

Towards Transferable Targeted 3D Adversarial Attack in the Physical World

Yao Huang^{1,3}, Yinpeng Dong^{2,3†}, Shouwei Ruan¹, Xiao Yang², Hang Su², Xingxing Wei^{1†}

¹ Institute of Artificial Intelligence, Beihang University, Beijing 100191, China

² Dept. of Comp. Sci. and Tech., Institute for AI, Tsinghua-Bosch Joint ML Center,

THBI Lab, BNRist Center, Tsinghua University, Beijing, 100084, China ³ RealAI

{y-huang, shouweiruan, xxwei}@buaa.edu.cn {dongyinpeng, suhangss}@email.tsinghua.edu.cn

{yangxiao19}@mails.tsinghua.edu.cn

Abstract

Compared with transferable untargeted attacks, transferable targeted adversarial attacks could specify the misclassification categories of adversarial samples, posing a greater threat to security-critical tasks. In the meanwhile, 3D adversarial samples, due to their potential of multi-view robustness, can more comprehensively identify weaknesses in existing deep learning systems, possessing great application value. However, the field of transferable targeted 3D adversarial attacks remains vacant. The goal of this work is to develop a more effective technique that could generate transferable targeted 3D adversarial examples, filling the gap in this field. To achieve this goal, we design a novel framework named **TT3D** that could rapidly reconstruct from few multi-view images into **Transferable Targeted 3D** textured meshes. While existing mesh-based texture optimization methods compute gradients in the high-dimensional mesh space and easily fall into local optima, leading to unsatisfactory transferability and distinct distortions, **TT3D** innovatively performs dual optimization towards both feature grid and Multi-layer Perceptron (MLP) parameters in the grid-based NeRF space, which significantly enhances black-box transferability while enjoying naturalness. Experimental results show that **TT3D** not only exhibits superior cross-model transferability but also maintains considerable adaptability across different renders and vision tasks. More importantly, we produce 3D adversarial examples with 3D printing techniques in the real world and verify their robust performance under various scenarios.

1. Introduction

Despite the unprecedented performance of deep neural networks (DNNs) in numerous tasks, including image classification [27] and object detection [13, 54], these models are

[†]Corresponding authors.

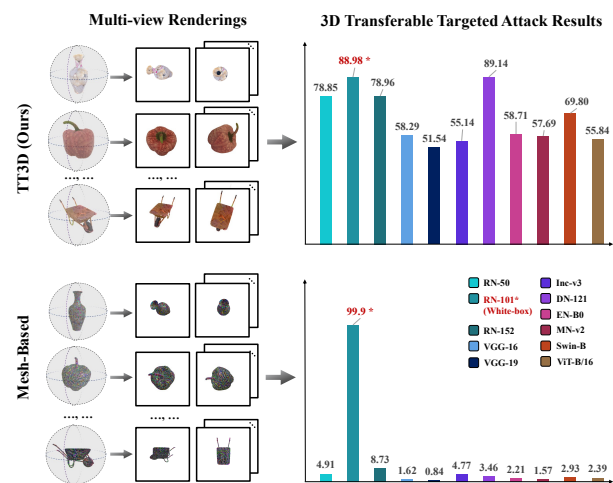


Figure 1. A comparison of transferable targeted attack performance in the 3D domain between our TT3D and the enhanced version of the typical mesh-based optimization method [49], as detailed in Sec. 4.2.1. The surrogate model is ResNet-101 [10] and we can see TT3D shows remarkable transferability.

vulnerable to adversarial examples [4, 5, 9, 20, 22, 29, 35, 39, 43–48, 50]. These adversarial examples are shown to have good transferability across models [4, 17, 50], emerging as a key concern since they can evade black-box models in practical applications. The majority of transferable attacks are untargeted, aiming to induce misclassification of the target model. However, transferable targeted attacks, which mislead DNNs to produce predetermined misclassifications, can lead to a more severe threat in real-world applications. Crafting such transferable targeted attacks is particularly challenging because they not only need to achieve a specific misclassification but also must avoid overfitting to ensure effective attacks across various models, which requires further investigation.

While transferable targeted adversarial attacks in the 2D domain have been extensively studied [15, 24, 42, 53], the

	AdvTurtle [1]	Zenget <i>al.</i> [52]	MeshAdv [49]	Oslundet <i>al.</i> [26]	FCA [40]	DTA [33]	ACTIVE [34]	TT3D (Ours)
Target Task	Classify	Classify	Classify	Classify	Detect	Detect	Detect+Segment*	Classify+Detect*+Caption*
Data	3D model	3D model	3D model	3D model	3D model	3D model	3D model	Few Multi-view images
3D Attack Type	Texture	Texture	Geometry/Texture	Texture	Texture	Texture	Texture	Texture+Geometry
Targeted	✓	✗	✓	✗	✗	✗	✗	✓
Transferability	✗	✗	✗	✓	✓	✓	✓	✓
Naturalness	✗	✗	✗	✗	✓	✗	✓	✓
Physical Attack	✓	✗	✗	✓	✗	✓	✓	✓

Table 1. A comparison among different 3D attack methods regarding target vision tasks, used data, 3D attack types, transferability, whether conducting targeted attacks, naturalness, and whether conducting attacks in the physical world. * represents the transferred task .

exploration of transferable targeted attacks in the 3D domain remains vacant. Compared with 2D adversarial attacks, 3D attacks possess greater practical value in the real world due to the potential of consistent attacks across various viewpoints. As illustrated in Tab. 1, existing 3D attack methods [1, 6, 33, 34, 40, 49, 52] could ensure multi-view effectiveness by modifying pre-existing meshes’ textures or geometry due to meshes’ 3D consistency. But, they struggle to simultaneously ensure transferability and targeted attack, especially for physical attacks. In addition, it is also difficult to maintain the naturalness of these adversarial examples.

Based on the above discussions, this paper aims to generate **transferable and natural 3D adversarial examples for physically targeted attacks**. However, there exist two challenges to meet our goal: (1) Prior mesh-based optimization methods involve direct alterations to vertex colors in the high-dimensional mesh space, easily falling into the overfitting and thus leading to unsatisfactory transferability. Thus, the first challenge is how to design an optimization method that deviates from the mesh space, avoiding overfitting for better transferability. (2) Existing methods struggle to simultaneously ensure attack performance and naturalness, whose attack process easily accompanies the extremely unnatural phenomena, such as visual anomalies in appearance or distortion or fragmentation of the 3D mesh, *etc.* Therefore, how to balance the attack performance and the visual naturalness is another challenge.

