
Adversarial Attacks on Black Box Video Classifiers: Leveraging the Power of Geometric Transformations

Shasha Li*, Abhishek Aich*, Shitong Zhu, M. Salman Asif, Chengyu Song,
Srikanth V. Krishnamurthy, Amit K. Roy-Chowdhury
University of California, Riverside, CA, USA

Abstract

When compared to the image classification models, black-box adversarial attacks against video classification models have been largely understudied. This could be possible because, with video, the temporal dimension poses significant additional challenges in gradient estimation. Query-efficient black-box attacks rely on effectively estimated gradients towards maximizing the probability of misclassifying the target video. In this work, we demonstrate that such effective gradients can be searched for by parameterizing the temporal structure of the search space with geometric transformations. Specifically, we design a novel iterative algorithm Geometric TRAnsformed Perturbations (GEO-TRAP), for attacking video classification models. GEO-TRAP employs standard geometric transformation operations to reduce the search space for effective gradients by searching for a small group of parameters that define these operations. This group of parameters describes the geometric progression of gradients, resulting in a reduced and structured search space. Our algorithm inherently leads to successful perturbations with surprisingly few queries. For example, adversarial examples generated from GEO-TRAP have better attack success rates with $\sim 73.55\%$ fewer queries compared to the state-of-the-art method for video adversarial attacks on the widely used Jester dataset. Overall, our algorithm exposes vulnerabilities of diverse video classification models and achieves new state-of-the-art results under black-box settings on two large datasets. Code is available here: <https://github.com/sli057/Geo-TRAP>

1 Introduction

Adversarial attacks are designed to expose vulnerabilities of Deep Neural Networks (DNNs). With real-world applications of video classification based on DNNs emerging [1–3], a key question that arises is “*what type of adversarial inputs can mislead, and thus render video classification networks vulnerable?*” Designing such adversarial attacks not only helps expose security flaws of DNNs, but can also potentially stimulate the design of more robust video classification models.

Adversarial attacks against image classification models have been studied in both *white-box* [4–8] and *black-box* [9–13] settings. In the white-box setting, an adversary has full access to the model under attack, including its parameters and training settings (hyper-parameters, training data, etc.) In the black-box setting, an adversary only has partial information about the victim model, such as the predicted labels of the model. In the case of video classification models, adversarial attacks in both white-box and black-box settings have garnered some interest [14–22], although the body of work here is more limited than the case of image classification models.

A common black-box attack paradigm is query-based, wherein the attacker can send queries to the victim model to collect the corresponding predicted labels, and thereby estimate the gradients needed

*Equal contribution. Corresponding author: Shasha Li (sli057@ucr.edu)

Table 1: **Comparison with state-of-the-art.** GEO-TRAP, compared to current black-box attack methods for videos, doesn’t train a different network to craft perturbations, and parameterizes the temporal dimension of videos in searching for effective perturbation directions.

Methods	WITHOUT training a “perturbation” network	CONSIDER temporal dimension?	PARAMETERIZE temporal dimension?
PATCHATTACK[19]	✗	✗	✗
HEURISTICATTACK [20]	✓	✗	✗
SPARSEATTACK [21]	✗	✗	✗
MOTION-SAMPLER ATTACK [22]	✓	✓	✗
GEO-TRAP (Ours)	✓	✓	✓

for curating the adversarial examples. Unlike static images, videos naturally include additional information from the temporal dimension. This high dimensionality (i.e., sequence of frames instead of one image) poses challenges to black-box adversarial attacks against video classification models; in particular, significantly more queries are typically needed for estimating the gradients for crafting adversarial samples [19–22]. [19] reduces the number of queries by adding perturbations on the patch level instead of at the pixel level; [20, 21] propose to add perturbations only on key pixels. [22] considers the intrinsic differences between images and videos (i.e., the temporal dimension), and proposes to use the optical-flow of clean videos as the motion prior for adversarial video generation. Similar to [22], we also explicitly consider the temporal dimension of video. However, rather than fixing the temporal search space using the motion prior of clean videos, we propose to parameterize the temporal structure of the space with geometric transformations. This results in a better structured and reduced search space, which allows us to generate successful attacks with much fewer queries in black-box settings than the state-of-the-art methods, including [22].

Contributions. In this paper, we propose a novel query-efficient black-box attack algorithm against video classification models. Due to the extra temporal dimension, generating video perturbations by searching for effective gradients remains a challenging task given the exceedingly large search space. These gradients are estimated by searching for ‘directions’ that maximize the probability of the victim model mis-classifying the crafted inputs. Our approach drastically reduces this large search space by defining this space with a small set of parameters that describe the geometric progression of gradients in the temporal dimension, resulting in a reduced and temporally structured search space. Conceptually, this parameterization of the temporal structure of the search space is performed using geometric transformations (e.g. affine transformations). We refer to our algorithm as Geometrically TRAnSformed Perturbations, or GEO-TRAP. Despite this surprisingly simple strategy, GEO-TRAP outperforms existing black-box video adversarial attack methods by significant margins ($\sim 1.8\%$ improvement in attack success rate with $\sim 73.55\%$ fewer queries for targeted attacks in comparison to the state-of-the-art [22] on the Jester dataset [23]).

2 Related Works

In this section, we review different black-box adversarial attacks strategies, and categorize our proposed method with respect to state-of-the-art black-box attacks designed for video classifiers.

In most real-world attacks, the adversary only has partial information about the victim models, such as the predicted labels. In such black-box settings, the adversary can first attack a local surrogate model and then transfer these attacks to the target victim model [24, 25], formally called as *transferability-based* black-box attack. Alternatively, they may estimate the adversarial gradient with zero-order optimization methods such as Finite Differences (FD) or Natural Evolution Strategies (NES) by querying the victim model [12, 13, 26], which is called *query-based* black-box attack. GEO-TRAP falls under the category of query-based black-box attacks (designed for videos).

