

# 3D-ANC: Adaptive Neural Collapse for Robust 3D Point Cloud Recognition

Yuanmin Huang<sup>1</sup>, Wenxuan Li<sup>1</sup>, Mi Zhang<sup>1†</sup>, Xiaohan Zhang<sup>1</sup>, Xiaoyu You<sup>2</sup>, Min Yang<sup>1†</sup>

<sup>1</sup>College of Computer Science and Artificial Intelligence, Fudan University, Shanghai, China

<sup>2</sup>School of Information Science and Engineering, East China University of Science and Technology, Shanghai, China  
yuanminhuang23@m.fudan.edu.cn, wxli24@m.fudan.edu.cn, mi\_zhang@fudan.edu.cn, xh\_zhang@fudan.edu.cn, xiaoyuyou@ecust.edu.cn, m\_yang@fudan.edu.cn

## Abstract

Deep neural networks have recently achieved notable progress in 3D point cloud recognition, yet their vulnerability to adversarial perturbations poses critical security challenges in practical deployments. Conventional defense mechanisms struggle to address the evolving landscape of multifaceted attack patterns. Through systematic analysis of existing defenses, we identify that their unsatisfactory performance primarily originates from an entangled feature space, where adversarial attacks can be performed easily. To this end, we present *3D-ANC*, a novel approach that capitalizes on the Neural Collapse (NC) mechanism to orchestrate discriminative feature learning. In particular, NC depicts where last-layer features and classifier weights jointly evolve into a simple equiangular tight frame (ETF) arrangement, establishing maximally separable class prototypes. However, leveraging this advantage in 3D recognition confronts two substantial challenges: (1) prevalent class imbalance in point cloud datasets, and (2) complex geometric similarities between object categories. To tackle these obstacles, our solution combines an ETF-aligned classification module with an adaptive training framework consisting of representation-balanced learning (RBL) and dynamic feature direction loss (FDL). *3D-ANC* seamlessly empowers existing models to develop disentangled feature spaces despite the complexity in 3D data distribution. Comprehensive evaluations state that *3D-ANC* significantly improves the robustness of models with various structures on two datasets. For instance, DGCNN’s classification accuracy is elevated from 27.2% to 80.9% on ModelNet40 – a 53.7% absolute gain that surpasses leading baselines by 34.0%.

**Extended version** — <https://arxiv.org/pdf/2511.07040>

## 1 Introduction

Point cloud data, which captures the 3D structures of objects and environments, is increasingly crucial in various applications, including autonomous driving, robotics, and 3D scene understanding (Hamdi, Giancola, and Ghanem 2021; Qiu, Anwar, and Barnes 2021; Abbasi et al. 2022; Duan et al. 2021). In recent years, deep neural networks (DNNs) have achieved remarkable success in point cloud recognition, largely due to their deep, non-linear architectures (Qi

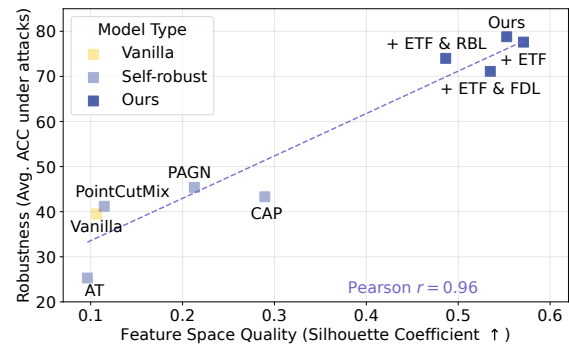


Figure 1: Adversarial robustness highly correlates with feature space quality. However, current models and self-robust defenses have poor feature disentanglement ability. *3D-ANC* significantly improves robustness with superior feature separability. Experiments are conducted on ModelNet40, PointNet. Detailed setting is described in Sec. 5.4.

et al. 2017a,b; Wang et al. 2019). However, alongside their widespread adoption, there has been growing concern over their vulnerability to *adversarial attacks*. As an attacker, one can perform imperceptible modifications to point clouds to drastically alter the predictions of a DNN. Given the significant risks posed to real-world applications, enhancing the adversarial robustness of point cloud recognition has emerged as a critical research area in 3D vision.

Since point cloud data consists of 3D coordinates of discrete points, this unique data format provides attackers with various perturbation strategies. Typical adversarial attacks for point clouds involve adding (Xiang, Qi, and Li 2019), deleting (Zheng et al. 2019), or shifting (Kim et al. 2021) points (i.e., altering the positions of existing points) within a point cloud. Additionally, attackers may employ generative models to transform a benign point cloud into an adversarial one (Hamdi et al. 2020; Zhou et al. 2020). More recent shape-invariant attacks (Huang et al. 2022; Lou et al. 2024) preserve the geometric structure of benign samples, thereby generating adversarial examples hard to detect. Given the diversity and sophistication of these attacks, the demand for robust defense mechanisms in point cloud recognition has become increasingly urgent.

<sup>†</sup>Corresponding authors: Mi Zhang and Min Yang.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Current defenses can be broadly classified into *input pre-processing* (Zhou et al. 2019; Sun et al. 2023) and *self-robust models* (Liu, Yu, and Su 2019; Sun et al. 2021). Input preprocessing focuses on mitigating anomaly patterns introduced by certain attacks, such as outlier points in a point cloud. While effective, this approach is limited in its ability to counter more advanced attacks, e.g., shape-invariant attacks (Lou et al. 2024), which tend to produce fewer outliers. In contrast, self-robust models aim to improve the model’s intrinsic robustness against potential adversarial threats. For example, techniques such as adversarial training (Zhang and Wang 2019) or self-supervised learning (Sun et al. 2021; Zhang et al. 2022) can be employed to develop robust models. Nonetheless, our pilot study (Sec. 3.1) reveals that existing self-robust models often struggle to extract well-separated sample features. This can lead to unsatisfying robustness under adversarial conditions, as benign samples can be easily perturbed across decision boundaries when the feature space overlaps between classes.

Faced with emerging attacks, a key challenge in developing effective self-robust models is to achieve a robustly disentangled feature space. To tackle this issue, we draw inspiration from the recently discovered *Neural Collapse (NC)* phenomenon (Papayan, Han, and Donoho 2020). It has been observed that, during the terminal phase of training on a balanced dataset, the last-layer features of same-class samples tend to collapse to their within-class mean. Simultaneously, the within-class means of all classes, along with the corresponding classifier vectors, converge to the vertices of a simplex equiangular tight frame (ETF). An ETF structure is depicted in Fig. 3(a), which maximizes the pairwise angular separation between vectors. As a result, NC offers a promising solution to a well-disentangled feature space: replacing the traditional learnable classification head with a non-learnable one satisfying ETF structure. This should enhance the adversarial robustness of a model inherently.

