

Local Aggressive Adversarial Attacks on 3D Point Cloud

Yiming Sun
Feng Chen *
Zhiyu Chen

Nanjing University of Posts and Telecommunications

1220045209@NJUPT.EDU.CN
CHENFENG1271@GMAIL.COM
1217053508@NJUPT.EDU.CN

Mingjie Wang

BNU-HKBU United International College

XIAOWANGIII@YAHOO.COM

Editors: Vineeth N Balasubramanian and Ivor Tsang

Abstract

Deep neural networks are found to be prone to adversarial examples which could deliberately fool the model to make mistakes. Recently, a few of works expand this task from 2D image to 3D point cloud by using global point cloud optimization. However, the perturbations of global point are not effective for misleading the victim model. First, not all points are important in optimization toward misleading. Abundant points account considerable distortion budget but contribute trivially to attack. Second, the multi-label optimization is suboptimal for adversarial attack, since it consumes extra energy in finding multi-label victim model collapse and causes instance transformation to be dissimilar to any particular instance. Third, the independent adversarial and perceptibility losses, caring misclassification and dissimilarity separately, treat the updating of each point equally without a focus. Therefore, once perceptibility loss approaches its budget threshold, all points would be stock in the surface of hypersphere and attack would be locked in local optimality. Therefore, we propose a local aggressive adversarial attacks (L3A) to solve above issues. Technically, we select a bunch of salient points, the high-score subset of point cloud according to gradient, to perturb. Then a flow of aggressive optimization strategies are developed to reinforce the unperceptive generation of adversarial examples toward misleading victim models. Extensive experiments on PointNet, PointNet++ and DGCNN demonstrate the state-of-the-art performance of our method against existing adversarial attack methods. Our code is available at <https://github.com/Chenfeng1271/L3A>.

Keywords: Point Clouds Processing, Adversarial Learning, Local Adversarial Attack

1. Introduction

Machine learning, especially deep learning, has achieved exceptional feats in autonomous driving, virtual reality, robotics etc. Among them, 3D sensor, such as LIDAR [Geiger et al. \(2013\)](#), which is a core ingredient of perceiving surrounding fine-grained environment, has drawn increasing attention from researchers. However, point cloud based deep learning methods [Zhou and Tuzel \(2018\)](#); [Liang et al. \(2019\)](#); [Shi et al. \(2019\)](#); [Mandikal et al. \(2018\)](#) are found to be vulnerable to adversarial examples where sometimes trivial movement of point cloud would mislead the deep learning model with high confidence.

1. The first authors are Yiming Sun and Feng Chen. The corresponding author is Zhiyu Chen

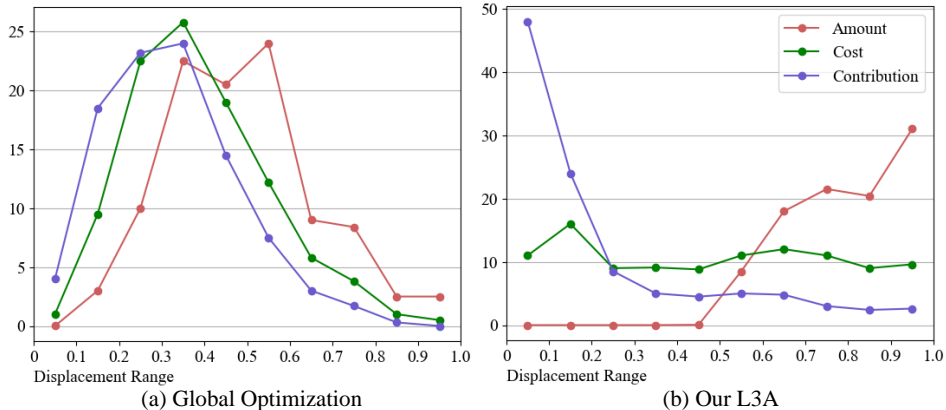


Figure 1: The statistic of amount, cost and contribution of points after attack using naive global optimization and our L3A. We divide the points in the adversarial example into 10 equal intervals according to the uniform displacement, and calculate the cost and contribution of the points. The cost of a point is defined as its displacement, and the contribution of a point is defined as the increase in classification confidence after its displacement is returned to zero. In this experiment, we set perturbation budget $\lambda=0.3$.

Naturally, to delve the weakness of deep neural networks (DNN) [Szegedy et al. \(2013\)](#); [Biggio et al. \(2013\)](#); [Athalye et al. \(2018\)](#) toward imperceptible point cloud perturbations, adversarial attack and defense are two important topics to research. Recently, while numerous approaches have been proposed to intensively study these two tasks in 2D image domain [Kurakin et al. \(2016b,a\)](#); [Madry et al. \(2017\)](#), due to the non-rigid and sparse attribute of point cloud, this expertise may not be directly applied to the 3D counterpart. The mainstream optimization-based approaches mainly use a pair of adversarial and distortion objectives to keep the generated point cloud misleading while imperceptible [Xiang et al. \(2019a\)](#); [Zhang et al. \(2019\)](#); [Liu et al. \(2020\)](#). However, this kind of approaches is always ineffective because of several reasons.

First, not all points in the point cloud are important to adversarial attack. Optimizing the whole points in point cloud is arduous. As shown in Figure 1 (a), We divide the points into 10 intervals based on their displacement distance, and count the costs and contributions of them. In the range of 0 to 0.3, over 40% of total points with small movement occupy considerable perceptibility budget while contribute trivially to attack. Besides, in the range of 0.6 to 1, even points move far away, they contribute only about 5% for attack in average. This makes the tradeoff of attack and imperceptibility poor.

