HHA: Hierarchical Hyperbolic Constraints for Imperceptible Point Cloud Attacks

Keke Tang[1*], Yu Liao[1*], Weilong Peng[1(✉)], Xiaofei Wang[2], Daizong Liu[3], Zhongyun Hua[4], Peican Zhu[5(✉)], and Zhihong Tian[1]

[1] Guangzhou University, Guangzhou, China
{tangbohutbh,liaoyu8896}@gmail.com, {wlpeng,tianzhihong}@gzhu.edu.cn
[2] University of Science and Technology of China, Hefei, China
wxf9545@mail.ustc.edu.cn
[3] Wuhan University, Wuhan, China
daizongliu@whu.edu.cn
[4] Harbin Institute of Technology, Shenzhen, China
huazhongyun@hit.edu.cn
[5] Northwest Polytechnical University, Xi'an, China
ericcan@nwpu.edu.cn

Abstract. Adversarial attacks on point clouds require effective constraints to ensure imperceptibility. However, existing methods often overlook the intrinsic hierarchical organization of 3D shapes, thereby limiting their ability to preserve structural coherence. In this paper, we propose HHA, a novel framework that generates hierarchy-aware adversarial perturbations by leveraging hyperbolic geometry. HHA first decomposes the input point cloud into semantic and geometric substructures to capture its multi-scale organization. Then, each substructure is embedded into hyperbolic space, where localized constraints limit distortion and maintain hierarchical consistency. This hyperbolic regularization keeps perturbations aligned with the underlying structure and thereby enhances imperceptibility. Extensive experiments validate that HHA produces adversarial point clouds with improved structural coherence and imperceptibility, outperforming state-of-the-art methods.

Keywords: Adversarial Attack · Point Clouds · Hyperbolic Space

1 Introduction

The concurrent advances in deep learning [15, 48] and depth-sensing technologies [12] have made it possible to address a wide range of 3D perception tasks using deep neural networks (DNNs) [39], leading to their widespread deployment in real-world applications such as autonomous driving and robotics. However, recent studies have shown that these models are vulnerable to adversarial attacks, where subtle perturbations to 3D point clouds can cause incorrect predictions [61], posing serious risks in safety-critical scenarios. These vulnerabilities have motivated extensive research on adversarial point cloud attacks, which

[*]Joint first authors. [(✉)]Corresponding authors.

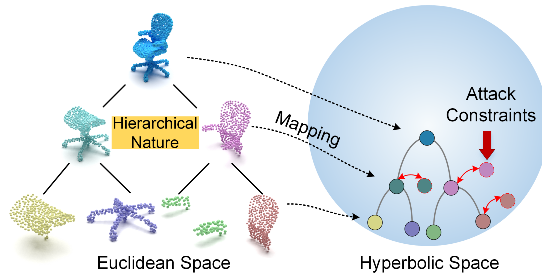

Fig. 1: Many 3D objects exhibit intrinsic hierarchical organization. Mapping them from Euclidean to hyperbolic space enables faithful modeling of such structure. Applying attack constraints in hyperbolic space guides perturbations to respect this hierarchy and results in adversarial point clouds with improved imperceptibility.

serve as essential tools for evaluating and improving the robustness of 3D deep models [19, 44, 55].

Ensuring imperceptibility is essential in adversarial attacks and has thus received significant attention. Prior methods aim to reduce perceptibility by enforcing geometric consistency, but most operate within Euclidean space. Early approaches rely on point-wise distance metrics such as the ℓ_2 norm, Chamfer distance, and Hausdorff distance [61]. More recent work introduces geometric priors, including curvature regularization [55], normal-guided perturbations [22], and tangent-aligned perturbation directions [11], to preserve local shape features. To overcome the limitations of purely Euclidean formulations, other methods adopt parametric representations to encode surface regularity, including global shape parameterizations and local structure-aware regularizations [40, 44]. These approaches typically operate on locally parameterized manifolds, which enhance surface coherence but still assume locally flat geometry and emphasize local consistency, therefore failing to capture the overall structural organization of 3D objects.

In practice, 3D shapes often exhibit intrinsic hierarchical organization, manifested as multi-scale geometry and part-level dependencies. Preserving this hierarchy is essential for producing adversarial perturbations that remain perceptually coherent and structurally plausible, yet most existing approaches do not model such relationships explicitly. Hyperbolic geometry, e.g., the Poincaré ball model, naturally encodes hierarchical relations through its negative curvature and exponential representational capacity [7], making it well suited for constraining multi-level structure. Motivated by this, we integrate hierarchical decomposition with localized hyperbolic regularization (see Fig. 1) so that perturbations are restricted at each semantic and geometric level, thereby improving structural fidelity under attack.

In this paper, we propose HHA, a hierarchical hyperbolic constraint framework for generating imperceptible adversarial point clouds. The primary objective is to preserve structural coherence, a property naturally represented in hyperbolic space through its negative curvature and exponential capacity for hierarchical relations. Specifically, HHA decomposes an input point cloud into multi-level semantic and geometric substructures. Each substructure is embedded into the Poincaré ball and subjected to localized hyperbolic regularization that limits distortion at that level. By enforcing level-wise constraints on predefined semantic and geometric groupings, HHA guides perturbations to remain faithful to the underlying multi-scale structure and thus improves perceptual plausibility. Extensive experiments on multiple datasets and diverse 3D DNN classifiers show that adversarial point clouds generated by HHA exhibit higher structural coherence and improved imperceptibility, outperforming those produced by state-of-the-art methods.

Overall, our contributions are summarized as follows:

- We highlight the importance of hierarchical structure in 3D point clouds and introduce hyperbolic geometry as a natural space for constraining adversarial perturbations.
- We propose an attack framework that decomposes point clouds into semantic and geometric parts and applies hyperbolic constraints to preserve multi-scale consistency.
- We show through extensive experiments that our method achieves higher imperceptibility across diverse datasets and models, surpassing existing approaches.

2 Related Work

Adversarial Attacks on 3D Point Clouds. Adversarial attacks on 3D point clouds are generally grouped into three types: addition-based attacks that introduce synthetic points to mislead classifiers [61]; deletion-based attacks that remove salient points critical for recognition [64, 68]; and perturbation-based attacks that alter the positions of existing points [38, 41–43, 45–47, 49–51, 54, 61]. This paper specifically focuses on perturbation-based approaches.

Early studies adapted 2D adversarial methods such as C&W [3] and FGSM [8] to 3D classifiers [22,24,61], while later ones employed generative modeling through latent perturbation [16] or GAN-based synthesis [69]. Across these approaches, imperceptibility is typically ensured by constraining perturbation magnitude and structural coherence.

Constraints for Imperceptible Adversarial Attacks. A common strategy to enhance imperceptibility is to constrain perturbations using geometric metrics such as the l_2 norm, Chamfer distance, and Hausdorff distance [61]. Further improvements incorporate priors such as curvature preservation [55, 65] and normal- or tangent-aligned perturbation directions [11, 22]. Beyond geometric constraints, spectral-domain regularization has also been explored to preserve intrinsic shape properties, where perturbations are analyzed or constrained in the graph spectral domain to maintain smoothness and structural integrity [23]. In addition, parametric surface representations have been employed to encode structural regularity, either through global shape modeling or localized structure-aware constraints [40, 44]. In contrast, this work imposes constraints from a hierarchical perspective by embedding multi-level substructures into hyperbolic space, which enables the preservation of both local fidelity and global structural coherence.

