

QD-PCQA: Quality-Aware Domain Adaptation for Point Cloud Quality Assessment

Guohua Zhang¹ Jian Jin² Meiqin Liu^{1*} Chao Yao³ Weisi Lin²

¹Beijing Jiaotong University ²Nanyang Technological University

³University of Science and Technology Beijing

24125207@bjtu.edu.cn, jianj008@gmail.com, mqliu@bjtu.edu.cn,

yaochao1986@gmail.com, wslin@ntu.edu.sg

Abstract

No-Reference Point Cloud Quality Assessment (NR-PCQA) still struggles with generalization, primarily due to the scarcity of annotated point cloud datasets. Since the Human Visual System (HVS) drives perceptual quality assessment independently of media types, prior knowledge on quality learned from images can be repurposed for point clouds. This insight motivates adopting Unsupervised Domain Adaptation (UDA) to transfer quality-relevant priors from labeled images to unlabeled point clouds. However, existing UDA-based PCQA methods often overlook key characteristics of perceptual quality, such as sensitivity to quality ranking and quality-aware feature alignment, thereby limiting their effectiveness. To address these issues, we propose a novel Quality-aware Domain adaptation framework for PCQA, termed QD-PCQA. The framework comprises two main components: i) a Rank-weighted Conditional Alignment (RCA) strategy that aligns features under consistent quality levels and adaptively emphasizes misranked samples to reinforce perceptual quality ranking awareness; and ii) a Quality-guided Feature Augmentation (QFA) strategy, which includes quality-guided style mixup, multi-layer extension, and dual-domain augmentation modules to augment perceptual feature alignment. Extensive cross-domain experiments demonstrate that QD-PCQA significantly improves generalization in NR-PCQA tasks.

1. Introduction

Point clouds, as 3D representations of objects or scenes [4], are widely used in applications such as Virtual Reality (VR) [41], Augmented Reality (AR) [9], 3D modeling [23], and autonomous driving [3]. However, point clouds are inevitably degraded during acquisition, processing, storage, and transmission [8]. Therefore, accurate Point Cloud

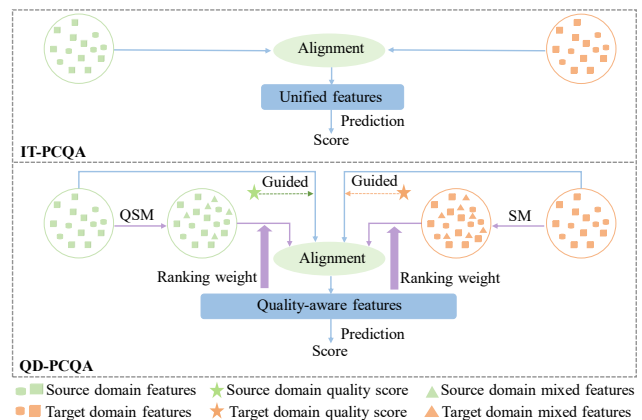


Figure 1. Comparison with IT-PCQA. IT-PCQA performs feature alignment via UDA, overlooking the characteristics of quality perception. In contrast, our QD-PCQA introduces two quality-aware strategies. The RCA strategy aligns features guided by quality scores of two domains, promoting quality-consistent adaptation. The QFA strategy enriches feature diversity through a multi-layer QSM and SM. Additionally, a rank-weighted module emphasizes misranked feature pairs to mitigate ranking bias.

Quality Assessment (PCQA) metrics are crucial.

PCQA metrics are commonly classified into Full-Reference (FR), Reduced-Reference (RR), and No-Reference (NR) methods. NR methods evaluate quality independently, without the original reference point cloud, whereas FR and RR methods rely on complete or partial reference information. Unfortunately, reference point clouds are often unavailable in practice due to transmission and storage constraints [31]. Therefore, NR methods are critical and are typically categorized into handcrafted and learning-based schemes [17, 30]. The former rely on manually designed features (*e.g.*, geometry distortion, color deviation, and statistical descriptors), while the latter are data-driven models [25, 39]. Benefiting from deep learning, data-driven methods achieve better performance and are more promis-

*Corresponding author: mqliu@bjtu.edu.cn

ing [13, 14, 16]. However, compared with traditional vision tasks such as image classification, object detection, and even Image Quality Assessment (IQA), the annotated PCQA data are scarce, which restricts generalization abilities [38].

Unsupervised Domain Adaptation (UDA) [34] provides effective solutions to address the generalization challenge caused by limited labeled data by transferring knowledge from a labeled source domain to an unlabeled target domain, which has demonstrated its effectiveness on typical vision tasks, especially for image classification tasks. Generally, a classical UDA method directly aligns feature representations across source and target domains to enable knowledge transfer without target labels [12]. Inspired by this, several works attempted to apply the UDA to the Quality Assessment (QA) tasks. For instance, Chen *et al.* [1] conducted the first attempt that utilized Maximum Mean Discrepancy (MMD) [28] to achieve UDA for the IQA task. Given that the Human Visual System (HVS) governs perceptual quality assessment independent of the specific media types (*e.g.*, image, point cloud, or 3D mesh), Yang *et al.* [38] proposed IT-PCQA, which transfers the capability of quality prediction from images to point clouds, benefiting from the availability of large-scale annotated image datasets. Specifically, the features of images and point clouds were directly aligned to a unified distribution through adversarial-based Domain Adversarial Neural Network (DANN) [7], and image labels were used for point cloud quality regression training to alleviate the problem of insufficient point cloud quality labels. Meanwhile, Lu *et al.* proposed the StyleAM [22], which introduced Style Mixup (SM) to augment source features and incorporates a style feature space into the alignment process, leading to further performance gains.

