

CLIP-PCQA: Exploring Subjective-Aligned Vision-Language Modeling for Point Cloud Quality Assessment

Yating Liu*, Yujie Zhang*, Ziyu Shan, Yiling Xu†

Cooperative Medianet Innovation Center, Shanghai Jiao Tong University, Shanghai, China
{Olivialyt, yujie19981026, shanziyu, yl.xu}@sjtu.edu.cn

Abstract

In recent years, No-Reference Point Cloud Quality Assessment (NR-PCQA) research has achieved significant progress. However, existing methods mostly seek a direct mapping function from visual data to the Mean Opinion Score (MOS), which is contradictory to the mechanism of practical subjective evaluation. To address this, we propose a novel language-driven PCQA method named CLIP-PCQA. Considering that human beings prefer to describe visual quality using discrete quality descriptions (*e.g.*, “excellent” and “poor”) rather than specific scores, we adopt a retrieval-based mapping strategy to simulate the process of subjective assessment. More specifically, based on the philosophy of CLIP, we calculate the cosine similarity between the visual features and multiple textual features corresponding to different quality descriptions, in which process an effective contrastive loss and learnable prompts are introduced to enhance the feature extraction. Meanwhile, given the personal limitations and bias in subjective experiments, we further convert the feature similarities into probabilities and consider the Opinion Score Distribution (OSD) rather than a single MOS as the final target. Experimental results show that our CLIP-PCQA outperforms other State-Of-The-Art (SOTA) approaches.

Code — <https://github.com/Olivialyt/CLIP-PCQA>

1 Introduction

Point clouds are essentially collections of points scattered in 3D space, where each point has the corresponding geometry coordinates and other attributes. With its widespread applications in various fields (*e.g.*, autonomous driving (Cui et al. 2021) and virtual/augmented reality (Park, Lepetit, and Woo 2008)), point clouds have become one of the most typical representations of 3D objects or scenes. However, many visual artifacts could be introduced into point clouds during processing stages, such as capturing, compression, transmission, and rendering. In order to ensure the Quality Of Experience (QoE), it is necessary to develop effective metrics that accurately measure the perceptual distortion, especially in the common No-Reference (NR) situation where “perfect” reference point clouds are not available.

*These authors contributed equally.

†Corresponding author.

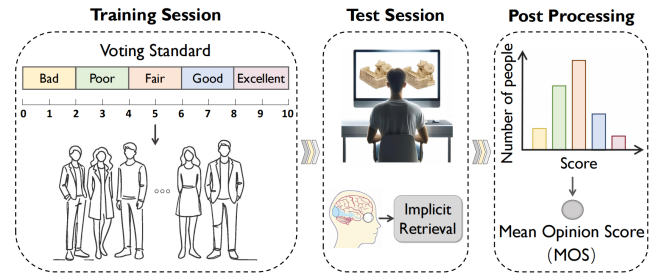


Figure 1: The process of subjective experiment. Participants first undergo a training process to understand the voting standard, then they are able to convert the quality descriptions into quantitative values. Note that multiple subjects may give diverse scores, which forms the score distribution.

In the past years, a variety of deep learning-based NR-PCQA methods have been proposed. These methods can be operated directly on 3D point clouds (Liu et al. 2023b; Shan et al. 2023; Tliba et al. 2023) or 2D projected images (Torlig et al. 2018; Liu et al. 2021; Mu et al. 2023). More recently, another family of methods (Zhang et al. 2022b; Xie et al. 2023) have designed multi-modal architecture to jointly utilize differentiated 2D and 3D information. These works, regardless of which modality they are based on, try to seek a direct mapping from visual data to quantitative scores. The common strategy is to use well-designed deep neural networks to extract effective quality features from input point clouds, and then regress the features into required scores.

Although promising results have been reported, the prevalent regression-based mapping strategy is somewhat contradictory to the mechanism of subjective evaluation. In fact, when stimulated by distorted samples, people’s immediate perception involves intrinsic descriptions such as “the quality of this sample is good” or “the distortion is annoying”. To acquire quantitative scores that are more convenient in practical applications, participants in subjective experiments need to undergo a training process in which they are informed about the voting standard (*i.e.*, the correspondence between quality/impairment descriptions and quantitative values). As shown in Figure 1, the five-grade categorical scales (Recommendation 2019) have been widely used. After the training session, subjects can use quantitative values

to represent their judgment by *implicit retrieval*, that is, finding the description that best matches the test sample. We argue that quality descriptions, such as “excellent” or “poor”, can better simulate human perception. Consequently, we can adopt a two-fold retrieval-based mapping strategy for quality prediction: i) first, find out which language description matches the test sample best; ii) second, represent the language descriptions using quantitative values.

Notably, direct retrieval for a single description may not generate the desired results. In subjective experiments, different subjects may have different understandings and preferences for the same sample, leading to diverse scores. To reduce personal bias, the final MOS is usually obtained by averaging the scores of multiple participants. The distribution of raw subjective opinion scores can provide richer information than MOS, such as individual diversity and uncertainty. Therefore, we further argue that predicting the Opinion Score Distribution (OSD) may be more consistent with subjective evaluation than predicting a scalar score.

Based on the above insights, we develop a novel NR-PCQA method named CLIP-PCQA, which uses a retrieval-based mapping strategy and considers score distributions as final targets. We first extract visual and textual features from point clouds and quality descriptions, where we resort to CLIP (Radford et al. 2021) that trains a visual encoder and a textual encoder on large-scale (400 million) image-text pairs. Specifically, to extract visual features, 3D point clouds are projected onto multiple image planes to obtain 2D color and depth maps, which are fed into two separate visual encoders and then simply fused. A contrastive loss is employed to help the model distinguish different samples and viewpoints. For the textual feature extraction, we place different quality-related adjectives (*e.g.*, “excellent” and “poor”) into a learnable template as prompts and then extract the corresponding features using the frozen textual encoder.