Deep Learning Models for 3D Point Clouds. Deep learning has substantially advanced 3D point cloud classification. Early approaches converted point sets into voxel grids and applied 3D CNNs [29], but were constrained by computational cost and limited resolution. The advent of PointNet and PointNet++ [33, 34] enabled direct, permutation-invariant processing of raw points and inspired many subsequent architectures. Later developments include point-wise convolutional [18, 56], graph-based [53, 66], and attention-driven Transformer models [57, 58, 67], followed by recent Mamba-based state-space designs [10] that further improve efficiency and scalability. This work focuses on crafting imperceptible adversarial perturbations that remain effective across these representative 3D classifiers.

Hyperbolic Space. Hyperbolic space, characterized by constant negative curvature, provides a natural geometry for embedding hierarchical data with low distortion [31]. It has been widely applied in various domains such as knowledge graph embedding [1] and hierarchical classification [62]. Recently, several works have extended hyperbolic modeling to point cloud analysis, demonstrating its advantages in capturing structural hierarchy and geometric relations. Representative applications include point cloud completion [20], segmentation [36], reconstruction from RGB-D images [17], structural analysis [30], and robust feature learning [2]. Inspired by these advances, this paper embeds multi-level point cloud substructures into hyperbolic space to constrain adversarial perturbations, thereby preserving both local fidelity and global structural consistency.

3 Problem Formulation

3.1 Preliminary on Adversarial Attacks

Let $\mathcal{P} \in \mathbb{R}^{n\times 3}$ be a clean point cloud with ground-truth label $y \in \{1, \dots, Z\}$, and let f be a point cloud classifier. The goal of adversarial attack is to generate a perturbed input $\mathcal{P}' = \mathcal{P} + \Delta$ such that $f(\mathcal{P}') \neq y$, while keeping the perturbation imperceptible. A standard objective is:

$$\min_{\Delta} \; L_{\mathrm{mis}}(f, \mathcal{P} + \Delta, y) + \lambda_1 D(\mathcal{P}, \mathcal{P} + \Delta), \tag{1}$$

where L_{mis} denotes the misclassification loss, typically the negated cross-entropy, and D represents the distortion metric such as the ℓ_2 norm or Chamfer distance. Although the overall distortion is penalized, the lack of structural constraints often results in distortions that undermine the coherence of the object's inherent hierarchical organization.

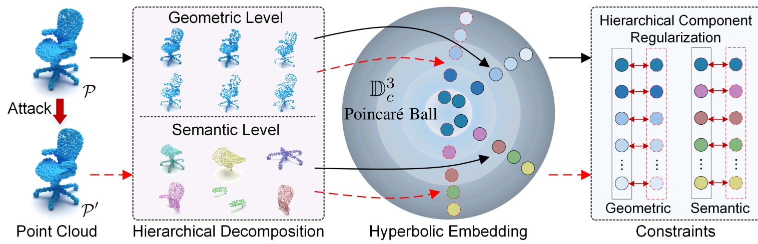

Fig. 2: Illustration of the hierarchical hyperbolic attack (HHA) framework. HHA decomposes a point cloud into semantic and geometric hierarchies, embeds each substructure in hyperbolic space, and imposes structure-aware constraints to guide imperceptible perturbations.

3.2 Adversarial Attacks in Hyperbolic Space

To promote structurally coherent adversarial point clouds, we introduce a hyperbolic regularization term that complements conventional distortion constraints. Unlike Euclidean metrics, which operate in flat geometry, hyperbolic space offers a natural model for hierarchical relationships, making it well-suited for structure-aware perturbation control.

Let $\mathcal{H} : \mathbb{R}^{n\times 3} \to \mathbb{H}^{n\times d}$ be a mapping from the Euclidean input into a hyperbolic space, e.g., the Poincaré ball [31]. The overall optimization objective becomes:

$$\min_{\Delta} \quad L_{\text{mis}}(f, \mathcal{P}+\Delta, y) + \lambda_1 D(\mathcal{P}, \mathcal{P}+\Delta) + \lambda_2 D_{\mathbb{H}}(\mathcal{H}(\mathcal{P}), \mathcal{H}(\mathcal{P}+\Delta)), \tag{2}$$

where $D_{\mathbb{H}}$ measures the discrepancy in hyperbolic space and λ_2 balances the contribution of this structural constraint.

This hyperbolic consistency term encourages perturbations that better align with the shape's hierarchical structure, thereby enhancing global coherence and imperceptibility.

4 Hierarchical Hyperbolic Attacks

In this section, we present the hierarchical hyperbolic attack (HHA) framework comprising three components: *hyperbolic embedding* for capturing global structural priors, *hierarchical decomposition* for extracting multi-scale geometric and semantic components, and *constraint-based optimization* for enforcing structure-aware perturbation control. Please refer to Fig. 2 for an illustration.

4.1 Hyperbolic Embedding of Point Clouds

Preliminaries on Hyperbolic Geometry. We adopt the d-dimensional Poincaré ball model $(\mathbb{D}_c^d, g^{\mathbb{D}})$ of hyperbolic space, defined as

$$\mathbb{D}_c^d = \left\{ x \in \mathbb{R}^d : \|x\| < \frac{1}{\sqrt{c}} \right\}, \tag{3}$$

where $c > 0$ controls the magnitude of negative curvature. The Riemannian metric is conformally equivalent to the Euclidean metric g^E, scaled by

$$g^{\mathbb{D}} = \lambda_c^2 g^E, \quad \lambda_c = \frac{2}{1 - c\|x\|^2}. \tag{4}$$

The geodesic distance between two points $x, y \in \mathbb{D}_c^d$ is defined as

$$d_{\mathbb{D}_c}(x, y) = \frac{2}{\sqrt{c}} \tanh^{-1}\left(\sqrt{c}\,\|-x \oplus_c y\|\right), \tag{5}$$

where $\oplus_c$ denotes the Möbius addition in the Poincaré ball [7,13], given by

$$x \oplus_c y = \frac{(1 + 2c\langle x, y\rangle + c\|y\|^2)x + (1 - c\|x\|^2)y}{1 + 2c\langle x, y\rangle + c^2\|x\|^2\|y\|^2}. \tag{6}$$

This operation generalizes Euclidean vector addition while preserving negative curvature, and the resulting distance naturally models hierarchical and tree-like structures. Note that as $c \to 0$, the model degenerates to Euclidean space, recovering standard vector addition and distance.

Exponential Embedding into Hyperbolic Space. For each point $p_i \in \mathbb{R}^3$ in $\mathcal{P}$, we define its embedding into hyperbolic space via the exponential map centered at the origin:

$$h(p_i) := \exp_c(p_i) = \tanh(\sqrt{c}\|p_i\|) \cdot \frac{p_i}{\sqrt{c}\|p_i\|}. \tag{7}$$

This yields the embedded representation $\mathcal{H}(\mathcal{P}) = \{h(p_i)\}_{i=1}^n$, where each point resides in $\mathbb{D}_c^3$ and carries curvature-aware distance properties. This embedding enables curvature-aware distance measurement in hyperbolic space, which provides a foundation for subsequent structure-aware constraints.

4.2 Hierarchical Decomposition of Point Clouds

To support fine-grained, localized regularization, the point cloud $\mathcal{P}$ is decomposed along two complementary axes: geometric structure and semantic functionality.

