

# Cross-Modal Transferable Adversarial Attacks from Images to Videos

Zhipeng Wei<sup>1,2</sup>, Jingjing Chen<sup>1,2†</sup>, Zuxuan Wu<sup>1,2</sup>, Yu-Gang Jiang<sup>1,2</sup>

<sup>1</sup>Shanghai Key Lab of Intelligent Information Processing,  
School of Computer Science, Fudan University

<sup>2</sup>Shanghai Collaborative Innovation Center on Intelligent Visual Computing

zpwei21@m.fudan.edu.cn, {chenjingjing, zxwu, ygj}@fudan.edu.cn

## Abstract

Recent studies have shown that adversarial examples hand-crafted on one white-box model can be used to attack other black-box models. Such cross-model transferability makes it feasible to perform black-box attacks, which has raised security concerns for real-world DNNs applications. Nevertheless, existing works mostly focus on investigating the adversarial transferability across different deep models that share the same modality of input data. The cross-modal transferability of adversarial perturbation has never been explored. This paper investigates the transferability of adversarial perturbation across different modalities, i.e., leveraging adversarial perturbation generated on white-box image models to attack black-box video models. Specifically, motivated by the observation that the low-level feature space between images and video frames are similar, we propose a simple yet effective cross-modal attack method, named as Image To Video (I2V) attack. I2V generates adversarial frames by minimizing the cosine similarity between features of pre-trained image models from adversarial and benign examples, then combines the generated adversarial frames to perform black-box attacks on video recognition models. Extensive experiments demonstrate that I2V can achieve high attack success rates on different black-box video recognition models. On Kinetics-400 and UCF-101, I2V achieves an average attack success rate of 77.88% and 65.68%, respectively, which sheds light on the feasibility of cross-modal adversarial attacks.

## 1. Introduction

Deep learning has achieved significant progress in a series of computer vision tasks, such as image recognition [3, 4, 12], video classification [41], object detection [24], and action recognition [2, 11], etc. However, recent studies have shown that deep neural networks (DNNs)

are highly vulnerable to adversarial examples [21, 28, 29], which are generated by adding small human-imperceptible perturbations that can lead to wrong predictions. The existence of adversarial examples has posed serious security threats for the application of DNNs, such as autonomous driving [25], face recognition [8], video analysis [5, 31, 34], etc. As a result, adversarial examples have attracted numerous research attentions in recent years.

It has been demonstrated in recent works [22, 33] that adversarial examples have the property of transferability, i.e., the adversarial example generated from one model can be used to attack other models. Such cross-model transferability makes it feasible to perform black-box attacks by leveraging adversarial examples handcrafted on white-box models. As a result, how to enhance the transferability of adversarial examples for efficient black-box attacks has attracted several research interests recently. These works either perform data augmentation [7, 20, 38], optimize gradient calculations [6, 20, 36], or disrupt common properties among different models [37], to avoid the generated adversarial samples being overfitted to white-box models. Nevertheless, all of these works require the white-box models and target black-box models to be homomodal, which share the same modal of input data. Transferability between heteromodal models has never been explored.

To bridge this gap, this paper investigates the cross-modal transferability of adversarial examples. Specifically, we explore the adversarial transferability between image models and video models by performing transfer-based black-box attacks on video models with image models pre-trained on ImageNet only. This is an extremely challenging setting since there are no white-box video models for generating video adversarial samples. There are two major obstacles in transferring adversarial perturbations generated on images models to attack video models. First, in addition to the domain gap between image and video data, video data contain additional temporal information, which leads to differences in the learned features between image models and video models. The difference makes it diffi-

<sup>†</sup>Correspondence to: Jingjing Chen.

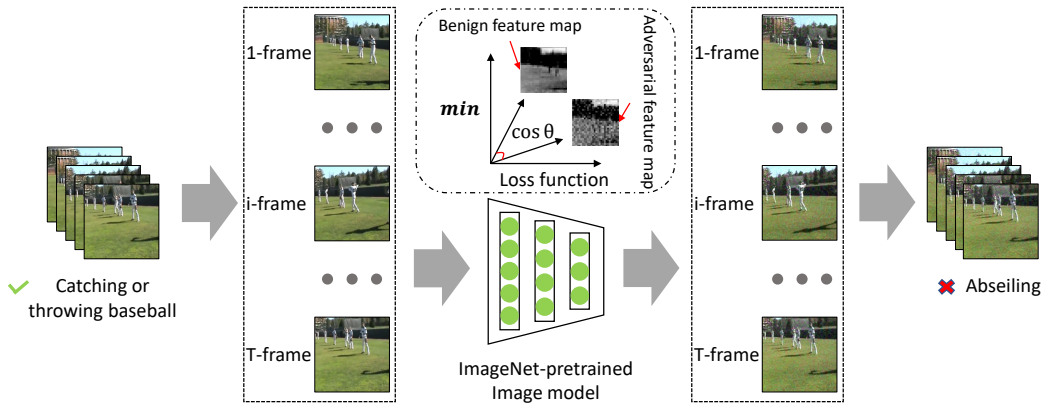


Figure 1. Overview of the proposed I2V attack. Given a video clip with a true label of “Catching or throwing baseball”, where each frame is input into the ImageNet-pretrained image model separately. Then the image model generates adversarial frames by minimizing the cosine similarity between features from adversarial and benign examples. As the image model and the video model share similar feature space, the generated video adversarial example can fool the video recognition models, and be misclassified as “Abseiling”.

cult to transfer the adversarial perturbations from images to videos. Second, existing transfer-based attacks on homomodal models (e.g., image models) are not applicable to the cross-modal attack scenarios. Unlike existing transfer-based attacks on images, where image labels are available for optimizing task-specific loss function (e.g., Cross Entropy loss) in the process of adversarial perturbations generation, in the cross-modal image to video attack, no labels are available for video frames.

