

Contrastive Pre-Training with Multi-View Fusion for No-Reference Point Cloud Quality Assessment

Ziyu Shan¹, Yujie Zhang¹, Qi Yang², Haichen Yang¹, Yiling Xu^{1*}

Jenq-Neng Hwang³, Xiaozhong Xu², Shan Liu²

¹Shanghai Jiao Tong University, ²Tencent, ³University of Washington

¹{shanziyu, yujie19981026, yanghaichen, yl.xu}@sjtu.edu.cn,

²{chinoyang, shanl}@tencent.com, ³hwang@uw.edu

Abstract

No-reference point cloud quality assessment (NR-PCQA) aims to automatically evaluate the perceptual quality of distorted point clouds without available reference, which have achieved tremendous improvements due to the utilization of deep neural networks. However, learning-based NR-PCQA methods suffer from the scarcity of labeled data and usually perform suboptimally in terms of generalization. To solve the problem, we propose a novel contrastive pre-training framework tailored for PCQA (CoPA), which enables the pre-trained model to learn quality-aware representations from unlabeled data. To obtain anchors in the representation space, we project point clouds with different distortions into images and randomly mix their local patches to form mixed images with multiple distortions. Utilizing the generated anchors, we constrain the pre-training process via a quality-aware contrastive loss following the philosophy that perceptual quality is closely related to both content and distortion. Furthermore, in the model fine-tuning stage, we propose a semantic-guided multi-view fusion module to effectively integrate the features of projected images from multiple perspectives. Extensive experiments show that our method outperforms the state-of-the-art PCQA methods on popular benchmarks. Further investigations demonstrate that CoPA can also benefit existing learning-based PCQA models.

1. Introduction

Point clouds have emerged as a prominent 3D multimedia representation in diverse scenarios, such as autonomous driving, digital museum, and immersive gaming [1, 21, 39].

*Corresponding author. This paper is supported in part by National Natural Science Foundation of China (62371290, U20A20185) and 111 project (BP0719010). The corresponding author is Yiling Xu (e-mail: yl.xu@sjtu.edu.cn).

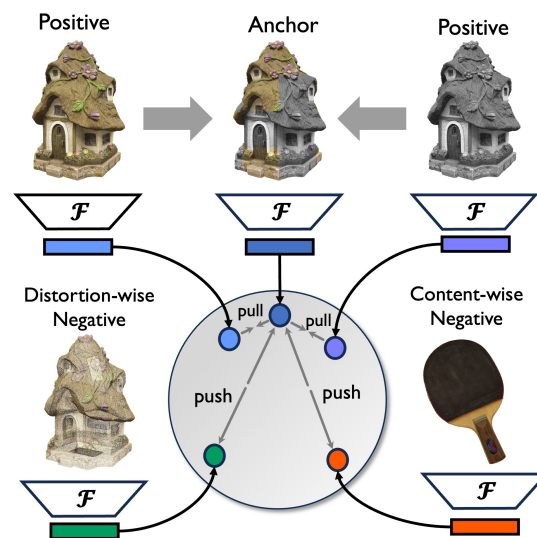


Figure 1. Illustration of our contrastive pre-training framework (CoPA). CoPA first generates anchor by randomly patch-mixing the projected images from a point cloud with different distortions, and then pre-trains the model by pulling positive samples closer to the anchor in the representation space, while pushing distortion-wise and content-wise negative samples apart.

These extensive applications stem from the rich information provided by point clouds (e.g., geometric coordinates, color, and opacity). Nevertheless, point clouds undergo various distortions at any stage of their operation cycle (e.g., acquisition, compression, and transmission) before being delivered to the terminals, leading to inevitable quality degradation and perceptual loss. To optimize the quality of experience in practical applications, point cloud quality assessment (PCQA) has become one of the most fundamental and challenging problems in both industry and academic area.

PCQA methods can be classified into full-reference (FR), reduced-reference (RR) and no-reference (NR) methods, hinging on the availability of high-quality reference point clouds. In this paper, we focus on NR-PCQA con-

sidering that in most cases pristine reference point clouds are not available [44].

Thanks to the tremendous improvements of deep learning networks, learning-based NR-PCQA methods have presented remarkable performance. However, the NR-PCQA problem is far from completely resolved because these data-driven methods suffer from the scarcity of labeled data. Most PCQA datasets [12, 13, 34] only provide hundreds of samples with labels (*i.e.*, mean opinion score (MOS)) due to the high cost of the annotation process, indicating that existing PCQA datasets are too small to train a model with good generalizability. Consequently, it limits the model performance with respect to cross-dataset evaluation.

Researchers have proposed several methods to address this problem. An intriguing way [36] is to take advantage of the rich subjective scores of natural images and infer point cloud quality through unsupervised adversarial domain adaptation. Nevertheless, the 2D-to-3D perceptual adaptation is considerably difficult due to significant discrepancies of content and distortion characteristics between natural images and point clouds. As a result, the learned universal encoder fails to extract adequately effective features to infer point cloud quality. Another common strategy [11, 24] is to employ distortion-related auxiliary tasks (*e.g.*, distortion type prediction) to equip models with the preliminary capability of recognizing distortion patterns. However, the models trained by these auxiliary tasks are not robust when adapted to new datasets with unseen distortion types.

Contrastive learning is another potential choice to promote the generalizability of NR-PCQA models because of its ability to leverage a large amount of unlabeled data. In the classical contrastive learning paradigm, for a given anchor target (often obtained through data augmentation), the model intends to learn representations by pulling similar samples (*i.e.*, positive samples) closer and pushing dissimilar ones (*i.e.*, negative samples) apart in the representation space. After this pre-training, the model is fine-tuned for different downstream tasks. However, commonly used anchor generation methods (*e.g.*, converting to grayscale images in image classification tasks) can introduce extra distortions and degrade the original perceptual quality, thus hindering the quality-aware representation learning in PCQA task. Furthermore, pre-trained models under the classical contrastive learning paradigm mainly focus on high-level semantic information [4, 9], while perceptual quality is determined by both high-level content and low-level distortion.