Second, multi-label optimization is not the optimal choice for adversarial attack. We observe that some examples after the attack are more likely to be identified for a particular wrong class. However, since the multi-label optimization attack always optimizes

all the wrong classes, there will be competition between them. This kind of contradictory competition not only affects the attack performance, but also makes the optimization of adversarial example unstable and finally disturbs the imperceptibility. Therefore, cultivating the adversarial attack toward one preferred wrong class is more useful for attack.

Third, adversarial loss and perceptibility loss have opposite purpose which isn't adjustable during traditional optimization. As shown in Eq 2, if perceptibility loss \mathcal{D} is closing to the budget λ , especially when λ is extremely small, each point would be bounded in the surface of hypersphere and sequentially limit the attack performance. The main reason is each point equally moves toward attack purpose and greedily occupying distortion budget and none wants to withdraw.

In this paper, we propose a novel local aggressive adversarial attacks (L3A) method which is more universal and can simultaneously address these issues. The basic idea of our method is to encourage the cost-effectiveness between attack and perceptibility for each point, i.e., points moving with large displacement would contribute more for attack and vice versa. First, we use the local salient points selection strategy. Salient points could be deemed as local salient area for classification whose transformation would be more effective for attack, instead of the whole point cloud. Then, to model the example preference as a prior knowledge, we propose a novel score loss and contractive loss to guide the adversarial examples to change towards the preferred wrong class and suppress the other weaker wrong classes in attacking process. Besides, withdrawing perturbation budget algorithm is designed to make the optimization more adjustable between attack and perceptibility to avoid it falling into local optimality.

Our key contributions can be summarized as:

- 1) We propose a novel local adversarial point cloud attack method named L3A which solves the problems existing in the global-optimization based attack.
- 2) We propose to enhance the cost-effectiveness of points between attack and perceptibility by activating the focus of optimization. By letting examples updating toward their preference and adjusting attack and indistinguishability purpose for each point, our adversarial examples could largely increase the final successful rate.
- 3) Extensive experiments, including elaborate ablation study, attack and defense on ModelNet40, show that our method performs favorably against existing state-of-the-art literature on PointNet, PointNet++ and DCGNN.

2. Related Work

Deep learning on point clouds. Recently, the availability of point cloud data accelerate the development of deep learning on 3D point cloud. For projection-based methods [Beltrán et al. \(2018\)](#); [Yang et al. \(2018\)](#); [Li et al. \(2016\)](#), the dimensionality of the point cloud is compressed by projection to specific plane, such as the front view, as pseudo-image to take advantages of 2D CNN expertise, but it would inevitably cause the loss of point cloud spatial information. Compared with the projection-based method, voxel-based [Su et al. \(2015\)](#), [Wu et al. \(2016\)](#) and [Qi et al. \(2016\)](#) divide the 3D space into equal-spaced 3D voxels and then encode the spatial information of each voxel, but they require intensive computation which is not conducive to real-time processing. Besides, PointNet [Qi et al. \(2017a\)](#) first propose to handle continuous raw point cloud directly and achieve great performance. Then

PointNet++ Qi et al. (2017b) further improves it by investing local information. Similar to PointNet++, DGCNN Wang et al. (2019) constructs *knn* graphs and applies EdgeConv to capture local geometric structures. Even though these methods achieve accurate prediction of classification and recognition on large scale dataset, such as ModelNet40 Wu et al. (2015b) and ShapeNet Chang et al. (2015), they are found to be prone to adversarial attack Geiger et al. (2013).

2D adversarial attack. Szegedy et al. (2013), Athalye et al. (2018) and Biggio et al. (2013) first reveal that the 2D image classification neural network could be attacked to make wrong prediction. This attack is always achieved by optimizing an objective that maximizes the prediction errors while restricting the perturbation magnitude under an L_p norm. Goodfellow et al. (2014), Kurakin et al. (2016b), Kurakin et al. (2016a) and Madry et al. (2017) are based on gradient descent to optimize the original image to achieve the goal of attack. Carlini and Wagner (2017) proposes a method of generating adversarial samples under different norms L_0 , L_2 and L_∞ . Brown et al. (2017), Thys et al. (2019) and Liu et al. (2018) no longer disturb the whole image, but generate an adversarial patch and paste it on a specific position of the image to achieve the purpose of the attack.

3D adversarial attack. Recently, inspired by the 2D image attack method, many researches attempt to extend the optimization-based attack to 3D point cloud. Xiang et al. (2019b) proposes an adversarial attack to add clusters to the point cloud. Liu et al. (2019), Xiang et al. (2019a) and Zhang et al. (2019) all propose the point cloud perturbation attack approach, however, these approaches depend on updating the whole point cloud without a focus, and a plethora of points make rare contribution to attack while take considerable perturbation magnitude. In order to solve this problem, it is necessary to select appropriate salient points in the point cloud. Zheng et al. (2019) and Wicker and Kwiatkowska (2019) try to shift salient local region toward inside which is thought as point removal, instead of perturbation. In this work, we propose an effective and flexible method for local adversarial perturbation to solve this issue.