Geometric Hierarchy. A coarse-to-fine spatial hierarchy $\{\mathcal{P}_m^{(g)}\}_{m=1}^{M_g}$ is constructed via recursive farthest point sampling (FPS). Starting from the entire shape $\mathcal{P}_1^{(g)} = \mathcal{P}$, each component $\mathcal{P}_m^{(g)}$ is recursively divided as follows:

- Sample a representative subset $\mathcal{P}_{m_1}^{(g)}$ by selecting half of the points in $\mathcal{P}_m^{(g)}$ using FPS;
- Define the complementary subset as $\mathcal{P}_{m_2}^{(g)} = \mathcal{P}_m^{(g)} \setminus \mathcal{P}_{m_1}^{(g)}$;
- Add both $\mathcal{P}_{m_1}^{(g)}$ and $\mathcal{P}_{m_2}^{(g)}$ to the component pool for further splitting.

The process continues until a maximum depth is reached or the number of points in a component falls below a threshold. This results in a binary decomposition tree containing M_g geometric components.

Semantic Hierarchy. A tree-structured semantic hierarchy $\{\mathcal{P}_m^{(s)}\}_{m=1}^{M_s}$ is constructed from part-level annotations:

– Assign a part label to each point in $\mathcal{P}$ using a pretrained segmentation network (e.g., PartSLIP [25]);
– Group points with the same label into initial components $\{\mathcal{P}_m^{(s)}\}_{m=1}^K$;
– Iteratively merge spatially adjacent components to form higher-level parts.

The merging process continues until no further merging is possible or a maximum level is reached, producing M_s semantic components.

4.3 Constraint-Based Optimization

Hierarchical Component Regularization. To preserve local structure, we impose distortion penalties on each component from both hierarchies. For each part $\mathcal{P}_m \in \{\mathcal{P}_m^{(s)}\} \cup \{\mathcal{P}_m^{(g)}\}$, the distortion in hyperbolic space is defined as:

$$
\begin{aligned}
D_{\mathbb{H}}(\mathcal{P}_m, \mathcal{P}'_m) = & \frac{1}{|\mathcal{P}_m|} \sum_{p_i \in \mathcal{P}_m} \min_{q_j \in \mathcal{P}'_m} d^2_{\mathbb{D}_c}\big(h(p_i), h(q_j)\big) \\
& + \frac{1}{|\mathcal{P}'_m|} \sum_{q_j \in \mathcal{P}'_m} \min_{p_i \in \mathcal{P}_m} d^2_{\mathbb{D}_c}\big(h(q_j), h(p_i)\big),
\end{aligned}
\tag{8}
$$

where $\mathcal{P}'_m = \mathcal{P}_m + \Delta_m$ is the adversarial counterpart of component $\mathcal{P}_m$, and $\Delta_m = \{\delta_i \mid p_i \in \mathcal{P}_m\}$ is the set of perturbations, with each δ_i being the perturbation applied to point p_i. This formulation adopts bidirectional nearest-neighbor Chamfer distance in hyperbolic space.
Unified Optimization Objective. The overall objective integrates the misclassification loss, general distortion penalty, and hierarchical hyperbolic regularization:

$$
\begin{aligned}
\min_{\Delta} \quad & L_{\text{mis}}(f, \mathcal{P} + \Delta, y) + \lambda_1 D(\mathcal{P}, \mathcal{P} + \Delta) \\
& + \lambda_2^{(s)} \frac{1}{M_s} \sum_{m=1}^{M_s} D_{\mathbb{H}}\big(\mathcal{P}_m^{(s)}, \mathcal{P}_m^{(s)} + \Delta_m^{(s)}\big) \\
& + \lambda_2^{(g)} \frac{1}{M_g} \sum_{m=1}^{M_g} D_{\mathbb{H}}\big(\mathcal{P}_m^{(g)}, \mathcal{P}_m^{(g)} + \Delta_m^{(g)}\big),
\end{aligned}
\tag{9}
$$

where $D_{\mathbb{H}}$ denotes the hyperbolic distance defined in Eq. 8 . Here, M_s and M_g represent the numbers of semantic and geometric components, and λ_1, $\lambda_2^{(s)}$, and $\lambda_2^{(g)}$ control the weights of the corresponding regularization terms.

The objective in Eq. 9 is optimized following the C&W iterative scheme [3].

5 Experimental Results

5.1 Experimental Setup

Implementation. The HHA framework is implemented in PyTorch [32], with hyperbolic operations supported by Geoopt [14]. The geometric hierarchy is built by recursively applying FPS until three levels are formed, while the semantic hierarchy is obtained using PartSLIP [25] to segment each object into 2–4 parts. The Poincaré ball model is adopted for hyperbolic space with curvature parameter $c = 1.0$. The distortion metric D is defined as CD plus 0.1 times HD. Regularization weights are set to $\lambda_1 = 1.0$, $\lambda_2^{(s)} = 0.1$, and $\lambda_2^{(g)} = 0.1$. Perturbation optimization is performed using the Adam optimizer ($\beta_1 = 0.9$, $\beta_2 = 0.999$) with a learning rate of 0.01, running 10 binary search rounds and 500 gradient steps. All experiments are conducted on a workstation equipped with dual 2.40 GHz CPUs, 128 GB RAM, and eight NVIDIA RTX 3090 GPUs.

Datasets. We evaluate our method on ModelNet40 [59], ShapeNet Part [4], and the real-world dataset ScanObjectNN [52]. All point clouds are uniformly sampled to 1,024 points as in [61]. Results on ShapeNet Part are included in the supplementary material.

In particular, ScanObjectNN is built from real-world indoor scene scans and features clutter, occlusions, and background noise, making it considerably more challenging than synthetic benchmarks.

Baseline Attack Methods. We compare our method with seven representative baselines covering diverse attack strategies: gradient-based PGD and IFGM [5], direction-based SI-Adv [11], the frequency-domain method AOF [21], the saliency-guided deformation method HIT-Adv [26], and geometry-aware optimization methods GeoA^3 [55] and LBC [40], a recent lattice-based attack. Together, these approaches span different perturbation spaces and design principles, providing a broad basis for evaluating the effectiveness of our method.

Victim 3D DNN Classifiers. We evaluate HHA against five representative point cloud DNN classifiers: PointNet [33], DGCNN [53], Point Transformer (PTv1) [67], PointMLP [28], and Mamba3D [10]. All models are trained following their original protocols.

Evaluation Setting and Metrics. For fair and consistent comparisons, each attack method is configured to reach its highest possible attack success rate (ASR), which measures the proportion of adversarial examples that cause misclassification by the target model. Under this setting of maximal adversarial effectiveness [44], we assess the imperceptibility of perturbations using six standard metrics: Chamfer distance (CD) [6], Hausdorff distance (HD) [37], l_2-norm (l_2), surface curvature (Curv), geometric regularity (GR) [55], and earth mover's distance (EMD) [35].

5.2 Comparison and Performance Analysis

Performance on ASR and Imperceptibility. Tab. 1 reports the results on ASR and imperceptibility metrics. Most attacks achieve nearly 100% ASR on both datasets, while HIT-Adv exhibits reduced ASR on PTv1 and PointMLP. Regarding imperceptibility, our method achieves the lowest distortion values across all metrics and victim models, with LBC typically being the second-best baseline. On PointNet with ScanObjectNN, it attains a CD of 0.474×10^{-4} and an EMD of 0.438×10^{-2}, markedly lower than any baseline. Consistent improvements are observed on DGCNN, PointMLP, PTv1, and Mamba3D, confirming

Table 1: Comparison on the perturbation sizes required by different methods to reach their highest achievable ASR. The evaluation is conducted across different DNN classifiers on ModelNet40 and ScanObjectNN.