However, adapting NC to point cloud recognition differs drastically from image classification under typical balanced settings. Specifically, two substantial challenges arise: (1) *the imbalanced nature of point cloud datasets*, which can severely distort the training of models with ETF classifiers, degrading performance on clean samples; and (2) *the complex inter-class geometric similarities*, where samples from distinct classes can share nearly identical geometries, causing feature overlap even when trained with ETF classifiers, leading to suboptimal robustness performance.

To address these challenges, we propose **3D-ANC**, an **adaptive** training framework designed to facilitate the integration of **NC** (an ETF classifier) into **3D** point cloud recognition. The framework comprises two key components: a *representation-balanced learning (RBL)* module and a *dynamic feature direction loss (FDL)*. Built upon an ETF classifier, RBL introduces learnable rotations to the ETF head, enabling adaptation to imbalanced class distributions while preserving ETF properties. Meanwhile, the dynamic FDL reinforces both intra-class compactness and inter-class repulsion, adaptively targeting geometrically similar classes that hinder feature disentanglement. Overall, 3D-ANC achieves a robustly disentangled feature space tailored

to point cloud data.

Our main contributions are summarized as follows:

- We introduce 3D-ANC, the first approach to leverage the NC phenomenon for robust point cloud recognition. By incorporating an ETF classifier within an adaptive training framework, 3D-ANC inherently preserves class-specific features and maximizes inter-class repulsion, significantly improving adversarial robustness.
- 3D-ANC is model-agnostic and computationally efficient, seamlessly integrating with existing point cloud recognition models across architectures by simply replacing the classification head with an ETF one.
- 3D-ANC achieves substantial improvements in adversarial robustness across three widely used models on two benchmark datasets. For example, 3D-ANC boosts the classification accuracy of DGCNN from 27.2% to 80.9%, outperforming the best baseline, which attains only 46.9%.

## 2 Related Work

### 2.1 Point Cloud Recognition

Point cloud recognition involves the task of categorizing objects represented as sets of points with 3D coordinates into distinct classes. This problem can be formally described as learning a mapping function  $F : \mathbf{X} \mapsto y$ , where  $\mathbf{X} \in \mathbb{R}^{N \times 3}$  represents a point cloud composed of  $N$  points in 3D space, and  $y \in \{1, 2, \dots, C\}$  denotes the ground truth label among  $C$  possible categories. In recent years, DNNs have shown great success in point cloud recognition due to their ability to extract complex, non-linear features (Lu and Shi 2020). This paper focuses on point-based models (Qi et al. 2017a,b). Unlike approaches that require pre-processing steps such as voxelization (Xie et al. 2020) or pillarization (Lang et al. 2019), point-based models directly take the 3D coordinates of point clouds as input. These models start by extracting latent features from each of the  $N$  points in the input point cloud  $\mathbf{X} \in \mathbb{R}^{N \times 3}$  using a non-linear feature extractor  $f$ , such as an MLP in PointNet/PointNet++ (Qi et al. 2017a,b), a GNN in DGCNN (Wang et al. 2019), a CNN in PointCNN (Li et al. 2018), or a Transformer in PCT (Guo et al. 2021). The extracted global feature  $\mathbf{h}$  is then passed through a fully-connected layer that serves as the classification head  $g$ . For a more detailed review, please refer to (Lu and Shi 2020).

### 2.2 Adversarial Examples

Adversarial attacks are designed to manipulate the output of DNN models by introducing small perturbations to the input data. Following attacks on image classification models (Szegedy et al. 2014; Goodfellow, Shlens, and Szegedy 2014), adversarial attacks also appear in point cloud recognition (Xiang, Qi, and Li 2019; Wen et al. 2020). The objective can be formulated as:

$$\max_{\eta} \mathcal{L}(F(\tilde{\mathbf{X}}, y)), \quad \text{s.t.}, \|\eta\|_p \leq \delta, \quad (1)$$

where  $F$  represents the model,  $\mathcal{L}$  is the classification loss function,  $\eta$ ,  $\tilde{\mathbf{X}} = \mathbf{X} \oplus \eta$ , and  $y$  denote the perturbation,




Figure 2: The t-SNE visualization of PointNet features on ModelNet40 test set under various defense schemes. For input preprocessing defense (DUP-Net and Diffusion), we use the feature of the preprocessed samples for visualization. The entangled feature space leaves room for adversarial attacks, i.e., samples can be perturbed to overlapped classes easily.

the adversarial example, and the ground truth label, respectively. Adversarial attacks in point cloud classification exhibit a wide range of strategies. With respect to the form of perturbations (i.e.,  $\oplus$  in Eq.1), options include shifting (Xiang, Qi, and Li 2019), adding (Liu, Yu, and Su 2020; Xiang, Qi, and Li 2019) or deleting (Wicker and Kwiatkowska 2019; Zheng et al. 2019) points, or transforming the entire point cloud (Hamdi et al. 2020; Zhou et al. 2020). Regarding the perturbation budget  $\|\eta\|_p$ , various distance metrics have been employed, such as  $L_0$ ,  $L_2$ , and  $L_\infty$ , with Chamfer Distance often used as a representative for  $L_2$  distance (Xiang, Qi, and Li 2019). Recently, shape-invariant attacks (Huang et al. 2022; Wen et al. 2020) have been proposed to maintain local or global geometric patterns of the point cloud during attacks, resulting in more imperceptible perturbations.

Adversarial defenses for point cloud classifiers have evolved in tandem with the development of attacks. Recent approaches can be broadly divided into input preprocessing (Zhou et al. 2019; Sun et al. 2023) and self-robust models (Liu, Yu, and Su 2019; Zhang and Wang 2019; Sun et al. 2021). Input preprocessing methods mitigate specific effects introduced by particular attacks. For instance, SOR (Zhou et al. 2019) filters out outlier points introduced by adversarial perturbations, while DUP-Net (Zhou et al. 2019) addresses inconsistent point density within a point cloud caused by such attacks. PointDP (Sun et al. 2023) employs diffusion models to purify the perturbed point cloud during the diffusion process. Conversely, self-robust models draw inspiration from defenses in image classification. For example, adversarial training methods (Madry et al. 2017) augment the training dataset with adversarial examples. Additionally, some models (Sun et al. 2021; Zhang et al. 2022; Ding et al. 2023) utilize self-supervised learning to enhance the classifiers in the presence of noisy inputs, thereby improving adversarial robustness.

Despite their effectiveness, both defense strategies share a common limitation: *they generalize poorly to unseen attacks with characteristics different from those they were designed to counter*. When confronted with novel attacks, input preprocessing techniques may fail to remove malicious points, and existing self-robust models may remain vulnerable. In contrast, our approach focuses on constructing a naturally well-disentangled feature space for point cloud recognition models. This design inherently equips our method with the capability to resist a wide range of attacks.