3D adversarial defense. As the counterpart of attack, defense technology is proposed to mitigate the harmful attack effect. L_p based attack is easy to create outlying points, which could be easily detected by point removal. Classical point cloud defense methods include the algorithm based on Simple Random Sampling (SRS) Zhou et al. (2019), which selects several points to remove in the perturbation point cloud with medium probability. Then, Statistical Outlier Removal (SOR) Rusu et al. (2008) is proposed by Rusu et al., which obtains better defense performance by removing outliers with *knn*. On this basis, DUP-Net Zhou et al. (2019) deletes out-of-surface points in the perturbation point cloud, and then uses interpolation method to repair the characteristics of the point cloud, so that the point cloud can be correctly classified. Moreover, IF-Defense Wu et al. (2020) generates an implicit function to obtain a surface recovery module to repair the disturbed point cloud.

3. Methodology

Optimization-based 3D adversarial attack could be referred as distributional attack, as shown in Eq 1.

$$\underset{\hat{P}}{\text{minimize}} \quad J(G(\theta, \hat{P}), y) \quad \text{s.t.} \quad \mathcal{D}(P, \hat{P}) < \lambda \quad (1)$$

where P and \hat{P} are the clean and adversarial examples, $G(\theta, P)$ is the victim DNN, J is the adversarial loss and \mathcal{D} is the perceptibility/dissimilar loss. $\mathcal{D}(P, \hat{P}) < \lambda$ is a regulation term to avoid overshifting where λ is perturbation budget. This equation could be reformulated as a penalty function in *Lagrange* way as Eq 2.

$$\mathcal{L}_{total} = J(G(\theta, \hat{P}), y) + \kappa_1 \mathcal{D}(P, \hat{P}) \quad (2)$$

Besides, we divide the point set into constant subset S_c and perturbation subset S_p where $P \cap \hat{P} = S_c$. Namely, all the following optimization would directly work on S_p .

3.1. Salient Point Selection

In order to find salient points in the point cloud, we take the gradient of cross entropy loss \mathcal{L}_{ce} in the victim model G as the instance-wise importance of points. According to gradient magnitude $\frac{\partial \mathcal{L}_{ce}}{\partial p_i}$ where $p_i \in P$, we select the top m vertexes and use group sampling (*knn*) of PointNet to further locate local point clusters where each cluster has n points in theory.

3.2. Score Loss and Contractive Loss

The primary goal of the point cloud optimization-based attack is making the victim model misclassify the examples. Inspired by the 2D adversarial methods, we first propose a basic loss to directly reduce the confidence of the victim model on the true label of the examples. As shown in Eq 3, we choose to use normalized estimation probability and ground truth as straightforward attack loss, where k enumerates the total classes, y is the ground truth for one-hot encoding and \hat{y} is the prediction probability.

$$\mathcal{L}_{base} = \sum_k y_k \hat{y}_k \quad (3)$$

Moreover, the final appearance and classification result of adversarial examples largely depends on instance-level preference, e.g., some examples of airplane may have a preference for being cars, but not all airplane instances have this preference. We rethink this instance-level preference as a prior knowledge and delve it in optimization. In multi-label optimization, the preferred wrong class would be engaged in competing with other non-preferred wrong classes, which may be detrimental to convergence and sequentially affect the final performance. Therefore, in Eq. 4, we introduce single-label optimization loss \mathcal{L}_{score} where the wrong class with highest confidence is selected as the explicit direction in overall optimization. Moreover, as shown in Eq. 5, to release the preferred wrong class from competition, we propose intra-class contractive loss \mathcal{L}_{cons} that not only encourages the confidence of preferred wrong class \hat{y}_{pre} to be larger, but also compresses the wealth of confidence occupied by non-preferred wrong classes $\hat{y}_{non-pre}$.

$$\mathcal{L}_{score} = 1 - \max_k (1 - y_k) \hat{y}_k \quad (4)$$

$$\begin{aligned} \mathcal{L}_{cons} = & \sum_k y_k \cdot \max\{\lambda - \hat{y}_{pre} + \hat{y}_k, 0\} \\ & + 1 - \frac{\hat{y}_{pre}}{\sum \hat{y}_{non-pre}} \end{aligned} \quad (5)$$

In early optimization, our score loss and contractive loss would model the instance preference in a static way and then quickly find and stabilize the optimization direction. Overall, the total attack loss J is composed by \mathcal{L}_{base} , \mathcal{L}_{score} and \mathcal{L}_{cons} , as shown in Eq 6 where κ_2 and κ_3 are the balance factors and set to 1 and 0.5 respectively.

$$J = \mathcal{L}_{base} + \kappa_2 \mathcal{L}_{score} + \kappa_3 \mathcal{L}_{cons} \quad (6)$$

3.3. Perceptibility Loss

There are several ways to measure the distance (perceptibility) of two point sets. We adopt three kinds of perceptibility loss, i.e., L_2 distance, Chamfer distance [Borgefors \(1988\)](#) and Hausdorff [Dubuisson and Jain \(1994\)](#) distance, as perceptibility loss \mathcal{D} . Given two point sets A and B , $\mathcal{D}_{chamfer}$, $\mathcal{D}_{Hausdorff}$ and \mathcal{D}_{L_2} could be represented as:

$$\begin{aligned} \mathcal{D}_{chamfer}(A, B) = & \frac{1}{|A|} \sum_{a \in A} \min_{b \in B} \|a - b\|_2 \\ & + \frac{1}{|B|} \sum_{b \in B} \min_{a \in A} \|a - b\|_2 \end{aligned} \quad (7)$$

$$\mathcal{D}_{Hausdorff}(A, B) = \max_{a \in A} \min_{b \in B} \|a - b\|_2 \quad (8)$$

$$\mathcal{D}_{L_2} = \sum_{i=1}^{|A|} \|A_i - B_i\|_2 \quad (9)$$

In our L3A, \mathcal{D}_{L_2} only calculate pair-wise points in S_p ($|S_p| \ll |P|$) for faster training speed. However, we find that this distance may lead to outliers due to the lack of information about the surrounding points. Chamfer and Hausdorff distances enumerate the whole point set to find the closest points which consumes $\mathcal{O}(|P| \times |S_p|)$. The difference is that the chamfer distance takes the average for all the nearest pairs but the Hausdorff distance takes the maximum for them. Instead of calculating the distance between the point pairs directly, Chamfer and Hausdorff distances are based on two complete point sets, thus the original shape of the point cloud can be retained as much as possible.