To solve these problems, we propose a novel **C**ontrastive pre-training framework tailored for **P**CQA (**CoPA**). CoPA has two main steps following the contrastive learning paradigm: 1) Anchor generation by patch mixing. CoPA applies random rotations to the point clouds to take advantage of the rich quality information of point clouds from multi-

ple perspectives, and then projects the point clouds impaired by different distortions into images. Subsequently, CoPA randomly mixes the local patches of the projected images to form a mixed image with multiple distortions, which is used as an anchor in the representation space, as illustrated in Fig. 1. The anchor generation process does not introduce additional distortions, thus preserving intrinsic quality information within the projected images. Moreover, due to the random rotation and patch mixing, CoPA can generate an unbounded number of anchors to form an extensive set of training pairs, allowing for a comprehensive pre-training of the network. 2) Content-wise and distortion-wise contrast. Inspired by the observation that perceptual quality is correlated with both content and distortion pattern [3, 23, 43]: for an anchor image, CoPA treats the projected images that form the anchor, as positive samples; images projected from different reference point clouds (content-wise) and images from the same reference point cloud but with unrelated distortions (distortion-wise) are regarded as negative samples. Following the contrastive learning paradigm, we finally obtain a pre-trained encoder, \mathcal{F} , that is expected to extract quality-aware features from projected images. The encoder \mathcal{F} is then incorporated into our model and fine-tuned with labeled data.

In fine-tuning stage, point clouds are first projected into multi-view images from different perspectives to mimic the subjective observation process. Then, a semantic-guided multi-view fusion module is proposed to integrate the multi-view quality-aware features generated by \mathcal{F} . Concretely, the multi-view images are stitched into a composed image, which is fed into a 2D backbone \mathcal{G} pre-trained on ImageNet [5] to extract the high-level semantic feature. Subsequently, the semantic feature guides the fusion of multi-view features through a cross-attention mechanism to obtain the final feature, followed by a quality score regression module to predict objective scores. We summarize the main contributions as follows:

- To tackle the challenge of label scarcity, we propose a contrastive pre-training framework (CoPA) tailored for PCQA. By generating anchor samples through local patch mixing and carefully designing positive/negative samples, CoPA enables the model to learn quality-aware representations to boost NR-PCQA performance.
- We propose a semantic-guided multi-view fusion module to effectively integrate the quality-aware features of the projected images from different perspectives in the fine-tuning stage.
- Extensive experiments show that the proposed method presents superior overall performance and generalization ability compared to the state-of-the-art NR-PCQA methods. Further investigation shows that CoPA can be integrated into other projection-based NR-PCQA methods with noticeable gains.

2. Related Works

2.1. No-Reference Point Cloud Quality Assessment

NR-PCQA aims to evaluate the perceptual quality of distorted point clouds without available references. Early NR-PCQA metrics typically train a regressor to obtain quality scores based on hand-crafted features. Zhang et al. [41] propose an NR-PCQA metric that projects point clouds into feature domains based on geometry and color and obtains the predicted MOS using a support vector machine (SVM). Zhou et al. [44] develop a general and efficient metric based on structure-guided resampling. Handcrafted features have explicit meanings and have shown decent evaluating ability. However, they are usually designed based on the still limited understanding of point cloud distortions, which are not comprehensive and thus are relatively limited in complex distortion environments.

With the boom of deep learning, many learning-based methods have been proposed for NR-PCQA. ResSCNN [13] employs a voxel-based sparse 3D-CNN for quality prediction. Fan et al. [7] transform point clouds into video sequences to infer the visual quality. PM-BVQA [27] uses an element-wise summation method to fuse the cross-modal and multi-scale features. Tliba et al. [29] propose a shallow model for interpolating compression-induced quality scores. PCQA-Graphpoint [30] employs a two-stream architecture to process geometry and color distortion, and combine features that leverage the attention mechanism. MVAT [18] aggregates the content and positional context via a multi-view aggregation transformer framework. ψ -Net [33] extracts geometric and color structural information from 3D patches by mapping the position vectors of neighboring points to weights. MM-PCQA [42] utilizes a novel multi-modal learning approach to assess the perceptual quality. However, these methods are all data-driven and thus suffer from the label paucity problem in PCQA.

To reduce dependence on labeled data, IT-PCQA [36] leverages the rich subjective scores of natural images and try to evaluate point cloud quality through unsupervised domain adaptation. Nevertheless, the domain gap hinders optimal training and leads to unsatisfactory performance. PQA-Net [11] and GPA-Net [24] both utilize distortion type prediction as an auxiliary task to integrate distortion-related knowledge into models. However, these methods ignore the correlation between distortions and present inferior performance for unseen distortion types.

2.2. Contrastive Learning

Contrastive learning aims to learn meaningful representations by maximizing the similarity between similar data samples and minimizing it between dissimilar ones [19, 38]. MoCo [9] uses a momentum update mechanism to maintain a long queue of negative samples for contrastive learning.

SimCLR [4] utilizes negative samples with a much larger batch size and constructs a rich family of data augmentations. However, most pre-trained models using these contrastive learning paradigms mainly focus on high-level semantic information, ignoring critical low-level distortions.

In image/video quality assessment scenarios, CONTRIQUE [15] utilizes a simple framework with quality-preserving augmentations to learn generalizable representations that perform well on synthetic and realistic distortions. Re-IQA [23] develops a holistic approach to assess the quality of images in the wild by leveraging the complementary content and image quality-aware features. QPT [43] supposes that the quality of patches from a distorted image should be similar, but vary from patches from different images and the same image with different degradations. CSPT [3] learns useful feature representation by using distorted video samples not only to formulate content-aware distorted instance contrasting but also to constitute an extra self-supervision signal for the distortion prediction task. However, when focusing on PCQA task, the significant volume of data brings challenges to the contrastive learning paradigm. To tackle this problem, we propose a projection-based contrastive pre-training framework considering the fact that point clouds are perceived from 2D perspective via projecting them on the screen in the subjective experiments of popular PCQA datasets [12, 13, 34]. Furthermore, after model pre-training, we fine-tune the model using a proposed multi-view fusion module to integrate the features from different perspectives.