To meet the above challenges, we design a novel framework called **TT3D** that could rapidly reconstruct a transferable targeted 3D adversarial textured mesh with guaranteed naturalness. To be specific, driven by the recent advancement in multi-view reconstruction techniques (*i.e.*, grid-based NeRF [23]), we first reconstruct an initial 3D mesh from its multi-view images, but instead of directly manipulating in the mesh space, we innovatively perform adversarial fine-tuning of appearance rendering parameters in the grid-based NeRF space. Such a method not only successfully avoids overfitting but also eliminates previous methods’ reliance on the pre-existing 3D mesh, greatly lowering the cost. The technical details can be found as follows:

Firstly, to enhance the transferability, we design a dual optimization strategy for adversarial fine-tuning. This op-

timization strategy simultaneously targets the parameters of the appearance feature grid and the corresponding MLP (feature grids and MLP are both components of grid-based NeRF, detailed in Sec. 3.1), thereby effectively perturbing at both the foundational feature level and the more advanced decision-making layers. Additionally, it’s worth mentioning that, though solely altering the geometric structure is challenging for achieving transferable targeted attacks, we observe that incorporating changes in geometric information under a texture optimization-focused method, can enhance the performance (shown in Sec. 4.2.1). Thus, we add the optimization of vertex coordinates into our methodology.

Secondly, to ensure the naturalness of the 3D adversarial textured mesh, we implement constraints on appearance and geometry during optimization. These include measures for visual consistency and vertex proximity, along with specialized loss functions to prevent mesh deformities and ensure surface smoothness. For 3D adversarial objects’ robustness in the real world, we also introduce an EOT that could effectively integrate various transformations in both 3D and 2D spaces to more realistically simulate real-world conditions. A comparison of our TT3D’s performance in transferable targeted attacks with the typical mesh-based optimization method is shown in Fig. 1.

The contributions of this paper are as follows:

- We propose a novel framework called **TT3D** for generating transferable targeted 3D adversarial examples, which is the first work to fill a critical gap in this field and eliminate the previous reliance on pre-existing 3D meshes, expanding the feasible region of 3D attacks.
- We design a dual optimization strategy within the grid-based NeRF space, which effectively perturbs at both the foundational feature level and the more sophisticated decision-making echelons of the neural network, and ensures the naturalness in the meanwhile.
- Experimental results demonstrate that our 3D adversarial object can easily be misclassified as a given label across various victim models, renders, and tasks. Moreover, we produce 3D adversarial objects in the real world by utilizing 3D printing technology and validate their robust performance under various settings, like different viewpoints and backgrounds.

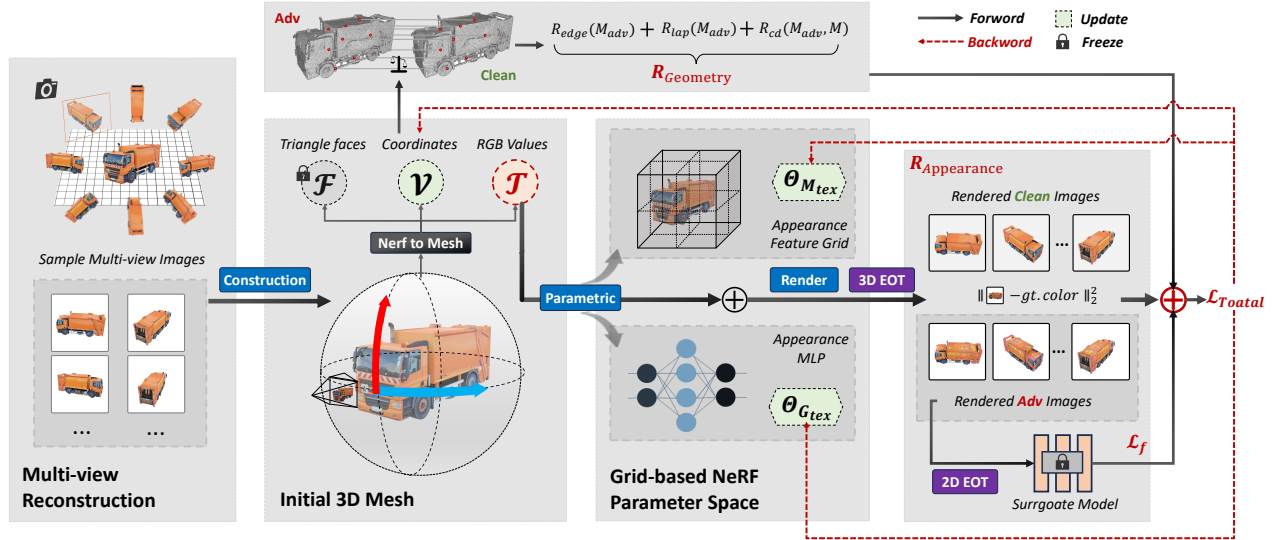


Figure 2. An overview of our TT3D framework. We first utilize 3D multi-view reconstruction technology, *i.e.*, grid-based NeRF with marching cubes techniques to obtain the initial clean 3D mesh. Then, we perform adversarial fine-tuning in the textual parameter space of grid-based NeRF instead of directly altering the texture \mathcal{T} , supplemented with geometric perturbations at vertex positions \mathcal{V} . To ensure the naturalness simultaneously, we add constraints to the distance between the 3D adversarial samples and the initial ones in terms of both texture and geometric structure when performing optimization .

2. Related Work

2.1. Transferable Targeted Adversarial Attack

Transferable targeted adversarial attacks in the 2D domain have been extensively studied. For instance, Li *et al.* [15] improve the transferability of targeted attacks by addressing ‘noise curing’ through self-adaptive gradient magnitudes and metric learning. Zhao *et al.* [53] demonstrate that simple logit loss-based attacks, without extensive resources, can achieve unexpected effectiveness in targeted transferability, challenging traditional, resource-intensive methods. The works of Wang *et al.* [42] and Naseer *et al.* [24] further advance the field, introducing sophisticated approaches for capturing class distributions and aligning image distributions. Moreover, Yang *et al.* [51] also propose a hierarchical generative network-based approach for crafting targeted transfer-based adversarial examples, showcasing the potential of generative models in this domain. However, the above works only focus on the 2D domain while the realm of transferable targeted 3D adversarial attacks remains vacant, posing a necessity for investigations.