After the feature extraction, to mimic the retrieval process, we calculate the cosine similarity between the visual feature and multiple textual features corresponding to different quality descriptions. These similarities are then fused into softmax probabilities. We treat the output softmax probabilities as the predicted quality score distribution and use the original opinion score distribution as supervisory information. Finally, by matching each quality description with a certain quantitative value and using the softmax probabilities as weights, we could obtain the quality score of point clouds. Our main contributions are summarized as follows:

- We propose a novel language-driven PCQA method named CLIP-PCQA, which utilizes the capacity of natural language to understand point cloud quality.
- We leverage quality descriptions to characterize people’s intuitive perception and align with subjective experiments by predicting the OSD. The entire method simultaneously and effectively simulates the mechanisms of HVS and the process of subjective experiments.
- We conduct comprehensive experiments on multiple benchmarks. Experimental results indicate that CLIP-PCQA achieves superior performance and further analyses reveal the model’s robustness under different settings.

2 Related Works

No-Reference Point Cloud Quality Assessment

NR-PCQA aims to evaluate the perceptual quality of distorted point clouds without available references. Existing NR methods can be roughly divided into point-based and projection-based methods. Point-based methods use 3D operations on raw 3D points without any transformation. ResS-CNN (Liu et al. 2023b) pioneered the first NR-PCQA method based on an end-to-end voxel-based sparse CNN. GPA-Net (Shan et al. 2023) uses a novel graph convolution kernel and a multi-task framework for quality regression, distortion type, and degree predictions. 3DTA (Zhu et al. 2024) employs a two-stage sampling method for efficient point cloud representation. Projection-based methods assess the quality based on the 2D projected images from point clouds. PQA-Net (Liu et al. 2021) is the first NR-PCQA method based on multi-view feature extraction and fusion. D³-PCQA (Liu et al. 2023c) establishes an intermediate description domain to bridge perceptual and quality domains. Fan *et al.* (Fan et al. 2024) pioneered a probabilistic architecture through a conditional variational autoencoder.

Some recent works propose to fuse multi-modal information for a comprehensive feature representation of point clouds. MM-PCQA (Zhang et al. 2022b) first fuses the texture features from 2D images and geometric information from 3D point clouds. MMF-PCQM (Wang, Gao, and Li 2023) utilizes a dual-branch fusion network and collaborative adversarial learning strategy for multi-modal learning.

Vision-Language Learning in Quality Assessment

Vision-Language (V-L) pre-training has achieved extraordinary success for a wide range of multi-modal tasks. By pre-training large-scale image-text pairs, CLIP (Radford et al. 2021) has raised a revolution in various vision fields (He et al. 2023; Gu et al. 2021; Luo et al. 2022; Wang, Xing, and Liu 2021; Zhang et al. 2022a), which proves the capacity of natural language to understand visual concepts. In the field of quality assessment, some IQA and VQA methods (Wang, Chan, and Loy 2023; Liu et al. 2023a; Wu et al. 2023a,b; Srinath et al. 2024; Agnolucci, Galteri, and Bertini 2024) have exploited CLIP for better quality prediction performance. Roughly, these works can be divided into two categories.

The first category is similar to the framework of CLIP and employs a pair of antonym prompts (positive description and negative description) as textual inputs. CLIP-IQA (Wang, Chan, and Loy 2023) and MaxVQA (Wu et al. 2023b) convert quality regression to a binary classification task and regard the softmax results as the quality scores. GRepQ (Srinath et al. 2024) and QualiCLIP (Agnolucci, Galteri, and Bertini 2024) artificially construct image groups with different quality, and then rank similarity or separate higher and lower quality groups based on antonym text prompts.

The second category incorporates the vision-language information from CLIP as one modality in multi-modal tasks to comprehensively assess quality. Ada-DQA (Liu et al. 2023a) constructs a diverse pool that encompasses a broad spectrum of quality-related factors and performs knowledge distillation, the features generated by CLIP and other pre-

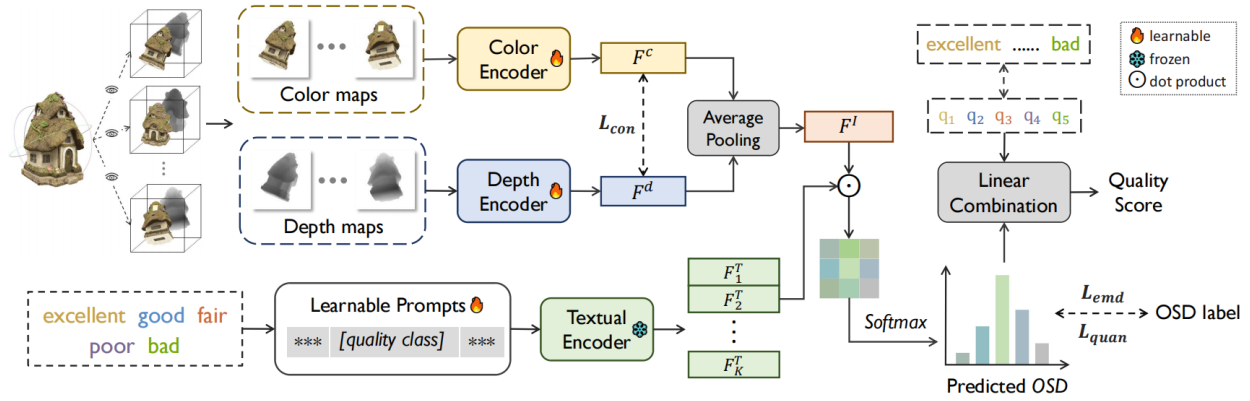


Figure 2: The proposed CLIP-PCQA framework, which includes two main parts: multi-modal feature extraction and vision-language alignment. We use the encoders in CLIP to extract features and then perform vision-language alignment using OSD.

trained models are aggregated. BVQI (Wu et al. 2023a) includes the spatial naturalness index, the temporal naturalness index, and the CLIP-based semantic affinity quality index, the three indices are aggregated for final prediction.

The spirit of our work is more similar to the first category. However, previous studies simply employ a pair of antonym prompts (e.g., “a good photo” and “a bad photo”), which fails to explore finer-grained descriptions in subjective experiments. To better simulate the process of subjective evaluation, our work merges the quality descriptions in five-grade voting scales into learnable prompts and considers score distributions as final targets. Furthermore, to better exploit the characteristics of point cloud data, we employ an effective contrastive loss to constrain the visual feature extraction.