3. Method

3.1. Overview

The goal of the proposed method is to pre-train an encoder that can effectively extract quality-aware features and then fine-tune it with labeled data to accurately predict the quality scores. Therefore, the framework of our model is divided into two stages: 1) In the first stage, we pre-train a quality-aware encoder \mathcal{F} using CoPA: each point cloud is first randomly rotated and projected into images. Then, we generate the anchor by local patch mixing and feed it into encoder \mathcal{F} , facilitating \mathcal{F} to learn quality-aware representations through both content-wise and distortion-wise contrast. 2) In the second stage, \mathcal{F} is fine-tuned with labeled data to predict quality scores: the labeled point clouds are first projected into multi-view images and then encoded by \mathcal{F} . Afterward, the multi-view features are fused to regress objective scores through a semantic-guided cross-attention operation, where the global semantic feature used to guide fusion is extracted by another 2D backbone \mathcal{G} pre-trained on ImageNet [5] for image classification.

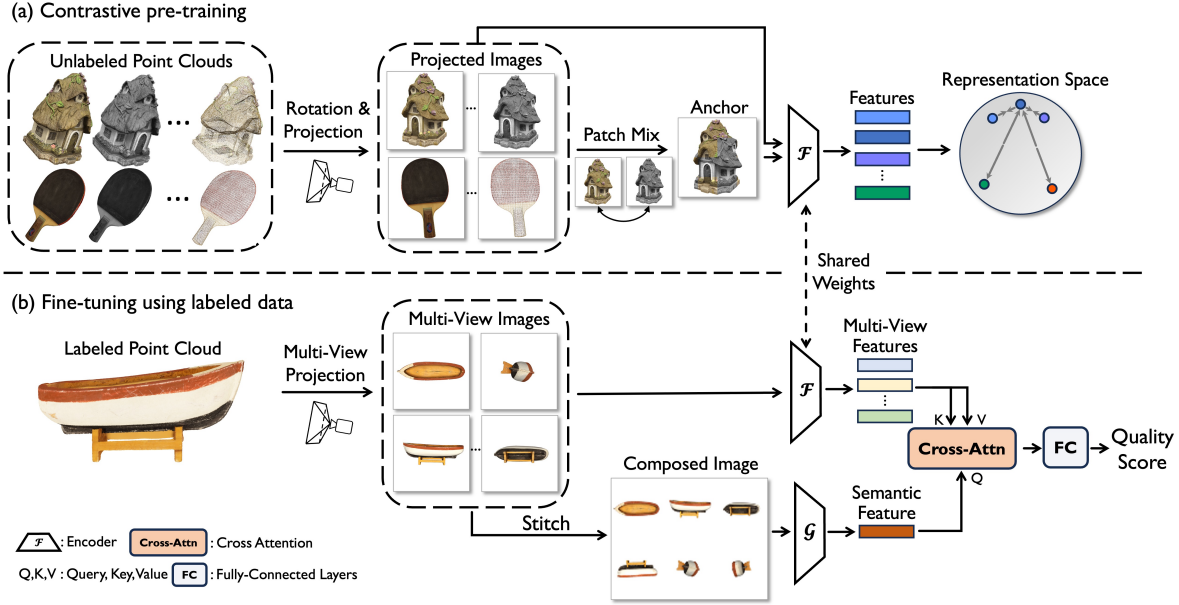


Figure 2. **Framework of the proposed method.** The framework mainly consists of two stages: (a) Contrastive pre-training. Unlabeled point clouds are projected into single-view images, and anchors are generated by local patch mixing. A single-view image encoder \mathcal{F} is pre-trained by pulling the defined positive samples to the anchors in the representation space and pushing the negative samples apart. (b) Fine-tuning with labeled data. The labeled point cloud is projected into multi-view images encoded by the pre-trained encoder \mathcal{F} . Then, the multi-view images are stitched to extract semantic features through a 2D backbone \mathcal{G} , which guides the fusion of multi-view images using the cross-attention mechanism. Finally, the quality score is regressed by the fully-connected layers.

3.2. Contrastive Pre-Training

Given N reference point clouds $p^{(1)}, \dots, p^{(N)}$ with varied contents, each point cloud $p^{(n)}$ is degraded with D types of distortion to form a distorted group $\{p_d^{(n)}\}_{d=1}^D$ that shares consistent content but different distortions.

Point Cloud Projection. Instead of performing contrastive learning directly in 3D space, we project point clouds into 2D images for two main advantages: 1) Pixel-to-pixel correspondence is established between each two projected images sampled from $\{p_d^{(n)}\}_{d=1}^D$, which facilitates the subsequent anchor generation. 2) Projected images have fixed and smaller data sizes compared to point clouds, which not only benefits parallel processing, but also increases batch size — a critical factor for contrastive learning [4].

To perform the projection, given a distorted point cloud $p_d^{(n)}$, we first apply random rotations \mathcal{R} to capture the visual information from different perspectives. For each rotation, the rotation degree is shared among all samples in $\{p_d^{(n)}\}_{d=1}^D$ to ensure the content consistency. Then we apply a geometric normalization operation \mathcal{N} to translate the rotated point cloud $\mathcal{R}(p_d^{(n)})$ to the origin and rescale it to the unit ball to achieve a consistent spatial scale. Finally, we render $\mathcal{N}(\mathcal{R}(p_d^{(n)}))$ into an image $x_d^{(n)} \in \mathbb{R}^{H \times W \times C}$ under a fixed viewing distance. Different rotation angles are applied for each point cloud to comprehensively capture

the visual information from multiple viewpoints.

Anchor Generation. Different from commonly used data augmentation methods [3, 43] (e.g., cropping, noise addition) that can bring unexpected distortions to affect the original quality information, we generate anchors by mixing local image patches to better preserve local distortion patterns, as illustrated in Fig. 1.

For two distorted projected images $x_{d_1}^{(n)}$ and $x_{d_2}^{(n)}$ that share the same reference content but different distortions, we first partition $x_{d_1}^{(n)}$ and $x_{d_2}^{(n)}$ into non-overlapping 16×16 patches following [6], and mix the patches using a random binary mask $\mathbf{M} \in \{0, 1\}^{H \times W}$. We define the mixing operation as follows:

$$\tilde{x}_{d_{1,2}}^{(n)} = \mathbf{M} \odot x_{d_1}^{(n)} + (\mathbf{1} - \mathbf{M}) \odot x_{d_2}^{(n)}, \quad (1)$$

where \mathbf{M} consists of 16×16 binary blocks, and the elements share the same value (0 or 1) within each binary block. $\mathbf{1}$ is a binary mask filled with ones and \odot denotes element-wise multiplication. \mathbf{M} is randomly generated so that numerous anchors and training pairs can be created.

