

Minimal Adversarial Examples for Deep Learning on 3D Point Clouds

Jaeyeon Kim¹ Binh-Son Hua^{2,3} Duc Thanh Nguyen⁴ Sai-Kit Yeung¹

¹Hong Kong University of Science and Technology

²VinAI Research, Vietnam

⁴Deakin University

³VinUniversity, Vietnam

Abstract

With recent developments of convolutional neural networks, deep learning for 3D point clouds has shown significant progress in various 3D scene understanding tasks, e.g., object recognition, semantic segmentation. In a safety-critical environment, it is however not well understood how such deep learning models are vulnerable to adversarial examples. In this work, we explore adversarial attacks for point cloud-based neural networks. We propose a unified formulation for adversarial point cloud generation that can generalise two different attack strategies. Our method generates adversarial examples by attacking the classification ability of point cloud-based networks while considering the perceptibility of the examples and ensuring the minimal level of point manipulations. Experimental results show that our method achieves the state-of-the-art performance with higher than 89% and 90% of attack success rate on synthetic and real-world data respectively, while manipulating only about 4% of the total points.

1. Introduction

Deep learning has shown great potentials in solving a wide spectrum of computer vision tasks. In life-crucial applications, one concern is that deep neural networks can be vulnerable to adversarial examples, a special kind of inputs that can fool the networks to make undesirable predictions. Several adversarial attack techniques have been proposed to generate such examples. In contrast, adversarial defense methods have been developed to detect and neutralise adversarial examples. Therefore, understanding how adversarial attacks and defenses operate is of great importance to make deep learning techniques more reliable and robust.

With the growing popularity of low-cost 3D sensors and light-field cameras, the community has started investigating the vulnerability of deep networks on 3D data, especially 3D point clouds [40, 19, 20, 43, 36]. However, existing works focus on common scenarios, such as generating adversarial point clouds by perturbing points. While such approaches have high attack success rates, the perturbations

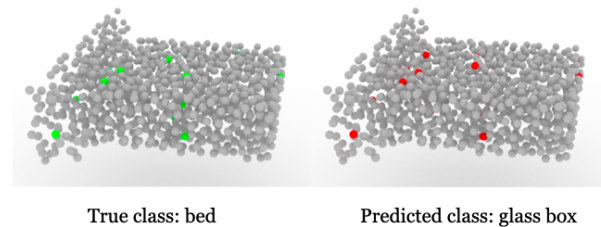


Figure 1: Left: input point cloud (1,024 points) classified correctly by PointNet [27] and 15 selected points (in green). Right: adversarial point cloud misclassified by PointNet and perturbed locations (in red).

are not imperceptible and can be identified easily by outlier detection or noise removal algorithms. In addition, existing adversarial attack methods do not perform optimally since all points in a point cloud are involved in the manipulation.

In this work, we study 3D adversarial attack in a more extreme but practical setting: how to generate an adversarial point cloud with minimum number of points perturbed from an original point cloud while maintaining the perceptibility of the original point cloud (see Figure 1). To address this problem, we propose a new formulation for adversarial point cloud generation that can be adapted to different attack strategies. The novelty of our work lies in both the research problem and proposed solution. Specifically, minimal 3D point cloud attack is an unexplored problem. We are also the first to propose a formulation that (i) considers both the perceptibility and optimality of adversarial samples, and (ii) generalises both point perturbation and addition in a unified framework. In summary, our contributions include:

- A new technique to generate minimal 3D adversarial point clouds;
- A unified formulation that generalises two adversarial point cloud generation strategies: point perturbation and point addition;
- A vulnerability analysis on the relation between the perturbed points found by our method and the concept of

critical points in PointNet [27];

- A benchmark of adversarial attacks on both synthetic and real-world 3D point clouds, which shows our method achieves consistent performance over both domains.

2. Related Work

3D Deep Learning. Recent availability of 3D datasets [39, 5, 13, 2, 7, 4, 34] has led to significant advances in deep learning on 3D data. In this domain, most existing works focus on designing convolution operations that enable convolutional neural networks to learn features directly from point clouds [27, 28, 18, 41, 14, 35, 31, 15, 47]. Several attempts have been made to create rotation invariant convolutions [29, 46, 26, 44]. Such convolutions allow scene understanding tasks, e.g., object recognition, semantic segmentation, to be trained directly with various input point clouds. Other improvements have also been found in the literature. For instance, Liu et al. [21] proposed an RNN-based model to extract the correlations of local areas. Yan et al. [42] extended the pointwise MLP network by adopting adaptive sampling to handle outliers and noise.

In this work, we adopt PointNet [27] as a target neural network to investigate important aspects of our method due to its popularity. However, we also verify our method with different deep learning models including PointNet++ [28], DGCNN [35], SpiderCNN [41], and PointASNL [42].

Adversarial Point Clouds. There exist studies on adversarial attacks and defenses for point cloud classification [40, 19, 20, 43, 48, 33, 12, 50, 17]. For instance, Xiang et al. [40] proposed an algorithm for point perturbation and addition based on the attack framework in [3] using Chamfer and Hausdorff distance. Liu et al. [19] extended the fast gradient sign method in [11] for constructing 3D adversarial examples using mesh and clipping norm. The basic ideas in these works follow previous adversarial attack techniques in 2D domain, which focus on how a point cloud should be perturbed to make an adversarial example. Readers are referred to [45] for a comprehensive review on adversarial attacks and defenses on images, and to [16] for a general analysis and proof of existence of adversarial examples.

It is also possible to mix newly added points with perturbed points in a point cloud to make adversarial examples [20]. Yuxin et al. [36] considered the consistency of local curvatures in a loss function to guide perturbed points lean towards object surfaces. Tsai et al. [33] incorporated the K-Nearest Neighbor loss in [40] to constrain adversarial samples to become physical objects. Zhao et al. [48] showed the vulnerability of isometry transformation by making perturbations on isometry matrix. Hadmi et al. [12] applied auto-encoders in perturbation to improve

the transferability of adversarial examples across networks. Lee et al. [17] added perturbation noise into the latent space of auto-encoders to maintain the shape of input point clouds. Zhou et al. [50] employed GAN [10] in generating adversarial point clouds with predefined target labels.