To address the aforementioned challenges and perform black-box attacks on video models, we propose a simple yet effective cross-modal attack method, named Image To Video (I2V) attack. Despite there being a domain gap between image data and video data, we observed that the intermediate features between images models and video models are similar to a certain extent. This motivates us to perturb the intermediate features of the ImageNet-pretrained image models to craft adversarial video frames for attacking video recognition models. To this end, the proposed I2V optimizes the adversarial perturbations by minimizing the cosine similarity of intermediate features between benign frames and the generated adversarial frames. The minimization of the cosine similarity makes the features extracted from adversarial video frames orthogonal to that from benign frames. Consequently, it will cause the adversarial video features to move away from the benign video features due to the feature similarity between image and video models. Figure 1 gives an overview of the proposed I2V attack method. I2V takes individual frames from a video clip as input to the image model and generates adversarial frames one by one. Then the generated adversarial frames are grouped into video adversarial examples according to the temporal information of the benign video clip. We briefly summarize our primary contributions as follows:

- We investigate the transferability of adversarial perturbations between image models and video models. Specially, we propose an I2V attack to boost the transferability of video adversarial examples generated from image models across different video recognition models. To the best of our knowledge, this is the first work on cross-modal transfer-based black-box attacks for video recognition models.
- We provide insightful analysis on the correlations of feature maps between image and video models. Based on this observation, I2V optimizes adversarial frames on perturbing feature maps of image models to boost the transferability across different video recognition models.
- We conduct empirical evaluations using six video recognition models trained with the Kinetic-400 dataset and UCF-101 dataset. Extensive experiments demonstrate that our proposed I2V helps to boost the transferability of video adversarial examples generated from image models.

## 2. Related Work

### 2.1. Transfer-based Attacks on Image Models

Prior works in generating adversarial examples with high transferability are based on white-box attacks, such as Fast Gradient Sign Method (FGSM) [10] and Basic Iterative Method (BIM) [19]. FGSM linearizes the loss function around the current parameters and performs a one-step update along with the gradient sign of the loss function with respect to inputs. BIM is the iterative version of FGSM and overfits the white-box model to generate stronger adversarial examples in attacking the white-box model. To

further improve the transferability of adversarial examples in attacking black-box models, several approaches are proposed recently. In general, there are three ways to improve transferability, which include data augmentation, gradient modification and common property disruption of discriminative features among different models. The main idea of data augmentation is to improve the generalization of adversarial examples and avoid overfitting the white-box model. For example, Diversity Input (DI) [38] attack conducts random resizing and padding to the input. Scale-invariant method (SIM) [20] applies the scale transformation to the input. Translation-invariant (TI) [7] attack performs horizontal and vertical shifts with a short distance to the input. The second way modifies gradients used for updating adversarial perturbations. For example, Momentum Iterative (MI) [6] attack integrates the momentum into the iterative process for stabilizing update directions. As an improved momentum method, Nesterov Accelerated Gradient (NAG) [20] can be also integrated into the BIM. Skip Gradient Method (SGM) [36] uses more gradients from the skip connections and emphasizes the gradients of shallow layers. The main idea of the third way focuses on disrupting the common property of classification among different models. For example, Attention-guided Transfer Attack (ATA) [37] prioritizes the corruption of critical features that are likely to be adopted by diverse architectures. Other transfer-based attacks such as Dispersion Reduction (DR) [23], Intermediate Level Attack (ILA) [13] have improved the transferability of adversarial examples in different tasks through perturbing feature maps. In contrast, the proposed I2V attack implements a cross-modal transfer-based attack through a correlation between the spatial features encoded between the image model and the video model.

## 2.2. Transfer-based Attacks on Video Recognition Models

There is much less work about transfer-based attacks on video models compared with transfer-based attacks on image models. Temporal Translation (TT) attack method [35] optimizes the adversarial perturbations over a set of temporal translated video clips for avoiding overfitting to the white-box model being attacked. Although TT achieves better results than transfer-based image attack methods, it increases the computational cost. Different from it, the proposed I2V attack achieves better performance without trained video models and is easy to perform.

## 2.3. Video Recognition Models

Video action recognition models have made significant progress in recent years. Previous studies [16, 40] adopt a 2D + 1D paradigm, where 2D CNNs are applied over per-frame input to extract features, followed by a 1D module (*e.g.*, RNNs) that integrates per-frame features. Current

studies use 3D CNNs to jointly capture the dynamic semantic of videos. For example, I3D [2] leverages ImageNet architecture designs and their parameters to encode spatio-temporal features by inflating the 2D convolution kernels into 3D. Non-local (NL) [30] network inserts a non-local operation into I3D for encoding long-range temporal dependencies between video frames. SlowFast [9] contrasts the visual tempos along the temporal axis, which involves a slow pathway and a fast pathway to capture spatial semantics and motion at fine temporal resolution respectively. Temporal Pyramid Network (TPN) [39] capture action instances at various tempos through a feature hierarchy architecture. In this paper, we use six representative video action recognition models for experiments, including NL, SlowFast, TPN with 3D Resnet-50 and Resnet-101 as backbones.