After local patch mixing, the anchor $x_{d_{1,2}}^{(n)}$ contains two different distortions d_1 and d_2 but shares the same reference content with $x_{d_1}^{(n)}$ and $x_{d_2}^{(n)}$. Therefore, we can assume that the anchor $x_{d_{1,2}}^{(n)}$ is “similar” to $x_{d_1}^{(n)}$ and $x_{d_2}^{(n)}$, which will be considered positive samples in the contrastive pre-training, as detailed in the following paragraph.

Content-wise and Distortion-wise Contrast. We enable the pre-trained model to be attentive to both high-level content and low-level distortion via carefully designed positive/negative samples. Given the anchor $\tilde{x}_{d_1,2}^{(n)}$, $x_{d_1}^{(n)}$ and $x_{d_2}^{(n)}$ are considered positive samples due to the consistent content and the correlated distortion, while other samples are treated as negative ones that can be further divided into two parts. First, the samples of the set $\{x_d^{(n)}|_{d \neq d_1, d_2}\}^D$ are denoted as distortion-wise negative samples due to the same reference content but different distortions. Second, distorted images $\{x_d^{(m)}|_{m \neq n, d=1}\}^{N,D}$ originating from other reference point clouds (*i.e.*, $p^{(m)}$) are denoted as content-wise negative samples. Notably, there may exist some corner samples that have different contents but present similar qualities, but they are negligible considering the small proportion as verified in [43].

In the pre-training process, given an input sample $x_d^{(n)}$ and the encoder \mathcal{F} , we denote the encoded feature as $f_d^{(n)} = \mathcal{F}(x_d^{(n)})/\|\mathcal{F}(x_d^{(n)})\|$, after a \mathcal{L}_∞ normalization. Then the distortion-wise and content-wise contrastive pre-training loss $\mathcal{L}_d, \mathcal{L}_c$ for the anchor feature $f_{d_1,2}^{(n)}$ and the corresponding positive/negative samples can be formulated as:

$$\mathcal{L}_d = -r \log \frac{\exp(f_{d_1,2}^{(n)} \cdot f_{d_1}^{(n)}/\tau)}{\sum_{d \neq d_1, d_2}^D \exp(f_{d_1,2}^{(n)} \cdot f_d^{(n)}/\tau)} - (1-r) \log \frac{\exp(f_{d_1,2}^{(n)} \cdot f_{d_2}^{(n)}/\tau)}{\sum_{d \neq d_1, d_2}^D \exp(f_{d_1,2}^{(n)} \cdot f_d^{(n)}/\tau)}, \quad (2)$$

$$\mathcal{L}_c = -r \log \frac{\exp(f_{d_1,2}^{(n)} \cdot f_{d_1}^{(n)}/\tau)}{\sum_{m \neq n}^N \sum_{d=1}^D \exp(f_{d_1,2}^{(n)} \cdot f_d^{(m)}/\tau)} - (1-r) \log \frac{\exp(f_{d_1,2}^{(n)} \cdot f_{d_2}^{(n)}/\tau)}{\sum_{m \neq n}^N \sum_{d=1}^D \exp(f_{d_1,2}^{(n)} \cdot f_d^{(m)}/\tau)} \quad (3)$$

where (\cdot) denotes dot-product, τ is a temperature hyper-parameter. And $r = \frac{H}{\sum_{h=1}^H \sum_{w=1}^W \mathbf{M}(h, w)}/HW$ represents the masking ratio. Intuitively, the more patches of the anchor come from $x_{d_1}^{(n)}$ (or $x_{d_2}^{(n)}$), the closer $x_{d_1}^{(n)}$ (or $x_{d_2}^{(n)}$) should be pulled toward the anchor in the representation space.

Overall, the contrastive pre-training loss function can be formulated as:

$$\mathcal{L}_{pre} = \lambda \mathcal{L}_d + (1-\lambda) \mathcal{L}_c \quad (4)$$

where λ is the weighting coefficient.

Considering the large scale of content-wise negative samples, it is non-trivial to extract quality-aware features

from samples with all contents at each iteration. Therefore, we adopt the momentum contrast strategy [9] to reduce computational consumption, where a dynamic queue is established to enlarge the available size for contrasting. Meanwhile, the encoder parameter $\theta_{\mathcal{F}}$ is updated smoothly using the weighted sum of its previous weights. Once the update of the encoder parameter has converged, the pre-trained single-view encoder \mathcal{F} can be used as a backbone in downstream tasks to extract quality-aware features.

3.3. Fine-Tuning with Labeled Data

After obtaining the pre-trained encoder \mathcal{F} , we next fine-tune it using labeled point clouds. To mimic the multi-view characteristic when observing 3D object, we first project each point cloud into images from different viewpoints and encode these projected images using \mathcal{F} . Then, a multi-view fusion module is proposed to take advantage of global semantic information to integrate the multi-view features, considering that different views do not contribute equally to quality decision.

Multi-View Projection. For a point cloud p with label q , we normalize it with $\mathcal{N}(\cdot)$ and render $\mathcal{N}(p)$ into multi-view images $\{x_i \in \mathbb{R}^{H \times W \times C}|_{i=1}^6\}$ from six perpendicular viewpoints (*i.e.*, along the positive and negative directions of x,y,z axes) with fixed viewing distance. The rendering configuration is identical to Sec. 3.2.

Semantic-Guided Multi-View Fusion. To fuse the multi-view features derived from six projected images, another 2D encoder \mathcal{G} pre-trained on ImageNet [5] (image classification task) is employed to extract global semantic feature. Following the practices in [36], we first stitch the multi-view images into a composed image $x_c \in \mathbb{R}^{2H \times 3W \times C}$, and next encode it with \mathcal{G} to obtain the semantic feature g . Then we utilize g to guide the fusion of $\{f_i|_{i=1}^6\}$ through a simple multi-headed cross-attention mechanism \mathcal{M} .

