

# Hide in Thicket: Generating Imperceptible and Rational Adversarial Perturbations on 3D Point Clouds

Tianrui Lou<sup>1</sup> Xiaojun Jia<sup>2,\*</sup> Jindong Gu<sup>3</sup> Li Liu<sup>4</sup> Siyuan Liang<sup>5</sup> Bangyan He<sup>6</sup> Xiaochun Cao<sup>1,\*</sup>

<sup>1</sup>Sun Yat-Sen University <sup>2</sup>Nanyang Technological University <sup>3</sup>University of Oxford

<sup>4</sup>National University of Defense Technology <sup>5</sup>National University of Singapore <sup>6</sup>Chinese Academy of Sciences

{loutianrui, jiaxiaojunqqaq, pandaliang521}@gmail.com jindong.gu@siemens.com

li.liu@oulu.fi hebangyan@iie.ac.cn caoxiaochun@mail.sysu.edu.cn

## Abstract

Adversarial attack methods based on point manipulation for 3D point cloud classification have revealed the fragility of 3D models, yet the adversarial examples they produce are easily perceived or defended against. The trade-off between the imperceptibility and adversarial strength leads most point attack methods to inevitably introduce easily detectable outlier points upon a successful attack. Another promising strategy, shape-based attack, can effectively eliminate outliers, but existing methods often suffer significant reductions in imperceptibility due to irrational deformations. We find that concealing deformation perturbations in areas insensitive to human eyes can achieve a better trade-off between imperceptibility and adversarial strength, specifically in parts of the object surface that are complex and exhibit drastic curvature changes. Therefore, we propose a novel shape-based adversarial attack method, HiT-ADV, which initially conducts a two-stage search for attack regions based on saliency and imperceptibility scores, and then adds deformation perturbations in each attack region using Gaussian kernel functions. Additionally, HiT-ADV is extendable to physical attack. We propose that by employing benign resampling and benign rigid transformations, we can further enhance physical adversarial strength with little sacrifice to imperceptibility. Extensive experiments have validated the superiority of our method in terms of adversarial and imperceptible properties in both digital and physical spaces. Our code is available at: <https://github.com/TRLou/HiT-ADV>.

## 1. Introduction

In recent years, numerous models based on Deep Neural Networks (DNNs) [7, 11, 31, 37, 48] have shown strong performance in a broad spectrum of tasks, including 3D data processing and analysis [13, 28, 30, 33, 53]. How-

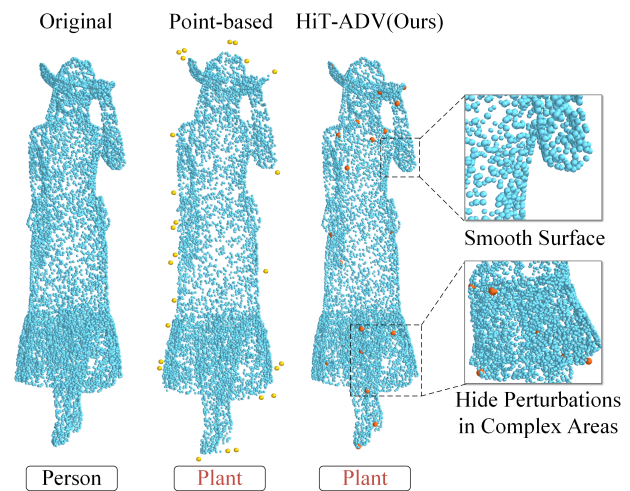


Figure 1. Comparison of successful adversarial examples generated by point-based attack method [24] and HiT-ADV. The orange points represent the central points of local shape deformations generated in HiT-ADV and the yellow points represent the outliers in point-based attack. HiT-ADV has no outlier points and exhibits a smooth surface. By concealing perturbations in complex areas, the deformation perturbations become difficult to perceive.

ever, it has been observed that meticulously crafted, imperceptible perturbations can lead DNNs to produce erroneous outputs [6, 8, 14, 15, 19, 20, 27, 42], highlighting their inherent vulnerability. This characteristic of DNNs significantly hinders their development in some safety-critical areas, e.g. autonomous driving [1, 4, 16, 36, 39] and robotic task planning [21, 38]. Numerous point cloud adversarial attacks have effectively compromised 3D point cloud models [10, 35, 43, 52], exposing their vulnerabilities. However, adversarial examples in point clouds are usually easily perceived by the human vision compared to those in other domains such as imaging, because point cloud data lacks RGB information to conceal perturbations. Thus, achieving

imperceptibility, a vital attribute of adversarial examples, remains a critical question for further investigation.

Existing adversarial attack methods for point clouds are primarily point-based, encompassing point movement [12, 23, 24, 35, 43], point addition [47], and point deletion [50]. However, as shown in Fig.1, point-based attack methods often result in noticeable outlier points and coarse surfaces. This directly makes the adversarial examples easily perceptible to the human eyes and prone to be defended by pre-processing defense methods [22, 46, 51]. Furthermore, when digital adversarial point clouds are converted into physical adversarial objects, retaining outliers causes noticeable shape distortion, while discarding them leads to a significant decrease in adversarial strength. While many prior studies have considered geometric features, using regularization terms to constrain optimization objectives in order to reduce outlier points [35, 43, 47], these methods face challenges in fundamentally eliminating them due to the trade-off between imperceptibility and adversarial strength. Another attack strategy is based on deformation [5, 25, 32, 49], which can effectively prevent outlier generation by reinforcing inter-point correlations. However, current deformation attack methods typically induce excessive and irrational perturbations, extremely compromising the imperceptibility of successful adversarial examples.

In fact, we find that the reason previous works have difficulty optimizing a satisfactory trade-off relationship is due to a misconception about imperceptibility. We argue that imperceptibility can be maintained with reasonable shape changes, not just by keeping the original shape exactly. Specifically, we find that the human eyes have varying sensitivity to different regions of an object. Concretely, they can confidently discern areas that should be smooth, but it is challenging to distinguish whether inherently complex surfaces have been tampered with. Therefore, concealing adversarial perturbations in complex surfaces can maintain the shape reasonable and the perturbations imperceptible.