3.4. Perturbation-budget Withdrawing Algorithm

It's worth noting that the attack loss in initial few iterations is extremely large which impetuously push the points to shift. However, after a while, the perturbation budget is almost consumed and following shift of points is limited on the surface of hypersphere, as shown in Eq. 1. Therefore, we aim to force some points which contribute trivially in

Algorithm 1 Perturbation-budget Withdrawing Algorithm (PWA)

Input: Attack loss J , perceptibility loss \mathcal{D} , perturbation subset S_p .**Parameter:** Total iteration number N , activation threshold h to use PWA, selection threshold w for filtering unimportant points, perceptibility budget λ , spatial position $x_i \in \mathbb{R}^3$ of $p_i \in S_p$, anchor position x'_i of p_i , learning rate lr .**Output:** Regulated gradient $\frac{\partial J}{\partial p_i}$ of points p_i and adjusted point coordinate x_i .

```

1: while  $0.2N < iteration < N$  do
2:   if  $\mathcal{D} > \lambda - h$  then
3:     record  $x'_i$ 
4:     if  $|\frac{\partial J}{\partial p_i}| < w$  then
5:        $x_i = x_i - (x_i - x'_i) \cdot lr \cdot \gamma |\frac{\partial J}{\partial p_i}|$ 
6:        $\frac{\partial \mathcal{L}_{total}}{\partial p_i} = 0$ .
7:     else
8:       pass
9:     end if
10:     $iteration = iteration + 1$ 
11:  else
12:     $iteration = iteration + 1$ 
13:  end if
14: end while
15: return gradient  $\frac{\partial J}{\partial p}$ , spatial position  $x$ 

```

current iteration to move backward to its anchor position x'_i to release some free budget for the moving of important points. With respect to \mathcal{D} , the anchor position of a point is remembered at latest distance calculation. For example, as shown in Eq. 9, \mathcal{D}_{L_2} uses corresponding points in clean and adversarial examples, the anchor position of point a_i is consistently to be b_i . While in $\mathcal{D}_{chamfer}$, each anchor location in clean example is calculated by its equation. As shown in the Perturbation-budget Withdrawing Algorithm (PWA) of Algorithm 1, we first push trivial points back with a minor step and then erase the gradient of them to avoid generating outlying points which would be easily removed by outlier defense. The next gradient-erasing step would provide more freedom in current iteration. Furthermore, the released budget by the moving-back step would work in next iteration. Essentially, this strategy is greedy and we will discuss the tradeoff between greedy attack and imperceptibility in Section 4.3.

4. Experimental Results

4.1. Experimental Setup

Dataset and victim model. We use the ModelNet40 Wu et al. (2015a) to evaluate the performance of our method, including training, testing the victim models and generating the adversarial examples. This dataset consists of 12,311 CAD models with 40 common object categories. These CAD models are divided into 9,843 for training and 2,468 for testing.

Table 1: Attack success rate of L3A using different perturbation loss on PointNet, PointNet++ and DGCNN.

Victim Model	λ	\mathcal{D}_{L_2}			$\mathcal{D}_{Chamfer}$			$\mathcal{D}_{Hausdorff}$		
		None	SOR	SRS	None	SOR	SRS	None	SOR	SRS
PointNet	0.001	99.7%	92.6%	99.6%	99.7%	92.7%	99.7%	99.7%	91.5%	99.7%
	0.0025	99.7%	92.7%	99.6%	99.7%	92.7%	99.7%	99.7%	92.4%	99.7%
	0.005	99.7%	92.7%	99.6%	99.7%	92.7%	99.7%	99.7%	92.4%	99.7%
PointNet++	0.001	97.0%	94.0%	88.5%	97.2%	94.0%	89.2%	97.0%	92.9%	90.0%
	0.0025	96.9%	93.2%	90.5%	97.3%	92.4%	89.2%	97.4%	94.4%	89.4%
	0.005	96.8%	93.3%	89.5%	97.0%	93.4%	89.7%	97.3%	92.9%	89.3%
DGCNN	0.001	98.8%	97.0%	97.0%	99.0%	96.0%	97.4%	98.9%	95.7%	97.5%
	0.0025	98.9%	95.8%	95.8%	99.0%	96.0%	97.1%	98.9%	96.0%	97.6%
	0.005	99.0%	96.0%	96.0%	99.0%	96.0%	97.7%	98.9%	96.4%	97.8%

The test split is used for evaluating the attack performance. Besides, we use PointNet [Qi et al. \(2017a\)](#), PointNet++ [Qi et al. \(2017b\)](#) and DGCNN [Wang et al. \(2019\)](#) as our victim models which are trained in official code with default hyperparameters.

Implementation Details. The point cloud contains 10,000 points, and we sample 1,024 points for training. The total epochs is set to 200. We use the Adam optimizer to optimize the objectives where the learning rate lr is set to 0.001, and the momentum β_1 and β_2 are set to 0.9 and 0.999 respectively. We warm up the initial optimization within the first n epochs (about 0.2 times of the total epochs). The rest epochs would follow the proposed strategy. Warm-up aims to stabilize the instance-level preference, and specifically find the preferred class.