However, these methods inherit feature alignment strategies from image classification tasks, which focus on semantic consistency but ignore quality information. This may lead to unexpected **quality-regardless feature alignment**, where features with similar semantics but different quality levels are incorrectly aligned. For example, mapping high-quality “tree” images in the source domain to low-quality “tree” images in the target domain. Such misalignment compromises the model’s ability to distinguish perceptual quality, limiting its effectiveness in quality assessment tasks. Besides, the strength of feature alignment is always unchanged, which should also be refined. Moreover, although StyleAM [22] improves performance through feature augmentation, it introduces several additional issues. First, **quality-regardless feature augmentation**: SM performs random interpolation of style features without considering quality information, resulting in augmented features that fail to represent perceptual quality effectively. Second, **layer-regardless feature augmentation**: Apply-

ing augmentation solely to the final layer overlooks the complementary nature of hierarchical features. Features from shallow layers are more sensitive to low-level distortions relevant to high-quality images, while deeper layers capture high-level semantics critical for assessing lower-quality samples. Ignoring this variation reduces the representational richness of the augmented features for quality assessment. Third, **augmentation imbalance**: As only features of the source domain are augmented, which widens the domain gap, making it easy to make the discriminator to distinguish the two domains, thereby weakening the adversarial feature learning of quality representations.

To overcome the above problems, we propose a novel QD-PCQA as shown in Figure 1, consisting of two main components: i) a **Rank-weighted Conditional Alignment (RCA)** strategy, which includes a quality-aware conditional module and a rank-weighted module. The former module introduces a quality-biased feature alignment method, which uses ground-truth quality scores from the source domain and pseudo quality scores from the target domain as condition information. By aligning features with similar quality levels, it addresses the quality-regardless feature alignment issue. The latter module assigns greater weights to misranked feature pairs to strengthen their alignment constraints, leading to more accurate correction of ranking bias. ii) a **Quality-guided Feature Augmentation (QFA)** strategy, which includes a Quality-guided Style Mixup (QSM) module, a multi-layer extension module, and a dual-domain augmentation module. The QSM module integrates a novel quality-guided selection and an SM operation. The selection operation matches samples with similar quality levels, ensuring that the subsequent SM produces quality-consistent and quality-aware augmentations to solve the quality-regardless feature augmentation issue. The multi-layer extension module is designed to address the issue of layer-regardless feature augmentation by leveraging the complementary nature of hierarchical features. To this end, it first categorizes input samples into high, medium, and low quality groups by a quality stratifier, and then embeds the QSM module into the shallow, middle, and deep layers of the network, respectively. The dual-domain augmentation module applies the QSM-based multi-layer extension module in the source domain and SM in the target domain integrated into the adversarial-based DANN [7], thereby solving the augmentation imbalance problem.

The contributions can be summarized as follows:

- We propose a novel domain adaptation quality assessment framework called QD-PCQA that leverages prior quality knowledge from images to predict point cloud quality.
- We develop the RCA strategy that aligns features under consistent quality levels to solve the issue of quality-regardless feature alignment, and adaptively emphasizes misranked samples to refine alignment strength, thereby

enhancing quality ranking sensitivity.

- We develop the QFA strategy that incorporates QSM, multi-layer feature integration, and dual-domain augmentation, enabling hierarchical quality representation and augmenting perceptual feature alignment.

2. Related Work

2.1. Point Cloud Quality Assessment

Early Point Cloud Quality Assessment (PCQA) research primarily focused on FR metrics to support Point Cloud Compression (PCC). MPEG standardized several FR metrics, including point-to-point [24], point-to-plane [32], and PSNRyuv [33]. To better align with human perception, later methods incorporated Human Visual System (HVS)-based features, such as structural similarity [37], curvature [11], and color brightness [26]. However, in the absence of reference point clouds, NR-PCQA has attracted increasing attention. Chetouani [5] proposed a Convolutional Neural Network (CNN)-based method using local geometric features. Liu *et al.* [18] enhanced representation through multi-view projections. Although 2D projections may introduce some distortions and the number and angles of views can affect quality representation, they are commonly used for subjective rating, and most PCQA metrics adopt projected images or videos as inputs [38]. Yang *et al.* further bridged 2D and 3D features by adaptively weighting projection blocks according to their relative contributions, showing through subjective experiments that point cloud quality can be effectively captured using six orthogonal views[38].

The above methods require large-scale training data from the same distribution as testing, limiting generalization to unseen scenarios. Hence, we propose a UDA-based method that transfers quality prediction capability from images to point clouds, considering that HVS drives perceptual assessment independently of media types.

2.2. Unsupervised Domain Adaptation

UDA aims to mitigate domain shift between labeled source and unlabeled target domains [7, 28]. Existing methods are typically categorized into marginal, conditional, and joint distribution alignment. Marginal alignment only focuses on global feature consistency using techniques such as MMD [28] or adversarial learning (*e.g.*, Domain Adversarial Neural Network (DANN) [7]). Joint alignment [20, 21] models feature-label dependencies, while conditional alignment further leverages label information, *e.g.*, using the Conditional Operator Discrepancy (COD) [35]. Recently, UDA has also been explored in quality assessment tasks. Chen *et al.* [1] first applied MMD [28] for feature alignment in IQA. StyleAM [22] introduced a style space into DANN-based adversarial alignment [7]. IT-PCQA [38] employed DANN [7] to unify image and point cloud features.

However, existing UDA-based PCQA methods often overlook the characteristics of quality perception, leading to sub-optimal results for NR-PCQA tasks. To address this issue, we propose the novel QD-PCQA.

3. Method

3.1. Data Preprocessing Module

We project 3D point cloud data onto the six vertical sides of a cube, generating multiple side views, following IT-PCQA [38] to share a common feature extractor. These views are then stitched into a multi-view image and, along with natural images, resized to 224×224 for unified input to the subsequent feature extractor.

3.2. QFA Strategy

The proposed Quality-guided Feature Augmentation (QFA) strategy enhances feature generalization by extracting domain-invariant and quality-aware representations to support both prediction and alignment. Built upon a modified ResNet-50 [10] backbone (excluding the original classification head), the QFA strategy integrates three key components, including Quality-guided Style Mixup (QSM), multi-layer extension, and dual-domain augmentation modules. To extract deep representations for subsequent prediction and alignment, we first obtain the feature $f_s^i \in \mathbb{R}^{C \times H \times W}$ from the source domain and $f_t^j \in \mathbb{R}^{C \times H \times W}$ from the target domain, formulated as:

$$f_s^i = G(x_s^i), \quad f_t^j = G(x_t^j), \quad (1)$$

where x_s^i denotes the i -th source domain sample and x_t^j denotes the j -th target domain sample; $G(\cdot)$ denotes the feature extractor; C , H and W denote the channel, height, and width of features, respectively.

