

Exploring the Devil in Graph Spectral Domain for 3D Point Cloud Attacks

Qianjiang Hu^{*}, Daizong Liu^{*}, and Wei Hu[✉]

Wangxuan Institute of Computer Technology, Peking University
No. 128, Zhongguancun North Street, Beijing, China
hqjpk@pku.edu.cn, dzliu@stu.pku.edu.cn, forhuwei@pku.edu.cn

Abstract. With the maturity of depth sensors, point clouds have received increasing attention in various applications such as autonomous driving, robotics, surveillance, *etc.*, while deep point cloud learning models have shown to be vulnerable to adversarial attacks. Existing attack methods generally add/delete points or perform point-wise perturbation over point clouds to generate adversarial examples in the data space, which may neglect the geometric characteristics of point clouds. Instead, we propose point cloud attacks from a new perspective—Graph Spectral Domain Attack (GSDA), aiming to perturb transform coefficients in the graph spectral domain that corresponds to varying certain geometric structure. In particular, we naturally represent a point cloud over a graph, and adaptively transform the coordinates of points into the graph spectral domain via graph Fourier transform (GFT) for compact representation. We then analyze the influence of different spectral bands on the geometric structure of the point cloud, based on which we propose to perturb the GFT coefficients in a learnable manner guided by an energy constraint loss function. Finally, the adversarial point cloud is generated by transforming the perturbed spectral representation back to the data domain via the inverse GFT (IGFT). Experimental results demonstrate the effectiveness of the proposed GSDA in terms of both imperceptibility and attack success rates under a variety of defense strategies. The code is available at <https://github.com/WoodwindHu/GSDA>.

Keywords: Point Cloud, Adversarial Attack, Graph Spectral Domain

1 Introduction

Deep Neural Networks (DNNs) are known to be vulnerable to adversarial examples [43,12], which are indistinguishable from legitimate ones by adding trivial perturbations, but lead to incorrect model prediction. Many efforts have been made into the attacks on the 2D image field [8,30,22,46], which often add imperceptible pixel-wise noise onto images to deceive the DNNs. Nevertheless, adversarial attacks on 3D point clouds—discrete representations of 3D scenes or

^{*} Q. Hu and D. Liu contributed equally to this work. [✉] Corresponding author: W. Hu. This work was supported by National Natural Science Foundation of China under Contract No. 61972009.

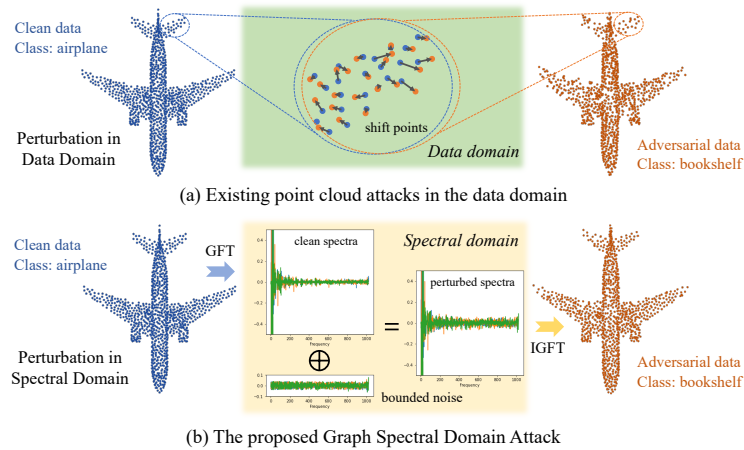


Fig. 1. (a) Existing point cloud attacks generally perturb point coordinates by shifting points in the data domain. (b) We explore perturbation in the graph spectral domain, leading to more imperceptible and effective adversarial examples. This is enlightened by spectral characteristics: most energies are concentrated in low-frequency components that represent the rough shape of point clouds while high-frequency components encode fine-grained details. When attacked with bounded spectral noise, spectra of the point cloud keeps similar patterns with the clean one, thus preserving geometric structures.

objects that consist of a set of points residing on irregular domains—are still relatively under-explored, which are however crucial in various safety-critical applications such as autonomous driving [4] and medical data analysis [41].

Existing 3D point cloud attacks [53,50,61,65,45,64,66,14] are all developed in the *data space*. Some of them [53,61,50,65] employ the gradient search method to identify critical points from point clouds and modify (add or delete) them to distort the most representative features for misclassification. Recently, more works [49,45,2,14,26,29,63] follow the C&W framework [12] to learn to perturb xyz coordinates of each point by gradient optimization in an end-to-end manner. Although the above two types of works achieve high attack success rates, the perturbed point clouds are often easily perceivable to humans, such as outliers and uneven distribution. This is because preserving geometric characteristics of point clouds is generally not considered yet in these methods.

To this end, we propose point cloud attacks from a new perspective—Graph Spectral Domain Attack (GSDA), aiming to exploit the elegant characterization of geometric structures in the graph spectral domain and thereby perturb graph transform coefficients that explicitly varies certain geometric structure. On the one hand, we provide graph spectral analysis of point clouds, which shows that the rough shape of point clouds is represented by low-frequency components while the fine details of objects are encoded in high-frequency components in general. With trivial perturbation in the spectral domain, the point cloud could retain the original rough shape with similar local details, as shown in Figure 1. On the other hand, the spectral characteristics of point clouds represent higher-

level and global information than point-to-point relations leveraged in previous works. That is, the spectral representation encodes more abstract and essential contexts for recognizing the point cloud.

Based on the above analysis, we develop a novel paradigm of Graph Spectral Domain Attack (GSDA) for point cloud attacks. In particular, different from images that are sampled on regular grids and typically transformed in the Discrete Cosine Transform (DCT) domain, point clouds reside on irregular domains with no ordering of points. Hence, we represent a point cloud over a graph naturally and adaptively, where each point is treated as a vertex and connected to its K nearest neighbors, and the coordinates of each point serve as the graph signal. Then, we compute the graph Laplacian matrix [39] that encodes the edge connectivity and vertex degree of the graph, whose eigenvectors form the basis of the Graph Fourier Transform (GFT) [15]. Because of the compact representation of point clouds in the GFT domain [19], we transform the coordinate signal of point clouds onto the spectral domain via the GFT, leading to transform coefficients corresponding to each spectral component. Next, we develop a learnable spectrum-aware perturbation approach to perturb the spectral domain with adversarial noise. In order to keep the energy balance among the whole frequency bands, we design an energy constraint function to restrict the perturbation size. Finally, we transform the perturbed GFT coefficients back to the data domain via the inverse GFT (IGFT) to produce the crafted point cloud. We iteratively optimize the adversarial loss function, and perform back-propagation to retrieve the gradient in the spectral domain for generating and updating the desired spectrum-aware perturbations. The point clouds reconstructed by the IGFT are taken as the adversarial examples.