4.2. Attack and Defense

We report the performance of our method with various settings on PointNet [Qi et al. \(2017a\)](#), PointNet++ [Qi et al. \(2017b\)](#) and DGCNN [Wang et al. \(2019\)](#) in Table 1. Our method could achieve promising results on all three victim models under no defense, SOR and SRS: L3A obtains 99.7%/97.1%/98.9% success rate on three model respectively without using defense strategy in average. Besides, after the defense of SOR/SRS, our model only decreases 7.2%/0.1%, 3.7%/4.8% and 2.8%/1.8 on three models averagely and all results are all over 92%, indicating our L3A is robust to defense. Moreover, we adjust the magnitude of λ from 0.001 to 0.005 which is extremely small, but the performance only has minor fluctuation.

Moreover, Figure 1 illustrates the cost-effectiveness comparison with and without using our method (i.e., L3A and global optimization with \mathcal{D}_{L_2}). Intuitively, we believe points moving far away from anchor and occupying large budget would contribute more to attack. However, in Figure 1 (a), over 46% percent of total points concentrate on 0.1-0.3 range where each point moves trivially, but they cost about 33% budget and only contribute 14% for attack. Besides, in the displacement range of 0.6 to 1.0 where per-point displacement is rigid, their contribution also decline to 3%, which is contrary to our intuition. Namely, this Gaussian-like distributions of three lines mean global optimization doesn't allocate reasonable budget to each point from low to high intervals. In Figure 1 (b), we also evaluate

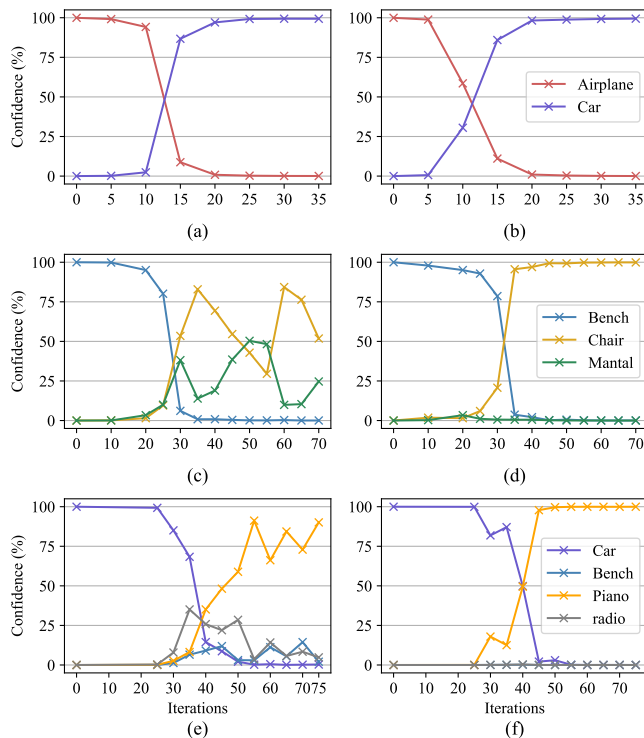


Figure 2: A comparison of optimization process of three examples (airplane, bench, car in three rows separately) with and without using \mathcal{L}_{score} and \mathcal{L}_{cons} . Subfigure (a), (c), (e) show three examples without using our method, the airplane converges to the best class, while the other two examples have poor convergence. Using our method, all examples could converge quickly to the best class as shown in Subfigure (b), (d), (f).

the cost-effectiveness of each points of our method. It shows that our method inhibits the cost of low range (0-0.3) and push over 78% points to locate in this range, which can avoid appreciable point transformation and keep original appearance as much as possible. Meanwhile, in the high range (0.6-1.0), the contribution of points in this range largely increases with stable cost, indicating the each further step the point made would effectively lead to mislead the victim model.

4.3. Single-label Attack and Withdrawing Algorithm

As shown in Figure 2, we report the convergence of optimizing airplane, bench, car toward adversarial examples by only using \mathcal{L}_{base} (left column) and using our \mathcal{L}_{score} and \mathcal{L}_{cons} (right column). Notably, our method could not only find the class-level updating direction toward specific class in Figure 2 (a & b), but also suppress all potential weak wrong classes when it has two strong competitive wrong classes or is harassed by considering multiple class

choice. Moreover, in the pair of Figure 2 (c & d) and (e & f), our method could faster the convergence of all examples and achieve higher confidence for preferred class.

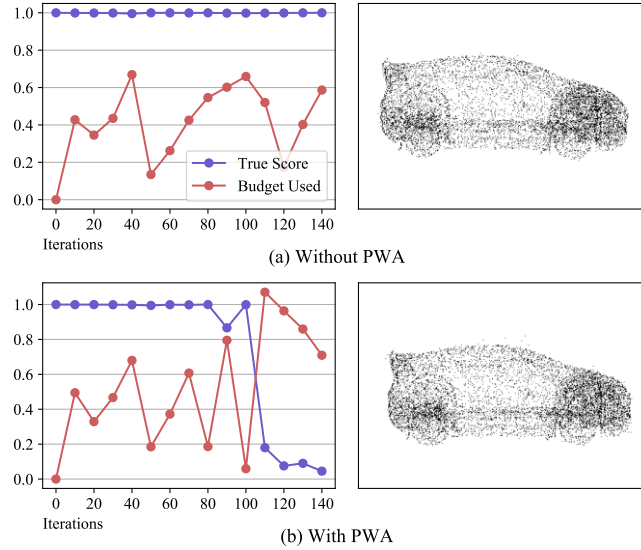


Figure 3: Attack results on car with and without using our PWA.