3 Proposed Method

In this study, we aim to develop a language-driven PCQA method that simulates the mechanism of subjective evaluation. The architecture of our network is shown in Figure 2. The proposed method consists of two main parts: multi-modal feature extraction and vision-language alignment. We will first formulate our PCQA model and then present details of different modules in the following sections.

Problem Formulation

Existing learning-based NR-PCQA approaches typically take point clouds or projected images as input x and seek a direct mapping function to MOS denoted by Q . This procedure can be formulated as:

$$\hat{Q} = r(\phi(x)), \quad (1)$$

where \hat{Q} indicates the predicted score, $\phi(\cdot)$ denotes the feature extraction operation while $r(\cdot)$ represents a regression head that maps the feature to a scalar value.

As aforementioned, this regression-based mapping strategy does not align well with the subjective assessment. Human beings tend to describe quality using discrete adjectives, and the training session in subjective experiments helps them establish the relation between quality descriptions and quantitative scores. During the test phase, different participants can retrieve different quality descriptions for

specific samples based on prior knowledge, and then convert them into quantitative values. To reduce personal limitations and bias, MOS is obtained by averaging the scores of multiple subjects. Therefore, given K quality descriptions $\{y_k\}_{k=1}^K$, we expect to predict the OSD $P = \{p_k\}_{k=1}^K$ for a specific sample, where p_k denotes the probability that subjects choose the k -th description. We model the quality evaluation process as follows:

$$\hat{p}_k = \text{softmax}(\psi(\phi_1(x), \phi_2(y_k))), \quad (2)$$

where $\phi_1(\cdot)$ and $\phi_2(\cdot)$ denote visual and textual feature extraction operations, respectively. $\psi(\cdot)$ calculates the similarities between visual features and textual features of different quality descriptions, which are then transformed into probabilities using the softmax operation. Note that we directly derive the probabilities by comparing visual and textual features and regard them as the uncertainty in subjective experiments. As a result, the process of subjective evaluation can be approximated as sampling from this distribution.

By employing the novel quality prediction strategy, we actually divide the PCQA task into two steps: multi-modal feature extraction and vision-language alignment. We use the visual and textual encoders in CLIP to extract multi-modal features and then perform vision-language alignment using the OSD. A contrastive loss and learnable prompts are introduced to enhance the feature extraction. The proposed quality prediction procedure is more consistent with the way how humans perceive point cloud quality and we will present the detailed pipeline of our method in the following.

Preprocessing

CLIP has a powerful capability to align visual data with language concepts. However, its visual encoder is pre-trained on 2D images while point clouds are sets of unordered points in 3D space. To convert point clouds into CLIP-accessible inputs, we project the 3D point clouds into multi-view images. To reduce information loss during the projection process, we obtain both color and depth maps from M viewpoints simultaneously, which is clearly shown in Figure 2. Color maps offer texture details about point clouds, while depth maps provide geometric information about the shape,

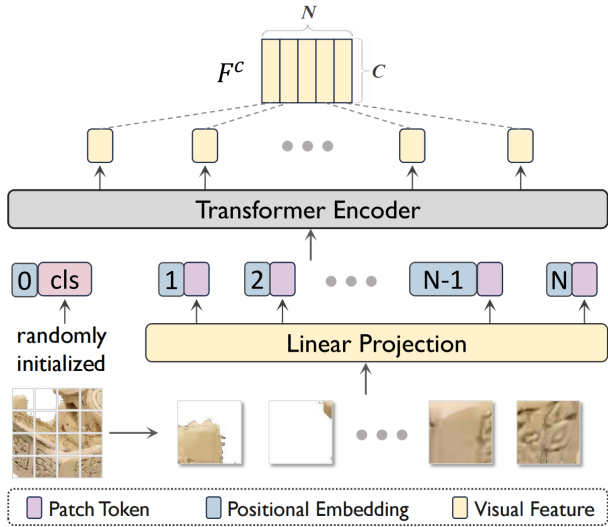


Figure 3: Illustration of the visual encoder architecture.

structure, and point distribution. Given a point cloud X , We denote its multi-view color and depth maps by $\{X_m^c \in \mathbb{R}^{H \times W \times 3} |_{m=1}^M\}$ and $\{X_m^d \in \mathbb{R}^{H \times W} |_{m=1}^M\}$ respectively, where m denotes the m -th viewpoint.

Multi-Modal Feature Extraction

Visual Feature Extraction We first intend to obtain visual features from point clouds. Note that previous studies usually freeze the visual encoder of CLIP to perform zero-shot prediction. However, features obtained from the frozen encoder can lead to a large performance gap for PCQA task because of the differences between point cloud projections and natural images. Therefore, we choose to finetune the visual encoder of CLIP in the training process.

The visual encoder in CLIP is capable of processing depth maps (Zhang et al. 2022a). Given the distinct characteristics of color and depth maps, we employ two pretrained Vision Transformer (ViT) from CLIP, denoted by ϕ_c and ϕ_d , which are independently updated during the training. To make the projected images compatible with the ViT, for X^c or X^d from an arbitrary perspective, we partition each image into N non-overlapping patches and feed the patches into the encoders. More specifically, we first flatten the patches and map them to C -dimension patch embeddings using a trainable linear projection. Together with a randomly initialized class token, the token sequence can be denoted as $z_0 = \{f_{cls}, f_1, \dots, f_N\} \in \mathbb{R}^{(N+1) \times C}$.

Then the position embeddings are added to all tokens to retain positional information. Subsequently, we input these tokens into the transformer consisting of L blocks. For instance, the procedure in the l -th block can be formulated as:

$$z_l = \text{MLP}(\text{LN}(z'_l)) + z'_l, z'_l = \text{MSA}(\text{LN}(z_{l-1})) + z_{l-1}, \quad (3)$$

where LN stands for layer normalization, MLP denotes the multi-layer perceptron and MSA denotes the multi-head self-attention.