In addition to creating adversarial examples by perturbing existing points or adding new points into point clouds, one can fool classification techniques by removing points from input data [49, 37, 43]. For instance, Zheng et al. [49] eliminated important points on saliency map of an input object. Matthew et al. [37] iteratively removed critical points from an input point cloud until a target classification technique failed to classify the point cloud.

In contrast to adversarial attacks, countermeasures for adversarial point clouds have been so far scarce. Typical defense approaches include outlier or salient point removal [19] and noise removal [43, 51]. Recently, Dong et al. [8] used relative position of each local part of a clean point cloud to global object center as adversarial indicator. Wu et al. [38] proposed a method to predict implicit functions capturing clean shapes of point clouds.

Minimal Adversarial Attacks. In 2D domain, there is a specific family of techniques that focus on perturbing a minimum number of pixels in adversarial attacks. For instance, Papernot et al. [25] perturbed pixels on saliency maps. Carlini et al. [3] extended this method and used ℓ_0 -norm optimisation to minimise the number of pixels to perturb. Recently, Modas et al. [22] and Croce et al. [6] focused on how to perturb a sparse set of pixels while still achieving good perceptibility. Local search and evolutionary algorithms were also applied to obtain sparse perturbations in [23, 32, 30].

In this paper, we also explore adversarial attacks that only manipulate a minimal set of points. However, unlike the above works, we propose a new formulation that is general and can be adapted to various adversarial point cloud generation strategies. In addition, we also consider the perceptibility of adversarial examples in our formulation.

3. Proposed Method

Our problem of interest can be stated as follows. Let $P = \{\mathbf{p}_1, \dots, \mathbf{p}_N\}$ be an input set of N points where each point \mathbf{p}_i is represented by a vector of its coordinates $\mathbf{p}_i = [p_{i,x}, p_{i,y}, p_{i,z}]^\top \in \mathbb{R}^3$. Let F denote a point-based neural network, e.g., PointNet [27], and $F_i(P)$ denote the probability that the point set P is classified into the i -th class. Ideally, if i^* is the true class label of the point cloud P , then $i^* = \operatorname{argmax}_i F_i(P)$. Let P' be an adversarial example generated from P . We aim to find P' that satisfies the following conditions:

- (i) The perceptibility of P is maintained, i.e., the gener-

ated point cloud P' should not much deform from P ;

- (ii) A minimum number of points in P are manipulated;
- (iii) $\operatorname{argmax}_i F_i(P) \neq \operatorname{argmax}_{i'} F_{i'}(P')$, i.e., P' and P are classified into different classes by the network F .

Note that generation of adversarial point clouds by compromising the attack success rate and perceptibility of adversarial samples has been investigated in the literature, e.g., [40]. However, our work differs from [40] in the following points. First, our focus is untargeted attack while that is targeted attack in [40]. Second, the number of manipulated points is not considered in [40], leading to extremely high numbers of points are manipulated (as shown in our experiments). Third, while critical points proposed in PointNet [27] are used to drive the solution in [40], we show that our method could reach those critical points yet occupy a small portion, proving the capability of our method of finding compact yet vulnerable point sets.

In the following, we present a new formulation to generate P' using ℓ_0 -norm optimisation and describe in details how our formulation can be applied to point perturbation and point addition.

3.1. Point Perturbation

3.1.1 Formulation

Given a point cloud P , we aim to find a minimal set of points that can be shifted to generate an adversarial point cloud P' to attack the network F . We express the selection of points in P for perturbation by a binary indication vector $\mathbf{a} = [a_1, \dots, a_N]^T \in \{0, 1\}^N$ where a_i is 1 if \mathbf{p}_i is selected, and 0 otherwise. Suppose that $E = \{\mathbf{e}_1, \dots, \mathbf{e}_N\}$ is the set of perturbations, in which $\mathbf{e}_i = [e_{i,x}, e_{i,y}, e_{i,z}]^T \in \mathbb{R}^3$ is the perturbation vector to be applied on \mathbf{p}_i to obtain \mathbf{p}'_i . Applying perturbation set E on the point cloud P results in an adversarial point cloud P' as

$$P' = \{\mathbf{p}'_i = \mathbf{p}_i + a_i \mathbf{e}_i \mid \mathbf{p}_i \in P\}. \quad (1)$$

The process of generation of P' can be formulated as,

$$\begin{aligned} \min_{\mathbf{a}, E} f(P, \mathbf{a}, E) &= \min_{\mathbf{a}, E} \{\lambda_1 \|\mathbf{a}\|_0 + \lambda_2 D(P, P')\} \\ \text{s.t.} \quad \operatorname{argmax}_i F_i(P) &\neq \operatorname{argmax}_{i'} F_{i'}(P') \end{aligned} \quad (2)$$

where $\|\mathbf{a}\|_0 = \#\{i : a_i \neq 0, i = 1, \dots, N\}$ is the ℓ_0 -norm of \mathbf{a} (i.e., the number of non-zero elements in \mathbf{a}) and $D(P, P')$ is some distance between P and P' .

The optimisation problem defined in Eq. (2) covers all the aforementioned conditions. In particular, the first term, $\|\mathbf{a}\|_0$ in the objective function $f(P, \mathbf{a}, E)$ imposes the quantity of selected points in the point selection process while the second term, $D(P, P')$ constrains the perceptibility of

the adversarial point cloud P' w.r.t. the original point cloud P . As will be explained later in this section, $D(P, P')$ can be defined using different distance metrics. The constraint $\operatorname{argmax}_i F_i(P) \neq \operatorname{argmax}_{i'} F_{i'}(P')$ ensures the generated point cloud P' can fool the network F , i.e., F would not classify P and P' into the same class.