## 3. Methodology

### 3.1. Preliminary

Given a video sample  $x \in \mathcal{X} \subset \mathbf{R}^{T \times H \times W \times C}$  with the true label  $y \in \mathcal{Y} = \{1, 2, \dots, K\}$ , where  $T, H, W, C$  denote the number of frames, height, width and channels respectively.  $K$  represents the number of classes. Let  $g$  denote the ImageNet-pretrained image model (*e.g.*, ResNet, VGG),  $f$  denote the video recognition model. We use  $f(x) : \mathcal{X} \rightarrow \mathcal{Y}$  to denote the prediction of the video recognition model for an input video. Thus, the proposed I2V attack aims to generate the adversarial example  $x_{adv} = x + \delta$  by  $g$ , which can fool the video model  $f$  into  $f(x_{adv}) \neq y$  without knowledge about  $f$ , where  $\delta$  denotes the adversarial perturbation. To ensure that the adversarial perturbation  $\delta$  is imperceptible, we restrict it by  $\|\delta\|_p \leq \epsilon$ , where  $\|\cdot\|_p$  denotes the  $L_p$  norm, and  $\epsilon$  is a constant of the norm constraint. We adopt  $L_\infty$  norm and untargeted adversarial attacks, which are commonly used in [7, 20, 36–38]. In a white-box setting, the objective of untargeted adversarial attacks can be formulated as follows:

$$\arg \max_{\delta} J(f(x + \delta), y), s.t. \|\delta\|_\infty \leq \epsilon, \quad (1)$$

where  $J$  is the loss function (*e.g.*, cross-entropy loss) of the video model  $f$ . However, in this paper, the adversary cannot access knowledge about  $f$ . The proposed I2V attack leverages adversarial examples generated from  $g$  to attack  $f$  in the black-box setting.

### 3.2. Correlation Analysis between Image and Video Models

Before introducing the proposed method, we firstly give an empirical analysis of the correlation between image and video models. It has been demonstrated in the prior work [15] that utilizing ImageNet-pretrained image models to

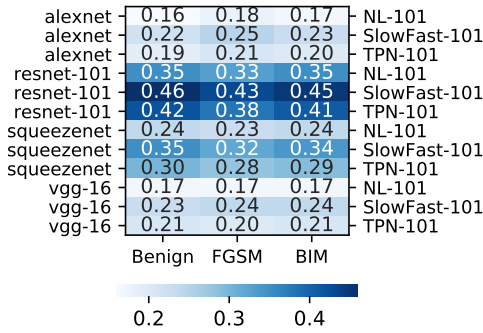



Figure 2. Cosine similarity analysis of intermediate features between image models and video models on benign examples and adversarial examples. The cosine similarities are calculated on the averaged intermediate features extracted from 400 randomly selected videos in Kinetics-400. Adversarial video samples are generated by FGSM and BIM, using video models as white-box models. Darker color represents a higher cosine similarity.


generate tentative perturbations assumes fewer queries for attacking black-box video recognition models. This basically suggests that *the intermediate features between images models and video models may be similar to a certain extent*. Hence perturbing intermediate feature maps of image models could affect that of video models. To verify this assumption, we analyze the intermediate features’ similarity of both benign and adversarial frames between image and video models with cosine similarity.

Figure 2 shows the cosine similarity of intermediate features between image and video models. The intermediate features are extracted from 400 randomly selected videos in Kinetics-400 and then averaged for calculating cosine similarity. For all video models, the intermediate features are extracted from the first 3D-Resnet block, while for different image models, the features are extracted from different intermediate layers, which are summarized in Table 2 (marked in red color). Here we choose different intermediate layers for different image models for the purpose of maximizing the similarity between image features and video features. From figure 2, it can be observed that, for both benign samples and adversarial samples, their intermediate layer features extracted from image models and video models are similar to a certain extent. This is mainly because the convolution operations in image models and video models are somehow similar. It is worthwhile to mention that the cosine similarities obtained from benign samples and adversarial samples are quite similar. This basically suggests that adversarial perturbations have little effect on the similarity of feature space between image and video models. Similar trends can be also observed when using other intermediate layers of video models.

To demonstrate that the adversarial perturbations on the



(a) Resnet-101



(b) NL-101

Figure 3. The magnitudes (y-axis) of channel-wise activation at the penultimate layer (2048 channels at x-axis) for both image and video models. The magnitude is calculated by the global average pooling at each channel. In each plot, the channel-wise magnitudes are averaged from randomly selected 400 videos of Kinetics-400 and displayed separately for benign and adversarial examples (generated by BIM). The 2048 channels are sorted in descending order of magnitude for benign examples.

feature maps are translatable between video and image models, we further compare the magnitude changes in the channel-wise activation of image and video models before and after adding the same adversarial perturbations to the video frames. The results are shown in Figure 3. As can be seen, the adversarial examples generated on NL-101 perturb the channel-wise activation magnitude not only in NL-101 but also in Resnet-101. Since each channel of features captures a specific pattern of the object and contributes differently to the final classification, the magnitude changes of image and video models are likely to result in wrong predictions, demonstrating possibilities in transferring adversarial perturbations between image and video models.

