

Verifiability and Predictability: Interpreting Utilities of Network Architectures for Point Cloud Processing

Wen Shen^{2,*}, Zhihua Wei^{2,*}, Shikun Huang², Binbin Zhang², Panyue Chen², Ping Zhao²,
Quanshi Zhang^{1,†}

¹Shanghai Jiao Tong University, Shanghai, China

²Tongji University, Shanghai, China

{wen.shen, zhihua.wei, hsk, 0206zbb, 2030793, zhaoping}@tongji.edu.cn, zqs1022@sjtu.edu.cn

Abstract

In this paper, we diagnose deep neural networks for 3D point cloud processing to explore utilities of different intermediate-layer network architectures. We propose a number of hypotheses on the effects of specific intermediate-layer network architectures on the representation capacity of DNNs. In order to prove the hypotheses, we design five metrics to diagnose various types of DNNs from the following perspectives, information discarding, information concentration, rotation robustness, adversarial robustness, and neighborhood inconsistency. We conduct comparative studies based on such metrics to verify the hypotheses. We further use the verified hypotheses to revise intermediate-layer architectures of existing DNNs and improve their utilities. Experiments demonstrate the effectiveness of our method. The code will be released when this paper is accepted.

1. Introduction

Recently, a series of works use deep neural networks (DNNs) for 3D point cloud processing and have achieved superior performance in various 3D tasks. However, traditional studies usually designed intermediate-layer architectures based on empiricism. Exploring and verifying utilities of each specific intermediate-layer architecture from the perspective of a DNN’s representation capacity still present significant challenges for state-of-the-art algorithms.

In this study, we aim to bridge the gap between the intermediate-layer architecture and its utilities. Table 1 lists three kinds of utilities considered in this study, including rotation robustness, adversarial robustness, and neighborhood inconsistency. Although there are many heuristic insights

*Wen Shen and Zhihua Wei have equal contributions.

†Quanshi Zhang is the corresponding author. He is with the John Hopcroft Center and the MoE Key Lab of Artificial Intelligence, AI Institute, at the Shanghai Jiao Tong University, China.

Verifiability

Intermediate-layer architectures	Rotation	Adversarial	Neighborhood
	robustness	robustness	consistency
(a) Modules of using information of local density to reweight features [40].	–	✓	–
(b) Modules of using information of local coordinates to reweight features [40].	✓	–	–
(c) Modules of concatenating multi-scale features [18].	–	✓	✓
(d) Modules of computing orientation-aware features [11].	✓	–	–

Predictability




Table 1. Illustration of the verified utilities of specific intermediate-layer architectures. “–” denotes that the utility has not been examined, instead of indicating non-existence of the utility. Please see Fig. 1 for architectural details.

on utilities of existing architectures for 3D point cloud processing, there does not exist a rigorous and quantitative verification of such insights.

Therefore, we propose a method to quantitatively diagnose the utilities of intermediate-layer network architectures, which will provide new insights into architectural design. This is a necessary step towards the deep learning with scientific rigour. Note that, utilities are not necessarily equivariant to advantages. For example, in most cases, the rotation robustness is supposed to be a good property. However, the rotation robustness sometimes requires a DNN not to encode rotation-sensitive but discriminative features.

This study focuses on two terms, *i.e.* verifiability and predictability. In terms of verifiability, we design new metrics to quantify utilities of existing intermediate-layer architectures to prove intuitive insights. In terms of predictability, we further use the verified insights to revise other networks to improve their utilities. Note that, the revision of intermediate-layer architectures generally does not change the depth of DNNs, so that we eliminate the influence of the depth change.

More specifically, we propose a few hypotheses of utilities of specific intermediate-layer architectures, as shown in Table 1. Theoretically, we could analyze specific intermediate-layer architectures *w.r.t.* all utilities. However, due to the limit of the page number, we only verify hypotheses with strong connection to human intuitions. We design and conduct comparative studies to verify these hypotheses. The verified hypotheses are further used to guide the architectural revision of existing DNNs to improve their utilities. The verified hypotheses can be summarized as follows.

- The specific module in [40], which uses the local density information to reweight features (Fig. 1 (a)), improves the adversarial robustness (Table 1 (a)).
- Another specific module in [40], which uses local 3D coordinates’ information to reweight features (Fig. 1 (b)), improves the rotation robustness (Table 1 (b)).
- The specific module in [25, 18], which extracts multi-scale features (Fig. 1 (c)), improves the adversarial robustness and the neighborhood consistency (Table 1 (c)). Neighborhood consistency measures whether a DNN assigns similar attention to neighboring points.
- The specific module in [11], which encodes the information of different orientations (Fig. 1 (d)), improves the rotation robustness (Table 1 (d)).

In order to verify the above hypotheses, we design the following five evaluation metrics and conduct a number of comparative experiments to quantify utilities of different intermediate-layer architectures.

1. *Information discarding and 2. information concentration*: Information discarding measures how much information of an input point cloud is forgotten during the computation of a specific intermediate-layer feature. From the perspective of information propagation, the forward propagation through layers can be regarded as a hierarchical process of discarding input information [31]. Ideally, a DNN is supposed to discard information that is not related to the task. Let us take the task of object classification for example. The information of foreground points is usually supposed to be related to the task, while that of background points is not related to the task and is discarded.

To this end, we further propose information concentration to measure the gap between the information related to the task and the information not related to the task. Information concentration can be used to evaluate a DNN’s ability to focus on points related to the task.

3. *Rotation robustness*: Rotation robustness measures whether a DNN will use the same logic to recognize the same object when a point cloud has been rotated by a random angle. In other words, if two point clouds have the same global shape but different orientations, the DNN is supposed to select the same regions/points to compute the intermediate-layer feature. Unlike images with rich color information, point clouds usually only use spatial contexts

for classification. Therefore, a well-trained DNN is supposed to have the rotation robustness.

4. *Adversarial robustness*: A reliable DNN is supposed to be robust to adversarial attacks.

5. *Neighborhood inconsistency*¹: Neighborhood inconsistency measures whether adjacent points have similar importance in the computation of an intermediate-layer feature. Adjacent points in a 3D object usually have similar shape contexts, so they are supposed to have similar importance. Therefore, ideally, a well-trained DNN should have a low value of neighborhood inconsistency.