3.2.1. QSM Module

To solve the issue of quality-regardless feature augmentation, the Quality-guided Style Mixup (QSM) module is proposed, which integrates a novel quality-guided selection and a Style Mixup (SM) operation. Unlike conventional SM that randomly mixes samples, the proposed selection operation leverages a Gaussian kernel [40] to adaptively pair source samples with similar quality scores, thereby guiding the SM operation to preserve quality consistency and enhance quality-aware feature augmentation.

Specifically, to ensure that source domain samples with similar perceptual quality are more likely to be selected for mixing, we define the selection probability of a candidate pair $(x_s^{i^*}, y_s^{i^*})$ given a sample (x_s^i, y_s^i) , formulated as:

$$P((x_s^{i^*}, y_s^{i^*}) | (x_s^i, y_s^i)) \propto \exp\left(-\frac{(y_s^i - y_s^{i^*})^2}{2\tau^2}\right), \quad (2)$$

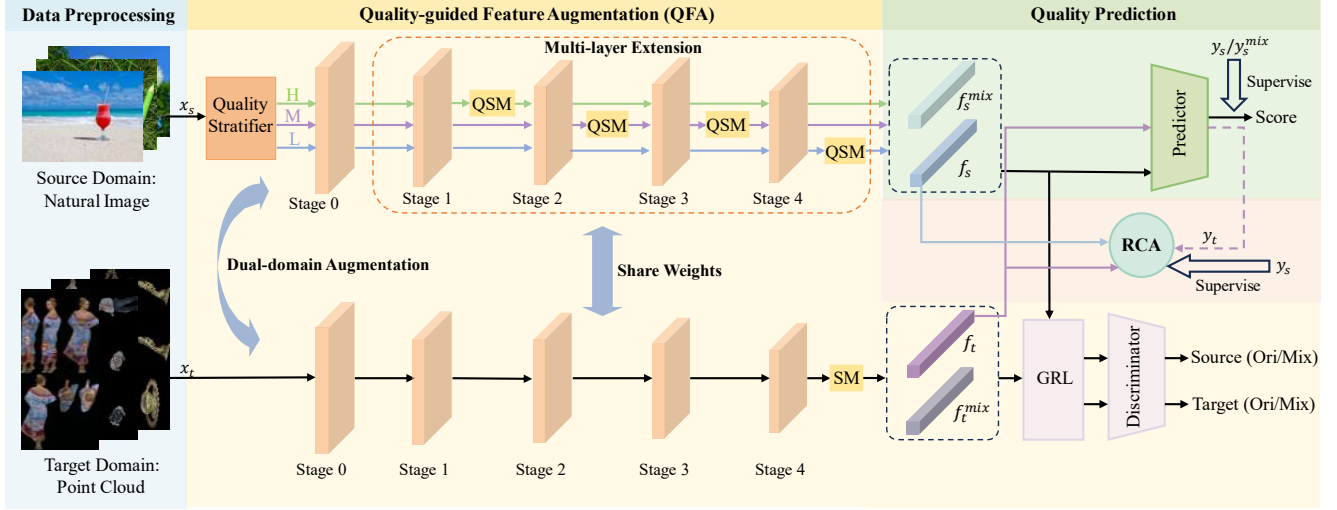


Figure 2. Architecture of the proposed QD-PCQA. Given the source-domain image x_s and target-domain point cloud x_t , both are cropped to the same size. Then processed by QFA for quality-guided feature augmentation and by RCA for rank-aware conditional alignment. Finally, the predictor outputs the final quality score.

where τ represents the bandwidth parameter of the Gaussian kernel, and y_s^i and $y_s^{i^*}$ represent the labels of source domain samples i and i^* , respectively.

After pairing, SM is applied to generate mixed features. Specifically, we compute the channel-wise mean $u(f)_c$ and standard deviation $\sigma(f)_c$ of a feature $f \in \mathbb{R}^{C \times H \times W}$, formulated as:

$$u(f)_c = \frac{1}{HW} \sum_{h=1}^H \sum_{w=1}^W f_{c,h,w},$$

$$\sigma(f)_c = \sqrt{\frac{1}{HW} \sum_{h=1}^H \sum_{w=1}^W (f_{c,h,w} - u(f)_c)^2}, \quad (3)$$

where $f_{c,h,w}$ denotes the feature at channel c , height h and width w ; H and W denote the height and width of the feature, respectively.

Given the feature styles $\langle u(f_s^i), \sigma(f_s^i) \rangle$ and $\langle u(f_s^{i^*}), \sigma(f_s^{i^*}) \rangle$ of two samples, along with their corresponding quality scores y_s^i and $y_s^{i^*}$, we construct a mixing vector λ to blend the feature styles and quality labels. The mixed feature styles $u(f)^{\text{mix}}$ and $\sigma(f)^{\text{mix}}$, as well as the mixed label y_s^{mix} , formulated as respectively:

$$u(f)^{\text{mix}} = \lambda u(f_s^i) + (1 - \lambda) u(f_s^{i^*}),$$

$$\sigma(f)^{\text{mix}} = \lambda \sigma(f_s^i) + (1 - \lambda) \sigma(f_s^{i^*}),$$

$$y_s^{\text{mix}} = \lambda y_s^i + (1 - \lambda) y_s^{i^*}, \quad (4)$$

where $\lambda \sim \text{Beta}(\alpha, \alpha)$ denotes sampled from the Beta distribution with hyperparameter $\alpha > 0$.