### 2.3 Neural Collapse

Neural Collapse (NC) was first introduced in (Papayan, Han, and Donoho 2020), describing how, in a well-trained classification model, the last-layer features collapse to their within-class centers. This collapse forms a simplex equiangular tight frame (ETF) structure, where the feature vectors align with the classifier’s weight vectors. Subsequent research has demonstrated that NC represents global optimality under balanced training with both cross-entropy (CE) loss (Fang et al. 2021; Ji et al. 2022; Lu and Steinerberger 2022) and mean squared error (MSE) loss (Han, Papayan, and Donoho 2022; Tirer and Bruna 2022; Zhou et al. 2022).

Further studies (Yang et al. 2022; Peifeng et al. 2023; Zhong et al. 2023; Thrapoulidis et al. 2022; Xie et al. 2023) have examined NC in the context of imbalanced training. Yang et al. (Yang et al. 2022) showed that using a fixed ETF classifier head can align features with the ETF structure, improving classification performance even with imbalanced datasets. Zhong et al. (Zhong et al. 2023) observed that while imbalanced distributions can disrupt the NC structure, initializing the classifier as a simplex ETF still benefits minor class differentiation.

Inspired by these findings, we aim to induce an equiangular separation among class centers to enhance feature disentanglement in the representation space, thereby improving the model’s robustness against adversarial attacks.

## 3 Pilot Study

### 3.1 Preliminary Results

We first conduct preliminary experiments to evaluate existing defenses. Given that adversarial attacks introduce subtle perturbations to point clouds—imperceptible to humans yet capable of deceiving model predictions—we hypothesize that the vulnerabilities of current defenses may originate from weaknesses in the feature space. To examine this hypothesis, we analyze the feature extraction capabilities of existing models and defenses, as illustrated in Fig. 2. The models are evaluated with the output of the feature extractor  $h = f(X)$  as the feature given benign samples.

As depicted in the figure, the vanilla PointNet faces challenges in distinguishing features across different classes. Although input preprocessing techniques (e.g., DUP-Net (Zhou et al. 2019), Diffusion (Sun et al. 2023)) and existing self-robust models (e.g., Adversarial Training (Liu,

Yu, and Su 2019), PointCutMix (Zhang et al. 2022)) offer some enhancement, *substantial overlap remains among inter-class features*. More visualizations of advanced models in Appendix D.5 exhibit similar trends. Now that *point cloud models struggle to extract well-separated features even from clean samples, adversarial examples can easily go across decision boundaries, posing considerable risk*. Such an observation motivates us to improve the feature extraction capabilities of existing point cloud recognition models.

### 3.2 Neural Collapse Phenomenon

We then introduce the Neural Collapse (NC) phenomenon. As explored in (Papayan, Han, and Donoho 2020), for neural networks trained on a balanced dataset, the NC phenomenon occurs in the last-layer features and the weight vectors of the classification head as training converges. Specifically, both the features and weight vectors collapse to the vertices of a simplex equiangular tight frame (ETF), as depicted in Fig. 3(a). Formally, a simplex ETF is defined as follows:

**Definition 1** (Simplex Equiangular Tight Frame). *A collection of vectors  $\mathbf{m}_i \in \mathbb{R}^d, i = 1, 2, \dots, K, d \geq K - 1$ , is said to be a simplex equiangular tight frame if:*

$$\mathbf{M} = \sqrt{\frac{K}{K-1}} \mathbf{R} \left( \mathbf{I}_K - \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^T \right), \quad (2)$$

where  $\mathbf{M} = [\mathbf{m}_1, \dots, \mathbf{m}_K] \in \mathbb{R}^{d \times K}$ ,  $\mathbf{R} \in \mathbb{R}^{d \times K}$  represents a rotation matrix satisfying  $\mathbf{R}^T \mathbf{R} = \mathbf{I}_K$ ,  $\mathbf{I}_K$  is the identity matrix, and  $\mathbf{1}_K$  is a vector of all ones.

This definition implies that a simplex ETF has the property where any pair of vectors exhibits the same maximal angular separation, i.e.,  $\mathbf{m}_i^T \mathbf{m}_j = -\frac{1}{K-1}$  for  $i \neq j$ , where  $i, j \in 1, \dots, K$ .

We now formally describe the NC phenomenon:

**NC1: Variability collapse.** The last-layer feature  $\mathbf{h}_{k,i}$  of any sample  $i$  from class  $k$  collapses to the within-class mean  $\bar{\mathbf{h}}_k = \text{Avg}_i \{\mathbf{h}_{k,i}\}$ . As a result, the within-class feature variability collapses to zero.

**NC2: Convergence to simplex ETF.** The normalized within-class means converge to a simplex ETF that satisfies Eq. 2, i.e.,  $\bar{\mathbf{h}}_k = (\bar{\mathbf{h}}_k - \bar{\mathbf{h}}_G) / \|\bar{\mathbf{h}}_k - \bar{\mathbf{h}}_G\|$  for  $k \in \{1, K\}$ , where  $\bar{\mathbf{h}}_G$  represents the global mean across all classes.

**NC3: Self-duality.** Classifier weight  $\mathbf{w}_k$  for class  $k$  aligns with the normalized class mean, i.e.,  $\bar{\mathbf{h}}_k = \mathbf{w}_k / \|\mathbf{w}_k\|$ .

**NC4: Nearest class center prediction.** The model’s prediction for any sample feature  $\mathbf{h}$  collapses to selecting the nearest class mean (center), i.e.,  $\text{argmax}_k \langle \mathbf{h}, \mathbf{w}_k \rangle = \text{argmin}_k \|\mathbf{h} - \bar{\mathbf{h}}_k\|$ .

## 4 The Proposed 3D-ANC

### 4.1 Neural Collapse for Robust Recognition

Preliminary results in Sec. 3.1 highlight the inadequacies of current models and defenses in effectively extracting discriminative features. Fortunately, the NC phenomenon offers an ideal framework for achieving a well-disentangled feature space, characterized by highly cohesive intra-class features (NC1) and distinct inter-class features (NC2).

Instead of allowing the model to gradually converge to NC during training, we propose to initialize the model’s classification head with a predefined neural collapse state, specifically a simplex Equiangular Tight Frame (ETF) head (Fig. 3 (a)). Formally, for a model  $F = f \circ g$  consisting of a feature extractor  $f(X) = \mathbf{h}$  and a classifier  $g(\mathbf{h})$ , we set  $g(\mathbf{h}) = \mathbf{h} \cdot \mathbf{W}$ , where  $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_K]$  represents a randomly initialized simplex ETF, as defined in Eq. 2.