Our main contributions are summarized as follows:

- We propose a novel paradigm of point cloud attacks—Graph Spectral Domain Attack (GSDA), which perturbs point clouds in the graph spectral domain to exploit the high-level spectral characterization of geometric structures. Such spectral approach marks the first significant departure from the current practice of point cloud attacks.
- We provide in-depth graph spectral analysis of point clouds, which enlightens our formulation of exploring destructive perturbation on appropriate frequency components. Based on this, we develop a learnable spectrum-aware perturbation approach to attack 3D models in an end-to-end manner.
- Extensive experiments show that the proposed GSDA achieves 100% of attack success rates in both targeted and untargeted settings with the least required perturbation size. We also demonstrate the imperceptibility of the GSDA compared to state-of-the-arts, as well as the robustness of the GSDA by attacking existing defense methods and implementing transfer-based attacks.

2 Related Works

3D Point Cloud Classification. Deep 3D point cloud learning has emerged in recent years, which has diverse applications in many fields, such as 3D object

classification [42,58], 3D scene segmentation [13,47,54], and 3D object detection in autonomous driving [4,56]. Among them, 3D object classification is the most fundamental yet important task, which learns representative information including both local details and global context of point clouds. Early works attempt to classify point clouds by adapting deep learning models in the 2D space [42,58]. In order to directly learn the 3D structure and address the unorderliness problem of point clouds, pioneering methods DeepSets [59] and PointNet [33] propose to achieve end-to-end learning on point cloud classification by formulating a general specification. PointNet++ [34] and other extensive works [9,27,57] are built upon PointNet to further capture the fine local structural information from the neighborhood of each point. Recently, some works focus on either designing special convolutions on the 3D domain [24,1,44,28] or developing graph neural networks [40,38,48,54,11] to improve point cloud learning. In this paper, we focus on PointNet [33], PointNet++ [34] and DGCNN [48] since these 3D models are extensively involved with practical 3D applications.

Adversarial Attack on 3D Point Clouds. Deep neural networks are vulnerable to adversarial examples, which has been extensively explored in the 2D image domain [32,31]. Recently, many works [53,50,61,65,45,64,66,14,25] adapt adversarial attacks into the 3D vision community, which can be divided into two categories: 1) point adding/dropping attack: Xiang *et al.* [53] proposed point generation attacks by adding a limited number of synthesized points/clusters/objects to a point cloud. Recently, more works [61,50,65] utilize gradient-guided attack methods to identify critical points from point clouds for deletion. 2) point perturbation attack: Previous point-wise perturbation attacks [49,45] learn to perturb xyz coordinates of each point by adopting the C&W framework [2] based on the Chamfer and Hausdorff distance with additional consideration of the benign distribution of points. Later works [14,26,29,63] further apply the iterative gradient method to achieve more fine-grained adversarial perturbation.

Spectral Methods for Point Clouds In the graph spectral domain, the rough shape of a point cloud will be encoded into low-frequency components, which is suitable for denoising point clouds. Rosman *et al.* proposed spectral point cloud denoising based on the non-local framework [36]. They group similar surface patches into a collaborative patch and perform shrinkage in the GFT domain by a low-pass filter, which leads to denoising of the 3D shape. Zhang *et al.* proposed a tensor-based method to estimate hypergraph spectral components and frequency coefficients of point clouds, which can be used to denoise 3D shapes [62]. On the contrary, high-frequency components often represent fine details of point clouds, which can be used to detect contours or process redundant information. Chen *et al.* proposed a high-pass filtering-based resampling method to highlight contours for large-scale point cloud visualization and extract key points for accurate 3D registration [3]. Sameera *et al.* proposed Spectral-GANs to generate high-resolution 3D point clouds, which takes the advantage of spectral representations for compact representation [35].

3 Graph Spectral Analysis for Point Clouds

In this section, we provide graph spectral analysis for point clouds, which lays the foundation for the proposed graph spectral domain attack in Sec. 4. We first discuss the benefit of graph spectral representations of point clouds and introduce how to transform point clouds onto the spectral domain in Sec. 3.1. Then, we analyze the roles of different spectral components with respect to geometric structures in Sec. 3.2.

3.1 Spectral Representations of Point Clouds

Signals can be compactly represented in the spectral domain, provided that the transformation basis characterizes principle components of the signals. For instance, images are often transformed onto the Discrete Cosine Transform (DCT) domain for compression and processing [5,23]. Different from images supported on regular grids, point clouds reside on irregular domains with no ordering of points, which hinders the deployment of traditional transformations such as the DCT. Though we may quantize point clouds onto regular voxel grids or project onto a set of depth images from multiple viewpoints, this would inevitably introduce quantization loss. Instead, graphs serve as a natural representation for irregular point clouds, which is accurate and structure-adaptive [19]. With an appropriately constructed graph that well captures the underlying structure, the Graph Fourier Transform (GFT) will lead to a compact representation of geometric data including point clouds in the spectral domain [37,16,18,17,60,55], which inspires new insights and understanding of point cloud attacks.

Formally, we represent a point cloud $\mathbf{P} = \{\mathbf{p}_i\}_{i=1}^n \in \mathbb{R}^{n \times 3}$ consisting of n points over a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}, \mathbf{A}\}$, which is composed of a vertex set \mathcal{V} of cardinality $|\mathcal{V}| = n$ representing points, an edge set \mathcal{E} connecting vertices, and an adjacency matrix \mathbf{A} . Each entry $a_{i,j}$ in \mathbf{A} represents the weight of the edge between vertices i and j , which often captures the similarity between adjacent vertices. Here, we construct an unweighted K -nearest-neighbor graph (K -NN graph), where each vertex is connected to its K nearest neighbors in terms of the Euclidean distance with weight 1. The coordinates of points in \mathbf{P} are treated as graph signals.

Prior to the introduction of the GFT, we first define the combinatorial graph Laplacian matrix [39] as $\mathbf{L} := \mathbf{D} - \mathbf{A}$, where \mathbf{D} is the *degree matrix*—a diagonal matrix where $d_{i,i} = \sum_{j=1}^n a_{i,j}$. Given real and non-negative edge weights in an undirected graph, \mathbf{L} is real, symmetric, and positive semi-definite [6]. Hence, it admits an eigen-decomposition $\mathbf{L} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$, where $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_n]$ is an orthonormal matrix containing the eigenvectors \mathbf{u}_i , and $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n)$ consists of eigenvalues $\{\lambda_1 = 0 \leq \lambda_2 \leq \dots \leq \lambda_n\}$. We refer to the eigenvalues as the *graph frequency/spectrum*, with a smaller eigenvalue corresponding to a lower graph frequency.

