [18], *etc.*) often exhibit weak transferability due to overfitting to the source model. Therefore, some studies attempted to alleviate such overfitting by introducing extra operations during the optimization to improve transferability, *e.g.*, random transformation [35], translation operation [6]. Recently, [36, 7, 23] performed attacks in the intermediate layers directly to enhance transferability. Instead of disturbing the output layer, these feature-level attacks maximize internal feature distortion and achieve higher transferability. However, existing methods generate adversarial examples by indiscriminately distorting features without aware of the intrinsic features of objects in the images, thus easily trapped into model-specific local optimum. Because classifiers tend to extract any available signal to maximize classification accuracy, even those imperceptible noises implied in the images [15], the model will learn extra “noisy” features together with intrinsic features of objects, while the “noisy” features are treated equally with object-related features to support the final decision, and such “noisy” features will be model-specific. Therefore, the adversarial examples crafted by existing methods tend to distort such model-specific features, thus overfitting to the source model and hindering the transferability of the adversarial examples.

This paper proposes a Feature Importance-aware Attack (FIA), which significantly improves the transferability of adversarial examples by disrupting the important object-aware features that are supposed to dominate the decision of different models. Against model-specific features, we introduce aggregate gradient, which will effectively suppress model-specific features while at the same time providing object-aware importance of the features. As illustrated in Fig. 1, compared to traditionally indiscriminate attacks, the adversarial image from the proposed FIA significantly defocuses the models, *i.e.*, failed to capture the important features of the object. Meanwhile, the models are misled to focus on those trivial areas. More specifically, random transformations (we adopt random pixel dropping) are first applied to the original images. Since the transformed images will preserve the spatial structure and texture but varying non-semantic details, the features from them will be consistent on the object-aware features but fluctuated on non-object (*i.e.*, model-specific “noisy”) features. With respect to these features, gradients are averaged to statistically suppress those fluctuated model-specific features. Meanwhile,

object-aware/important features are preserved to guide the generation of more transferable adversarial examples since the feature importance is highly correlated to objects of interest and consistent across different models.

Our main contributions are summarized as follows.

- We propose Feature Importance-aware Attack (FIA) that enhances the transferability of adversarial examples by disrupting the critical object-aware features that dominate the decision of different models.
- We analyze the rationale behind the relatively low transferability of existing works, *i.e.*, overfitting to model-specific “noisy” features, against which we introduce aggregate gradient to guide the generation of more transferable adversarial examples.
- Extensive experiments on diverse classification models demonstrate the superior transferability of adversarial examples generated by the proposed FIA as compared to state-of-the-art transferable attacking methods.

2. Related Work

Since Szegedy *et al.* [30] demonstrated the existence of adversarial examples, many adversarial attack algorithms [10, 18, 3, 33, 14, 2] have been proposed to discuss the vulnerability of neural networks. In this work, we mainly focus on the transfer-based attacks which utilize the transferability of adversarial examples to perform black-box attack, *i.e.*, using adversarial examples crafted on one model to attack other models.

There are many works aiming to improve the transferability of adversarial examples. Dong *et al.* [5] adopted momentum in iterative algorithms to stabilize updates and avoid poor local optimum, while Lin *et al.* [20] introduced Nesterov accelerated gradient to further enhance transferability. Xie *et al.* [35] applied image transformation (*i.e.*, randomly resizing and padding) to the inputs at each iteration to improve transferability. Gao *et al.* [8] crafted patch-wise noise instead of pixel-wise noise to enhance the transferability of adversarial examples. Dong *et al.* [6] proposed translation-invariant attack by optimizing perturbations over translated images which leads to higher transferability against defense models.

Instead of disrupting the output layer, several works proposed to attack internal features. Zhou *et al.* [36] first demonstrated that maximizing the feature distance between natural images and their adversarial examples in the intermediate layers can boost the transferability. Naseer *et al.* [24] also concluded that neural representation distortion does not suffer from the overfitting problem, and can exhibit cross-architecture, cross-dataset and cross-task transferability. Huang *et al.* [13] fine-tuned an existing adversarial example for greater transferability by increasing its perturba-

