

# Fine-Grained Image Quality Assessment for Perceptual Image Restoration

Xiangfei Sheng<sup>\*1</sup>, Xiaofeng Pan<sup>\*1</sup>, Zhichao Yang<sup>1</sup>, Pengfei Chen<sup>1</sup>, Leida Li<sup>1,2†</sup>

<sup>1</sup>School of Artificial Intelligence, Xidian University

<sup>2</sup>State Key Lab. of Electromechanical Integrated Manufacturing of High-Performance Electronic Equipments, Xidian University

xiangfeisheng@gmail.com, {panxf@stu., yangzhichao@stu., chenpengfei@, ldli@}xidian.edu.cn

## Abstract

Recent years have witnessed remarkable achievements in perceptual image restoration (IR), creating an urgent demand for accurate image quality assessment (IQA), which is essential for both performance comparison and algorithm optimization. Unfortunately, the existing IQA metrics exhibit inherent weakness for IR task, particularly when distinguishing fine-grained quality differences among restored images. To address this dilemma, we contribute the first-of-its-kind fine-grained image quality assessment dataset for image restoration, termed **FGRestore**, comprising 18,408 restored images across six common IR tasks. Beyond conventional scalar quality scores, FGRestore was also annotated with 30,886 fine-grained pairwise preferences. Based on FGRestore, a comprehensive benchmark was conducted on the existing IQA metrics, which reveal significant inconsistencies between score-based IQA evaluations and the fine-grained restoration quality. Motivated by these findings, we further propose **FGResQ**, a new IQA model specifically designed for image restoration, which features both coarse-grained score regression and fine-grained quality ranking. Extensive experiments and comparisons demonstrate that FGResQ significantly outperforms state-of-the-art IQA metrics.

**Homepage** — <https://sxfly99.github.io/FGResQ-Home/>

## Introduction

Perceptual image restoration (IR) stands as a cornerstone in low-level computer vision, aiming to recover high-quality images from their degraded observations while maintaining perceptual fidelity (Potlapalli et al. 2023; Luo et al. 2023). Recent years have witnessed remarkable breakthroughs in this field, largely driven by the evolution of generative models. These advances have enabled IR algorithms to achieve unprecedented visual quality, creating an urgent demand for accurate image quality assessment (IQA) methods that can reliably evaluate and compare restored images. Such assessment capabilities are essential not only for performance benchmarking across different IR algorithms but also for guiding algorithmic optimization.

<sup>\*</sup>These authors contributed equally.

<sup>†</sup>Corresponding author

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

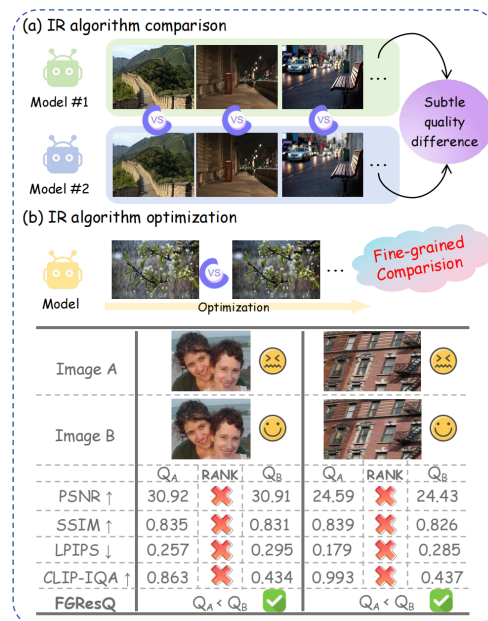


Figure 1: Illustration of the **fine-grained challenge in IQA for image restoration**. Both IR algorithm comparison and optimization processes require distinguishing subtle quality differences between restored images. Existing IQA metrics fail to provide correct rankings for fine-grained image pairs. (Best viewed zoomed in.)

