

Article

Less Adaptation, More Transfer: Spectral View Randomization for 3D Point Cloud Transfer Attacks

Yang Gao, Jingyi Liu, Hongjia Liu, Haoran Li and Jian Xu *

Software College, Northeastern University, Shenyang 110169, China; gaoy@mail.neu.edu.cn (Y.G.); liujingyi@stumail.neu.edu.cn (J.L.); liuhongjia@stumail.neu.edu.cn (H.L.); lihaoran@stumail.neu.edu.cn (H.L.)

* Correspondence: xuj@mail.neu.edu.cn

Abstract

Point cloud perception is important in autonomous driving, robotics, and other security-critical 3D systems, yet learned point cloud classifiers remain vulnerable to transferable adversarial perturbations. A central difficulty in transfer-based black-box attacks is surrogate overfitting: an update that is highly effective on an accessible source model may not generalize to an unknown target architecture. We introduce SpecEOT, a source-agnostic and graph-spectral expectation-over-transformation attack. A fixed graph Fourier transform (GFT) basis is constructed from each clean point cloud. At every optimization iteration, each non-identity view independently samples a frequency band and a perturbation sign from uniform distributions; the resulting view gradients are averaged with equal weights and used to update the adversarial point cloud through projected Adam ascent. We evaluate the stochastic method over repeated seeds, extend the ablation to two source architectures, and analyze the interaction between band count and randomization strength while reporting computational cost and assessing robustness to Gaussian jitter and point dropout. SpecEOT achieves strong transferability on ModelNet40 and ShapeNet.

Keywords: 3D point cloud adversarial attacks; transferable adversarial examples; graph Fourier transform; expectation over transformation

1. Introduction

Point clouds are widely used as compact 3D representations in autonomous driving, robotics, augmented reality, and other security- and safety-critical systems [1]. Deep point cloud classifiers achieve strong accuracy on standard benchmarks [2,3] but remain vulnerable to adversarial perturbations [4]. In this work, an attacker is an entity with access to a labeled point cloud and full white-box access to an accessible source classifier, but without knowledge of or query access to the deployed target classifier. Examples include a malicious supplier, an adversary controlling a compromised perception pipeline, or a security evaluator using a locally trained surrogate. The attacker crafts an untargeted adversarial point cloud against the source model and transfers the same point cloud directly to the target. This transfer-based black-box setting is widely used to study attacks when the deployed architecture, parameters, and gradients are unavailable [5–7]. Thus, the security question is not just whether the source classifier is fooled but whether the perturbation remains effective across heterogeneous target architectures.

A central difficulty in transfer-based attacks is surrogate overfitting [8]. Gradients that are highly effective on the source model may exploit source-specific decision boundaries and generalize poorly to target models. Prior point cloud attacks address transferability



Academic Editor: Christos Bouras

Received: 25 June 2026

Revised: 21 July 2026

Accepted: 22 July 2026

Published: 24 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

through autoencoder-assisted objectives, perturbation factorization, input transformations, or compact shared representations [8–13]. In parallel, graph-spectral analysis provides a geometry-aware coordinate system for irregular point sets [14–16]. GFT-based attacks such as SAAO generate structured spectral transformations rather than unstructured coordinate noise [17]. These two lines of work motivate an explicit connection to expectation over transformation (EOT) [18–20]: instead of adapting an attack to one selected spectral transformation, it is possible to optimize the expected source loss over a distribution of spectral views.

When multiple spectral views are available, an important question remains: should their contributions be selected or weighted using signals from the source model? Views with larger source losses, stronger source gradient alignment, or apparently reliable source spectral responses may seem preferable; however, all of these criteria are measured on the surrogate. Therefore, we treat the possibility that source-conditioned weighting can re-amplify surrogate-specific directions as an empirical hypothesis rather than a theoretical guarantee. SpecEOT evaluates the complementary design: it uniformly samples structured views and aggregates them with equal weights, so no source-preferred view receives special treatment.

Figure 1 summarizes the motivating comparison among a vanilla attack, equally weighted graph-spectral EOT, and source-conditioned adaptive weighting. We additionally measure the source–target gradient gap, defined as the alignment between the realized update and the source gradient minus its mean alignment with the target-model gradients. A smaller value indicates a less source-dominated update. Because target gradients are unavailable to the attacker, this quantity is used only after attack generation as a diagnostic.




Figure 1. Motivating observation for SpecEOT. (A) Source-conditioned selection may re-amplify surrogate-specific directions, whereas uniform spectral EOT averages gradients over structured graph-spectral views. (B) Target-wise transfer ASR from PointNet to DGCNN, PCT, and PointNet++. SpecEOT improves transferability across target architectures. (C) Source–target gradient gap. SpecEOT yields a smaller gap, indicating a less source-dominated update. Target gradients are used only for post hoc diagnosis, never for attack generation.

Motivated by this observation, we propose *SpecEOT*, a source-agnostic graph-spectral EOT framework for transferable 3D point cloud attacks. Given a clean point cloud, SpecEOT constructs a fixed spectral basis from its κ -nearest-neighbor graph. During each attack iteration, SpecEOT independently and uniformly samples one spectral band and one sign for every non-identity view, multiplicatively perturbs the selected Fourier coefficients, reconstructs the view, and projects it into the same ℓ_∞ ball as the adversarial point cloud. The view gradients are averaged with equal weights, normalized per sample, and used in projected Adam ascent.

SpecEOT combines three established ingredients in a deliberately minimal way: GFT supplies structured frequency-localized transformations; EOT supplies a Monte Carlo gradient estimator over those transformations; and projected gradient optimization enforces

the perturbation budget [17,18,21]. The contribution is not a new transform or optimizer in isolation, but the finding that independent uniform spectral sampling and equal aggregation can be preferable to source-conditioned spectral adaptation for transfer. We support this claim empirically and avoid interpreting the diagnostic evidence as a universal theoretical guarantee.

Our contributions are summarized as follows:

- We identify a security-relevant failure mode in transfer-based 3D point cloud attacks: source-conditioned adaptive view weighting can reinforce surrogate-specific update directions and hurt black-box transferability.
- We propose SpecEOT, a minimal graph-spectral EOT attack using a fixed GFT basis, independently sampled uniform band/sign views, equal gradient aggregation, and projected Adam ascent without CD/HD optimization losses.
- We conduct extensive evaluations on ModelNet40 and ShapeNet across PointNet, PointNet++, DGCNN, and PCT, showing that SpecEOT improves black-box transferability over competitive point cloud attack baselines.