The verified hypotheses are then applied to existing DNNs to revise their intermediate-layer architectures and improve their utilities. Note that this study aims to verify some insights about intermediate-layer architectures in the scenario of object classification, in order to improve utilities of existing DNNs. The classification accuracy is reported in supplementary materials.

Note that in comparative studies, unnecessarily complex intermediate-layer architectures usually bring in additional uncertainty, which will prevent our experiments from obtaining reliable and rigorous results. Therefore, we conduct experiments on simple-yet-classic intermediate-layer architectures.

Contributions of our study are summarized as follows. (1) We propose a few hypotheses on utilities of specific intermediate-layer architectures. (2) We design five metrics to conduct comparative studies to verify these hypotheses, which provide new insights into architectural utilities. (3) It is proved that the verified hypotheses can be used to revise existing DNNs to improve their utilities.

2. Related work

Deep learning on 3D Point Cloud: Recently, many approaches use DNNs for 3D point cloud processing and have exhibited superior performance in various 3D tasks [24, 33, 36, 43, 42, 8, 37, 13, 30, 17, 29]. PointNet [24] was a pioneer in this direction, which used a max pooling layer to aggregate all individual point features into a global feature. However, such architecture fell short of capturing local features. PointNet++ [25] hierarchically used PointNet as a local descriptor to extract contextual information. Some studies [11, 39, 13, 19] further improved the networks’ ability to capture local geometric features. Some studies used graph convolutional neural networks for 3D point cloud processing [32, 39]. Others focused on the correlations between different regions of the 3D point cloud [18] or interaction between points [49]. In comparison, our study focuses on the utility analysis of intermediate-layer network architectures for point cloud processing.

¹Values of the rotation robustness and the neighborhood consistency are negative numbers. For intuitive comparisons, we showed results of rotation non-robustness and neighborhood inconsistency in Section 5.

Visualization or diagnosis of representations: The most intuitive way to interpret DNNs is the visualization of visual patterns corresponding to a feature map or the network output [44, 22, 6, 51], such as gradient-based methods [7, 28], and the estimation of the saliency map [27, 20, 12, 24, 50]. In comparison, our study aims to explore the utility of intermediate-layer network architectures by diagnosing the information-processing logic of DNNs.

Quantitative evaluation of representations: In the field of explainable AI, explaining the capacity of representations has attracted increasing research attention. Some studies aimed to disentangle features of a DNN into quantifiable and interpretable feature components [48, 47, 46]. Some studies quantified the representation similarity to help understand the neural networks [9, 14, 23, 26]. Li *et al.* [15] quantitated the importance of different feature dimensions to guide model compression. Zhang *et al.* [45] quantitated the significance of interactions among multiple input variables of the DNN. Other studies explained the representation capacity of DNNs [4, 16, 38, 45]. The information-bottleneck theory [35, 31, 3] explained the trade-off between the information compression and the discrimination power of features in a neural network. Achille and Soatto [1] designed an information Dropout layer and quantified the information transmitted through it. Ma *et al.* [21] presented a method to calculate the entropy of the input information. Inspired by [21], we propose five metrics to diagnose feature representations of different DNNs and explore utilities of different intermediate-layer network architectures.

3. Metrics to Diagnose Networks

3.1. Preliminaries: quantification of entropy-based information discarding

We extend the method of calculating the entropy of the input information, which is proposed in [21], as the technical foundation. Based on this, a number of new metrics are designed to diagnose the DNN. The method quantifies the discarding of the input information during the layerwise forward propagation by computing the entropy of the input information given the specific feature of an intermediate layer. Given a point cloud X , let $f = h(X)$ denote the feature of a specific intermediate layer. It is assumed that f' represents the same object concept² as f when f' satisfies $\|f' - f\|^2 < \epsilon$, where feature $f' = h(X')$, $X' = X + \delta$. δ denotes a random noise. Given a specific feature, the conditional entropy of the input information is computed, when the input represents a specific object concept. *I.e.* we calculate entropy $H(X')$, *s.t.* $\|f' - f\|^2 < \epsilon$. It is assumed that X' follows a Gaussian distribution $X' \sim \mathcal{N}(X, \Sigma = \text{diag}[\sigma_1^2, \sigma_2^2, \dots])$.

²In this study, the concept of an object is referred to as a small range of features that represent the same object instance.

Σ measures the maximum perturbation added to X following the maximum-entropy principle, which subjects to $\|f' - f\|^2 < \epsilon$. Considering the assumption of the i.i.d. dimensions of X' , the overall entropy $H(X')$ can be decomposed into point-wise entropies.

$$\max_{\sigma = [\sigma_1, \sigma_2, \dots]^T} H(X'), \quad \text{s.t. } \|h(X') - f\|^2 < \epsilon^3, \quad (1)$$

where $H(X') = \sum_i H_i$; $H_i = \log \sigma_i + \frac{1}{2} \log(2\pi e)$ denotes the entropy of the i -th point. H_i quantifies how much information of the i -th point can be discarded, when the feature $h(X')$ is required to represent the concept of the target object.

3.2. Five metrics

Metric 1, information discarding: The information discarding is defined as $H(X')$ in Eqn. (1). The information discarding is measured at the point level, *i.e.* H_i , which quantifies how much information of the i -th point is discarded during the computation of an intermediate-layer feature. The point with a lower value of H_i is regarded more important in the computation of the feature.

Metric 2, information concentration: The information concentration is based on the metric of information discarding. The information concentration is used to analyze a DNN's ability to maintain the input information related to the task, and discard redundant information unrelated to the task. *E.g.*, in the task of object classification, background points are usually supposed not to be related to the task and are therefore more likely to be discarded by the DNN. Let $\Lambda^{\text{foreground}}$ denote the set of points in the foreground object in the point cloud X , and let $\Lambda^{\text{background}}$ denote the set of points in the background. Information concentration can be computed as the relative background information discarding *w.r.t.* foreground information discarding.

$$\mathbb{E}_{i \in \Lambda^{\text{background}}} [H_i] - \mathbb{E}_{i \in \Lambda^{\text{foreground}}} [H_i], \quad (2)$$

where a higher value of information concentration indicates that the DNN concentrates more on the foreground points during the computation of the feature.