Finally, we could obtain the last-layer output of all tokens $z_L \in \mathbb{R}^{(N+1) \times C}$. To better account for multiple local distortions across the entire image, we use the feature map of all patches with the dimension of $N \times C$ instead of the class token as in CLIP.

We illustrate the whole feature extraction process for an arbitrary color map in Figure 3. Mathematically, given the color map and depth map of the m -th viewpoint (i.e., X_m^c and X_m^d), the output features can be formulated as:

$$\begin{aligned} F_m^c &= \{f_{n,m}^c\}_{n=1}^N = \phi_c(X_m^c) \in \mathbb{R}^{N \times C}, \\ F_m^d &= \{f_{n,m}^d\}_{n=1}^N = \phi_d(X_m^d) \in \mathbb{R}^{N \times C}, \end{aligned} \quad (4)$$

where $f_{n,m}^c \in \mathbb{R}^C$ (or $f_{n,m}^d \in \mathbb{R}^C$) denote the output token corresponding to the n -th color or depth patch of the m -th viewpoint. To merge the information of different patches and viewpoints, the final visual feature is calculated as follows:

$$F^I = \frac{1}{2MN} \sum_{n=1}^N \sum_{m=1}^M (f_{n,m}^c + f_{n,m}^d) \in \mathbb{R}^C. \quad (5)$$

To enable the model to distinguish different samples and viewpoints, we introduce a contrastive scheme. Specifically, the color and depth maps from the same view of a specific sample are considered positive pairs, while others are treated as negative ones. We first flattened the features in Equation (4) into $\tilde{F}_m^c \in \mathbb{R}^{NC}$ and $\tilde{F}_m^d \in \mathbb{R}^{NC}$, respectively. Then the contrastive loss is calculated as follows:

$$\begin{aligned} L_{con} &= \frac{1}{2BM} \sum_{b=1}^B \sum_{m=1}^M (l_{con}(\tilde{F}_{m,b}^c, \tilde{F}_{m,b}^d) + l_{con}(\tilde{F}_{m,b}^d, \tilde{F}_{m,b}^c)), \\ l_{con}(i, j) &= -\log \frac{\exp(\cos(i, j)/\tau_1)}{\sum_{k=1}^{BM} \mathbf{1}_{[k \neq i]} \exp(\cos(i, k)/\tau_1)}, \end{aligned} \quad (6)$$

where $\tilde{F}_{m,b}^c$ (or $\tilde{F}_{m,b}^d$) denotes the flattened color (or depth) feature corresponding to the m -th viewpoint of the b -th sample in a mini-batch. B is the batch size, $\mathbf{1}_{[k \neq i]}$ is an indicator function evaluating to 1 if $k \neq i$, $\cos(\cdot)$ calculates the cosine similarity between two features and τ_1 is temperature. Obviously, the contrastive loss minimizes the distance between color and depth features from the same view of a specific sample, while maximizing the distance of other pairs, which helps the network better characterize the point cloud data.

Textual Feature Extraction The textual encoder in CLIP is capable of projecting textual descriptions into the embedding space close to the features of the paired images. To take advantage of this semantic analysis capability, we convert different quality-related adjectives into prompts and freeze the textual encoder to extract the corresponding features.

Concretely, to simulate the training session in subjective experiments, we adopt the widely used five quality-related adjectives from ITU-R BT.500 (Recommendation 2019) (i.e., “excellent”, “good”, “fair”, “poor” and “bad”) as K quality level tokens and obtain their embeddings as $A = \{A_k\}_{k=1}^K$. Considering that fixed prompts (e.g., “the quality of this image is excellent” or “an image of excellent quality”) may restrict the space of model optimization, we introduce the unified learnable prompt version following

CoOp (Zhou et al. 2022). More specifically, we insert the quality level tokens into the middle of the template sentence to acquire the quality-related prompts $Y = \{Y_k\}_{k=1}^K$, that is

$$Y_k = [V]_1 \dots [V]_{W/2} [A_k] [V]_{W/2+1} \dots [V]_W, \quad (7)$$

where each $[V]$ is a learnable context token with the same dimension as word embedding, and W is a hyper-parameter specifying the number of context tokens.

Subsequently, these prompts will pass through the frozen textual encoder ϕ_t to obtain a list of features as:

$$F^T = \{F_k^T\}_{k=1}^K \in \mathbb{R}^{K \times C}, F_k^T = \phi_t(Y_k), \quad (8)$$

where F_k^T denotes the textual feature corresponding to the k -th quality description.

Vision-Language Alignment

The aforementioned analysis argues that participants in subjective experiments tend to first match the distorted samples with quality descriptions and then represent the quality descriptions using quantitative values. To mimic such a process, we first calculate the cosine similarities between image features and multiple textual features corresponding to different quality descriptions:

$$\pi_k = \frac{F^I \cdot (F_k^T)^T}{\|F^I\| \|F_k^T\|}. \quad (9)$$

The larger the similarity, the more likely the corresponding quality description is to be retrieved. Given the existence of personal limitations and bias in subjective experiments, we further transform the similarities into probabilities. Concretely, we leverage a softmax function to convert the similarity values into probabilities to acquire the predicted OSD $\hat{P} = \{\hat{p}_k\}_{k=1}^K$ as:

$$\hat{p}_k = \frac{\exp(\pi_k)}{\sum_{i=1}^K \exp(\pi_i)}, \quad (10)$$

where we treat the output softmax probabilities as the predicted quality score distribution. Generally, we regard that the derived probabilities indicate the uncertainty in subjective experiments. As a result, the process of subjective evaluation can be approximated as sampling from this distribution. We expect \hat{P} to approximate the ground truth OSD during the training phase.