According to NC3 and NC4, in an NC-convergent model, the feature  $\mathbf{h}_k$  of class  $k$  aligns with both the classification vector  $\mathbf{w}_k$  and the normalized class mean  $\bar{\mathbf{h}}_k$ . To encourage this alignment, we introduce a dot product loss defined as:

$$\mathcal{L}_{dot}(\mathbf{h}, \mathbf{W}) = \frac{1}{2\sqrt{E_W E_H}} \left( \mathbf{w}_k^T \mathbf{h} - \sqrt{E_W E_H} \right)^2, \quad (3)$$

where  $\mathbf{h}$  denotes the feature of a sample belonging to class  $k$ , and  $\mathbf{w}_k$  is the classification vector for class  $k$ . The terms  $E_W$  and  $E_H$  are predefined  $L_2$ -norm constraints, ensuring that  $\|\mathbf{h}\|^2 \leq E_H$  and  $\|\mathbf{w}_k\|^2 \leq E_W$ .

To guarantee a thoroughly disentangled feature space, the matrix  $\mathbf{W}$  remains fixed during training. By enforcing the simplex ETF head, the feature extractor is compelled to learn features that inherently align with the classifier vector  $\mathbf{w}_k$  for each class  $k$ . This design encourages the feature extractor to focus on class-specific characteristics within each point cloud, thereby enhancing the model’s robustness.

### 4.2 Adaptive Training Framework

While Yang et al. (Yang et al. 2022) suggested that directly training a model with a fixed simplex ETF head can lead to effective convergence in image classification, we find this approach challenging to implement for point cloud data.

Unlike image datasets, point cloud datasets often exhibit complex distribution characteristics that hinder direct application of ETF-based training. For instance, the widely used ModelNet40 dataset (Wu et al. 2015) *not only suffers from significant class imbalance but also presents substantial geometric similarity between certain classes (Fig. 3 (left))*. Specifically, the number of samples per class varies considerably: the class *chair* contains  $\sim 900$  samples, whereas the class *bowl* has over  $10\times$  fewer. Moreover, as highlighted in our case study in Sec. 5.5, some classes exhibit strong geometric resemblance, such as *night stand* and *dresser*, or *desk* and *table*, *vase* and *flower pot*.

These intrinsic imbalances and similarities hinder models from learning well-disentangled features, as dominant classes may bias feature learning, while geometrically similar classes tend to cluster in feature space. The fixed ETF head exacerbates these issues, as its global orientation is fixed once initialized, preventing adaptation to the underlying data distribution. To address these challenges, we propose an adaptive training framework.

**Representation-Balanced Learning (RBL)** The classification matrix  $\mathbf{W}$  in a fixed ETF head is untrainable. Its orientation is dictated by the rotation matrix  $\mathbf{R}$  in Eq. 2. However, the orientation can critically influence a model’s generalization performance (Peifeng et al. 2023). Our ablation

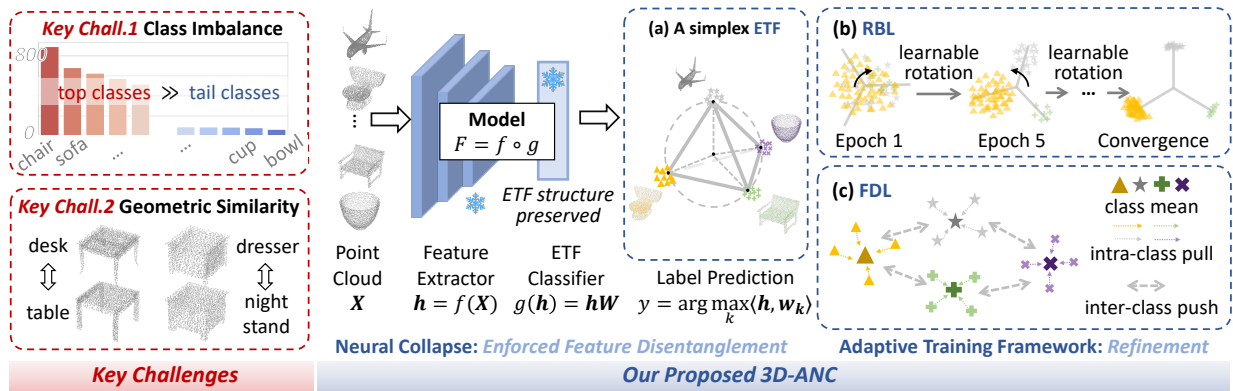


Figure 3: Overview of the proposed 3D-ANC. (Left): Key challenges in point cloud recognition to form a well-separated feature space: (1) class imbalance and (2) geometric similarity between certain classes. (Right): (a) Our 3D-ANC enforces feature disentanglement via ETF classifier  $g$ . To perform fine-grained refinement towards feature separation for robust point cloud recognition, we further propose an adaptive training framework comprising (b) Representation-Balanced Learning and (c) dynamic Feature Direction Loss, which progressively adapts the training with ETF classifier.

study (Sec. 5.3) shows that using a fixed ETF head harms classification accuracy on benign samples.

To address this issue, RBL enables the ETF head to be learnable, albeit under specific constraints. RBL was initially introduced to enhance the clean performance of models employing a fixed ETF head in long-tailed scenarios (Peifeng et al. 2023). In this approach, the rotation matrix  $\mathbf{R}$  is updated during training, but is constrained to remain an orthogonal matrix, preserving the ETF properties (Fig. 3 (b)). This constraint can be enforced with standard optimization algorithms such as SGD or Adam. Consequently, both the feature extractor and the classifier jointly optimize feature representations and global orientation, reaching balanced clean performance on imbalanced point cloud datasets. RBL also forms more accurate feature representations, which is crucial for FDL below.

**Dynamic Feature Direction Loss (FDL)** While RBL enhances feature alignment, it does not fully address the challenges in complex point cloud datasets. In particular, point clouds often share overlapping geometric features between classes, making it challenging for ETF and RBL to achieve adequate feature disentanglement.

To better handle these characteristics, we introduce the Feature Direction Loss (FDL), which dynamically adjusts the direction of sample features during training. FDL encourages each sample to align more closely with its within-class mean  $\bar{\mathbf{h}}_k$ . Additionally, it promotes repulsion by pushing sample features away from the mean of the closest non-ground-truth class  $\bar{\mathbf{h}}_{k'}$ , where  $k' \neq k$  (Fig. 3 (c)).

Formally, the dynamic FDL is defined as:

$$\mathcal{L}_{FDL}(\mathbf{h}, \bar{\mathbf{h}}_k, \bar{\mathbf{h}}_{k'}) = -\frac{\mathbf{h}^T \bar{\mathbf{h}}_k}{\|\mathbf{h}\| \|\bar{\mathbf{h}}_k\|} + \frac{\mathbf{h}^T \bar{\mathbf{h}}_{k'}}{\|\mathbf{h}\| \|\bar{\mathbf{h}}_{k'}\|}, \quad (4)$$

where  $\|\cdot\|$  denotes the  $L_2$  norm,  $\bar{\mathbf{h}}_k, \bar{\mathbf{h}}_{k'}$  are updated at every epoch. This formulation clusters intra-class features while enhancing inter-class repulsion by reducing similarity with non-ground-truth class means.