Finally, the mixed style feature f_s^{mix} is formulated as:

$$f_s^{\text{mix}} = \sigma(f)^{\text{mix}} \frac{f_s^i - u(f_s^i)}{\sigma(f_s^i)} + u(f)^{\text{mix}}. \quad (5)$$

3.2.2. Multi-Layer Extension Module

To solve the issue of layer-regardless feature augmentation, the multi-layer extension module is proposed to determine at which layer to use the QSM based on the quality scores of the source domain samples. As illustrated in Figure 2, we introduce a quality stratifier that categorizes input source domain samples into high, medium, and low quality groups. It follows a quantile-based quality classification method [19]. Based on this stratification, the multi-layer extension module embeds the QSM module into the shallow, middle, and deep layers of the network, respectively, considering the complementary nature of hierarchical features to solve the layer-regardless feature augmentation issue. Specifically, QSM is selectively applied at different network stages based on image quality. Stage 1 is used for high-quality samples, stages 2–3 for medium-quality samples, and stage 4 for low-quality samples.

Existing methods suffer from layer-indiscriminate feature augmentation because they overlook the hierarchical complementarity of features in quality assessment [2]. Shallow layers are more sensitive to low-level distortions (e.g., slight blur or blocking artifacts) typical of high-quality samples, while deeper layers capture high-level semantics crucial for assessing severely degraded ones [42]. Consequently, the reliance on different feature layers varies with distortion severity and perceived quality: high-quality samples depend more on shallow features, whereas low-quality

samples rely on deeper representations. To overcome this, we design a quality-aware hierarchical augmentation mechanism that applies QSM according to the feature-level relevance to perceptual quality, rather than uniformly across layers.

3.2.3. Dual-Domain Augmentation Module

To address the augmentation imbalance, a dual-domain augmentation module is designed to enable feature augmentation in both domains, rather than only in the source domain. The multi-layer extension module is applied to the source domain, while SM is applied to target features after stage 4, as target quality scores are unavailable. The module enriches feature representations and promotes greater overlap between the source and target feature distributions, facilitating the learning of shared, domain-invariant representations and thus narrowing the domain gap. Moreover, it increases the difficulty of the domain discriminator, encouraging the feature extractor to learn more robust and invariant features. Since the target domain lacks labels, we perform uniform feature mixing at the final layer without mixing labels, aiming solely to enhance representation. The mixed features from both domains are finally fed into a DANN [7] through a Gradient Reversal Layer (GRL) for adversarial training.

We can optimize the discriminator D with a binary cross entropy loss \mathcal{L}_D , which is formulated as:

$$\mathcal{L}_D = -\frac{1}{n_s} \sum_{i=1}^{n_s} \log(1 - D(f_s^i)) - \frac{1}{n_t} \sum_{j=1}^{n_t} \log(D(f_t^j)), \quad (6)$$

where n_s and n_t denote the number of source domain and target domain samples, f_s^i and f_t^j are the feature representations from the source and target domains, and $D(\cdot)$ is the domain discriminator.

3.3. RCA Strategy

To solve the issue of quality-regardless feature alignment, we design a Rank-weighted Conditional Alignment (RCA) strategy, which includes a quality-aware conditional module and a rank-weighted module. The quality-aware conditional module is a feature alignment with quality-aware bias, which uses the true quality scores from the source domain and pseudo quality scores from the target domain as conditions to align features with a consistent quality level. The rank-weighted module assigns greater weights to misranked feature pairs to strengthen their alignment constraints, leading to more accurate correction of ranking bias.

Specifically, the RCA strategy is built on the statistical metric of COD [35], which measures local distribution differences in a fine-grained manner. However, COD [35] assigns equal weights to all sample pairs, ignoring the role of quality ranking in perceptual alignment. To address this issue, RCA introduces a rank-weighted module to highlight

sample pairs that exhibit ranking bias in cross-domain prediction. Compared to other sample pairs, these positions are more likely to reflect domain differences, thus increasing their weight prompts the model to prioritize correcting quality-perceptual alignment errors. The following loss function \mathcal{L}_R is designed to further optimize the feature extractor G formulated as:

$$\begin{aligned} \mathcal{L}_R = & \text{tr} \left(\mathbf{K}_Y^{tt} (\mathbf{K}_Y^{tt} + \epsilon \mathbf{I})^{-1} \tilde{\mathbf{K}}_X^{tt} (\mathbf{K}_Y^{tt} + \epsilon \mathbf{I})^{-1} \right) \\ & + \text{tr} \left(\mathbf{K}_Y^{ss} (\mathbf{K}_Y^{ss} + \epsilon \mathbf{I})^{-1} \tilde{\mathbf{K}}_X^{ss} (\mathbf{K}_Y^{ss} + \epsilon \mathbf{I})^{-1} \right) \\ & - 2 \text{tr} \left(\mathbf{K}_Y^{ts} (\mathbf{K}_Y^{ss} + \epsilon \mathbf{I})^{-1} \tilde{\mathbf{K}}_X^{st} (\mathbf{K}_Y^{tt} + \epsilon \mathbf{I})^{-1} \right), \end{aligned} \quad (7)$$

where $\text{tr}(\cdot)$ denotes the trace. \mathbf{K}_Y^{tt} and \mathbf{K}_Y^{ss} are the label kernel matrices for the target and source domains, and \mathbf{K}_Y^{ts} is the cross-domain label kernel matrix. $(\mathbf{K} + \epsilon \mathbf{I})^{-1}$ denotes the Tikhonov-regularized inverse [35] of \mathbf{K} for numerical stability, where \mathbf{I} is the identity matrix and $\epsilon > 0$ is a regularization constant. $\tilde{\mathbf{K}}_X^{tt}$, $\tilde{\mathbf{K}}_X^{ss}$, and $\tilde{\mathbf{K}}_X^{st}$ are the feature kernel matrices constructed with our proposed rank-weighted strategy. For example, $\tilde{\mathbf{K}}_X^{st}(i, j)$ denotes the rank-weighted feature kernel value between the i -th source sample and the j -th target sample, which is computed as follows:

$$\tilde{\mathbf{K}}_X^{st}(i, j) = k(f_s^i, f_t^j) \cdot (1 + \mathbf{W}^{st}(i, j)), \quad (8)$$

where $k(f_s^i, f_t^j)$ denotes the kernel function (*i.e.*, Gaussian kernel) to measure the similarity between source-domain feature f_s^i and target-domain feature f_t^j . $\mathbf{W}^{st}(i, j)$ denotes the rank-weight matrix element, to capture the contribution of the sample pair (i, j) , formulated as:

$$\mathbf{W}^{st}(i, j) = \max(0, -(\hat{y}_s^i - \hat{y}_t^j) \cdot \text{sign}(y_s^i - y_t^j)), \quad (9)$$

where \hat{y}_s^i denotes the predicted value for the i -th sample in the source domain; \hat{y}_t^j denotes the predicted value for the j -th sample in the target domain; y_s^i represents the true label of the i -th sample in the source domain; y_t^j represents the pseudo label of the j -th sample in the target domain; $\text{sign}(\cdot)$ denotes the sign function, which returns 1 for positive inputs, -1 for negative inputs, and 0 for zero.