In practical application scenarios, we prefer the model to output an overall quality score for a given point cloud sample, which benefits the quality-oriented optimization. Therefore, in the test phase, we quantify the K quality-related descriptions into certain quantitative values. For example, ["excellent", "good", "fair", "poor" and "bad"] are projected to quality scores $q = \{q_k\}_{k=1}^K = [5, 4, 3, 2, 1]$ for databases with MOS = 5 as the best quality. Using the softmax probabilities as weights, the final predicted score \hat{Q} is calculated by a linear combination of the predicted OSD and the score values of different quality levels:

$$\hat{Q} = \sum_{k=1}^K \hat{p}_k * q_k, \text{ for } k = 1, \dots, K. \quad (11)$$

Loss Function

Based on the existing findings (Hoßfeld, Heegaard, and Varela 2015), subjective uncertainty can be expressed by the Cumulative Distribution Function (CDF). Therefore, to constrain the network optimization, we leverage the Earth Mover's Distance (EMD) following (Hou, Yu, and Samaras 2016) to measure the discrepancy between the predicted OSD $\hat{P} = \{\hat{p}_k\}_{k=1}^K$ and the ground truth OSD $P = \{p_l\}_{l=1}^L$ as:

$$L_{emd} = \left(\frac{1}{K} \sum_{k=1}^K |\text{CDF}_{\hat{P}}(q_k) - \text{CDF}_P(q_k)|^2 \right)^{1/2}, \quad (12)$$

where L represents the number of score options available to participants in the original scoring. $\text{CDF}(\cdot)$ returns the cumulative probability and can be formulated as:

$$\text{CDF}_P(q_k) = \Pr(Q \leq q_k). \quad (13)$$

Following (Gao et al. 2022), we further introduce a quantile-based loss function. Quantiles refer to numerical points that divide the probability distribution range of a random variable into several equal parts, and θ -quantile denotes the specific score whose CDF equals θ . Since $\text{CDF}_{\hat{P}}$ and CDF_P are discrete, we first use a linear interpolation algorithm to convert them into continuous distributions. Then the quantile-based loss function can be calculated as:

$$L_{quan} = \frac{1}{J} \sum_{j=1}^J |s_{\hat{P}}(\theta_j) - s_P(\theta_j)|, \quad (14)$$

where $\theta_j \in \theta = [\theta_1, \theta_2, \dots, \theta_J]$ and J is the number of quantiles we calculated. $s_{\hat{P}}(\theta_j)$ and $s_P(\theta_j)$ represent the predicted θ_j -quantile and the ground-truth θ_j -quantile.

The final loss function is the combination of three losses:

$$Loss = L_{emd} + \alpha L_{quan} + \beta L_{con}, \quad (15)$$

where α and β are hyper-parameters used to control the scale of the three loss functions.

4 Experiments

Databases and Evaluation Metrics

Databases. To illustrate the effectiveness of our method, we employ three benchmarks with available raw opinion scores: SJTU-PCQA (Yang et al. 2020a), LS-PCQA Part I (Liu et al. 2023b) and BASICS (Ak et al. 2024). SJTU-PCQA includes 9 native point clouds and 378 distorted point clouds disturbed by 7 types of distortion under 6 levels. LS-PCQA comprises 85 native point clouds, and a total of 930 labeled point clouds generated by 31 distortion types. BASICS consists of 1494 distorted point clouds impaired by 4 types of compression, which are generated from 75 references.

Evaluation Metrics. Three widely adopted evaluation metrics are employed to quantify the level of agreement between the predicted quality scores and the MOSs: Pearson Linear Correlation Coefficient (PLCC), Spearman Rank Correlation Coefficient (SRCC), and Root Mean Squared Error (RMSE). Since there exists a misalignment in the value range between the predicted scores and the MOSs, we introduce a common four-parameter logistic function (Ankowiak et al. 2000) to align their ranges.

| Type | Method | SJTU-PCQA (Yang et al. 2020a) | | | LS-PCQA (Liu et al. 2023b) | | | BASICS (Ak et al. 2024) | | |
|------|---------------------------------------|-------------------------------|-----------------|-------------------|----------------------------|-----------------|-------------------|-------------------------|-----------------|-------------------|
| | | PLCC \uparrow | SRCC \uparrow | RMSE \downarrow | PLCC \uparrow | SRCC \uparrow | RMSE \downarrow | PLCC \uparrow | SRCC \uparrow | RMSE \downarrow |
| FR | MSE-p2po (Mekuria et al. 2016) | 0.896 | 0.810 | 1.046 | 0.427 | 0.301 | 0.744 | 0.849 | 0.774 | 0.559 |
| | MSE-p2pl (Mekuria et al. 2016) | 0.783 | 0.696 | 1.471 | 0.454 | 0.286 | 0.734 | 0.895 | 0.836 | 0.468 |
| | PSNR $_{yuv}$ (Torlig et al. 2018) | 0.764 | 0.762 | 1.513 | 0.527 | 0.482 | 0.699 | 0.617 | 0.592 | 0.835 |
| | PCQM (Meynet et al. 2020) | 0.836 | 0.874 | 2.378 | 0.208 | 0.426 | 0.789 | 0.891 | 0.808 | 0.478 |
| | PointSSIM (Alexiou and Ebrahimi 2020) | 0.751 | 0.725 | 1.539 | 0.225 | 0.164 | 0.804 | 0.780 | 0.731 | 0.657 |
| | MPED (Yang et al. 2022) | 0.903 | 0.900 | 1.000 | 0.613 | 0.609 | 0.646 | 0.807 | 0.712 | 0.615 |
| | GraphSIM (Yang et al. 2020b) | 0.896 | 0.874 | 1.040 | 0.358 | 0.331 | 0.767 | 0.895 | 0.813 | 0.465 |
| | TCDM (Zhang et al. 2023) | 0.952 | 0.929 | 0.713 | 0.438 | 0.413 | 0.739 | 0.874 | 0.757 | 0.505 |
| NR | PQA-Net (Liu et al. 2021) | 0.898 | 0.875 | 1.340 | 0.607 | 0.595 | 0.697 | 0.655 | 0.387 | 0.797 |
| | ResSCNN (Liu et al. 2023b) | 0.889 | 0.880 | 0.878 | 0.648 | 0.620 | 0.615 | 0.391 | 0.352 | 0.975 |
| | MM-PCQA (Zhang et al. 2022b) | 0.939 | 0.910 | 0.805 | 0.644 | 0.605 | 0.621 | 0.793 | 0.738 | 0.628 |
| | GMS-3DQA (Zhang et al. 2024) | 0.916 | 0.886 | 0.931 | 0.666 | 0.645 | 0.606 | 0.895 | 0.807 | 0.472 |
| | 3DTA (Zhu et al. 2024) | 0.953 | 0.931 | 0.706 | 0.613 | 0.604 | 0.645 | 0.907 | 0.825 | 0.437 |
| | CLIP-PCQA (ours) | 0.956 | 0.936 | 0.693 | 0.755 | 0.736 | 0.533 | 0.932 | 0.872 | 0.382 |