### 4.3 Overall Framework

The overall framework of our proposed 3D-ANC is illustrated in Fig. 3 (right). The training objective is defined as:

$$\mathcal{L} = \mathcal{L}_{dot}(\mathbf{h}, \mathbf{W}) + \lambda \cdot \mathcal{L}_{FDL}(\mathbf{h}, \bar{\mathbf{h}}_k, \bar{\mathbf{h}}_{k'}), \quad (5)$$

where  $\lambda$  is a hyperparameter that controls the weight of  $\mathcal{L}_{FDL}$ . To ensure that  $\mathcal{L}_{FDL}$  provides accurate guidance, it is only incorporated after several warm-up epochs. Detailed pseudocode is provided in Appendix B.1.

During training, sample features undergo global rotation by RBL, while FDL further aligns them with their respective class means. RBL dynamically provides corrected orientations for ETF, where the adapted sample features is used by FDL to refine the feature space. This dual adjustment enhances the model’s capability to distinguish between classes, even in scenarios of high inter-class geometric similarity. Combining the strengths of the ETF head with the adaptive training framework, we construct a robust model that naturally promotes well-disentangled feature extraction, addressing the unique characteristics of point cloud data.

Importantly, 3D-ANC is model-agnostic, which can be integrated with any point cloud recognition model. By simply replacing the classifier with a simplex ETF head and training the model using the adaptive training framework, robust models can be developed with minimal computational overhead. Additionally, 3D-ANC can be combined with input preprocessing defenses to further enhance robustness.

### 4.4 Robustness Analysis

We further analyze how 3D-ANC produces a robust feature extractor characterized by cohesive intra-class features and distinct inter-class features. The analysis is based on the gradient of loss functions with respect to sample feature  $\mathbf{h}$ .

We start with  $\mathcal{L}_{dot}$ , whose gradient is given by:

$$\nabla_{\mathbf{h}} \mathcal{L}_{dot}(\mathbf{h}, \mathbf{W}) = (\cos \angle(\mathbf{h}, \mathbf{w}_k) - 1) \mathbf{w}_k. \quad (6)$$

Thus,  $\mathcal{L}_{dot}$  pulls the feature  $\mathbf{h}$  to align with the predefined classification vector  $\mathbf{w}_k$  in the ETF head. According to NC3,  $\mathbf{w}_k$  shares the same direction as the mean feature of class  $k$ .

Defense	PointNet						DGCNN						PCT					
	IFGM	G3-Pert	G3-Add	Drop	AdvPC	Avg.	IFGM	G3-Pert	G3-Add	Drop	AdvPC	Avg.	IFGM	G3-Pert	G3-Add	Drop	AdvPC	Avg.
Vanilla	1.6	64.4	48.7	65.0	17.6	39.5	0.0	34.6	27.4	69.6	4.3	27.2	7.4	64.7	56.5	82.9	26.2	47.5
SOR <sup>†</sup>	23.6	63.4	50.9	69.0	20.9	45.6	11.9	35.2	25.2	71.6	13.2	31.4	38.3	63.7	54.9	79.0	32.2	53.6
DUP-Net <sup>†</sup>	24.8	63.0	50.4	67.8	22.0	45.6	32.1	37.1	20.7	49.2	13.7	30.5	45.1	62.5	52.2	73.9	36.9	54.1
Diffusion <sup>†</sup>	33.2	64.4	49.2	69.1	23.5	47.9	63.4	39.8	29.5	68.6	33.2	46.9	62.2	64.7	57.9	<b>84.1</b>	38.4	61.5
AT <sup>°</sup>	15.1	48.1	22.0	36.3	4.9	25.3	22.7	28.6	30.3	59.1	2.5	28.6	5.1	59.2	57.4	79.8	5.3	41.4
PAGN <sup>°</sup>	24.1	63.5	50.3	69.5	19.5	45.4	10.5	34.9	21.5	71.2	6.1	28.8	49.5	52.1	52.5	75.1	33.9	52.6
PointCutMix <sup>°</sup>	32.4	45.8	44.2	73.6	9.8	41.2	37.4	43.6	43.2	63.3	19.5	41.4	58.3	41.9	47.8	74.2	23.1	49.0
CAP <sup>°</sup>	25.5	65.4	49.9	59.2	16.3	43.3	2.0	58.5	36.4	78.7	15.4	38.2	45.7	47.7	43.4	76.0	13.8	45.3
Ours <sup>°</sup>	<b>80.2</b>	<b>79.8</b>	<b>75.6</b>	<b>80.4</b>	<b>77.9</b>	<b>78.8</b>	<b>84.3</b>	<b>84.2</b>	<b>69.9</b>	<b>84.6</b>	<b>81.3</b>	<b>80.9</b>	<b>75.2</b>	<b>81.1</b>	<b>72.6</b>	82.2	<b>75.5</b>	<b>77.3</b>

Table 1: Classification accuracy (%) with defenses (input preprocessing<sup>†</sup>, self-robust<sup>°</sup>) on ModelNet40. Best values in bold.

Attack	Vanilla	SOR <sup>†</sup>	DUP-Net <sup>†</sup>	Diffusion <sup>†</sup>	AT <sup>°</sup>	PAGN <sup>°</sup>	PointCutMix <sup>°</sup>	CAP <sup>°</sup>	Ours
KNN	30.2	36.6	37.2	<u>50.0</u>	4.8	36.1	27.8	32.0	<b>79.9</b>
GeoA3	8.1	51.3	51.0	<b>72.2</b>	1.3	49.2	54.7	49.3	<u>55.3</u>
SI	18.7	35.5	34.1	<u>45.6</u>	4.6	40.2	44.2	29.4	<b>53.0</b>
HiT	41.9	51.4	52.8	51.9	35.8	48.5	<b>73.7</b>	51.5	<u>66.1</u>
Avg.	24.7	43.7	43.8	<u>54.9</u>	11.6	43.5	50.1	40.6	<b>63.6</b>

Table 2: Classification accuracy (%) of shape-invariant attacks against defenses on ModelNet40, PointNet. Best values in bold, second-best values underlined.