3.1.2 Perceptibility

There are several ways to realise the perceptibility $D(P, P')$ in Eq. (2). If we assume the correspondence between each point $\mathbf{p}_i \in P$ and its perturbed point $\mathbf{p}'_i \in P'$ defined in Eq. (1) is maintained, then we can define $D(P, P')$ using the Euclidean distances between \mathbf{p}_i and \mathbf{p}'_i as,

$$D_{Euclidean}(P, P') = \frac{1}{N} \sum_{i=1}^N (a_i \|\mathbf{e}_i\|_2). \quad (3)$$

However, such correspondences are not always well defined, e.g., when the number of points changes in the case of point addition, making Euclidean distance not a valid choice. We further propose to use Chamfer distance and Hausdorff distance to measure perceptibility. Specifically, we can define $D(P, P')$ as

$$\begin{aligned} D_{Chamfer}(P, P') &= \max \left\{ \frac{1}{|P|} \sum_{\mathbf{p}_i \in P} \min_{\mathbf{p}'_j \in P'} \|\mathbf{p}_i - \mathbf{p}'_j\|_2, \right. \\ &\quad \left. \frac{1}{|P'|} \sum_{\mathbf{p}'_j \in P'} \min_{\mathbf{p}_i \in P} \|\mathbf{p}'_j - \mathbf{p}_i\|_2 \right\} \end{aligned} \quad (4)$$

or

$$\begin{aligned} D_{Hausdorff}(P, P') &= \max \left\{ \max_{\mathbf{p}_i \in P} \left\{ \min_{\mathbf{p}'_j \in P'} \|\mathbf{p}_i - \mathbf{p}'_j\|_2 \right\}, \right. \\ &\quad \left. \max_{\mathbf{p}'_j \in P'} \left\{ \min_{\mathbf{p}_i \in P} \|\mathbf{p}'_j - \mathbf{p}_i\|_2 \right\} \right\}. \end{aligned} \quad (5)$$

As shown in Eq. (4)-(5), both Chamfer distance and Hausdorff distance do not require the same of number of points in the point clouds P and P' . Hence, they can be adapted easily to different point generation methods, e.g., point addition as presented in Section 3.2.

3.1.3 Relaxed Formulation

To solve the constrained optimisation problem in Eq. (2), we convert it into an unconstrained optimisation problem using a Lagrange multiplier-like form as:

$$\begin{aligned} \min_{\mathbf{a}, E} f(P, \mathbf{a}, E) \\ = \min_{\mathbf{a}, E} \{\lambda_1 \|\mathbf{a}\|_0 + \lambda_2 D(P, P') + h(P')\} \end{aligned} \quad (6)$$

where, like [40], we define,

$$h(P') = \max \left\{ 0, F_{i^*}(P') - \max_{i^* \neq i'} F_{i'}(P') \right\} \quad (7)$$

where i^* is the true class label of the point set P .

Since the problem in Eq. (6) is NP-hard in general [24], we further relax it as ℓ_1 -norm optimisation [1, 9] as:

$$\begin{aligned} & \min_{\hat{\mathbf{a}}, E} f(P, \hat{\mathbf{a}}, E) \\ & = \min_{\hat{\mathbf{a}}, E} \{ \lambda_1 \|\hat{\mathbf{a}}\|_1 + \lambda_2 D(P, P') + h(P') \} \end{aligned} \quad (8)$$

where $\hat{\mathbf{a}} = [\hat{a}_1, \dots, \hat{a}_N]^\top \in [0, 1]^N$, and $\|\hat{\mathbf{a}}\|_1 = \sum_{i=1}^N \hat{a}_i$ is the ℓ_1 -norm of $\hat{\mathbf{a}}$.

To solve Eq. (8), we apply the iterative gradient method in [11, 19]. Since the final aim is to obtain a binary vector for \mathbf{a} , we randomly initialise $\hat{\mathbf{a}}$ with near-binary values, i.e., \hat{a}_i is randomly set to either 0.0001 or 0.9999. Near-binary values are used to give \hat{a}_i chances to turn into 1 (or 0) if \mathbf{p}_i is selected (or otherwise). Finally, we achieve the final solution for \mathbf{a} as,

$$a_i = \begin{cases} 0, & \text{if } \hat{a}_i = 0 \\ 1, & \text{otherwise.} \end{cases} \quad (9)$$

Note that we aim to find solutions for both \mathbf{a} and \mathbf{e} . Given \mathbf{a} obtained from Eq. (9), we only consider perturbations e_i if $a_i = 1$.

The optimisation problem defined in Eq. (8) formulates our proposed adversarial generation method. This formulation is general and can be adapted conveniently to other adversarial generation strategies, e.g., point addition.

3.2. Point Addition

In addition to point perturbation, we can generate an adversarial example P' by extending P with a minimum number of additional points. We show that our proposed formulation in Eq. (8) can also be applied in this task. Specifically, suppose that there are no more than K points added to the original point cloud P . We can construct a new point set \tilde{P} including all the points in P and K new points. These K new points can be generated by randomly choosing K points in P and adding them to \tilde{P} . We note that this way of construction of \tilde{P} does not change the perceptibility of P as $D(P, \tilde{P}) = 0$ for either the Chamfer distance or Hausdorff distance used to define $D(P, \tilde{P})$. In addition, both P and \tilde{P} are treated equally by the network F , i.e., $\forall i, F_i(P) = F_i(\tilde{P})$, as the geometric structure of the point clouds remains unchanged.