2. Related Work

The introduction positions SpecEOT at the intersection of transfer attacks, graph-spectral point cloud transformations, and EOT. This section provides the technical background for each component and clarifies which elements are inherited from prior work and which design choice is evaluated here.

2.1. Adversarial Attacks on 3D Point Clouds

Adversarial vulnerability has been extensively studied in image classification, where small input perturbations can induce incorrect predictions [22–24]. Similar risks arise in point cloud perception. Because point clouds are unordered, irregular, and geometry-dependent, attacks include point perturbation, point addition or deletion, geometry-aware optimization, and generative approaches [2,11,25]. Early methods such as 3D-ADV typically shift existing points or add adversarial points [26], while geometry-aware attacks introduce local smoothness, CD, HD, or nearest-neighbor regularization to control distortion [9,27]. SpecEOT differs in that CD and HD are evaluation metrics rather than optimization terms, and its hard feasibility constraint is coordinate-wise ℓ_∞ projection.

2.2. Transfer-Based Black-Box Attacks

In practical deployments, an attacker may be unable to access target parameters or gradients. Transfer-based black-box attacks instead optimize on an accessible surrogate and evaluate the same adversarial example on unseen targets without queries. Their central difficulty is surrogate overfitting: a perturbation can exploit a source-specific boundary and transfer poorly. In the point cloud domain, AdvPC combines classification and autoencoder objectives [9]; PF-Attack factorizes perturbations during optimization [8]; ANF learns transferable perturbations with a normalizing flow [11]; and recent MAT-ADV and COSA studies investigate transformation- or representation-based transfer mechanisms [12,13]. SpecEOT shares the goal of suppressing source-specific adaptation but uses a fixed model-independent distribution of graph-spectral views rather than source-conditioned selection or target queries.

2.3. Graph Spectral Transformations for Point Cloud Attacks

Graph spectral analysis provides a geometry-aware coordinate system for irregular point sets. A graph Fourier transform (GFT) decomposes point coordinates into modes ordered by the spectrum of a graph Laplacian [14,16]. SAAO demonstrates that graph-spectral transformations can be used in a transferable point cloud attack [17].

The proposed SpecEOT inherits the GFT representation but changes the optimization rule: it draws bands and signs independently from a fixed uniform distribution, and estimates the corresponding EOT gradient with equal weights [18].

3. Methodology

Building on the preceding literature, we now define the threat objective, fixed graph basis, randomized transformation distribution, EOT gradient estimator, and projected optimizer used by SpecEOT.

3.1. Overview

The overall pipeline of SpecEOT is illustrated in Figure 2. Given a clean point cloud, SpecEOT constructs a fixed graph Fourier basis from its κ -nearest-neighbor graph and partitions the ordered spectrum into equal-count bands. At every attack iteration and for every non-identity view, a band and a sign are sampled independently from uniform distributions. Multiplicative spectral views are reconstructed, projected into the attack budget, and passed to the source model together with the exact identity view. Equal aggregation estimates the spectral EOT gradient, which is normalized per sample and used in projected Adam ascent. Neither target-model information nor a geometry loss is used during attack generation.




Figure 2. Overview of SpecEOT. A fixed GFT basis is constructed from the clean point cloud, randomized band/sign views are sampled independently at each iteration, and the identity and randomized view gradients are aggregated equally before projected Adam ascent.

3.2. Problem Formulation

Let $P \in \mathbb{R}^{N \times 3}$ denote a clean point cloud with ground-truth label y , where N is the number of input points. Given a white-box source classifier f_s , the goal of an untargeted transfer-based attack is to generate an adversarial point cloud P^{adv} that fools the source model and transfers to unknown target models. The adversarial point cloud is constrained by an ℓ_∞ perturbation budget:

$$P^{\text{adv}} = P + \delta, \quad \|\delta\|_\infty \leq \epsilon. \quad (1)$$

Equivalently, P^{adv} is restricted to the following feasible set:

$$\mathcal{B}_\epsilon(P) = \{X \in \mathbb{R}^{N \times 3} \mid \|X - P\|_\infty \leq \epsilon\}. \quad (2)$$

The projection onto this feasible set is implemented by coordinate-wise clipping:

$$\Pi_{\mathcal{B}_\epsilon(P)}(X) = P + \text{clip}(X - P, -\epsilon, \epsilon). \quad (3)$$

A direct source-model attack maximizes the cross-entropy loss $\mathcal{L}_{\text{CE}}(f_s(X), y)$ with respect to the input point coordinates. However, the resulting perturbation may overfit to the local decision boundary of the source classifier and transfer poorly to target architectures. To reduce such source-specific overfitting, we propose *SpecEOT*, a graph-spectral expectation-over-transformation framework that constructs a source-agnostic distribution of randomized graph-spectral views and optimizes the attack loss averaged over these views.

3.3. Fixed Graph Spectral Basis

For each clean point cloud P , we construct a κ -nearest-neighbor graph over its Euclidean XYZ coordinates, with $\kappa = 16$. The adjacency matrix is binary, and is symmetrized by taking the union of directed neighbor edges. Let $A \in \mathbb{R}^{N \times N}$ denote the resulting symmetric adjacency matrix and D the degree matrix. We use the unnormalized graph Laplacian $L = D - A$. Its eigendecomposition is $L = U\Lambda U^\top$, where $U = [u_1, \dots, u_N] \in \mathbb{R}^{N \times N}$ contains the graph Fourier basis vectors and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_N)$ contains the ordered graph frequencies, with $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N$.

SpecEOT computes the graph-spectral basis only once from the clean point cloud and keeps it fixed throughout the attack. This provides a stable geometry-aware coordinate system for all adversarial updates and spectral views. For any point cloud $X \in \mathbb{R}^{N \times 3}$ represented in this basis, the graph Fourier transform and its inverse are expressed as $\Phi(X) = U^\top X$ and $X = U\Phi(X)$, respectively.

Keeping the basis fixed avoids repeatedly changing the spectral coordinate system during optimization and ensures that all random spectral views are generated from the same geometry-induced reference frame.

3.4. Equal-Count Spectral Band Partition

We divide the graph spectrum into B disjoint frequency bands. Splitting the eigenvalue range uniformly can produce highly imbalanced bands when the spectrum is non-uniformly distributed. Therefore, SpecEOT uses an equal-count partition over the ordered spectral indices. Specifically, let $q_b = \lfloor bN/B \rfloor$ for $b = 0, \dots, B$. The b -th spectral band is defined as

$$\mathcal{I}_b = \{q_{b-1} + 1, \dots, q_b\}, \quad b = 1, \dots, B. \quad (4)$$

The corresponding binary band mask is denoted by $M_b \in \{0, 1\}^{N \times 1}$, where

$$[M_b]_i = \begin{cases} 1, & i \in \mathcal{I}_b, \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

When applied to spectral coefficients in $\mathbb{R}^{N \times 3}$, M_b is broadcast to all coordinate channels. This equal-count design ensures that each band contributes a comparable number of graph Fourier modes and avoids degenerate or nearly empty spectral bands.