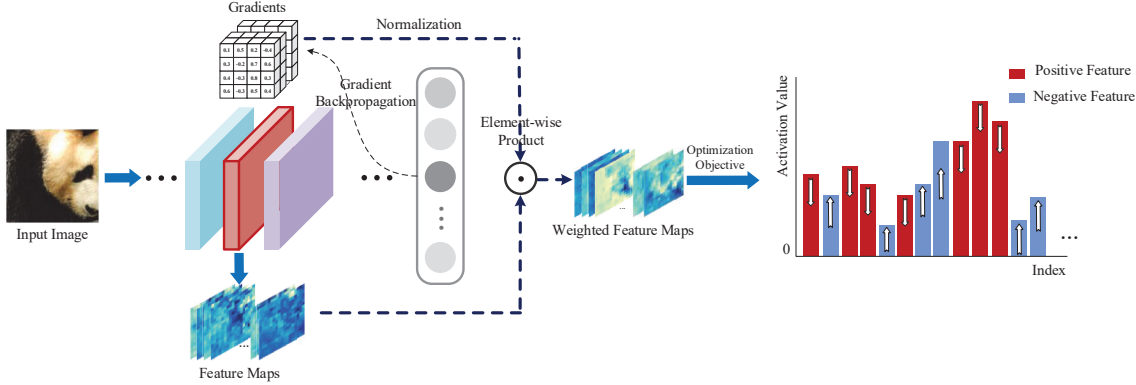


Figure 2. Overview of Feature Importance-aware Attack (FIA). Given an input image, feature maps are extracted from an intermediate layer (red block) of the source classification-model. Then, the gradients back propagated from the output to the feature maps are calculated to serve as the feature importance. After element-wise product between the feature maps and normalized gradients (*i.e.*, feature importance), the weighted feature maps could be optimized by specifically suppressing positive/important features and promoting negative/trivial features, achieving higher transferable adversarial examples.

tion on a pre-specified layer from the source model. Ganesan *et al.* [7] proposed a principled manner to inflict severe damage to feature representation, hence obtaining higher transferability. Our proposed method also falls into this category, and the key difference is that our method considers the feature importance and disrupts the critical object-aware features that dominate the decision across different models, while existing methods distort features indiscriminately.

3. Preliminaries

Assume a classification model $f_\theta : x \mapsto y$, where x and y denotes the clean image and true label, respectively, and θ indicates the parameters of the model. We aim to generate an adversarial example $x^{adv} = x + \epsilon$, which is distorted by carefully designed perturbation ϵ but will mislead the classifier, *i.e.*, $f_\theta(x^{adv}) \neq y$. Typically, ℓ_p -norm is commonly adopted to regularize the perturbation. Therefore, the generation of adversarial examples can be formulated as an optimization problem as shown in the following.

$$\arg \max_{x^{adv}} J(x^{adv}, y), \text{ s.t. } \|x - x^{adv}\|_p \leq \epsilon, \quad (1)$$

where the loss function $J(\cdot, \cdot)$ measures the distance between true and predicted labels (*i.e.*, cross-entropy), and $p = \infty$ in this work. Many methods have been proposed to solve the above optimization problem, *e.g.*, Fast Gradient Sign Method (FGSM) [10], Basic Iterative Method (BIM) [18], Momentum Iterative Method (MIM) [5], *etc.* However, optimizing Eq. 1 requires explicit access to the parameters of f_θ , while this is impractical in black-box attacking. Therefore, a feasible solution is to optimize on an analogous model f_ϕ (*i.e.*, the source model) with accessible parameters ϕ , thus generating highly transferable adversarial examples to attack the target model f_θ .




Figure 3. Visualization of feature maps and corresponding gradients at the layer Conv3_3 of VGG16. Raw features and gradients are obtained from the original images, and aggregate features and gradients are obtained from multiple transforms (*i.e.*, random pixel dropping) of the original images.

4. Feature Importance-aware Attack

Empirical studies from most DNN-based classifiers, *e.g.*, Inception [29] and ResNet [11], have shown that deep models tend to extract semantic features, which are object-aware discriminative and thus effectively boost classification accuracy. Intuitively, disrupting those object-aware features that dominate the decision of all models could benefit the transferability of adversarial examples. However, different networks also extract exclusive features to better fit themselves to the data domain, which results in model-specific feature representation. Without awareness of such characteristics, existing adversarial attacks tend to craft adversarial examples by indiscriminately distorting features against a source model, thus trapped into model-specific local optimum and significantly degrading the transferability. Therefore, avoiding such local optimum is a key to transferability. More specifically, the generation of adversarial examples needs to be guided by model-agnostic critical features from the source model, which is referred to as feature importance. Fig. 2 overviews the proposed feature importance-aware transferable attack (FIA), where aggregate gradient (detailed in Section 4.1) can effectively avoid local optimum and represent transferable feature importance. Then,




Figure 4. Illustration of the aggregate gradient. The gradients are obtained from multiple random masked images, and the final aggregate gradient (*i.e.*, feature importance) is represented by averaging these gradients.

aggregate gradient serves as weights in the optimization to distort important features (discussed in Section 4.2), *i.e.*, reducing features with positive weights and increasing those corresponding to negatives.

4.1. Feature Importance by Aggregate Gradient

For simplicity, let f denote the source model, and the feature maps from the k -th layer are expressed as $f_k(x)$. Since feature importance is proportional to how the features contribute to the final decision, an intuitive strategy is to obtain the gradient w.r.t. $f_k(x)$ as written in the following.

$$\Delta_k^x = \frac{\partial l(x, t)}{\partial f_k(x)}, \quad (2)$$

where $l(\cdot, \cdot)$ denotes the logit output with respect to the true label t . However, the raw gradient Δ_k^x would carry model-specific information. As shown in Fig. 3, the raw gradient maps and raw feature maps are both visually noisy, *i.e.*, pulses and large gradient on non-object regions, which may be caused by the model-specific solution space.