Similarly, we also construct a vector $\tilde{\mathbf{a}} = [\tilde{a}_1, \dots, \tilde{a}_{N+K}]^\top \in [0, 1]^{N+K}$ and a perturbation set $\tilde{E} = \{\tilde{\mathbf{e}}_1, \dots, \tilde{\mathbf{e}}_{N+K}\}$ by extending \mathbf{a} and E with K new elements and solve the optimisation problem in

Eq. (8) with a new objective function $f(\tilde{P}, \tilde{\mathbf{a}}, \tilde{E})$. The vector $\tilde{\mathbf{a}}$ is initialised as follows, \tilde{a}_i is set to 0 for $i \in \{1, \dots, N\}$, and to a random value in $\{0.0001, 0.9999\}$ for $i \in \{N+1, \dots, N+K\}$. Furthermore, during the optimisation process, we fix $\tilde{a}_i = 0, \forall i = 1, \dots, N$, i.e., original points in P will not be changed. Finally, the adversarial point cloud P' is obtained by including points $\tilde{\mathbf{p}}_i \in \tilde{P}$ such that $\tilde{a}_i = 1$.

4. Experiments and Results

4.1. Experiment Setup

Datasets. We experimented our method on ModelNet40 [39] and ScanObjectNN [34] dataset. ModelNet40 is a benchmark dataset for classification of 3D CAD models. It consists of 9,843 models for training and 2,468 models for testing. We followed the experimental setup in [27] to sample the surfaces of the models in ModelNet40 uniformly and normalised points into a unit cube. ScanObjectNN is an object dataset from real-world indoor scans including practical challenges such as view occlusions and object partiality. It has 15,000 objects organised in five challenging variants, e.g., objects with background, translated objects, rotated objects, and scaled objects. In our experiments, we used ‘OBJ_BG’, the most challenging variant including objects with background. We followed Uy et al. [34] to normalise the point clouds containing background using mean and furthest point distance.

Implementation Details. We adopted PointNet [27] as a test base network to conduct important experiments. To adapt with point clouds of varying sizes, we modified the max-pooling operator and batch norm accordingly. We used Adam optimiser with learning rate of 0.01. For our adversarial attack algorithm, we performed exhaustive search for the parameters in Eq. (8) and empirically set them as $\lambda_1 = 0.15$ and $\lambda_2 = 50$. Iterative gradient method in [19] with 250 iterations was employed to solve Eq. (8).

4.2. Evaluation and Comparison

We evaluated our method based on attack success rate, perceptibility of adversarial examples, and average number of manipulated points. The perceptibility of adversarial examples was measured using Chamfer and Hausdorff distance. We report the performance of our method on ModelNet40 and ScanObjectNN dataset in Table 1 & 2 respectively.

We also compared our method with existing methods. In particular, we tested the method by Xiang et al. [40], which also generated adversarial examples by point perturbation and addition. Since this method aimed for targeted attack, for fair comparison, we altered it to untargeted attack. We re-implemented the method by Liu et al. [20] which made attacks in two manners: adversarial sink (i.e.,

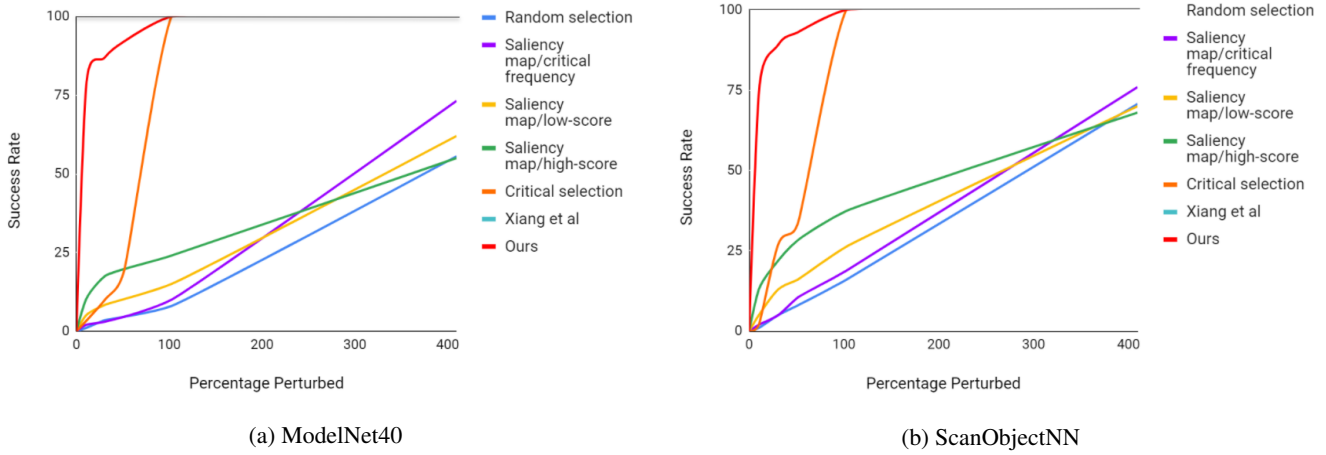


Figure 2: Performance trend of our adversarial attack and existing methods.