Given three common sets of inputs: query set Q , key set K , and value set V , we define the \mathcal{M} as:

$$\begin{aligned} \mathcal{M}(Q, K, V) &= (\Gamma_1 \oplus \Gamma_2 \cdots \oplus \Gamma_{N_\Gamma})W, \\ \Gamma_\mu &= \mathcal{A}(QW_\mu^Q, KW_\mu^K, VW_\mu^V)|_{\mu=1}^{N_\Gamma}, \\ \mathcal{A}(Q, K, V) &= \text{softmax} \left(\frac{QK^T}{\sqrt{d_f}} \right) V \end{aligned} \quad (5)$$

where Γ_μ is the μ -th head, \mathcal{A} denotes attention function. W, W^Q, W^K, W^V are learnable linear mappings and d_f is a scaling factor. Then the multi-view fusion can be formulated as:

$$F = \mathcal{M}(g, \{f_i|_{i=1}^6\}, \{f_i|_{i=1}^6\}) \quad (6)$$

where F is the fused feature. In Eq. (6), intuitively, the semantic feature acts as a query and computes similarity with each f_i to find the semantically active viewpoints and fuse their quality-aware features attentively.

3.4. Quality Regression and Loss Function

After multi-view fusion, the final feature F is fed into two-fold fully-connected layers to regress the predicted quality score \hat{q} . Inspired by [42], our loss function includes two parts: mean square error (MSE) and rank error. The MSE optimizes the model to improve the prediction accuracy, which is formulated as:

$$\mathcal{L}_{mse} = \frac{1}{B} \sum_b (\hat{q}_b - q_b)^2 \quad (7)$$

where \hat{q}_b is the predicted quality score of b -th sample in a mini-batch with the size of B , and q_b is the corresponding ground truth MOS.

To better recognize quality differences for the point clouds with close MOSs, we use a differential ranking loss [25] to model the ranking relationship between \hat{q} and q :

$$\begin{aligned} \mathcal{L}_{rank} &= \frac{1}{B^2} \sum_{i=1}^B \sum_{j=1}^B \mathcal{L}_{rank}^{ij}, \\ \mathcal{L}_{rank}^{ij} &= \max(0, |q_i - q_j| - e(q_i, q_j) \cdot (\hat{q}_i - \hat{q}_j)), \\ e(q_i, q_j) &= \begin{cases} 1, & q_i \geq q_j, \\ -1, & q_i < q_j, \end{cases} \end{aligned} \quad (8)$$

Then the overall loss function for fine-tuning can be calculated as the sum of MSE loss and ranking loss:

$$\mathcal{L}_{fine} = \alpha \mathcal{L}_{mse} + (1 - \alpha) \mathcal{L}_{rank} \quad (9)$$

where the hyper-parameter α is to balance the two losses.

4. Experiments

4.1. Datasets and Evaluation Metrics

Datasets. Our experiments are based on three commonly used PCQA datasets, including LS-PCQA [13], SJTU-PCQA [34], and WPC [12]. The pre-training is based on LS-PCQA, which is a large-scale PCQA dataset and contains 104 pristine reference point clouds and 24,024 distorted point clouds, and each reference point cloud is impaired with 33 types of distortions (*e.g.*, V-PCC, G-PCC) under 7 levels. For a fair comparison, the model is fine-tuned on all three datasets separately using labeled data, where SJTU-PCQA includes 9 reference point clouds and 378 distorted samples impaired with 7 types of distortions (*e.g.*, color noise, downsampling) under 6 levels, while WPC contains 20 reference point clouds and 740 distorted sampled disturbed by 5 types of distortions (*e.g.*, compression, gaussian noise).

Evaluation Metrics. Three widely adopted evaluation metrics are employed to quantify the level of agreement between predicted quality scores and MOSs: Spearman rank

order correlation coefficient (SROCC), Pearson linear correlation coefficient (PLCC), and root mean square error (RMSE). To ensure consistency between the value ranges of the predicted scores and subjective values, nonlinear Logistic-4 regression is used to align their ranges.

4.2. Implementation Details

Our experiments are performed using PyTorch [20] on NVIDIA 3090 GPUs. All point clouds are rendered into projected images with a spatial resolution of 512×512 by PyTorch3D [22]. Following the contrastive learning paradigm, the experiment is performed in two stages:

Pre-Training. We employ a Swin-T [14] as encoder \mathcal{F} to be pre-trained for 200 epochs. The momentum parameter updating follows the configuration in [9]. We use the SGD optimizer [26] with weight decay 0.0001, momentum of 0.95, and batch size of 128. The learning rate is 0.005 and decayed by 0.2 every 10 epochs. λ and τ are set to 0.3 and 0.2. Each point cloud is randomly rotated 6 times before projection, and the masking ratio r is bounded to $[0.25, 0.75]$ to ensure that local patches are sufficiently mixed and mitigate the influence of projected images' backgrounds.

Fine-Tuning. We use a ResNet50 [8] pre-trained on ImageNet [5] as encoder \mathcal{G} to extract semantic features. The encoded features from \mathcal{F} and \mathcal{G} are projected to a channel size of 1024 by one-layer MLPs. The multi-headed cross-attention employs 8 heads and d_f is empirically set to 1024. We use the Adam optimizer [10] with weight decay of 0.0001 and batch size of 16. The hidden dimension of fully-connected layers is set to 512. The learning rate is initialized with 0.003 and decayed by 0.9 every 5 epochs. For LS-PCQA, the model is trained for 20 epochs, while 150 epochs for SJTU-PCQA and WPC. The hyper-parameter α is set to 0.5.

Considering the limited dataset scale, in the fine-tuning stage, 5-fold cross validation is used for SJTU-PCQA and WPC to reduce content bias. Take SJTU-PCQA for example, in each fold, the dataset is split into train-test with ratio 7:2 according to the reference point clouds, where the performance on testing set with minimal training loss is recorded and then averaged across five folds to get the final result. A similar procedure is repeated for WPC where the train-test ratio is 4:1. As for the large-scale LS-PCQA, it is split into train-val-test with a ratio around 8:1:1 (no content overlap exists). The result on the testing set with the best validation performance is recorded. Note that the pre-training is only conducted on the training set of LS-PCQA.

4.3. Comparison with State-of-the-art Methods

15 state-of-the-art PCQA methods are selected for comparison, including 10 FR-PCQA and 5 NR-PCQA methods. The FR-PCQA methods include MSE-p2point [16], Hausdorff-p2point [16], MSE-p2plane [28], Hausdorff-

Table 1. Performance results on the LS-PCQA [13], SJTU-PCQA [34] and WPC [12] databases. “P” and “I” stand for the method is based on the point cloud and image modality, respectively. “↑”/“↓” indicates that larger/smaller is better. The best performance results are marked in **RED** and the second results are marked in **BLUE** for both FR-PCQA and NR-PCQA methods. “FT” indicates fine-tuning.