Based on this observation, we propose a novel imperceptible shape-based attack, called HiT-ADV. Firstly, we introduce a two-stage attack region search module to search multiple regions that are vulnerable to attack and insensitive to the human vision. Secondly, we conduct an iterative attack using Gaussian kernel functions to calculate deformation perturbations. Correspondingly, we propose novel imperceptibility regularization terms to constrain deformation perturbations, continuously refining the shape of adversarial point clouds to ensure their imperceptibility. Finally, we analyze the main challenge in extending to physical attack, which is the vanishing of digital adversarial strength during the physical attack process. The vanished digital adversarial strength primarily comes from outlier points, non-shape-altering point movements, and adversarial rigid transformations. To tackle this issue, we perform benign resampling

and benign rigid transformations before every iteration to ensure our adversarial properties are mainly from physical deformation. Through this way, HiT-ADV can be effectively extended to physical attack. Comprehensive experiments on two benchmarks with classifiers in two representative DNN architectures validate that HiT-ADV is imperceptible and effective, and outperforms the state-of-the-art methods. Additionally, we validate the effectiveness of HiT-ADV in physical scenarios and visualize the attack results.

Our main contributions are in three aspects:

- We propose an imperceptible shape-based attack framework. By concealing shape perturbations in complex surface areas, we achieve a favorable trade-off between adversarial strength and imperceptibility.
- We propose an optimization method that effectively suppresses the digital adversarial strength of adversarial examples, enhancing their performance in the physical world.
- We validate the superiority of our framework to the state-of-the-art point cloud attack methods via extensive experiments. Besides, we print the adversarial examples generated by HiT-ADV and use the re-scanned point clouds to successfully attack the classification model, to verify the feasibility of our method when extended to the physical world.

## 2. Related Work

**DNN Models for 3D Classification.** Point cloud is a promising method of 3D data representation, effectively capturing the geometric information of real-world objects, and can be extensively acquired through LiDAR or RGB-D cameras. However, their irregular and unordered nature poses challenges for direct utilization. Early research efforts attempted to transform them into structured formats, such as voxels, multi-view images, or meshes, for subsequent tasks. Until recently, PointNet [28] innovatively tackled the irregularity and unordered nature of point cloud data by employing multi-layer perceptrons (MLP) and max-pooling layer, facilitating end-to-end classification and segmentation directly on point clouds. Due to the convenience and high performance of PointNet, numerous research efforts were inspired to propose various DNN architectures for direct point cloud learning. Examples include PointNet++ [29], which utilizes MLPs to learn hierarchical point cloud features, DGCNN [40] with graph neural networks, PCT [9] with the transformer architecture, and so forth [17, 26, 34, 44].

**3D Adversarial Attack.** The vulnerability of 3D DNN models mirrors that of the 2D image models [6], where a slight perturbation, imperceptible to the human eye, can easily lead to erroneous model classification. Liu et al. [24] first extended the gradient-based adversarial attack, FGM [6], to the point cloud domain under an  $l_2$

norm constraint. Xiang et al. [47] introduced the C&W attack framework [2] in point cloud attacks, producing superior adversarial examples by shifting point coordinates. Moreover, they proposed achieving adversarial properties by adding points, clusters, and objects. Besides, Zheng et al. [50] proposed that deleting a small number of points with high saliency can effectively cause misclassification in point cloud classification models.

Since point clouds only contain coordinate information and lack the RGB three-channel data of images to conceal adversarial perturbations, coupled with the human eye’s high sensitivity to outliers in adversarial point clouds, many researchers have been dedicated to making adversarial perturbations more imperceptible. Tsai et al. [35] improved the C&W attack framework by introducing a KNN regularization term, resulting in a tighter point cloud surface and suppressing the generation of outlier points. Wen et al. [43] introduced geometry-aware regularization constraints, which encompass maintaining local curvature consistency and ensuring a uniform surface on the adversarial point cloud. Huang et al. [12] first transformed the coordinate system onto the tangent plane and then introduced perturbations on this plane to generate shape-invariant point clouds.

These point-based methods above enhanced imperceptibility through geometric perspectives while struggling to eliminate all outliers and obtain smooth surfaces. Thus, many studies have attempted to propose shape-based attack methods. Zhang et al. [49] and Miao et al. [5] suggested directly attacking mesh data to generate sufficiently smooth results, utilizing edge length regularization and Gaussian curvature regularization respectively. Liu et al. [25] attempted to conduct shape-based attacks by adding new features to objects, like adversarial sticks. Additionally, they proposed using the Gaussian kernel for sinking operations to achieve point cloud shape deformation. Tang et al. [32] proposed to adversarially stretch the latent variables in an auto-encoder, which can be decoded as smooth adversarial point clouds. While these shape-based methods eliminate outlier points and ensure smoothness, they either face challenges in data acquisition or are easily perceptible. In contrast, our proposed shape-based attack HiT-ADV is both smooth and imperceptible, and these characteristics enable it to be easily extended to the physical scenarios.

### 3. Preliminaries

We focus on adversarial attacks in the context of point cloud classification tasks. We denote  $\mathcal{D} = \{\mathcal{P}_i, y_i\}_{i=1}^n$  as the training dataset of a group of point clouds. We use  $\mathcal{P}_i \in \mathbb{R}^{m \times 3}$  to represent a single point cloud to be classified in  $\mathcal{D}$ , which contains a set of unordered point sets, denoted as  $\{\mathbf{p}_j\}$ , and  $y_i$  is the corresponding label. Each point  $\mathbf{p}_j \in \mathbb{R}^3$  in this set contains three-dimensional coordinate information. The goal of the adversarial attack is to

introduce a small perturbation  $\delta_i$  to the point cloud  $\mathcal{P}_i$ , such that the classifier  $\mathcal{F}$  makes an incorrect classification of  $\mathcal{P}_i$ , i.e.  $\mathcal{F}(\mathcal{P}_i + \delta_i) \neq y_i$ . Note that the situation discussed in this paper is assumed to be untargeted adversarial attack by default. For the visualization results of targeted attack, please refer to the appendix.