3.5. Band-Randomized Spectral View Generation

At iteration t , let P_t^{adv} denote the current adversarial point cloud, initialized as $P_0^{\text{adv}} = P$. SpecEOT generates K total views. The first is an exact identity view:

$$P_t^{(0)} = P_t^{\text{adv}}. \quad (6)$$

The remaining $K - 1$ views are randomized graph-spectral transformations.

At every iteration, we independently sample a frequency-band index and a perturbation sign for each randomized view $k = 1, \dots, K - 1$:

$$b_{t,k} \sim \text{Unif}\{1, \dots, B\}, \quad s_{t,k} \sim \text{Unif}\{-1, +1\}. \quad (7)$$

Given the spectral coefficients of the current adversarial point cloud, $\Phi_t = U^\top P_t^{\text{adv}}$, we then perturb only the sampled spectral band:

$$\tilde{\Phi}_t^{(k)} = \Phi_t \odot (\mathbf{1} + s_{t,k} \rho M_{b_{t,k}}), \quad (8)$$

where ρ controls the strength of spectral randomization and $\odot$ denotes element-wise multiplication.

The corresponding spectral view is reconstructed using the inverse graph Fourier transform and projected back into the feasible perturbation region:

$$P_t^{(k)} = \Pi_{B_\epsilon(P)}(U \tilde{\Phi}_t^{(k)}), \quad k = 1, \dots, K - 1. \quad (9)$$

Thus, every spectral view satisfies the same ℓ_∞ constraint as the final adversarial point cloud.

The transformation in Equation (8) defines a structured graph-spectral view distribution. A fresh pair $(b_{t,k}, s_{t,k})$ is drawn for every non-identity view at every iteration; no deterministic cycling is used in the full method. In minibatched execution, the draw for a given (t, k) is shared across batch elements, while each point cloud applies it in its own fixed basis.

3.6. Spectral EOT Objective and Gradient Aggregation

SpecEOT optimizes the source-model classification loss under the structured graph-spectral view distribution. At iteration t , the EOT objective is estimated using the K generated views:

$$\mathcal{L}_{\text{EOT}}(P_t^{\text{adv}}, y) = \frac{1}{K} \sum_{k=0}^{K-1} \mathcal{L}_{\text{CE}}(f_s(P_t^{(k)}), y). \quad (10)$$

The implementation optimizes this classification objective alone. Let

$$G_t = \nabla_{P_t^{\text{adv}}} \mathcal{L}_{\text{EOT}} \quad (11)$$

denote the equal-weight Monte Carlo gradient estimate. For independently sampled views with finite per-view gradient variance, the variance of their sample mean decreases proportionally to $1/K$. This is the standard Monte Carlo effect underlying EOT; it does not guarantee that transfer ASR will increase monotonically with K . The identity view, correlations induced by a shared input and basis, finite optimization, and transformation strength all affect the realized attack.

We normalize G_t independently for each point cloud before Adam ascent. The two pre-specified candidates were unit- ℓ_2 and mean-absolute normalization. Based on development set performance, we selected unit- ℓ_2 normalization:

$$\hat{G}_t^{(n)} = \frac{G_t^{(n)}}{\|G_t^{(n)}\|_2 + \epsilon_g}, \quad \epsilon_g = 10^{-12}, \quad (12)$$

where the norm spans all $3N$ coordinates.

3.7. Projected Optimization

Starting from $P_0^{\text{adv}} = P$, SpecEOT performs projected Adam ascent. Let

$$\Delta_t = \text{AdamAscend}(\hat{G}_t; \alpha) \quad (13)$$

denote the Adam ascent step computed from the normalized gradient $\hat{G}_t$, with learning rate α . The adversarial point cloud is updated by

$$P_{t+1}^{\text{adv}} = \Pi_{\mathcal{B}_\epsilon(P)}(P_t^{\text{adv}} + \Delta_t). \quad (14)$$

After T iterations, P_T^{adv} is returned. Coordinate-wise projection is applied to every spectral view and after every optimizer update, so both intermediate views and the final example satisfy $\|P - P^{\text{adv}}\|_\infty \leq \epsilon$.

4. Experiments

We first define the threat model, then report results on transferability, ablation, robustness, computational efficiency, and algorithmic diagnostics.

4.1. Experimental Settings

4.1.1. Threat Model and Attack Protocol

We adopt an untargeted transfer-based black-box threat model. The attacker has access to the clean point cloud, its ground-truth label, and a single source classifier. The target architectures, parameters, logits, predictions, and gradients remain inaccessible during attack generation, and the attacker issues no queries to any target model. An adversarial point cloud is optimized solely against the source classifier and evaluated without modification on each target model.

Our canonical protocol sets $\epsilon = 0.10$, following [13], and $\kappa = 16$. It employs an unnormalized graph Laplacian, equal-count frequency bands, multiplicative spectral views, a single exact identity view, equal gradient aggregation, and Adam ascent. At each iteration, a frequency band and a perturbation sign are sampled independently for every non-identity view.

The implementation uses Python 3.9 and the PyTorch 2.4.0 ecosystem, and all experiments were conducted on a Ubuntu server equipped with two NVIDIA L20 GPUs.

4.1.2. Datasets and Models

We evaluate SpecEOT on the ModelNet40 and ShapeNet datasets. ModelNet40 comprises 12,311 CAD models distributed across 40 object categories, split into 9843 training samples and 2468 test samples. ShapeNet contains 16,881 shapes spanning 16 categories. Following standard practice, each object is represented by 1024 uniformly sampled points, with only the 3D spatial coordinates provided as input.

We select four representative point cloud classifiers: PointNet [28], PointNet++ [29], DGCNN [30], and PCT [31]. For each transfer scenario, one classifier acts as the source model while the remaining three serve as unseen target models.

4.1.3. Baselines and Metrics

We compare SpecEOT against five established point-cloud attacks: 3D-ADV [26], KNN attack [27], AdvPC [9], PF-Attack [8], and SAAO [17]. These baselines encompass various strategies, including coordinate perturbation, geometry-aware regularization, autoencoder-assisted transfer, perturbation factorization, and graph-spectral designs. Because the official

implementation of SAAO is unavailable, we re-implement it using the same source and target protocols and perturbation budget.