Whilst several white-box attacks have been proposed for video classification models [14–18], black-box video attacks are relatively under explored. PATCHATTACK (V-BAD)[19] is the first to propose a black-box video attack framework which uses a hybrid attack strategy of first generating initial perturbations for each video frame by attacking a local image classifier, and then updating the perturbations by querying the victim model. Compared to PATCHATTACK[19], GEO-TRAP does not require training a local classifier. PATCHATTACK[19] crafts video perturbations by treating each frame as a separate image, but reduces the search space of the gradient estimation by morphing the

perturbations in patches/partitions. However, its attack performance has been shown to be inferior to that of a more recent approach [22] (discussed below). HEURISTICATTACK [20] uses a query-based attack strategy, and reduces the search space by generating adversarial perturbations only on heuristically selected key frames and salient regions. SPARSEATTACK [21] reduces the search space by adding perturbations only on key frames using a reinforcement learning based framework. MOTION-SAMPLER ATTACK [22] proposed a query-based attack strategy that utilizes a motion excited sampler to obtain *motion-aware* perturbation prior by using the optical-flow of the clean video. This motion-aware prior reduces the search space for gradients resulting in fewer queries. Similar to [22] but different from [19–21], GEO-TRAP explicitly considers the temporal dimension of video in order to search for effective gradients. However unlike [22], GEO-TRAP does not fix the temporal structure of the search space using a *pre-computed fixed* motion prior, but parameterize it with simple geometric transformations. These black-box video attack methods are summarized in Table 1.

3 Attacking via Geometrically TRAnsformed Perturbations (GEO-TRAP)

Notation. We denote the tuple of a video clip and its corresponding label as (\mathbf{x}, y) , which represents a data-point in the distribution \mathcal{X} . Each video sample $\mathbf{x} \in \mathbb{R}^{T \times H \times W \times C}$ has T frames of H height, W width, and C channels. We denote the victim video classification model as $\mathbf{f}_\theta : \mathcal{X} \rightarrow \mathcal{Y}$, where θ represents the model’s parameters learned from the training subset of \mathcal{X} , via a mapping to the label space \mathcal{Y} . We further assume \mathcal{X} consists of videos from $|\mathcal{Y}| = K$ categories. To make the perturbations imperceptible to humans, we impose the perturbation budget ρ_{\max} with the $\|\cdot\|_p$ norm. Throughout this paper, we consider $\|\cdot\|_\infty$ norm following [15, 19, 22] (the method can be extended to $p = 1, 2$ norms). To constrain $\|\cdot\|_\infty$ of perturbation below a budget ρ_{\max} , we use the `clip`(\cdot) function to keep the perturbation pixel value in $[-\rho_{\max}, \rho_{\max}]$. The function `sign`(\cdot) extracts the sign of given input variable. The superscript i , throughout the paper, denotes the iteration i . The subscript t denotes the frame index. For clarity, we represent vectors/tensors with the bold font and scalars with the regular font.

Problem Statement. We consider the scenario of attacking a standard video classification model using a query-based paradigm under **black-box settings** (assuming no access to θ nor the training subset of \mathcal{X}). Specifically, we aim to craft perturbed videos \mathbf{x}_{adv} with imperceptible differences from \mathbf{x} , in order to alter the decision of the target model \mathbf{f}_θ via multiple queries to guide the gradient estimation. This problem can be mathematically formulated as follows.

$$\underset{\mathbf{x}_{\text{adv}}}{\text{argmin}} \quad \mathcal{L}(\mathbf{f}_\theta(\mathbf{x}_{\text{adv}}), y) \quad \text{s.t.} \quad \|\mathbf{x}_{\text{adv}} - \mathbf{x}\|_\infty \leq \rho_{\max} \quad (1)$$

$\mathcal{L}(\mathbf{f}_\theta(\mathbf{x}_{\text{adv}}), y)$ is the objective function, capturing the similarity between the classifier’s output and the ground truth label y , and varies with different attack goals (targeted or untargeted). The challenge is to obtain \mathbf{x}_{adv} with as few queries as possible by estimating gradient $\mathbf{g}^* = \nabla_{\mathbf{x}_{\text{adv}}} \mathcal{L}(\mathbf{f}_\theta(\mathbf{x}_{\text{adv}}), y)$, which is unknown in the considered black-box setting.

Overview of GEO-TRAP. We propose a novel iterative video perturbation framework that follows the principle of the Basic Iterative Method [27] in order to fool \mathbf{f}_θ under $\|\cdot\|_\infty$ norm as follows.

$$\mathbf{x}_{\text{adv}}^{(0)} = \mathbf{x}, \quad \mathbf{x}_{\text{adv}}^{(i)} = \text{clip}(\mathbf{x}_{\text{adv}}^{(i-1)} - h \text{sign}(\mathbf{g}^{(i)})) \quad (2)$$

where h is a hyperparameter and $\mathbf{g}^{(i)}$ is the gradient estimated by querying the black-box victim model at the i^{th} iteration using our proposed GEO-TRAP algorithm. As shown in (2), effective perturbations rely on the guidance of the gradient $\mathbf{g}^{(i)}$. Therefore, efficiently estimating $\mathbf{g}^{(i)}$ is at the core of GEO-TRAP for successfully subverting video classifiers. We execute the following two steps in each iteration to estimate $\mathbf{g}^{(i)}$.

1. For any input video $\mathbf{x}^{(i)}$, a random noise tensor $\mathbf{r}_{\text{frame}} \in \mathbb{R}^{H \times W \times C}$ and a set of geometric transformation parameters $\Phi_{\text{warp}} \in \mathbb{R}^{T \times D}$ are chosen with each element sampled from a standard normal distribution. D represents the number of parameters needed for the geometric transformation of a single frame (details are provided in Section 3.2). In this setup, our search space for estimating $\mathbf{g}^{(i)}$ consists of $\mathbf{r}_{\text{frame}}$ and Φ_{warp} .
2. We then warp $\mathbf{r}_{\text{frame}}$ with Φ_{warp} to get the candidate direction $\pi = [\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_T] \in \mathbb{R}^{T \times H \times W \times C}$ (see Algorithm 2.1 TRANS-WARP). π is then employed to compute a gradient estimator Δ by querying the black-box victim model with a standard gradient estimation algorithm (see Algorithm 2 GRAD-EST). The gradient estimator Δ is then used to update $\mathbf{g}^{(i)}$.