For any graph signal $\mathbf{x} \in \mathbb{R}^n$ residing on the vertices of \mathcal{G} , its GFT coefficient vector $\hat{\mathbf{x}} \in \mathbb{R}^n$ is defined as [15]:

$$\hat{\mathbf{x}} = \phi_{\text{GFT}}(\mathbf{x}) = \mathbf{U}^\top \mathbf{x}. \quad (1)$$

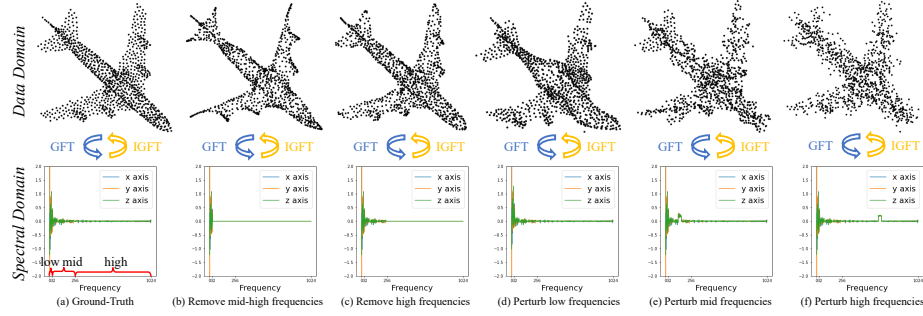


Fig. 2. Graph spectral analysis for 3D point clouds. We take an example from the *airplane* object to investigate the roles of difference frequency bands in the graph spectral domain. (a) Ground-Truth; (b) Remove mid-high frequencies; (c) Remove high frequencies; (d) Perturb low frequencies; (e) Perturb mid frequencies; (f) Perturb high frequencies.

The inverse GFT (IGFT) follows as:

$$\mathbf{x} = \phi_{\text{IGFT}}(\hat{\mathbf{x}}) = \mathbf{U}\hat{\mathbf{x}}. \quad (2)$$

Since \mathbf{U} is an orthonormal matrix, both GFT and IGFT operations are lossless.

3.2 Analysis in Graph Spectral Domain

When an appropriate graph is constructed that captures the geometric structure of point clouds well, the low-frequency components of the corresponding GFT characterize the *rough shape* of point clouds, while the high-frequency components represent *fine details or noise* (i.e., large variations such as geometric contours) in general.

To analyze such characteristics, we provide a toy experiment to investigate the roles of difference frequency bands in the graph spectral domain. We randomly take a clean point cloud *airplane* from the ModelNet40 dataset [51], and sample it into 1024 points as an example point cloud \mathbf{P} . We construct a K -NN graph with $K = 10$ over the point cloud, and perform the GFT on the three coordinate signals of each point in \mathbf{P} . The resulting transform coefficient vectors $\phi_{\text{GFT}}(\mathbf{P})$ are presented in Figure 2(a). We see that, $\phi_{\text{GFT}}(\mathbf{P})$ has larger amplitudes at lower-frequency components and much smaller amplitudes at higher-frequency components, demonstrating that most information is concentrated in low-frequency components.

To further investigate how each frequency band contributes to the geometric structure of the point cloud in the data domain, we first divide the whole spectral domain into three bands: low-frequency band $B_l \in [0, \lambda_l)$, mid-frequency band $B_m \in [\lambda_l, \lambda_h)$, and high-frequency band $B_h \in [\lambda_h, \lambda_{\max}]$, where $\lambda_l, \lambda_h, \lambda_{\max}$ bound the three bands. While the division of frequency bands is not established, we propose an appropriate division by the distribution of energy—squared sum of transform coefficients. In this example, we compute that the lowest 32 frequencies account for almost 90% of energy, while the lowest 256 frequencies account for

almost 97% of energy. Based on this observation, we set $\lambda_l = \lambda_{32} = 1.19$ and $\lambda_h = \lambda_{256} = 13.93$ in this instance.

Next, we study the influence of each frequency band on the geometric structure by removing different bands. As shown in Figure 2(b), when GFT coefficients in both mid- and high-frequency bands are assigned 0, the point cloud reconstructed with only low-frequency components exhibits the rough shape of the original object. By adding more information from the mid-frequency band, the reconstructed point cloud has richer local contexts in Figure 2(c), but still lacks fine-grained details such as the engines of the airplane. To summarize, each frequency band is crucial to represent different aspects of the geometric structure of a point cloud.

Then here is a key question: what are the results of attacking different frequency bands? We investigate into this by separately perturbing each frequency band with a large perturbation size, *i.e.*, we perturb consecutive 32 frequencies in each frequency band by adding noise of 0.2 on each frequency. As shown in Figure 2(d-f), attacking low-frequency components introduces deformation in the coarse shape but remains smoothness of the surface. When attacking mid-frequency components, the shape of the object becomes much rougher. In comparison, perturbing high-frequency components loses local details and induces noise and outliers, though the silhouette of the shape is preserved to some degree thanks to the clean lower-frequency components.

Inspired by the above properties of point clouds in the graph spectral domain, we summarize several insights for developing effective spectral domain attacks:

- Each frequency band represents the geometric structure of the point cloud from different perspectives (*e.g.*, low-frequency components represent the basic shape while high-frequency components encode fine-grained details). Perturbing only one frequency band may result in the corresponding distortion in the data domain.
- Although a large perturbation size can ensure a high success attack rate, it may severely change the spectral characteristics, thus leading to perceptible deformations in the data domain.

Based on the above insights, a desirable spectral domain attack for point clouds need to not only perform the perturbation among appropriate frequencies for striking a balance, but also restrict the perturbation size for preserving the original spectral characteristics.

4 Graph Spectral Domain Attack

4.1 The Formulation

Given a clean point cloud $\mathbf{P} = \{\mathbf{p}_i\}_{i=1}^n \in \mathbb{R}^{n \times 3}$ where each point $\mathbf{p}_i \in \mathbb{R}^3$ is a vector that contains the (x, y, z) coordinates, a well-trained classifier $f(\cdot)$ can predict its accurate label $y = f(\mathbf{P}) \in \mathbb{Y}$, $\mathbb{Y} = \{1, 2, 3, \dots, c\}$ that represents its object class best, where c is the number of classes. The goal of point cloud attacks on classification is to perturb point cloud \mathbf{P} into an adversarial one \mathbf{P}' , so that $f(\mathbf{P}') = y'$ (targeted attack) or $f(\mathbf{P}') \neq y$ (untargeted attack), where $y' \in \mathbb{Y}$ but $y' \neq y$.