Since early pseudo-labels are unreliable due to the unconstrained model, directly applying RCA may lead to incorrect alignment. Therefore, we propose a two-stage training strategy to introduce the reliable pseudo-label after the model is constrained and improve the robustness of RCA, as detailed in Section 3.5.

3.4. Quality Prediction Module

We employ a quality predictor P with two Fully-Connected layers to regress quality scores from the aligned features.

Given a set of distorted source images with features and quality labels, denoted as $\{x_s^i | f_s^i, y_s^i\}_{i=1}^{n_s}$, where $f_s^i =$

$G(x_s^i)$, the loss function \mathcal{L}_P of the prediction network is formulated as:

$$\mathcal{L}_P = \frac{1}{n_s} \sum_i^{n_s} (\hat{y}_s^i - y_s^i)^2, \quad (10)$$

where y_s^i denotes the quality labels of the distorted samples and $\hat{y}_s^i = P(f_s^i)$ denotes the predicted quality score of samples in the source domain.

Note that during training, to enable RCA, pseudo-labels for target domain samples are required. Therefore, target features $f_t^j = G(x_t^j)$ are also passed through the predictor P for quality estimation. At test time, the quality score of a target sample x_t^j is predicted in the same way as during pseudo-label y_t generation.

3.5. Training Strategy

After determining the network structure, we adopt a two-stage training strategy to progressively learn domain-invariant features and stabilize quality prediction. Due to the large domain gap, pseudo-labels are unreliable in the early stage. Thus, we first train a DANN-based model without pseudo-labels to achieve initial feature alignment. Once the model stabilizes, we introduce the pseudo-label-dependent RCA strategy to further refine cross-domain alignment and quality regression. As shown in Table 6 and Figure 3, this strategy yields more stable convergence and better performance. Notably, the RCA strategy is applied solely to the original source and target features to ensure accurate distribution alignment.

To enable stochastic application of the QSM and SM modules, we define two loss formulations:

The mixed-sample loss, denoted as \mathcal{L}_{all}^{mix} , is computed when QSM and SM are applied, formulated as:

$$\begin{aligned} \mathcal{L}_{all}^{mix} = & \lambda_1 \mathcal{L}_P(\hat{y}_s^{mix}, y_s^{mix}) + \lambda_2 \mathcal{L}_D(f_s^{mix}, f_t^{mix}) \\ & + \lambda_3 \mathcal{L}_R(y_s, y_t, f_s, f_t), \end{aligned} \quad (11)$$

The original-sample loss, denoted as \mathcal{L}_{all}^{orig} , is used when no mixing is applied, formulated as:

$$\begin{aligned} \mathcal{L}_{all}^{orig} = & \lambda_1 \mathcal{L}_P(\hat{y}_s, y_s) + \lambda_2 \mathcal{L}_D(f_s, f_t) \\ & + \lambda_3 \mathcal{L}_R(y_s, y_t, f_s, f_t) \end{aligned} \quad (12)$$

where λ_1 , λ_2 , and λ_3 denote loss weights; where we use \mathcal{L}_{all}^{mix} if $p > 0.5$, and \mathcal{L}_{all}^{orig} otherwise. p represents a draw from a uniform distribution over $[0, 1]$, indicating that QSM and SM are applied with a probability of 0.5.

4. Experiments

4.1. Experimental Setup

4.1.1. Dataset and Evaluation Metrics

We evaluate our method on four datasets, including TID2013 [27] and KADID-10k [15] for natural images, and

SJTU-PCQA [36] and WPC [29] for point clouds. We adopt Pearson’s Linear Correlation Coefficient (PLCC), Spearman Rank Order Correlation Coefficient (SROCC), Kendall Rank Order Correlation Coefficient (KROCC), and Root Mean Square Error (RMSE) as metrics. Better performance is indicated by higher PLCC, SROCC, and KROCC, and lower RMSE.

4.1.2. Implementation Details

All experiments are implemented in PyTorch on an NVIDIA RTX 4090 GPU. Natural image datasets serve as source domains, while point cloud datasets are split into 75% for training and 25% for testing as target domains. We adopt a modified ResNet-50 [10] backbone pretrained on ImageNet [6]. The domain discriminator D consists of two Fully-Connected(FC) layers with ReLU activations and one FC layer with Sigmoid. The predictor uses two FC-ReLU layers followed by a Sigmoid function. Input images are randomly cropped and flipped during training, and center-cropped during testing. We use the SGD optimizer with momentum 0.9, weight decay 5×10^{-4} , a batch size of 36, and 30,000 iterations. DANN [7] is applied during the first 5,000 iterations (warm-up), followed by joint optimization with RCA for the remaining iterations. For QSM, α is determined via ablation studies, set to 1 for TID2013 [27] and 0.5 for KADID-10k [15] to augment style diversity. Finally, we empirically set $\lambda_1 = \lambda_2 = \lambda_3$ to balance the three loss terms, the Gaussian kernel bandwidth to $\tau = 5 \times 10^{-2}$ for stable optimization, and the Tikhonov regularization [35] parameter to $\epsilon = 10^{-3}$. Samples are divided into three quality levels (low, medium, high) based on the 33% and 67% quantiles of subjective scores.