Currently, IR algorithm evaluation still predominantly relies on traditional reference-based metrics such as PSNR and SSIM (Wang et al. 2004), which assess restoration quality by measuring the similarity between the restored images and ground-truth images, even though their inconsistency with human perceptual judgment has been increasingly observed by the research community (Jinjin et al. 2020). Recent efforts have introduced no-reference IQA methods (Xu et al. 2024; Zhou et al. 2025), including CLIP-IQA (Wang, Chan, and Loy 2023) and other learning-based approaches (Ke et al. 2021), which prove particularly valuable in real-world restoration scenarios where ground truth images are unavailable. However, a critical oversight persists: these methods have been primarily designed for image quality assessment, without specific consideration for the fine-grained evalua-

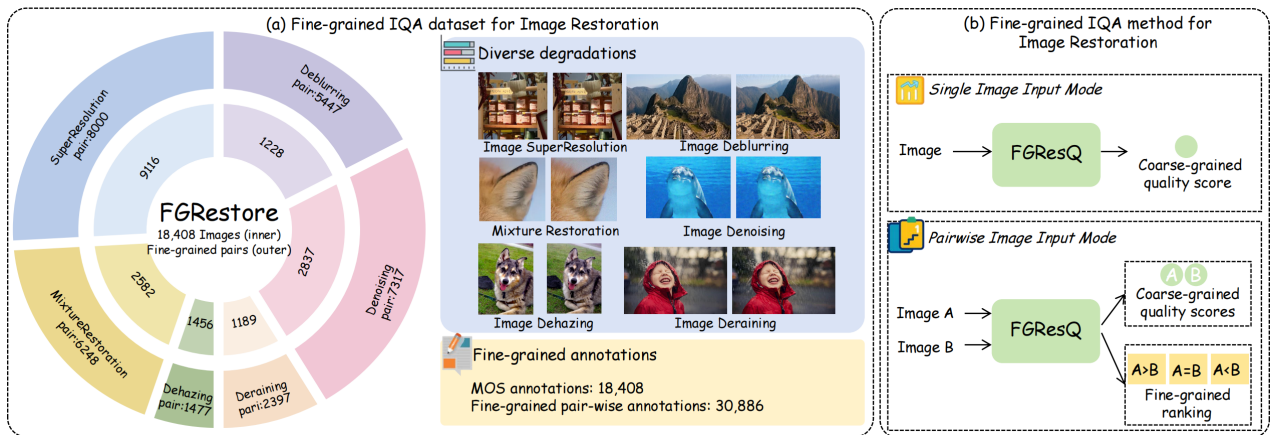


Figure 2: Overview of our method. (a) **FGRestore** provides comprehensive fine-grained quality annotations for multiple IR tasks. (b) **FGResQ** enables both coarse-grained quality scoring and fine-grained quality ranking capabilities.

tion requirements inherent in IR tasks.

A critical yet underexplored challenge in IR evaluation lies in its inherently **fine-grained nature**. As illustrated in Figure 1, IR algorithm comparison and optimization typically involve distinguishing between images with subtle quality differences—scenarios where even state-of-the-art IQA methods struggle to provide reliable assessments. This observation leads us to fundamentally rethink existing evaluation paradigms for IR tasks and raises a crucial question:

*Can existing score-based IQA methods objectively capture fine-grained quality differences in restored images?*

To investigate this question, we first conducted comprehensive computational analyses on established IQA datasets for image restoration (Jinjin et al. 2020; Shi et al. 2019). Our findings consistently reveal a significant performance gap: while existing IQA methods achieve reasonable results in coarse-grained quality assessment, they consistently underperform in fine-grained scenarios, with significant inconsistencies between predicted scores and human perceptual judgments. This observation highlights the inadequacy of current evaluation protocols for capturing the nuanced quality differences which are crucial in IR applications. Moreover, when quality differences between images become subtle, both human and IQA models encounter significant challenges in providing reliable absolute quality scores. In contrast, pairwise comparison (Yang et al. 2024b,a) presents a more reliable alternative for fine-grained assessment, as humans demonstrate superior consistency in relative quality judgments compared to absolute scoring.