Fig. 3. The overall pipeline of the proposed GSDA. Given a clean point cloud, we first construct a K -NN graph and transform the point cloud into the spectral domain via the GFT. Then, we perturb the GFT coefficients in a learnable manner with a specifically designed restrictive function over the spectral energy. Subsequently, we transform the perturbed spectral signals back to the data domain via the IGFT. Finally, we take the reconstructed point cloud as the adversarial example and feed it into the target 3D model for attack.

We propose a novel GSDA attack that aims to learn destructive yet imperceptible perturbations in the spectral domain for generating adversarial point clouds. Formally, we formulate the proposed GSDA as the following optimization problem:

$$\begin{aligned} \min_{\Delta} \mathcal{L}_{adv}(\mathbf{P}', \mathbf{P}, y), \text{ s.t. } \|\phi_{GFT}(\mathbf{P}') - \phi_{GFT}(\mathbf{P})\|_p < \epsilon, \\ \text{where } \mathbf{P}' = \phi_{IGFT}(\phi_{GFT}(\mathbf{P}) + \Delta), \end{aligned} \quad (3)$$

where $\mathcal{L}_{adv}(\mathbf{P}', \mathbf{P}, y)$ is the adversarial loss and Δ is the perturbation. In the imposed constraint, ϵ aims to restrict the perturbation size in the spectral domain, which preserves the original spectral characteristics so that the resultant adversarial point cloud \mathbf{P}' is visually indistinguishable from its clean version \mathbf{P} . We adopt the l_2 -norm in this equation.

To back-propagate the gradient in a desired direction for optimizing the perturbation learning, we define our adversarial loss $\mathcal{L}_{adv}(\mathbf{P}', \mathbf{P}, y)$ as follows:

$$\mathcal{L}_{adv}(\mathbf{P}', \mathbf{P}, y) = \mathcal{L}_{class}(\mathbf{P}', y) + \beta \cdot \mathcal{L}_{reg}(\mathbf{P}', \mathbf{P}), \quad (4)$$

where $\mathcal{L}_{class}(\mathbf{P}', y)$ is to promote the misclassification of point cloud \mathbf{P}' . $\mathcal{L}_{reg}(\mathbf{P}', \mathbf{P})$ is a regularization term that minimizes the distance between \mathbf{P}' and \mathbf{P} to guide perturbation in appropriate frequencies. β is a penalty parameter controlling the regularization term.

Specifically, $\mathcal{L}_{class}(\mathbf{P}', y)$ is formulated as a cross-entropy loss as follows:

$$\mathcal{L}_{class}(\mathbf{P}', y) = \begin{cases} -\log(p_{y'}(\mathbf{P}')), & \text{for targeted attack,} \\ \log(p_y(\mathbf{P}')), & \text{for untargeted attack,} \end{cases} \quad (5)$$

where $p(\cdot)$ is the softmax functioned on the output of the target model, *i.e.*, the probability with respect to adversarial class y' or clean class y . By minimizing this loss function, our GSDA optimizes the spectral perturbation Δ to mislead the target model $f(\cdot)$.

Besides, to strike a balance of perturbing different frequency bands, we utilize both Chamfer distance loss [10] and Hausdorff distance loss [20] as the

$\mathcal{L}_{reg}(\mathbf{P}', \mathbf{P})$ function. This reflects the imperceptibility in the data domain for updating the frequency perturbation Δ .

4.2 The Algorithm

Based on the problem formulation, we develop an efficient and effective algorithm for the GSDA model. As shown in Figure 3, the proposed GSDA attack is composed of four steps: Firstly, the GSDA transforms the clean point cloud \mathbf{P} from the data domain to the graph spectral domain via the GFT operation ϕ_{GFT} . Then, the GSDA perturbs the GFT coefficients through our designed perturbation strategy imposed with an energy constraint function as in Eq. (3). Next, we convert the perturbed spectral signals back to the data domain via the IGFT operation ϕ_{IGFT} for constructing the adversarial point cloud \mathbf{P}' . Finally, we optimize the adversarial loss function to iteratively update the desired perturbations added in the spectral domain. In the following, we elaborate on each module in order.

Transform onto Spectral Domain. Given a clean point cloud \mathbf{P} , we employ the GFT to transform \mathbf{P} onto the graph spectral domain. Specifically, we first construct a K -NN graph on the whole point cloud, and then compute the graph Laplacian matrix \mathbf{L} . Next, we perform eigen-decomposition to acquire the orthonormal eigenvector matrix \mathbf{U} , which serves as the GFT basis. The GFT coefficients $\phi_{\text{GFT}}(\mathbf{P})$ is then obtained by:

$$\phi_{\text{GFT}}(\mathbf{P}) = \mathbf{U}^\top \mathbf{P}, \quad (6)$$

where $\phi_{\text{GFT}}(\mathbf{P}) \in \mathbb{R}^{n \times 3}$ corresponds to the transform coefficients of the x-, y-, z-coordinate signals.

Perturbation in the Graph Spectral Domain. We deploy a trainable perturbation Δ to perturb the spectral representation of \mathbf{P} : $\phi_{\text{GFT}}(\mathbf{P}) + \Delta$. Further, to restrict the perturbation size among appropriate frequencies for enhancing the imperceptibility, instead of the constraint of the entire energy in Eq. (3), we constrain the perturbation Δ on each frequency in a valid range $[\epsilon_{\min}, \epsilon_{\max}]$ in order to keep the trend of spectral characteristics, *i.e.*, low-frequency components with large magnitudes and high-frequency components with small magnitudes. In particular, we define the perturbation size of each frequency as:

$$\Delta / \phi_{\text{GFT}}(\mathbf{P}) \in [\epsilon_{\min}, \epsilon_{\max}], \quad (7)$$

which maintains a certain ratio to each frequency with $\epsilon_{\max} = -\epsilon_{\min}$. In order to adjust the perturbation Δ adaptively and further improve the success rate of the proposed spectral attack, we formulate the attack process as an optimization problem by leveraging the gradients of the target 3D model $f(\cdot)$ via backward-propagation. We update the perturbation Δ with the gradients, and learn the perturbation Δ as:

$$\begin{aligned} \Delta' &\leftarrow \Delta - lr \cdot \partial_{\Delta}(\mathcal{L}_{adv}(\mathbf{P}', \mathbf{P}, y)), \\ \text{s.t. } \Delta / \phi_{\text{GFT}}(\mathbf{P}) &\in [\epsilon_{\min}, \epsilon_{\max}], \end{aligned} \quad (8)$$

where lr is the learning rate.