In order to obtain more imperceptible adversarial point clouds  $\mathcal{P}' = \mathcal{P} + \delta$ , our attack process adopts a gradient-based optimization algorithm, following [47], and can be formulated as:

$$\min \mathcal{L}_{cls}(\mathcal{P}') + \lambda * \mathcal{L}_{dis}(\mathcal{P}, \mathcal{P}'), \quad (1)$$

where  $\mathcal{L}_{cls}$  is an adversarial loss function measuring the successful attack probability and  $\mathcal{L}_{dis}$  is a regularization term constraining the distance between benign point cloud  $\mathcal{P}$  and the adversarial one  $\mathcal{P}'$ .  $\lambda$  is a hyper-parameter, automatically optimized to an approximate optimal value through binary search, which can balance the trade-off relationship between the two terms above. Concretely,  $\mathcal{L}_{cls}$  can be formulated as:

$$\mathcal{L}_{cls} = \max \{-\kappa, \mathcal{Z}_t(\mathcal{P}') - \max_{j \neq t} \mathcal{Z}_j(\mathcal{P}')\}, \quad (2)$$

where  $\kappa \geq 0$  represents a margin threshold, and  $\mathcal{Z}$  represents the output logits of  $\mathcal{F}(\mathcal{P}')$ , and  $t$  is the true label class.

**Problem Analysis.** The point-based adversarial perturbation  $\delta \in \mathbb{R}^3$  is calculated from the back-propagated gradient of the loss in Eq.1. Even though some methods, such as [5, 35, 43, 47], attempt to eliminate outliers and ensure overall imperceptibility by modifying the regularization terms  $\mathcal{L}_{dis}$ , due to the lack of hard correlation between the  $\delta_j$  of different points, there will always be a small number of points escaping the regularization term constraint when the optimization objective  $\mathcal{L}_{cls}$  becomes difficult. Please refer to the visualization result of Fig.3. Furthermore, excessively demanding the shape of adversarial point clouds to mirror the original, while also lacking outlier points, severely restricts the search space, leading to poor adversarial strength. As compared to point-based attack methods, shape-based attacks, such as [25], often achieve smooth surfaces and eliminate outliers because they rigidly constrain the relationship between the perturbations  $\delta_j$  of different points. However, shape-based attacks often require a larger search space. Therefore, ensuring the imperceptibility of shape-based attacks remains a critical challenge that requires resolution. We have an intuition that the human eyes have different sensitivities to changes in different parts of a 3D object. Deformations occurring in areas with complex surfaces are less likely to be perceived by human eyes because people do not have an absolute prior cognition of the specific shape of these areas. Based on this

intuition, we propose an imperceptible shape-based attack method HiT-ADV, which will be introduced in detail in the next section.

## 4. Methodology

In this section, we introduce the specific framework and implementation of HiT-ADV, as well as methods to further suppress digital adversarial strength. Please refer to Fig.2 for the demonstration of our framework.

### 4.1. Saliency and Imperceptibility Score

The design objective of this module is to locate areas in the object that are both likely to affect the classification result and imperceptible to human vision. Note that when a region undergoes deformation, the center of this region often suffers the most drastic perturbations. Therefore, we can effectively search suitable central points for deformation perturbations in the clean point cloud  $\mathcal{P}_i$  using a saliency and imperceptibility score (SI score)  $\mathcal{S}$ . The definition of saliency score  $\mathcal{S}_1$  follows [50], which measures the contribution of each point to the correct classification. As discussed in Section 3, points located in complex surface areas should be evaluated with higher imperceptibility scores  $\mathcal{S}_2$ . Therefore, we propose to formulate  $\mathcal{S}_2$  as follow:

$$\begin{aligned} \mathcal{S}_2(\mathbf{p}_j; \mathcal{P}) &= \sqrt{\frac{1}{k} \sum_{\mathbf{q} \in \mathcal{N}_{\mathbf{p}_j}} (\mathcal{C}(\mathbf{q}; \mathcal{P}) - \mu)^2}, \\ \mu &= \frac{1}{k} \sum_{\mathbf{q} \in \mathcal{N}_{\mathbf{p}_j}} \mathcal{C}(\mathbf{q}; \mathcal{P}), \end{aligned} \quad (3)$$

where  $\mathcal{N}_{\mathbf{p}_j}$  represents the  $k$  nearest neighbor of point  $\mathbf{p}_j$ . And  $\mathcal{C}$  is the local curvature of  $\mathbf{p}_j$  whose formulation follows [43]:

$$\mathcal{C}(\mathbf{p}_j; \mathcal{P}) = \frac{1}{k} \sum_{\mathbf{q} \in \mathcal{N}_{\mathbf{p}_j}} | \langle (\mathbf{q} - \mathbf{p}_j) / \|\mathbf{q} - \mathbf{p}_j\|_2, \mathbf{n}_{\mathbf{p}_j} \rangle |, \quad (4)$$

where  $\mathbf{n}_{\mathbf{p}_j}$  denotes the normal vector of the local surface around  $\mathbf{p}_j$ , and  $\langle \cdot, \cdot \rangle$  denotes inner product.  $\mathcal{S}_2$  evaluates the standard deviation of local surface curvature, describing the intensity of local curvature variation. We can normalize  $\mathcal{S}_1$  and  $\mathcal{S}_2$  to the range  $[0, 1]$  and add them together to obtain the SI score  $\mathcal{S} = \mathcal{S}_1 + \alpha \cdot \mathcal{S}_2$ , where  $\alpha$  is set to 1 by default.

### 4.2. Two-stage Attack Region Searching Module

To perform adversarial deformation on different local regions, we first partition the point cloud  $\mathcal{P}_j$ . Concretely, we obtain  $n$  points by farthest point sampling (FPS) and gather  $k$  nearest neighbor points around them. From these regions, we select the points with the highest SI scores as their central points. To further enhance imperceptibility, we design the second-stage selection to choose the final deformation

regions, by selecting  $\tilde{n}$  central points, denoted as  $\{\tilde{\mathbf{p}}_i\}_{i=1}^{\tilde{n}}$ , with the highest SI scores from the  $n$  candidate points in the previous stage. Through this mechanism, we ensure that the deformation perturbations can be concealed in complex surface areas.