2.2. 3D Adversarial Attack

Since Athalye *et al.* [1] verified the existence of 3D adversarial examples, plenty of generation methods [26, 33, 34, 40, 49, 52] have emerged. Among them, except for MeshAdv [49], which has attempted to change the mesh geometry for 3D attacks but still fails to perform transferable targeted attacks, the rest methods [26, 33, 34, 40, 52], including MeshAdv (optional), are based on modifying the

texture of the 3D mesh. We can split texture-based attack methods into two categories. For methods like [26, 49, 52], they directly alter vertex colors in the high-dimensional mesh space, leading to unavoidable overfitting. On the other hand, works like [33, 34, 40] optimize a texture map and render it onto a specified 3D model (*e.g.*, a car), enhancing transferability to some extent. However, their transferability is limited to non-targeted attacks in object detection, lacking comparability. Moreover, most of these attacks are confined to the digital realm, and physical attacks often come with visual unnaturalness. To address these issues, our work, without the reliance on 3D models, can effectively perform an 3D attack that not only ensures attack performance (*i.e.*, transferable targeted 3D attacks feasible even in real-world scenarios), but also could meet demands for naturalness.

3. Methodology

In this section, we detail the proposed TT3D framework for generating transferable targeted 3D adversarial examples, as presented in Fig. 2. TT3D adopts grid-based NeRF [23] to model 3D objects from multi-view images and perform optimization. As below, we first introduce the background knowledge of grid-based NeRF and then present the problem formulation and our optimization strategy.

3.1. Preliminary

Given a set of multi-view images, vanilla NeRF [21] encodes a real-world object into a continuous volumetric radiance field $F : (\mathbf{x}, \mathbf{d}) \rightarrow (\mathbf{c}, \sigma)$. F is approximated by a multi-layer perceptron (MLP), which takes a 3D location

$\mathbf{x} \in \mathbb{R}^3$ and a unit-norm viewing direction $\mathbf{d} \in \mathbb{R}^3$ as inputs, and outputs a volume density $\sigma \in \mathbb{R}^+$ and emitted RGB color $\mathbf{c} \in [0, 1]^3$. Despite its efficacy, vanilla NeRF often faces challenges in terms of computational efficiency.

In addressing the limitations of traditional methods, grid-based NeRF techniques [8, 23, 32, 38] employ two distinct 3D grids, namely G_{geo} and G_{tex} , to explicitly represent geometric and textural information of objects. This structured approach, utilizing feature grids, maps a 3D point \mathbf{x} to corresponding feature vectors \mathbf{f}_{geo} and \mathbf{f}_{tex} for both geometry and texture, which are expressed as:

$$\mathbf{f}_{geo}(\mathbf{x}) = G_{geo}(\mathbf{x}; \Theta_{G_{geo}}), \mathbf{f}_{tex}(\mathbf{x}) = G_{tex}(\mathbf{x}; \Theta_{G_{tex}}), \quad (1)$$

where $\Theta_{G_{geo}}$ and $\Theta_{G_{tex}}$ denote the sets of feature vectors within G_{geo} and G_{tex} , respectively.

To derive the final rendering attributes, *i.e.*, volume density σ and emitted color \mathbf{c} , these feature vectors \mathbf{f}_{geo} and \mathbf{f}_{tex} are then processed through two shallow MLPs: M_{geo} and M_{tex} , formulated as:

$$\sigma \leftarrow M_{geo}(\mathbf{f}_{geo}(\mathbf{x}); \Theta_{M_{geo}}), \mathbf{c} \leftarrow M_{tex}(\mathbf{f}_{tex}(\mathbf{x}); \Theta_{M_{tex}}), \quad (2)$$

where $\Theta_{M_{geo}}$ and $\Theta_{M_{tex}}$ are weights of M_{geo} and M_{tex} .

In this paper, by leveraging grid-based NeRF [23] with Marching Cubes [19], we first efficiently reconstruct 3D mesh as an accurate geometry and continue adversarial fine-tuning within the parameter space of grid-based NeRF.

3.2. Problem Formulation

For the 3D adversarial attack task, we aim to develop an effective method that can generate transferable targeted 3D adversarial samples and maintain their visual naturalness. Similar to [1], we also propose to craft 3D adversarial examples as textured mesh, which can fully leverage 3D printing techniques for physically realizable adversarial attacks.

Specifically, we define the reconstructed mesh representation of the 3D object as $\mathcal{M} = (\mathcal{V}, \mathcal{T}, \mathcal{F})$, where $\mathcal{V} \in \mathbb{R}^{n \times 3}$ is the xyz coordinates of n vertices, $\mathcal{T} \in \mathbb{R}^{n \times 3}$ is the rgb value of vertices, and $\mathcal{F} \in \mathbb{Z}^{m \times 3}$ is the set of m triangle faces which encodes each triangle with the indices of vertices. Here, we do not pose changes to the mesh topology \mathcal{F} to prevent serious distortions of geometric structures, and instead, we turn to \mathcal{V} and \mathcal{T} .

In this paper, we focus on the challenging transferable targeted 3D attacks against image classification models under different views. Given a surrogate classifier $f: \mathcal{X} \rightarrow \mathcal{Y}$, the goal of the attack is to generate a 3D adversarial example $\mathcal{M}_{adv} = (\mathcal{V}^*, \mathcal{T}^*, \mathcal{F})$ for the original one \mathcal{M} with color and vertex perturbations. Then, with a differentiable rendering function [12] and a random viewpoint \mathbf{v} , we render the \mathcal{M}_{adv} into the corresponding 2D image $\hat{I}_{\mathbf{v}}(\mathcal{M}_{adv})$, which can be misclassified by other common classifiers as the target class $y^* (\neq y)$. In general, the perturbation should be

small to make our 3D adversarial example natural in the physical world. Thus, the optimization problem of crafting 3D adversarial examples can be formulated as

$$\min_{\mathcal{V}^*, \mathcal{T}^*} \mathbb{E}_{\mathbf{v} \in V} L_f(\hat{I}_{\mathbf{v}}(\mathcal{M}_{adv}), y^*) + \beta \cdot \mathcal{R}(\mathcal{M}_{adv}, \mathcal{M}), \quad (3)$$

$$\text{where } \hat{I}_{\mathbf{v}}(\mathcal{M}_{adv}) = \mathcal{S}(\mathcal{V}^*, \mathcal{T}^*, \mathcal{F}, \mathbf{v}),$$

where $\hat{I}_{\mathbf{v}}(\mathcal{M}_{adv})$ represents the rendered image of \mathcal{M}_{adv} under the viewpoint \mathbf{v} and V is the feasible distribution. L_f is the cross-entropy loss [2] that facilitates the misclassification of $\hat{I}_{\mathbf{v}}(\mathcal{M}_{adv})$ to y^* , \mathcal{R} is the regularization for minimizing a perceptibility distance between \mathcal{M}_{adv} and \mathcal{M} and β is a balancing hyperparameter between these two losses.