Next, we analyze the loss function  $\mathcal{L}_{FDL}$ :

$$\nabla_{\mathbf{h}} \mathcal{L}_{FDL} = \nabla_{\text{pull}} + \nabla_{\text{push}}, \quad (7)$$

$$\nabla_{\text{pull}} = \frac{(\mathbf{h}^T \bar{\mathbf{h}}_k) \mathbf{h}}{\|\mathbf{h}\|^3 \|\bar{\mathbf{h}}_k\|} - \frac{\bar{\mathbf{h}}_k}{\|\mathbf{h}\| \|\bar{\mathbf{h}}_k\|}, \quad (8)$$

$$\nabla_{\text{push}} = \frac{\bar{\mathbf{h}}_{k'}}{\|\mathbf{h}\| \|\bar{\mathbf{h}}_{k'}\|} - \frac{(\mathbf{h}^T \bar{\mathbf{h}}_{k'}) \mathbf{h}}{\|\mathbf{h}\|^3 \|\bar{\mathbf{h}}_{k'}\|}. \quad (9)$$

Hence,  $\mathcal{L}_{FDL}$  promotes the alignment of feature  $\mathbf{h}$  with the class  $k$  mean, while simultaneously pushing it away from the closest non-ground-truth class  $k'$  feature mean. The degree of repulsion is adaptively scaled by the similarity between the two classes, thereby strengthening the separation of challenging class pairs.

Since  $\bar{\mathbf{h}}_k$  and  $\bar{\mathbf{h}}_{k'}$  are derived from actual sample features extracted by model,  $\mathcal{L}_{FDL}$  provides a practical complement to  $\mathcal{L}_{dot}$ : it balances the ideal simplex ETF structure with a more flexible yet well-disentangled feature organization. Notably, the effectiveness of this complement depends on accurate  $\bar{\mathbf{h}}_k$  and  $\bar{\mathbf{h}}_{k'}$ . To address potential biases in feature distribution, the global rotation by RBL mitigates imbalances, indirectly enhancing feature disentanglement. In summary, adversarial robustness is initially enhanced by the ETF classifier and is further reinforced by  $\mathcal{L}_{FDL}$  with RBL, which explicitly separates features that are closely clustered. Further details are provided in Appendix A.

## 5 Experiments

In this section, we address the following research questions:

**RQ1:** How robust is 3D-ANC against both standard adversarial attacks and shape-invariant attacks?

**RQ2:** How do the ETF head and the adaptive training framework (RBL, FDL) contribute to 3D-ANC?

**RQ3:** How does feature disentanglement correlate with robustness? To what extent does 3D-ANC improve the feature disentanglement capabilities of existing models?

### 5.1 Experimental Settings

**Models and Datasets.** We evaluate three point cloud classification models with different architectures: PointNet (Qi et al. 2017a), DGCNN (Wang et al. 2019), Point Cloud Transformer (PCT) (Guo et al. 2021) using ModelNet40 (Wu et al. 2015) (with 9,843 training, 2,468 test samples) and ShapeNet (Chang et al. 2015) (with 35,708 training, 15,429 test samples). Each object is sampled to 2,048 points and normalized to a unit sphere. For efficiency, we randomly select 2,732 samples from ShapeNet’s test set, with a maximum of 60 samples per class.

**Baseline Attacks.** We test nine attacks: IFGM (Liu, Yu, and Su 2019), G3-Add (Xiang, Qi, and Li 2019), G3-Pert (Xiang, Qi, and Li 2019), Drop (Zheng et al. 2019), AdvPC (Hamdi et al. 2020), KNN (Tsai et al. 2020), GeoA3 (Wen et al. 2020), ShapeInvariant (SI) (Huang et al. 2022), and HiT (Lou et al. 2024). We consider untargeted attacks, i.e., the attacker aims at suppressing the prediction of the ground truth label.

**Baseline Defenses.** We include three input-preprocessing defenses: SOR (Zhou et al. 2019), DUP-Net (Zhou et al. 2019), and PointDP (Diffusion) (Sun et al. 2023), and four self-robust models: AT (Liu, Yu, and Su 2019), PAGN (Liang et al. 2022), PointCutMix(K) (Zhang et al. 2022), and CAP (Ding et al. 2023).

**Evaluation Metrics.** We use classification accuracy (ACC) as evaluation metric of robustness. We use Silhouette Coefficient (SC) to quantify the quality of the model feature space. More details are available in Appendix B.2.

### 5.2 Robustness against Attacks

To address **RQ1**, we first evaluate the robustness of 3D-ANC against various attacks. The results for ModelNet40 are presented in Table 1, while the results for ShapeNet are provided in Appendix D. The *Vanilla* row represents models with no defense. The *Avg.* columns summarize the average results for each row.

Table 1 shows that *existing defenses struggle with unseen attacks*. Among input preprocessing defenses, DUP-Net performs poorly against Drop attacks, as its reconstruction capabilities cannot restore missing parts of a point cloud, re-

Component	Clean	IFGM	G3-Pert	G3-Add	Drop	AdvPC	Avg.
Vanilla	86.2	1.6	64.4	48.7	65.0	17.6	39.5
+ ETF	-0.6	<u>+76.9</u>	+13.9	<u>+25.0</u>	<u>+14.2</u>	<b>+60.6</b>	<u>+38.1</u>
+ ETF & RBL	<b>+3.7</b>	<u>+66.7</u>	<u>+14.4</u>	+24.4	+7.9	+59.2	+34.5
+ ETF & FDL	-1.4	+66.5	+9.9	+22.7	+4.9	+54.3	+31.7
+ All (Ours)	<u>+0.9</u>	<b>+78.5</b>	<b>+15.4</b>	<b>+27.0</b>	<b>+15.4</b>	<u>+60.3</u>	<b>+39.3</b>

Table 3: Classification accuracy (%) of attacks against 3D-ANC with different components on ModelNet40, PointNet. Best values in bold, second-best values underlined.

sulting in lowest ACC on both DGCNN and PCT. For self-robust models, AT fails to generalize to unseen attacks like AdvPC, leading to an ACC of 2.5% on DGCNN. Similar performance drops share with PointCutMix and CAP. In contrast, *3D-ANC excels due to the inherently well-disentangled feature modeling provided by NC*. It achieves state-of-the-art performance against most attacks across three models. For example, the average ACC of DGCNN increases substantially from 27.2% to 80.9%, surpassing second-best baseline by 34.0%.

Additionally, we evaluate shape-invariant attacks on PointNet in Table 2. *The results echo that existing defenses struggle with unseen adversarial patterns*. For instance, Diffusion performs well against attacks with subtle perturbations along the object surface, such as GeoA3, yet fails to counteract the hidden noise beneath the surface introduced by HiT. Conversely, while PointCutMix is robust against HiT, its self-supervised learning cannot survive more severe perturbations. On the contrary, *3D-ANC achieves best overall robustness by enforcing the robustly disentangled feature space*. Results for DGCNN, PCT are in Appendix D.2.

### 5.3 Ablation Study

To address **RQ2**, we systematically investigate the impact of each component in 3D-ANC using PointNet as follows.