### 4.3. Shape-based Attack Method based on Gaussian Kernel Function

For the obtained central points, we add random perturbations  $\delta$  to them as initialization. Then we need to diffuse the perturbations, calculating the size and direction of perturbations for each point around the central points. To make the adversarial point cloud surface smooth, we use the classical Gaussian kernel function to calculate the adversarial perturbations. Given a central point  $\tilde{\mathbf{p}}$  and an original point  $\mathbf{p}$ , we can formulate the Gaussian kernel function as:

$$Gauss(\tilde{\mathbf{p}}, \mathbf{p}) = e^{-\frac{\|\tilde{\mathbf{p}} - \mathbf{p}\|_2^2}{2\sigma^2}}, \quad (5)$$

where  $\sigma$  is the hyper-parameter of bandwidth controlling the shape of the Gaussian kernel function. Note that the magnitude and direction of the final perturbation of an individual point should be collectively calculated by Gaussian kernel functions of multiple different deformation central points. Thus we use Nadaraya-Watson kernel regression [41] to calculate the adversarial point  $\mathbf{p}'_j$ :

$$\mathbf{p}'_j = \frac{\sum_{i=1}^{\tilde{n}} (\mathbf{p}_j + \delta_i) \cdot Gauss(\tilde{\mathbf{p}}_i, \mathbf{p}_j)}{\sum_{i=1}^{\tilde{n}} Gauss(\tilde{\mathbf{p}}_i, \mathbf{p}_j)} \quad (6)$$

Both  $\delta$  and  $\sigma$  can affect the generation of the adversarial point cloud, so their values will be continuously updated during the iterative optimization of the attack process.

To constrain the imperceptibility of the adversarial examples, we use three regularization terms. The first one is denoted as kernel loss  $\mathcal{L}_{ker}$ , formulated as:

$$\mathcal{L}_{ker} = \|\delta\|_2 + \|a - \sigma\|_2, \quad (7)$$

where  $a$  is a hyper-parameter and  $\sigma$  is clipped to ensure it is smaller than  $a$ .  $\mathcal{L}_{ker}$  requires the perturbations on central points to be as small as possible, and the Gaussian kernel function to be as flat as possible to prevent drastic local spikes. Further, we fine-tune the deformation perturbations based on the surface complexity, allowing more severe perturbations to be concealed in more complex areas. The second regularization term describes the relationship between the  $\sigma$  of the Gaussian kernel and the standard deviation of curvature at the central points  $\tilde{\mathbf{p}}$  by cosine similarity:

$$\mathcal{L}_{hide} = \frac{\sigma \cdot \mathcal{C}_{std}(\tilde{\mathbf{p}}_i; \mathcal{P})}{\|\sigma\|_2 \times \|\mathcal{C}_{std}(\tilde{\mathbf{p}}_i; \mathcal{P})\|_2}, \quad (8)$$

where  $\mathcal{C}_{std}(\tilde{\mathbf{p}}_i; \mathcal{P})$  is equally defined as  $\mathcal{S}_2$ , and both  $\mathcal{C}_{std}$  and  $\sigma$  are normalized to the range  $[0, 1]$ . This regularization



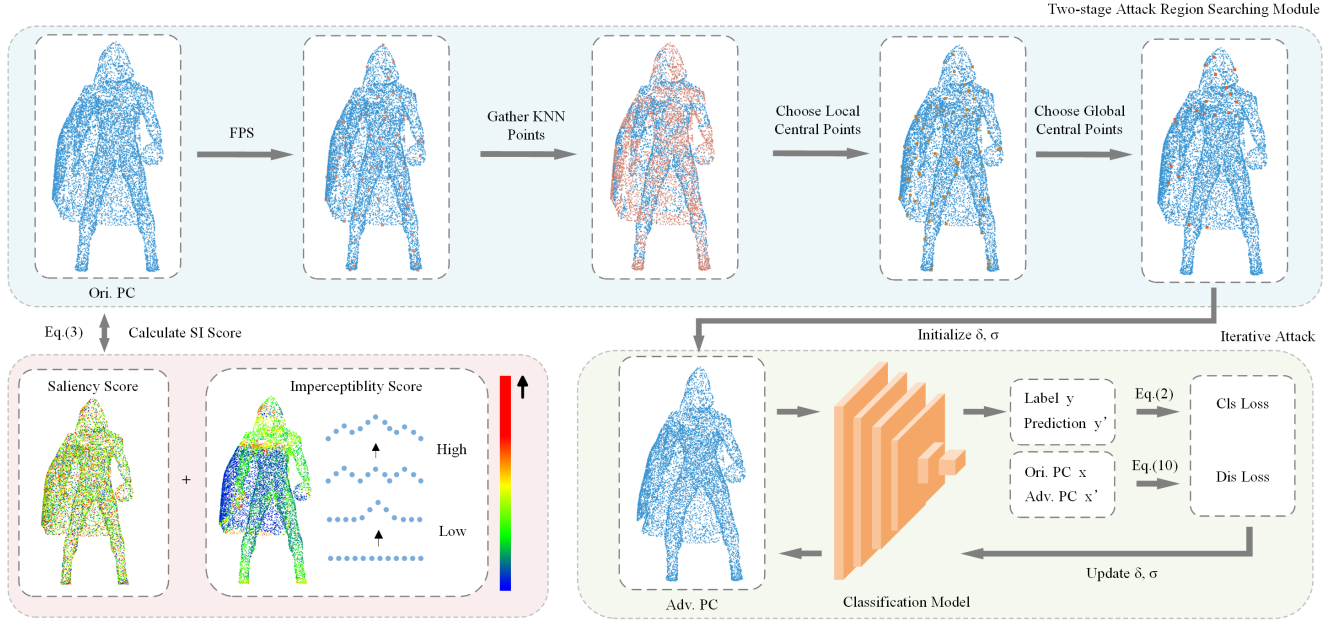


Figure 2. Demonstration of the framework of HiT-ADV. For clean point cloud samples, HiT-ADV first calculates the SI score for each point, and we color them in this figure according to their ranking. Subsequently, we employ a two-stage attack region search to locate the global central points of imperceptible regions. Finally, we iteratively attack using multiple Gaussian kernel functions and propose a new distance loss for constraint.