To suppress model-specific information, we propose the aggregate gradient, which aggregates gradients from randomly transformed x as shown in Fig. 4. The transformation is supposed to distort image details but preserve the spatial structure and general texture. Since semantically object-aware/important features/gradients are robust to such transformation but model-specific ones are vulnerable to the transforms, those robust/transerable features/gradients will be highlighted after aggregation, while the others would be neutralized. In this paper, we adopt random pixel dropping (*i.e.*, randomly mask) with the probability p_d . Therefore, the aggregate gradient can be expressed in the following.

$$\bar{\Delta}_k^x = \frac{1}{C} \sum_{n=1}^N \Delta_k^{x \odot M_{p_d}^n}, \quad M_{p_d} \sim \text{Bernoulli}(1 - p_d), \quad (3)$$

where the M_{p_d} is a binary matrix with the same size as x , and \odot denotes the element-wise product. The normalizer C is obtained by ℓ_2 -norm on the corresponding summation term. The ensemble number N indicates the number

Algorithm 1: Feature Importance-aware Attack

Input: The original clean image x , classification model f , intermediate layer k , drop probability p_d and ensemble number N in aggregate gradient, max perturbation ϵ , and the number of iteration T .

Output: The adversarial image x^{adv} .

Initialize: $\Delta = 0, g_0 = 0, \mu = 1, \alpha = \epsilon/T$

Obtain aggregate gradient:

for $n = 0$ to $N - 1$ **do**

$$\Delta = \Delta + \Delta_k^{x \odot M_{p_d}^n}$$

$$\Delta = \frac{\Delta}{\|\Delta\|_2}$$

Construct optimization objective:

$$\mathcal{L}(x^{adv}) = \sum (\Delta \odot f_k(x^{adv}))$$

Update x^{adv} by momentum iterative method:

for $t = 0$ to $T - 1$ **do**

$$g_{t+1} = \mu \cdot g_t + \frac{\nabla_x \mathcal{L}(x_t^{adv})}{\|\nabla_x \mathcal{L}(x_t^{adv})\|_1}$$

$$x_{t+1}^{adv} = \text{Clip}_{x, \epsilon} \{x_t^{adv} - \alpha \cdot \text{sign}(g_{t+1})\}$$

return x_T^{adv}

of random masks applied to the input x . Fig. 4 illustrates the process of gradient aggregation. The aggregate gradient $\bar{\Delta}_k^x$ highlights the regions of robust and critical object-aware features that can guide the adversarial example towards the more transferable direction. Fig. 3 visualizes the aggregate gradient, which is cleaner and focuses more on the objects as compared to the raw gradient, providing better feature importance in the transferable perspective.

4.2. Attack Algorithm

Utilizing the aforementioned feature importance (*i.e.*, aggregate gradient $\bar{\Delta}_k^x$), we design the loss function Eq. 4 to guide the generation of adversarial example x^{adv} by explicitly suppress important features. For simplicity, we denote $\bar{\Delta}_k^x$ as Δ in the rest of this paper.

$$\mathcal{L}(x^{adv}) = \sum (\Delta \odot f_k(x^{adv})). \quad (4)$$

Intuitively, the important features will yield relatively higher intensity in Δ , which indicates the efforts of correcting the features to approach the true label, and the sign of Δ provides the correcting direction. The objective of generating transferable adversarial examples is to decrease important features with positive Δ and increase those corresponding to negative Δ . Therefore, it is straightforward to achieve this goal by minimizing Eq. 4. Finally, substitute the Eq. 4 into Eq. 1, we get the proposed objective for

Table 1. Success rate of different attacks against normally trained models. The first column shows source models, and the first row lists target models. FIA is our method and FIA+PIDIM is the combination of FIA and PIDIM. “*” indicates white-box attack since the target model is the source model, and the best results are highlighted in bold.