Note that most widely used benchmark datasets for point cloud classification only contain foreground objects. Therefore, we generate a new dataset, where each point cloud contains both the foreground object and the background. In this new dataset, the background is composed of points that are irrelevant to the foreground. We will introduce details in Section 5.

Metric 3, rotation robustness: Given the two point clouds with the same shape but different orientations, the

³We follow the [21] to slightly adjust the value of λ to make the learned σ satisfy that $\mathbb{E}_{f'} [\|f' - f\|^2]$ is about twice the inherent variance of intermediate-layer features subject to a small input noise. λ is a hyperparameter defined in [21].




Figure 1. Illustration of the specific intermediate-layer architectures. Please see texts in Section 4.1 for architectural details.

rotation robustness is proposed to measure whether a DNN uses similar subsets of two point clouds to compute the intermediate-layer feature. Let $X_1 = \theta_1(X)$ and $X_2 = \theta_2(X)$ denote the point clouds that have the same shape but different orientations, where θ_1 and θ_2 denote two different rotation operations. To quantify the similarity of the attention on the two point clouds, we compute the Jensen-Shannon divergence between the distributions of information discarding in Eqn. (1), δ_1 and δ_2 w.r.t. $X'_1 = \theta_1(X + \delta_1)$, $X'_2 = \theta_2(X + \delta_2)$. *I.e.* we measure whether the DNN ignores similar sets of points to compute features of the two point clouds.

$$JSD(\delta_1 || \delta_2), \quad \text{s.t.} \quad \begin{cases} ||h(X'_1) - h(X_1)||^2 < \epsilon \\ ||h(X'_2) - h(X_2)||^2 < \epsilon \end{cases} \quad (3)$$

where $JSD(\delta_1 || \delta_2)$ measures the dissimilarity between information distributions over the two point clouds.

The rotation non-robustness¹ is defined as the average of the dissimilarity of attention on any two point clouds with different orientations, *i.e.* $\mathbb{E}_{\forall \theta_1, \theta_2} [JSD(\delta_1 || \delta_2)]$. In this study, we use the variational-approximation-based method in [10] to approximate the Jensen-Shannon divergence.

Metric 4, adversarial robustness: We use the method in [34] to perform adversarial attacks. The objective is

$$\min ||\epsilon||_2^2, \quad \text{s.t.} \quad C(X + \epsilon) = \hat{l} \neq l^*, \quad (4)$$

where $C(\cdot)$ is the predicted label; l^* is the correct label of X ; \hat{l} is a target incorrect label. In this study, we perform targeted adversarial attacks against all incorrect classes. We use the average of $||\epsilon||_2$ over all incorrect classes to measure the adversarial robustness.

Metric 5, neighborhood inconsistency¹: The neighborhood inconsistency is proposed to evaluate a DNN’s ability to assign similar attention to neighboring points during the computation of an intermediate-layer feature. Ideally, for a DNN, except for special points (*e.g.* those on the edge), most neighboring points in a small region of a point cloud usually have similar shape contexts, so they are supposed to make similar contributions to the classification and receive similar attention, *i.e.* low neighborhood inconsistency. Let $\mathbf{N}(i)$ denote a set of K nearest points of the i -th point. We define the neighborhood inconsistency as the difference between the maximum and minimum point-wise information

discarding within $\mathbf{N}(i)$.

$$\mathbb{E}_i [\max_{j \in \mathbf{N}(i)} H_j - \min_{j \in \mathbf{N}(i)} H_j]. \quad (5)$$

4. Hypotheses and Comparative Study

4.1. Overview of intermediate-layer architectures

• **Notation:** Let $x_i \in \mathbb{R}^3$ denote the i -th point, $i = 1, 2, \dots, n$; let $\mathbf{N}(i)$ denote a set of K nearest points of x_i ; let $\mathbf{F}_i \in \mathbb{R}^{d \times K}$ denote intermediate-layer features of neighboring points $\mathbf{N}(i)$, where each column of \mathbf{F}_i represents the feature of a specific point in $\mathbf{N}(i)$.

• **Architecture 1, features reweighted by the information of the local density:** Architecture 1 focuses on the use of the local density information to reweight features [40]. As shown in Fig. 1 (a), for each point x_i , Architecture 1 uses the local density w.r.t. neighboring points of x_i to compute $\mathbf{W}^{\text{H1}} \in \mathbb{R}^{K \times K}$, which reweights intermediate-layer features \mathbf{F}_i .

$$\mathbf{F}'_i = \mathbf{F}_i \text{diag}[\mathbf{W}^{\text{H1}}], \quad \mathbf{W}^{\text{H1}} = \text{mlp}(\text{density}(\mathbf{N}(i))), \quad (6)$$

where $\text{diag}[\mathbf{W}^{\text{H1}}]$ transforms the vector \mathbf{W}^{H1} into a diagonal matrix; $\text{density}(\mathbf{N}(i))$ is a vector representing the density of neighboring points in $\mathbf{N}(i)$; mlp is a two-layer perceptron network.

• **Architecture 2, features reweighted by the information of local coordinates:** As shown in Fig. 1 (b), for each point x_i , Architecture 2 uses the information of local 3D coordinates to compute $\mathbf{W}^{\text{H2}} \in \mathbb{R}^{M \times K}$ to reweight intermediate-layer features \mathbf{F}_i .

$$\mathbf{F}'_i = \mathbf{F}_i (\mathbf{W}^{\text{H2}})^\top, \quad \mathbf{W}^{\text{H2}} = \text{mlp}(\{x_j | j \in \mathbf{N}(i)\}), \quad (7)$$

where the mlp is a single-layer perceptron network.