We evaluate performance using the attack success rate (ASR), defined as the percentage of adversarial examples that induce misclassification. White-box ASR is measured on the source model, whereas black-box ASR is measured on the unseen target models. Unless otherwise stated, results produced by our implementations are averaged over at least three runs with different random seeds.

4.2. Transferability Analysis

We evaluate transferability on ModelNet40 and ShapeNet under the ℓ_∞ perturbation budget $\epsilon = 0.10$. Tables 1 and 2 report the transfer ASR matrices. Across both datasets, SpecEOT achieves competitive or superior black-box transferability in most source–target pairs. The ShapeNet results show a similar overall trend, with SpecEOT attaining the highest black-box ASR in several challenging cross-architecture settings. These findings indicate that uniform graph-spectral randomization can improve transferability across datasets and model families. Because SpecEOT uses no target-model information during attack generation, these gains cannot be attributed to target-specific optimization. Performance nevertheless varies across architectures, highlighting the difficulty of achieving uniformly strong transferability in 3D adversarial attacks.

Table 1. Transferability on ModelNet40 under the ℓ_∞ budget $\epsilon = 0.10$. Underlined values indicate white-box attack success rates on the source model. Bold values highlight the highest black-box transferability among the evaluated methods.

Source Model	Method	PointNet	PointNet++	DGCNN	PCT
PointNet	3D-Adv	<u>100.00 ± 0.00</u>	1.25 ± 0.32	1.38 ± 0.28	1.36 ± 0.41
	KNN	<u>90.15 ± 1.82</u>	5.35 ± 0.74	6.54 ± 0.85	5.62 ± 0.63
	AdvPC	<u>99.75 ± 0.12</u>	8.52 ± 1.05	5.41 ± 0.77	8.95 ± 1.12
	PF-Attack	<u>84.12 ± 2.45</u>	26.85 ± 2.11	21.35 ± 1.87	39.45 ± 3.05
	SAAO	<u>99.62 ± 0.35</u>	47.15 ± 2.94	42.85 ± 2.76	36.88 ± 2.48
	SpecEOT	<u>94.56 ± 2.33</u>	51.34 ± 3.12	65.46 ± 1.80	45.00 ± 5.41
PointNet++	3D-Adv	2.85 ± 0.54	<u>100.00 ± 0.00</u>	1.75 ± 0.35	2.24 ± 0.42
	KNN	9.42 ± 1.15	<u>100.00 ± 0.00</u>	11.65 ± 1.25	25.15 ± 2.01
	AdvPC	10.35 ± 1.22	<u>94.15 ± 1.45</u>	9.45 ± 1.08	23.85 ± 1.95
	PF-Attack	38.15 ± 2.85	<u>93.85 ± 1.38</u>	33.15 ± 2.64	49.85 ± 3.14
	SAAO	48.85 ± 3.10	<u>99.35 ± 0.45</u>	52.15 ± 3.25	51.65 ± 3.42
	SpecEOT	43.33 ± 1.80	<u>100.00 ± 0.00</u>	54.17 ± 4.77	49.79 ± 1.80
DGCNN	3D-Adv	1.45 ± 0.25	1.85 ± 0.31	<u>100.00 ± 0.00</u>	1.05 ± 0.18
	KNN	10.65 ± 1.12	9.45 ± 1.05	<u>100.00 ± 0.00</u>	13.52 ± 1.42
	AdvPC	14.85 ± 1.35	21.65 ± 1.75	<u>92.55 ± 1.54</u>	22.45 ± 1.88
	PF-Attack	36.15 ± 2.45	29.25 ± 2.15	<u>85.35 ± 2.30</u>	40.85 ± 2.75
	SAAO	41.55 ± 2.85	49.35 ± 3.12	<u>100.00 ± 0.00</u>	41.45 ± 2.80
	SpecEOT	43.96 ± 2.51	44.58 ± 2.51	<u>100.00 ± 0.00</u>	50.54 ± 1.33
PCT	3D-Adv	1.41 ± 0.22	1.65 ± 0.35	2.45 ± 0.41	<u>100.00 ± 0.00</u>
	KNN	3.75 ± 0.65	3.55 ± 0.58	4.25 ± 0.72	<u>100.00 ± 0.00</u>
	AdvPC	6.45 ± 0.85	10.45 ± 1.15	6.15 ± 0.82	<u>100.00 ± 0.00</u>
	PF-Attack	19.15 ± 1.65	31.15 ± 2.25	19.65 ± 1.70	<u>95.75 ± 1.25</u>
	SAAO	31.75 ± 2.15	34.85 ± 2.45	37.85 ± 2.55	<u>98.25 ± 0.85</u>
	SpecEOT	26.46 ± 4.86	59.79 ± 2.51	84.79 ± 3.80	<u>100.00 ± 0.00</u>

Table 2. Transferability on ShapeNet under the ℓ_∞ budget $\epsilon = 0.10$. Underlined values denote white-box attack success rates on the source model. Bold values highlight the highest black-box transferability among the evaluated methods.