To ensure the feasibility of our 3D adversarial example in the real world, we also introduce an EOT scheme (as detailed in Sec. 3.5) to help bridge the gap from “digital” to “physical”. However, the mesh-based optimization by following the objective function (3.2) needs to calculate gradients in high-dimensional mesh space due to thousands of points in each 3D object. It will easily distort the natural textural appearance of the textured mesh and is trapped into the overfitting with unsatisfactory transferability.

3.3. Dual Optimization in NeRF Parameter Space

In this section, we aim to deviate from the existing mesh-based optimization regime and perform the optimization trajectory in the parameter space of grid-based NeRF [23] as a regularization for escaping from overfitting.

Specifically, as introduced in Sec. 3.1, grid-based NeRF stores appearance feature vectors in a structured feature grid $G_{tex}(\cdot; \Theta_{G_{tex}})$ and processes them into emitted colors through a shallow MLP $M_{tex}(\cdot; \Theta_{M_{tex}})$. Actually, the clean textural parameter set \mathcal{T} , representing the colors of all vertices in the initial 3D textured mesh \mathcal{M} , are estimated based on the aforementioned mechanism, expressed as follows:

$$\mathcal{T} = \{\mathbf{c} | \mathbf{c} = M_{tex}(\mathbf{f}_{tex}(\mathbf{x}); \Theta_{M_{tex}}), \forall \mathbf{x} \in \mathcal{V}\}, \quad (4)$$

where $\mathbf{f}_{tex}(\mathbf{x})$ is the textual feature vector of \mathbf{x} and mapped by $G_{tex}(\cdot; \Theta_{G_{tex}})$ following Eq. (1).

Existing mesh-based optimization methods, which primarily perform direct alterations to the vertex colors in \mathcal{T} , struggle with the computing complexity, unsatisfactory transferability, and naturalness inherent in explicit mesh spaces. To prevent this, our method innovatively utilizes the previously established color estimation mechanism, *i.e.*, Eq. (4) to perform adversarial fine-tuning in the space of grid-based NeRF, indirectly generating adversarial textures. Moreover, we undertake a dual optimization, concurrently targeting both the parameters of the appearance feature grid, $\Theta_{G_{tex}}$, and those of the corresponding MLP, $\Theta_{M_{tex}}$, which effectively integrates the manipulation of detailed textural features contained within the grids with the nuanced processing capabilities inherent in the MLPs.

In essence, such a dual optimization strategy facilitates the embedding of adversarial perturbations at both the foundational feature level and the more sophisticated decision-making echelons of the neural network, thus yielding more transferable 3D adversarial examples while also ensuring naturalness. It is also noteworthy that though merely optimizing the geometric shape proves challenging for achieving transferable targeted 3D attacks, optimizing the texture in conjunction with adjustments to the geometric shape (*i.e.*, mesh vertex coordinates) can enhance transferability to a certain extent, as verified in Sec. 4.2.1. However, fine-tuning the geometric parameters in the grid-based NeRF space will lead to significant deformation, and ensuring naturalness by restricting the distance to the original mesh’s geometry in such a space would require repeatedly using the marching cubes [19] technique to convert geometric parameters into mesh during optimization. This incurs a substantial time cost. Therefore, we make a compromise that explicitly modifies the mesh vertex coordinates.

To sum up, we target the grid-based NeRF’s inner parameters: $(\Theta_{G_{tex}}, \Theta_{M_{tex}})$ instead of vertex colors \mathcal{T} in appearance, accompanying the optimization to mesh’s vertex coordinates, thus the objective function (3.2) becomes:

$$\min_{\mathcal{V}^*, \Theta_{G_{tex}}^*, \Theta_{M_{tex}}^*} \mathbb{E}_{\mathbf{v} \in V} L_f(\hat{I}_{\mathbf{v}}(\mathcal{M}_{adv}), y^*) + \beta \cdot \mathcal{R}(\mathcal{M}_{adv}, \mathcal{M}). \quad (5)$$

3.4. Regularization for Naturalness

To ensure the perceptible naturalness of 3D adversarial examples, we impose constraints to both appearance and geometry. For the constraint of appearance, we evaluate the disparity between images rendered under various viewpoints from the adversarial textured mesh \mathcal{M}_{adv} and those rendered from the initial clean textured mesh \mathcal{M} , which is achieved by calculating the squared distance between these two sets of images. This appearance-related constraint, denoted as R_{rgb} , ensures that the adversarial examples, while effective, do not deviate significantly in appearance from the original mesh and can be formulated as follows:

$$R_{rgb}(\mathcal{M}_{adv}, \mathcal{M}) = \frac{1}{N} \sum_{\mathbf{v} \in V} \left| \hat{I}_{\mathbf{v}}(\mathcal{M}_{adv}) - \hat{I}_{\mathbf{v}}(\mathcal{M}) \right|^2, \quad (6)$$

where N is the sampled number for one epoch.

Then, referring to the constraint in geometry, we incorporate the Chamfer distance R_{cd} to make vertices of \mathcal{M}_{adv} not far from their initial positions in \mathcal{M} , the Laplacian Smoothing Loss R_{lap} [25] for preventing self-intersecting, and the Mesh Edge Length Loss R_{edge} [41] for the smoothness of the mesh surface. Above all, the whole regularization R could be represented as follows:

$$R(\mathcal{M}_{adv}, \mathcal{M}) = \lambda_1 R_{rgb}(\mathcal{M}_{adv}, \mathcal{M}) + \lambda_2 R_{cd}(\mathcal{M}_{adv}, \mathcal{M}) + \lambda_3 R_{lap}(\mathcal{M}_{adv}) + \lambda_4 R_{edge}(\mathcal{M}_{adv}), \quad (7)$$

where $\lambda_1, \lambda_2, \lambda_3$ and λ_4 are hyperparameters that represent weights of R_{rgb}, R_{cd}, R_{lap} , and R_{edge} , respectively.