• **Architecture 3, multi-scale features:** Architecture 3 focuses on the use of multi-scale contextual information [25, 18]. As illustrated in Fig. 1 (c), $\{\mathbf{F}_i^{\text{scale}=K_1}, \dots, \mathbf{F}_i^{\text{scale}=K_T}\}$ denote features that are extracted using contexts of x_i at different scales, $\mathbf{F}_i^{\text{scale}=K_t} \in \mathbb{R}^{d \times K_t}$. Each specific context w.r.t. x_i is composed of K_t nearest neighboring points around x_i . Then, $f_{i,(\text{scale}=K_t)}^{\text{upper}} \in \mathbb{R}^D$ in the upper layer is computed using $\mathbf{F}_i^{\text{scale}=K_t}$. Architecture 3

concatenates these multi-scale features to obtain f_i^{upper} .

$$f_i^{\text{upper}} = \text{concat} \left\{ \begin{array}{l} f_{i,(\text{scale}=K_1)}^{\text{upper}} \\ f_{i,(\text{scale}=K_2)}^{\text{upper}} \\ \dots \\ f_{i,(\text{scale}=K_T)}^{\text{upper}} \end{array} \right\}, \quad (8)$$

$$f_{i,(\text{scale}=K_t)}^{\text{upper}} = g(\mathbf{F}_i^{\text{scale}=K_t}),$$

where $g(\cdot)$ is a function for feature extraction. Details about this function are introduced in [24], which are also summarized in supplementary materials⁴.

• **Architecture 4, orientation-aware features:** Architecture 4 focuses on the use of orientation information [11]. As illustrated in Fig. 1 (d), for each point x_i , $\mathbf{F}_i^{\text{oe}} \in \mathbb{R}^{d \times O}$ denotes the feature of x_i , which encodes the information of various orientations, where O is the number of orientations. Architecture 4 uses \mathbf{F}_i^{oe} to compute the orientation-aware feature $f_i^{\text{oe}} \in \mathbb{R}^d$.

$$f_i^{\text{oe}} = \text{Conv}^{\text{oe}}(\mathbf{F}_i^{\text{oe}}), \quad (9)$$

where Conv^{oe} is a special convolution operator. Details about this operator and the computation of f_i^{oe} are introduced in [11]³.

4.2. Four hypotheses and comparative study design

Hypothesis 1: Architecture 1 designed by [40] (in Fig. 1 (a)) increases adversarial robustness.

This hypothesis is based on the observation that PointConv [40] has strong adversarial robustness, which may stem from Architecture 1. To verify this, we construct two versions of the PointConv for comparison, *i.e.* one with Architecture 1 and the other without Architecture 1.

To obtain the PointConv without Architecture 1, we remove⁵ all the modules of Architecture 1 from the original network (see the footnote⁶), which are located behind the 2-nd, 5-th, 8-th, 11-th, and 14-th nonlinear transformation layers. The global architecture of PointConv is introduced in [40]³.

$$f_i^{\text{upper}} = \text{mlp}(\mathbf{F}_i) \text{diag}[\mathbf{W}^{\text{H}_1}] \implies f_i^{\text{upper}} = \text{mlp}(\mathbf{F}_i), \quad (10)$$

where f_i^{upper} is the feature in the upper layer; $\text{diag}[\mathbf{W}^{\text{H}_1}]$ transforms the vector \mathbf{W}^{H_1} into a diagonal matrix.

⁴For the convenience of readers to quickly understand relevant technologies in original papers, we summarize these relevant technologies in supplementary materials.

⁵Note that removing or adding modules of these specific intermediate-layer network architectures generally has no effects on the depth of DNNs, so that we eliminate the influence of changes in DNNs' depth.

⁶The PointConv for classification is revised from the code for segmentation released by [40].

Hypothesis 2: Architecture 2 designed by [40] (in Fig. 1 (b)) increases rotation robustness.

This hypothesis is proposed based on the observation that PointConv [40] has strong rotation robustness, which may stem from Architecture 2. To verify this, we construct two versions PointConv for comparison, *i.e.* one with Architecture 2 and the other without Architecture 2.

To obtain the PointConv without Architecture 2, we remove⁴ all the modules of Architecture 2, which are located before the 3-rd, 6-th, 9-th, 12-th, and 15-th nonlinear transformation layers. The global architecture of PointConv is introduced in [40]³.

$$f_i^{\text{upper}} = \text{mlp}(\mathbf{F}_i)(\mathbf{W}^{\text{H}_2})^\top \implies f_i^{\text{upper}} = \text{mlp}(\mathbf{F}_i). \quad (11)$$

Hypothesis 3: Architecture 3 used in [25, 18] (in Fig. 1 (c)) increases adversarial robustness and neighborhood consistency.

This hypothesis is inspired by [25, 18], which encodes multi-scale contextual information. To verify this hypothesis, we construct⁴ three versions of Point2Sequence for comparison. The baseline network of Point2Sequence concatenates features of 4 different scales to compute the feature in the upper layer, $\{f_{i,(\text{scale}=K_1)}^{\text{upper}}, f_{i,(\text{scale}=K_2)}^{\text{upper}}, f_{i,(\text{scale}=K_3)}^{\text{upper}}, f_{i,(\text{scale}=K_4)}^{\text{upper}}\}$. In this study, we set $K_1 = 128$, $K_2 = 64$, $K_3 = 32$, and $K_4 = 16$. The first network extracts features with three different scales, $\{f_{i,(\text{scale}=K_1)}^{\text{upper}}, f_{i,(\text{scale}=K_2)}^{\text{upper}}, f_{i,(\text{scale}=K_3)}^{\text{upper}}\}$, and the second one extracts features with two different scales, $\{f_{i,(\text{scale}=K_1)}^{\text{upper}}, f_{i,(\text{scale}=K_2)}^{\text{upper}}\}$. The global architecture of Point2Sequence is introduced in [18]³.

Hypothesis 4: Architecture 4 designed by [11] (in Fig. 1 (d)) increases rotation robustness.

This hypothesis is proposed based on the observation that PointSIFT [11] exhibits strong rotation robustness. It may be because Architecture 4 ensures that features contain information from various orientations. To verify this hypothesis, we construct two versions of the PointSIFT for comparisons, *i.e.* one with Architecture 4 and the other without Architecture 4.

To get the PointSIFT without Architecture 4, we remove all modules of Architecture 4 from the original network (see the footnote⁷), which are located before the 1-st, 3-rd, 5-th, and 7-th nonlinear transformation layers. The global architecture of PointSIFT is introduced in [11]³.

⁷The PointSIFT for classification is revised from the code for segmentation released by [11].