Source Model	Method	PointNet	PointNet++	DGCNN	PCT
PointNet	3D-Adv	<u>100.00 ± 0.00</u>	1.25 ± 0.15	1.55 ± 0.21	0.22 ± 0.08
	KNN	<u>99.25 ± 0.45</u>	17.15 ± 1.25	18.95 ± 1.35	3.55 ± 0.45
	AdvPC	<u>100.00 ± 0.00</u>	18.55 ± 1.42	14.35 ± 1.15	7.45 ± 0.85
	PF-Attack	<u>62.85 ± 2.45</u>	31.25 ± 1.85	23.55 ± 1.65	17.25 ± 1.25
	SAAO	<u>100.00 ± 0.00</u>	43.15 ± 2.85	42.05 ± 2.65	28.85 ± 1.95
	SpecEOT	<u>91.83 ± 3.61</u>	44.38 ± 2.15	45.11 ± 2.77	3.12 ± 1.01
PointNet++	3D-Adv	2.45 ± 0.35	<u>100.00 ± 0.00</u>	1.65 ± 0.25	0.42 ± 0.12
	KNN	6.15 ± 0.75	<u>100.00 ± 0.00</u>	5.45 ± 0.65	1.85 ± 0.35
	AdvPC	26.55 ± 1.85	<u>99.55 ± 0.35</u>	24.15 ± 1.65	20.45 ± 1.45
	PF-Attack	34.15 ± 2.15	<u>79.45 ± 2.55</u>	22.65 ± 1.55	20.85 ± 1.65
	SAAO	56.85 ± 3.15	<u>100.00 ± 0.00</u>	52.85 ± 2.95	30.15 ± 2.25
	SpecEOT	55.58 ± 3.88	<u>100.00 ± 0.00</u>	54.58 ± 3.61	45.21 ± 1.80
DGCNN	3D-Adv	2.65 ± 0.45	0.55 ± 0.15	<u>100.00 ± 0.00</u>	0.11 ± 0.05
	KNN	37.85 ± 2.45	33.95 ± 2.15	<u>100.00 ± 0.00</u>	11.55 ± 1.15
	AdvPC	23.45 ± 1.65	37.25 ± 2.35	<u>98.75 ± 0.55</u>	19.85 ± 1.45
	PF-Attack	39.45 ± 2.55	39.85 ± 2.65	<u>72.25 ± 2.85</u>	29.85 ± 1.85
	SAAO	24.05 ± 1.75	32.65 ± 2.15	<u>99.45 ± 0.35</u>	25.15 ± 1.95
	SpecEOT	37.29 ± 1.80	34.46 ± 1.33	<u>100.00 ± 0.00</u>	29.38 ± 1.00
PCT	3D-Adv	0.48 ± 0.15	0.18 ± 0.05	2.25 ± 0.35	<u>100.00 ± 0.00</u>
	KNN	0.62 ± 0.18	0.75 ± 0.22	2.15 ± 0.32	<u>100.00 ± 0.00</u>
	AdvPC	18.25 ± 1.35	27.65 ± 1.85	29.75 ± 1.95	<u>99.65 ± 0.25</u>
	PF-Attack	27.15 ± 1.75	29.85 ± 2.05	20.85 ± 1.55	<u>90.55 ± 1.85</u>
	SAAO	31.15 ± 2.15	23.55 ± 1.65	30.45 ± 2.25	<u>97.95 ± 0.85</u>
	SpecEOT	34.12 ± 2.97	36.04 ± 1.80	27.71 ± 1.80	<u>62.88 ± 5.41</u>

4.3. Ablation Study

To isolate the contribution of each algorithmic component, we perform a two-source ablation study on ModelNet40 under the ℓ_∞ budget $\epsilon = 0.10$. Table 3 compares the full SpecEOT pipeline against three variants using PointNet and DGCNN as source models.

Table 3. Two-source ablation study on ModelNet40 at $\epsilon = 0.10$. Underlined values denote white-box attack success rates on the source model. Bold values highlight the highest black-box transferability among the evaluated methods.

Source	Variant	PointNet	PointNet++	DGCNN	PCT
PointNet	Full	<u>94.56 ± 2.33</u>	51.34 ± 3.12	65.46 ± 1.80	45.00 ± 5.41
	Cyclic sampling	<u>98.28 ± 2.40</u>	45.62 ± 0.91	63.38 ± 1.00	41.52 ± 0.13
	w/o identity	<u>92.06 ± 2.64</u>	45.46 ± 7.86	51.42 ± 1.80	42.25 ± 0.79
	w/o gradient normalization	<u>90.30 ± 3.11</u>	48.33 ± 3.61	61.38 ± 9.49	38.83 ± 2.51
DGCNN	Full	43.96 ± 2.51	44.58 ± 2.51	<u>100.00 ± 0.00</u>	50.54 ± 1.33
	Cyclic sampling	21.88 ± 0.63	38.75 ± 0.69	<u>100.00 ± 0.00</u>	44.38 ± 0.89
	w/o identity	25.00 ± 6.25	45.62 ± 3.12	<u>100.00 ± 0.00</u>	41.67 ± 3.61
	w/o gradient normalization	26.04 ± 1.80	44.38 ± 9.38	<u>100.00 ± 0.00</u>	48.21 ± 6.12

The full configuration achieves the best overall transfer performance across the evaluated target architectures. It independently samples frequency bands and perturbation signs at each iteration, includes an exact identity view in EOT aggregation, and applies per-sample ℓ_2 gradient normalization before the Adam update.

Replacing independent sampling with a fixed cyclic schedule (cyclic sampling) reduces transferability for several source–target pairs. This result suggests that drawing

fresh band and sign samples at every iteration helps reduce overfitting to a deterministic transformation sequence and promotes more transferable adversarial structures.

Removing the identity view (w/o identity) reduces black-box transferability in five of the six source–target pairs. The identity view anchors the optimization to the exact adversarial point cloud, helping the aggregated EOT gradient remain relevant to the input evaluated at test time. Without this anchor, the attack relies entirely on randomized spectral views, which can weaken cross-model generalization.

Finally, disabling gradient normalization (w/o gradient normalization) lowers black-box ASR in all six source–target pairs. Per-sample normalization stabilizes the scale of the aggregated Monte Carlo gradient estimate and supports reliable projected Adam updates within the ℓ_∞ boundary.

4.4. Hyperparameter Analysis

We conduct a staged hyperparameter search on ModelNet40 to establish a reliable default configuration for SpecEOT. Table 4 reports the mean ASR across three unseen target models (DGCNN, PointNet++, and PCT), the corresponding standard deviation, and the per-batch wall-clock runtime for each candidate value.

Table 4. Hyperparameter search results. Attack success rate (ASR) is averaged across three unseen target models (DGCNN, PointNet++, and PCT). Runtime denotes the wall-clock time required to process a single batch of 32 examples. These values were obtained from independent runs conducted for the staged hyperparameter search, and are not repeated measurements of the canonical configuration reported in Table 1.

Stage	Candidate	Mean ASR (%)	Std. (pp)	Runtime (s)
ρ	0.05	48.78	1.20	23.97
	0.075	50.51	0.60	23.73
	0.10	50.86	0.60	23.37
	0.125	51.90	1.59	23.50
	0.15	52.25	1.80	25.03
	0.20	53.99	3.94	23.97
Views	2	46.82	0.60	5.77
	4	49.25	1.04	9.53
	8	53.99	3.94	23.37
	12	53.42	2.76	45.83
	16	50.29	1.80	75.60
Bands	4	50.42	2.08	46.43
	8	53.42	2.76	46.07
	16	50.64	1.59	46.40
	32	48.90	1.59	46.27
Steps	30	30.97	3.18	28.37
	50	53.42	2.76	45.80
	70	46.93	3.94	63.90
	100	53.49	0.60	91.27
Learning rate	0.001	13.75	3.61	91.17
	0.002	42.15	1.59	91.67
	0.003	53.49	0.60	91.87
	0.005	53.26	2.62	91.40
Gradient norm	Mean absolute	53.26	2.62	92.90
	ℓ_2	53.69	1.59	91.93

For the spectral perturbation strength ρ , moderate values yield stable transfer gains while maintaining nearly constant computational overhead. Therefore, we select a value that balances mean attack performance and variability across runs.