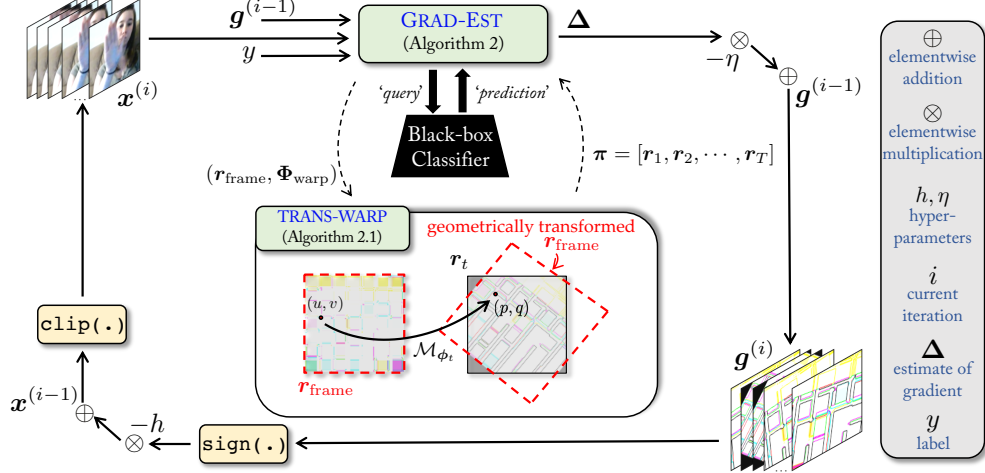


Figure 1: **Overview of GEO-TRAP.** GEO-TRAP is a black-box attack algorithm guided by the key observation that strong gradients $\mathbf{g}^{(i)}$ can be computed by finding better gradient search direction candidates π . We propose to search each frame of the directions \mathbf{r}_t by warping a randomly sampled $\mathbf{r}_{\text{frame}}$ using a geometric transformation \mathcal{M}_{ϕ_t} ; different \mathbf{r}_t in π are warped by the same $\mathbf{r}_{\text{frame}}$, thus have geometric progression among frames.

The overall attack strategy is summarized in Algorithm 1 and pictorially illustrated in Figure 1. Since in Step 2 above, the gradient estimation (i.e., GRAD-EST) procedure includes the geometric transformation strategy (i.e., TRANS-WARP), we will next describe GRAD-EST and then move on to TRANS-WARP. For the simplicity of exposition, we drop the superscript i and shorten the loss function to $\mathcal{L}(\mathbf{f}_{\theta}(\mathbf{x}_{\text{adv}}), y)$ to \mathcal{L} (as the model parameters θ remain unchanged) in rest of this section.

3.1 GEO-TRAP Gradient Estimation (GRAD-EST)

Let $\mathbf{g}^* = \nabla_{\mathbf{x}} \mathcal{L}$ be the ideal value of the gradient of \mathcal{L} at \mathbf{x} , required to create \mathbf{x}_{adv} in (2). To find an efficient estimator \mathbf{g} for \mathbf{g}^* , a (new) surrogate loss $\ell(\mathbf{g}) = -\langle \mathbf{g}^*, \mathbf{g} \rangle$ is defined such that the estimator \mathbf{g} has a sufficiently large inner product with the actual gradient \mathbf{g}^* (\mathbf{g} is normalized to a unit vector; we ignore the normalization operation for ease of explanation). The loss function definition and the algorithm to estimate \mathbf{g} follow [13].

As \mathbf{g}^* is unknown in the black-box setting, this surrogate loss function can be estimated as

$$\ell(\mathbf{g}) = -\langle \mathbf{g}^*, \mathbf{g} \rangle = -\langle \nabla_{\mathbf{x}} \mathcal{L}, \mathbf{g} \rangle \approx -\frac{\mathcal{L}(\mathbf{x} + \epsilon \mathbf{g}, y) - \mathcal{L}(\mathbf{x}, y)}{\epsilon}. \quad (3)$$

To iteratively estimate \mathbf{g} , we need to, in turn, estimate the gradient of $\ell(\mathbf{g})$, i.e., $\Delta = \nabla_{\mathbf{g}} \ell(\mathbf{g})$. With antithetic sampling [28], Δ can be estimated as

$$\Delta = \frac{\ell(\mathbf{g} + \delta \pi) - \ell(\mathbf{g} - \delta \pi)}{\delta} \pi, \quad (4)$$

where δ is a small number adjusting the magnitude of the loss variation and $\pi \in \mathbb{R}^{T \times H \times W \times C}$ is a random candidate direction. Our core contribution lies in the fact that instead of randomly sampling π in the search space [13], we reduce the search dimensionality by warping a randomly sampled tensor $\mathbf{r}_{\text{frame}} \in \mathbb{R}^{H \times W \times C}$ with another randomly sampled geometric (e.g., affine) transformation parameter tensor $\Phi_{\text{warp}} \in \mathbb{R}^{T \times D}$ to get π . The search space is then reduced from $T \times H \times W \times C$ to $(H \times W \times C) + (T \times D)$ and D is a relatively small number, $D \ll H \times W \times C$. With $\mathbf{w}_1 = \mathbf{g} + \delta \pi$ and $\mathbf{w}_2 = \mathbf{g} - \delta \pi$ and combining (3) with (4), we get

$$\Delta = \frac{\mathcal{L}(\mathbf{x} + \epsilon \mathbf{w}_2, y) - \mathcal{L}(\mathbf{x} + \epsilon \mathbf{w}_1, y)}{\epsilon \delta} \pi. \quad (5)$$

Note that by querying the victim model \mathbf{f}_{θ} with $\mathbf{x} + \epsilon \mathbf{w}_1$, we are able to retrieve the value of $\mathcal{L}(\mathbf{x} + \epsilon \mathbf{w}_1, y)$; similarly we can obtain the value of $\mathcal{L}(\mathbf{x} + \epsilon \mathbf{w}_2, y)$ ($\mathcal{L}(\cdot)$ is defined following [17]). In summary, we estimate Δ with these two queries to the victim model. The resulting algorithm for estimating gradient of $\nabla_{\mathbf{g}} \ell$ or Δ for consequently estimating \mathbf{g} is shown in Algorithm 2. Eventually at every iteration, we use Δ to update \mathbf{g} by applying a one-step gradient descent as $\mathbf{g} \leftarrow \mathbf{g} - \eta \Delta$, where η is a hyperparameter to update \mathbf{g} . This updated \mathbf{g} is later used to obtain \mathbf{x}_{adv} using (2).