Model	Attack	ModelNet40							ScanObjectNN						
		ASR (%)	CD (10^{-4})	HD (10^{-2})	l_2	GR	Curv (10^{-2})	EMD (10^{-2})	ASR (%)	CD (10^{-4})	HD (10^{-2})	l_2	GR	Curv (10^{-2})	EMD (10^{-2})
PointNet	PGD	100	26.582	10.777	2.804	0.406	5.681	5.705	100	20.118	7.314	2.722	0.307	7.641	4.993
	IFGM	100	8.982	9.859	1.195	0.368	1.209	1.537	100	5.915	6.297	0.993	0.258	1.830	1.463
	GeoA^3	100	4.869	0.524	1.423	0.215	0.348	2.415	100	3.425	0.655	1.355	0.128	0.574	0.493
	AOF	100	18.063	1.204	1.896	0.161	3.719	3.575	100	5.842	0.484	0.924	0.159	3.455	1.814
	SI-Adv	100	2.855	2.259	0.779	0.183	0.276	0.783	100	1.120	1.122	0.400	0.116	2.154	0.493
	HIT-Adv	100	389.907	13.346	9.900	0.126	4.387	30.876	100	135.931	4.232	4.503	0.092	2.744	13.999
	LBC	100	0.897	0.306	**0.269**	0.121	0.215	0.473	100	0.879	**0.092**	0.398	0.091	0.587	0.551
	Ours	100	**0.603**	**0.279**	0.334	**0.119**	**0.164**	**0.361**	100	**0.474**	0.108	**0.379**	**0.089**	**0.523**	**0.438**
DGCNN	PGD	100	28.517	8.772	2.853	0.347	5.915	5.927	100	19.288	1.030	2.766	0.145	7.963	5.091
	IFGM	100	10.701	7.148	1.622	0.363	2.849	3.777	100	4.821	0.842	0.891	0.120	3.349	2.256
	GeoA^3	100	13.017	2.059	1.589	0.169	1.700	2.279	100	9.328	0.886	1.508	0.213	4.337	2.241
	AOF	100	25.229	1.249	2.351	0.161	4.325	4.959	100	7.973	0.375	1.084	0.107	4.262	2.513
	SI-Adv	100	8.911	1.811	1.348	0.136	1.317	2.908	100	3.007	0.555	0.752	0.096	1.276	1.526
	HIT-Adv	100	315.570	10.189	8.863	0.096	4.032	27.682	100	81.896	3.149	3.602	0.093	2.320	11.176
	LBC	100	2.854	0.334	**0.572**	0.120	0.641	1.427	100	2.716	0.289	**0.541**	0.115	1.407	1.353
	Ours	100	**2.189**	**0.278**	1.157	**0.092**	**0.494**	**1.266**	100	**1.308**	**0.157**	0.876	**0.090**	**1.116**	**0.933**
PTv1	PGD	100	24.521	1.724	2.584	0.181	5.548	5.589	100	17.785	1.484	2.636	0.136	7.483	4.899
	IFGM	100	7.544	5.563	1.082	0.520	0.900	1.704	100	2.778	0.464	0.819	0.107	2.278	1.592
	GeoA^3	100	7.815	3.296	0.973	0.199	1.618	2.357	100	6.452	0.830	1.050	0.122	4.259	2.132
	AOF	100	31.840	1.164	3.065	0.194	6.943	5.581	100	11.957	0.370	1.029	0.106	4.220	2.432
	SI-Adv	100	13.456	3.044	1.806	0.186	2.334	3.385	100	6.833	1.186	1.221	0.115	2.253	2.090
	HIT-Adv	45.2	130.731	4.940	4.767	**0.095**	1.232	9.689	54.1	37.566	1.682	2.265	0.093	3.028	6.773
	LBC	100	5.188	0.482	0.564	0.128	0.725	1.311	100	4.076	0.174	0.813	0.109	2.789	1.647
	Ours	100	**2.698**	**0.461**	**0.553**	0.127	**0.603**	**1.169**	100	**2.389**	**0.106**	**0.731**	**0.090**	**2.151**	**1.388**
PointMLP	PGD	100	26.396	3.834	2.776	0.236	6.139	5.777	100	20.499	1.065	2.565	0.150	7.897	5.086
	IFGM	100	6.431	2.304	0.918	0.189	1.463	2.115	100	2.711	0.487	0.677	0.098	2.313	1.450
	GeoA^3	100	11.209	3.103	2.789	0.251	1.951	2.524	100	7.619	0.758	0.940	0.119	4.001	1.985
	AOF	100	24.630	1.318	2.327	0.164	4.871	4.896	100	8.358	0.248	0.906	0.102	4.100	2.324
	SI-Adv	100	11.859	2.205	1.661	0.174	1.435	3.619	100	3.222	0.610	0.627	0.095	1.095	1.243
	HIT-Adv	48.9	214.755	5.768	6.647	0.115	3.265	20.762	54.2	45.232	2.134	2.816	0.099	2.240	8.450
	LBC	100	3.922	0.419	0.920	0.117	1.024	1.770	100	2.076	**0.074**	0.697	0.094	2.211	1.338
	Ours	100	**1.811**	**0.411**	**0.882**	**0.115**	**0.453**	**0.983**	100	**0.979**	0.097	**0.611**	**0.090**	**0.928**	**0.754**
Mamba3D	PGD	100	28.407	3.170	2.763	0.229	6.084	5.872	100	21.271	1.371	2.604	0.159	8.108	5.162
	IFGM	100	6.075	1.652	0.854	0.179	1.827	1.979	100	4.902	1.521	0.962	0.146	3.211	1.935
	GeoA^3	100	8.833	1.772	1.019	0.255	1.959	2.043	100	5.748	0.899	1.083	0.177	3.588	2.014
	AOF	100	18.073	0.711	1.619	0.137	3.413	4.033	100	8.238	0.262	0.967	0.103	3.983	2.505
	SI-Adv	100	5.685	1.722	0.965	0.152	0.617	1.940	100	2.836	0.887	0.868	0.127	**0.833**	1.256
	HIT-Adv	84.7	171.139	7.310	5.786	0.122	2.733	17.675	77.6	110.132	3.317	4.038	0.122	2.410	12.606
	LBC	100	2.987	0.327	0.718	0.136	0.489	1.927	100	2.894	0.286	0.827	0.111	0.865	1.196
	Ours	100	**1.321**	**0.289**	**0.667**	**0.114**	**0.390**	**0.766**	100	**1.383**	**0.204**	**0.804**	**0.093**	0.968	**0.883**

strong generalization. Among baselines, GeoA^3, SI-Adv, and AOF show moderate imperceptibility in certain cases, whereas HIT-Adv introduces large displacements with metric values orders of magnitude higher. These results demonstrate that our approach generates adversarial examples with superior imperceptibility while maintaining attack strength.