### 3.3. Image To Video (I2V) Attack

Based on the above observations, we propose the Image To Video (I2V) attack, which generates video adversarial examples from an ImageNet-pretrained image model, to boost the transferability of hetero-modal models and attack video models in the black-box setting. By perturbing intermediate features of image models, I2V generates adversarial examples to disturb intermediate features of black-box video models with high probability. In particular, I2V optimizes the  $i$ -th adversarial frame by:

$$\arg \min_{\delta} \text{CosSim}(g_l(x^i + \delta), g_l(x^i)), \text{ s.t. } \|\delta\|_{\infty} \leq \epsilon, \quad (2)$$

---

**Algorithm 1: Image to Video (I2V) attack.**

---

**Input** : A video example  $x$ , image model  $g$ .  
**Parameter**: Perturbation budget  $\epsilon$ , iteration number  $I$ , step size  $\alpha$ , the number of layer  $l$ .  
**Output** : The adversarial example  $x_{adv}$ .

```
1 for  $i = 1$  to  $T$  do
2    $x^i =$  the  $i$ -th frame of  $x$ 
3    $\delta_0^i = (\frac{0.01}{255})^{H \times W \times C}$ 
4   for  $j = 0$  to  $I - 1$  do
5     Update  $\delta_j^i$  by Adam optimizer:
6      $\delta_{j+1}^i = ADAM(\delta_j^i, \alpha, CosSim(g_l(x^i + \delta_j^i), g_l(x^i)))$ 
7   end
8   Project  $x_{adv}^i$  to the  $\epsilon$ -ball of  $x^i$ :
9    $x_{adv}^i = clip_{x^i, \epsilon}(x^i + \delta_I^i)$ 
10 end
11 return  $x_{adv} = (x_{adv}^1, \dots, x_{adv}^i, \dots, x_{adv}^T)$ 
```

---

where  $g_l(x^i)$  denotes the intermediate feature map of the  $l$ -th layer with respect to  $x^i$  in the image model,  $x^i \in \mathbf{R}^{H \times W \times C}$  denotes the  $i$ -th frame of  $x$ , the function  $CosSim$  calculates the cosine similarity between  $g_l(x^i + \delta)$  and  $g_l(x^i)$ .

In this way, the minimization of the cosine similarity makes it possible to optimize adversarial examples with features that are orthogonal to ones of benign examples. Consider that  $g_l(x^i)$  is the output of the penultimate layer and let  $W = (W_1, \dots, W_y, \dots, W_K)$  denote the weight of the classification layer, thus  $W_y$  and  $g_l(x^i)$  are highly aligned for making a true prediction. With minimizing  $CosSim(g_l(x^i + \delta), g_l(x^i)) = \frac{g_l(x^i + \delta)^T g_l(x^i)}{\|g_l(x^i + \delta)\| \cdot \|g_l(x^i)\|}$ , we can get minimizing  $g_l(x^i + \delta)^T g_l(x^i)$  if  $g_l(x^i + \delta)$  and  $g_l(x^i)$  have unit length. Due to the high alignment between  $W_y$  and  $g_l(x^i)$ , the minimization of the cosine similarity induces that the value of  $W_y \cdot g_l(x^i + \delta)$  decreases a lot to fool the image model  $g$  into making error predictions. Based on the similarity of feature space between image and video models, the generated adversarial examples  $x_{adv} = (x_{adv}^1, \dots, x_{adv}^i, \dots, x_{adv}^T)$  may fool video models with high probability by perturbing video intermediate features.

Following [32], we initialize the adversarial perturbations  $\delta$  with a small constant value  $\frac{0.01}{255}$  and use the Adam optimizer [17] to solve the Equation 2 and updates  $\delta_j^i$ . Algorithm 1 illustrates the generation of adversarial examples of the proposed I2V attack. Where  $I$  denotes the iteration number of Adam optimizer,  $clip_{x^i, \epsilon}$  denote to project  $x^i + \delta_I^i$  to the vicinity of  $x^i$  for meeting  $\|\delta_I^i\|_\infty \leq \epsilon$ . In the end, I2V attack combines all generated adversarial frames  $x_{adv}^i$  into a video adversarial examples  $x_{adv}$ .

Model	UCF-101	Kinetics-400
NL-50	81.26	75.17
NL-101	82.21	75.81
SlowFast-50	85.25	76.66
SlowFast-101	86.10	76.95
TPN-50	87.13	78.90
TPN-101	90.28	79.70

Table 1. Top-1 validation accuracy(%) of video recognition models on UCF-101 and Kinetics-400.

### 3.4. Attacking an Ensemble of Models

MIFGSM [6] shows that attacking an ensemble of models can boost the transferability of generated adversarial examples. When a generated example remains adversarial over an ensemble of models, it may transfer to attack other models. Based on this, we propose to use multiple ImageNet-pretrained image models to perform the I2V attack, named ENS-I2V, which optimizes  $i$ -th adversarial frame by:

$$\arg \min_{\delta} \sum_{n=1}^N CosSim(g_l^n(x^i + \delta), g_l^n(x^i)), s.t. \|\delta\|_\infty \leq \epsilon, \quad (3)$$

where  $N$  is the number of used image models,  $g_l^n(\cdot)$  returns the intermediate features of the  $l$ -th layer in the  $n$ -th image model. The intermediate features of the adversarial frames generated by ENS-I2V are orthogonal to the ensemble of features from benign examples, thus ENS-I2V allows the generation of highly transferable adversarial examples.