The view count results indicate that a moderate K provides the best computational trade-off. Too few views provide a limited Monte Carlo estimate of the EOT objective, whereas additional views substantially increase computation without proportional performance gains.

The results for the number of equal-count frequency bands B favor moderate granularity; both coarser and finer partitions reduce overall transferability.

For the number of optimization steps T and the Adam learning rate α , we select settings that provide stable convergence within a practical computational budget. Finally, unit- ℓ_2 gradient normalization marginally outperforms mean-absolute normalization, and is adopted as the default.

4.5. Spectral Diversity: $B \times \rho$ Interaction

The number of frequency bands B and the spectral perturbation strength ρ jointly control the diversity of the randomized graph-spectral views sampled during each iteration (see Equation (8)). To analyze this interaction, we perform a controlled grid search while keeping all other hyperparameters fixed.

Table 5 reports the average transfer ASR and a quantitative spectral-view diversity metric for each (B, ρ) pair. The diversity metric is the mean pairwise Euclidean distance among the generated views in the graph Fourier domain, normalized by the ℓ_2 norm of the clean point cloud.

The results show a clear joint pattern. For every fixed band count, increasing ρ increases the diversity score of the spectral view because stronger multiplicative perturbations generate more varied views. When B is large, however, each band contains fewer Fourier modes, so even a moderate ρ produces lower overall diversity. Conversely, combining a small B with a large ρ can overperturb broad low- or mid-frequency regions and reduce the structural consistency needed for effective black-box transfer. Across the nine reported configurations, spectral view diversity has a Pearson correlation of $r = 0.484$ and a Spearman rank correlation of $\rho_s = 0.567$ with mean transfer ASR. These descriptive coefficients indicate a moderate but non-monotonic association; given the small number of configurations, we do not interpret them as a significance test or a causal guarantee.

Table 5. Interaction between the number of frequency bands (B) and spectral perturbation strength (ρ). Transfer ASR is averaged over three target models. These values were obtained from independent runs conducted for the band-strength interaction analysis, and are not repeated measurements of the canonical configuration reported in Table 1.

B	ρ	Transfer ASR (%)	Spectral-View Diversity
4	0.05	53.26 $\pm$ 2.41	0.021404 $\pm$ 0.000301
4	0.10	53.04 $\pm$ 2.76	0.041344 $\pm$ 0.000588
4	0.15	52.65 $\pm$ 0.60	0.058766 $\pm$ 0.000844
8	0.05	51.88 $\pm$ 1.04	0.012590 $\pm$ 0.000077
8	0.10	53.61 $\pm$ 4.21	0.024428 $\pm$ 0.000147
8	0.15	53.31 $\pm$ 1.20	0.034969 $\pm$ 0.000214
16	0.05	49.10 $\pm$ 1.59	0.007215 $\pm$ 0.000148
16	0.10	52.92 $\pm$ 1.80	0.014069 $\pm$ 0.000285
16	0.15	51.18 $\pm$ 1.20	0.020204 $\pm$ 0.000411

4.6. Controlled Input Corruptions

We evaluate robustness by applying controlled input corruptions before attack generation. For Gaussian jitter, coordinate-wise noise is added and clipped to ± 0.02 . For point dropout, 10% of the points are removed and replaced by duplicates of retained points, ensuring that all models process exactly 1024 coordinates. We restrict the evaluation to examples that remain correctly classified by the source and all target models after corruption. Table 6 reports the corresponding transfer success rates. These evaluations are

controlled algorithmic stress tests rather than comprehensive simulations of real-world LiDAR acquisition.

Table 6. Transferability under controlled input corruptions on ModelNet40, using PointNet as the source model.

Condition	DGCNN	PCT	PointNet++
No corruption	65.46	45.00	51.34
Gaussian jitter	50.51	47.53	40.76
10% point dropout with replication	58.01	42.68	48.22

4.7. Computational Efficiency

To assess computational cost, we measure the wall-clock runtime and peak GPU memory required to generate a batch of adversarial examples. Table 7 compares SpecEOT with SAAO using a batch size of $n = 32$. Although SpecEOT processes more spectral views, its total runtime and per-sample processing time are substantially lower than those of SAAO, at the cost of moderately higher peak memory usage.

The primary computational overhead in SpecEOT stems from the graph Fourier transform and its inverse, which are applied to each randomized view. Because the fixed spectral basis is computed only once per point cloud, the eigendecomposition is not repeated during optimization. Under the evaluated implementation and hardware, SpecEOT therefore provides a favorable runtime–memory trade-off relative to our SAAO implementation.

Table 7. Wall-clock runtime and peak memory usage comparison. Measurements reflect the attack generation cost for a single batch of 32 examples on identical server hardware.

Method	n	ϵ	Steps	Views	Total (s)	s/Sample	Peak MB
SpecEOT	32	0.10	100	12	91.00	2.8439	1661.9
SAAO	32	0.10	100	8	272.52	8.5164	1391.2

4.8. Visualization

Figure 3 compares clean and adversarial point clouds that were initially classified correctly by all tested models. Visual inspection alone does not establish imperceptibility: the ℓ_∞ budget provides a hard coordinate-wise bound, while the post hoc CD and HD values summarize the resulting geometric distortion.




Figure 3. Qualitative comparison of adversarial examples generated by SpecEOT using a PointNet source model on ModelNet40 at $\epsilon = 0.10$. Chamfer distance (CD) and Hausdorff distance (HD) are reported as post hoc distortion metrics because SpecEOT does not explicitly optimize them during generation.

4.9. Case Study

To examine the graph-spectral characteristics of the update direction, we compare the normalized spectral energy of the SpecEOT update against the gradient spectra of three unseen targets (PointNet++, DGCNN, and PCT) for two representative ModelNet40 examples. Figure 4 shows that the SpecEOT update and the target-gradient mean share a broad concentration in low-to-mid-frequency bands, although their agreement varies across samples and individual bands.




Figure 4. Graph-spectral energy profiles for two representative ModelNet40 examples. The green curve shows the SpecEOT update direction obtained from the PointNet source model. The dashed curve shows the mean gradient energy profile of the three unseen target models (PointNet++, DGCNN, and PCT), and the shaded region spans their minimum-to-maximum values in each spectral band. This figure is a cross-architecture diagnostic for the selected examples rather than an estimate of seed-wise uncertainty.

This observation is consistent with the hypothesis that averaging gradients across structured spectral views can preserve components shared across architectural variations. Because the figure contains only two selected examples, we use it as an illustrative diagnostic rather than statistical evidence. The multi-seed $B \times \rho$ analysis in Table 5, including the reported diversity–ASR correlations, provides the complementary quantitative assessment.