Table 1: Benchmarking results of SOTA methods on the SJTU-PCQA, LS-PCQA and BASICS databases. “ \uparrow ”/“ \downarrow ” indicates that larger/smaller is better. Best in **red** and the second in **blue**.

| Train on | SJTU-PCQA | | LS-PCQA | |
|------------------|--------------|--------------|--------------|--------------|
| | LS-PCQA | BASICS | SJTU-PCQA | BASICS |
| ResSCNN | 0.300 | 0.092 | 0.472 | 0.078 |
| MM-PCQA | 0.150 | 0.198 | 0.728 | 0.445 |
| GMS-3DQA | 0.366 | 0.575 | 0.811 | 0.457 |
| 3DTA | 0.288 | 0.633 | 0.671 | 0.448 |
| CLIP-PCQA | 0.326 | 0.685 | 0.873 | 0.570 |

Table 2: Evaluation of cross-database generalization. The training and test are all performed on the complete database. PLCC are reported, best in **red** and the second in **blue**.

Implementation Details

Database Split. We partition the databases according to content (reference point clouds) and k-fold cross-validation is used for training. Specifically, 9-fold cross-validation is applied for SJTU-PCQA following (Zhang et al. 2022b), and we adopt a 5-fold cross-validation both for LS-PCQA and BASICS. For each fold, the test performance with minimal training loss is recorded and the average result across all folds is recorded to alleviate randomness.

Network Details. We use the basic version of Vision Transformer (Dosovitskiy et al. 2020) with 16×16 patch embeddings (namely ViT-B/16) in CLIP-Visual as our visual encoders. The textual encoder is also the pre-trained transformer in CLIP-Text. The number of projection views $M = 6$ and the images are randomly cropped into $224 \times 224 \times 3$ as inputs. We set the number of context tokens W as 16. For the loss function, we set $\theta = [0.25, 0.50, 0.75]$. α is set to $1/K$ and β is set to 0.08.

Training Strategy. The initial learning rate is set as $4e-6$ and the model is trained for 50 epochs with a default batch size of 16. We use the Adam optimizer (Kingma and Ba 2014) with a weight decay of $1e-4$. Depending on the raw score ranges of different databases, we evenly divide them into five thresholds as the quantitative values q . For example, we set $q = [5, 4, 3, 2, 1]$ for LS-PCQA and




Figure 4: Weight matrix visualization. Darker colors indicate higher values, which mean greater correlations.

$q = [10, 8, 6, 4, 2]$ for SJTU-PCQA, respectively.

Performance Comparison

Overall Performance. For the sake of comprehensively investigating the prediction performance of CLIP-PCQA, 13 SOTA methods are selected for comparison, including 8 FR-PCQA methods and 5 NR-PCQA methods. We strictly retrain the baselines with available code, using the same database split to ensure a fair comparison. The experimental results are illustrated in Table 1. Benefiting from data-driven deep learning approaches, many existing NR methods outperform traditional FR methods. By introducing vision-language alignment and distribution prediction, our method CLIP-PCQA yields the best performance across three databases. Notably, our method significantly surpasses the second-best results on the LS-PCQA (exceeding 0.089 in PLCC and 0.091 in SRCC). Furthermore, while the performance of other methods varies greatly across different databases, our approach is robust to databases of different domains or various point cloud contents and distortions.

Cross-database Evaluation. To test the generalizability of our method, we further perform cross-database evaluation. We train the compared methods on SJTU-PCQA and

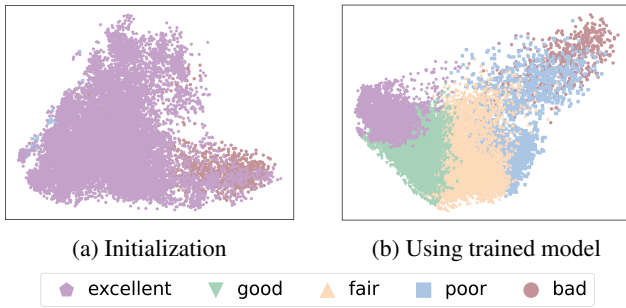


Figure 5: PCA of the visual features on complete LS-PCQA.

| Index | Modality | PLCC \uparrow | SRCC \uparrow | RMSE \downarrow |
|-------|-----------------------------|-----------------|-----------------|-------------------|
| (1) | Color + Text | 0.738 | 0.716 | 0.550 |
| (2) | Depth + Text | 0.536 | 0.395 | 0.694 |
| (3) | Color + Depth | 0.688 | 0.652 | 0.585 |
| (4) | Color + Depth + Text | 0.755 | 0.736 | 0.533 |

Table 3: Contributions of different modalities on LS-PCQA.

LS-PCQA, and test the trained model on the other databases. The results are reported in Table 2, where our CLIP-PCQA achieves the best generalization performance in most cases. Since the contents are limited in SJTU-PCQA while abundant in LS-PCQA, almost all methods perform poorly when trained on SJTU-PCQA and tested on LS-PCQA. It is worth noting that our method exceeds the other counterparts by a noticeable margin when trained on LS-PCQA and generalized to SJTU-PCQA. Our cross-database result of 0.873 even presents competitive performance compared to the intra-database results in Table 1, which indicates that our method learned more general quality-related features when trained on databases with richer contents and distortions.

Visualization

To validate that our model has learned the correspondence between the quality-related descriptions and the distortion degrees, we provide qualitative results in this section.

First, we visualize the output softmax probabilities of different quality levels in Figure 4. It can be observed that samples with more severe distortions exhibit a stronger correlation with text descriptions indicating lower quality, which confirms that quality information is indeed encapsulated within the textual features.