Visualization. Fig. 3 shows adversarial point clouds generated by different attack methods against PointNet. PGD introduces dense scattered outliers, while AOF also yields irregular points and noticeable distortions. IFGM and GeoA^3 produce fewer outliers but still cause local structural inconsistencies. SI-Adv and LBC perform better, though some displaced points remain visible. HIT-

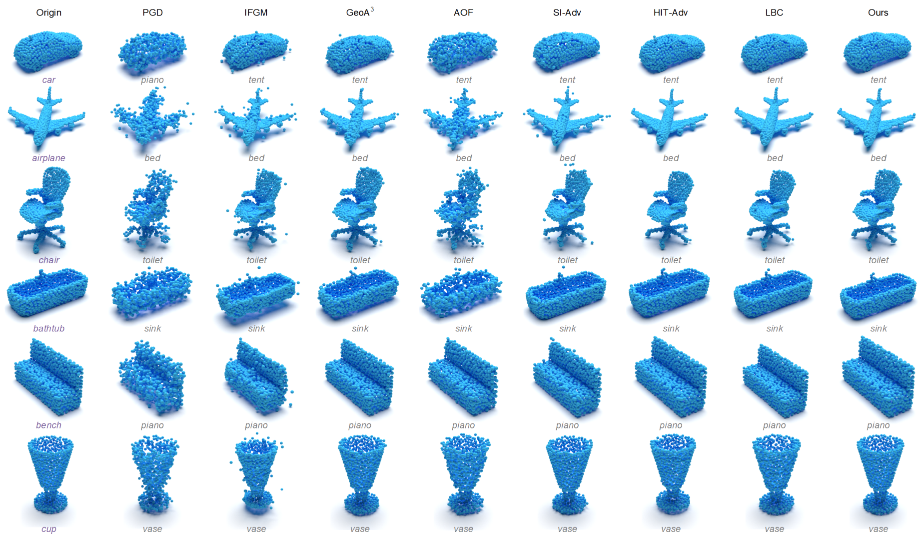

Fig. 3: Visualization of original and adversarial point clouds generated by various attack methods targeting PointNet on ModelNet40. Ground-truth and predicted labels are shown below each example in purple and gray, respectively.

Table 2: Comparison of semantic structural coherence (SSC) and geometric structural coherence (GSC) across different attack methods and DNN classifiers on ModelNet40 and ScanObjectNN. Lower values indicate better structural preservation.

Data	Model	SSC(10^{-1})								GSC(10^{-1})							
		PGD	IFGM	GeoA^3	AOF	SI-Adv	HIT-Adv	LBC	Ours	PGD	IFGM	GeoA^3	AOF	SI-Adv	HIT-Adv	LBC	Ours
ModelNet40	PointNet	2.18	0.62	0.34	1.33	0.29	6.66	0.43	**0.11**	2.52	0.83	0.51	1.61	0.47	7.53	1.09	**0.23**
	DGCNN	2.24	1.28	0.89	1.95	1.07	5.83	0.58	**0.40**	2.59	1.53	1.11	2.27	1.30	6.59	1.24	**0.61**
	PTv1	2.20	0.72	0.70	1.47	1.17	3.61	0.64	**0.38**	2.55	0.94	0.82	1.74	1.42	4.12	0.88	**0.47**
	PointMLP	2.21	0.93	0.86	1.82	1.28	5.73	0.61	**0.49**	2.55	1.16	1.09	2.13	1.53	6.48	0.97	**0.68**
	Mamba3D	2.19	0.90	0.82	1.70	1.38	4.07	0.71	**0.54**	2.53	1.12	1.04	1.98	1.64	4.64	1.00	**0.63**
ScanObjectNN	PointNet	1.86	0.54	0.50	0.60	0.44	3.09	0.35	**0.08**	2.13	0.66	2.93	0.72	2.15	3.49	1.54	**0.15**
	DGCNN	1.87	0.89	0.82	1.00	0.54	2.63	0.47	**0.29**	2.14	1.04	0.96	1.16	0.65	2.99	0.98	**0.38**
	PTv1	1.82	0.55	0.86	0.92	0.75	1.93	0.57	**0.34**	2.09	0.66	0.94	1.08	0.89	2.19	1.01	**0.36**
	PointMLP	1.77	0.56	0.75	1.16	0.56	4.49	0.48	**0.32**	2.02	0.67	0.87	1.34	0.67	5.04	0.92	**0.42**
	Mamba3D	1.84	0.75	0.60	0.81	1.13	3.31	0.56	**0.43**	2.19	0.88	0.72	0.93	1.32	3.74	0.89	**0.46**

Adv differs from other methods in that its perturbations often apply a near-rigid global shift to the entire point cloud, making the results appear visually coherent despite producing large distortions when measured by CD and HD (see Tab. 1). This unusual behavior has also been consistently observed in [19]. In contrast, our method preserves both local geometry and global arrangement, producing adversarial examples that remain visually similar to the originals and achieve substantially lower metric distortion, fully consistent with quantitative evaluations.

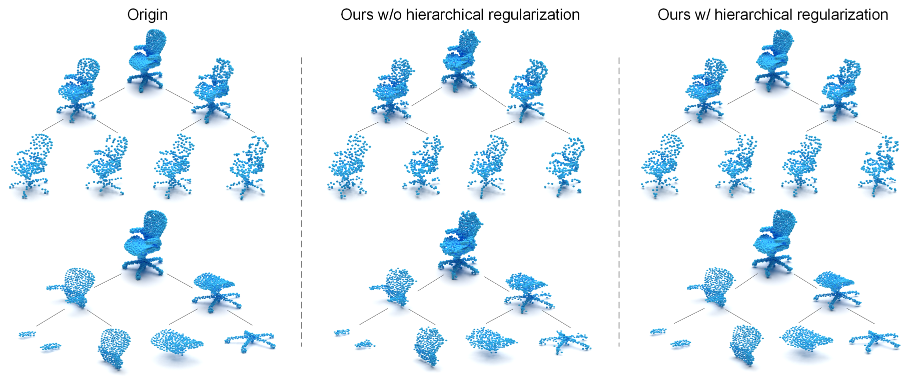

Fig. 4: Tree-structured visualization of geometric and semantic hierarchies for the original point cloud and adversarial point clouds generated by HHA with and without hierarchical regularization.

Performance on Structural Coherence. To further assess the ability of HHA to maintain structural hierarchy, we evaluate the change in structural coherence before and after attacks. Specifically, we compute the average value of $D_{\mathbb{H}}$ (Eqn. 9) across all semantic and geometric components, defining the resulting metrics as *semantic structural coherence* (SSC) and *geometric structural coherence* (GSC), where lower values indicate better preservation. As shown in Tab. 2, HHA consistently achieves the lowest SSC and GSC across all datasets and classifiers, confirming its effectiveness in preserving hierarchical structure and explaining its superior imperceptibility.

5.3 Ablation Studies and Other Analysis

Importance of Geometric/Semantic Hierarchical Regularization. The ablation results in Tab. 3 analyze the individual and joint contributions of geometric (G) and semantic (S) hierarchical regularization within the HHA framework. Applying either regularization alone effectively reduces distortion compared with the baseline configuration, with the semantic term showing a slightly stronger effect across most metrics. When both regularizations are applied together, distortions measured by CD, HD, and EMD are further reduced, consistently achieving the best imperceptibility on both ModelNet40 and ScanObjectNN. These results demonstrate that geometric and semantic hierarchical regularizations are complementary and that their joint integration is crucial for improving the imperceptibility of adversarial point clouds.

The visualization in Fig. 4 compares the hierarchical decompositions of adversarial point clouds generated with and without the proposed regularization. Without hierarchical constraints, perturbations introduce noticeable inconsistencies, especially in finer components, resulting in irregular substructures. In contrast, applying hierarchical regularization preserves the integrity of both geometric and semantic hierarchies, maintaining the relationships between parts

Table 3: Ablation study on the effect of geometric (G) and semantic (S) hierarchical regularization in HHA, evaluated on imperceptibility metrics when attacking PointNet.