4.10. Context from Recent Baselines

For additional numerical context, Table 8 lists attack success rates for two recent transfer-based methods, MAT-ADV [12] and COSA [13], under the ℓ_∞ budget $\epsilon = 0.18$. The MAT-ADV and COSA values are taken directly from the original papers because official implementations were unavailable; N/A indicates a source–target pair that was not reported. This is a cross-paper reference rather than a controlled comparison. Differences in model checkpoints, preprocessing pipelines, evaluation subsets, and possibly the definition of the ASR denominator prevent a rigorous ranking. The table is included only to illustrate the range of results published under the respective original protocols, and no claim of absolute superiority is made.

Table 8. Cross-paper context on ModelNet40 at $\epsilon = 0.18$. MAT-ADV and COSA values are taken from their respective publications, whereas SpecEOT values are obtained using our implementation. Protocol differences preclude a rigorous ranking; N/A means not reported. Underlined values denote white-box attack success rates on the source model.

Source	Method	PointNet	PointNet++	DGCNN	PCT
PointNet	MAT-ADV	<u>81.57</u>	24.96	24.78	N/A
	COSA	<u>97.70</u>	N/A	55.10	48.8
	SpecEOT	<u>96.88</u>	56.25	71.88	50.33

Table 8. Cont.

Source	Method	PointNet	PointNet++	DGCNN	PCT
DGCNN	MAT-ADV	24.13	24.34	<u>97.14</u>	N/A
	COSA	39.70	N/A	<u>97.60</u>	45.7
	SpecEOT	31.25	43.75	<u>100</u>	53.12
PCT	MAT-ADV	N/A	N/A	N/A	N/A
	COSA	35.50	N/A	58.60	<u>97.70</u>
	SpecEOT	28.12	62.5	81.25	<u>100</u>
PointNet++	MAT-ADV	21.58	<u>91.74</u>	28.26	N/A
	COSA	N/A	N/A	N/A	N/A
	SpecEOT	28.12	<u>100</u>	53.12	40.62

5. Discussion and Conclusions

Mechanism of source-agnostic spectral randomization. The empirical results support a specific bounded claim. Under the evaluated object-classification protocol, stronger source-conditioned adaptation does not guarantee stronger transferability. Loss- or gradient-based view weighting may emphasize directions favored by the surrogate model. In contrast, uniform sampling and equal EOT aggregation average the gradient over a predefined model-independent transformation distribution. The cross-architecture spectral diagnostic and multi-seed diversity analysis are consistent with a less source-dominated update direction. However, this evidence does not imply that equal weighting is universally optimal. In particular, the $1/K$ Monte Carlo variance reduction associated with EOT improves estimation of a fixed expectation but does not intrinsically guarantee cross-architecture generalization.

Efficiency and deployment scope. SpecEOT requires K source-model forward and backward passes per iteration, together with an eigendecomposition of the $N \times N$ graph Laplacian for each input. Using a fixed basis avoids recomputing the decomposition during optimization. Nevertheless, the dense matrix construction used here is unsuitable for strict real-time pipelines or large-scale LiDAR scenes. The reported wall-clock and memory measurements characterize only the current 1024-point implementation, and do not imply real-time readiness. Future scaling efforts could explore sparse eigensolvers, localized graph blocks, and batched approximate spectral bases.

Limitations and future work. Our evaluation is limited to 1024-point object classification on ModelNet40 and ShapeNet. The controlled Gaussian jitter and point dropout tests represent only two isolated sensor corruptions, and do not capture non-uniform LiDAR density, range-dependent noise, motion distortion, adverse weather, severe occlusion, background clutter, or reconstruction artifacts. Because the spectral basis is fixed from the initial input graph, severe topology changes caused by real-world corruptions may alter the structural meaning of the frequency bands and reduce the effectiveness of the sampled views. Furthermore, we have not evaluated large outdoor scenes, object detection, semantic segmentation, targeted attacks, defense mechanisms, or physical-world deployment. Consequently, references to autonomous transport and mobile robotics provide security motivation rather than evidence of deployment readiness. Future work should investigate dynamic spectral bases, variable point count processing, realistic sensor simulation, detection-specific loss functions, and physical acquisition constraints.

Conclusion. We present SpecEOT, a source-agnostic graph-spectral EOT framework that combines a fixed GFT basis, independent uniform band-and-sign sampling, equal gradient aggregation, and projected Adam ascent. Experiments on ModelNet40 and ShapeNet show that this deliberately simple design achieves competitive or superior black-box transferability across many source–target pairs. The two-source ablation, hyperparameter

analysis, controlled corruption tests, and computational measurements further characterize where the method is effective and where limitations remain. Overall, the results support the bounded conclusion that reducing source-conditioned spectral adaptation can improve transferability under the evaluated 3D object classification protocol.

Author Contributions: Conceptualization, Y.G. and J.L.; methodology, H.L. (Haoran Li); validation, H.L. (Hongjia Liu); formal analysis, Y.G.; resources, J.X. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded in part by the National Natural Science Foundation of China under Grant 62372096 and the Fundamental Research Funds for the Central Universities under Grant N2417008.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The ModelNet40 and ShapeNet benchmark datasets are available from their original providers. The attack implementation and experiment scripts are available at <https://github.com/a772316182/SpecEOT> (accessed on 2 July 2026).

Conflicts of Interest: The authors declare no conflicts of interest.