3.5. Physical Attack

To realize physically feasible 3D adversarial examples, we need to ensure their robustness against complex transformations in the physical world, including 3D rotations, affine projections, color discrepancies, *etc.* A prevalent technique that we utilize here is the Expectation Over Transformation (EOT) algorithm [1], which optimizes the adversarial example across a distribution of varying transformations in both 2D and 3D spaces. Specifically, we effectively merge diverse 3D transformations, such as pose, distance, and viewpoint shifts, during rendering with 2D transformations such as contrast and blurring on rendered images. With EOT, the rendered images for optimization in Sec. 3.3 can be as below:

$$\hat{I}_{\mathbf{v}}(\mathcal{M}_{adv}) = t(\mathcal{S}(\mathcal{V}^*, \mathcal{T}^*, \mathcal{F}, \rho(\mathbf{v}))), \quad (8)$$

where $t \in T, \rho \in \mathcal{Q}, \mathbf{v} \in V$,

where t and ρ represent the randomly sampled transformations from T and \mathcal{Q} . T and \mathcal{Q} are the sets of various transformations in the 2D space and the 3D space, respectively.

4. Experiments

In this section, we present the experimental results both in the digital world and the physical world to demonstrate the effectiveness of the proposed method.

4.1. Experiment Settings

Datasets. We use the IM3D [28] dataset in our experiments, which contains 1K typical 3D objects from 100 ImageNet [3] categories. For our experimental validation, we choose 100 objects randomly from 30 categories.

Victim models. Two typical classifier models, ResNet-101 [10] and DenseNet-121 [11] are chosen as surrogate models to perform attacks. Other models to test transferability consist of the CNN-based ResNet-50 [10], ResNet-152 [10], VGG-16 [31], VGG19 [31], Inception-v3 [36], EfficientNet-B0 [37], MobileNet-V2 [30] and the Transformer-based Swin-B [18], ViT-B/16 [7]. For testing the cross-render transferability, we choose two commercial rendering software: MeshLab and Blender. For testing the cross-task transferability, we choose two other tasks: zero-shot detection [16] and image caption [14] due to their non-limitation to types of objects.

Evaluation metrics. To quantitatively evaluate the effectiveness of our proposed TT3D, we measure the attack success rate (ASR). Specifically, given that our method is a targeted attack, we judge by whether rendered images can be predicted as the target label by victim models. For each 3D object, we render 100 images using random rendering

Source Model	Methods	Victim Model										
		RN-50	RN-101	RN-152	VGG-16	VGG-19	Inc-v3	DN-121	EN-B0	MN-v2	Swin-B	ViT-B/16
ResNet-101	Mesh-based [49] (enhanced)	4.91	99.90*	8.73	1.62	0.84	4.77	3.46	2.21	1.57	2.93	2.39
	MLP-only	40.48	56.90*	43.87	26.29	22.92	27.81	56.15	25.26	23.11	28.78	35.37
	Grid-only	64.85	80.89*	70.49	42.89	37.29	51.60	80.59	42.54	40.24	50.07	48.66
	MLP+Grid	78.85	88.98*	78.96	58.29	51.54	55.14	89.14	58.71	57.69	69.80	55.84
DenseNet-121	Mesh-based [49] (enhanced)	1.10	2.23	2.55	1.72	1.78	2.76	99.12*	3.91	1.39	2.74	1.68
	MLP-only	24.58	17.68	23.18	19.30	15.69	21.68	58.04*	17.95	13.96	17.70	22.66
	Grid-only	45.41	31.46	39.51	30.76	29.39	37.87	89.62*	25.20	21.39	29.37	28.37
	MLP+Grid	70.98	51.61	59.59	57.59	50.25	43.81	96.74*	40.59	45.06	55.15	36.49

Table 2. The ASR(%) of 3d adversarial examples generated by different methods including the enhanced mesh-based, the mlp-only, the grid-only, and our dual optimization method under random viewpoints against ResNet-50 (RN-50), ResNet-101 (RN-101), ResNet-152 (RN-152), VGG-16, VGG-19, Inception-v3 (Inc-v3), DenseNet-121 (DN-121), EfficientNet-B0 (EN-B0), MobileNet-v2 (MN-v2), Swin-B, and ViT-B/16. The adversarial examples are learned against the surrogate models ResNet-101 and DenseNet-121.




Figure 3. Visual examples of original objects, our TT3D and mesh-based optimization methods under random viewpoints.

parameters. The ASR for an object is calculated as the proportion of successfully attacked images in 100 rendered images. The final ASR is determined by averaging the ASRs across all reconstructed adversarial objects.

TT3D hyperparameters. β in Sec. 3.3 is 10^3 , the epochs are 250, λ_1 is 1, λ_2 is 3000, λ_3 is 10^{-3} , λ_4 is 10^{-2} .

4.2. Attack performance in the Digital World

4.2.1 Basic Results

Effectiveness of the proposed method. To verify the effects of our proposed method, we compare the attack performance with different methods. Among them, for a fair comparison with mesh-based optimization method [49], we enhance it to optimize both the vertex colors and the coordinates. Specifically, this method is based on the classical method MeshAdv [49] and to further accentuate the limitations of this method, we put no limitations on the color modifications. Tab. 2 shows the attack success rates (%) in various common classifier models, where we find that mesh-based optimization exhibits nearly no transferability due to its tendency to overfit. In contrast, the dual optimization targeting both MLP and Grid parameters demonstrates a substantial improvement in transferability. Then, to explore the effects of different parameters in grid-based NeRF space, we also test the attack performance of MLP-only optimization and Grid-only optimization. Experimental results show that MLP-only optimization shows reasonable success, indicating its effectiveness in manipulating model-specific features. However, Grid-only optimization presents better transferability, likely due to its ability to capture spatial relationships more effectively. The combination

of both, as evidenced by the highest success rates across victim models in Tab. 2, suggests that this dual optimization method effectively integrates detailed textural features and complex decision-making processes, proving the necessity and superiority of the dual optimization strategy.

Better visual naturalness. In TT3D, while achieving excellent attack effectiveness, our method also exhibits visual naturalness compared to mesh-based optimization methods. As illustrated in Fig. 3, 3D objects optimized using mesh-based methods display non-sensical noise for appearance. In contrast, our TT3D, benefiting from targeting on both the foundation feature level and decision layers, could generate adversarial perturbations that are more semantically informative and show greater resemblance to the original ones. More specific details, including quantitative metrics, are provided in the *Supplementary Materials*.