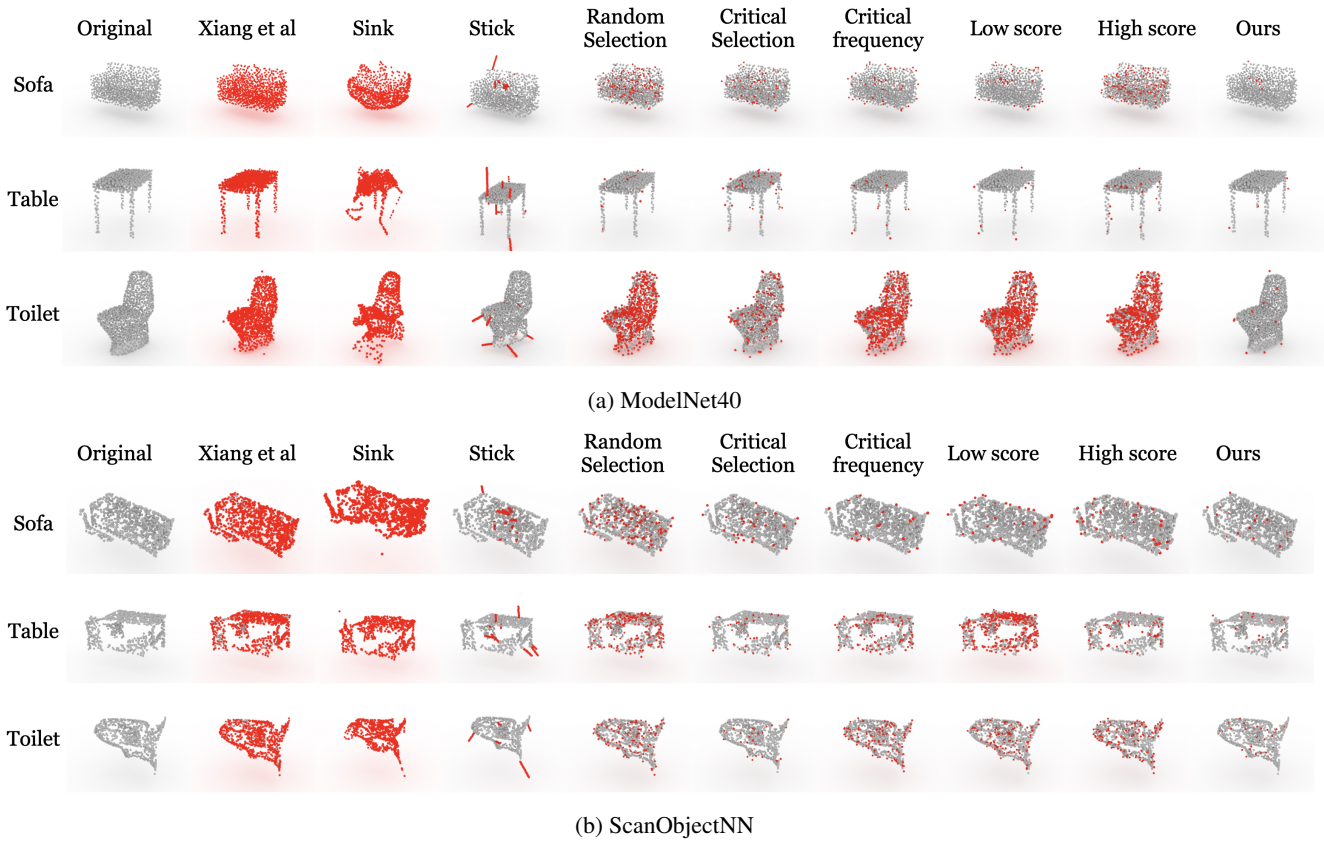


Figure 3: Adversarial examples generated by our method and existing methods. Red points represent perturbed points. On ScanObjectNN, our method is able to create examples that are indistinguishable from realistic noise.

pulling points towards a sink point) and adversarial stick (i.e., resampling points on added sticks). In addition, we evaluated the method by Wicker and Kwiatkowska [37], and by Zheng et al. [49]. Since both [37] and [49] created adversarial examples by removing points, for fair comparison,

we applied their point selection strategies but then replaced the point removal by point perturbation and point addition. In detail, the authors in [37] selected points either randomly or from a critical point set determined by PointNet [27]. In [49], selected points were sampled from saliency maps

of input point clouds in three different ways: (i) critical frequency (i.e., points frequently chosen by the max-pooling operator in PointNet [27]), (ii) low-score (i.e., points having small loss gradient for a target network), and (iii) high-score (i.e., point having large loss gradient). We will further discuss the “critical points” in Section 4.3. We used 250 iterations for [40] and [20]. For other baseline methods, we perturbed the objects until a successful attack is made.

As shown in Table 1 & 2, compared with other methods, our method achieves the highest success rate yet lowest number of processing points for both point perturbation and point addition, and on both ModelNet40 and ScanObjectNN. Specifically, our method uses only 4% of the total input points (1,024) to reach > 89% and > 90% of success rate on ModelNet40 and ScanObjectNN respectively. We notice that our method performs consistently (in terms of both the attack rate and the number of points) on both synthetic and real-world datasets. The adversarial sink and adversarial stick in [20] respectively take the second place w.r.t. the success rate in point perturbation on ModelNet40 and ScanObjectNN. However, both of them require great numbers of points, especially the adversarial sink. The method in [40] is ranked third for its success rate but also incurs heavy point manipulations. In addition, we observe this method often generated obvious outliers, which could be detected easily by outlier removal methods. The saliency map-based attack method [49] and the one in [37] with random point selection require roughly 10× larger point sets than our method while achieving much lower success rates. The method in [37] with critical point selection shows relatively small number of points (though still more than our method), but the success rate is well below par. To further explore the performance trend, we plot the success rate over different numbers of points selected for point perturbation in Figure 2. As shown in the graphs, our method significantly outperforms existing ones with a small set of points.

Experimental results also show that our method generates high-imperceptibility adversarial samples, evident by their Chamfer and Hausdorff distances to input point clouds (see Table 1 & 2). We empirically observe that Hausdorff distance results in less outlier points than Chamfer distance. Figure 3 qualitatively compares adversarial samples generated by our method and others. As shown, our method less likely produces outliers. This is due to the use of object perceptibility and minimal point set in our formula. We observe that, adversarial examples created from real-world data (see Figure 3(b)) are neither noticeable in perception nor distinguishable from common noise.

4.3. Vulnerability Analysis

Qi et al. [27] proposed a notion called “critical points” that characterise the shape of a point cloud. These points can be identified from the max pooling layer in PointNet [27]. It

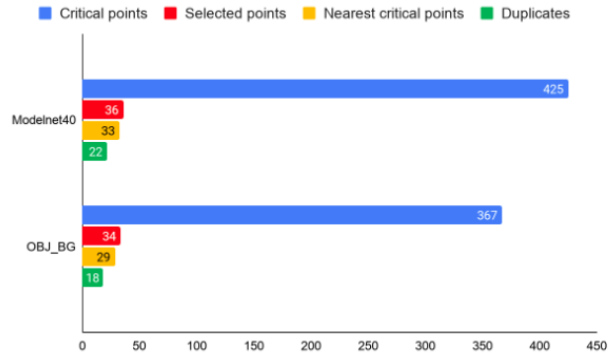


Figure 4: Our selected points vs critical points.