Table 2: Ablation study of our L3A on PointNet, PointNet++ and DGCNN using score loss, contractive loss, PWA and local salient point selection.

Group	Method				Victim Model		
	\mathcal{L}_{score}	\mathcal{L}_{cons}	PWA	Local	PN	PN++	DGCNN
1	×	×	×	×	98.3%	96.1%	100.0%
2	✓				98.8%	97.2%	100.0%
3		✓			100.0%	99.9%	100.0%
4			✓		98.3%	96.3%	100.0%
5				✓	97.2%	94.0%	98.0%
6	✓	✓			100.0%	99.9%	100.0%
7		✓	✓		100.0%	99.9%	100.0%
8			✓	✓	97.2%	94.5%	98.2%
9	✓	✓	✓		100%	100%	100%
10		✓	✓	✓	99.3%	99.9%	99.3%
11	✓		✓	✓	98.2%	96.7%	99.4%
12	✓	✓		✓	99.6%	97.1%	98.9%
13	✓	✓	✓	✓	99.3%	98.9%	99.9%

In addition, the adversarial loss of some examples may not further decrease, because these examples are regulated by perceptibility loss \mathcal{D} and trapped in local optimality. As

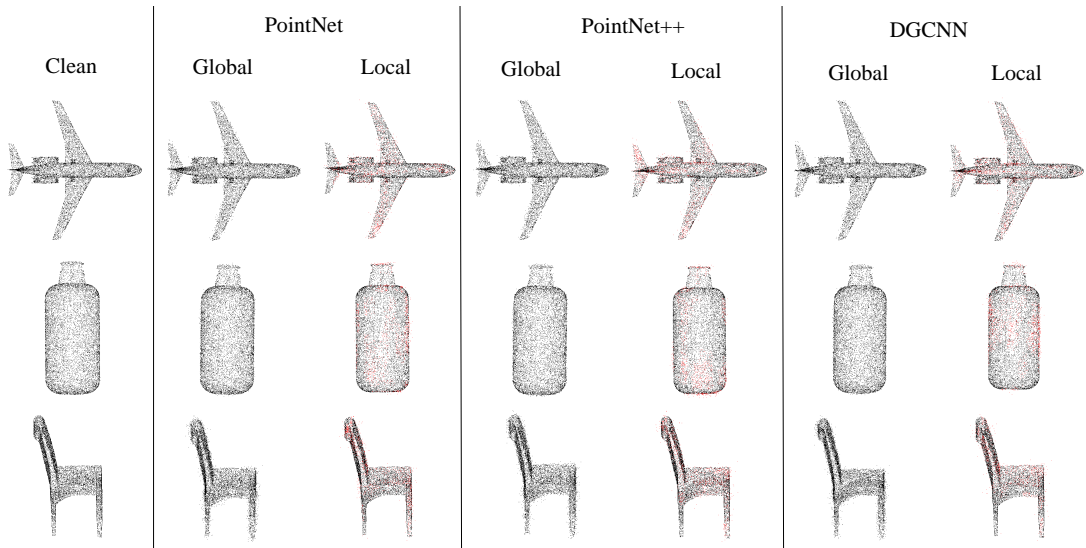


Figure 4: Visualization of adversarial examples of PointNet, PointNet++, and DGCNN models which are attacked by global optimization and our L3A. The points selected in the local point selection are marked red in the figure.

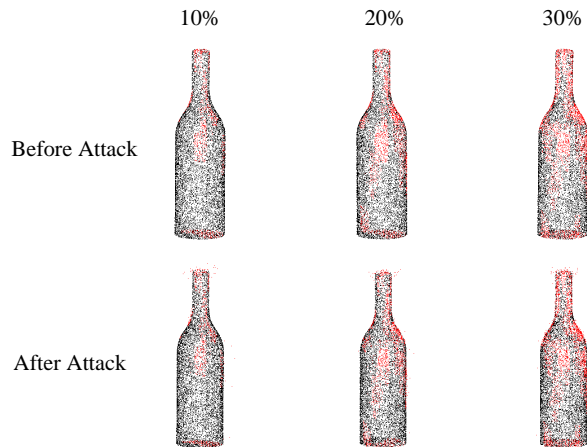


Figure 5: Visualization of different proportions of selected points. Salient points marked in red are selected at 10%, 20%, and 30% ratios in PointNet model.

shown in Figure 3 (a), the budget is consumed in early stage. At 50 iteration, the budget is released due to self-adjustment, but following optimization still fail to make the attack succeeded. Our PWA solves this issue by encouraging a part of unimportant points to

Table 3: The effect of point ratio on performance. m represents the number of selected vertexes, and n represents the number of point of cluster around the vertexes using the knn .

m	n	Proportion	PointNet	PointNet++	DGCNN
30	30	25.06%	95.5%	89.9%	95.2%
30	40	25.92%	95.5%	90.5%	95.1%
30	50	26.56%	95.5%	90.1%	94.9%
40	30	32.17%	96.2%	91.7%	97.3%
40	40	32.28%	96.2%	92.0%	97.4%
40	50	33.53%	96.3%	91.9%	97.0%
50	30	37.51%	96.8%	92.8%	98.4%
50	40	39.22%	96.8%	92.6%	98.4%
50	50	39.53%	96.8%	92.9%	98.5%

release budget to escape this local optimality. Figure 3 (b) shows that at 90 iteration, our PWA force points to release about 0.7 budget in a sudden, and pending the confidence could fall down to 0.08. Although our PWA algorithm is a greedy strategy, it does not heavily disturb perceptibility, since the PWA doesn't trigger continuously but is always in the withdraw-adjustment cycle. This gives all the points a chance to be adjusted to a better position in the next iteration. As shown in the right part of Figure 3, even after heavy budget withdrawing, the appearance of point cloud is still reasonable like a car.