	Attack	Inc-v3	Inc-v4	IncRes-v2	Res-50	Res-152	Vgg-16	Vgg-19
Inc-v3	MIM	100.0%*	41.6%	38.8%	33.1%	29.7%	38.6%	38.3%
	DIM	99.6%*	64.6%	59.6%	40.7%	36.3%	47.6%	46.4%
	PIM	97.9%*	55.8%	51.5%	53.3%	46.3%	61.6%	60.5%
	PIDIM	98.1%*	70.5%	66.4%	61.8%	56.3%	57.7%	56.0%
	NRDM	98.2%*	68.1%	59.9%	47.5%	37.1%	49.9%	50.5%
	FDA	98.7%*	71.9%	66.3%	48.4%	37.1%	51.9%	53.0%
	FIA	98.3%*	83.5%	80.6%	70.4%	64.9%	71.4%	73.3%
	FIA+PIDIM	98.8%*	87.8%	85.7%	79.7%	74.4%	82.4%	84.1%
IncRes-v2	MIM	60.2%	52.5%	99.3%*	40.1%	36.1%	46.9%	43.8%
	DIM	75.2%	71.3%	97.1%*	50.9%	43.7%	51.5%	51.4%
	PIM	66.8%	62.9%	99.6%*	56.2%	50.8%	64.4%	63.5%
	PIDIM	80.5%	78.0%	98.5%*	56.6%	50.1%	62.5%	62.6%
	NRDM	65.0%	55.5%	75.3%*	53.0%	44.0%	49.4%	49.1%
	FDA	64.3%	53.7%	76.0%*	50.2%	41.0%	53.5%	53.9%
	FIA	81.1%	77.5%	89.2%*	71.8%	68.9%	71.4%	71.4%
	FIA+PIDIM	84.2%	79.7%	91.6%*	79.0%	78.4%	80.6%	79.9%
Res-152	MIM	57.2%	48.2%	45.7%	90.6%	99.8%*	72.8%	72.9%
	DIM	80.3%	72.2%	72.6%	95.0%	99.9%*	88.4%	88.0%
	PIM	66.0%	56.4%	51.1%	92.3%	100.0%*	83.2%	82.5%
	PIDIM	82.2%	76.6%	77.0%	96.7%	99.8%*	91.2%	89.9%
	NRDM	64.5%	59.1%	51.2%	87.7%	95.4%*	79.3%	79.3%
	FDA	60.7%	52.3%	48.0%	85.0%	95.3%*	75.0%	75.0%
	FIA	85.3%	81.1%	77.8%	96.8%	99.5%*	91.5%	91.5%
	FIA+PIDIM	90.3%	85.9%	85.6%	98.2%	99.5%*	95.8%	95.7%
Vgg-16	MIM	80.3%	81.1%	74.6%	89.3%	84.4%	100.0%*	96.5%
	DIM	87.2%	87.0%	80.9%	92.0%	87.8%	99.8%*	98.9%
	PIM	84.1%	82.0%	75.6%	91.1%	85.9%	100.0%*	98.9%
	PIDIM	89.1%	89.5%	84.7%	93.8%	90.8%	99.9%*	98.8%
	NRDM	73.6%	72.8%	57.1%	77.5%	73.0%	93.2%*	91.1%
	FDA	76.1%	76.7%	64.0%	81.7%	78.7%	95.7%*	95.7%
	FIA	95.7%	95.6%	92.3%	97.3%	95.3%	99.8%*	99.6%
	FIA+PIDIM	97.6%	97.5%	93.8%	98.2%	96.4%	99.8%*	99.8%

feature importance-aware transferable adversarial attack.

$$\arg \min_{x^{adv}} \mathcal{L}(x^{adv}), \text{ s.t. } \|x - x^{adv}\|_{\infty} \leq \epsilon. \quad (5)$$

There are many existing gradient-based attack methods aiming to solve the above objective function Eq. 5, e.g., BIM [18], MIM [5], etc. Given the superior performance of MIM, we adopt this method to solve Eq. 5, and the details are shown in Algorithm 1.

4.3. Drawback of Related Attacks

It is worth to further emphasize the advantage of feature importance awareness over related feature-based attacks, i.e., NRDM [24] and FDA [7]. For better illustration, their loss functions are written in Eq. 6 and 7, respectively

$$\mathcal{L}_{NRDM} = \|f_k(x^{adv}) - f_k(x)\|_2, \quad (6)$$

where ℓ_2 -norm is adopted to simply measure the feature distortion, which indiscriminately disturbs the features.

$$\mathcal{L}_{FDA} = \log(\mathcal{D}(f_k(x^{adv}) | f_k(x) < C_k(h, w))) - \log(\mathcal{D}(f_k(x^{adv}) | f_k(x) > C_k(h, w))) \quad (7)$$

where $\mathcal{D}(\cdot)$ is the ℓ_2 -norm, and $C_k(h, w)$ represents the mean activation values across channels.

From the objective functions, NRDM simply optimizes the feature distortion between original images and adversarial images without any constraints. For FDA, although it introduced a similar idea of utilizing feature activation to guide the optimization, i.e., features that support the ground truth should be suppressed, while those that do not support the ground truth should be enhanced. However, FDA uses the mean across channels as the distinguishable criterion, which cannot effectively avoid model-specific information. By contrast, the proposed FIA provides more intrinsic feature importance by the aggregate gradient, and hence achieving higher transferable adversarial examples. Quantitative comparison in the experimental evaluations demonstrate the superior performance of the proposed FIA.

5. Experimental Evaluation

5.1. Experiment Setup

Dataset: For fair comparison, we follow the previous works [6, 8] to use the ImageNet-compatible dataset [1],

Table 2. Success rate of different attacks against defense models. The first column shows source models, and the first row lists target models. FIA is our method and FIA+PITIDIM is the combination of FIA and PITIDIM. The best results are highlighted in bold.