4.2.2 Cross-render Transferability

As illustrated in Fig. 4, we observe that there exist variations in the rendering results of different renderers. To further validate the effectiveness of our method, we conduct transferability tests across various renderers, including two commercial-grade rendering software. As shown in Tab. 3, our 3D adversarial examples still perform well, even when faced with completely unknown rendering systems. This demonstrates the robustness of our method, as there is no significant performance degradation across different rendering environments, highlighting the broad applicability and resilience of our adversarial examples.

Zero-shot detector. We utilize the state-of-the-art model of Liu *et al.* [16] for zero-shot detection. In our tests, we

Source Model	Render	Victim Model										
		RN-50	RN-101	RN-152	VGG-16	VGG-19	Inc-v3	DN-121	EN-B0	MN-v2	Swin-B	ViT-B/16
ResNet-101	Nvdiffrast	78.85	88.98*	78.96	58.29	51.54	55.14	89.14	58.71	57.69	69.80	55.84
	Meshlab	53.89	60.74*	50.18	45.43	40.00	43.33	51.20	56.56	45.50	54.65	45.80
	Blender	60.03	83.88*	51.15	39.47	40.29	49.34	60.17	36.68	43.59	50.00	48.19
DenseNet-121	Nvdiffrast	70.98	51.61	59.59	57.59	50.25	43.81	96.74*	40.59	45.06	55.15	36.49
	Meshlab	46.34	40.33	39.43	43.24	38.00	34.00	74.67*	43.04	29.43	44.85	33.75
	Blender	47.66	39.06	30.94	42.34	40.00	31.25	88.75*	29.84	26.09	39.53	26.09

Table 3. The ASR(%) of 3d adversarial examples with different renders against RN-50, RN-101, RN-152, VGG-16, VGG-19, Inc-v3, DN-121, EN-B0, MN-v2, Swin-B, and ViT-B/16. The adversarial examples are learned against surrogate models RN-101 and DN-121.

Source Model	Methods	Victim Model										
		RN-50	RN-101	RN-152	VGG-16	VGG-19	Inc-v3	DN-121	EN-B0	MN-v2	Swin-B	ViT-B/16
ResNet-101	Tex	76.96	91.60*	78.52	52.00	49.51	53.31	72.96	55.73	53.09	64.40	54.50
	Tex+Geo	78.85	88.98*	78.96	58.29	51.54	55.14	89.14	58.71	57.69	69.80	55.84
DenseNet-121	Tex	65.71	46.13	52.13	46.60	40.85	36.44	94.28*	37.49	38.72	47.62	35.58
	Tex+Geo	70.98	51.61	59.59	57.59	50.25	43.81	96.74*	40.59	45.06	55.15	36.49

Table 4. The ASR(%) of tex-only and tex+geo(sub) optimization methods against RN-50, RN-101, RN-152, VGG-16, VGG-19, Inc-v3, DN-121, EN-B0, MN-v2, Swin-B, and ViT-B/16. The adversarial examples are learned against the surrogate models RN-101 and DN-121.



Figure 4. Comparison between rendering results of different renders, including differential rendering library Nvdiffrast, commercial software Meshlab, and Blender.

eliminate background distractions by using a plain white background, focusing solely on the detector’s ability to identify the target label in the rendered adversarial sample images. A detection threshold greater than 0.5 is considered successful. Experimental results reveal that even in this setting, our method achieved a success rate of 76.94%, as evidenced by some successful examples shown in Fig. 5.

4.2.3 Cross-task Transferability

To verify the transferability of our adversarial examples across different tasks, we select zero-shot detection and image caption for testing due to their broad applicability.

Image caption. For image captioning, we employ the BLIP model from Li et al. [14], still maintaining the white background setup. We present the rendered adversarial images to the BLIP model and deem the test successful if the generated caption includes the target label. Due to testing costs, we randomly selected three viewpoints for 100 different 3D objects. The results demonstrate that our method still achieves a success rate of 32.33% in this task, as evidenced by some successful examples shown in Fig. 5.

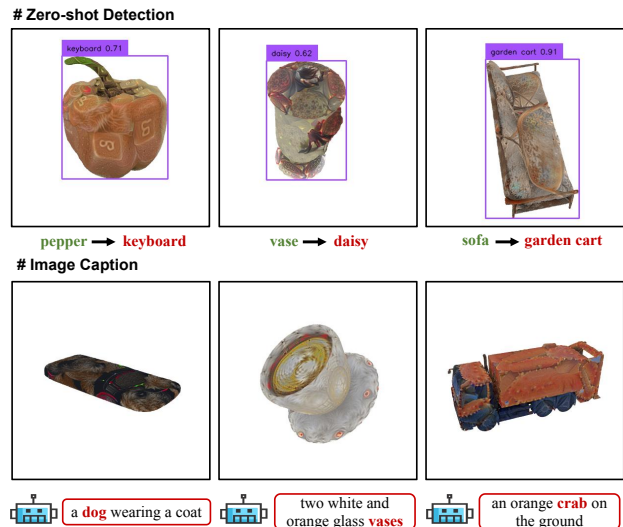


Figure 5. The prediction examples of our 3D adversarial examples targeting zero-shot detection and image caption tasks. Green and red text represent the clean label and the target label, respectively.

4.3. Additional Results and Ablation Study

The effects of additional vertex optimization. To validate the additional vertex optimization mentioned in Sec. 3.2 that could serve as an auxiliary means to further enhance texture-focused optimization performance, we conduct an ablation study on it. As shown in Tab. 4, it is evident that adding the geometric changes of the mesh vertex coordinates while optimizing texture features in the grid-based parameter space can, to a certain extent, enhance the transferability of the attack. Additionally, we find that solely altering the geometric structure with requirements for natu-

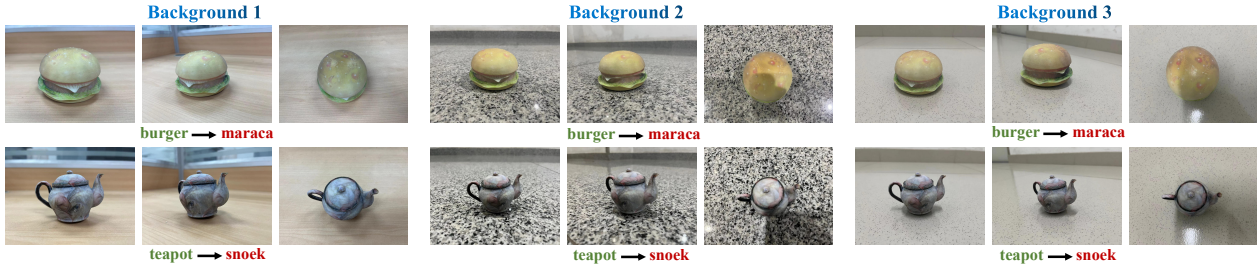


Figure 6. Visual examples of printed 3D adversarial objects towards the target label under different backgrounds (B-1, B-2, B-3) and viewpoints in the physical world. The first row is learned against ResNet-101 and the second row is learned against DenseNet-121.