Ref	Modal	Methods	LS-PCQA [13]			SJTU-PCQA [34]			WPC [12]		
			SROCC↑	PLCC↑	RMSE↓	SROCC↑	PLCC↑	RMSE↓	SROCC↑	PLCC↑	RMSE↓
FR	P	MSE-p2po	0.325	0.528	0.158	0.783	0.845	0.122	0.564	0.557	0.188
	P	HD-p2po	0.291	0.488	0.163	0.681	0.748	0.156	0.106	0.166	0.222
	P	MSE-p2pl	0.311	0.498	0.160	0.703	0.779	0.149	0.445	0.491	0.199
	P	HD-p2pl	0.291	0.478	0.163	0.617	0.661	0.177	0.344	0.380	0.211
	P	PSNR-yuv	0.548	0.547	0.155	0.704	0.715	0.165	0.563	0.579	0.186
	P	PointSSIM	0.180	0.178	0.183	0.735	0.747	0.157	0.453	0.481	0.200
	P	PCQM	0.439	0.510	0.152	0.864	0.883	0.112	0.750	0.754	0.150
	P	GraphSIM	0.320	0.281	0.178	0.856	0.874	0.114	0.679	0.693	0.165
	P	MS-GraphSIM	0.389	0.348	0.174	0.888	0.914	0.096	0.704	0.718	0.159
NR	P	MPED	0.659	0.671	0.138	0.898	0.915	0.096	0.656	0.670	0.169
	I	PQA-Net	0.588	0.592	0.202	0.659	0.687	0.172	0.547	0.579	0.189
	I	IT-PCQA	0.326	0.347	0.224	0.539	0.629	0.218	0.422	0.468	0.221
	P	GPA-Net	0.592	0.619	0.186	0.878	0.886	0.122	0.758	0.769	0.162
	P	ResSCNN	0.594	0.624	0.172	0.834	0.863	0.153	0.735	0.752	0.177
	P+I	MM-PCQA	0.581	0.597	0.189	0.876	0.898	0.109	0.761	0.774	0.149
	P	CoPA+FT	0.613	0.636	0.161	0.897	0.913	0.092	0.779	0.785	0.144

p2plane [28], PSNR-yuv [31], PointSSIM [2], PCQM [17], GraphSIM [35], MS-GraphSIM [40], and MPED [37]. The NR-PCQA methods include PQA-Net [11], IT-PCQA [36], GPA-Net [24], ResSCNN [13], and MM-PCQA [42]. For a comprehensive comparison, we conduct the experiment in four aspects: 1) We compare the prediction accuracy performance, following the cross-validation configuration in Sec. 4.2. 2) We conduct a visualization analysis for the quality-aware representation learned by CoPA and demonstrate the superiority of our model in terms of distinguishing different distortion patterns. 3) We report the results of cross-dataset evaluation for the NR-PCQA methods to verify the generalization performance of our model. 4) We compare the performance of the NR-PCQA methods with fewer MOS labels for training (fine-tuning for our model).

Comparison of Prediction Accuracy. The prediction accuracy of all the selected methods on LS-PCQA, SJTU-PCQA and WPC is presented Tab. 1. From the table, we have the following observations: 1) Our model outperforms all the NR-PCQA methods on all three datasets. For example, our model improves the best performance by about 3.20% of SROCC (0.613 vs. 0.594), 1.9% of PLCC (0.636 vs. 0.624), and 6.4% of RMSE (0.161 vs. 0.172) on LS-PCQA. 2) Our method presents competitive performance compared to FR-PCQA methods despite the inaccessibility to the reference information. 3) Our model demonstrates robust performance across the three datasets, regardless of the variations in dataset scale, content, and distortion types.

Visualization Analysis. We visualize the learned representations of PQA-Net, GPA-Net, our model without pre-training and the complete model (*i.e.*, CoPA + fine-tuning) on a unified testing set of a certain fold in SJTU-PCQA.

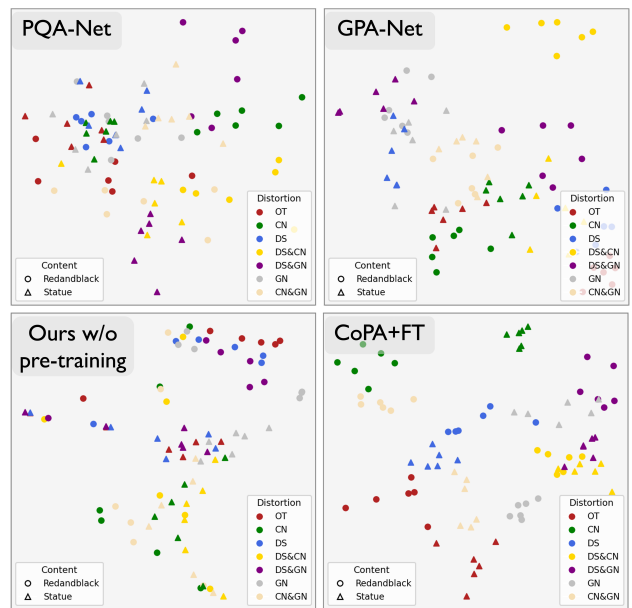


Figure 3. T-SNE embedding of the representation spaces of PQA-Net, GPA-Net, our model without pre-training, and our complete model on testing set of SJTU-PCQA. The scattered points are color and shape encoded according to distortion type and content.

Specifically, we use t-SNE algorithm [32] to embed the representations into 2D feature space, and the visualization results are illustrated in Fig. 3. We can see that: 1) Compared to PQA-Net and GPA-Net, our model achieves better clustering results for different distortions, from which we can conclude that the proposed method is able to learn more distortion-discriminative representations than distort-

Table 2. Cross-dataset evaluation for NR-PCQA methods. Training and testing are both conducted on complete datasets. Results of PLCC are reported.