Algorithm 1: GEO-TRAP: Query-based Iterative attack for Video Classifiers

Input : video \mathbf{x} , corresponding label y , step-size η for updating the gradient, step-size h for updating adversarial video.

Output : adversarial video \mathbf{x}_{adv}

```

1 Initialize:  $\mathbf{x}^{(0)} = \mathbf{x}, \mathbf{g}^{(0)} = \mathbf{0}, i = 1$ 
2 while  $\operatorname{argmax}(\mathbf{f}_\theta(\mathbf{x}^{(i)})) = y$  do
3    $\Delta = \text{GRAD-EST}(\mathbf{x}^{(i-1)}, \mathbf{g}^{(i-1)}, y)$  /* Gradient Estimation */
4    $\mathbf{g}^{(i)} \leftarrow \mathbf{g}^{(i-1)} - \eta \Delta$ 
5    $\mathbf{x}^{(i)} \leftarrow \text{clip}(\mathbf{x}^{(i-1)} - h \text{sign}(\mathbf{g}^{(i)}))$ 
6    $i \leftarrow i + 1$ 
7 end
8 return  $\mathbf{x}_{\text{adv}} = \mathbf{x}^{(i)}$ 

```

3.2 Noise Warping using Geometric Transformation (TRANS-WARP)

To tackle the challenge of the high-dimensionality of the search space, we propose to parameterize the search space with a single random noise tensor $\mathbf{r}_{\text{frame}} \in \mathbb{R}^{H \times W \times C}$ and a sequence of geometric transformations $\Phi_{\text{warp}} \in \mathbb{R}^{T \times D}$. Apart from the reduction of the search space of gradient estimation, our geometric transformation provides a temporal structure to π , which we discuss next.

At every iteration, $\pi = [\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_T]$ represents the candidate direction for Δ . These directions $\mathbf{r}_t \in \mathbb{R}^{H \times W \times C}$ are used to compute Δ in order to update gradient \mathbf{g} . To obtain π , we use a sequence of transformation vectors $\Phi_{\text{warp}} = [\phi_1, \phi_2, \dots, \phi_T]$ where $\phi_t \in \mathbb{R}^D$. The dimensionality D , chosen by the attacker, can vary depending on the transformation type that is populated from ϕ_t , e.g., $D = 6$ for affine transformation. We take affine transformation as an example to describe the warping process. We start by randomly sampling $\mathbf{r}_{\text{frame}}$ and the sequence of ϕ_t along with initializing each element in the sequence of \mathbf{r}_t with zero in **every** iteration. TRANS-WARP then computes \mathbf{r}_t by warping $\mathbf{r}_{\text{frame}}$ using the parameters in $\phi_t = [\phi_{11}^t, \phi_{12}^t, \phi_{13}^t, \phi_{21}^t, \phi_{22}^t, \phi_{23}^t] \in \mathbb{R}^6$ of $\Phi_{\text{warp}} \in \mathbb{R}^{T \times 6}$ as follows. For all C channels, let (p, q) and (u, v) be the target and source coordinates in \mathbf{r}_t and $\mathbf{r}_{\text{frame}}$, respectively. \mathbf{r}_t (for all channels) is computed as

$$\mathbf{r}_t(p, q) \leftarrow \mathbf{r}_{\text{frame}}(u, v), \quad 1 \leq p, u \leq H, 1 \leq q, v \leq W. \quad (6)$$

Location (p, q) is computed using the affine transform matrix \mathcal{M}_{ϕ_t} created with ϕ_t in homogeneous coordinates [29] as shown below. t is dropped for simplicity.

$$\begin{pmatrix} p \\ q \\ 1 \end{pmatrix} = \mathcal{M}_{\phi} \begin{pmatrix} u \\ v \\ 1 \end{pmatrix} = \begin{bmatrix} \phi_{11} & \phi_{12} & \phi_{13} \\ \phi_{21} & \phi_{22} & \phi_{23} \\ 0 & 0 & 1 \end{bmatrix} \begin{pmatrix} u \\ v \\ 1 \end{pmatrix} \quad (7)$$

We compactly denote this warping operation in (6) and (7) with $\mathbf{r}_t = \mathcal{T}(\mathbf{r}_{\text{frame}}, \phi_t)$. Affine transformation allows translation, rotation, scaling, and skew to be applied to $\mathbf{r}_{\text{frame}}$ to get each \mathbf{r}_t . Therefore, the sequence of \mathbf{r}_t have affine geometric progression among its temporal dimension. Other examples of geometric transformations may be more constrained, such as the similarity transformation \mathcal{M}_{ϕ}^S (that allows translation, dilation (uniform scale) and rotation with $D = 4$) and translation-dilation $\mathcal{M}_{\phi}^{\text{TD}}$ (that allows translation and uniform dilation with $D = 3$) as shown below.

$$[\phi_{11}, \phi_{12}, \phi_{13}, \phi_{23}] \rightarrow \mathcal{M}_{\phi}^S = \begin{bmatrix} \phi_{11} & \phi_{12} & \phi_{13} \\ -\phi_{12} & \phi_{11} & \phi_{23} \\ 0 & 0 & 1 \end{bmatrix}, [\phi_{11}, \phi_{13}, \phi_{23}] \rightarrow \mathcal{M}_{\phi}^{\text{TD}} = \begin{bmatrix} \phi_{11} & 0 & \phi_{13} \\ 0 & \phi_{11} & \phi_{23} \\ 0 & 0 & 1 \end{bmatrix} \quad (8)$$

4 What Makes GEO-TRAP Effective?

Potent iterative algorithms should rely on few queries for crafting successful perturbations for time efficiency. To minimize the number of queries, iterative algorithms need to find strong gradients in their early iterations. As discussed earlier, videos inherently incur a larger search space due to the temporal dimension and thus, pose challenges in searching for effective gradients. In this section, we provide empirical evidence to show that by parameterizing the temporal dimension, GEO-TRAP finds better gradients, in general, than previous works. We use three baselines in this analysis.