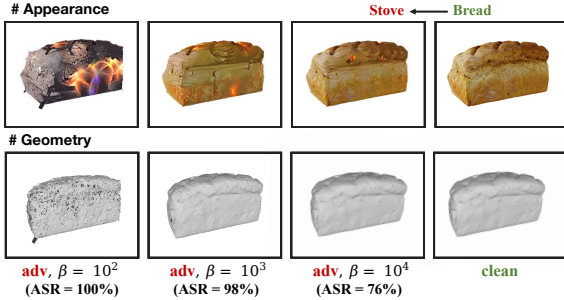


Figure 7. An visual example of the change in appearance and geometry with β and its corresponding attack success rate. naturalness, hardly achieves successful targeted attacks with arbitrary objectives. Therefore, we have not included results for solely optimization geometric structures in Tab. 4.

The effects of β in \mathcal{R} . β as a hyperparameter to adjust the regularization \mathcal{R} responsible for naturalness, has a significant impact on both naturalness and attack performance. Thus, to fully explore its impacts, we conduct an ablation study on the variation of β . As illustrated in Fig. 7, when β increases from lower to higher values (from 10^2 , 10^3 to 10^4), the distance from both the appearance and geometric structure of the adversarial sample to the original 3D object becomes closer, making the perturbation more concealed, but accompanying an observed decline in the overall success rate of the attack (detailed in the *Supplementary Materials*). To strike a balance between attack performance and naturalness, we ultimately choose a β value of 10^3 . It is also noteworthy that, regardless of the parameter setting, our adversarial examples feature textures not in the form of noise, but as semantically meaningful textures. This approach is more natural compared to explicit optimization in mesh-based optimization methods, as shown in Fig. 3.

4.4. Attack Performance in the Physical World

To verify the feasibility of TT3D in the physical world, we first utilize 3D printing techniques to print 3D adversarial objects. Then, considering that the backgrounds in the physical world are diverse, unlike the white background of the optimization process, we not only validate the digital-to-physical transferability of TT3D under different viewpoints, but also add three randomly selected, distinct backgrounds

B-1, B-2, and B-3 to evaluate its robustness against backgrounds. For the number of testing objects, we print 20 3D objects, including 10 for the surrogate ResNet-101 and 10 for the surrogate DenseNet-121. The specific process involves 1) placing the 3D adversarial objects on the surface; 2) slowly circling them with a smartphone for about 360° , excluding the bottom part. which lasts approximately 20 seconds per object in each setting, capturing 10 frames per second, resulting in a total of 200 frames; 3) calculating the attack success rate, which is determined by the proportion of successful frames. The experimental results are presented in Fig. 6 and Tab. 5, demonstrating the robust effectiveness of our method in various scenarios from the physical world. Further experiments on transferability tests in the physical world are available in *Supplementary Materials*.

	B-1	B-2	B-3
ResNet-101	81.30	76.45	91.55
DenseNet-121	79.30	83.25	84.50

Table 5. The ASR(%) of printed adversarial meshes against the RN-101 and DN-121 under different backgrounds: B-1, B-2, B-3 and various viewpoints in the real world.

5. Conclusion

In this paper, we propose a novel 3D attack framework **TT3D** that could rapidly reconstruct multi-view images into transferable targeted 3D adversarial examples, effectively filling the gap in 3D transferable targeted attacks. Besides, TT3D leverages a dual optimization strategy in the grid-based NeRF space, which significantly improves the black-box transferability meanwhile ensuring visual naturalness. Extensive experiments further demonstrate TT3D’s strong cross-model transferability and adaptability in various renders and vision tasks, with real-world applicability confirmed through 3D printing techniques.

Acknowledgement

This work was supported in part by the Project of the National Natural Science Foundation of China (Nos. 62076018, U2341228, 62276149) and in part by the Fundamental Research Funds for the Central Universities.

References

- [1] Anish Athalye, Logan Engstrom, Andrew Ilyas, and Kevin Kwok. Synthesizing robust adversarial examples. In *International conference on machine learning*, pages 284–293. PMLR, 2018. 2, 3, 4, 5
- [2] Pieter-Tjerk De Boer, Dirk P Kroese, Shie Mannor, and Reuven Y Rubinfeld. A tutorial on the cross-entropy method. *Annals of operations research*, 134:19–67, 2005. 4
- [3] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 5
- [4] 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
- [5] Yinpeng Dong, Shouwei Ruan, Hang Su, Caixin Kang, Xingxing Wei, and Jun Zhu. Viewfool: Evaluating the robustness of visual recognition to adversarial viewpoints. *Advances in Neural Information Processing Systems*, 35: 36789–36803, 2022. 1
- [6] Yinpeng Dong, Jun Zhu, Xiao-Shan Gao, et al. Isometric 3d adversarial examples in the physical world. *Advances in Neural Information Processing Systems*, 35:19716–19731, 2022. 2
- [7] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 5
- [8] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022. 4
- [9] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 1
- [10] 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 (CVPR)*, 2016. 1, 5
- [11] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 5
- [12] Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. Modular primitives for high-performance differentiable rendering. *ACM Transactions on Graphics*, 39(6), 2020. 4
- [13] Bo Li, Xiaoyang Xie, Xingxing Wei, and Tang Wenting. Ship detection and classification from optical remote sensing images: A survey. *Chinese Journal of Aeronautics*, 34(3):145–163, 2021. 1
- [14] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. 5, 7
- [15] Maosen Li, Cheng Deng, Tengjiao Li, Junchi Yan, Xinbo Gao, and Heng Huang. Towards transferable targeted attack. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 641–649, 2020. 1, 3
- [16] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. 5, 6
- [17] Yanpei Liu, Xinyun Chen, Chang Liu, and Dawn Song. Delving into transferable adversarial examples and black-box attacks. In *International Conference on Learning Representations*, 2017. 1
- [18] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 5
- [19] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pages 347–353. 1998. 4, 5
- [20] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017. 1
- [21] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021. 3
- [22] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2574–2582, 2016. 1
- [23] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. 2, 3, 4
- [24] Muzammal Naseer, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Fatih Porikli. On generating transferable targeted perturbations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7708–7717, 2021. 1, 3
- [25] Andrew Nealen, Takeo Igarashi, Olga Sorkine, and Marc Alexa. Laplacian mesh optimization. In *Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia and Southeast Asia*, pages 381–389, 2006. 5
- [26] Scott Oslund, Clayton Washington, Andrew So, Tingting Chen, and Hao Ji. Multiview robust adversarial stickers for