Train	Test	PQA-Net	GPA-Net	ResSCNN	MM-PCQA	CoPA+FT
LS	SJTU	0.342	0.556	0.546	0.581	0.644
LS	WPC	0.266	0.433	0.466	0.454	0.516
WPC	SJTU	0.235	0.553	0.572	0.612	0.643
SJTU	WPC	0.220	0.418	0.269	0.409	0.533

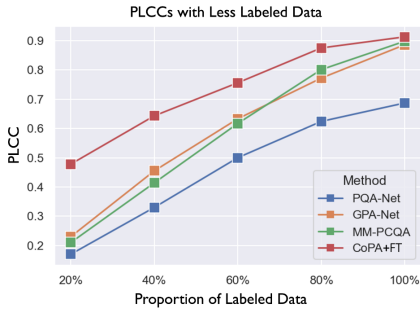


Figure 4. PLCCs of the NR-PCQA methods with less labeled data on SJTU-PCQA. Our pre-trained model outperforms the compared methods by a large margin when the training data is limited.

tion type prediction task. 2) Compared to the proposed network without pre-training, the complete model apparently achieves better discrimination performance, which demonstrates the effectiveness of contrastive pre-training.

Cross-Dataset Evaluation. The cross-dataset evaluation is conducted to test the generalizability of the NR-PCQA methods when encountering various data distribution. In Tab. 2, we mainly train the compared models on the complete LS-PCQA and test the trained model on the complete SJTU-PCQA and WPC, and the result with minimal training loss is recorded. The procedure is repeated for the mutual cross-dataset evaluation between SJTU-PCQA and WPC. Notably, considering that LS-PCQA shares some reference point clouds with SJTU-PCQA, we remove these groups of point clouds from LS-PCQA to avoid information leakage. From Tab. 2, we can see that the performance of cross-dataset evaluation is relatively low due to the huge variation with respect to both distortion types and contents. However, our method still outperforms the NR-PCQA methods by a large margin, demonstrating the superior generalizability of CoPA.

Performance with Less Labeled Data. To test the generalization performance of the NR-PCQA models with less labeled data, we compare their performances with different proportions of labeled data. The result is illustrated in Fig. 4. We can see that our model outperforms the compared NR-PCQA methods, and the margin is larger with less training data. Therefore, we can conclude that the pre-training framework can significantly reduce the reliance on annotated data, thus demonstrating its potential to address the issue of label scarcity for the NR-PCQA task.

Table 3. Ablation on components of our model on SJTU-PCQA. ‘✓’ or ‘✗’ means the setting is preserved or discarded. ‘Dis.’ and ‘Con.’ indicates distortion-wise and content-wise contrast.

Index	Pre-training	Multi-view fusion	Fine-tuning loss	SROCC	PLCC
①	✓	✓	✓	0.897	0.913
②	✗	✓	✓	0.806	0.893
③	Dis. only	✓	✓	0.856	0.887
④	Con. only	✓	✓	0.889	0.901
⑤	✓	Max Pooling	✓	0.841	0.872
⑥	✓	Average Pooling	✓	0.810	0.847
⑦	✓	✓	MSE only	0.889	0.910

Table 4. PLCCs of using CoPA to pre-train NR-PCQA methods.

Dataset	PQA-Net	PQA-Net+CoPA	MM-PCQA	MM-PCQA+CoPA
SJTU	0.687	0.749	0.898	0.910
WPC	0.579	0.911	0.774	0.778

4.4. Ablation Studies

To study the effectiveness of our proposed method, we further investigate the individual contribution of different components and then analyze the effects of the pre-training framework on other NR-PCQA models.

Ablation on Different Components. In Tab. 3, we report the results on SJTU-PCQA with the conditions of dropping some components. From Tab. 3, we have the following observations: 1) Seeing ① and ② - ④, the pre-training framework brings the most significant improvements when considering both distortion-wise and content-wise contrast. 2) Seeing ① and ⑤ - ⑥, the semantic-guided multi-view fusion performs better than the symmetric fusion strategies (*i.e.* max and average pooling). 3) Seeing ① and ⑦, the performance is close, which demonstrates the robustness of our model using different fine-tuning loss functions.

Contrastive pre-training on other NR-PCQA Methods. In Tab. 4, we employ the contrastive pre-training framework in Sec. 3.2 to optimize other projection-based NR-PCQA models (some details like feature dimensions are accordingly modified). We can see that the pre-training framework can fit different backbones and bring noticeable gain for existing projection-based NR-PCQA methods.

5. Conclusion

In this paper, we propose a novel no-reference point cloud quality method based on contrastive learning. Utilizing local patch mixing to generate anchors without impairing the distortion patterns, the proposed CoPA learns quality-aware representations through both content-wise and distortion-wise contrasts. Moreover, in the fine-tuning stage, we use the semantic-guided multi-view fusion to attentively integrate the quality-aware features from different perspectives. Experiments show that our model presents competitive and generalizable performance compared to the state-of-the-art NR-PCQA methods.