4.2. Performance Evaluation

Because images and point cloud projections share the same 2D format, image-to-image IQA transfer methods can also be applied to image-to-point cloud transfer. Therefore, we compare the proposed QD-PCQA with three categories of reproduced baselines: image-to-image transfer IQA methods (e.g., StyleAM [22], Chen’s [1]), general UDA methods adapted for image-to-point cloud transfer (e.g., DANN [7], COD [35]), and image-to-point cloud transfer methods such as IT-PCQA [38]. “No Adapt” denotes our reproduced baseline using ResNet-50 [10] as the feature extractor, followed by a quality predictor.

As shown in Table 1, we compare QD-PCQA with representative baselines trained on TID2013 [27] and evaluated on SJTU-PCQA [36] and WPC [29]. QD-PCQA consistently achieves state-of-the-art performance. For example, in the TID2013 [27] \rightarrow SJTU-PCQA [36] setting, it achieves a PLCC of 0.842 and an RMSE of 1.358, surpassing IT-PCQA [38] by 21.5% in PLCC and reducing RMSE by 16.4%. On TID2013 [27] \rightarrow WPC [29], it reaches a PLCC of 0.563, outperforming DANN [7] by 73.2% and

Table 1. Performance of training on TID2013 [27] and testing on SJTU-PCQA [36] or WPC [29].

Methods	Model	TID2013 [27] → SJTU-PCQA [36]				TID2013 [27] → WPC [29]			
		PLCC↑	SROCC↑	KROCC↑	RMSE↓	PLCC↑	SROCC↑	KROCC↑	RMSE↓
StyleAM [22]	I-to-I	0.612	0.519	0.352	1.883	0.352	0.299	0.265	21.323
Chen’s [1]	I-to-I	0.706	0.632	0.453	1.667	0.386	0.302	0.293	21.102
DANN [7]	I-to-PC	0.596	0.512	0.489	1.965	0.325	0.296	0.245	21.364
COD [35]	I-to-PC	0.712	0.611	0.492	1.735	0.426	0.396	0.320	20.994
No Adapt	I-to-PC	0.548	0.444	0.321	1.995	0.320	0.296	0.221	21.423
IT-PCQA [38]	I-to-PC	0.693	0.636	0.493	1.624	0.429	0.403	0.323	20.867
QD-PCQA (Ours)	I-to-PC	0.842	0.753	0.566	1.358	0.563	0.572	0.462	20.835

Table 2. Performance of training on KADID-10k [15] and testing on SJTU-PCQA [36] or WPC [29].

Methods	Model	KADID-10k [15] → SJTU-PCQA [36]				KADID-10k [15] → WPC [29]			
		PLCC↑	SROCC↑	KROCC↑	RMSE↓	PLCC↑	SROCC↑	KROCC↑	RMSE↓
StyleAM [22]	I-to-I	0.698	0.589	0.450	1.712	0.378	0.320	0.267	21.303
Chen’s [1]	I-to-I	0.643	0.568	0.406	1.824	0.402	0.312	0.303	20.991
DANN [7]	I-to-PC	0.602	0.534	0.446	1.812	0.324	0.284	0.256	21.385
COD [35]	I-to-PC	0.721	0.612	0.496	1.703	0.423	0.395	0.315	20.965
No Adapt	I-to-PC	0.552	0.462	0.392	1.965	0.321	0.297	0.254	21.403
IT-PCQA [38]	I-to-PC	0.703	0.641	0.512	1.606	0.432	0.402	0.386	20.945
QD-PCQA (Ours)	I-to-PC	0.843	0.724	0.528	1.305	0.553	0.534	0.468	20.794




Figure 3. t-SNE visualization of the aligned representations learned from the source and target domains. Label values are denoted by color gradients. Ten label values are selected from the range of the variable for visualization.

IT-PCQA [38] by 31.2%, demonstrating strong robustness under challenging cross-domain conditions. Similarly, Table 2 shows results for baselines trained on KADID-10k [15]. QD-PCQA consistently leads across evaluations. For KADID-10k [15] → SJTU-PCQA [36], it achieves a PLCC of 0.843 and RMSE of 1.305, surpassing IT-PCQA [38] by

19.9% in PLCC and reducing RMSE by 18.7%.

These results align with the intrinsic differences in dataset distortions. SJTU-PCQA [36] primarily contains low-level texture and color degradations, often as single-type or binary distortions (e.g., color noise, downsampling), which are similar to pixel-level degradations in image datasets such as TID2013 [27]. This similarity facilitates cross-domain transfer, explaining the generally higher performance on SJTU-PCQA [36]. In contrast, WPC [29] exhibits high-level semantic distortions caused by G-PCC compression, which differ significantly from image-domain degradations and pose greater challenges for cross-domain alignment. To address these challenges, our QFA strategy effectively enriches the source domain with diverse samples, facilitating shared feature learning and narrowing the domain gap. Meanwhile, the multi-layer extension module enhances texture details for high-quality samples and strengthens semantic robustness for low-quality ones. As shown in Figure 3, QD-PCQA aligns source and target domain features effectively, confirming its generalization.

4.3. Ablation Study

This section assesses the effectiveness of each component using TID2013 [27] as the source and SJTU-PCQA [36] as the target domain.

4.3.1. QSM Module

We compare three style mixing strategies: No Style Mixup, Style Mixup, and our proposed QSM. As shown in Table 3, introducing style mixing improves cross-domain generalization, while QSM further boosts performance by guiding the mixing according to quality similarity, ensuring perceptual consistency in the augmented features. Moreover, the ablation on the Beta parameter α in Table 4 shows that $\alpha = 1$ achieves the best results, highlighting that the mixing strength directly affects performance by balancing style diversification and quality structure preservation.

Table 3. Ablation study on our QSM module.

Methods	PLCC \uparrow	SROCC \uparrow	KROCC \uparrow	RMSE \downarrow
No Style Mixup	0.792	0.676	0.484	1.543
Style Mixup	0.820	0.711	0.507	1.392
QSM	0.842	0.753	0.566	1.358

Table 4. Ablation study on the Beta distribution parameter α , which controls the concentration of the mixing coefficient λ .