Dataset	G	S	CD (10^{-4})	HD (10^{-2})	l_2	GR	Curv (10^{-2})	EMD (10^{-2})
ModelNet40	✓		0.686	0.301	0.397	0.122	0.182	0.431
		✓	0.621	0.279	0.359	0.127	0.168	0.375
	✓	✓	**0.603**	**0.279**	**0.334**	**0.119**	**0.164**	**0.361**
ScanObjectNN	✓		0.510	0.154	0.405	0.095	0.611	0.523
		✓	0.498	0.121	0.385	0.093	0.528	0.444
	✓	✓	**0.474**	**0.108**	**0.379**	**0.089**	**0.523**	**0.438**

Table 4: Effect of the hyperbolic curvature c on imperceptibility (CD/HD) for attacking PointNet on ModelNet40. We sweep c in hyperbolic space and include a Euclidean control by replacing the hyperbolic distance in the hierarchical component regularization with Euclidean distance. Lower values indicate better imperceptibility.

Curvature c	Euclidean	0.01	0.1	0.3	0.5	1.0	1.2	1.5
CD (10^{-4})	1.034	0.805	0.733	0.652	0.681	**0.603**	0.726	0.757
HD (10^{-2})	0.511	0.330	0.328	0.293	0.314	**0.279**	0.317	0.334

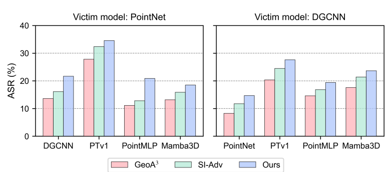

Fig. 5: Transferability of adversarial point clouds generated by GeoA^3, SI-Adv, and HHA in attacking PointNet and DGCNN trained on ModelNet40.

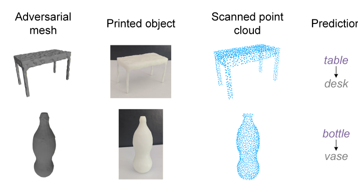

Fig. 6: Physical-world evaluation of HHA for attacks on PointNet (ModelNet40). Adversarial point clouds are reconstructed into meshes, 3D-printed, and re-scanned for classification.

across scales. These observations visually confirm that the proposed constraints effectively enforce structure-aware perturbations, aligning with improvements observed in quantitative imperceptibility metrics.

Sensitivity to the Hyperbolic Curvature c. We study how the curvature c influences imperceptibility for attacking PointNet on ModelNet40. As shown in Tab. 4, increasing c from 0.01 to 1.0 reduces both CD and HD, indicating improved structural preservation under moderate negative curvature. When c exceeds 1.0, both metrics increase, suggesting that overly large curvature weakens the constraint and leads to larger distortion.

Hierarchical Regularization in Hyperbolic vs. Euclidean Space. To isolate the effect of the underlying geometry, we apply the same hierarchical component regularization in both hyperbolic and Euclidean spaces. The Euclidean counterpart corresponds to replacing the hyperbolic distance with Euclidean distance, which is consistent with the $c \to 0$ limit of hyperbolic geometry. As

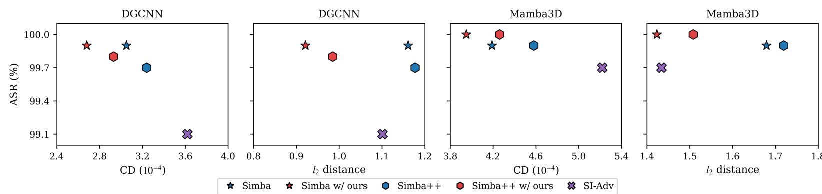

Fig. 7: The trade-off between ASR and imperceptibility, measured by CD and l_2, for different black-box methods in attacking DGCNN and Mamba3D on ModelNet40.

Table 5: The average time required by different methods to generate an adversarial point cloud to attack PointNet on ModelNet40.

Attack	PGD	IFGM	GeoA^3	AOF	SI-Adv	HIT-Adv	LBC	Ours
Time (s)	0.80	0.82	12.88	13.52	3.27	37.76	13.26	12.04

Table 6: Effect of different segmentation networks on HHA when attacking PointNet on ScanObjectNN.

Method	Time	mIoU(%)	CD(10^{-4})	HD(10^{-2})
PartSLIP	**∼6 s**	58.6	0.474	**0.108**
P^3-SAM	∼9 s	**63.4**	**0.469**	0.111

shown in Tab. 4, the Euclidean setting results in noticeably higher CD and HD compared to the hyperbolic formulation, while $c = 1.0$ achieves the lowest distortion. This comparison highlights the advantage of hyperbolic geometry for hierarchy-aware regularization.

Transferability and Black-box Performance. To ensure fairness, the transferability study considers only methods demonstrating strong imperceptibility in Tab. 1 and Fig. 3. The results in Fig. 5 indicate that adversarial point clouds generated by HHA achieve higher ASR on unseen models compared with GeoA^3 and SI-Adv, confirming its superior cross-model transfer capability. For black-box scenarios, Fig. 7 shows that incorporating HHA into SimBA [9] and SimBA++ [63] retains their high ASR while substantially lowering distortion. These findings verify that HHA not only performs effectively in white-box settings but also enhances the quality of attacks in black-box configurations.

Physical-World Transferability. We evaluate whether HHA transfers to the physical world using a reconstruction→3D printing→re-scanning pipeline when attacking PointNet on ModelNet40. As shown in Fig. 6, the adversarial effect persists after fabrication and re-acquisition: the re-scanned point clouds are still misclassified (e.g., table→desk and bottle→vase). This confirms that HHA is not limited to digital perturbations and can survive practical sources of distortion introduced by meshing, printing artifacts, and sensor noise, indicating strong physical-world transferability.

Computational Overhead. Tab. 5 reports the average runtime required to generate a single adversarial example on PointNet using ModelNet40. Although HHA incorporates hierarchical hyperbolic constraints, it maintains a runtime comparable to GeoA^3, demonstrating its practical efficiency.

Effect of Different Semantic Segmenters. To assess the impact of the segmentation backbone on HHA performance, we compare a stronger segmenter,

Table 7: Performance of HHA with different hyperbolic models, e.g., the Lorentz model (L) and the Poincaré ball model (P), evaluated on multiple classifiers on ModelNet40.

Model	CD (10^{-4})		HD (10^{-2})		SSC		GSC	
	L	P	L	P	L	P	L	P
PointNet	2.108	**0.603**	1.023	**0.279**	0.027	**0.011**	0.044	**0.023**
DGCNN	4.566	**2.189**	0.963	**0.278**	0.081	**0.040**	0.095	**0.061**
PTv1	5.136	**2.698**	1.205	**0.461**	0.063	**0.038**	0.106	**0.047**
PointMLP	3.982	**1.811**	1.004	**0.411**	0.084	**0.049**	0.107	**0.068**
Mamba3D	3.053	**1.321**	0.805	**0.289**	0.080	**0.054**	0.107	**0.063**

Table 8: Comparison of ASR for different methods in attacking PointNet on ModelNet40 and ScanObjectNN under various defense mechanisms.

	ModelNet40						ScanObjectNN					
Attack	-	SOR	SRS	DUP-Net	IF-Defense	AT	-	SOR	SRS	DUP-Net	IF-Defense	AT
IFGM	100	60.3	58.3	23.2	15.8	23.4	100	44.6	42.4	42.5	21.7	28.2
GeoA3	100	37.9	43.2	**56.7**	13.6	52.0	100	39.3	31.2	41.1	10.4	47.1
AOF	100	61.7	58.6	54.4	27.0	28.6	100	50.2	41.5	44.2	19.3	22.8
SI-Adv	100	47.1	**68.1**	33.1	38.4	30.5	100	42.5	24.8	39.7	49.7	45.5
HIT-Adv	100	32.8	31.9	34.8	8.5	79.8	100	47.6	**52.2**	46.3	12.1	80.3
LBC	100	11.7	18.4	43.3	45.2	81.4	100	18.3	37.5	46.7	50.2	72.6
Ours	100	**68.4**	62.3	55.1	**75.7**	**93.1**	100	**50.6**	30.4	**48.7**	**62.3**	**88.2**

P^3-SAM [27], with PartSLIP [25]. As shown in Tab. 6, P^3-SAM provides slightly more accurate part segmentation and yields marginal improvements in imperceptibility (lower CD and HD) when attacking PointNet on ScanObjectNN. These differences are minor, suggesting that HHA is largely insensitive to segmentation quality and remains stable even when using simpler or noisier semantic hierarchies.