Algorithm 2: GRAD-EST($\mathbf{x}^{(i-1)}, \mathbf{g}^{(i-1)} \in \mathbb{R}^{T \times H \times W \times C}, y$) \rightarrow Estimate $\Delta = \nabla_{\mathbf{g}} \ell(\mathbf{g}) \in \mathbb{R}^{T \times H \times W \times C}$

Input : video $\mathbf{x}^{(i)}$, label y , gradient estimator $\mathbf{g}^{(i-1)}$, δ for loss variation, ϵ for approximation.

Output : estimation of $\Delta = \nabla_{\mathbf{g}} \ell(\mathbf{g})$

- 1 **Sample** $\mathbf{r}_{\text{frame}} \in \mathbb{R}^{H \times W \times C}$, $\Phi_{\text{warp}} \in \mathbb{R}^{T \times D}$ (each element from a normal distribution $\mathcal{N}(0, 1)$)
 - 2 $\pi = \text{TRANS-WARP}(\mathbf{r}_{\text{frame}}, \Phi_{\text{warp}})$ /* Use Geometric Transformations */
 - 3 $\mathbf{w}_1 = \mathbf{g}^{(i-1)} + \delta\pi$, $\mathbf{w}_2 = \mathbf{g}^{(i-1)} - \delta\pi$
 - 4 $L_1 = \mathcal{L}(\mathbf{x}^{(i-1)} + \epsilon\mathbf{w}_2, y)$, $L_2 = \mathcal{L}(\mathbf{x}^{(i-1)} + \epsilon\mathbf{w}_1, y)$ /* Query victim model twice */
 - 5 $\Delta = (L_2 - L_1)\pi/\epsilon\delta$
 - 6 **return** Δ
-

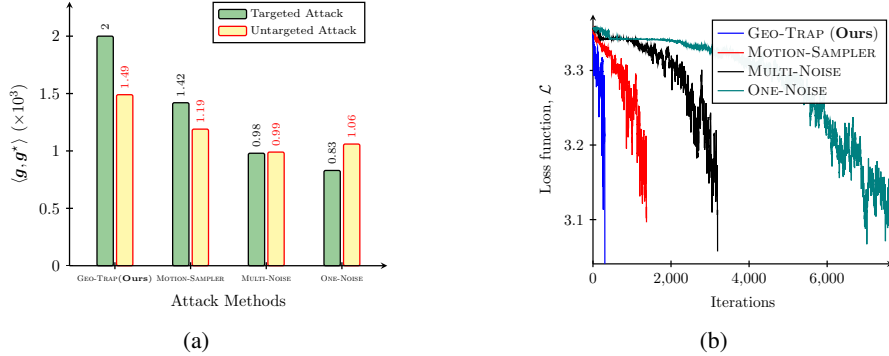


Figure 2: **Gradient Analysis of GEO-TRAP.** (a) GEO-TRAP’s high query-efficiency is a direct implication of good quality gradient estimation (for both targeted and untargeted attack), shown here with higher cosine similarity with \mathbf{g}^* compared to other methods. (b) Better quality of estimated gradients by GEO-TRAP results in a successful attack with fewer queries compared to other attacks.

- **MULTI-NOISE ATTACK** [13] which computes search directions \mathbf{r}_t separately for each frame by sampling each element of \mathbf{r}_t from a standard normal distribution, resulting in a search space dimension of $T \times H \times W \times C$. It does not explicitly consider the temporal dimension; temporal progression in any arbitrary direction is possible between a sequence of perturbation frames.
- **ONE-NOISE ATTACK** which computes \mathbf{r}_1 by sampling each element from a standard normal distribution and applies the same \mathbf{r}_1 across all $\mathbf{r}_t (t = 1, 2, \dots, T)$. ONE-NOISE ATTACK reduces the search space but completely ignores the temporal dimension when generating the perturbation.
- **MOTION-SAMPLER ATTACK** [22] which uses the optical flow of the original video \mathbf{x} to warp $\mathbf{r}_{\text{frame}}$ to get each \mathbf{r}_t . It reduces the search space by using the motion prior of \mathbf{x} as the temporal progression between perturbation frames. In contrast, rather than fixing the temporal search space using a motion prior, GEO-TRAP parameterizes the temporal structure of the space with Φ_{warp} .

We measure the gradient estimation quality by calculating the cosine similarity between the ground truth \mathbf{g}^* and the estimated gradient \mathbf{g} following [19] for the aforementioned baselines. For each attack, we average over 1000 randomly selected videos with their cosine similarity values in the first attack iteration. We choose the first iteration because the initial \mathbf{g}^* is the same for the different attack methods, ensuring a fair comparison. As shown in Figure 2(a), our proposed method for estimating the gradients, yields \mathbf{g} of the best quality for both untargeted and targeted attacks among all evaluated approaches. This leads to faster loss convergence / few queries as shown in Figure 2(b). We validate such trends with different loss functions and more datasets in the Supplementary Material.

The empirical results validate that by carefully considering the temporal dimension and parameterizing the temporal structure of the gradient search space with geometric transformations, GEO-TRAP finds better gradients. GEO-TRAP and MOTION-SAMPLER ATTACK [22] are better than the other two; the reason could be that temporally structured perturbations are more likely to disrupt the motion context of videos. However, the gradients estimated by MOTION-SAMPLER ATTACK [22] are not as effective as our proposed approach; the reason could be that the motion-prior of the clean video does not necessarily represent the temporal behavior of effective video perturbations. By allowing flexibility of the temporal progression while maintaining only a minimally sufficient space through its geometric parameterization, GEO-TRAP generates effective temporally structured perturbations. Note that one could use other, potentially better ways to parameterize the temporal progression of the video perturbation; this is regarded as future works.

Algorithm 2.1: TRANS-WARP($\mathbf{r}_{\text{frame}} \in \mathbb{R}^{H \times W \times C}$, $\Phi_{\text{warp}} \in \mathbb{R}^{T \times D}$) \rightarrow Estimate $\pi \in \mathbb{R}^{T \times H \times W \times C}$

Input : noise tensor $\mathbf{r}_{\text{frame}}$, warp tensors Φ_{warp} , transformation operation $\mathcal{T}_{\phi}(\cdot)$.