	# of added modules	Locations of added modules	
(a)	Add Architecture 1 to PointNet++ for adversarial robustness	3	behind the 3-rd, 6-th, and 9-th nonlinear transformation layers
	Add Architecture 1 to Point2Sequence for adversarial robustness	1	behind the last nonlinear transformation layer
	Add Architecture 1 to RSCNN for adversarial robustness	2	behind the 2-nd and 5-th nonlinear transformation layers
(b)	Add Architecture 2 to PointNet++ for rotation robustness	3	behind the 3-rd, 6-th, and 9-th nonlinear transformation layers
	Add Architecture 2 to Point2Sequence for rotation robustness	1	behind the last nonlinear transformation layer
	Add Architecture 2 to RSCNN for rotation robustness	2	behind the 2-nd and 5-th nonlinear transformation layers
(c)	Add Architecture 3 to PointNet++ for adversarial robustness & neighborhood consistency	2	for 1-st to 3-rd nonlinear transformation layers, $K_1=16$, $K_2=128^*$ for 4-th to 6-th nonlinear transformation layers, $K_1=32$, $K_2=128^*$
	Add Architecture 3 to RSCNN for adversarial robustness & neighborhood consistency	2	for 1-st to 3-rd nonlinear transformation layers, $K_1=16$, $K_2=32$ for 4-th to 6-th nonlinear transformation layers, $K_1=16$, $K_2=48^*$
	Add Architecture 4 to PointNet++ for rotation robustness	1	before the 7-th nonlinear transformation layer
(d)	Add Architecture 4 to Point2Sequence for rotation robustness	1	before the 14-th nonlinear transformation layer
	Add Architecture 4 to RSCNN for rotation robustness	1	before the 7-th nonlinear transformation layer

* K_1 and K_2 are hyper-parameters of added modules of Architecture 3, which have been introduced in Eqn. (8).

Table 2. Adding⁴ specific intermediate-layer architectures to existing DNNs to improve utilities.

4.3. Comparative study for the improvement of utilities of existing DNNs

In this section, we further prove that the verified four hypotheses can be used to revise existing intermediate-layer network architectures in order to improve their utilities. We apply our method to three benchmark DNNs, including PointNet++ [25], Point2Sequence [18], and RSCNN [19].

In Section 4.2, we remove specific intermediate-layer architectures from original DNNs. Actually, if we take the DNN without the specific intermediate-layer architecture as the original one (e.g. the PointConv without Architecture 1) and take the real original DNN as the revised one (e.g. the PointConv with Architecture 1), then it is naturally proved that the verified hypotheses can be used to revise DNNs to improve their utilities.

Nevertheless, in this section, we further prove that these specific intermediate-layer architectures can improve utilities of other DNNs.

Architecture 1 designed by [40] is added to PointNet++, Point2Sequence, and RSCNN in order to improve their adversarial robustness.

For each network, we construct two versions for comparison, *i.e.* one with Architecture 1 and the other without Architecture 1.

Table 2 (a) shows details about how to obtain DNNs with Architecture 1. Global architectures of PointNet++, Point2Sequence, and RSCNN are introduced in [25]³, [18]³, and [19]³.

Eqn. (12) shows how to add Architecture 1 behind a nonlinear transformation layer.

$$f_i^{\text{upper}} = \text{mlp}(\mathbf{F}_i) \implies f_i^{\text{upper}} = \text{mlp}(\mathbf{F}_i) \text{diag}[\mathbf{W}^{\text{H}_1}], \quad (12)$$

where \mathbf{W}^{H_1} denotes the formula of Architecture 1, which has been introduced in Eqn. (6).

Architecture 2 designed by [40] is added to PointNet++, Point2Sequence, and RSCNN in order to improve their rotation robustness.

Table 2 (b) shows details about how to obtain DNNs with Architecture 2. Global architectures of PointNet++, Point2Sequence, and RSCNN are introduced in [25]³, [18]³, and [19]³.

Eqn. (13) shows how to add Architecture 2 behind a nonlinear transformation layer.

$$f_i^{\text{upper}} = \text{mlp}(\mathbf{F}_i) \implies f_i^{\text{upper}} = \text{mlp}(\mathbf{F}_i)(\mathbf{W}^{\text{H}_2})^\top, \quad (13)$$

where \mathbf{W}^{H_2} denotes the formula of Architecture 2, which has been introduced in Eqn. (7).

Architecture 3 used in [18] is added to PointNet++ and RSCNN in order to improve their adversarial robustness and neighborhood consistency.

Table 2 (c) shows details about how to obtain DNNs with Architecture 3.

Architecture 4 designed by [11] is added to PointNet++, Point2Sequence, and RSCNN in order to improve their rotation robustness.

Table 2 (d) shows details about how to obtain DNNs with Architecture 4.

5. Experiments

To demonstrate the broad applicability of our method, we applied our method to diagnose seven widely used DNNs, including PointNet [24], PointNet++ [25], PointConv [40], DGCNN [39], PointSIFT [11], Point2Sequence [18], and RSCNN [19]. These DNNs were trained using three benchmark datasets, including the ModelNet40

Models	Information discarding	Information concentration	Rotation non-robustness ¹	Adversarial robustness	Neighborhood inconsistency ¹
PointNet	-8370.42	1.089	8.000	1.994	3.127
PointNet++	-8166.16	1.625	7.093	2.504	3.409
PointConv	-8766.00	0.865	4.875	2.878	3.781
DGCNN	-9187.51	1.336	5.985	2.421	1.449
PointSIFT	-8415.92	0.079	3.931	2.839	2.423
Point2Sequence	-8328.34	1.321	9.506	2.526	3.184
RSCNN	-8009.52	2.220	3.645	2.314	3.585

Table 3. Quantification of the representation capacity of different DNNs on the ModelNet40 dataset.

dataset [41], the ShapeNet⁸ dataset [2], the 3D MNIST [5] dataset.