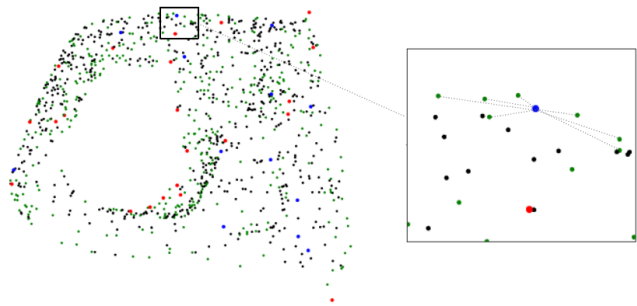


Figure 5: Adversarial example of a sink. Critical points are highlighted in green. Our selected points are marked in red (if they are identical to critical points) and in blue (if they are found within the 5 nearest points of a critical point).

is also indicated in [27] that, critical points play role as an indicator for object recognition. Therefore, modification of critical points may lead to wrong classification results.

We found that our adversarial generation algorithm could somehow reach those critical points. To confirm this, we measured the coincidence of our selected points and PointNet’s critical points. Specifically, we counted the duplicates in the two point sets, and the number of our selected points found within the 5 nearest points of a critical point. Recall that our algorithm randomly initialises the selected points (i.e., the vector \hat{a} in Eq. (8)). Figure 4 provides the numerical data of this experiment. It is shown that about 50% of the selected points are identical to the critical points and 80% of the selected points are close to the critical points. Figure 5 visualises our selected points and critical points. As shown, our selected points are close to critical points yet occupy a small portion, proving the capability of our method of finding compact yet vulnerable point sets.

We also experimented our method with two additional initialisation schemes for the selected points: critical points-based initialisation and all point-based initialisation (i.e., taking all points in a point cloud to initialise the vector \hat{a}). Table 3 reports the performance of various initialisation

	Success Rate	Chamfer Distance	Hausdorff Distance	# Points
Xiang et al. [40]	85.9	1.77×10^{-4}	2.38×10^{-2}	967
Adversarial sink [20]	88.3	7.65×10^{-3}	1.92×10^{-1}	1024
Adversarial stick [20]	83.7	4.93×10^{-3}	1.49×10^{-1}	210
Random selection [37]	55.56	7.47×10^{-4}	2.49×10^{-3}	413
Critical selection [37]	18.99	1.15×10^{-4}	9.39×10^{-3}	50
Saliency map/critical frequency [49]	63.15	5.72×10^{-4}	2.50×10^{-3}	303
Saliency map/low-score [49]	55.97	6.47×10^{-4}	2.50×10^{-3}	358
Saliency map/high-score [49]	58.39	7.52×10^{-4}	2.48×10^{-3}	424
Ours	89.38	1.55×10^{-4}	1.88×10^{-2}	36

(a) Point Perturbation

	Success Rate	Chamfer Distance	Hausdorff Distance	# Points
Xiang et al. [40]	73.59	7.98×10^{-3}	5.46×10^{-2}	200
Random selection [37]	43.90	2.16×10^{-4}	2.49×10^{-3}	121
Critical selection [37]	47.64	2.05×10^{-4}	2.50×10^{-3}	118
Saliency map/critical frequency [49]	45.13	2.13×10^{-4}	2.49×10^{-3}	118
Saliency map/low-score [49]	60.96	1.64×10^{-4}	2.50×10^{-3}	89
Saliency map/high-score [49]	41.06	2.27×10^{-4}	2.49×10^{-3}	128
Ours	89.01	1.53×10^{-4}	1.98×10^{-2}	38

(b) Point Addition

Table 1: Attack performance to PointNet on ModelNet40.

	Success Rate	Chamfer Distance	Hausdorff Distance	# Points
Xiang et al. [40]	81.32	1.13×10^{-4}	1.74×10^{-2}	959
Adversarial sink [20]	78.7	1.37×10^{-3}	9.81×10^{-2}	1023
Adversarial stick [20]	87.5	5.18×10^{-3}	1.67×10^{-1}	210
Random selection [37]	63.72	6.10×10^{-4}	2.50×10^{-3}	340
Critical selection [37]	47.99	1.99×10^{-4}	2.69×10^{-2}	70
Saliency map/critical frequency [49]	66.9	4.69×10^{-4}	2.50×10^{-3}	265
Saliency map/low-srop [49]	63.81	5.49×10^{-4}	2.50×10^{-3}	306
Saliency map/high-srop [49]	66.82	6.16×10^{-4}	2.47×10^{-3}	350
Ours	91.72	1.12×10^{-4}	1.15×10^{-2}	34

(a) Point Perturbation

	Success Rate	Chamfer Distance	Hausdorff Distance	# Points
Xiang et al [40]	69.26	6.07×10^{-3}	4.71×10^{-2}	200
Random selection [37]	60.05	1.77×10^{-4}	2.50×10^{-3}	97
Critical selection [37]	57.44	1.76×10^{-4}	2.50×10^{-3}	98
Saliency map/critical frequency [49]	59.63	1.79×10^{-4}	2.49×10^{-3}	97
Saliency map/low-score [49]	60.96	1.64×10^{-4}	2.50×10^{-3}	90
Saliency map/high-score [49]	57.87	1.87×10^{-4}	2.50×10^{-3}	103
Ours	90.44	1.08×10^{-4}	1.10×10^{-2}	38

(b) Point Addition

Table 2: Attack performance to PointNet on ScanObjectNN.

schemes. As shown in the results, the critical points-based initialisation results in the least number of points but in-

curs the lowest success rate. Moreover, this scheme deteriorates the perceptibility of adversarial examples (as shown in

	Success Rate	Chamfer Distance	Hausdorff Distance	# Points
Critical points	88.04	9.53×10^{-4}	8.28×10^{-3}	29
Using all points	91.12	1.15×10^{-4}	8.71×10^{-3}	60
Ours (random points)	90.16	1.09×10^{-4}	9.88×10^{-3}	37

Table 3: Attack performance with different initialisation strategies on ScanObjectNN.