Output : candidate directions $\pi = [\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_T]$.

```

1 Initialize  $\pi = \emptyset$ 
2 for  $t = [1, 2, \dots, T]$  do
3    $\phi_t = \Phi_{\text{warp}}[t]$ 
4    $\mathbf{r}_t = \mathcal{T}(\mathbf{r}_{\text{frame}}, \phi_t)$  /* Warping Operation */
5    $\pi \leftarrow$  append  $\mathbf{r}_t$ 
6 end
7 return  $\pi$ 

```

5 Experiments

Datasets. Following previous work like [15], we use the human action recognition dataset UCF-101 [30] and the hand gesture recognition dataset 20BN-JESTER (Jester) [23] to validate our attacks. *UCF-101* includes 13320 videos from 101 human action categories (e.g., applying lipstick, biking, blow drying hair, cutting in the kitchen). Given the diversity it provides, we consider the dataset to validate the feasibility of our attacks on *coarse-grained* actions. *Jester*, on the other hand, includes hand gesture videos that are recorded by crowd-sourced workers performing 27 kinds of gestures (e.g., sliding hand left, sliding two fingers left, zooming in with full hand, zooming out with full hand). The appearance of different hand gestures is similar; it is the motion information that matters in the video classification. We use this dataset to validate our attack with regard to *fine-grained* actions.

Baselines. Among the four state-of-the-art black-box video attack methods [19–22] described in Section 2, we use [21, 22] as baselines for following reasons. Our first baseline is MOTION-SAMPLER ATTACK [22], which has been shown to outperform PATCHATTACK [19], ONE-NOISE and MULTI-NOISE attacks (introduced in Section. 4). Our second baseline is HEURISTICATTACK [20]. We note that SPARSEATTACK [21] is not included in our analysis as we couldn’t replicate their results.

Attack Settings. We consider four state-of-the-art video classification models representing diverse methodologies of learning from videos, i.e., C3D [31], SlowFast [32], TPN [33] and I3D [34], as our black-box victim models to attack. More details about the four video models are provided in the Supplementary Materials. For UCF-101, we randomly select one video from each category following the setting in [19, 22]. For Jester, since the number of categories is small, we randomly select four videos from each category. All attacked videos are correctly classified by the black-box model. For targeted attack, a random target class is chosen for each video. The maximum noise value ρ_{max} is 10 pixel values (out of 255) following [15, 35, 36]. We provide more results for different ρ_{max} in Supplementary Material. Note that since the perturbation generated by HEURISTICATTACK [20] is sparse and thus more imperceptible, we do not impose a perturbation budget on it. We set the maximum query limit to $Q = 60,000$ for untargeted attack and $Q = 200,000$ for targeted attack. The other hyper-parameters, i.e., ϵ , δ , η , and h take the same values as mentioned in [22]. Unless otherwise specified, a translation-dilation transformation (with $D = 3$) is used for our attack method.

Metrics. Following [19, 22], we evaluate GEO-TRAP, in terms of (a) Success Rate (SR), i.e., the total success rate in attacking within query and perturbation budgets; and (b) Average Number of Queries (ANQ) i.e., the average total queries from attacks for all videos (including failed ones).

5.1 Comparison to State-of-the-Art

Untargeted Attack. We report the untargeted attack performance of our attack method and the baseline methods in Table 2. We observe that, in general, GEO-TRAP requires fewer average number of queries when attacking different black-box victim models: on average over 45 % fewer queries than MOTION-SAMPLER ATTACK [22]. At the same time, GEO-TRAP yields higher attack success rates: on average about 6% higher than HEURISTICATTACK [20]. When attacking SlowFast model on the Jester dataset, GEO-TRAP achieves 100% successful rate with only 521 queries while the baseline methods need at least 1906 queries. We also observe that the TPN model is more robust towards black-box attacks compared with the other three video recognition models.




Figure 3: **Visualization of Perturbations and Perturbed Video.** We visualize the generated perturbations and perturbed video for GEO-TRAP and other baselines for UCF-101 (*left*) and Jester (*right*) datasets for untargeted attack against SlowFast classifier with $\rho_{\max} = 10/255$.

Visualization. We show two visualizations of adversarial frames on Jester and UCF-101 in Fig. 3. We observe that the generated adversarial frames have little difference from the clean ones but can lead to a failed classification. Also, our attack method could lead to *sparse* perturbations in the spatial and temporal dimension as the perturbations are sometimes zoomed out (thus get very small), and sometimes are translated out of the sight with choice of geometric transformation. More examples are in the Supplementary Material.

Targeted Attack. We report the targeted attack performance of our method and the baseline methods in Table. 3. We observe that in some cases, HEURISTICATTACK [20] requires fewer number of queries than GEO-TRAP, but its attack success rates are pretty low in those cases. For example, when attacking the TPN model on the Jester dataset, although HEURISTICATTACK [20] requires only 12k average number of queries, its attack success rate is less than half of ours, 44.4% v.s. 92.6%. GEO-TRAP consistently yields higher attack success rates, on average over 30% higher than HEURISTICATTACK [20] and over 8% higher than MOTION-SAMPLER ATTACK [22]. In addition, in most cases, GEO-TRAP requires fewer average number of queries than the two baseline attacks, on average over 45 % fewer queries than MOTION-SAMPLER ATTACK [22]. The targeted attack performance further validates the effectiveness of our method.