Implementation details: When computing the information discarding, we bounded each dimension of $\sigma = [\sigma_1, \sigma_2, \dots]^T$ within 0.08 for fair comparison between different DNNs. It was because when processing a point cloud, some widely used operations (*e.g.* the $g(\cdot)$ operation in Eqn. (8)) would randomly and completely discard the information of some points. This resulted in that the learned σ corresponding to these points could be infinite in theory. For the computation of rotation robustness, during the training and testing phases, each point cloud was rotated by random angles. For the computation of neighborhood inconsistency, we used k -NN search to select 16 neighbors for each point.

To analyze the information concentration of DNNs, we generated a new dataset that contained both the foreground objects and the background, since most widely used benchmark datasets for point cloud classification only contain foreground objects. Specifically, for each sample (*i.e.* the foreground object) in the ModelNet40, we generated the background as follows. First, we randomly sampled a set of 500 points from point clouds, which had different labels from the foreground object. Second, we resized this set of points to the density of the foreground object. Finally, we randomly located it around the foreground object. *The dataset will be released when this paper is accepted.*

The entropy-based method [21] quantified the layer-wise information discarding. This method assumed the feature space of the concept of a specific object satisfied $\|f' - f\|^2 < \epsilon$, where $f = h(X)$, $f' = h(X')$, $X' = X + \delta$. δ denotes a random noise. For point cloud processing, each dimension of the intermediate-layer feature is computed using the context of a specific point x_i . However, adding noise to a point cloud will change the context of each point. In order to extend the entropy-based method to point cloud processing, we selected the same set of points as the contexts *w.r.t.* x_i and x'_i , so as to generate a convincing evaluation⁹.

Comparisons of the representation capacity of DNNs: As shown in Table 3, we measured the proposed metrics

⁸The ShapeNet dataset for classification is converted from the ShapeNet part segmentation dataset.

⁹Detailed discussions are presented in supplementary materials.

for the fully connected layer close to the network output¹⁰, which had 512 hidden units. We measured adversarial robustness by performing adversarial attacks over all incorrect classes. We found that PointNet++ and RSCNN had relatively higher values of information discarding. PointNet++ and RSCNN concentrated more on the foreground object. PointConv, DGCNN, PointSIFT, and RSCNN performed well in rotation robustness. PointConv and PointSIFT exhibited higher adversarial robustness. DGCNN and PointSIFT exhibited lower neighborhood inconsistency.

Verifying hypotheses of utilities of specific network intermediate-layer architectures: As shown in Table 4, the proposed four hypotheses had been verified. Architecture 1 improved the utility of adversarial robustness. One possible reason was that Architecture 1 considered distances between each point and its neighbors during the computing of densities, which increased the difficulty of adversarial attacks. Architecture 3 also improved the utility of adversarial robustness. We found that the utility of adversarial robustness increased with the scale number of features. The reason may be that concatenating features with different scales enhanced the representation capacity, so that it was more challenging to conduct adversarial attacks. Architecture 2 and Architecture 4 improved the utility of rotation robustness. The reason may be that both Architecture 2 and Architecture 4 extracted contextual information from coordinates of each point’s neighbors using non-linear transformations. Such contextual information improved rotation robustness. Besides, networks with Architecture 3 usually had lower neighborhood inconsistency than those without Architecture 3. DNNs that extracted features from contexts of more scales usually exhibited lower neighborhood inconsistency. One possible reason was that extracting multi-scale features enhanced connections between neighboring points.

Improving utilities of existing DNNs: In this experiment, we aimed to prove that utilities of existing DNNs could be improved by using the verified hypotheses to guide the architectural revision. To this end, we conducted comparative studies as designed in Table 2. As shown in Table 5, adding specific intermediate-layer architectures to existing DNNs improved their utilities. Specifically, adding modules of Architecture 1 improved the utility of adversarial robustness of PointNet++, Point2Sequence, and RSCNN. Adding modules of Architecture 2 improved the utility of rotation robustness of PointNet++, Point2Sequence, and RSCNN. Adding modules of Architecture 3 improved utilities of adversarial robustness and neighborhood consistency of both PointNet++ and RSCNN. Adding modules of Architecture 4 improved the utility of rotation robustness of PointNet++, Point2Sequence, and RSCNN.

Relationship between utilities and accuracy: Note that

¹⁰Theoretically, features of any layer can be used to analyze the representation capacity of DNNs.

		ModelNet40 dataset			ShapeNet dataset			3D MNIST dataset		
		w/	w/o	Δ	w/	w/o	Δ	w/	w/o	Δ
Hypothesis 1	Keep/remove Architecture 1 from PointConv for adversarial robustness	2.878	2.629	0.249	2.407	2.271	0.136	2.737	2.530	0.207
Hypothesis 2	Keep/remove Architecture 2 from PointConv for rotation non-robustness	4.875	5.066	0.191	4.470	5.368	0.898	6.650	8.019	1.369
Hypothesis 3	Keep/remove Architecture 3 from Point2Sequence (4 vs. 3 scales) for adversarial robustness	2.526	2.521	0.005	2.520	2.514	0.006	2.468	2.479	-0.011
	Keep/remove Architecture 3 from Point2Sequence (4 vs. 2 scales) for adversarial robustness		2.513	0.013		2.488	0.032		2.460	0.008
	Keep/remove Architecture 3 from Point2Sequence (4 vs. 3 scales) for neighborhood inconsistency	3.184	3.332	0.148	2.815	2.992	0.177	3.097	3.342	0.245
	Keep/remove Architecture 3 from Point2Sequence (4 vs. 2 scales) for neighborhood inconsistency		3.148	-0.036		2.947	0.132		3.431	0.334
Hypothesis 4	Keep/remove Architecture 4 from PointSIFT for rotation non-robustness	3.931	7.274	3.343	3.678	6.223	2.545	6.308	5.619	-0.689

Table 4. Verifying hypotheses of utilities of specific intermediate-layer network architectures. The column Δ denotes the increase of the utility of the network with the specific architecture *w.r.t.* the network without the specific architecture⁴. In particular, for adversarial robustness, Δ was calculated as the adversarial robustness of the network w/ the specific architecture minus the adversarial robustness of the network w/o the specific architecture. For rotation non-robustness and neighborhood inconsistency, Δ was calculated as the rotation non-robustness/neighborhood inconsistency of the network w/o the specific architecture minus the rotation non-robustness/neighborhood inconsistency of the network w/ the specific architecture. $\Delta > 0$ indicates that the corresponding hypothesis has been verified. Experimental results show that the proposed four hypotheses were verified.