	Attack	Adv-Inc-v3	Adv-IncRes-v2	Ens3-Inc-v3	Ens4-Inc-v3	Ens-IncRes-v2
Inc-v3	MIM	22.9%	17.5%	15.4%	15.8%	7.8%
	DIM	26.0%	24.5%	17.8%	20.8%	10.0%
	TIM	32.0%	26.4%	30.1%	32.5%	22.4%
	PIM	34.3%	30.2%	33.3%	38.4%	26.2%
	TIDIM	40.7%	37.1%	40.8%	42.3%	30.4%
	PITIDIM	41.6%	33.9%	43.1%	47.3%	31.4%
	FIA	54.5%	54.9%	43.9%	42.0%	23.5%
	FIA+PITIDIM	64.8%	59.0%	62.5%	63.2%	50.9%
IncRes-v2	MIM	25.5%	29.9%	21.4%	22.7%	12.5%
	DIM	33.1%	42.9%	30.5%	29.7%	19.0%
	TIM	40.0%	43.5%	39.5%	41.5%	38.4%
	PIM	39.0%	35.3%	39.4%	42.2%	32.8%
	TIDIM	50.4%	55.7%	50.1%	49.5%	48.1%
	PITIDIM	53.8%	55.2%	54.7%	54.5%	50.6%
	FIA	54.9%	56.8%	46.9%	44.7%	37.4%
	FIA+PITIDIM	55.1%	52.9%	54.9%	56.2%	50.6%
Res-152	MIM	36.9%	34.8%	36.2%	37.4%	22.0%
	DIM	54.3%	54.6%	53.3%	50.4%	33.5%
	TIM	41.5%	37.5%	43.1%	47.6%	34.1%
	PIM	40.7%	38.9%	46.9%	51.8%	38.8%
	TIDIM	52.4%	48.6%	57.5%	61.1%	46.3%
	PITIDIM	51.9%	49.0%	58.6%	64.8%	47.9%
	FIA	70.1%	66.7%	61.4%	60.3%	41.7%
	FIA+PITIDIM	66.3%	62.5%	69.6%	72.7%	61.4%
Vgg-16	MIM	64.3%	61.1%	64.3%	64.3%	45.0%
	DIM	69.9%	66.2%	70.3%	67.8%	49.9%
	TIM	52.8%	46.2%	55.1%	55.3%	41.6%
	PIM	51.9%	43.2%	50.2%	56.3%	39.9%
	TIDIM	59.1%	48.2%	59.6%	60.3%	47.9%
	PITIDIM	51.0%	44.6%	55.6%	60.7%	43.1%
	FIA	87.8%	86.3%	85.6%	86.0%	70.8%
	FIA+PITIDIM	74.7%	71.4%	77.3%	80.1%	67.0%

which consists of 1000 images and was used for NIPS 2017 adversarial competition.

Target Models: The proposed FIA is validated on twelve state-of-the-art classification models, including seven normally trained models* and five adversarially trained models† (i.e., defense models). The normally trained models are Inception-V3 (Inc-v3) [29], Inception-V4 (Inc-v4) [28], Inception-ResNet-V2 (IncRes-v2) [28], ResNet-V1-50 (Res-50) [11], ResNet-V1-152 (Res-152) [11], VGG16 (Vgg-16) [26], and VGG19 (Vgg-19) [26]. With adversarial training [19, 32], the corresponding defense models are Adv-Inc-v3, Adv-IncRes-v2, Ens3-Inc-v3, Ens4-Inc-v3, and Ens-IncRes-v2.

Baseline Attacks: To demonstrate the effectiveness of the proposed FIA, we compare it to diverse state-of-the-art attack methods, e.g., MIM [5], DIM [35], TIM [6], PIM [8], as well as combined versions of these methods i.e., TIDIM [6], PIDIM [8] and PITIDIM [8]. In addition, recent feature-level attacks are also involved, i.e., NRDM [24], FDA [7].

Parameter Settings: In all experiments, the maximum perturbation $\epsilon = 16$, the iteration $T = 10$, and the step size

$\alpha = \epsilon/T = 1.6$ (recap Algorithm 1). For the baseline attacks, the transform probability is 0.7 in DIM, and the kernel size is 15 for TIM. Since settings of PIM will vary with target models and the ways of method combination, we will specifically detail its settings (i.e., the amplification factor β , project factor γ , and project kernel size k_w) in each related experiment. In the proposed FIA, the drop probability $p_d = 0.3$ when attacking normally trained models and $p_d = 0.1$ when attacking defense models, and the ensemble number $N = 30$ in aggregate gradient. For feature-level attacks, we choose the same layer, i.e., *Mixed_5b* for Inc-V3, *Conv3_3* for Vgg-16, *Conv_4a* for InRes-V2, and the last layer of second block for Res-152.

5.2. Comparison of Transferability

To quantitatively compare the transferability between the proposed FIA and the baselines, we choose Inc-v3, IncRes-v2, Res-152, Vgg-16 as the source model, respectively, and attack the other normally trained models (Table 1) and defense models (Table 2). Please note that TIM is not included in Table 1 because it is designed for defense models. For qualitative comparison, please refer to the supplementary.