## 4. Experiments

### 4.1. Experimental Setting

**Dataset.** We evaluate our approach using UCF-101 [27] and Kinetics-400 [2] datasets, which are widely used datasets for video recognition. UCF-101 consists of 13,320 videos from 101 actions. Kinetics-400 contains approximately 240,000 videos from 400 human actions.

**ImageNet-pretrained image models.** We perform our proposed methods on four ImageNet-pretrained image models: Alexnet [18], Resnet-101 [12], Squeezenet 1.1 [14] and Vgg-16 [26]. Where Squeezenet 1.1 has 2.4x less computation and slightly fewer parameters than SqueezeNet 1.0, without sacrificing accuracy. These four models are commonly used in the image classification.

**Video recognition models.** Our proposed methods are evaluated on three different architectures of video recognition models: Non-local (NL) [30], SlowFast [9], TPN [39]. NL, SlowFast and TPN use 3D Resnet-50/101 as the backbone. We train these video models from scratch with Kinetics-400 and fine-tune them on UCF-101. For Kinetics-




Figure 4. AASR (%) of the I2V attack with various step sizes and iteration numbers.

400, we skip every other frame from randomly selected 64 consecutive frames into constructing input clips. For UCF-101, we use 32 consecutive frames as input clips. The spatial size of the input is  $224 \times 224$ . Table 1 summarizes top-1 validation accuracy of these six models on UCF-101 and Kinetics-400.

**Attack setting.** In our experiments, we use the Attack Success Rate (ASR) to evaluate the attack performance, which is the rate of adversarial examples that are successfully misclassified by the black-box video recognition model. Thus higher ASR means better adversarial transferability. If not specifically stated, average ASR (AASR) is the average ASR over all black-box video models. Following [7, 38], we randomly sample one video, which is correctly classified by all models, from each class to conduct our experiments, and set the norm constraint  $\epsilon = 16$ .

## 4.2. Ablation Study

We first investigate the effects of step size  $\alpha$ , iteration number  $I$  and different attacked layers  $l$  of image models. The evaluations are conducted on video models trained on Kinetics-400.

**Step size and iteration number.** Equation 2 is solved by the Adam optimizer, which can be affected by the step size  $\alpha$  and iteration number  $I$ . Figure 4 shows the results of using Block-2 of Resnet-101 as the perturbed layer of the image model with different step sizes and iteration numbers. It can be seen that smaller  $\alpha$  and  $I$  have poorer AASR because of under-fitting. While larger  $\alpha$  can achieve better AASR with a smaller  $I$ . To achieve the best performance, we adopt  $\alpha = 0.005$  and  $I = 60$  in subsequent experiments.

**Intermediate layer selection.** For each image model, we select four layers from bottom to top layer (as shown in Table 2) to craft the adversarial perturbations. Figure 5 shows the results of performing attack on different layers. Attacking the middle layers (layer 2 or layer 3) of image models is better than attacking bottom or top layers. Based on the results, we attack the middle layers, which are marked in Red in Table 2 for each model.

Model	Layer 1	Layer 2	Layer 3	Layer 4
Alexnet	ReLU-1	ReLU-2	ReLU-3	ReLU-5
Resnet-101	Block-1	Block-2	Block-3	Block-4
Squeezenet	Fire-1	Fire-3	Fire-5	Fire-8
Vgg-16	ReLU-1	ReLU-5	ReLU-9	ReLU-13

Table 2. Intermediate layer selection. For each image model, the selected intermediate layer used for crafting adversarial perturbations in the proposed I2V attack is in Red.




Figure 5. AASR (%) of the I2V attack with various attack layers  $l$  in different image models.

## 4.3. Performance Comparison

Since transferability between hetero-modal models has never been explored, we compare our proposed I2V and ENS-I2V attacks with DR [23], which is originally proposed to enhance cross-task transferability. DR minimizes the standard deviation of intermediate features to degrade the recognizability of images. We extend DR to optimize adversarial examples on the Imagenet-pretrained image models and use the same setting as the I2V attack.

The results of attacking UCF-101 and Kinetics-400 datasets are shown in Table 3 and 4, respectively. From the results, we have the following observations. First, the proposed I2V and ENS-I2V attacks achieve much higher ASR than DR by a large margin. For example, compared with DR, I2V can boost AASR of more than 63.33% and 42.51% for UCF-101 and Kinetics-400 separately. Second, I2V using Resnet-101 as the white-box image model outperforms all the other I2V attacks, which suggests that 2D Resnet101 and 3D Resnet-101 in the backbone of video models share more similar feature space than other 2D image models. Third, ENS-I2V further improves the average AASR to 65.68% against UCF101 and 77.88% against Kinetics-400. Which demonstrates the validity of attacking an ensemble of image models. In general, our method, which considers minimizing the cosine similarity between features from adversarial and benign examples, consistently outperforms DR. These experiments validate the effectiveness of the proposed I2V attack.