	ModelNet40 dataset			ShapeNet dataset			3D MNIST dataset		
	added	ori.	Δ	added	ori.	Δ	added	ori.	Δ
Add Architecture 1 to PointNet++ for adversarial robustness	2.519	2.504	0.015	2.496	2.437	0.059	2.427	2.352	0.075
Add Architecture 1 to Point2Sequence for adversarial robustness	2.544	2.526	0.018	2.500	2.520	-0.020	2.475	2.468	0.007
Add Architecture 1 to RSCNN for adversarial robustness	2.342	2.314	0.028	2.337	2.227	0.110	2.283	2.279	0.004
Add Architecture 2 to PointNet++ for rotation non-robustness	3.845	7.093	3.248	2.921	5.929	3.008	3.143	6.531	3.388
Add Architecture 2 to Point2Sequence for rotation non-robustness	4.963	9.506	4.543	4.017	7.451	3.434	4.354	6.890	2.536
Add Architecture 2 to RSCNN for rotation non-robustness	3.993	3.645	-0.348	4.685	3.460	-1.225	3.391	3.456	0.065
Add Architecture 3 to PointNet++ for adversarial robustness	3.010	2.504	0.506	2.987	2.437	0.550	2.604	2.352	0.252
Add Architecture 3 to RSCNN for adversarial robustness	2.350	2.314	0.036	2.342	2.279	0.063	2.332	2.227	0.105
Add Architecture 3 to PointNet++ for neighborhood inconsistency	3.010	3.409	0.399	3.288	3.352	0.064	3.480	3.541	0.062
Add Architecture 3 to RSCNN for neighborhood inconsistency	3.497	3.585	0.088	3.167	3.478	0.311	3.397	3.928	0.531
Add Architecture 4 to PointNet++ for rotation non-robustness	6.191	7.093	0.902	4.898	5.929	1.031	3.513	6.531	3.018
Add Architecture 4 to Point2Sequence for rotation non-robustness	6.005	9.506	3.501	8.385	7.451	-0.934	9.494	6.933	-2.561
Add Architecture 4 to RSCNN for rotation non-robustness	2.424	3.645	1.221	2.697	3.460	0.763	3.555	3.456	-0.099

Table 5. Improving utilities of existing DNNs by adding modules of specific intermediate-layer architectures. The column “added” denotes the utility of the network which the specific architecture was added to⁴. The column “ori.” denotes the utility of the original network. The column Δ denotes the improvement of the utility of the network which the specific architecture was added to *w.r.t.* the original network. In particular, for adversarial robustness, Δ was calculated as the value of the “added” column minus the value of the “ori.” column. For rotation non-robustness and neighborhood inconsistency, Δ was calculated as the value of the “ori.” column minus the value of the “added” column. $\Delta > 0$ indicates that the specific architecture improves the utility of the DNN. Experimental results show that the verified hypotheses could be used to revise existing DNNs to improve their utilities.

this study focused on the verification of utilities of specific network architectures and the architectural revision of existing DNNs to improve their utilities, instead of the classification accuracy. Nevertheless, we provided the classification accuracy of different versions of DNNs in supplementary materials. *Experimental results show that adding a specific architecture to existing DNNs has effects on accuracy. We have detailedly discussed the relationship between utilities and accuracy in supplementary materials.*

6. Conclusion

In this paper, we have verified a few hypotheses of the utility of four specific intermediate-layer network architectures for 3D point cloud processing. Comparative studies are conducted to prove the utility of the specific architectures, including rotation robustness, adversarial robustness, and neighborhood inconsistency. In preliminary experi-

ments, we have verified that Architecture 2 and Architecture 4 mainly improve the rotation robustness; Architecture 1 and Architecture 3 have positive effects on adversarial robustness; Architecture 3 usually alleviates the neighborhood inconsistency. These verified hypotheses have further been used to revise existing DNNs to improve their utilities.

Considering that unnecessarily complex intermediate-layer architectures will bring in uncertainty of experiments, we only verify utilities of simple network architectures *w.r.t.* the object classification. More generic hypotheses about utilities of other tasks (*e.g.* segmentation and reconstruction) need to be verified in the future.

Acknowledgments This work is partially supported by the National Nature Science Foundation of China (No. 61906120, U19B2043, 61976160).