Inverse Transform onto Data Domain. After obtaining the perturbed spectral representations, we apply the IGFT to convert the perturbed signals from the spectral domain back to the data domain as:

$$\mathbf{P}' = \phi_{\text{IGFT}}(\phi_{\text{GFT}}(\mathbf{P}) + \mathbf{\Delta}) = \mathbf{U}(\phi_{\text{GFT}}(\mathbf{P}) + \mathbf{\Delta}), \quad (9)$$

where \mathbf{P}' is the crafted adversarial point cloud.

5 Experiments

5.1 Dataset and 3D Models

Dataset. We adopt the point cloud benchmark ModelNet40 [51] dataset in all the experiments. This dataset contains 12,311 CAD models from 40 most common object categories in the world. Among them, 9,843 objects are used for training and the other 2,468 for testing. As in previous works [33], we uniformly sample $n = 1,024$ points from the surface of each object, and re-scale them into a unit ball. For adversarial point cloud attacks, we follow [53,49] and randomly select 25 instances for each of 10 object categories in the ModelNet40 testing set, which can be well classified by the classifiers of interest.

3D Models. We select three commonly used networks in 3D computer vision community as the victim models, *i.e.*, PointNet [33], PointNet++ [34], and DGCNN [48]. We train them from scratch, and the test accuracy of each trained model is within 0.1% of the best reported accuracy in their original papers. We generate the adversarial point clouds on each of them, and further explore their transferability among these three models.

5.2 Implementation Details

Experimental Settings. For generating the adversarial examples, we update the frequency perturbation $\mathbf{\Delta}$ with 500 iterations. We use Adam optimizer [21] to optimize the objective of our proposed GSDA attack in Eq. (3) with a fixed learning rate $lr = 0.01$, and the momentum is set as 0.9. We set $K = 10$ to build a K -NN graph. The penalty β in Eq. (4) is initialized as 10 and adjusted by 10 runs of binary search [30]. The weights of Chamfer distance loss [10] and Hausdorff distance loss [20] in the regularization term are set to 5.0 and 0.5, respectively. Since the targeted attack is more challenging than the untargeted attack, we focus on the targeted attack in the experiments. All experiments are implemented on a single NVIDIA RTX 2080Ti GPU.

Evaluation Metrics. To quantitatively evaluate the effectiveness of our proposed GSDA attack, we measure by the attack success rate, which is the ratio of successfully fooling a 3D model. Besides, to measure the perturbation size of different attackers, we adopt four evaluation metrics: (1) Data domain: l_2 -norm distance $\mathcal{D}_{\text{norm}}$ [7] which measures the square root of the sum of squared shifting distance, Chamfer distance \mathcal{D}_c [10] which measures the average squared distance between each adversarial point and its nearest original point, Hausdorff distance \mathcal{D}_h [20] which measures the maximum squared distance between each adversarial point and its nearest original point and is thus sensitive to outliers; (2) Spectral domain: perturbed energy $\mathcal{E}_{\Delta} = \|\phi_{\text{GFT}}(\mathbf{P}') - \phi_{\text{GFT}}(\mathbf{P})\|_2$.

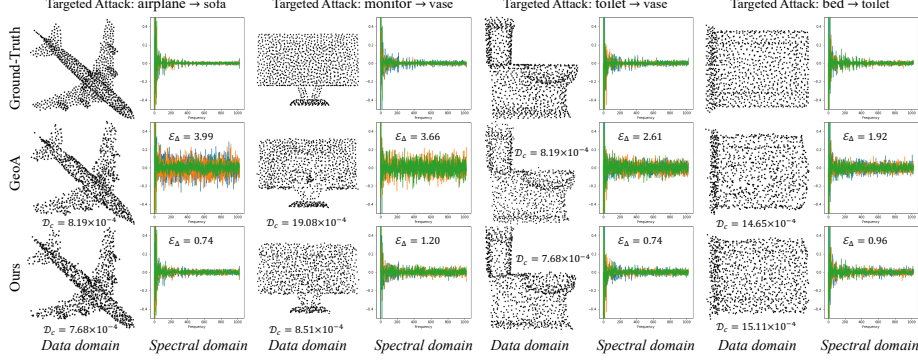


Fig. 4. Visualization of the generated adversarial examples in both the data and spectral domains. Specifically, we compare our GSDA with GeoA via evaluation metrics of the perturbation budget \mathcal{D}_c in the data domain and the perturbed energy \mathcal{E}_Δ in the spectral domain.

5.3 Evaluation on Our GSDA Attack

Table 1. Comparative results on the perturbation sizes of different methods in the data domain for adversarial point clouds. The results of 3D-ADV are borrowed from its original paper. ^p, ^c, ^o denote the variant of adversarial point, adversarial cluster and adversarial object, respectively.

Attack Model	Methods	Success Rate	Perturbation Size		
			\mathcal{D}_{norm}	\mathcal{D}_c	\mathcal{D}_h
PointNet	FGSM	100%	0.7936	0.1326	0.1853
	3D-ADV ^p	100%	0.3032	0.0003	0.0105
	3D-ADV ^c	92.1%	-	0.1652	-
	3D-ADV ^o	81.9%	-	0.1321	-
	GeoA	100%	0.4385	0.0064	0.0175
	Ours	100%	0.1741	0.0007	0.0031
PointNet++	FGSM	100%	0.8357	0.1682	0.2275
	3D-ADV ^p	100%	0.3248	0.0005	0.0381
	GeoA	100%	0.4772	0.0198	0.0357
	Ours	100%	0.2072	0.0081	0.0248
DGCNN	FGSM	100%	0.8549	0.189	0.2506
	3D-ADV ^p	100%	0.3326	0.0005	0.0475
	GeoA	100%	0.4933	0.0176	0.0402
	Ours	100%	0.2160	0.0104	0.1401

Quantitative results. In order to fairly compare our GSDA attack with existing methods, we perform four adversarial attacks, namely FGSM [61], 3D-ADV [53], GeoA [49] and ours, and measure the perturbation in the data domain with three evaluation metrics when these methods reach 100% of attack success rate. Specifically, we implement these attacks on three 3D models PointNet, PointNet++ and DGCNN. Corresponding results are shown in Table 1. We see that, our GSDA generates adversarial point clouds with almost the lowest perturbation sizes in all evaluation metrics on three attack models.