Attack	Image Model	Black-box Video Model					
		NL-101	NL-50	SlowFast-101	SlowFast-50	TPN-101	TPN-50
DR [23]	Alexnet	29.70	26.73	19.80	25.74	8.91	13.86
	Resnet-101	14.85	23.76	18.81	27.72	15.84	20.79
	Squeezenet	12.87	24.75	12.87	15.84	4.95	13.86
	Vgg-16	14.85	29.70	13.86	22.77	6.93	15.84
I2V	Alexnet	50.49	53.46	35.64	43.56	28.71	44.50
	Resnet-101	<b>71.28</b>	60.39	50.49	57.42	<b>61.38</b>	71.28
	Squeezenet	43.56	54.45	36.63	37.62	23.76	35.64
	Vgg-16	43.56	48.51	28.71	39.60	21.78	32.67
ENS-I2V	Ensemble	<b>71.28</b>	<b>76.23</b>	<b>56.43</b>	<b>62.37</b>	52.47	<b>75.24</b>

Table 3. ASR (%) against video recognition models on UCF-101.

Attack	Image Model	Black-box Video Model					
		NL-101	NL-50	SlowFast-101	SlowFast-50	TPN-101	TPN-50
DR [23]	Alexnet	22.00	31.50	43.00	41.75	31.00	39.00
	Resnet-101	25.50	37.25	49.00	52.25	41.50	42.75
	Squeezenet	17.00	25.00	37.00	36.50	24.25	29.50
	Vgg-16	16.75	23.00	36.75	35.75	23.75	29.00
I2V	Alexnet	44.00	54.75	61.50	59.50	59.75	69.50
	Resnet-101	56.25	64.50	74.75	<b>77.00</b>	<b>87.25</b>	<b>90.25</b>
	Squeezenet	37.75	51.00	62.50	60.25	55.50	58.50
	Vgg-16	39.00	46.25	57.75	59.00	59.00	70.50
ENS-I2V	Ensemble	<b>65.00</b>	<b>72.25</b>	<b>79.75</b>	76.50	85.75	88.00

Table 4. ASR (%) against video recognition models on Kinetics-400.

#### 4.4. Comparing against Stronger Baselines

We further compare the proposed I2V attack against several existing transfer-based attacks that are designed for homomodal models (e.g., images models or video models). It’s worthwhile to mention that the comparisons are unfair because the existing transfer-based attacks require the white-box video recognition models to generate the adversarial perturbations. For the comparison, several transfer-based attacks, such as FGSM [10], BIM [19], MI [6], DI [38], TI [7], SIM [20], SGM [36], TAP [42], ATA [37], and TT [35] are used as baselines. For these baselines, NL-101, SlowFast-101 and TPN-101 are used as the white-box models. It has been illustrated in ILA [13], the transferability of generated adversarial examples can be further enhanced through the proposed fine-tuning methods ILAP and ILAF. Compared to ILAP, ILAF achieves better performance by maintaining the existing adversarial direction and increases the magnitude of feature perturbations under  $L_2$  norm [13]. Therefore, for the compared baseline methods, we use ILAF to fine-tune the generated adversarial examples.

Figure 6 shows the comparison results. From the results, we have the following observations. First, despite the comparisons being unfair since our method does not require any

white-box video models, the proposed ENS-I2V still performs much better than ILAF for most cases. As shown in Figure 6(a)(b)(d)(f), on both UCF-101 and Kinetics-400, ENS-I2V exceeds ILAF by a large margin when using NL-101 or SlowFast-101 as the white-box model. Second, our I2V attack performs worse than baseline methods when they use TPN-101 as the white-box model on Kinetics-400 (Figure 6f). This may be because that Kinetics-400 contains richer motion information than UCF-101 and such motion information are unlikely to be well captured by image models. On the contrary, by fusing multi-layer features, TPN-101 can better capture the motion information. As a result, disrupting motion information (Figure 6f) can achieve better performances.

#### 4.5. Discussion

To experimentally demonstrate the effectiveness of the optimized object function (Equation 2), we investigate the changes in the cosine similarity of the adversarial image/video features to the benign image/video features by increasing the iteration number. Pearson Correlation Coefficient (PCC) [1] is used to measure the linear correlation of cosine similarity trends computed from image and video models. Figure 7 shows the PCC analysis of cosine