References

- [1] Alessandro Achille and Stefano Soatto. Information dropout: Learning optimal representations through noisy computation. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2897–2905, 2018. 3
- [2] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 7
- [3] Hao Cheng, Dongze Lian, Shenghua Gao, and Yanlin Geng. Evaluating capability of deep neural networks for image classification via information plane. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 168–182, 2018. 3
- [4] Xu Cheng, Zhefan Rao, Yilan Chen, and Quanshi Zhang. Explaining knowledge distillation by quantifying the knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12925–12935, 2020. 3
- [5] D de la Iglesia Castro. 3d mnist dataset. <https://www.kaggle.com/daavoo/3d-mnist/version/13>, 2016. 7
- [6] Alexey Dosovitskiy and Thomas Brox. Inverting visual representations with convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4829–4837, 2016. 3
- [7] Ruth Fong and Andrea Vedaldi. Net2vec: Quantifying and explaining how concepts are encoded by filters in deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8730–8738, 2018. 3
- [8] Matheus Gadelha, Rui Wang, and Subhransu Maji. Multiresolution tree networks for 3d point cloud processing. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 103–118, 2018. 2
- [9] Akhilesh Gotmare, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. A closer look at deep learning heuristics: Learning rate restarts, warmup and distillation. *arXiv preprint arXiv:1810.13243*, 2018. 3
- [10] John R Hershey and Peder A Olsen. Approximating the kullback leibler divergence between gaussian mixture models. In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP'07*, volume 4, pages IV–317. IEEE, 2007. 4
- [11] Mingyang Jiang, Yiran Wu, Tianqi Zhao, Zelin Zhao, and Cewu Lu. Pointsift: A sift-like network module for 3d point cloud semantic segmentation. *arXiv preprint arXiv:1807.00652*, 2018. 1, 2, 5, 6
- [12] Pieter-Jan Kindermans, Kristof T Schütt, Maximilian Alber, Klaus-Robert Müller, Dumitru Erhan, Been Kim, and Sven Dähne. Learning how to explain neural networks: Patternnet and patternattribution. *arXiv preprint arXiv:1705.05598*, 2017. 3
- [13] Artem Komarichev, Zichun Zhong, and Jing Hua. A-cnn: Annularly convolutional neural networks on point clouds. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [14] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. *arXiv preprint arXiv:1905.00414*, 2019. 3
- [15] Yuchao Li, Shaohui Lin, Baochang Zhang, Jianzhuang Liu, David Doermann, Yongjian Wu, Feiyue Huang, and Rongrong Ji. Exploiting kernel sparsity and entropy for interpretable cnn compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2800–2809, 2019. 3
- [16] Ruofan Liang, Tianlin Li, Longfei Li, Jing Wang, and Quanshi Zhang. Knowledge consistency between neural networks and beyond. *arXiv preprint arXiv:1908.01581*, 2019. 3
- [17] Weiping Liu, Jia Sun, Wanyi Li, Ting Hu, and Peng Wang. Deep learning on point clouds and its application: A survey. *Sensors*, 19(19):4188, 2019. 2
- [18] Xinhai Liu, Zhizhong Han, Yu-Shen Liu, and Matthias Zwicker. Point2sequence: Learning the shape representation of 3d point clouds with an attention-based sequence to sequence network. *arXiv preprint arXiv:1811.02565*, 2018. 1, 2, 4, 5, 6
- [19] Yongcheng Liu, Bin Fan, Shiming Xiang, and Chunhong Pan. Relation-shape convolutional neural network for point cloud analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8895–8904, 2019. 2, 6
- [20] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, pages 4765–4774, 2017. 3
- [21] Haotian Ma, Yinqing Zhang, Fan Zhou, and Quanshi Zhang. Quantifying layerwise information discarding of neural networks. *arXiv preprint arXiv:1906.04109*, 2019. 3, 7
- [22] Aravindh Mahendran and Andrea Vedaldi. Understanding deep image representations by inverting them. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5188–5196, 2015. 3
- [23] Ari Morcos, Maithra Raghu, and Samy Bengio. Insights on representational similarity in neural networks with canonical correlation. In *Advances in Neural Information Processing Systems*, pages 5727–5736, 2018. 3
- [24] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 652–660, 2017. 2, 3, 5, 6
- [25] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems*, pages 5099–5108, 2017. 2, 4, 5, 6
- [26] Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Sohl-Dickstein. Svcca: Singular vector canonical correlation analysis for deep learning dynamics and interpretability. In *Advances in Neural Information Processing Systems*, pages 6076–6085, 2017. 3
- [27] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Why should i trust you?: Explaining the predictions of any

- classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144. ACM, 2016. 3
- [28] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 618–626, 2017. 3
- [29] Wen Shen, Binbin Zhang, Shikun Huang, Zhihua Wei, and Quanshi Zhang. 3d-rotation-equivariant quaternion neural networks. *arXiv preprint arXiv:1911.09040*, 2019. 2
- [30] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointcnn: 3d object proposal generation and detection from point cloud. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–779, 2019. 2
- [31] Ravid Shwartz-Ziv and Naftali Tishby. Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810*, 2017. 2, 3
- [32] Martin Simonovsky and Nikos Komodakis. Dynamic edge-conditioned filters in convolutional neural networks on graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3693–3702, 2017. 2
- [33] Hang Su, Varun Jampani, Deqing Sun, Subhransu Maji, Evangelos Kalogerakis, Ming-Hsuan Yang, and Jan Kautz. Splatnet: Sparse lattice networks for point cloud processing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2530–2539, 2018. 2
- [34] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013. 4
- [35] Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000. 3
- [36] Diego Valsesia, Giulia Fracastoro, and Enrico Magli. Learning localized generative models for 3d point clouds via graph convolution. 2018. 2
- [37] Weiyue Wang, Ronald Yu, Qiangui Huang, and Ulrich Neumann. Sgpn: Similarity group proposal network for 3d point cloud instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2569–2578, 2018. 2
- [38] Xin Wang, Jie Ren, Shuyun Lin, Xiangming Zhu, Yisen Wang, and Quanshi Zhang. A unified approach to interpreting and boosting adversarial transferability. *arXiv preprint arXiv:2010.04055*, 2020. 3
- [39] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *arXiv preprint arXiv:1801.07829*, 2018. 2, 6
- [40] Wenxuan Wu, Zhongang Qi, and Fuxin Li. Pointconv: Deep convolutional networks on 3d point clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9621–9630, 2019. 1, 2, 4, 5, 6
- [41] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. 7
- [42] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 206–215, 2018. 2
- [43] Lequan Yu, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng. Pu-net: Point cloud upsampling network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2790–2799, 2018. 2
- [44] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014. 3
- [45] Hao Zhang, Sen Li, Yinchao Ma, Mingjie Li, Yichen Xie, and Quanshi Zhang. Interpreting and boosting dropout from a game-theoretic view. *arXiv preprint arXiv:2009.11729*, 2020. 3
- [46] Quanshi Zhang, Ruiming Cao, Feng Shi, Ying Nian Wu, and Song-Chun Zhu. Interpreting cnn knowledge via an explanatory graph. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. 3
- [47] Quanshi Zhang, Ying Nian Wu, and Song-Chun Zhu. Interpretable convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8827–8836, 2018. 3
- [48] Quanshi Zhang, Yu Yang, Haotian Ma, and Ying Nian Wu. Interpreting cnns via decision trees. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6261–6270, 2019. 3
- [49] Hengshuang Zhao, Li Jiang, Chi-Wing Fu, and Jiaya Jia. Pointweb: Enhancing local neighborhood features for point cloud processing. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [50] Tianhang Zheng, Changyou Chen, Junsong Yuan, Bo Li, and Kui Ren. Pointcloud saliency maps. *arXiv preprint arXiv:1812.01687*, 2019. 3
- [51] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Object detectors emerge in deep scene cnns. *arXiv preprint arXiv:1412.6856*, 2014. 3