Motivated by the above insights, we introduce **FGRestore**, the first fine-grained IQA dataset specifically for image restoration tasks, as shown in Figure 2(a). Our dataset comprises 18,408 restored images spanning six common IR tasks: image super-resolution, deblurring, denoising, dehazing, deraining, and mixture restoration. Beyond conventional Mean Opinion Score (MOS) annotations, **FGRestore** incorporates 30,886 fine-grained preference annotations for restored image pairs with subtle quality differences. Based on **FGRestore**, we further propose **FGResQ**, a unified fine-grained IQA model tailored for IR evaluation (Figure 2(b)).

Our approach employs degradation-aware feature learning that incorporates human-level degradation knowledge from Vision-Language Models (VLMs) into the proposed IQA framework, enabling both accurate quality scoring and fine-grained quality ranking across diverse IR tasks.

Our main contributions are threefold:

- We construct **FGRestore**, the first fine-grained IQA dataset specifically designed for IR tasks, containing 18,408 restored images across six common IR categories with both MOS and 30,886 fine-grained pairwise preference annotations.
- We propose **FGResQ**, a unified fine-grained IQA model that incorporates degradation-aware feature learning and demonstrates state-of-the-art performance in both coarse-grained score regression and fine-grained ranking.
- Through comprehensive computational analysis on existing IR datasets and extensive experiments, we reveal fundamental limitations of current score-based IQA methods in capturing fine-grained quality differences, providing valuable insights for future IR evaluation research.

## Related Work

### Image Restoration Quality Assessment Dataset

To evaluate IR algorithms, numerous datasets have been proposed. These include datasets dedicated to single restoration types, such as the MDD13 (Liu et al. 2013) dataset for deblurring, the IVC-Dehazing (Ma, Liu, and Wang 2015) dataset for dehazing, the IVIPC-DQA (Wu et al. 2019) dataset for deraining, and the QADS (Zhou et al. 2019) and SRIQA-Bench (Chen et al. 2025) datasets for super-resolution. Additionally, there are datasets that encompass distortions from various algorithms, such as PIPAL (Jinjin et al. 2020), which includes distortion types like denoising, super-resolution, and super-resolution with post-denoising, KADID-10k (Lin, Hosu, and Saupé 2019), which contains various image distortions including denoising degradations and DiffQA (Chen et al. 2025), which accounts for degradation artifacts induced by diffusion-based image enhance-

Dataset	Year	Task	Alg.	Annotation Type	Number (Img/Pairs/Ann.)
MDD13	2013	1	5	MOS (PC)	1,200/0/13,592
IVC-Dehazing	2015	1	8	MOS (ACR)	200/0/4,800
QADS	2019	1	21	MOS (PC)	1,260/0/126,000
SISRSet	2019	1	8	MOS (PC)	260/0/32,000
IVIPC-DQA	2019	1	6	MOS (5-level)	2,136/0/27,192
KADID-10k	2019	1	1	MOS(DCR)	10,125/0/303,750
exBeDDE	2020	1	10	MOS (PC)	1,670/0/18,380
PIPAL	2020	3	40	MOS (Elo)	29,000/0/1.13M
RealSRQ	2022	1	10	MOS (PC)	1,620/0/65,400
DiffQA	2025	1	1	RANK	177,319/177,319/537,624
SRIQA-Bench	2025	1	10	RANK	1100/5500/55000
FGRestore	2025	6	108	MOS+RANK	18,408/30,886/45,318

Table 1: Comparison with the previous datasets. MOS: Mean Opinion Score; PC: pairwise comparison; ACR: Absolute Category Rating; DCR: Degradation Category Rating; 5-level: 5-point quality scale. Number show Images/Pairs/Annotations.

ment. A brief comparison of FGRestore and previous IR evaluation datasets is illustrated in Table 1.

### Image Quality Assessment

Existing image restoration (IR) algorithms typically employ full-reference image quality assessment (FR-IQA) metrics for performance evaluation, such as the conventional PSNR and SSIM (Wang et al. 2004), as well as deep learning-based metrics like LPIPS (Zhang et al. 2018a) and DISTS (Ding et al. 2020). To overcome the flawed "perfect reference" assumption in traditional FR-IQA, which is challenged by imaging system limitations and superior generative methods, A-FINE (Chen et al. 2025) proposes a generalized model that adaptively assesses both the fidelity and naturalness of test images.