	Success Rate	Chamfer Distance	Hausdorff Distance	# Points
Pointnet [27]	89.38	1.55×10^{-4}	1.88×10^{-2}	36
Pointnet++ [28]	88.76	5.10×10^{-4}	3.55×10^{-2}	59
DGCNN [35]	62.16	7.78×10^{-4}	3.54×10^{-2}	107
SpiderCNN [41]	89.92	4.71×10^{-4}	4.16×10^{-2}	57
PointASNL [42]	72.16	2.21×10^{-4}	2.03×10^{-2}	45

(a) Modelnet40

	Success Rate	Chamfer Distance	Hausdorff Distance	# Points
Pointnet [27]	91.72	1.12×10^{-4}	1.15×10^{-2}	34
Pointnet++ [28]	94.05	2.81×10^{-4}	2.06×10^{-2}	50
DGCNN [35]	66.46	6.80×10^{-4}	2.94×10^{-2}	103
SpiderCNN [41]	86.27	3.23×10^{-4}	2.18×10^{-2}	57
PointASNL [42]	55.88	2.07×10^{-4}	1.30×10^{-2}	47

(b) ScanObjectNN

Table 4: Attack performance to various network architectures on Modelnet40 and ScanObjectNN.

the Chamfer distances). Utilising all points for initialisation shows the opposite, i.e., more points are selected but high success rate is achieved. Our initialisation scheme compromises all the criteria, i.e., adversarial examples are created with low number of points, high attack rate, and acceptable perceptibility. In addition, random initialisation can be applied to other networks which do not support critical points.

4.4. Attack Performance to other Architectures

We also applied our adversarial example generation algorithm to attack existing point cloud architectures other than PointNet. Those architectures include PointNet++ [28], DGCNN [35], SpiderCNN [41], and PointASNL [42].

We report the attack performance of our method to these network architectures in Table 4. Amongst all the models, DGCNN [35] is shown to be the most robust one, which requires the greatest number of points to be fooled while maintaining low attack success rate. PointNet [27], on the other hand, appears to be the most fragile model, which can be fooled easily with less points to achieve high success rate. The remaining models can be attacked by slightly different numbers of points. Table 4 also shows that the attack performance of our method to a network architecture is consistent across both synthetic and real-world datasets.

5. Conclusion

In this paper, we propose a unified formulation for minimal adversarial 3D point clouds generation that can generalise two attack strategies including point perturbation and point addition. We experimented our method on benchmark datasets, and showed that existing point cloud neural networks, e.g., PointNet, are vulnerable to attacks that perturb only 4% of the points in a point cloud to reach more than 89% and 90% of success rate on synthetic and real-world data respectively. These results pose a challenge in developing countermeasures to defend against such attacks.

With increasingly more 3D data used in consumer devices, we envision that adversarial attack and defense for point clouds will become diverse, making this topic worthy for future research. For example, in this paper, we only investigated attacks by perturbing point cloud coordinates. For real-world point clouds, color could be another vulnerable channel for adversarial attacks. Besides, it is important to study how to create adversarial point clouds in physical world. The results of our work show that such a task could be practical as only a few percentages of an input point cloud need to be modified.

Acknowledgment. Sai-Kit Yeung was partially supported by an internal grant from HKUST (R9429).

References

- [1] Edoardo Amaldi and Viggo Kann. On the approximability of minimizing nonzero variables or unsatisfied relations in linear systems. *Theoretical Computer Science*, 209(1–2):237–260, 1998. 4
- [2] Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3d semantic parsing of large-scale indoor spaces. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [3] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy (SP)*, 2017. 2
- [4] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Habber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *International Conference on 3D Vision (3DV)*, 2017. 2
- [5] Angel Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qi-Xing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. Shapenet: An information-rich 3d model repository. In *arXiv:1512.03012*, 2015. 2
- [6] Francesco Croce and Matthias Hein. Sparse and imperceivable adversarial attacks. In *International Conference on Computer Vision (ICCV)*, 2019. 2
- [7] Angela Dai, Angel X Chang, Manolis Savva, Maciej Habber, Thomas Funkhouser, and Matthias Niessner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [8] Xiaoyi Dong, Dongdong Chen, Hang Zhou, Gang Hua, Weiming Zhang, and Nenghai Yu. Self-robust 3d point recognition via gather-vector guidance. In *Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [9] David L. Donoho and Michael Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via L1 minimization. *Proceedings of the National Academy of Sciences*, 100(5):2197–2202, 2003. 4
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NIPS)*, 2014. 2
- [11] Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations (ICLR)*, 2015. 2, 4
- [12] Abdullah Hamdi, Sara Rojas, Ali Thabet, and Bernard Ghanem. Advpc: Transferable adversarial perturbations on 3d point clouds. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [13] Binh-Son Hua, Quang-Hieu Pham, Duc Thanh Nguyen, Minh-Khoi Tran, Lap-Fai Yu, and Sai-Kit Yeung. Scenenn: A scene meshes dataset with annotations. In *International Conference on 3D Vision (3DV)*, 2016. 2
- [14] Binh-Son Hua, Minh-Khoi Tran, and Sai-Kit Yeung. Pointwise convolutional neural network. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [15] Qiangui Huang, Weiyue Wang, and Ulrich Neumann. Recurrent slice networks for 3d segmentation on point clouds. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [16] Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. In *Advances in Neural Information Processing Systems*, volume 32, 2019. 2
- [17] Kibok Lee, Zhuoyuan Chen, Xinchun Yan, Raquel Urtasun, and Ersin Yumer. Shapeadv: Generating shape-aware adversarial 3d point clouds. *arXiv preprint arXiv:2005.11626*, 2020. 2
- [18] Yangyan Li, Rui Bu, Mingchao Sun, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. *Advances in Neural Information Processing Systems (NIPS)*, 2018. 2
- [19] Daniel Liu, Ronald Yu, and Hao Su. Extending adversarial attacks and defenses to deep 3d point cloud classifiers. In *IEEE International Conference on Image Processing (ICIP)*, 2019. 1, 2, 4
- [20] Daniel Liu, Ronald Yu, and Hao Su. Adversarial point perturbations on 3d objects. In *European Conference on Computer Vision (ECCV) Workshop on Adversarial Robustness in the Real World*, 2020. 1, 2, 4, 6, 7
- [21] Xinhai Liu, Zhizhong Han, Yu-Shen Liu, and Matthias Zwicker. Point2sequence: Learning the shape representation of 3d point clouds with an attention-based sequence to sequence network. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2019. 2
- [22] Apostolos Modas, Seyed-Mohsen Moosavi-Dezfooli, and Pascal Frossard. Sparsefool: A few pixels make a big difference. In *Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [23] Nina Narodytska and Shiva Prasad Kasiviswanathan. Simple black-box adversarial perturbations for deep networks. *arXiv:1612.06299*, 2016. 2
- [24] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, 24(2):227–234, 1995. 4
- [25] Nicolas Papernot, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Berkay Celik, and Ananthram Swami. The limitations of deep learning in adversarial settings. In *IEEE European Symposium on Security and Privacy*, 2016. 2
- [26] Adrien Poulénard, Marie-Julie Rakotosaona, Yann Ponty, and Maks Ovsjanikov. Effective rotation-invariant point cnn with spherical harmonics kernels. *International Conference on 3D Vision (3DV)*, 2019. 2
- [27] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 3, 4, 5, 6, 8
- [28] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems (NIPS)*, 2017. 2, 8
- [29] Yongming Rao, Jiwen Lu, and Jie Zhou. Spherical fractal convolutional neural networks for point cloud recognition. In *Computer Vision and Pattern Recognition (CVPR)*, 2019. 2