Attacking Normally Trained Models. We follow the set-

* <https://github.com/tensorflow/models/tree/master/research/slim>

† https://github.com/tensorflow/models/tree/archive/research/adv_imagenet_models

Table 3. Success rate of different attacks against defense models when using an ensemble model which contains Res-50, Res-152, Vgg-16, and Vgg-19. The best results are highlighted in bold.

Attack	Adv- Inc-v3	Adv- IncRes-v2	Ens3- Inc-v3	Ens4- Inc-v3	Ens- IncRes-v2
MIM	71.0%	69.6%	70.6%	71.3%	51.7%
DIM	81.2%	83.6%	81.6%	79.8%	65.9%
TIM	68.0%	63.9%	70.7%	72.7%	60.7%
PIM	72.1%	66.4%	76.3%	79.3%	68.2%
TIDIM	74.4%	68.3%	75.9%	77.8%	67.1%
PITIDIM	72.1%	66.4%	76.3%	79.3%	68.2%
FIA	90.9%	90.0%	88.0%	88.4%	75.8%

tings in [8], *i.e.*, $\beta = 10$, $\gamma = 16$ for PIM, and $\beta = 2.5$, $\gamma = 2$ for PIDIM which is the combined version of PIM and DIM. The project kernel size $k_w = 3$ for the two methods. As shown in Table 1, our method significantly outperforms the other methods for the transferable attack, improving the success rate by 9.5% on average. In particular, the success rate of FIA against each of the normally trained models of Vgg-16, while the other methods may drop to around 60% in the transferable attack. Our method can be easily adapted to other methods to further improve the transferability, *e.g.*, FIA+PIDIM is the combination of FIA and PIDIM ($\beta = 2.5$, $\gamma = 2$ and $k_w = 3$), which results in 1% \sim 10% increase of success rate. Finally, compared to other feature-level attacks, *i.e.*, NRDM and FDA, FIA performs the best in ALL cases (white-box attack and black-box attack), which demonstrates the effectiveness of the proposed aggregate gradient in locating critical object-aware features that dominate the decision across different models.

The results in Table 1 also give insight into relation between model complexity and transferability, *i.e.*, less complicated/smaller models tend to yield higher transferable adversarial examples (on the premise that the models should achieve similar classification accuracy). For instance, it gets the highest success rate when using Vgg-16 as the source model. Intuitively, those larger models (*e.g.*, IncRes-v2 and Res-152) provide a larger search space, making it more difficult to avoid local optimum.

Attacking Defense Models Since defense models are trained adversarially thus showing strong robustness to adversarial examples. In Table 2, the other feature-level attacks are not listed because of their bad performance. In PIM, we follow the suggestion from [8] to remove the momentum term since it may hinder the performance of attacking defense models. The settings of PIM and its combinations are $\beta = 10$, $\gamma = 16$, and $k_w = 7$. As shown in Table 2, FIA or the correspondingly combined version FIA+PITIDIM ($\beta = 2.5$, $\gamma = 2.0$, and $k_w = 7$) outperforms other methods. In most cases, FIA+PITIDIM and FIA rank the top two, and FIA+PITIDIM would further improve the transferability as compared to FIA. In average, our approach increases the success rate by about 12.8% in

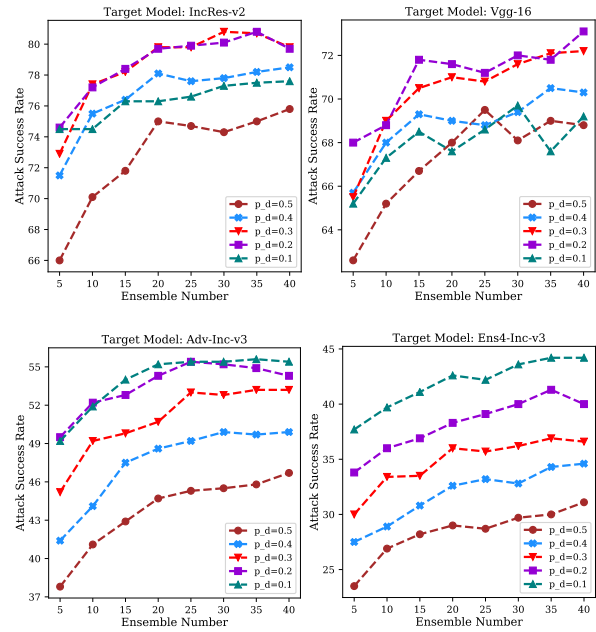


Figure 5. Effect of the drop probability and ensemble number on attack success rate. The adversarial examples are generated by FIA with different parameter setting against the source model Inc-v3. The drop probability changes from 0.1 to 0.5 with the step of 0.1, and the ensemble number varies from 5 to 40 with the step of 5. The top row are success rates of attacking two normally trained models IncRes-v2 and Vgg-16, and the bottom are results of attacking two defense models Adv-Inc-v3 and Ens4-Inc-v3.

attacking defense models as compared to the other methods.