References

- [1] Mohamed Afham, Isuru Dissanayake, Dinithi Dissanayake, Amaya Dharmasiri, Kanchana Thilakarathna, and Ranga Rodrigo. Crosspoint: Self-supervised cross-modal contrastive learning for 3d point cloud understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9902–9912, 2022. 1
- [2] Evangelos Alexiou and Touradj Ebrahimi. Towards a point cloud structural similarity metric. In *ICMEW*, pages 1–6, 2020. 7
- [3] Pengfei Chen, Leida Li, Jinjian Wu, Weisheng Dong, and Guangming Shi. Contrastive self-supervised pre-training for video quality assessment. *IEEE Transactions on Image Processing*, 31:458–471, 2021. 2, 3, 4
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 2, 3, 4
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255. Ieee, 2009. 2, 3, 5, 6
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 4
- [7] Yu Fan, Zicheng Zhang, Wei Sun, Xiongkuo Min, Ning Liu, Quan Zhou, Jun He, Qiyuan Wang, and Guangtao Zhai. A no-reference quality assessment metric for point cloud based on captured video sequences. In *2022 IEEE 24th International Workshop on Multimedia Signal Processing (MMSP)*, pages 1–5. IEEE, 2022. 3
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [9] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 2, 3, 5, 6
- [10] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [11] Qi Liu, Hui Yuan, Honglei Su, Hao Liu, Yu Wang, Huan Yang, and Junhui Hou. Pqa-net: Deep no reference point cloud quality assessment via multi-view projection. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(12):4645–4660, 2021. 2, 3, 7
- [12] Qi Liu, Honglei Su, Zhengfang Duanmu, Wentao Liu, and Zhou Wang. Perceptual quality assessment of colored 3d point clouds. *IEEE Transactions on Visualization and Computer Graphics*, 2022. 2, 3, 6, 7
- [13] Yipeng Liu, Qi Yang, Yiling Xu, and Le Yang. Point cloud quality assessment: Dataset construction and learning-based no-reference metric. *ACM Transactions on Multimedia Computing, Communications and Applications*, 19(2s): 1–26, 2023. 2, 3, 6, 7
- [14] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 6
- [15] Pavan C Madhusudana, Neil Birkbeck, Yilin Wang, Balu Adsumilli, and Alan C Bovik. Image quality assessment using contrastive learning. *IEEE Transactions on Image Processing*, 31:4149–4161, 2022. 3
- [16] R Mekuria, Z Li, C Tulvan, and P Chou. Evaluation criteria for point cloud compression. *ISO/IEC MPEG*, (16332), 2016. 6
- [17] Gabriel Meynet, Yana Nehmé, Julie Digne, and Guillaume Lavoué. Pqcm: A full-reference quality metric for colored 3d point clouds. In *QoMEX*, pages 1–6, 2020. 7
- [18] Baoyang Mu, Feng Shao, Xiongli Chai, Qiang Liu, Hangwei Chen, and Qiuping Jiang. Multi-view aggregation transformer for no-reference point cloud quality assessment. *Displays*, 78:102450, 2023. 3
- [19] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 3
- [20] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 6
- [21] 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. 1
- [22] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:2007.08501*, 2020. 6
- [23] Avinab Saha, Sandeep Mishra, and Alan C Bovik. Re-iqu: Unsupervised learning for image quality assessment in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5846–5855, 2023. 2, 3
- [24] Ziyu Shan, Qi Yang, Rui Ye, Yujie Zhang, Yiling Xu, Xiaozhong Xu, and Shan Liu. Gpa-net: No-reference point cloud quality assessment with multi-task graph convolutional network. *IEEE Transactions on Visualization and Computer Graphics*, 2023. 2, 3, 7
- [25] Wei Sun, Xiongkuo Min, Wei Lu, and Guangtao Zhai. A deep learning based no-reference quality assessment model for ugc videos. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 856–865, 2022. 6
- [26] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum

- in deep learning. In *International conference on machine learning*, pages 1139–1147. PMLR, 2013. 6
- [27] Wen-xu Tao, Gang-yi Jiang, Zhi-di Jiang, and Mei Yu. Point cloud projection and multi-scale feature fusion network based blind quality assessment for colored point clouds. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 5266–5272, 2021. 3
- [28] Dong Tian, Hideaki Ochimizu, Chen Feng, Robert Cohen, and Anthony Vetro. Geometric distortion metrics for point cloud compression. In *IEEE ICIP*, pages 3460–3464, 2017. 6, 7
- [29] Marouane Tliba, Aladine Chetouani, Giuseppe Valenzise, and Frédéric Dufaux. Representation learning optimization for 3d point cloud quality assessment without reference. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 3702–3706. IEEE, 2022. 3
- [30] Marouane Tliba, Aladine Chetouani, Giuseppe Valenzise, and Frédéric Dufaux. Pcqa-graphpoint: Efficient deep-based graph metric for point cloud quality assessment. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023. 3
- [31] Eric M Torlig, Evangelos Alexiou, Tiago A Fonseca, Ricardo L de Queiroz, and Touradj Ebrahimi. A novel methodology for quality assessment of voxelized point clouds. In *Applications of Digital Image Processing XLI*, pages 174–190, 2018. 7
- [32] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9 (11), 2008. 7
- [33] Jian Xiong, Sifan Wu, Wang Luo, Jinli Suo, and Hao Gao. ψ -net: Point structural information network for no-reference point cloud quality assessment. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023. 3
- [34] Qi Yang, Hao Chen, Zhan Ma, Yiling Xu, Rongjun Tang, and Jun Sun. Predicting the perceptual quality of point cloud: A 3d-to-2d projection-based exploration. *IEEE Transactions on Multimedia*, 23:3877–3891, 2020. 2, 3, 6, 7
- [35] Qi Yang, Zhan Ma, Yiling Xu, Zhu Li, and Jun Sun. Inferring point cloud quality via graph similarity. *IEEE transactions on pattern analysis and machine intelligence*, 44(6):3015–3029, 2020. 7
- [36] Qi Yang, Yipeng Liu, Siheng Chen, Yiling Xu, and Jun Sun. No-reference point cloud quality assessment via domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21179–21188, 2022. 2, 3, 5, 7
- [37] Qi Yang, Yujie Zhang, Siheng Chen, Yiling Xu, Jun Sun, and Zhan Ma. Mped: Quantifying point cloud distortion based on multiscale potential energy discrepancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5):6037–6054, 2022. 7
- [38] Rui Ye, Zhenyang Ni, Chenxin Xu, Jianyu Wang, Siheng Chen, and Yonina C Eldar. Fedfm: Anchor-based feature matching for data heterogeneity in federated learning. *IEEE Transactions on Signal Processing*, 2023. 3
- [39] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8552–8562, 2022. 1
- [40] Yujie Zhang, Qi Yang, and Yiling Xu. Ms-graphsim: Inferring point cloud quality via multiscale graph similarity. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 1230–1238, 2021. 7
- [41] Zicheng Zhang, Wei Sun, Xiongkuo Min, Tao Wang, Wei Lu, and Guangtao Zhai. No-reference quality assessment for 3d colored point cloud and mesh models. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(11):7618–7631, 2022. 3
- [42] Zicheng Zhang, Wei Sun, Xiongkuo Min, Quan Zhou, Jun He, Qiyuan Wang, and Guangtao Zhai. Mm-pcqa: Multi-modal learning for no-reference point cloud quality assessment. *arXiv preprint arXiv:2209.00244*, 2022. 3, 6, 7
- [43] Kai Zhao, Kun Yuan, Ming Sun, Mading Li, and Xing Wen. Quality-aware pre-trained models for blind image quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22302–22313, 2023. 2, 3, 4, 5
- [44] Wei Zhou, Qi Yang, Qiuping Jiang, Guangtao Zhai, and Weisi Lin. Blind quality assessment of 3d dense point clouds with structure guided resampling. *arXiv preprint arXiv:2208.14603*, 2022. 2, 3