- [30] Lukas Schott, Jonas Rauber, Matthias Bethge, and Wieland Brendel. Towards the first adversarially robust neural network model on MNIST. In *International Conference on Learning Representations (ICLR)*, 2019. 2
- [31] Yiru Shen, Chen Feng, Yaoqing Yang, and Dong Tian. Mining point cloud local structures by kernel correlation and graph pooling. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [32] Jiawei Su, Danilo Vasconcellos Vargas, and Kouichi Sakurai. One pixel attack for fooling deep neural networks. *IEEE Transactions on Evolutionary Computation*, 23(5):828–841, 2019. 2
- [33] Tzungyu Tsai, Kaichen Yang, Tsung-Yi Ho, and Yier Jin. Robust adversarial objects against deep learning models. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2020. 2
- [34] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Duc Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *International Conference on Computer Vision (ICCV)*, 2019. 2, 4
- [35] 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*, 38(5):1–12, 2019. 2, 8
- [36] Yuxin Wen, Jiehong Lin, Ke Chen, and Kui Jia. Geometry-aware generation of adversarial and cooperative point clouds. In *arXiv:1902.10899*, 2019. 1, 2
- [37] Matthew Wicker and Marta Kwiatkowska. Robustness of 3d deep learning in an adversarial setting. In *Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 5, 6, 7
- [38] 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
- [39] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. 2, 4
- [40] Chong Xiang, Charles R. Qi, and Bo Li. Generating 3d adversarial point clouds. In *Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 3, 4, 6, 7
- [41] Yifan Xu, Tianqi Fan, Mingye Xu, Long Zeng, and Yu Qiao. Spidercnn: Deep learning on point sets with parameterized convolutional filters. In *European Conference on Computer Vision (ECCV)*, 2018. 2, 8
- [42] Xu Yan, Chaoda Zheng, Zhen Li, Sheng Wang, and Shuguang Cui. Pointasnl: Robust point clouds processing using nonlocal neural networks with adaptive sampling. In *Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 8
- [43] Jiancheng Yang, Qiang Zhang, Rongyao Fang, Bingbing Ni, Jinxian Liu, and Qi Tian. Adversarial attack and defense on point sets. In *arXiv:1902.10899*, 2019. 1, 2
- [44] Yang You, Yujing Lou, Qi Liu, Yu-Wing Tai, Lizhuang Ma, Cewu Lu, and Weiming Wang. Pointwise rotation-invariant network with adaptive sampling and 3d spherical voxel convolution. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2020. 2
- [45] Xiaoyong Yuan, Pan He, Qile Zhu, and Xiaolin Li. Adversarial examples: Attacks and defenses for deep learning. *IEEE Transactions on Neural Networks and Learning Systems*, 30(9):2805–2824, 2019. 2
- [46] Zhiyuan Zhang, Binh-Son Hua, David W. Rosen, and Sai-Kit Yeung. Rotation invariant convolutions for 3d point clouds deep learning. In *International Conference on 3D Vision (3DV)*, 2019. 2
- [47] Zhiyuan Zhang, Binh-Son Hua, and Sai-Kit Yeung. Shellnet: Efficient point cloud convolutional neural networks using concentric shells statistics. In *International Conference on Computer Vision (ICCV)*, 2019. 2
- [48] Yue Zhao, Yuwei Wu, Caihua Chen, and Andrew Lim. On isometry robustness of deep 3d point cloud models under adversarial attacks. In *Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [49] Tianhang Zheng, Changyou Chen, Junsong Yuan, Bo Li, and Kui Ren. Pointcloud saliency maps. In *International Conference on Computer Vision (ICCV)*, 2019. 2, 5, 6, 7
- [50] 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 *Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [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 *International Conference on Computer Vision (ICCV)*, 2019. 2