Since the relatively lower transferability in attacking defense models, we further improve transferability by generating adversarial examples on an ensemble of models [21], which prevents adversarial examples from falling into local optimum of a single model. Following the setting in [8], *i.e.*, $\beta = 5$ and $\gamma = 8$ for PIM and its combined versions, the results are shown in Table 3, where all methods are improved, and our method still outperforms the others.

5.3. Effect of Parameters in Aggregate Gradient

There are three parameters, *i.e.*, drop probability p_d , ensemble number N , and layer k (Eq. 3), which affects the performance of the proposed FIA. For the first two parameters, we adopt Inc-v3 as the source model and modify p_d from 0.1 to 0.5 with the step of 0.1. For each p_d , N is iterated from 5 to 40 with the step of 5. Fig. 5 illustrates the effect of p_d and N by attacking IncRes-v2, Vgg-16, Adv-Inc-V3, and Ens4-Inc-v3, where the drop probability and ensemble number affect the success rate in a roughly consistent manner across different target models.

More specifically, the drop probability p_d plays an important role in affecting the success rate, and such effect tends to be consistent across different target models. A large p_d (*e.g.*, 0.5) will destroy important structural information




Figure 6. Effect of layer choice on attack success rate. Different layers from the source models (*i.e.*, Inc-v3 and Vgg-16) are selected to generate adversarial examples, whose success rates are reported against different target models.

of an image, thus significantly reducing the success rate. Therefore, an optimal p_d for attacking normally trained models is between 0.2 and 0.3, and it should be around 0.1 if attacking defense models. For the ensemble number N , a larger N tends to yield a higher success rate but will saturate gradually. Finally, we determine the ensemble number $N = 30$, drop probability $p_d = 0.3$ against normally trained models, and $p_d = 0.1$ against defense models.

Feature-level attacks are significantly affected by the choice of feature layer k since early layers of DNNs may be working to construct a basic feature set which is usually data-specific, and further layers may process these extracted features to maximize the classification accuracy of the model which makes the features become model-specific [16]. Therefore, early layers have not learned salient features and semantic concepts of the true classes, and later layers are model-specific that should be avoided in transferable attacks. By contrast, middle layers have well-separated class representations and they are not highly correlated to the model architecture, thus middle layers are the best choice to be attacked for better transferability. The same conclusion can be drawn from Fig. 6, which reports the success rate of attacking different target models by adversarial examples optimized on different source model layers. Based on this conclusion, we first select a few middle layers for each source model and determine the final attacked layer according to the empirical results.

5.4. Ablation Study

The key of the proposed FIA is the aggregate gradient Δ , which significantly boosts the transferability as demonstrated in the aforementioned results. To highlight the contribution of aggregate gradient, we conduct the ablation study to compare the performance of objectives with and without aggregate gradient. We construct three objective functions as express in the following, where \mathcal{L}_1 optimizes the feature divergence without constraints like most of the baseline methods, and \mathcal{L}_3 is equivalent to our proposed loss (Eq. 4). The \mathcal{L}_2 uses non-aggregate gradient Δ_{clean} , *i.e.*,




Figure 7. Effect of aggregate gradient on attack success rate. \mathcal{L}_1 optimizes feature distortion without gradient guidance, \mathcal{L}_2 uses raw gradient, and \mathcal{L}_3 adopts aggregate gradient.

gradient from the original clean image. Fig. 7 shows the success rate using the three losses, respectively.

$$\mathcal{L}_1 = \sum |f_k(x) - f_k(x^{adv})|, \quad (8)$$

$$\mathcal{L}_2 = \sum (\Delta_{clean} \odot (f_k(x) - f_k(x^{adv}))), \quad (9)$$

$$\mathcal{L}_3 = \sum (\Delta \odot (f_k(x) - f_k(x^{adv}))). \quad (10)$$

The proposed loss \mathcal{L}_3 outperforms the others by a large margin in all case, indicating effectiveness of the proposed aggregate gradient.

6. Conclusion

In this work, we proposed a Feature Importance-aware Attack (FIA) to generate highly transferable adversarial examples, whose effectiveness are demonstrated in our exhaustive experiments. The proposed FIA explores the feature importance through aggregate gradient across various of classification models and introduces such transferable information into the search of adversarial examples. Consequently, the optimization process is guided towards disrupting the critical object-aware features that dominate the decision of the models, thus gaining remarkable transferability. We conducted extensive experiments to demonstrate the superior performance of FIA as compared to those state-of-the-art methods, and our method can serve as a benchmark for evaluating the robustness of various models.

Acknowledgments

This work was supported by National Natural Science of China (Grants No. 62122066, U20A20182, 61872274, U20A20178, 62032021, and 62072395), National Key R&D Program of China (Grant No. 2020AAA0107705), the Fundamental Research Funds for the Central Universities (Grant No. 2021QNA5016).