Effect of Different Hyperbolic Models. We also investigate the impact of employing different hyperbolic models within HHA. The results in Tab. 7 show that HHA with the Poincaré ball model consistently achieves lower distortion than with the Lorentz model. This advantage is attributed to the conformal property of the Poincaré ball, which better preserves local geometric relations and enhances the structural consistency enforced by hierarchical regularization.

Robustness Under Defenses. We evaluate robustness against five representative defenses, including statistical outlier removal (SOR), simple random sampling (SRS), DUP-Net [70], IF-Defense [60], and adversarial training (AT) [24]. As shown in Tab. 8, our method consistently maintains higher ASR under these defenses on both ModelNet40 and ScanObjectNN. These results demonstrate that the proposed hierarchical hyperbolic perturbations are more resistant to common denoising-based and training-based defenses than prior attacks.

6 Conclusion

This paper has proposed HHA, a hierarchical hyperbolic constraint framework for generating imperceptible adversarial point clouds. By embedding point cloud substructures into hyperbolic space and imposing constraints at both semantic and geometric levels, the method effectively preserves multi-scale structural consistency while maintaining strong attack capability. Extensive experiments demonstrate the effectiveness of HHA in improving imperceptibility and superiority over state-of-the-art methods. Future work will explore defense-oriented applications of hierarchical hyperbolic constraints, including robustness enhancement and adversarial detection under realistic deployment settings.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (62472117, 62572400, U2436208, 62572150, 62372129), the Guangdong Basic and Applied Basic Research Foundation (2025A1515010157, 2024A1515012064, 2026B1515020033), the CCF-NetEase ThunderFire Innovation Research Funding (CCF-Netease 202514), the Science and Technology Projects in Guangzhou (2025A03J0137, 2024B0101010002), the Project of Guangdong Key Laboratory of Industrial Control System Security (2024B1212020010), and the High-Quality Talent Training Program – Graduate Student Development Project of the Graduate School, Guangzhou University.

References

1. Balazevic, I., Allen, C., Hospedales, T.: Multi-relational poincaré graph embeddings. In: NeurIPS. vol. 32 (2019)
2. Bi, J., Wu, Q., Qian, J., Luo, L., Yang, J.: Dual manifold regularization steered robust representation learning for point cloud analysis. In: AAAI. vol. 39, pp. 1844–1852 (2025)
3. Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. In: S&P. pp. 39–57 (2017)
4. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)
5. Dong, X., Chen, D., Zhou, H., Hua, G., Zhang, W., Yu, N.: Self-robust 3d point recognition via gather-vector guidance. In: CVPR. pp. 11513–11521 (2020)
6. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: CVPR. pp. 605–613 (2017)
7. Ganea, O., Bécigneul, G., Hofmann, T.: Hyperbolic neural networks. In: NeurIPS. vol. 31 (2018)
8. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: ICLR (2015)
9. Guo, C., Gardner, J., You, Y., Wilson, A.G., Weinberger, K.: Simple black-box adversarial attacks. In: ICML. pp. 2484–2493 (2019)

10. Han, X., Tang, Y., Wang, Z., Li, X.: Mamba3d: Enhancing local features for 3d point cloud analysis via state space model. In: ACM Multimedia. pp. 4995–5004 (2024)
11. Huang, Q., Dong, X., Chen, D., Zhou, H., Zhang, W., Yu, N.: Shape-invariant 3d adversarial point clouds. In: CVPR. pp. 15335–15344 (2022)
12. Keselman, L., Iselin Woodfill, J., Grunnet-Jepsen, A., Bhowmik, A.: Intel realsense stereoscopic depth cameras. In: CVPR Workshops. pp. 1–10 (2017)
13. Khrulkov, V., Mirvakhabova, L., Ustinova, E., Oseledets, I., Lempitsky, V.: Hyperbolic image embeddings. In: CVPR. pp. 6418–6428 (2020)
14. Kochurov, M., Karimov, R., Kozlukov, S.: Geoopt: Riemannian optimization in pytorch. arXiv preprint arXiv:2005.02819 (2020)
15. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. Nature **521**(7553), 436–444 (2015)
16. Lee, K., Chen, Z., Yan, X., Urtasun, R., Yumer, E.: Shapeadv: Generating shape-aware adversarial 3d point clouds. arXiv preprint arXiv:2005.11626 (2020)
17. Li, W., Yang, Z., Han, W., Man, H., Wang, X., Fan, X.: Hyperbolic-constraint point cloud reconstruction from single rgb-d images. In: AAAI. vol. 39, pp. 4959–4967 (2025)
18. Li, Y., Bu, R., Sun, M., Wu, W., Di, X., Chen, B.: Pointcnn: Convolution on χ-transformed points. In: NeurIPS. pp. 820–830 (2018)
19. Li, Z., Du, X., Lei, N., Chen, L., Wang, W.: Nopain: No-box point cloud attack via optimal transport singular boundary. In: CVPR. pp. 3492–3502 (2025)
20. Lin, F., Yue, Y., Hou, S., Yu, X., Xu, Y., Yamada, K.D., Zhang, Z.: Hyperbolic chamfer distance for point cloud completion. In: ICCV. pp. 14595–14606 (2023)
21. Liu, B., Zhang, J., Zhu, J.: Boosting 3d adversarial attacks with attacking on frequency. IEEE Access **10**, 50974–50984 (2022)
22. Liu, D., Hu, W.: Imperceptible transfer attack and defense on 3d point cloud classification. IEEE TPAMI **45**(4), 4727–4746 (2023). `https://doi.org/10.1109/TPAMI.2022.3193449`
23. Liu, D., Hu, W., Li, X.: Point cloud attacks in graph spectral domain: When 3d geometry meets graph signal processing. IEEE TPAMI **46**(5), 3079–3095 (2023)
24. Liu, D., Yu, R., Su, H.: Extending adversarial attacks and defenses to deep 3d point cloud classifiers. In: 2019 IEEE International Conference on Image Processing (ICIP). pp. 2279–2283. IEEE (2019)
25. Liu, M., Zhu, Y., Cai, H., Han, S., Ling, Z., Porikli, F., Su, H.: Partslip: Low-shot part segmentation for 3d point clouds via pretrained image-language models. In: CVPR. pp. 21736–21746 (2023)
26. 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: CVPR. pp. 24326–24335 (2024)
27. Ma, C., Li, Y., Yan, X., Xu, J., Yang, Y., Wang, C., Zhao, Z., Guo, Y., Chen, Z., Guo, C.: P3-sam: Native 3d part segmentation. arXiv preprint arXiv:2509.06784 (2025)
28. 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: ICLR (2022)
29. Maturana, D., Scherer, S.: Voxnet: A 3d convolutional neural network for real-time object recognition. In: IROS. pp. 922–928 (2015)
30. Montanaro, A., Valsesia, D., Magli, E.: Rethinking the compositionality of point clouds through regularization in the hyperbolic space. In: NeurIPS. vol. 35, pp. 33741–33753 (2022)