5.2 Different Geometric Transformations in TRANS-WARP

As discussed in Section 3.2, different kinds of geometric transformations could be used in the TRANS-WARP function. In addition to the translation-dilation transformation ($\mathcal{M}_{\phi}^{\text{TD}}$ in (8), $D = 3$) employed throughout the paper, we report the performance of GEO-TRAP with two other different geometric

Table 2: **Untargeted Attacks.** GEO-TRAP demonstrates highly successful untargeted attacks (high Success Rate (SR)) with fewer queries (low Average Number of Queries (ANQ))

Datasets	Methods	Black-box Video Classifiers							
		C3D		SlowFast		TPN		I3D	
		ANQ (\downarrow)	SR (\uparrow)	ANQ (\downarrow)	SR (\uparrow)	ANQ (\downarrow)	SR (\uparrow)	ANQ (\downarrow)	SR (\uparrow)
Jester	HEURISTICATTACK [20]	4699	99.0%	3572	98.1%	4679	82.0%	4248	98.1%
	MOTION-SAMPLER ATTACK [22]	4549	99.0%	1906	100%	6269	91.3%	3029	99.4%
	GEO-TRAP (Ours)	1602	100%	521	100%	3315	92.4%	1599	100%
UCF-101	HEURISTICATTACK [20]	5206	70.2%	3507	87.2%	6539	71.8%	6949	84.7%
	MOTION-SAMPLER ATTACK [22]	14336	81.6%	4673	97.2%	20369	75.8%	7400	94.4%
	GEO-TRAP (Ours)	11490	86.2%	1547	98.8%	17716	76.1%	4887	97.4%

Table 3: **Targeted Attacks.** GEO-TRAP demonstrates highly successful targeted attacks (high Success Rate (SR)) with fewer queries (low Average Number of Queries (ANQ))

Datasets	Methods	Black-box Video Classifiers							
		C3D		SlowFast		TPN		I3D	
		ANQ (↓)	SR (↑)	ANQ (↓)	SR (↑)	ANQ (↓)	SR (↑)	ANQ (↓)	SR (↑)
Jester	HEURISTICATTACK [20]	15595	46.3%	30768	98.1%	12006	44.4%	31088	77.8%
	MOTION-SAMPLER ATTACK [22]	26704	98.2%	33087	100%	63721	80.9%	39037	90.7%
	GEO-TRAP (Ours)	6198	100%	7788	100%	41294	92.6%	19542	98.2%
UCF-101	HEURISTICATTACK [20]	26741	29.0%	22152	61.4%	71828	36.4%	92244	43.7%
	MOTION-SAMPLER ATTACK [22]	100467	71.1%	57126	86.0%	151409	31.6%	96498	59.6%
	GEO-TRAP (Ours)	71820	85.8%	21878	95.0%	141629	40.0%	76708	74.6%

transformations, i.e., similarity transformation (\mathcal{M}_ϕ^S in (8), $D = 4$) and affine transformation (\mathcal{M}_ϕ in (7), $D = 6$). Figure 4 shows the untargeted attack performance on Jester with these different geometric transformations (more results are available in the Supplementary Material). We observe that the transformation with fewer degrees of freedom (DOF) (translation-dilation transformation) tends to require fewer queries while having the same or higher attack success rates (the attack success rates are available in the Supplementary Material). We believe that $D = 3$ provides enough temporal flexibility to disrupt the motion context of the videos; additional degrees of freedom seemingly increase the search space unnecessarily, resulting in more queries.

6 Conclusion

Black-box adversarial attacks on video classifiers is a challenging problem that has been largely understudied. In this work, we demonstrate that searching for effective gradients in a reduced but structured search space for crafting perturbations leads to highly successful attacks with fewer queries compared to state-of-the-art attack strategies. In particular, we propose a novel iterative algorithm that employs Geometric transformations to parameterize and reduce the search space, for estimating gradients that maximize the probability of mis-classification of the perturbed video. This simple and novel strategy exposes the vulnerability of widely used video classifiers. For instance, GEO-TRAP decreases average query numbers by 64.78%, 72.66% and 47.21% to attack C3D, SlowFast, and I3D, respectively, for almost 100% success rate in untargeted attacks.

7 Broader Impact

In this work, by leveraging geometric transformations for effective gradient estimations, we propose a highly query-efficient adversarial attack on video classification models which demonstrates state-of-the-art results. As more and more safety-critical systems (e.g., perceptual modules in autonomous vehicles) nowadays rely on video models, we are hopeful that our work, in addition to future research (including designing sophisticated video generative models as in [37–39] for distribution-driven attacks [40, 41]), can eventually help build sufficiently robust video models to best avoid malicious sub-versions. On one hand, we believe our algorithm could allow further research in adversarial robustness and data augmentation strategies of deep vision models. It should also give a direction to researchers to design counter defense methodologies [42–46]. On the other hand, it highlights a key drawback of different video classifiers which will allow adversaries to design more sophisticated attacks, both in white-box and black-box settings. Addressing such fallacies in designing deep neural networks is of utmost importance before introducing them in real-world scenarios.

Acknowledgement. The authors would like to thank Dr. Cliff Wang of US Army Research Office for his extensive comments and input on this work. This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Agreement No. HR00112090096. Approved for public release; distribution is unlimited.

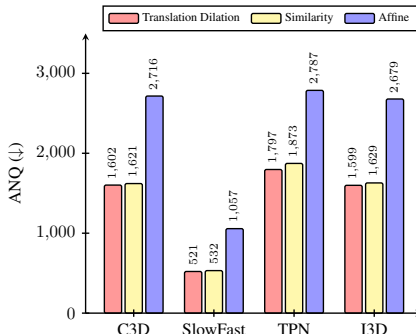


Figure 4: **Performance with different \mathcal{M}_ϕ .** GEO-TRAP results in best performance when \mathcal{M}_ϕ is set as translation-dilation operation.