term loss results in smaller  $\sigma$  values for the Gaussian kernels at deformation center points located on complex surfaces and larger  $\sigma$  values for smoother surfaces. The final regularization term is the typical Chamfer loss  $\mathcal{L}_{cha}$ :

$$\mathcal{L}_{cha} = \frac{1}{m} \sum_{\mathbf{p} \in \mathcal{P}} \min_{\mathbf{p}' \in \mathcal{P}'} \|\mathbf{p} - \mathbf{p}'\|_2^2 + \frac{1}{m'} \sum_{\mathbf{p}' \in \mathcal{P}'} \min_{\mathbf{p} \in \mathcal{P}} \|\mathbf{p}' - \mathbf{p}\|_2^2 \quad (9)$$

$\mathcal{L}_{dis}$  in Eq. 1 can be expressed with 3 hyper-parameters as:

$$\mathcal{L}_{dis} = \lambda_1 \mathcal{L}_{ker} + \lambda_2 \mathcal{L}_{hide} + \lambda_3 \mathcal{L}_{cha}, \quad (10)$$

#### 4.4. Suppressing Digital Adversarial Strength

To further enhance the physical adversarial strength of the generated adversarial point clouds by HiT-ADV, we should avoid adversarial strength obtained by rigid transformations and non-shape-altering point movements. Thus before each iteration of attack, we need to minimize the adversarial strength above as much as possible. Concretely, we use the MaxOT [5] optimization paradigm to suppress adversarial strength from rigid transformations:

$$\tilde{\mathcal{T}} = \operatorname{argmax}_{\mathcal{T}} \mathcal{L}_{cls}(\mathcal{T}(\mathcal{P}'), y) \quad (11)$$

The rigid transformation can be formulated as:

$$\mathcal{T}(\mathcal{P}') = \mathbf{S}(\mathbf{R}(\mathcal{P}')) + \mathbf{T}, \quad (12)$$

where  $\mathbf{S}, \mathbf{R}, \mathbf{T}$  are scaling, rotation and translation matrices respectively. Note that the benign resampling process is non-differentiable and cannot be directly incorporated into the MaxOT optimization paradigm. Therefore, we choose to upsample before rigid transformation and then use farthest point sampling to downsample and match the number of points with the original  $\mathcal{P}$  after the iteration. Through these two steps, we can effectively suppress digital adversarial strength during the iteration process and the overall optimization objective is reformulated as:

$$\min \mathcal{L}_{cls}(\tilde{\mathcal{T}}(\mathcal{P}')) + \lambda * \mathcal{L}_{dis}(\tilde{\mathcal{T}}(\mathcal{P}), \tilde{\mathcal{T}}(\mathcal{P}')) \quad (13)$$

## 5. Experiments

### 5.1. Implementation

We implement our framework and reproduce all the DNN models with PyTorch, and report the results on a workstation with an Intel Xeon Gold 6226R CPU@2.90Hz and 64 GB of memory using a single RTX 3090 GPU.

### 5.2. Experimental Setup

**Datasets.** Our experiment utilize two public point cloud datasets, ModelNet40 [45] and ShapeNet Part [3], to demonstrate the superiority of HiT-ADV. The ModelNet40 dataset comprises 12,311 CAD models across 40 primary object classes, split into 9,843 objects for training and 2,468

Model	Method	ModelNet40				ShapeNet Part			
		ASR(%) $\uparrow$	CSD $\downarrow$	Uniform $\downarrow$	KNN(*1e-3) $\downarrow$	ASR(%) $\uparrow$	CSD $\downarrow$	Uniform $\downarrow$	KNN(*1e-3) $\downarrow$
PointNet	IFGSM( $l_{inf}$ )	99.01	2.6341	0.3134	0.8558	95.20	2.8831	0.3035	0.8076
	IFGM( $l_2$ ) [24]	99.68	1.6502	0.3173	0.7886	97.49	2.3401	0.2404	0.5898
	3D-ADV [47]	<b>100.00</b>	0.9500	0.2965	0.7355	<b>100.00</b>	2.1603	0.2103	0.5322
	KNN [35]	99.68	2.3133	0.3695	<b>0.6077</b>	96.40	2.5067	0.2723	<b>0.5098</b>
	GeoA <sup>3</sup> [43]	<b>100.00</b>	1.7112	0.2919	0.7174	<b>100.00</b>	3.7291	0.2658	0.6892
	SI-ADV [12]	99.32	1.4601	0.3059	1.5979	96.38	2.8476	0.2945	0.9217
	AOF [23]	<b>100.00</b>	2.8213	0.3809	1.2519	99.85	3.3172	0.3601	1.3110
	MeshAttack [49]	96.64	2.2161	0.3008	0.7971	97.24	3.0671	0.3444	0.7798
	Eps-iso [5]	97.17	2.2514	0.3053	0.7949	98.01	2.6093	0.2908	0.7457
HiT-ADV (Ours)	<b>100.00</b>	<b>0.4709</b>	<b>0.2883</b>	0.7447	<b>100.00</b>	<b>0.9810</b>	<b>0.1874</b>	0.5272	
DGCNN	IFGSM( $l_{inf}$ )	98.96	2.7243	0.3112	0.8891	98.46	3.3550	0.2094	0.7085
	IFGM( $l_2$ ) [24]	98.96	2.0193	0.2983	0.6846	99.51	2.6947	0.2080	0.6627
	3D-ADV [47]	<b>100.00</b>	1.0206	0.2919	0.8677	<b>100.00</b>	2.3788	0.1886	0.6837
	KNN [35]	96.15	3.2322	0.3987	<b>0.6080</b>	98.09	3.4718	0.2372	0.6349
	GeoA <sup>3</sup> [43]	99.71	1.0286	0.2887	0.8651	<b>100.00</b>	2.9907	0.1895	0.7180
	SI-ADV [12]	96.08	2.3211	0.5557	1.5992	96.43	3.0193	0.2456	0.6824
	AOF [23]	97.09	2.3148	0.3226	1.0211	98.54	3.2227	0.2444	0.8970
	MeshAttack [49]	<b>100.00</b>	1.0411	0.3192	0.8366	99.43	3.1836	0.2232	0.7092
	Eps-iso [5]	<b>100.00</b>	1.0430	0.3001	0.8377	<b>100.00</b>	3.0629	0.2171	0.7108
	HiT-ADV (Ours)	<b>100.00</b>	<b>0.1946</b>	<b>0.2867</b>	0.8374	<b>100.00</b>	<b>0.6629</b>	<b>0.1870</b>	<b>0.6325</b>

Table 1. Comparison of different point cloud attack methods on the ModelNet40 and ShapeNet datasets targeting an undefended PointNet and DGCNN classification model, including adversarial attack success rate (ASR), curvature standard deviation difference (CSD), uniform metric, and KNN distance.