No-reference IQA (NR-IQA) methods have been proposed to assess image quality without reference images. Early NR-IQA methods, such as NIQE(Mittal, Soundararajan, and Bovik 2012), primarily rely on hand-crafted Natural Scene Statistics (NSS) features. With the advancement of deep learning, many works have turned to neural networks for NR-IQA. For instance, MetalQA(Zhu et al. 2020)leverages meta-learning to enhance model generalization. Additionally, methods such as CLIP-IQA(Wang, Chan, and Loy 2023)leverage text prompts to enhance perceptual capabilities(Sheng et al. 2023; Yang et al. 2025; Li et al. 2025). Given the powerful learning capabilities of large multimodality models (LMMs)(Sheng et al. 2025a,b), emerging work, such as Q-Align(Wu et al. 2023), has begun to apply these models to quality assessment. Most recently, methods represented by Co-Instruct(Wu et al. 2024)have extended LMMs to compare multiple images. Furthermore, VisualQuality-R1(Wu et al. 2025) introduces a novel methodology for IQA that leverages LMMs within a reinforcement learning-to-rank framework.

Although these datasets and methods have contributed significantly to the IQA community, most methods are score-based and cannot effectively handle fine-grained perceptual differences brought by image restoration algorithms. Recent works have explored fine-grained quality assess-

Type	Method	[0.0,0.2] SRCC	[0.2,0.4] SRCC	[0.4,0.6] SRCC	[0.6,0.8] SRCC	[0.8,1.0] SRCC	Overall SRCC
FR	PSNR	0.323	0.082	0.209	0.161	0.072	0.422
	SSIM	0.293	0.108	0.258	0.254	0.049	0.530
	LPIPS	-0.034	0.077	0.325	0.287	0.124	0.612
	DISTS	0.168	0.159	0.310	0.242	0.165	0.585
NR	NIQE	-0.126	-0.002	0.107	0.001	0.080	0.153
	IL-NIQE	-0.235	-0.098	0.126	0.128	0.054	0.289
	BRISQUE	-0.142	0.025	0.125	0.035	0.131	0.185
	DB-CNN	-0.157	0.321	0.353	0.330	-0.016	0.636
	HyperIQA	0.100	0.274	0.314	0.292	0.032	0.584
	MetalQA	0.037	0.160	0.204	0.174	-0.101	0.423
	LIQE	-0.232	0.053	0.175	0.299	0.107	0.479
	CLIP-IQA	-0.152	0.211	0.238	0.293	0.071	0.530
	Q-Align	0.230	0.301	0.337	0.213	0.178	0.418
	DeQA-Score	0.568	0.676	0.623	0.516	0.350	0.747

Table 2: Performance comparison across different MOS ranges on PIPAL dataset.

ment in specific domains. Zhang et al. (Zhang et al. 2018b; Zhang, Lin, and Huang 2021) introduced fine-grained assessment for compressed images, focusing on distinguishing compression artifacts within similar quality levels. However, these prior works primarily target image compression, whereas our work specifically addresses the unique challenges of fine-grained quality assessment in image restoration tasks, where the evaluation focus shifts from compression artifacts to perceptual restoration quality differences.

### Preliminary Validation Analysis

In image restoration tasks, both algorithm comparison and optimization frequently involve evaluating images with subtle quality differences. Algorithm comparison requires distinguishing between restoration results with marginal quality differences, while parameter optimization involves incremental quality changes that demand sensitive assessment methods to identify optimal configurations. To investigate whether existing IQA methods can objectively capture fine-grained quality differences in IR task, we conduct a comprehensive computational analysis on established IQA datasets. Specifically, we evaluate state-of-the-art IQA methods to assess their fine-grained discrimination capabilities. Our evaluation encompasses both full-reference (FR) and no-reference (NR) methods. FR methods include PSNR, SSIM (Wang et al. 2004), LPIPS (Zhang et al. 2018a), and DISTS (Ding et al. 2020). NR methods include traditional approaches such as NIQE (Mittal, Soundararajan, and Bovik 2012), IL-NIQE (Zhang, Zhang, and Bovik 2015), and BRISQUE (Mittal, Moorthy, and Bovik 2012), as well as deep learning-based methods including DB-CNN (Zhang et al. 2020), HyperIQA (Su et al. 2020), MetalQA (Zhu et al. 2020), LIQE (Zhang et al. 2023), CLIP-IQA (Wang, Chan, and Loy 2023), Q-Align (Wu et al. 2023), and DeQA-Score (You et al. 2025).