References

- [1] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3d convolutional neural networks for automatic human action recognition, January 1 2013. US Patent 8,345,984.
- [2] Wongun Choi, Yuanqing Lin, Yu Xiang, and Silvio Savarese. Subcategory-aware convolutional neural networks for object detection, May 8 2018. US Patent 9,965,719.
- [3] Abhishek Aich, Meng Zheng, Srikrishna Karanam, Terrence Chen, Amit K Roy-Chowdhury, and Ziyang Wu. Spatio-temporal representation factorization for video-based person re-identification. *arXiv preprint arXiv:2107.11878*, 2021.
- [4] 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.
- [5] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and Harnessing Adversarial Examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [6] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 39–57. IEEE, 2017.
- [7] Dawn Song, Kevin Eykholt, Ivan Evtimov, Earlene Fernandes, Bo Li, Amir Rahmati, Florian Tramèr, Atul Prakash, and Tadayoshi Kohno. Physical adversarial examples for object detectors. In *12th {USENIX} Workshop on Offensive Technologies ({WOOT} 18)*, 2018.
- [8] Chaowei Xiao, Ruizhi Deng, Bo Li, Fisher Yu, Mingyan Liu, and Dawn Song. Characterizing adversarial examples based on spatial consistency information for semantic segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 217–234, 2018.
- [9] Nicolas Papernot, Patrick McDaniel, and Ian Goodfellow. Transferability in machine learning: from phenomena to black-box attacks using adversarial samples. *arXiv preprint arXiv:1605.07277*, 2016.
- [10] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical black-box attacks against machine learning. In *Proceedings of the 2017 ACM on Asia conference on computer and communications security*, pages 506–519, 2017.
- [11] Arjun Nitin Bhagoji, Warren He, Bo Li, and Dawn Song. Practical black-box attacks on deep neural networks using efficient query mechanisms. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 154–169, 2018.
- [12] Pin-Yu Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui Hsieh. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *Proceedings of the 10th ACM workshop on artificial intelligence and security*, pages 15–26, 2017.
- [13] Andrew Ilyas, Logan Engstrom, and Aleksander Madry. Prior convictions: Black-box adversarial attacks with bandits and priors. *arXiv preprint arXiv:1807.07978*, 2018.
- [14] Xingxing Wei, Jun Zhu, Sha Yuan, and Hang Su. Sparse adversarial perturbations for videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8973–8980, 2019.
- [15] Shasha Li, Ajaya Neupane, Sujoy Paul, Chengyu Song, Srikanth V Krishnamurthy, Amit K Roy-Chowdhury, and Ananthram Swami. Stealthy adversarial perturbations against real-time video classification systems. In *NDSS*, 2019.
- [16] Zhikai Chen, Lingxi Xie, Shanmin Pang, Yong He, and Qi Tian. Appending adversarial frames for universal video attack. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3199–3208, 2021.
- [17] Roi Pony, Itay Naeh, and Shie Mannor. Over-the-air adversarial flickering attacks against video recognition networks. *arXiv preprint arXiv:2002.05123*, 2020.
- [18] Shao-Yuan Lo and Vishal M Patel. Multav: Multiplicative adversarial videos. *arXiv preprint arXiv:2009.08058*, 2020.
- [19] 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.

- [20] 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*, pages 12338–12345, 2020.
- [21] Huanqian Yan, Xingxing Wei, and Bo Li. Sparse black-box video attack with reinforcement learning. *arXiv preprint arXiv:2001.03754*, 2020.
- [22] Hu Zhang, Linchao Zhu, Yi Zhu, and Yi Yang. Motion-excited sampler: Video adversarial attack with sparked prior. In *European Conference on Computer Vision*. Springer, 2020.
- [23] Joanna Materzynska, Guillaume Berger, Ingo Bax, and Roland Memisevic. The jester dataset: A large-scale video dataset of human gestures. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019.
- [24] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into transferable adversarial examples and black-box attacks. *ICLR*, 2017.
- [25] 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.
- [26] Andrew Ilyas, Logan Engstrom, Anish Athalye, and Jessy Lin. Black-box adversarial attacks with limited queries and information. In *International Conference on Machine Learning*, pages 2137–2146. PMLR, 2018.
- [27] Alexey Kurakin, Ian Goodfellow, Samy Bengio, et al. Adversarial examples in the physical world, 2017.
- [28] Hongyu Ren, Shengjia Zhao, and Stefano Ermon. Adaptive antithetic sampling for variance reduction. In *International Conference on Machine Learning*, pages 5420–5428. PMLR, 2019.
- [29] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521540518, second edition, 2004.
- [30] 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.
- [31] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015.
- [32] 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.
- [33] 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.
- [34] 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.
- [35] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. Universal adversarial perturbations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1765–1773, 2017.
- [36] Konda Reddy Mopuri, Utkarsh Ojha, Utsav Garg, and R Venkatesh Babu. Nag: Network for adversary generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 742–751, 2018.
- [37] Abhishek Aich, Akash Gupta, Rameswar Panda, Rakib Hyder, M Salman Asif, and Amit K Roy-Chowdhury. Non-adversarial video synthesis with learned priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6090–6099, 2020.
- [38] Akash Gupta, Abhishek Aich, and Amit K. Roy-Chowdhury. Alanet: Adaptive latent attention network for joint video deblurring and interpolation. In *Proceedings of the 28th ACM International Conference on Multimedia*, page 256–264, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379885.

- [39] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. MoCoGAN: Decomposing motion and content for video generation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1526–1535, 2018.
- [40] Muzammal Naseer, Salman H Khan, Harris Khan, Fahad Shahbaz Khan, and Fatih Porikli. Cross-Domain Transferability of Adversarial Perturbations. *arXiv preprint arXiv:1905.11736*, 2019.
- [41] Omid Poursaeed, Isay Katsman, Bicheng Gao, and Serge Belongie. Generative Adversarial Perturbations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4422–4431, 2018.
- [42] Xuwang Yin, Soheil Kolouri, and Gustavo K Rohde. Gat: Generative adversarial training for adversarial example detection and robust classification. In *International Conference on Learning Representations*, 2019.
- [43] Yisen Wang, Xingjun Ma, James Bailey, Jinfeng Yi, Bowen Zhou, and Quanquan Gu. On the convergence and robustness of adversarial training. In *ICML*, volume 1, page 2, 2019.
- [44] Shasha Li, Shitong Zhu, Sudipta Paul, Amit Roy-Chowdhury, Chengyu Song, Srikanth Krishnamurthy, Ananthram Swami, and Kevin S Chan. Connecting the dots: Detecting adversarial perturbations using context inconsistency. In *European Conference on Computer Vision*, pages 396–413. Springer, 2020.
- [45] Xueping Wang, Shasha Li, Min Liu, Yaonan Wang, and Amit K Roy-Chowdhury. Multi-expert adversarial attack detection in person re-identification using context inconsistency. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.
- [46] Mingjun Yin, Shasha Li, Zikui Cai, Chengyu Song, M Salman Asif, Amit K Roy-Chowdhury, and Srikanth V Krishnamurthy. Exploiting multi-object relationships for detecting adversarial attacks in complex scenes. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.