- arbitrary objects in the physical world. *Journal of Computational and Cognitive Engineering*, 1(4):152–158, 2022. [2](#), [3](#)
- [27] Waseem Rawat and Zenghui Wang. Deep convolutional neural networks for image classification: A comprehensive review. *Neural computation*, 29(9):2352–2449, 2017. [1](#)
- [28] Shouwei Ruan, Yinpeng Dong, Hang Su, Jianteng Peng, Ning Chen, and Xingxing Wei. Towards viewpoint-invariant visual recognition via adversarial training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4709–4719, 2023. [5](#)
- [29] Shouwei Ruan, Yinpeng Dong, Hang Su, Jianteng Peng, Ning Chen, and Xingxing Wei. Improving viewpoint robustness for visual recognition via adversarial training. *arXiv preprint arXiv:2307.11528*, 2023. [1](#)
- [30] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. [5](#)
- [31] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [5](#)
- [32] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2022. [4](#)
- [33] Naufal Suryanto, Yongsu Kim, Hyoeun Kang, Harashta Tatimma Larasati, Youngyeo Yun, Thi-Thu-Huong Le, Hunmin Yang, Se-Yoon Oh, and Howon Kim. Dta: Physical camouflage attacks using differentiable transformation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15305–15314, 2022. [2](#), [3](#)
- [34] Naufal Suryanto, Yongsu Kim, Harashta Tatimma Larasati, Hyoeun Kang, Thi-Thu-Huong Le, Yoonyoung Hong, Hunmin Yang, Se-Yoon Oh, and Howon Kim. Active: Towards highly transferable 3d physical camouflage for universal and robust vehicle evasion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4305–4314, 2023. [2](#), [3](#)
- [35] 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](#)
- [36] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016. [5](#)
- [37] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019. [5](#)
- [38] Jiaxiang Tang, Hang Zhou, Xiaokang Chen, Tianshu Hu, Er-rui Ding, Jingdong Wang, and Gang Zeng. Delicate textured mesh recovery from nerf via adaptive surface refinement. *arXiv preprint arXiv:2303.02091*, 2023. [4](#)
- [39] Simen Thys, Wiebe Van Ranst, and Toon Goedemé. Fooling automated surveillance cameras: adversarial patches to attack person detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019. [1](#)
- [40] Donghua Wang, Tingsong Jiang, Jialiang Sun, Weien Zhou, Zhiqiang Gong, Xiaoya Zhang, Wen Yao, and Xiaoqian Chen. Fca: Learning a 3d full-coverage vehicle camouflage for multi-view physical adversarial attack. In *Proceedings of the AAAI conference on artificial intelligence*, pages 2414–2422, 2022. [2](#), [3](#)
- [41] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *Proceedings of the European conference on computer vision (ECCV)*, pages 52–67, 2018. [5](#)
- [42] Zhibo Wang, Hongshan Yang, Yunhe Feng, Peng Sun, Hengchang Guo, Zhifei Zhang, and Kui Ren. Towards transferable targeted adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20534–20543, 2023. [1](#), [3](#)
- [43] Xingxing Wei, Ying Guo, and Jie Yu. Adversarial sticker: A stealthy attack method in the physical world. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3): 2711–2725, 2023. [1](#)
- [44] Xingxing Wei, Ying Guo, Jie Yu, and Bo Zhang. Simultaneously optimizing perturbations and positions for black-box adversarial patch attacks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(7):9041–9054, 2023.
- [45] Xingxing Wei, Yao Huang, Yitong Sun, and Jie Yu. Unified adversarial patch for cross-modal attacks in the physical world. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4445–4454, 2023.
- [46] Xingxing Wei, Yao Huang, Yitong Sun, and Jie Yu. Unified adversarial patch for visible-infrared cross-modal attacks in the physical world. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [47] Xingxing Wei, Jie Yu, and Yao Huang. Physically adversarial infrared patches with learnable shapes and locations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12334–12342, 2023.
- [48] Xingxing Wei, Jie Yu, and Yao Huang. Infrared adversarial patches with learnable shapes and locations in the physical world. *International Journal of Computer Vision*, pages 1–17, 2023. [1](#)
- [49] Chaowei Xiao, Dawei Yang, Bo Li, Jia Deng, and Mingyan Liu. Meshadv: Adversarial meshes for visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6898–6907, 2019. [1](#), [2](#), [3](#), [6](#)
- [50] 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 *Pro-*

ceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019. [1](#)

- [51] Xiao Yang, Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Boosting transferability of targeted adversarial examples via hierarchical generative networks. In *European Conference on Computer Vision*, pages 725–742. Springer, 2022. [3](#)
- [52] Xiaohui Zeng, Chenxi Liu, Yu-Siang Wang, Weichao Qiu, Lingxi Xie, Yu-Wing Tai, Chi-Keung Tang, and Alan L Yuille. Adversarial attacks beyond the image space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4302–4311, 2019. [2](#), [3](#)
- [53] Zhengyu Zhao, Zhuoran Liu, and Martha Larson. On success and simplicity: A second look at transferable targeted attacks. *Advances in Neural Information Processing Systems*, 34:6115–6128, 2021. [1](#), [3](#)
- [54] Zhong-Qiu Zhao, Peng Zheng, Shou-tao Xu, and Xindong Wu. Object detection with deep learning: A review. *IEEE transactions on neural networks and learning systems*, 30(11):3212–3232, 2019. [1](#)