α	PLCC \uparrow	SROCC \uparrow	KROCC \uparrow	RMSE \downarrow
0.1	0.834	0.732	0.512	1.374
0.2	0.838	0.749	0.532	1.364
0.5	0.840	0.752	0.556	1.361
1.0	0.842	0.753	0.566	1.358

4.3.2. Multi-Layer Extension Module

We compare Style Mixup (SM) applied at different network stages and our proposed multi-layer extension. As shown in Table 5, the multi-layer extension consistently performs best, indicating that leveraging hierarchical features in a complementary way is more effective for PCQA.

Table 5. Ablation study on SM insertion in different layers.

Methods	PLCC \uparrow	SROCC \uparrow	KROCC \uparrow	RMSE \downarrow
@stage1	0.818	0.698	0.493	1.412
@stage4	0.815	0.701	0.496	1.381
@stage2,3	0.820	0.711	0.507	1.364
@Ours	0.842	0.753	0.566	1.358

4.3.3. RCA Strategy

We evaluate three distribution alignment methods: MMD [28], COD [35], and our Rank-weighted Conditional Alignment (RCA). As shown in Table 6, COD, which leverages quality priors for conditional alignment, outperforms the

marginal MMD baseline, highlighting the benefit of quality-aware alignment based on scores. Compared with COD, our RCA further improves performance, boosting SROCC by over 5.9%, demonstrating the effectiveness of incorporating misranking-aware weighting.

Table 6. Ablation study on our RCA strategy.

Methods	PLCC \uparrow	SROCC \uparrow	KROCC \uparrow	RMSE \downarrow
MMD	0.803	0.695	0.507	1.432
COD	0.820	0.711	0.517	1.396
RCA	0.842	0.753	0.566	1.358

4.3.4. Dual-Domain Augmentation Module

This section evaluates the generalization of the dual-domain augmentation module. As shown in Table 7, DA(W) and DA(S) denote the module’s performance on the WPC [29] and SJTU-PCQA [36] datasets, respectively. Although a performance drop is observed on the unseen WPC [29] dataset, this is expected because WPC [29] exhibits distortion types and perceptual characteristics that differ markedly from other datasets, resulting in a larger domain gap. Nevertheless, the dual-domain approach still surpasses single-domain mixing, showing stronger robustness under severe cross-dataset shifts.

Table 7. Ablation study on dual-domain augmentation module.

Methods	PLCC \uparrow	SROCC \uparrow	KROCC \uparrow	RMSE \downarrow
W/o DA(W)	0.257	0.281	0.189	21.961
Ours(W)	0.303	0.296	0.201	21.959
W/o DA(S)	0.834	0.726	0.543	1.364
Ours(S)	0.842	0.753	0.566	1.358

5. Conclusion

In this paper, we propose a novel domain adaptation quality assessment framework named QD-PCQA, which leverages prior quality knowledge from images to predict point cloud quality. This framework consists of two key components: a Rank-weighted Conditional Alignment (RCA) strategy that adaptively aligns features based on quality-level consistency, and a Quality-guided Feature Augmentation (QFA) strategy that performs hierarchical, dual-domain, quality-aware style mixup. These components jointly preserve perceptual quality sensitivity and enhance generalization under domain shifts. Overall, QD-PCQA provides new insights into intrinsic correlations between different media and consistently delivers robust, competitive results.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (62372036, 62120106009, U22A2022, 62332017), and the Ministry of Education, Singapore (Tier 1 RG103/24).

References

- [1] Baoliang Chen, Haoliang Li, Hongfei Fan, and Shiqi Wang. No-reference screen content image quality assessment with unsupervised domain adaptation. *IEEE Transactions on Image Processing*, 2020. 2, 3, 6, 7
- [2] Chaofeng Chen, Jiadi Mo, Jingwen Hou, Haoning Wu, Liang Liao, Wenxiu Sun, Qiong Yan, and Weisi Lin. TopIQ: A top-down approach from semantics to distortions for image quality assessment. *IEEE Transactions on Image Processing*, 33:2404–2418, 2024. 4
- [3] Siheng Chen, Baoan Liu, Chen Feng, Carlos Vallespi-Gonzalez, and Carl Wellington. 3D point cloud processing and learning for autonomous driving: Impacting map creation, localization, and perception. In *Proceedings of the IEEE Signal Processing Magazine*, 38(1):68–86, 2020. 1
- [4] Tianang Chen, Jian Jin, Shilv Cai, Zhuangzi Li, and Weisi Lin. MUGSQA: Novel multi-uncertainty-based gaussian splatting quality assessment method, dataset, and benchmarks. *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, 2025. 1
- [5] Aladine Chetouani, Maurice Quach, Giuseppe Valenzise, and Frédéric Dufaux. Deep learning-based quality assessment of 3D point clouds without reference. In *Proceedings of the IEEE International Conference on Multimedia & Expo Workshops*, pages 1–6, 2021. 3
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 6
- [7] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario March, and Victor Lempitsky. Domain-adversarial training of neural networks. *Journal of machine learning research*, 17(59):1–35, 2016. 2, 3, 5, 6, 7
- [8] Wei Gao, Shangkun Sun, Huiming Zheng, Yuyang Wu, Hua Ye, and Yongchi Zhang. OpenDMC: An open-source library and performance evaluation for deep-learning-based multi-frame compression. In *Proceedings of the ACM International Conference on Multimedia*, pages 9685–9688, 2023. 1
- [9] Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bennamoun. Deep learning for 3D point clouds: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(12):4338–4364, 2020. 1
- [10] 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. 3, 6
- [11] Zhouyan He, Gangyi Jiang, Zhidi Jiang, and Mei Yu. Towards a colored point cloud quality assessment method using colored texture and curvature projection. In *Proceedings of the IEEE International Conference on Image Processing*, pages 1444–1448, 2021. 3
- [12] Janguang Jiang, Yang Shu, Jianmin Wang, and Mingsheng Long. Transferability in deep learning: A survey. *arXiv preprint arXiv:2201.05867*, 2022. 2
- [13] Shuo Jin, Meiqin Liu, Chao Yao, Chunyu Lin, and Yao Zhao. Kernel dimension matters: To activate available kernels for real-time video super-resolution. In *Proceedings of the ACM International Conference on Multimedia*, pages 8617–8625, 2023. 2
- [14] Shuo Jin, Siyue Yu, Bingfeng Zhang, Mingjie Sun, Yi Dong, and Jimin Xiao. Feature purification matters: Suppressing outlier propagation for training-free open-vocabulary semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20291–20300, 2025. 2
- [15] Hanhe Lin, Vlad Hosu, and Dietmar Saupe. KADID-10k: A large-scale artificially distorted iqa database. In *Proceedings of the International Conference on Quality of Multimedia Experience*, pages 1–3, 2019. 6, 7
- [16] Meiqin Liu, Shuo Jin, Chao Yao, Chunyu Lin, and Yao Zhao. Temporal consistency learning of inter-frames for video super-resolution. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(4):1507–1520, 2022. 2
- [17] Meiqin Liu, Chenming Xu, Chao Yao, Chunyu Lin, and Yao Zhao. JNMR: Joint non-linear motion regression for video frame interpolation. *IEEE Transactions on Image Processing*, 32:5283–5295, 2023. 1
- [18] 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. 3
- [19] Xialei Liu, Joost Van De Weijer, and Andrew D Bagdanov. RankIQA: Learning from rankings for no-reference image quality assessment. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1040–1049, 2017. 4
- [20] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *Proceedings of International Conference on Machine Learning*, pages 2208–2217, 2017. 3
- [21] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. *Advances in Neural Information Processing Systems*, 31, 2018. 3
- [22] Yiting Lu, Xin Li, Jianzhao Liu, and Zhibo Chen. StyleAM: Perception-oriented unsupervised domain adaptation for no-reference image quality assessment. *IEEE Transactions on Multimedia*, 27:2043–2058, 2024. 2, 3, 6, 7
- [23] Rufael Mekuria, Kees Blom, and Pablo Cesar. Design, implementation, and evaluation of a point cloud codec for tele-immersive video. *IEEE Transactions on Circuits and Systems for Video Technology*, 27(4):828–842, 2016. 1