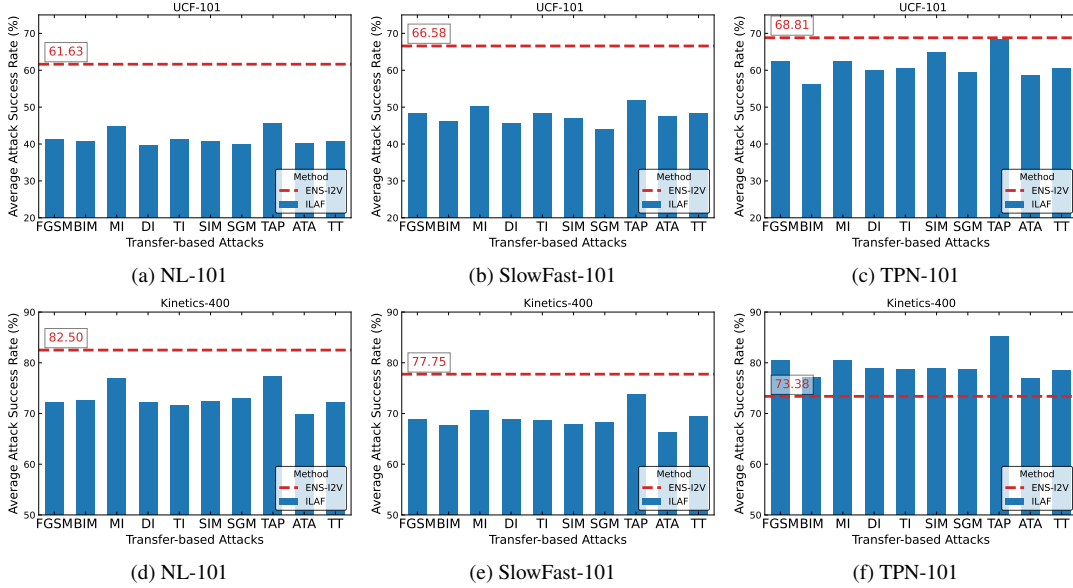


Figure 6. AASR (%) against video recognition models with fine-tuning attack methods. The top and bottom rows are the results on UCF-101 and Kinetics-400 respectively. The three columns use the NL-101, SlowFast-101, and TPN-101 models as white-box models respectively. AASR is calculated by averaging ASR over black-box video models that have a different architecture from the white-box model. Red dashed lines denote the performance of our proposed ENS-I2V.

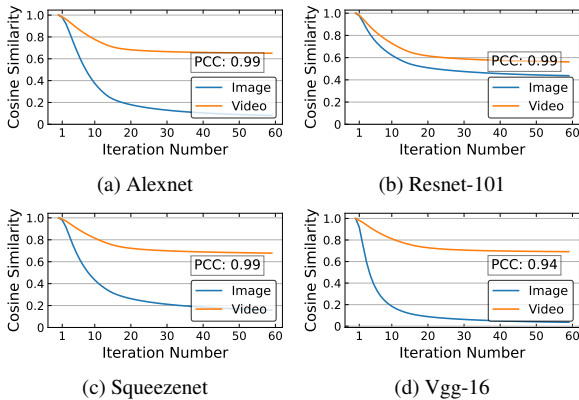


Figure 7. Pearson correlation coefficient (PCC) analysis between cosine similarity trends computed from image and video models. The used intermediate layers of image models are summarized in Tab. 2 (marked in red). NL-101 uses the second 3D-Resnet block.

similarity trends using 4 image models and a video model (NL-101). As can be seen, all PCCs are close to 1, which implies an exact positive linear relationship between the directional changes of image and video intermediate features. It suggests that minimizing the cosine similarity of image models can make intermediate features of video adversarial examples generated from Imagenet-pretrained models orthogonal to their benign video features. Similar trends are observed when using other video models.

## 5. Conclusion

In this paper, we identify the existence of a similar feature space between image and video models, which can be leveraged to generate adversarial examples from image models to attack black-box video models. More specifically, we proposed the Image To Video (I2V) attack, which optimizes adversarial frames on the ImageNet-pretrained image model by minimizing the cosine similarity between features from adversarial and benign examples for perturbing intermediate feature space. Besides, we proposed ENS-I2V to attack an ensemble of image models for boosting transferability. The results indicate that cross-modal adversarial transferability occurs even across image and video domains. In the future, we will combine temporal information of videos into image models to further boost transferability. Code is available at <https://github.com/zhipeeng-wei/Image-to-Video-I2V-attack>.

## 6. Acknowledgments

This project was supported by National Key R&D Program of China (No. 2021ZD0112804) and in part by NSFC project (#62032006), Science and Technology Commission of Shanghai Municipality Project (20511101000). Y.-G. Jiang was sponsored in part by “Shuguang Program” supported by Shanghai Education Development Foundation and Shanghai Municipal Education Commission (No. 20SG01).