Method	PointNet					DGCNN				
	No Defense	SRS	SOR	DUP-Net	AT	No Defense	SRS	SOR	DUP-Net	AT
IFGSM	99.01	77.37	60.68	43.57	49.09	98.96	67.31	82.14	61.10	72.06
IFGM [24]	99.68	81.81	33.49	21.49	21.65	98.96	58.62	72.54	59.77	42.58
3D-ADV [47]	<b>100.00</b>	24.28	24.25	14.03	60.98	<b>100.00</b>	35.55	43.08	52.33	74.82
KNN [35]	99.68	84.72	28.21	21.27	11.39	96.15	74.26	79.88	66.25	52.41
GeoA3 [43]	<b>100.00</b>	40.66	18.76	13.65	64.35	99.71	<b>76.28</b>	44.61	53.00	77.30
SI-ADV [12]	99.32	58.32	50.20	56.73	36.21	96.08	72.17	71.78	73.02	51.99
AOF [23]	<b>100.00</b>	73.82	86.20	63.27	27.20	97.09	50.81	63.58	60.24	33.62
MeshAttack [49]	96.64	25.67	18.01	18.46	24.21	<b>100.00</b>	32.12	45.66	53.28	74.29
Eps-iso [5]	97.17	26.16	19.79	19.61	30.23	<b>100.00</b>	31.36	47.28	54.33	78.96
HiT-ADV (Ours)	<b>100.00</b>	<b>84.95</b>	<b>87.86</b>	<b>72.27</b>	<b>99.76</b>	<b>100.00</b>	66.82	<b>83.98</b>	<b>92.37</b>	<b>100.00</b>

Table 2. Comparison results of ASR (%) for different attack methods on models with and without applied defense methods.

objects for testing. The ShapeNet Part dataset includes 16,881 pre-aligned shapes from 16 categories, with 12,137 objects designated for training and 2,874 objects for testing.

**Point Cloud Classifiers.** We select the two most representative point cloud classification models for experimentation: the MLP-based PointNet [28] and the graph-based DGCNN [40]. Before starting the following experiments, we train both classification models using clean examples according to their original papers.

**Attack Methods for Comparison.** We select nine state-of-the-art point cloud adversarial attack methods as our baseline for comparison, including IFGSM( $l_{inf}$ ), IFGM( $l_2$ ) [24], 3D-ADV [47], KNN [35], GeoA<sup>3</sup> [43], SI-ADV [12], AOF [23], MeshAttack [49] and Eps-iso [5].

**Evaluation Metric.** We assess the efficiency of adversarial attacks through the attack success rate (ASR), which represents the percentage of generated adversarial examples that can deceive the classifiers. Additionally, we evaluate the imperceptibility of these attack methods by three distance metrics, including one fully related metric curvature standard deviation distance (CSD), and two no related metrics, i.e. uniform metric [18] and KNN distance [35]. The CSD metric intuitively compares the consistency of  $C_{std}$  between clean point clouds and adversarial point clouds. Besides providing a macro-level comparison of the similarity, it also evaluates whether the attack methods can generate perturbations reasonably based on the surface complexity. CSD can be formulated as:

$$CSD = \|C_{std}(\mathcal{P}) - C_{std}(\mathcal{P}')\|_2 \quad (14)$$




Figure 3. Visualization of original and adversarial point clouds generated by different adversarial attack methods for attacking PointNet.

### 5.3. Performance Analysis

**Quantitative Results on Imperceptibility.** In order to fairly compare the imperceptibility of various point cloud adversarial attack methods, we adjust the settings of all methods so that they could achieve attack success rates of over 96% on classification models without defense. Under these settings, we compare their three imperceptibility metrics, and the results are shown in Tab.1. The results show that HiT-ADV can achieve superior performance in CSD and Uniform, while maintaining 100% attack success rates across all models and datasets. Although HiT-ADV gets a bit lower KNN metrics compared to methods that use only KNN distance regularization, its performance is still comparable to other point cloud methods. It demonstrates that HiT-ADV can generate high-quality adversarial point clouds while achieving better imperceptibility by hiding the perturbations in areas that are insensitive to human vision.

**Attack Performance against Defense.** To evaluate the adversarial strength of our proposed deformation attack method, we conducted tests on classification models with four different defense methods, i.e. SRS, SOR, DUP-Net [51] and adversarial training (AT) [24]. For experiments using the three pre-processing defense methods SRS, SOR, and DUP-Net, we generate adversarial examples on a model without defense and then calculate the ASR on a model with defense. For AT, we adapt the PGD-AT model with the  $l_2$  norm constraint and the budget is set as 1. Please refer to Tab.2 for comparative results. All the detailed settings for the attack methods in Tab.2 are the same as in Tab.1.

It’s worth noting that the adversarial strength of HiT-ADV comes from deformations in the overall shape, which is why we perform well on three preprocessing defense methods. We also demonstrate that point-based AT models, while generally effective against point-based attacks, are not effective at defending against shape-based attacks. This experiment shows that the attack method we proposed further reveals the vulnerability of DNN models. The question of how to effectively enhance robustness against shape-based attacks requires further research.

**Visualization.** The visual comparison of the adversarial examples generated by each method is shown in Fig. 3. It is quite evident that, aside from HiT-ADV, other methods more or less have obvious outliers. It is an unavoidable common problem of point-based attack methods because it is difficult to generate adversarial examples that are sufficiently adversarial and free of outliers with soft regularization terms in the limited search space constrained by the original shape. Furthermore, our shape-based perturbation is also sufficiently natural to fully ensure imperceptibility. Note that the Eps-iso method, which constrains outliers with relatively strong edge length loss and Laplacian loss, claims to be a shape-based attack. However, this method firstly does not relax the search space, which still allows outliers to escape constraints, and secondly, the excessive aggregation of point clouds results in the loss of object surface details. These problems lead to unsatisfactory visualization results when point clouds are used as input for this method.