- [24] Rufael Mekuria, Zhu Li, Christian Tulvan, and Phil Chou. Evaluation criteria for point cloud compression. *ISO/IEC MPEG*, w16332, Geneva, Switzerland, 2016., 2016. 3
- [25] Lili Meng, Jian Jin, Yingnan Wang, Miaohui Wang, Guosheng Lin, Cheng Liang, Jiande Sun, Huaxiang Zhang, and Weisi Lin. Boosting the no-reference image quality assessment via low-quality pseudo references. *IEEE Transactions on Circuits and Systems for Video Technology*, pages 1–14, 2026. 1
- [26] Gabriel Meynet, Yana Nehmé, Julie Digne, and Guillaume Lavoué. PCQM: A full-reference quality metric for colored 3D point clouds. In *Proceedings of the International Conference on Quality of Multimedia Experience*, pages 1–6, 2020. 3
- [27] Nikolay Ponomarenko, Lina Jin, Oleg Ieremeiev, Vladimir Lukin, Karen Egiazarian, Jaakko Astola, Benoit Vozel, Kacem Chehdi, Marco Carli, Federica Battisti, et al. Image database TID2013: Peculiarities, results and perspectives. *Signal Processing: Image Communication*, 30:57–77, 2015. 6, 7
- [28] Artem Rozantsev, Mathieu Salzmann, and Pascal Fua. Beyond sharing weights for deep domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(4):801–814, 2018. 2, 3, 8
- [29] Honglei Su, Zhengfang Duanmu, Wentao Liu, Qi Liu, and Zhou Wang. Perceptual quality assessment of 3D point clouds. In *Proceedings of the IEEE International Conference on Image Processing*, pages 3182–3186, 2019. 6, 7, 8
- [30] Qi Tang, Yao Zhao, Meiqin Liu, and Chao Yao. SeeClear: Semantic distillation enhances pixel condensation for video super-resolution. *Advances in Neural Information Processing Systems*, 37:134902–134926, 2024. 1
- [31] Qi Tang, Yao Zhao, Meiqin Liu, and Chao Yao. A review of video super-resolution algorithms based on deep learning. *Acta Automatica Sinica*, 51(7):1480–1524, 2025. 1
- [32] Dong Tian, Hideaki Ochimizu, Chen Feng, Robert Cohen, and Anthony Vetro. Geometric distortion metrics for point cloud compression. In *Proceedings of the IEEE International Conference on Image Processing*, pages 3460–3464, 2017. 3
- [33] 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. 3
- [34] Qingsong Xu, Yilei Shi, Xin Yuan, and Xiao Xiang Zhu. Universal domain adaptation for remote sensing image scene classification. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–15, 2023. 2
- [35] Hao Ran Yang, Chuan Xian Ren, and You Wei Luo. COD: Learning conditional invariant representation for domain adaptation regression. In *Proceedings of European Conference on Computer Vision*, pages 108–125, 2024. 3, 5, 6, 7, 8
- [36] 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. 6, 7, 8
- [37] 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. 3
- [38] 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, 6, 7
- [39] Renyu Yang, Jian Jin, Lili Meng, Meiqin Liu, Yilin Wang, Balu Adsumilli, and Weisi Lin. Scaling audio-visual quality assessment dataset via crowdsourcing. *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, 2026. 1
- [40] Huaxiu Yao, Yiping Wang, Linjun Zhang, James Y Zou, and Chelsea Finn. C-Mixup: Improving generalization in regression. In *Proceedings of Advances in Neural Information Processing Systems*, pages 3361–3376, 2022. 3
- [41] Guohua Zhang, Jian Jin, Meiqin Liu, Chao Yao, Weisi Lin, and Yao Zhao. GT-PCQA: geometry-texture decoupled point cloud quality assessment with mllm. *arXiv preprint arXiv:2603.14951*, 2026. 1
- [42] Weiyu Zhou, Yongqing Yang, Tao Hu, Pu Hui, Jian Jin, Yu Cao, Qingsen Yan, and Yanning Zhang. High dynamic range imaging via spatial-frequency interaction. *IEEE Transactions on Circuits and Systems for Video Technology*, pages 1–12, 2026. 4