In addition, we use Principal Component Analysis (PCA) to embed visual features into 2D representation space, which is illustrated in Figure 5. Specifically, to intuitively demonstrate the clustering results of the samples with different quality, we assign each sample to the quality level category with the highest probability in the softmax output. Since PCQA is essentially a continuous regression task rather than a classification task, it is nearly impossible to distinctly classify all samples into five categories. From Figure 5, we can observe that distorted samples within the same category are clustered in the feature space. Moreover, the adjacent categories exhibit continuity, indicating the lack of clear bound-

| Type | Position | PLCC \uparrow | SRCC \uparrow | RMSE \downarrow |
|------------------|---------------|-----------------|-----------------|-------------------|
| fixed | - | 0.729 | 0.696 | 0.559 |
| | begin | 0.728 | 0.702 | 0.559 |
| learnable | middle | 0.755 | 0.736 | 0.533 |
| | end | 0.742 | 0.719 | 0.548 |

Table 4: Ablation on the prompt type and the position of quality level token on LS-PCQA.

| L_{emd} | L_{quan} | L_{con} | PLCC \uparrow | SRCC \uparrow | RMSE \downarrow |
|-----------|------------|-----------|-----------------|-----------------|-------------------|
| ✓ | | | 0.734 | 0.717 | 0.553 |
| | ✓ | | 0.649 | 0.606 | 0.592 |
| ✓ | ✓ | | 0.743 | 0.721 | 0.548 |
| ✓ | | ✓ | 0.745 | 0.722 | 0.542 |
| ✓ | ✓ | ✓ | 0.755 | 0.736 | 0.533 |

Table 5: Ablation of the loss function on LS-PCQA.

aries between different quality levels.

Ablation Studies

Contributions of Modalities. We conduct experiments to validate the effectiveness of three types of input and report the results in Table 3. Seeing (1), (2) and (4), combining color and depth maps enhances the model’s ability to gain better quality representations. Since color maps include more visual information, the performance deteriorates when only depth maps are used. Moreover, we assess the performance of removing the text branch and using a two-fold MLP for regression on visual features in (3). Comparing (3) and (4), we can conclude that incorporating language significantly improves performance.

Prompt Design. To validate the effectiveness of learnable prompts, we perform experiments with fixed prompt “The quality of this image is [CLASS]”. Additionally, we test the performance with different positions of the class token and report the results in Table 4. We can observe that appropriate learnable prompts outperform fixed prompts and alleviate the need for meticulous prompt design. Moreover, positioning the quality level token in the middle performs best, where bidirectional contextual information helps the model better understand the semantics of the class token.

Loss Function. We explore how the three loss functions contribute to our model in Table 5. It can be seen that the predictive performance using only EMD or quantile-based loss is inferior to using both. Additionally, the contrastive loss L_{con} helps the model further improve performance.

5 Conclusion

In this paper, to simulate the HVS mechanism and reflect subjective diversity and uncertainty, we propose a novel language-driven method CLIP-PCQA, which predicts quality according to discrete quality descriptions and quality score distribution. Experiments demonstrate that our method yields the SOTA performance, and we believe our study could shed light on future research.

Acknowledgments

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

References

- Agnolucci, L.; Galteri, L.; and Bertini, M. 2024. Quality-Aware Image-Text Alignment for Real-World Image Quality Assessment. *arXiv preprint arXiv:2403.11176*.
- Ak, A.; Zerman, E.; Quach, M.; Chetouani, A.; Smolic, A.; Valenzise, G.; and Le Callet, P. 2024. BASICS: Broad Quality Assessment of Static Point Clouds in a Compression Scenario. *IEEE Transactions on Multimedia*.
- Alexiou, E.; and Ebrahimi, T. 2020. Towards a point cloud structural similarity metric. In *2020 IEEE International Conference on Multimedia & Expo Workshops*.
- Antkowiak, J.; Baina, T. J.; Baroncini, F. V.; Chateau, N.; FranceTelecom, F.; Pessoa, A. C. F.; Colonnese, F. S.; Contin, I. L.; Caviedes, J.; and Philips, F. 2000. Final report from the video quality experts group on the validation of objective models of video quality assessment march 2000. *Final report from the video quality experts group on the validation of objective models of video quality assessment march*.
- Cui, Y.; Chen, R.; Chu, W.; Chen, L.; Tian, D.; Li, Y.; and Cao, D. 2021. Deep learning for image and point cloud fusion in autonomous driving: A review. *IEEE Transactions on Intelligent Transportation Systems*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Fan, S.; Guo, Z.; Gao, W.; and Li, G. 2024. Uncertainty-aware No-Reference Point Cloud Quality Assessment. *arXiv preprint arXiv:2401.08926*.
- Gao, Y.; Min, X.; Zhu, Y.; Li, J.; Zhang, X.-P.; and Zhai, G. 2022. Image quality assessment: From mean opinion score to opinion score distribution. In *Proceedings of the 30th ACM International Conference on Multimedia*.
- Gu, X.; Lin, T.-Y.; Kuo, W.; and Cui, Y. 2021. Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*.
- He, W.; Jamonnak, S.; Gou, L.; and Ren, L. 2023. CLIP-S4: Language-Guided Self-Supervised Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Höfßfeld, T.; Heegaard, P. E.; and Varela, M. 2015. QoE beyond the MOS: Added value using quantiles and distributions. In *2015 Seventh International Workshop on Quality of Multimedia Experience*.
- Hou, L.; Yu, C.-P.; and Samaras, D. 2016. Squared earth mover's distance-based loss for training deep neural networks. *arXiv preprint arXiv:1611.05916*.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Liu, H.; Wu, M.; Yuan, K.; Sun, M.; Tang, Y.; Zheng, C.; Wen, X.; and Li, X. 2023a. Ada-DQA: Adaptive Diverse Quality-aware Feature Acquisition for Video Quality Assessment. In *Proceedings of the 31st ACM International Conference on Multimedia*.
- Liu, Q.; Yuan, H.; Su, H.; Liu, H.; Wang, Y.; Yang, H.; and Hou, J. 2021. PQA-Net: Deep no reference point cloud quality assessment via multi-view projection. *IEEE transactions on circuits and systems for video technology*.
- Liu, Y.; Yang, Q.; Xu, Y.; and Yang, L. 2023b. Point cloud quality assessment: Dataset construction and learning-based no-reference metric. *ACM Transactions on Multimedia Computing, Communications and Applications*.
- Liu, Y.; Yang, Q.; Zhang, Y.; Xu, Y.; Yang, L.; Xu, X.; and Liu, S. 2023c. Once-training-all-fine: No-reference point cloud quality assessment via domain-relevance degradation description. *arXiv preprint arXiv:2307.01567*.
- Luo, H.; Ji, L.; Zhong, M.; Chen, Y.; Lei, W.; Duan, N.; and Li, T. 2022. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*.
- Mekuria, R.; Li, Z.; Tulvan, C.; and Chou, P. 2016. Evaluation criteria for point cloud compression. *ISO/IEC MPEG*.
- Meynet, G.; Nehmé, Y.; Digne, J.; and Lavoué, G. 2020. PCQM: A full-reference quality metric for colored 3D point clouds. In *2020 Twelfth International Conference on Quality of Multimedia Experience*.
- Mu, B.; Shao, F.; Chai, X.; Liu, Q.; Chen, H.; and Jiang, Q. 2023. Multi-view aggregation transformer for no-reference point cloud quality assessment. *Displays*.
- Park, Y.; Lepetit, V.; and Woo, W. 2008. Multiple 3d object tracking for augmented reality. In *2008 7th IEEE/ACM International Symposium on Mixed and Augmented Reality*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*.
- Recommendation, I.-R. B. 2019. Methodologies for the subjective assessment of the quality of television images.
- Shan, Z.; Yang, Q.; Ye, R.; Zhang, Y.; Xu, Y.; Xu, X.; and Liu, S. 2023. Gpa-net: No-reference point cloud quality assessment with multi-task graph convolutional network. *IEEE Transactions on Visualization and Computer Graphics*.
- Srinath, S.; Mitra, S.; Rao, S.; and Soundararajan, R. 2024. Learning Generalizable Perceptual Representations for Data-Efficient No-Reference Image Quality Assessment. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*.
- Tliba, M.; Chetouani, A.; Valenzise, G.; and Dufaux, F. 2023. Quality evaluation of point clouds: a novel no-reference approach using transformer-based architecture. *arXiv preprint arXiv:2303.08634*.