## References

- [1] Theodore Wilbur Anderson. An introduction to multivariate statistical analysis. Technical report, Wiley New York, 1962. [7](#)
- [2] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. [1](#), [3](#), [5](#)
- [3] Jingjing Chen, Liangming Pan, Zhipeng Wei, Xiang Wang, Chong-Wah Ngo, and Tat-Seng Chua. Zero-shot ingredient recognition by multi-relational graph convolutional network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10542–10550, 2020. [1](#)
- [4] Jingjing Chen, Bin Zhu, Chong-Wah Ngo, Tat-Seng Chua, and Yu-Gang Jiang. A study of multi-task and region-wise deep learning for food ingredient recognition. *IEEE Transactions on Image Processing*, 30:1514–1526, 2020. [1](#)
- [5] Kai Chen, Zhipeng Wei, Jingjing Chen, Zuxuan Wu, and Yu-Gang Jiang. Attacking video recognition models with bullet-screen comments. *ArXiv*, abs/2110.15629, 2021. [1](#)
- [6] Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9185–9193, 2018. [1](#), [3](#), [5](#), [7](#)
- [7] Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Evading defenses to transferable adversarial examples by translation-invariant attacks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4312–4321, 2019. [1](#), [3](#), [6](#), [7](#)
- [8] Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Chaowei Xiao, Atul Prakash, Tadayoshi Kohno, and Dawn Song. Robust physical-world attacks on deep learning visual classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1625–1634, 2018. [1](#)
- [9] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019. [3](#), [5](#)
- [10] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. [2](#), [7](#)
- [11] Yanbin Hao, Zi-Niu Liu, Hao Zhang, Bin Zhu, Jingjing Chen, Yu-Gang Jiang, and Chong-Wah Ngo. Person-level action recognition in complex events via tsd-tsm networks. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4699–4702, 2020. [1](#)
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [1](#), [5](#)
- [13] Qian Huang, Isay Katsman, Horace He, Zeqi Gu, Serge Belongie, and Ser-Nam Lim. Enhancing adversarial example transferability with an intermediate level attack. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4733–4742, 2019. [3](#), [7](#)
- [14] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016. [5](#)
- [15] Linxi Jiang, Xingjun Ma, Shaoxiang Chen, James Bailey, and Yu-Gang Jiang. Black-box adversarial attacks on video recognition models. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 864–872, 2019. [3](#)
- [16] Yu-Gang Jiang, Zuxuan Wu, Jun Wang, Xiangyang Xue, and Shih-Fu Chang. Exploiting feature and class relationships in video categorization with regularized deep neural networks. *IEEE transactions on pattern analysis and machine intelligence*, 40(2):352–364, 2017. [3](#)
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [5](#)
- [18] Alex Krizhevsky. One weird trick for parallelizing convolutional neural networks. *arXiv preprint arXiv:1404.5997*, 2014. [5](#)
- [19] Alexey Kurakin, Ian Goodfellow, Samy Bengio, et al. Adversarial examples in the physical world, 2016. [2](#), [7](#)
- [20] Jiadong Lin, Chuanbiao Song, Kun He, Liwei Wang, and John E Hopcroft. Nesterov accelerated gradient and scale invariance for adversarial attacks. *arXiv preprint arXiv:1908.06281*, 2019. [1](#), [3](#), [7](#)
- [21] Aishan Liu, Xianglong Liu, Jiaxin Fan, Yuqing Ma, Anlan Zhang, Huiyuan Xie, and Dacheng Tao. Perceptual-sensitive gan for generating adversarial patches. In *33rd AAAI Conference on Artificial Intelligence*, 2019. [1](#)
- [22] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into transferable adversarial examples and black-box attacks. *arXiv preprint arXiv:1611.02770*, 2016. [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 the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 940–949, 2020. [3](#), [6](#), [7](#)
- [24] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28:91–99, 2015. [1](#)
- [25] Mahmood Sharif, Sruti Bhagavatula, Lujio Bauer, and Michael K Reiter. Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proceedings of the 2016 acm sigsac conference on computer and communications security*, pages 1528–1540, 2016. [1](#)
- [26] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [5](#)
- [27] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. [5](#)
- [28] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus.

- Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013. [1](#)
- [29] Jiakai Wang, Zixin Yin, Pengfei Hu, Xianglong Liu, Renshuai Tao, Haotong Qin, Aishan Liu, and Dacheng Tao. Defensive patches for robust recognition in the physical world. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022. [1](#)
- [30] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018. [3](#), [5](#)
- [31] Xingxing Wei, Huanqian Yan, and Bo Li. Sparse black-box video attack with reinforcement learning. *IJCV*, 2022. [1](#)
- [32] Xingxing Wei, Jun Zhu, Sha Yuan, and Hang Su. Sparse adversarial perturbations for videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8973–8980, 2019. [5](#)
- [33] Zhipeng Wei, Jingjing Chen, Micah Goldblum, Zuxuan Wu, Tom Goldstein, and Yu-Gang Jiang. Towards transferable adversarial attacks on vision transformers. *ArXiv*, abs/2109.04176, 2021. [1](#)
- [34] Zhipeng Wei, Jingjing Chen, Xingxing Wei, Linxi Jiang, Tat-Seng Chua, Fengfeng Zhou, and Yu-Gang Jiang. Heuristic black-box adversarial attacks on video recognition models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12338–12345, 2020. [1](#)
- [35] Zhipeng Wei, Jingjing Chen, Zuxuan Wu, and Yu-Gang Jiang. Boosting the transferability of video adversarial examples via temporal translation, 2021. [3](#), [7](#)
- [36] Dongxian Wu, Yisen Wang, Shu-Tao Xia, James Bailey, and Xingjun Ma. Skip connections matter: On the transferability of adversarial examples generated with resnets. *arXiv preprint arXiv:2002.05990*, 2020. [1](#), [3](#), [7](#)
- [37] Weibin Wu, Yuxin Su, Xixian Chen, Shenglin Zhao, Irwin King, Michael R Lyu, and Yu-Wing Tai. Boosting the transferability of adversarial samples via attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1161–1170, 2020. [1](#), [3](#), [7](#)
- [38] Cihang Xie, Zhishuai Zhang, Yuyin Zhou, Song Bai, Jianyu Wang, Zhou Ren, and Alan L Yuille. Improving transferability of adversarial examples with input diversity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2730–2739, 2019. [1](#), [3](#), [6](#), [7](#)
- [39] Ceyuan Yang, Yinghao Xu, Jianping Shi, Bo Dai, and Bolei Zhou. Temporal pyramid network for action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 591–600, 2020. [3](#), [5](#)
- [40] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4694–4702, 2015. [3](#)
- [41] Hao Zhang, Yanbin Hao, and Chong-Wah Ngo. Token shift transformer for video classification. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 917–925, 2021. [1](#)
- [42] Wen Zhou, Xin Hou, Yongjun Chen, Mengyun Tang, Xiangqi Huang, Xiang Gan, and Yong Yang. Transferable adversarial perturbations. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 452–467, 2018. [7](#)