Note that, as \mathcal{D}_c measures the average squared distance between each adversarial point and its nearest original point, attacking by adding a few points in 3D-ADV^p has a natural advantage in terms of \mathcal{D}_c because most of the distance is equal to 0. However, it induces much larger distortions than ours on the other two metrics.

Table 2. Defense by dropping different ratios of points via SOR.

Method	Attack success rate (%) defense via SOR						
	0%	1%	2%	5%	10%	15%	20%
GeoA	100	83.47	70.56	52.61	31.58	18.62	11.71
Ours	100	91.87	89.91	85.38	72.53	50.00	27.78

Table 3. The attack success rate (%) on PointNet model by various attacks under defense.

Attack	No Defense	SRS	DUP-Net	IF-Defense
FGSM	100%	9.68%	4.38%	4.80%
3D-ADV	100%	22.53%	15.44%	13.70%
GeoA	100%	67.61%	59.15%	38.72%
Ours	100%	81.03%	68.98%	50.26%

Overall, this demonstrates that our generated point clouds are less distorted quantitatively. Besides, for each attack method, it takes larger perturbation sizes to successfully attack PointNet++ and DGCNN than to attack PointNet, which indicates that PointNet++ and DGCNN are harder to attack.

Visualization results. We also provide some visualization results of generated adversarial point clouds and corresponding spectral coefficients for comparison. As shown in Figure 4, we compare the visualization of both GeoA and ours on four examples under the setting of targeted attacks. Since GeoA implements perturbation in the data domain while we conduct it in the spectral domain, we provide the perturbation budgets in both two domains for fair comparison. Figure 4 shows that our GSDA attack has less perturbation \mathcal{E}_Δ than GeoA in the spectral domain as we develop an energy constraint function for optimization. By reflecting the perturbation in the data domain, our adversarial examples are more imperceptible than those of GeoA in both local details and distributions. Quantitatively, we achieve smaller perturbation budget \mathcal{D}_c in the data domain. To conclude, our GSDA attack is more effective and imperceptible.

5.4 Analysis on Robustness of Our GSDA Attack

Attacking the Defenses. To further examine the robustness of our proposed GSDA attack, we employ several 3D defenses to investigate whether our attack is still effective. Specifically, we employ the PointNet model with the following defense methods: Statistical Outlier Removal (SOR) [67], Simple Random Sampling (SRS) [61], DUP-Net defense [67] and IF-Defense [52]. Table 2 shows that across a range of dropping ratios via the defense SOR, the attack success rates drop a lot. However, the performance of our GSDA attack decays much slower than that of GeoA, validating that our attack is much more robust. We also report the results with other defenses in Table 3. We observe that FGSM and 3D-ADV attacks have low success rates under all the defenses, which is because they often lead to uneven local distribution and outliers. Besides, GeoA achieves relatively higher attack success rates, since it utilizes a geometry-aware loss function to constrain the similarity in curvature and thus has fewer outliers. In comparison, our attack achieves the highest success rates than all other attacks under all defenses since the trivial perturbation in the spectral domain reflects less noise in the data domain, thus enhancing the robustness.

Transferability of Adversarial Point Clouds. To investigate the transferability of our proposed GSDA attack, we craft adversarial point clouds on normally trained models and test them on all the three 3D models we consider. The success rates, which are the misclassification rates of the corresponding models on adversarial examples, are shown in Table 4. The left, middle and right parts present the three models we attack: PointNet, PointNet++ and DGCNN; the

Table 4. The attack success rate (%) of transfer-based attacks.

Attacks	PointNet	PointNet++	DGCNN	PointNet	PointNet++	DGCNN	PointNet	PointNet++	DGCNN
FGSM	100%	3.99%	0.63%	3.16%	100%	5.57%	3.59%	7.21%	100%
3D-ADV	100%	8.45%	1.28%	6.63%	100%	10.98%	6.82%	13.53%	100%
GeoA	100%	11.59%	2.59%	9.47%	100%	19.77%	12.46%	24.24%	100%
Ours	100%	11.51%	8.39%	10.89%	100%	30.84%	32.31%	84.49%	100%

columns of the table present the models we test. We observe that although our GSDA attack is not tailored for transferability, it has relatively higher success rates of transfer-based attacks than others. This is because our perturbation in the spectral domain not only keeps spectral characteristics but also reflects trivial noise in the data domain, which is thus more robust. A further improvement of the transferability could be developed by employing an additional adversarial learning mechanism with the Auto-Encoder reconstruction as in [14] (left as the future works). The comparatively slightly lower transferability of our GSDA in this table may be related with special properties of certain 3D models. To summarize, this intrinsic property makes it possible to design black-box defense against such adversarial instances.

5.5 Ablation Study

Table 5. Sensitivity analysis on the number K .

Number K	Success Rate	\mathcal{D}_c	\mathcal{D}_h	\mathcal{E}_Δ
$K=5$	100%	0.0007	0.0031	3.4679
$K=10$	100%	0.0007	0.0031	3.4510
$K=20$	100%	0.0007	0.0030	3.4617
$K=40$	100%	0.0007	0.0030	3.4688

Table 6. Comparison of different spectral representations.

Method	Attack success rate (%)				\mathcal{D}_c	\mathcal{D}_h	\mathcal{E}_Δ
	defense via SOR						
	0%	5%	10%	20%			
Ours-1DDCT	100.00	6.84	3.24	0.62	0.0006	0.0037	28.7834
Ours-3DDCT	93.06	37.48	25.17	19.60	0.0029	0.0415	19.4804
Ours-GFT	100.00	85.38	72.53	27.78	0.0007	0.0031	3.4510

Choice of Spectral Representation. We first evaluate the benefit of spectral representation in the GFT domain. Since the DCT is widely used in the 2D field to transform images onto the spectral domain, we compare the spectral domain attack performances on different spectral representations of the DCT and GFT. We implement 1D DCT and 3D DCT for comparison. For 1D DCT, we regard coordinates as the signal and perform 1D DCT on it. As to 3D DCT, we voxelize the pointcloud and perform 3D DCT on coordinates signal. As shown in Table 6, we see that our attack with the GFT is more robust to SOR defense and has much lower perturbation size in the spectral domain than the attack with the DCT. The main reason is that point clouds are unordered, which makes it challenging to capture the correlations among points by the DCT. In contrast, the GFT well captures the underlying structure of point clouds via the appropriate graph construction.