4.4. Ablation Study

We evaluate each component of our method, and the results are shown in the Table 2. Compared the group 1 to 4, the proposed \mathcal{L}_{score} , \mathcal{L}_{cons} and PWA could bring 0.5%/1.1%, 1.7%/3.8% and 0/0.3% improvements on PointNet/PointNet++ respectively in global optimization case. However, when only using the local salient point selection, i.e., in group 5, the performance meets 1.1%/2.1%/2.0% decrease on three models. This means only using local points to optimization may not a good choice. In addition, we further assess \mathcal{L}_{score} , \mathcal{L}_{cons} and PWA in local optimization case. As shown in group 5 and 8, solely using PWA would bring 0.5%/0.2% gains on PointNet++ and DGCNN, and similarly, the fully armed setting could achieve further improvement. It is worthy to note that the best performance of this table uses variants with global optimization settings, i.e., group 6 and 7, instead of our fully armed setting. However, our method is only 0.7%/1.0%/0.1% lower than them and most points don't move after attack, which is more similar to original examples in visual appearance. We also visualize the point cloud in Figure 4 where selected points using our method are marked in red. Our L3A keeps most points motionless and drives the pending points to move surrounding the original outlook. Therefore, our method could largely attack the victim model while transforming the point cloud indistinguishably.

We also elaborate on the influence of selection point ratios to the ultimate results. In Table 3, we select different combinations of cluster and sample number m and n . It is obvious that with the increase of proportion, the performance of our method on three victim

Table 4: Comparison between our method and some global and local attack methods on PointNet, PointNet++ and DGCNN.

	ours	global					local		
	L3A	RP	GP	Chamfer	EoTPG	Xiang et al.	Zheng et al.	Kim et al.	Nudge
Pointnet	99.3%	12.4%	89.8%	100.0%	40.5%	100%	55.7%	89.4%	97.4%
Pointnet++	98.9%	9.4%	98.8%	97.7%	87.8%	-	41.5%	88.8%	-
DGCNN	99.9%	11.6%	34.5%	99.6%	23.3%	-	35.8%	62.2%	69.5%

models accordingly rise. Moreover, even using similar but reverse setting, such as $n = 30$ & $m = 40$ and $n = 40$ & $m = 30$, the proportion of them is different, since most cluster would concentrate on salient region and therefore the overlapped points would increase with cluster size. As shown in Figure 5, we visualize the selected points in bottle. Most of selected points locate in the mouth, upper part and bottom where are the salient region.

4.5. Comparison of Other Methods

As shown in Table 4, we first compare the performance of our L3A with existing global-optimization-based attack methods, including random perturbation (RP), gradient projection attack (GP) Liu et al. (2020), Chamfer attack Liu et al. (2020), EoTPG Zhang et al. (2019) and Xiang et al. Xiang et al. (2019b) on PointNet, PointNet++ and DGCNN, and then compare our method with some methods based on local optimization including Zheng et al. Zheng et al. (2019), Kim et al. Kim et al. (2020) and Nudge Zhao et al. (2020). In the global optimization method, our L3A achieves the best performance on PointNet++ and DGCNN and is just 0.7% lower than Chamfer and Xiang et al. on PointNet. Moreover, it outperforms Chamfer and EoTPG with 1.2% and 11.1% gain on PointNet++ and with 0.3% and 76.6% gain on DGCNN respectively. In the local optimization method, our L3A achieves the best performance on PointNet, PointNet++ and DGCNN. Moreover, Kim et.al. and Nudge are concurrent local attack work as our method, but our L3A is more cost-effective and achieves much better attack performance than them.

5. Conclusion

In this paper, we analyze existing weaknesses in global-optimization-based attack and further propose a novel local aggressive adversarial attack to solve them. To encourage the cost-effectiveness of points between attack and perceptibility, a salient subset of point cloud is selected to optimize toward the direction of preferred wrong class. Besides, to avoid falling in local optimality, PWA could adjust the attack and perceptibility purpose of each point by releasing reasonable budget from insignificant points. Extensive experiments demonstrate the state-of-the-art performance of our method on PointNet, PointNet++ and DGCNN.

Acknowledge

This work was supported by the National Natural Science Foundation of China (Nos. 62076139, 61702280), the National Postdoctoral Program for Innovative Talents (No. BX20180146),

China Postdoctoral Science Foundation (No. 2019M661901), Jiangsu Planned Projects for Postdoctoral Research Funds (No. 2019K024), Zhejiang Lab (No. 2021KF0AB05), 1311 Talent Program of Nanjing University of Posts and Telecommunications, and Natural Science Foundation of Jiangsu Province (No. BK20170900).