Figure 4. The visualization of the physical adversarial attack process using HiT-ADV, which incorporates benign resampling and benign rigid transformations, is shown in a sequence from left to right: the mesh reconstructed from the adversarial point cloud, the physical adversarial sample produced by 3D printing, the point cloud obtained from rescanning and classification results comparison.

**Results of Physical Attack.** In this experiment, we employ the optimization methods discussed in Sec. 4.4 to further enhance the physical adversarial strength. To validate the effectiveness of adversarial examples generated by HiT-ADV in physical scenarios, we conduct a physical attack experiment and visualize the intermediate results of each stage in Fig. 4. The overall physical attack process includes reconstructing into a mesh, 3D printing as an adversarial object, and re-scanning to form the input point cloud. As previously analyzed, in the physical attack process, the digital adversarial strength stemming from outlier points, non-shape-altering point movements, and adversarial rigid transformations can not survive. However, HiT-ADV still achieves successful physical attacks, owing to its adversarial strength derived purely from the deformation of object.

**Ablation Study.** We conduct an ablation experiment focusing on the three regularization terms in  $\mathcal{L}_{dis}$ , with the results shown in Table 1. It is apparent that using all three regularization terms simultaneously can reduce three kinds of distance metrics without sacrificing adversarial strength. When any regularization term is absent in  $\mathcal{L}_{dis}$ , all three distance metrics increase to varying degrees, indicating a decrease in imperceptibility, which proves the necessity of employing all three regularization terms. Additionally,  $\mathcal{L}_{hide}$  notably impacts CSD, indicating its efficacy in hiding deformation perturbations in complex areas. For more ablation experiments, please refer to the appendix section.

$\mathcal{L}_{ker}$	$\mathcal{L}_{cha}$	$\mathcal{L}_{hide}$	ASR $\uparrow$	CSD $\downarrow$	Uniform $\downarrow$	KNN $\downarrow$
	✓	✓	<b>100.00</b>	0.6145	0.3542	0.8196
✓		✓	<b>100.00</b>	0.5020	0.3412	0.7814
✓	✓		<b>100.00</b>	0.8828	0.2945	0.7729
✓	✓	✓	<b>100.00</b>	<b>0.4709</b>	<b>0.2883</b>	<b>0.7447</b>

Table 3. The results of ablation experiment on different regularization term constraints in  $\mathcal{L}_{dis}$ . The KNN distances in the table also need to be multiplied by  $10^{-3}$ .

## 6. Conclusion

In this paper, we investigate and analyze the inherent issue of poor imperceptibility in point cloud adversarial attacks. We devise an imperceptible shape-based point cloud adversarial attack method called HiT-ADV, which conceals deformation perturbations in complex surface areas. Furthermore, we propose to suppress digital adversarial strength by benign rigid transformations and benign resampling. Our work demonstrates that the robustness of existing point cloud models remain vulnerable. We hope to inspire efforts to enhance true robustness.

**Acknowledgment.** Supported in part by the National Key R&D Program of China (Grant No. 2022ZD0118100), in part by National Natural Science Foundation of China (No. 62025604), in part by Shenzhen Science and Technology Program (Grant No. KQTD20221101093559018), in part by Guangdong Major Project of Basic and Applied Basic Research (Grant No. 2023B0303000010).



## References

- [1] Yulong Cao, Chaowei Xiao, Dawei Yang, Jing Fang, Ruigang Yang, Mingyan Liu, and Bo Li. Adversarial objects against lidar-based autonomous driving systems. *arXiv preprint arXiv:1907.05418*, 2019. [1](#)
- [2] 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. [3](#)
- [3] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. [5](#)
- [4] Yao Deng, Xi Zheng, Tianyi Zhang, Chen Chen, Guannan Lou, and Miryung Kim. An analysis of adversarial attacks and defenses on autonomous driving models. In *2020 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, pages 1–10. IEEE, 2020. [1](#)
- [5] 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](#), [3](#), [5](#), [6](#)
- [6] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. [1](#), [2](#)
- [7] Jindong Gu. Interpretable graph capsule networks for object recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1469–1477, 2021. [1](#)
- [8] Jindong Gu, Hengshuang Zhao, Volker Tresp, and Philip HS Torr. Segpgd: An effective and efficient adversarial attack for evaluating and boosting segmentation robustness. In *European Conference on Computer Vision*, pages 308–325. Springer, 2022. [1](#)
- [9] Meng-Hao Guo, Jun-Xiong Cai, Zheng-Ning Liu, Tai-Jiang Mu, Ralph R Martin, and Shi-Min Hu. Pct: Point cloud transformer. *Computational Visual Media*, 7(2):3, 2021. [2](#)
- [10] Bangyan He, Jian Liu, Yiming Li, Siyuan Liang, Jingzhi Li, Xiaojun Jia, and Xiaochun Cao. Generating transferable 3d adversarial point cloud via random perturbation factorization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 764–772, 2023. [1](#)
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [1](#)
- [12] Qidong Huang, Xiaoyi Dong, Dongdong Chen, Hang Zhou, Weiming Zhang, and Nenghai Yu. Shape-invariant 3d adversarial point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15335–15344, 2022. [2](#), [3](#), [6](#)
- [13] Xiaoshui Huang, Guofeng Mei, Jian Zhang, and Rana Abbas. A comprehensive survey on point cloud registration. *arXiv preprint arXiv:2103.02690*, 2021. [1](#)
- [14] Xiaojun Jia, Xingxing Wei, Xiaochun Cao, and Hassan Foroosh. Comdefend: An efficient image compression model to defend adversarial examples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6084–6092, 2019. [1](#)
- [15] Xiaojun Jia, Xingxing Wei, Xiaochun Cao, and Xiaoguang Han. Adv-watermark: A novel watermark perturbation for adversarial examples. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 1579–1587, 2020. [1](#)
- [16] Jesse Levinson, Jake Askeland, Jan Becker, Jennifer Dolson, David Held, Soeren Kammel, J Zico Kolter, Dirk Langer, Oliver Pink, Vaughan Pratt, et al. Towards fully autonomous driving: Systems and algorithms. In *2011 IEEE intelligent vehicles symposium (IV)*, pages 163–168. IEEE, 2011. [1](#)
- [17] Jiaxin Li, Ben M Chen, and Gim Hee Lee. So-net: Self-organizing network for point cloud analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9397–9406, 2018. [2](#)
- [18] Ruihui Li, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng. Pu-gan: A point cloud upsampling adversarial network. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. [6](#)
- [19] Siyuan Liang, Xingxing Wei, Siyuan Yao, and Xiaochun Cao. Efficient adversarial attacks for visual object tracking. In *CEuropean Conference on Computer Vision*, 2020. [1](#)
- [20] Siyuan Liang, Baoyuan Wu, Yanbo Fan, Xingxing Wei, and Xiaochun Cao. Parallel rectangle flip attack: A query-based black-box attack against object detection. *arXiv preprint arXiv:2201.08970*, 2022. [1](#)
- [21] Nan Lin, Yuxuan Li, Keke Tang, Yujun Zhu, Xiayu Zhang, Ruolin Wang, Jianmin Ji, Xiaoping Chen, and Xinming Zhang. Manipulation planning from demonstration via goal-conditioned prior action primitive decomposition and alignment. *IEEE Robotics and Automation Letters*, 7(2):1387–1394, 2022. [1](#)
- [22] Aishan Liu, Shiyu Tang, Siyuan Liang, Ruihao Gong, Boxi Wu, Xianglong Liu, and Dacheng Tao. Exploring the relationship between architectural design and adversarially robust generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4096–4107, 2023. [2](#)
- [23] Binbin Liu, Jinlai Zhang, and Jihong Zhu. Boosting 3d adversarial attacks with attacking on frequency. *IEEE Access*, 10:50974–50984, 2022. [2](#), [6](#)
- [24] Daniel Liu, Ronald Yu, and Hao Su. Extending adversarial attacks and defenses to deep 3d point cloud classifiers. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 2279–2283. IEEE, 2019. [1](#), [2](#), [6](#), [7](#)
- [25] Daniel Liu, Ronald Yu, and Hao Su. Adversarial shape perturbations on 3d point clouds. In *Computer Vision—ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 88–104. Springer, 2020. [2](#), [3](#)
- [26] Yongcheng Liu, Bin Fan, Shiming Xiang, and Chunhong Pan. Relation-shape convolutional neural network for point cloud analysis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8895–8904, 2019. [2](#)
- [27] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learn-