Sensitivity on the Number K . As shown in Table 5, we investigate whether the adversarial effects vary with respect to different settings of the number K in the K -NN graph. Specifically, we implement $K = 5, 10, 20, 40$ to perform the proposed GSDA attack on PointNet model and report the corresponding perturbation budgets when achieving 100% of attack success rate. We see that

the attack performance is insensitive to the number K since our attack with different K 's requires similar perturbation budgets $\mathcal{D}_c, \mathcal{D}_h$ in the data domain, as well as in the spectral domain measured by \mathcal{E}_Δ .




Fig. 5. Attacking performance when applying different ϵ_{max} restrict, in terms of success rate, Chamfer distance and Hausdorff distance results between adversarial point clouds and the originals.

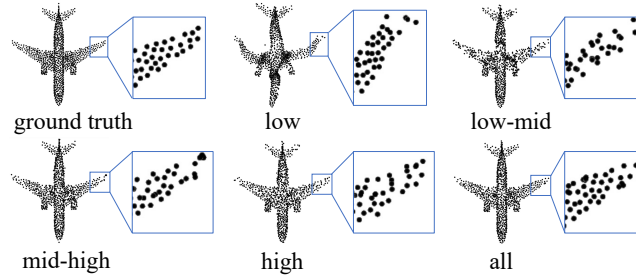


Fig. 6. Visualization of attacks on specific frequency bands.

Sensitivity on the Perturbation Size $[\epsilon_{min}, \epsilon_{max}]$. As shown in Figure 5, when we relax the perturbation constraint $[\epsilon_{min}, \epsilon_{max}]$ in the spectral domain, the success rates of our GSDA increase and achieve almost 100% of success rate when $\epsilon_{max} = -\epsilon_{min} = 3.0$, meanwhile \mathcal{D}_c and \mathcal{D}_h achieve an good performance. **Attacks on Specific Frequency Bands.** We also perform an ablation study on how attacking specific frequency bands (*i.e.*, low, low-mid, mid-high, high) affects point clouds in the data domain. As shown in Figure 6, only attacking low or high frequency band results in distortion of the rough shape or local details, which is consistent with our analysis in Sec. 3. Instead, attacking appropriate positions among the entire spectra achieves the most imperceptible results.

6 Conclusion

We propose a novel paradigm of point cloud attacks—Graph Spectral Domain Attack (GSDA), which explores insightful spectral characteristics and performs perturbation in the spectral domain to generate adversarial examples with geometric structures well preserved. Extensive experiments show the vulnerability of popular 3D models to the proposed GSDA and demonstrate the robustness of our adversarial point clouds. Such spectral methodology blazes a new path of developing attacks in 3D deep learning. Future promising directions include developing effective defense methods against such spectral domain attacks.

References

1. Atzmon, M., Maron, H., Lipman, Y.: Point convolutional neural networks by extension operators. arXiv preprint arXiv:1803.10091 (2018)
2. Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. In: 2017 IEEE Symposium on Security and Privacy (SP). pp. 39–57 (2017)
3. Chen, S., Tian, D., Feng, C., Vetro, A., Kovačević, J.: Fast resampling of three-dimensional point clouds via graphs. *IEEE Transactions on Signal Processing* **66**(3), 666–681 (2017)
4. 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 (CVPR). pp. 1907–1915 (2017)
5. Choi, J., Han, B.: Task-aware quantization network for jpeg image compression. In: European Conference on Computer Vision (ECCV). pp. 309–324 (2020)
6. Chung, F.R., Graham, F.C.: Spectral graph theory. No. 92, American Mathematical Soc. (1997)
7. Cortes, C., Mohri, M., Rostamizadeh, A.: L2 regularization for learning kernels. arXiv preprint arXiv:1205.2653 (2012)
8. Dong, Y., Liao, F., Pang, T., Su, H., Zhu, J., Hu, X., Li, J.: Boosting adversarial attacks with momentum. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9185–9193 (2018)
9. Duan, Y., Zheng, Y., Lu, J., Zhou, J., Tian, Q.: Structural relational reasoning of point clouds. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 949–958 (2019)
10. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 605–613 (2017)
11. Gao, X., Hu, W., Qi, G.J.: Grapher: Unsupervised learning of graph transformation equivariant representations via auto-encoding node-wise transformations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7163–7172 (2020)
12. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572 (2014)
13. Graham, B., Engelcke, M., Van Der Maaten, L.: 3d semantic segmentation with submanifold sparse convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9224–9232 (2018)
14. Hamdi, A., Rojas, S., Thabet, A., Ghanem, B.: Advpc: Transferable adversarial perturbations on 3d point clouds. In: European Conference on Computer Vision (ECCV). pp. 241–257 (2020)
15. Hammond, D.K., Vandergheynst, P., Gribonval, R.: Wavelets on graphs via spectral graph theory. *Appl. Comput. Harmonic Anal.* **30**(2), 129–150 (2011)
16. Hu, W., Cheung, G., Li, X., Au, O.: Depth map compression using multi-resolution graph-based transform for depth-image-based rendering. In: Proc. IEEE Int. Conf. Image Process. pp. 1297–1300 (2012)
17. Hu, W., Cheung, G., Ortega, A.: Intra-prediction and generalized graph Fourier transform for image coding. *IEEE Signal Process. Lett.* **22**(11), 1913–1917 (2015)
18. Hu, W., Cheung, G., Ortega, A., Au, O.C.: Multiresolution graph Fourier transform for compression of piecewise smooth images. *IEEE Trans. Image Process.* **24**(1), 419–433 (2015)