References

- Anish Athalye, Logan Engstrom, Andrew Ilyas, and Kevin Kwok. Synthesizing robust adversarial examples. In *ICML*, pages 284–293, 2018.
- Jorge Beltrán, Carlos Guindel, Francisco Miguel Moreno, Daniel Cruzado, Fernando Garcia, and Arturo De La Escalera. Birdnet: a 3d object detection framework from lidar information. In *ICITS*, pages 3517–3523, 2018.
- Battista Biggio, Iginio Corona, Davide Maiorca, Blaine Nelson, Nedim Šrndić, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. Evasion attacks against machine learning at test time. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 387–402, 2013.
- Gunilla Borgefors. Hierarchical chamfer matching: A parametric edge matching algorithm. *TPAMI*, 10(6):849–865, 1988.
- Tom B Brown, Dandelion Mané, Aurko Roy, Martín Abadi, and Justin Gilmer. Adversarial patch. *arXiv preprint arXiv:1712.09665*, 2017.
- Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *IEEE Symposium on Security and Privacy*, pages 39–57, 2017.
- A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, and H. Su. Shapenet: An information-rich 3d model repository. *Computer Science*, 2015.
- M-P Dubuisson and Anil K Jain. A modified hausdorff distance for object matching. In *ICCV*, pages 566–568, 1994.
- Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- Jaeyeon Kim, Binh-Son Hua, Duc Thanh Nguyen, and Sai-Kit Yeung. Minimal adversarial examples for deep learning on 3d point clouds. *arXiv preprint arXiv:2008.12066*, 2020.
- Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236*, 2016a.
- Alexey Kurakin, Ian Goodfellow, Samy Bengio, et al. Adversarial examples in the physical world, 2016b.

- Bo Li, Tianlei Zhang, and Tian Xia. Vehicle detection from 3d lidar using fully convolutional network. *arXiv preprint arXiv:1608.07916*, 2016.
- Ming Liang, Bin Yang, Yun Chen, Rui Hu, and Raquel Urtasun. Multi-task multi-sensor fusion for 3d object detection. In *CVPR*, pages 7345–7353, 2019.
- Daniel Liu, Ronald Yu, and Hao Su. Extending adversarial attacks and defenses to deep 3d point cloud classifiers. In *ICIP*, pages 2279–2283, 2019.
- Daniel Liu, Ronald Yu, and Hao Su. Adversarial shape perturbations on 3d point clouds. In *ECCV*, pages 88–104. Springer, 2020.
- Xin Liu, Huanrui Yang, Ziwei Liu, Linghao Song, Hai Li, and Yiran Chen. Dpatch: An adversarial patch attack on object detectors. *arXiv preprint arXiv:1806.02299*, 2018.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- Priyanka Mandikal, KL Navaneet, Mayank Agarwal, and R Venkatesh Babu. 3d-lmnet: Latent embedding matching for accurate and diverse 3d point cloud reconstruction from a single image. *arXiv preprint arXiv:1807.07796*, 2018.
- Charles R Qi, Hao Su, Matthias Nießner, Angela Dai, Mengyuan Yan, and Leonidas J Guibas. Volumetric and multi-view cnns for object classification on 3d data. In *CVPR*, pages 5648–5656, 2016.
- Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, pages 652–660, 2017a.
- Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017b.
- Radu Bogdan Rusu, Zoltan Csaba Marton, Nico Blodow, Mihai Dolha, and Michael Beetz. Towards 3d point cloud based object maps for household environments. *Robotics and Autonomous Systems*, 56(11):927–941, 2008.
- Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointcnn: 3d object proposal generation and detection from point cloud. In *CVPR*, pages 770–779, 2019.
- Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *ICCV*, pages 945–953, 2015.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- Simen Thys, Wiebe Van Ranst, and Toon Goedemé. Fooling automated surveillance cameras: adversarial patches to attack person detection. In *CVPRW*, pages 1–6, 2019.

- 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 TOG*, 38(5):1–12, 2019.
- Matthew Wicker and Marta Kwiatkowska. Robustness of 3d deep learning in an adversarial setting. In *CVPR*, pages 11767–11775, 2019.
- Jiajun Wu, Chengkai Zhang, Tianfan Xue, William T Freeman, and Joshua B Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. *arXiv preprint arXiv:1610.07584*, 2016.
- Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*, pages 1912–1920, 2015a.
- Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*, pages 1912–1920, 2015b.
- 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.
- Chong Xiang, Charles R Qi, and Bo Li. Generating 3d adversarial point clouds. In *CVPR*, pages 9136–9144, 2019a.
- Chong Xiang, Charles R Qi, and Bo Li. Generating 3d adversarial point clouds. In *CVPR*, pages 9136–9144, 2019b.
- Bin Yang, Wenjie Luo, and Raquel Urtasun. Pixor: Real-time 3d object detection from point clouds. In *CVPR*, pages 7652–7660, 2018.
- Qiang Zhang, Jiancheng Yang, Rongyao Fang, Bingbing Ni, Jinxian Liu, and Qi Tian. Adversarial attack and defense on point sets. *arXiv preprint arXiv:1902.10899*, 2019.
- Yiren Zhao, Iliia Shumailov, Robert Mullins, and Ross Anderson. Nudge attacks on point-cloud dnns. *arXiv preprint arXiv:2011.11637*, 2020.
- Tianhang Zheng, Changyou Chen, Junsong Yuan, Bo Li, and Kui Ren. Pointcloud saliency maps. In *ICCV*, pages 1598–1606, 2019.
- Hang Zhou, Kejiang Chen, Weiming Zhang, Han Fang, Wenbo Zhou, and Nenghai Yu. Dup-net: Denoiser and upsampler network for 3d adversarial point clouds defense. In *ICCV*, pages 1961–1970, 2019.
- Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *CVPR*, pages 4490–4499, 2018.