- ing models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017. **1**
- [28] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. **1, 2, 6**
- [29] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. **2**
- [30] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointcnn: 3d object proposal generation and detection from point cloud. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 770–779, 2019. **1**
- [31] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. **1**
- [32] Keke Tang, Jianpeng Wu, Weilong Peng, Yawen Shi, Peng Song, Zhaoquan Gu, Zhihong Tian, and Wenping Wang. Deep manifold attack on point clouds via parameter plane stretching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2420–2428, 2023. **2, 3**
- [33] Liyao Tang, Yibing Zhan, Zhe Chen, Baosheng Yu, and Dacheng Tao. Contrastive boundary learning for point cloud segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8489–8499, 2022. **1**
- [34] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6411–6420, 2019. **2**
- [35] Tzungyu Tsai, Kaichen Yang, Tsung-Yi Ho, and Yier Jin. Robust adversarial objects against deep learning models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 954–962, 2020. **1, 2, 3, 6**
- [36] James Tu, Mengye Ren, Sivabalan Manivasagam, Ming Liang, Bin Yang, Richard Du, Frank Cheng, and Raquel Urtasun. Physically realizable adversarial examples for lidar object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13716–13725, 2020. **1**
- [37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. **1**
- [38] Sai Vemprala and Ashish Kapoor. Adversarial attacks on optimization based planners. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9943–9949. IEEE, 2021. **1**
- [39] Xupeng Wang, Mumuxin Cai, Ferdous Sohel, Nan Sang, and Zhengwei Chang. Adversarial point cloud perturbations against 3d object detection in autonomous driving systems. *Neurocomputing*, 466:27–36, 2021. **1**
- [40] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (tog)*, 38(5):1–12, 2019. **2, 6**
- [41] Geoffrey S Watson. Smooth regression analysis. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 359–372, 1964. **4**
- [42] Xingxing Wei, Siyuan Liang, Ning Chen, and Xiaochun Cao. Transferable adversarial attacks for image and video object detection. *arXiv preprint arXiv:1811.12641*, 2018. **1**
- [43] Yuxin Wen, Jiehong Lin, Ke Chen, CL Philip Chen, and Kui Jia. Geometry-aware generation of adversarial point clouds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6):2984–2999, 2020. **1, 2, 3, 4, 6**
- [44] Wenxuan Wu, Zhongang Qi, and Li Fuxin. Pointconv: Deep convolutional networks on 3d point clouds. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 9621–9630, 2019. **2**
- [45] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. **5**
- [46] Ziyi Wu, Yueqi Duan, He Wang, Qingnan Fan, and Leonidas J Guibas. If-defense: 3d adversarial point cloud defense via implicit function based restoration. *arXiv preprint arXiv:2010.05272*, 2020. **2**
- [47] Chong Xiang, Charles R Qi, and Bo Li. Generating 3d adversarial point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9136–9144, 2019. **2, 3, 6**
- [48] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *Proceedings of the British Machine Vision Conference 2016*. British Machine Vision Association, 2016. **1**
- [49] Jinlai Zhang, Lyujie Chen, Binbin Liu, Bo Ouyang, Qizhi Xie, Jihong Zhu, Weiming Li, and Yanmei Meng. 3d adversarial attacks beyond point cloud. *Information Sciences*, 633:491–503, 2023. **2, 3, 6**
- [50] Tianhang Zheng, Changyou Chen, Junsong Yuan, Bo Li, and Kui Ren. Pointcloud saliency maps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1598–1606, 2019. **2, 3, 4**
- [51] Hang Zhou, Kejiang Chen, Weiming Zhang, Han Fang, Wenbo Zhou, and Nenghai Yu. Dup-net: Denoiser and up-sampler network for 3d adversarial point clouds defense. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1961–1970, 2019. **2, 7**
- [52] Hang Zhou, Dongdong Chen, Jing Liao, Kejiang Chen, Xiaoyi Dong, Kunlin Liu, Weiming Zhang, Gang Hua, and Nenghai Yu. Lg-gan: Label guided adversarial network for flexible targeted attack of point cloud based deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10356–10365, 2020. **1**
- [53] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4490–4499, 2018. **1**