**ETF** (1) The ETF classifier *significantly enhances robustness* by enforcing feature orthogonality. The average ACC of **ETF** improves by 38.1% compared to **Vanilla**. (2) However, the fixed global orientation limits adaptability on imbalanced data, resulting in a 0.6% drop in clean ACC, suggesting *a trade-off between robustness and generalization*.

**RBL** (1) RBL facilitates ETF by allowing learnable rotation, enabling ETF to better align with the sample features and thereby *preventing clean performance degradation*, as seen from the 3.7% improvement from **ETF** to **ETF & RBL**. (2) However, it also introduces *instability in feature consistency*, which weakens adversarial robustness.

**FDL** FDL promotes intra-class clustering and inter-class repulsion, yet *its effectiveness depends on the feature quality*.

(1) **ETF** vs. **ETF & FDL**: We observe that with misaligned features, FDL leads to suboptimal feature directions, reducing both clean ACC and robustness. (2) **ETF & RBL** vs. **All**: Conversely, when RBL provides better-aligned features, FDL effectively enhances class separation, thereby improving robustness. (3) With more advanced architectures like DGCNN and PCT, where feature space is more structured, FDL further improves both clean ACC and robustness. Results and visualizations can be found in Appendix D.3, D.5.




Figure 4: Case study for highly similar classes.

### 5.4 Quality of Feature Space

To address **RQ3**, we analyze the correlation between feature space quality and adversarial robustness, as depicted in Fig. 1. Specifically, we quantify feature quality using *Silhouette Coefficient (SC)*, which captures intra-class compactness and inter-class separability. Robustness is measured by average ACC in Table 1. The results indicate that a higher-quality feature space substantially enhances robustness, as adversarial perturbations struggle to shift samples across well-separated decision boundaries. Our 3D-ANC can effectively improve the feature quality of existing models.

### 5.5 More Results

**Case Study** Fig. 8 illustrates that the feature distributions of certain classes frequently overlap, as exemplified by the orange and cyan classes in the lower center of the **All** subplot. We further show how inter-class geometric similarity can challenge the training of a raw NC model.

Specifically, in Fig. 4, we examine two pairs of classes: *desk vs. table* and *night stand vs. dresser*. Given the high similarity between samples in these class pairs, the feature extractor naturally encounters difficulties in distinguishing them, a task that even humans might struggle with.

From Fig. 8, 3D-ANC makes significant efforts to help the feature extractor differentiate between these classes, leading to partially separated feature strips. However, *such inherent dataset flaws limit defenses from achieving optimal robustness with ideal feature separation*.

**Efficiency** 3D-ANC replaces the classifier with an ETF head, which introduces almost no additional cost compared to other defenses, as detailed in Appendix D.4.

## 6 Conclusion

In summary, our work introduces Neural Collapse for robust point cloud recognition. Our analysis reveals that existing models and defenses suffer from inadequate feature disentanglement, which renders them susceptible to adversarial attacks. To this end, we propose 3D-ANC, which replaces model classifier with an ETF head to achieve feature disentanglement. Recognizing the unique challenges posed by class imbalance and complex geometric similarities in point cloud data, we further propose an adaptive training framework that integrates RBL and dynamic FDL. The framework effectively boosts robustness, surpassing current defenses. Moreover, 3D-ANC is model-agnostic and efficient, allowing for integration with any existing model architecture.

## Acknowledgements

We would like to thank the anonymous reviewers for their insightful comments that helped improve the quality of the paper. We also appreciate the researchers who contributed to the collection and curation of public datasets. This work was supported in part by the National Natural Science Foundation of China (62472096, 62172104, 62172105, 62102093, 62102091, 62302101, 62202106, 6250074640). Min Yang is a faculty of the Shanghai Institute of Intelligent Electronics & Systems and Engineering Research Center of Cyber Security Auditing and Monitoring, Ministry of Education, China.

## References

- Abbasi, R.; Bashir, A. K.; Alyamani, H. J.; Amin, F.; Doh, J.; and Chen, J. 2022. Lidar point cloud compression, processing and learning for autonomous driving. *IEEE Transactions on Intelligent Transportation Systems*, 24(1): 962–979.
- Carlini, N.; and Wagner, D. 2017. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*, 39–57. Ieee.
- Chang, A. X.; Funkhouser, T.; Guibas, L.; Hanrahan, P.; Huang, Q.; Li, Z.; Savarese, S.; Savva, M.; Song, S.; Su, H.; et al. 2015. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*.
- Ding, D.; Jiang, E.; Huang, Y.; Zhang, M.; Li, W.; and Yang, M. 2023. CAP: Robust Point Cloud Classification via Semantic and Structural Modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12260–12270.
- Duan, H.; Wang, P.; Huang, Y.; Xu, G.; Wei, W.; and Shen, X. 2021. Robotics dexterous grasping: The methods based on point cloud and deep learning. *Frontiers in Neuro-robotics*, 15: 658280.
- Fang, C.; He, H.; Long, Q.; and Su, W. J. 2021. Exploring deep neural networks via layer-peeled model: Minority collapse in imbalanced training. *Proceedings of the National Academy of Sciences*, 118(43): e2103091118.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2014. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.
- Guo, M.-H.; Cai, J.-X.; Liu, Z.-N.; Mu, T.-J.; Martin, R. R.; and Hu, S.-M. 2021. Pct: Point cloud transformer. *Computational Visual Media*, 7: 187–199.
- Hamdi, A.; Giancola, S.; and Ghanem, B. 2021. Mvtn: Multi-view transformation network for 3d shape recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1–11.
- Hamdi, A.; Rojas, S.; Thabet, A.; and Ghanem, B. 2020. Advpc: Transferable adversarial perturbations on 3d point clouds. In *European Conference on Computer Vision*, 241–257. Springer.
- Han, X. Y.; Pappas, V.; and Donoho, D. L. 2022. Neural Collapse Under MSE Loss: Proximity to and Dynamics on the Central Path. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Huang, Q.; Dong, X.; Chen, D.; Zhou, H.; Zhang, W.; and Yu, N. 2022. Shape-invariant 3D Adversarial Point Clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15335–15344.
- Ji, W.; Lu, Y.; Zhang, Y.; Deng, Z.; and Su, W. J. 2022. An Unconstrained Layer-Peeled Perspective on Neural Collapse. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Kim, J.; Hua, B.-S.; Nguyen, T.; and Yeung, S.-K. 2021. Minimal adversarial examples for deep learning on 3d point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7797–7806.
- Lang, A. H.; Vora, S.; Caesar, H.; Zhou, L.; Yang, J.; and Beijbom, O. 2019. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12697–12705.
- Li, Y.; Bu, R.; Sun, M.; Wu, W.; Di, X.; and Chen, B. 2018. Pointcnn: Convolution on x-transformed points. *Advances in neural information processing systems*, 31.
- Liang, Q.; Li, Q.; Nie, W.; and Liu, A.-A. 2022. PAGN: perturbation adaption generation network for point cloud adversarial defense. *Multimedia Systems*, 1–9.
- Liu, D.; Yu, R.; and Su, H. 2019. Extending adversarial attacks and defenses to deep 3d point cloud classifiers. In *2019 IEEE International Conference on Image Processing (ICIP)*, 2279–2283. IEEE.
- Liu, D.; Yu, R.; and Su, H. 2020. Adversarial shape perturbations on 3d point clouds. In *European Conference on Computer Vision*, 88–104. Springer.
- Lou, T.; Jia, X.; Gu, J.; Liu, L.; Liang, S.; He, B.; and Cao, X. 2024. Hide in Thicket: Generating Imperceptible and Rational Adversarial Perturbations on 3d Point Clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24326–24335.
- Lu, H.; and Shi, H. 2020. Deep learning for 3D point cloud understanding: a survey. *arXiv preprint arXiv:2009.08920*.
- Lu, J.; and Steinerberger, S. 2022. Neural collapse under cross-entropy loss. *Applied and Computational Harmonic Analysis*, 59: 224–241.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2017. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*.
- Pappas, V.; Han, X. Y.; and Donoho, D. L. 2020. Prevalence of Neural Collapse During The Terminal Phase Of Deep Learning Training. *Proceedings of the National Academy of Sciences*, 117(40): 24652–24663.
- Peifeng, G.; Xu, Q.; Wen, P.; Yang, Z.; Shao, H.; and Huang, Q. 2023. Feature Directions Matter: Long-tailed Learning Via Rotated Balanced Representation.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017a. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 652–660.

- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30.
- Qiu, S.; Anwar, S.; and Barnes, N. 2021. Geometric back-projection network for point cloud classification. *IEEE Transactions on Multimedia*, 24: 1943–1955.
- Rusu, R. B.; Marton, Z. C.; Blodow, N.; Dolha, M.; and Beetz, M. 2008. Towards 3D point cloud based object maps for household environments. *Robotics and Autonomous Systems*, 56(11): 927–941.
- Sun, J.; Cao, Y.; Choy, C. B.; Yu, Z.; Anandkumar, A.; Mao, Z. M.; and Xiao, C. 2021. Adversarially Robust 3d Point Cloud Recognition Using Self-Supervisions. *Advances in Neural Information Processing Systems*, 34: 15498–15512.
- Sun, J.; Wang, J.; Nie, W.; Yu, Z.; Mao, Z.; and Xiao, C. 2023. A Critical Revisit of Adversarial Robustness in 3d Point Cloud Recognition with Diffusion-driven Purification. In *International Conference on Machine Learning*, 33100–33114. PMLR.
- Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I.; and Fergus, R. 2014. Intriguing properties of neural networks. In *2nd International Conference on Learning Representations, ICLR 2014*.
- Thrapoulidis, C.; Kini, G. R.; Vakilian, V.; and Behnia, T. 2022. Imbalance trouble: Revisiting neural-collapse geometry. *Advances in Neural Information Processing Systems*, 35: 27225–27238.
- Tirer, T.; and Bruna, J. 2022. Extended unconstrained features model for exploring deep neural collapse. In *International Conference on Machine Learning*, 21478–21505. PMLR.
- Tsai, T.; Yang, K.; Ho, T.-Y.; and Jin, Y. 2020. Robust adversarial objects against deep learning models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 954–962.
- Wang, Y.; Sun, Y.; Liu, Z.; Sarma, S. E.; Bronstein, M. M.; and Solomon, J. M. 2019. Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)*, 38(5): 1–12.
- Wen, Y.; Lin, J.; Chen, K.; Chen, C. P.; and Jia, K. 2020. Geometry-aware generation of adversarial point clouds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Wicker, M.; and Kwiatkowska, M. 2019. Robustness of 3d deep learning in an adversarial setting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11767–11775.
- Wu, Z.; Song, S.; Khosla, A.; Yu, F.; Zhang, L.; Tang, X.; and Xiao, J. 2015. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1912–1920.
- Xiang, C.; Qi, C. R.; and Li, B. 2019. Generating 3d adversarial point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9136–9144.
- Xie, J.; Zheng, Z.; Gao, R.; Wang, W.; Zhu, S.-C.; and Wu, Y. N. 2020. Generative VoxelNet: learning energy-based models for 3D shape synthesis and analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Xie, L.; Yang, Y.; Cai, D.; and He, X. 2023. Neural collapse inspired attraction–repulsion-balanced loss for imbalanced learning. *Neurocomputing*, 527: 60–70.
- Yang, Y.; Chen, S.; Li, X.; Xie, L.; Lin, Z.; and Tao, D. 2022. Inducing Neural Collapse In Imbalanced Learning: Do We Really Need A Learnable Classifier At The End of Deep Neural Network? *Advances in Neural Information Processing Systems*, 35: 37991–38002.
- Yu, L.; Li, X.; Fu, C.-W.; Cohen-Or, D.; and Heng, P.-A. 2018. Pu-net: Point cloud upsampling network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2790–2799.
- Zhang, H.; and Wang, J. 2019. Defense Against Adversarial Attacks Using Feature Scattering-based Adversarial Training. In *NeurIPS*.
- Zhang, J.; Chen, L.; Ouyang, B.; Liu, B.; Zhu, J.; Chen, Y.; Meng, Y.; and Wu, D. 2022. Pointcutmix: Regularization Strategy for Point Cloud Classification. *Neurocomputing*, 505: 58–67.
- Zheng, T.; Chen, C.; Yuan, J.; Li, B.; and Ren, K. 2019. Pointcloud saliency maps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1598–1606.
- Zhong, Z.; Cui, J.; Yang, Y.; Wu, X.; Qi, X.; Zhang, X.; and Jia, J. 2023. Understanding imbalanced semantic segmentation through neural collapse. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 19550–19560.
- Zhou, H.; Chen, D.; Liao, J.; Chen, K.; Dong, X.; Liu, K.; Zhang, W.; Hua, G.; and Yu, N. 2020. Lg-gan: Label guided adversarial network for flexible targeted attack of point cloud based deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10356–10365.
- Zhou, H.; Chen, K.; Zhang, W.; Fang, H.; Zhou, W.; and Yu, N. 2019. Dup-net: Denoiser and upsampler network for 3d adversarial point clouds defense. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1961–1970.
- Zhou, J.; Li, X.; Ding, T.; You, C.; Qu, Q.; and Zhu, Z. 2022. On the optimization landscape of neural collapse under mse loss: Global optimality with unconstrained features. In *International Conference on Machine Learning*, 27179–27202. PMLR.