References

- [1] Nips17 adversarial attacks and defenses competition. https://github.com/cleverhans-lab/cleverhans/tree/master/cleverhans_v3.1.0/examples/nips17_adversarial_competition. 5
- [2] Wieland Brendel, Jonas Rauber, and Matthias Bethge. Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2018. 1, 2
- [3] Nicholas Carlini and David A. Wagner. Towards evaluating the robustness of neural networks. In *Proceedings of IEEE Symposium on Security and Privacy (SP)*, pages 39–57, 2017. 1, 2
- [4] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, 2018. 1
- [5] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9185–9193, 2018. 2, 3, 5, 6
- [6] Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Evading defenses to transferable adversarial examples by translation-invariant attacks. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4312–4321, 2019. 2, 5, 6
- [7] Aditya Ganesan, Vivek B. S., and Venkatesh Babu Radhakrishnan. FDA: feature disruptive attack. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)*, pages 8068–8078, 2019. 2, 3, 5, 6
- [8] Lianli Gao, Qilong Zhang, Jingkuan Song, Xianglong Liu, and Heng Tao Shen. Patch-wise attack for fooling deep neural network. In *Proceedings of European Conference on Computer Vision (ECCV)*, volume 12373, pages 307–322, 2020. 2, 5, 6, 7
- [9] Ross B. Girshick. Fast R-CNN. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)*, pages 1440–1448, 2015. 1
- [10] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2015. 1, 2, 3
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 3, 6
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *Proceedings of European Conference on Computer Vision (ECCV)*, pages 630–645, 2016. 1
- [13] Qian Huang, Isay Katsman, Zeqi Gu, Horace He, Serge J. Belongie, and Ser-Nam Lim. Enhancing adversarial example transferability with an intermediate level attack. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)*, pages 4732–4741, 2019. 2
- [14] Andrew Ilyas, Logan Engstrom, Anish Athalye, and Jessy Lin. Black-box adversarial attacks with limited queries and information. In *Proceedings of International Conference on Machine Learning (ICML)*, pages 2142–2151, 2018. 1, 2
- [15] Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. In *Proceedings of Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 125–136, 2019. 2
- [16] Nathan Inkawhich, Kevin J. Liang, Lawrence Carin, and Yiran Chen. Transferable perturbations of deep feature distributions. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2020. 8
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Proceedings of Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 1106–1114, 2012. 1
- [18] Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. Adversarial examples in the physical world. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2017. 1, 2, 3, 5
- [19] Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. Adversarial machine learning at scale. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2017. 6
- [20] Jiadong Lin, Chuanbiao Song, Kun He, Liwei Wang, and John E. Hopcroft. Nesterov accelerated gradient and scale invariance for adversarial attacks. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2020. 2
- [21] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into transferable adversarial examples and black-box attacks. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2017. 2, 7
- [22] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3431–3440, 2015. 1
- [23] 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 *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 937–946, 2020. 2
- [24] Muzammal Naseer, Salman H Khan, Shafin Rahman, and Fatih Porikli. Task-generalizable adversarial attack based on perceptual metric. *arXiv preprint arXiv:1811.09020*, 2018. 2, 5, 6
- [25] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In *Proceedings of Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 91–99, 2015. 1

- [26] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2015. [6](#)
- [27] Yi Sun, Xiaogang Wang, and Xiaoou Tang. Deeply learned face representations are sparse, selective, and robust. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2892–2900, 2015. [1](#)
- [28] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A. Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Proceedings of AAAI Conference on Artificial Intelligence*, pages 4278–4284, 2017. [6](#)
- [29] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, 2016. [3](#), [6](#)
- [30] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2014. [1](#), [2](#)
- [31] Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1701–1708, 2014. [1](#)
- [32] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian J. Goodfellow, Dan Boneh, and Patrick D. McDaniel. Ensemble adversarial training: Attacks and defenses. In *Proceedings of International Conference on Learning Representations (ICLR)*, 2018. [6](#)
- [33] Jonathan Uesato, Brendan O’Donoghue, Pushmeet Kohli, and Aaron van den Oord. Adversarial risk and the dangers of evaluating against weak attacks. In *Proceedings of International Conference on Machine Learning (ICML)*, pages 5032–5041, 2018. [1](#), [2](#)
- [34] Chaowei Xiao, Bo Li, Jun-Yan Zhu, Warren He, Mingyan Liu, and Dawn Song. Generating adversarial examples with adversarial networks. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3905–3911, 2018. [1](#)
- [35] Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, and Alan L. Yuille. Improving transferability of adversarial examples with input diversity. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2730–2739, 2019. [1](#), [2](#), [6](#)
- [36] Wen Zhou, Xin Hou, Yongjun Chen, Mengyun Tang, Xiangqi Huang, Xiang Gan, and Yong Yang. Transferable adversarial perturbations. In *Proceedings of European Conference on Computer Vision (ECCV)*, pages 471–486, 2018. [2](#)