31. Nickel, M., Kiela, D.: Poincaré embeddings for learning hierarchical representations. In: NeurIPS. vol. 30 (2017)
32. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: NeurIPS. pp. 8026–8037 (2019)
33. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: CVPR. pp. 652–660 (2017)
34. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++ deep hierarchical feature learning on point sets in a metric space. In: NeurIPS. pp. 5105–5114 (2017)
35. Rubner, Y., Tomasi, C., Guibas, L.J.: The earth mover's distance as a metric for image retrieval. IJCV **40**, 99–121 (2000)
36. Sur, T., Mukherjee, S., Rahaman, K., Chaudhuri, S., Khan, M.H., Banerjee, B.: Hyperbolic uncertainty-aware few-shot incremental point cloud segmentation. In: CVPR. pp. 11810–11821 (2025)
37. Taha, A.A., Hanbury, A.: Metrics for evaluating 3d medical image segmentation: analysis, selection, and tool. BMC medical imaging **15**, 1–28 (2015)
38. Tang, K., Cao, Z., Peng, W., Wang, X., Zhu, P., Tian, Z.: Anchor field consistency for imperceptible adversarial attacks on 3d point clouds. In: ICASSP. pp. 13782–13786 (2026)
39. Tang, K., Chen, Y., Peng, W., Zhang, Y., Fang, M., Wang, Z., Song, P.: Reppvconv: attentively fusing reparameterized voxel features for efficient 3d point cloud perception. The Visual Computer **39**(11), 5577–5588 (2023)
40. 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: AAAI. vol. 39, pp. 20814–20822 (2025)
41. Tang, K., Du, Z., Peng, W., Wang, X., Zhu, P., Liu, L., Tian, Z.: Cage-based deformation for transferable and undefendable point cloud attack. arXiv preprint arXiv:2507.00690 (2025)
42. Tang, K., Gao, Y., Peng, W., Wang, X., Fang, M., Zhu, P.: Transferable and undefendable point cloud attacks via medial axis transform. Computer Aided Geometric Design p. 102515 (2026)
43. Tang, K., Hao, T., Wang, X., Peng, W., Zhang, D., Zhu, P., Tian, Z.: Less is more: Sparse and cooperative perturbation for point cloud attacks. In: AAAI. vol. 40, pp. 9430–9438 (2026)
44. Tang, K., He, X., Peng, W., Wu, J., Shi, Y., Liu, D., Zhou, P., Wang, W., Tian, Z.: Manifold constraints for imperceptible adversarial attacks on point clouds. In: AAAI. vol. 38, pp. 5127–5135 (2024)
45. Tang, K., Huang, L., Peng, W., Liu, D., Wang, X., Ma, Y., Liu, L., Tian, Z.: Flat: Flux-aware imperceptible adversarial attacks on 3d point clouds. In: ECCV. pp. 198–215 (2024)
46. Tang, K., Ke, W., Peng, W., Wang, X., Du, Z., Wu, Z., Zhu, P., Tian, Z.: Imperceptible adversarial attacks on point clouds guided by point-to-surface field. In: ICASSP. pp. 1–5 (2025)
47. 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 preprint arXiv:2601.23102 (2026)
48. Tang, K., Ma, Y., Miao, D., Song, P., Gu, Z., Tian, Z., Wang, W.: Decision fusion networks for image classification. TNNLS **36**(3), 3890–3903 (2025)

49. Tang, K., Shi, Y., Lou, T., Peng, W., He, X., Zhu, P., Gu, Z., Tian, Z.: Rethinking perturbation directions for imperceptible adversarial attacks on point clouds. IEEE Internet of Things Journal **10**(6), 5158–5169 (2022)
50. Tang, K., Wang, Z., Peng, W., Huang, L., Wang, L., Zhu, P., Wang, W., Tian, Z.: Symattack: Symmetry-aware imperceptible adversarial attacks on 3d point clouds. In: MM. pp. 3131–3140 (2024)
51. Tang, K., Wu, J., Peng, W., Shi, Y., Song, P., Gu, Z., Tian, Z., Wang, W.: Deep manifold attack on point clouds via parameter plane stretching. In: AAAI. vol. 37, pp. 2420–2428 (2023)
52. Uy, M.A., Pham, Q.H., Hua, B.S., Nguyen, T., Yeung, S.K.: Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In: ICCV. pp. 1588–1597 (2019)
53. Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph cnn for learning on point clouds. ACM TOG (Proc. of SIGGRAPH) **38**(5), 1–12 (2019)
54. Wang, Z., Peng, W., Wang, L., Wu, Z., Zhu, P., Tang, K.: Eia: Edge-aware imperceptible adversarial attacks on 3d point clouds. In: MMM (2025)
55. Wen, Y., Lin, J., Chen, K., Chen, C.P., Jia, K.: Geometry-aware generation of adversarial point clouds. IEEE TPAMI **44**(6), 2984–2999 (2022)
56. Wu, W., Qi, Z., Fuxin, L.: Pointconv: Deep convolutional networks on 3d point clouds. In: CVPR. pp. 9621–9630 (2019)
57. 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: CVPR. pp. 4840–4851 (2024)
58. Wu, X., Lao, Y., Jiang, L., Liu, X., Zhao, H.: Point transformer v2: Grouped vector attention and partition-based pooling. In: NeurIPS. vol. 35, pp. 33330–33342 (2022)
59. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: CVPR. pp. 1912–1920 (2015)
60. Wu, Z., Duan, Y., Wang, H., Fan, Q., Guibas, L.J.: If-defense: 3d adversarial point cloud defense via implicit function based restoration. arXiv preprint arXiv:2010.05272 (2020)
61. Xiang, C., Qi, C.R., Li, B.: Generating 3d adversarial point clouds. In: CVPR. pp. 9136–9144 (2019)
62. Xiong, B., Cochez, M., Nayyeri, M., Staab, S.: Hyperbolic embedding inference for structured multi-label prediction. In: NeurIPS. vol. 35, pp. 33016–33028 (2022)
63. Yang, J., Jiang, Y., Huang, X., Ni, B., Zhao, C.: Learning black-box attackers with transferable priors and query feedback. In: NeurIPS. vol. 33, pp. 12288–12299 (2020)
64. Yang, J., Zhang, Q., Fang, R., Ni, B., Liu, J., Tian, Q.: Adversarial attack and defense on point sets. arXiv preprint arXiv:1902.10899 (2019)
65. Zhang, J., Gu, W., Huang, Y., Jiang, Z., Wu, W., Lyu, M.R.: Curvature-invariant adversarial attacks for 3d point clouds. In: AAAI. vol. 38, pp. 7142–7150 (2024)
66. Zhao, H., Jiang, L., Fu, C.W., Jia, J.: Pointweb: Enhancing local neighborhood features for point cloud processing. In: CVPR. pp. 5565–5573 (2019)
67. Zhao, H., Jiang, L., Jia, J., Torr, P.H., Koltun, V.: Point transformer. In: ICCV. pp. 16259–16268 (2021)
68. Zheng, T., Chen, C., Yuan, J., Li, B., Ren, K.: Pointcloud saliency maps. In: ICCV. pp. 1598–1606 (2019)
69. Zhou, H., Chen, D., Liao, J., Chen, K., Dong, X., Liu, K., Zhang, W., Hua, G., Yu, N.: Lg-gan: Label guided adversarial network for flexible targeted attack of point cloud based deep networks. In: CVPR. pp. 10356–10365 (2020)

70. Zhou, H., Chen, K., Zhang, W., Fang, H., Zhou, W., Yu, N.: Dup-net: Denoiser and upsampler network for 3d adversarial point clouds defense. In: ICCV. pp. 1961–1970 (2019)