References

- Chen, X.; Ma, H.; Wan, J.; Li, B.; Xia, T. Multi-view 3d object detection network for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*; IEEE: New York, NY, USA, 2017.
- Guo, Y.; Wang, H.; Hu, Q.; Liu, H.; Liu, L.; Bennamoun, M. Deep learning for 3d point clouds: A survey. *Proc. IEEE Trans. Pattern Anal. Mach. Intell.* **2020**, *43*, 4338–4364. [[CrossRef](#)]
- Ma, X.; Qin, C.; You, H.; Ran, H.; Fu, Y. Rethinking Network Design and Local Geometry in Point Cloud: A Simple Residual MLP Framework. In *Proceedings of the International Conference on Learning Representations 2022*; OpenReview: Alameda, CA, USA, 2022.
- Liu, R.; Lam, K.Y.; Zhou, W.; Wu, S.; Zhao, J.; Hu, D.; Gong, M. STBA: Towards Evaluating the Robustness of DNNs for Query-Limited Black-box Scenario. In *Proceedings of the IEEE Transactions on Multimedia 2025*; IEEE: New York, NY, USA, 2025.
- Yang, M.; Liu, D.; Tang, K.; Zhou, P.; Chen, L.; Chen, J. Hiding Imperceptible Noise in Curvature-Aware Patches for 3D Point Cloud Attack. In *Proceedings of the European Conference on Computer Vision 2024*; Springer: Berlin/Heidelberg, Germany, 2024.
- Lou, T.; Jia, X.; Gu, J.; Liu, L.; Liang, S.; He, B.; Cao, X. Hide in thicket: Generating imperceptible and rational adversarial perturbations on 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition 2024*; IEEE: New York, NY, USA, 2024.
- Tang, K.; Du, Z.; Peng, W.; Wang, X.; Liu, D.; Liu, L.; Tian, Z. Imperceptible 3D Point Cloud Attacks on Lattice-based Barycentric Coordinates. In *Proceedings of the AAAI Conference on Artificial Intelligence 2025*; AAAI: Washington, DC, USA, 2025.
- He, B.; Liu, J.; Li, Y.; Liang, S.; Li, J.; Jia, X.; Cao, X. Generating transferable 3d adversarial point cloud via random perturbation factorization. In *Proceedings of the AAAI Conference on Artificial Intelligence 2023*; AAAI: Washington, DC, USA, 2023.
- Hamdi, A.; Rojas, S.; Thabet, A.; Ghanem, B. Advpc: Transferable adversarial perturbations on 3d point clouds. In *Proceedings of the European Conference on Computer Vision 2020*; Springer: Berlin/Heidelberg, Germany, 2020.
- Liu, D.; Hu, W. Imperceptible transfer attack and defense on 3d point cloud classification. In *Proceedings of the IEEE Transactions on Pattern Analysis and Machine Intelligence*; IEEE: New York, NY, USA, 2022.
- Chen, H.; Zhao, S.; Yang, X.; Yan, H.; He, Y.; Xue, H.; Qian, F.; Su, H. Anf: Crafting transferable adversarial point clouds via adversarial noise factorization. In *Proceedings of the IEEE Transactions on Big Data*; IEEE: New York, NY, USA, 2024.
- Tang, K.; Gao, Y.; Peng, W.; Wang, X.; Fang, M.; Zhu, P. Transferable and undefendable point cloud attacks via medial axis transform. *Comput. Aided Geom. Des.* **2026**, *125*, 102515. [[CrossRef](#)]
- Tang, K.; Liu, X.; Peng, W.; Wang, X.; Liu, D.; Zhu, P.; Lu, C.; Tian, Z. Rethinking Transferable Adversarial Attacks on Point Clouds from a Compact Subspace Perspective. *arXiv* **2026**, arXiv:2601.23102.
- Chung, F.R. *Spectral Graph Theory*; American Mathematical Society: Providence, RI, USA, 1997; Volume 92.
- Rahaman, N.; Baratin, A.; Arpit, D.; Draxler, F.; Lin, M.; Hamprecht, F.; Bengio, Y.; Courville, A. On the spectral bias of neural networks. In *Proceedings of the International Conference on Machine Learning*; PMLR: New York, NY, USA, 2019.
- Hu, Q.; Liu, D.; Hu, W. Exploring the devil in graph spectral domain for 3d point cloud attacks. In *Proceedings of the European Conference on Computer Vision*; Springer: Berlin/Heidelberg, Germany, 2022.

17. Hu, S.; Liu, D.; Hu, W. Improving the Transferability of Point Cloud Attack via Spectral-Aware Admix and Optimization Designs. In *Proceedings of the Artificial Neural Networks and Machine Learning*; Springer: Cham, Switzerland, 2026.
18. Athalye, A.; Engstrom, L.; Ilyas, A.; Kwok, K. Synthesizing Robust Adversarial Examples. In *Proceedings of the International Conference on Machine Learning*; PMLR: New York, NY, USA, 2018.
19. Wang, T.; Ying, Z.; Li, Q. Boost adversarial transferability by uniform scale and mix mask method. *arXiv* **2023**, arXiv:2311.12051.
20. Zheng, M.; Wu, K.; Fan, Y.; Huang, R.; Wu, B. Seeking flat minima over diverse surrogates for improved adversarial transferability: A theoretical framework and algorithmic instantiation. In *Proceedings of the IEEE Transactions on Pattern Analysis and Machine Intelligence*; IEEE: New York, NY, USA, 2026.
21. Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; Vladu, A. Towards deep learning models resistant to adversarial attacks. In *Proceedings of the International Conference on Learning Representations*; OpenReview: Alameda, CA, USA, 2018.
22. Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I.; Fergus, R. Intriguing properties of neural networks. In *Proceedings of the International Conference on Learning Representations*; OpenReview: Alameda, CA, USA, 2014.
23. Goodfellow, I.; Shlens, J.; Szegedy, C. Explaining and harnessing adversarial examples. In *Proceedings of the International Conference on Learning Representations*; OpenReview: Alameda, CA, USA, 2015.
24. Carlini, N.; Wagner, D. Towards evaluating the robustness of neural networks. In *Proceedings of the 2017 IEEE Symposium on Security and Privacy (sp)*; IEEE: New York, NY, USA, 2017.
25. Wen, Y.; Lin, J.; Chen, K.; Chen, C.P.; Jia, K. Geometry-aware generation of adversarial point clouds. In *Proceedings of the IEEE Transactions on Pattern Analysis and Machine Intelligence*; IEEE: New York, NY, USA, 2020.
26. Xiang, C.; Qi, C.R.; Li, B. Generating 3d adversarial point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*; IEEE: New York, NY, USA, 2019.
27. Tsai, T.; Yang, K.; Ho, T.Y.; Jin, Y. Robust adversarial objects against deep learning models. In *Proceedings of the AAAI Conference on Artificial Intelligence 2020*; AAAI: Washington, DC, USA, 2020.
28. Qi, C.R.; Su, H.; Mo, K.; Guibas, L.J. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*; IEEE: New York, NY, USA, 2017.
29. Qi, C.R.; Yi, L.; Su, H.; Guibas, L.J. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates Inc.: Red Hook, NY, USA, 2017.
30. Wang, Y.; Sun, Y.; Liu, Z.; Sarma, S.E.; Bronstein, M.M.; Solomon, J.M. Dynamic graph cnn for learning on point clouds. *ACM Trans. Graph.* **2019**, *38*, 146. [[CrossRef](#)]
31. Wu, X.; Jiang, L.; Wang, P.S.; Liu, Z.; Liu, X.; Qiao, Y.; Ouyang, W.; He, T.; Zhao, H. Point transformer v3: Simpler faster stronger. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*; IEEE: New York, NY, USA, 2014.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.