- Torlig, E. M.; Alexiou, E.; Fonseca, T. A.; de Queiroz, R. L.; and Ebrahimi, T. 2018. A novel methodology for quality assessment of voxelized point clouds. In *Applications of Digital Image Processing XLI*.
- Wang, J.; Chan, K. C.; and Loy, C. C. 2023. Exploring clip for assessing the look and feel of images. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Wang, J.; Gao, W.; and Li, G. 2023. Applying collaborative adversarial learning to blind point cloud quality measurement. *IEEE Transactions on Instrumentation and Measurement*.
- Wang, M.; Xing, J.; and Liu, Y. 2021. Actionclip: A new paradigm for video action recognition. *arXiv preprint arXiv:2109.08472*.
- Wu, H.; Liao, L.; Wang, A.; Chen, C.; Hou, J.; Sun, W.; Yan, Q.; and Lin, W. 2023a. Towards robust text-prompted semantic criterion for in-the-wild video quality assessment. *arXiv preprint arXiv:2304.14672*.
- Wu, H.; Zhang, E.; Liao, L.; Chen, C.; Hou, J.; Wang, A.; Sun, W.; Yan, Q.; and Lin, W. 2023b. Towards explainable in-the-wild video quality assessment: a database and a language-prompted approach. In *Proceedings of the 31st ACM International Conference on Multimedia*.
- Xie, W.; Wang, K.; Ju, Y.; and Wang, M. 2023. pmbqa: Projection-based blind point cloud quality assessment via multimodal learning. In *Proceedings of the 31st ACM International Conference on Multimedia*.
- Yang, Q.; Chen, H.; Ma, Z.; Xu, Y.; Tang, R.; and Sun, J. 2020a. Predicting the perceptual quality of point cloud: A 3d-to-2d projection-based exploration. *IEEE Transactions on Multimedia*.
- Yang, Q.; Ma, Z.; Xu, Y.; Li, Z.; and Sun, J. 2020b. Inferring point cloud quality via graph similarity. *IEEE transactions on pattern analysis and machine intelligence*.
- Yang, Q.; Zhang, Y.; Chen, S.; Xu, Y.; Sun, J.; and Ma, Z. 2022. MPED: Quantifying point cloud distortion based on multiscale potential energy discrepancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Zhang, R.; Guo, Z.; Zhang, W.; Li, K.; Miao, X.; Cui, B.; Qiao, Y.; Gao, P.; and Li, H. 2022a. Pointclip: Point cloud understanding by clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Zhang, Y.; Yang, Q.; Zhou, Y.; Xu, X.; Yang, L.; and Xu, Y. 2023. TCDM: Transformational Complexity Based Distortion Metric for Perceptual Point Cloud Quality Assessment. *IEEE Transactions on Visualization and Computer Graphics*.
- Zhang, Z.; Sun, W.; Min, X.; Zhou, Q.; He, J.; Wang, Q.; and Zhai, G. 2022b. MM-PCQA: Multi-modal learning for no-reference point cloud quality assessment. *arXiv preprint arXiv:2209.00244*.
- Zhang, Z.; Sun, W.; Wu, H.; Zhou, Y.; Li, C.; Chen, Z.; Min, X.; Zhai, G.; and Lin, W. 2024. Gms-3dqa: Projection-based grid mini-patch sampling for 3d model quality assessment. *ACM Transactions on Multimedia Computing, Communications and Applications*.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022. Learning to prompt for vision-language models. *International Journal of Computer Vision*.
- Zhu, L.; Cheng, J.; Wang, X.; Su, H.; Yang, H.; Yuan, H.; and Korhonen, J. 2024. 3DTA: No-Reference 3D Point Cloud Quality Assessment with Twin Attention. *IEEE Transactions on Multimedia*.