19. Hu, W., Pang, J., Liu, X., Tian, D., Lin, C.W., Vetro, A.: Graph Signal Processing for geometric data and beyond: Theory and applications. *IEEE Transactions on Multimedia* (2021)
20. Huttenlocher, D.P., Klanderman, G.A., Rucklidge, W.J.: Comparing images using the hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **15**(9), 850–863 (1993)
21. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
22. Kurakin, A., Goodfellow, I., Bengio, S.: Adversarial machine learning at scale. *arXiv preprint arXiv:1611.01236* (2016)
23. Li, M., Zuo, W., Gu, S., Zhao, D., Zhang, D.: Learning convolutional networks for content-weighted image compression. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 3214–3223 (2018)
24. Li, Y., Bu, R., Sun, M., Wu, W., Di, X., Chen, B.: Pointcnn: Convolution on x-transformed points. *Advances in Neural Information Processing Systems (NIPS)* **31**, 820–830 (2018)
25. Liu, D., Hu, W.: Imperceptible transfer attack and defense on 3d point cloud classification. *arXiv preprint arXiv:2111.10990* (2021)
26. 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 (2019)
27. Liu, Y., Fan, B., Meng, G., Lu, J., Xiang, S., Pan, C.: Densepoint: Learning densely contextual representation for efficient point cloud processing. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 5239–5248 (2019)
28. Liu, Y., Fan, B., Xiang, S., Pan, C.: Relation-shape convolutional neural network for point cloud analysis. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 8895–8904 (2019)
29. Ma, C., Meng, W., Wu, B., Xu, S., Zhang, X.: Efficient joint gradient based attack against sor defense for 3d point cloud classification. In: *Proceedings of the 28th ACM International Conference on Multimedia*. pp. 1819–1827 (2020)
30. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083* (2017)
31. Moosavi-Dezfooli, S.M., Fawzi, A., Fawzi, O., Frossard, P.: Universal adversarial perturbations. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1765–1773 (2017)
32. Moosavi-Dezfooli, S.M., Fawzi, A., Frossard, P.: Deepfool: a simple and accurate method to fool deep neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2574–2582 (2016)
33. 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 Conference on Computer Vision and Pattern Recognition (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. *Advances in Neural Information Processing Systems (NIPS)* (2017)
35. Ramasinghe, S., Khan, S., Barnes, N., Gould, S.: Spectral-gans for high-resolution 3d point-cloud generation. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 8169–8176. *IEEE* (2020)
36. Rosman, G., Dubrovina, A., Kimmel, R.: Patch-collaborative spectral point-cloud denoising. In: *Computer Graphics Forum*. vol. 32, pp. 1–12. *Wiley Online Library* (2013)

37. Shen, G., Kim, W.S., Narang, S.K., Ortega, A., Lee, J., Wey, H.: Edge-adaptive transforms for efficient depth map coding. In: Proc. Picture Coding Symp. pp. 566–569 (2010)
38. Shen, Y., Feng, C., Yang, Y., Tian, D.: Mining point cloud local structures by kernel correlation and graph pooling. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4548–4557 (2018)
39. Shuman, D.I., Narang, S.K., Frossard, P., Ortega, A., Vandergheynst, P.: The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Process. Mag.* **30**(3), 83–98 (2013)
40. Simonovsky, M., Komodakis, N.: Dynamic edge-conditioned filters in convolutional neural networks on graphs. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3693–3702 (2017)
41. Singh, S.P., Wang, L., Gupta, S., Goli, H., Padmanabhan, P., Gulyás, B.: 3d deep learning on medical images: a review. *Sensors* **20**(18), 5097 (2020)
42. Su, H., Maji, S., Kalogerakis, E., Learned-Miller, E.: Multi-view convolutional neural networks for 3d shape recognition. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 945–953 (2015)
43. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R.: Intriguing properties of neural networks. arXiv preprint arXiv:1312.6199 (2013)
44. Thomas, H., Qi, C.R., Deschaud, J.E., Marcotegui, B., Goulette, F., Guibas, L.J.: Kpconv: Flexible and deformable convolution for point clouds. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 6411–6420 (2019)
45. 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. vol. 34, pp. 954–962 (2020)
46. Tu, C.C., Ting, P., Chen, P.Y., Liu, S., Zhang, H., Yi, J., Hsieh, C.J., Cheng, S.M.: Autozoom: Autoencoder-based zeroth order optimization method for attacking black-box neural networks. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 742–749 (2019)
47. Wang, W., Yu, R., Huang, Q., Neumann, U.: Sgpn: Similarity group proposal network for 3d point cloud instance segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2569–2578 (2018)
48. Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (TOG)* **38**(5), 1–12 (2019)
49. Wen, Y., Lin, J., Chen, K., Chen, C.P., Jia, K.: Geometry-aware generation of adversarial point clouds. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* (2020)
50. Wicker, M., Kwiatkowska, M.: Robustness of 3d deep learning in an adversarial setting. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11767–11775 (2019)
51. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1912–1920 (2015)
52. 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)

53. Xiang, C., Qi, C.R., Li, B.: Generating 3d adversarial point clouds. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9136–9144 (2019)
54. Xu, Q., Sun, X., Wu, C.Y., Wang, P., Neumann, U.: Grid-gcn for fast and scalable point cloud learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5661–5670 (2020)
55. Xu, Y., Hu, W., Wang, S., Zhang, X., Wang, S., Ma, S., Guo, Z., Gao, W.: Predictive generalized graph Fourier transform for attribute compression of dynamic point clouds. arXiv preprint arXiv:1908.01970 (2019)
56. Yang, B., Wang, J., Clark, R., Hu, Q., Wang, S., Markham, A., Trigoni, N.: Learning object bounding boxes for 3d instance segmentation on point clouds. arXiv preprint arXiv:1906.01140 (2019)
57. Yang, J., Zhang, Q., Ni, B., Li, L., Liu, J., Zhou, M., Tian, Q.: Modeling point clouds with self-attention and gumbel subset sampling. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3323–3332 (2019)
58. Yu, T., Meng, J., Yuan, J.: Multi-view harmonized bilinear network for 3d object recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 186–194 (2018)
59. Zaheer, M., Kottur, S., Ravanbakhsh, S., Poczos, B., Salakhutdinov, R.R., Smola, A.J.: Deep sets. *Advances in Neural Information Processing Systems (NIPS)* **30** (2017)
60. Zhang, C., Florencio, D., Loop, C.: Point cloud attribute compression with graph transform. In: *Proc. IEEE Int. Conf. Image Process.* pp. 2066–2070 (2014)
61. Zhang, Q., Yang, J., Fang, R., Ni, B., Liu, J., Tian, Q.: Adversarial attack and defense on point sets. arXiv preprint arXiv:1902.10899 (2019)
62. Zhang, S., Cui, S., Ding, Z.: Hypergraph spectral analysis and processing in 3d point cloud. *IEEE Transactions on Image Processing* **30**, 1193–1206 (2020)
63. Zhang, Y., Liang, G., Salem, T., Jacobs, N.: Defense-pointnet: Protecting pointnet against adversarial attacks. In: 2019 IEEE International Conference on Big Data (Big Data). pp. 5654–5660 (2019)
64. Zhao, Y., Wu, Y., Chen, C., Lim, A.: On isometry robustness of deep 3d point cloud models under adversarial attacks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1201–1210 (2020)
65. Zheng, T., Chen, C., Yuan, J., Li, B., Ren, K.: Pointcloud saliency maps. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 1598–1606 (2019)
66. 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: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10356–10365 (2020)
67. 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: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 1961–1970 (2019)