Specifically, we partition the quality score range into several intervals to examine how IQA methods perform when evaluating images with similar quality levels, which is particularly relevant for fine-grained assessment scenarios in image restoration. Table 2 presents the performance comparison across different MOS ranges on the PIPAL dataset (Jinjin et al. 2020). While most IQA methods achieve reasonable overall SRCC values on the whole dataset, their performance dramatically deteriorates when evaluated within

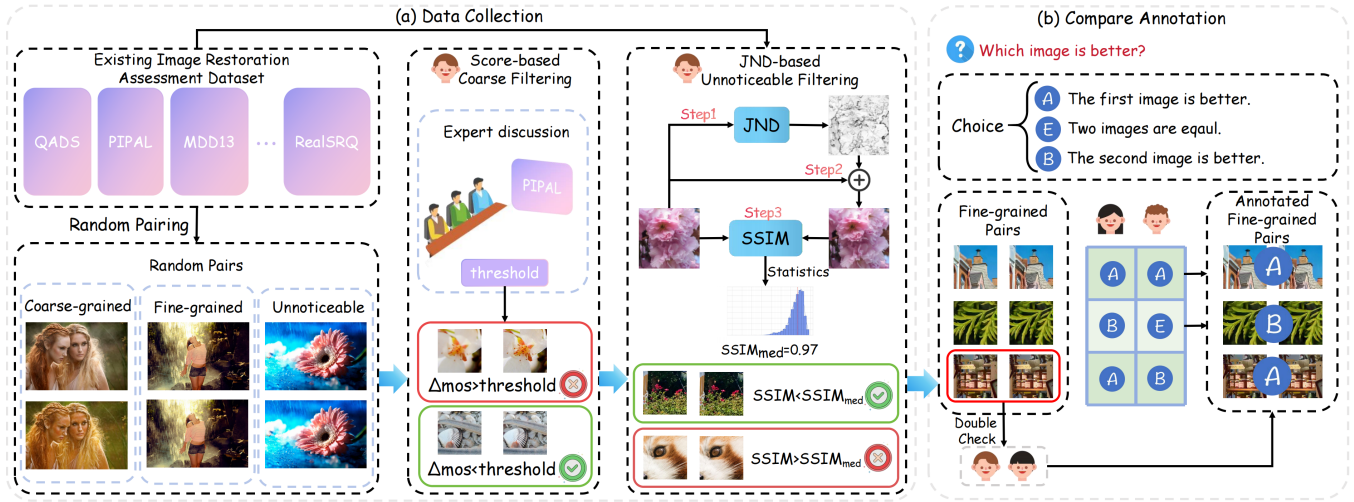


Figure 3: Overview of the FGRestore dataset construction.

narrow quality ranges. We also conducted similar analyses on other IQA datasets, which consistently demonstrate the same pattern. More results are provided in the *supplementary material*. These findings provide strong evidence for the inadequacy of current score-based IQA approaches in fine-grained scenarios, motivating the development of our fine-grained evaluation framework for image restoration.

## FGRestore Dataset

Motivated by the findings revealed in our preliminary analysis, we construct FGRestore, the first comprehensive fine-grained image quality assessment dataset for perceptual image restoration evaluation. The dataset construction pipeline is illustrated in Figure 3, which involves collecting images from multiple datasets, filtering image pairs through score-based and JND-based criteria to retain fine-grained pairs, and conducting systematic subjective studies for pairwise preference annotation.

### Image and Pair Collection

We collect images from multiple restoration-specific datasets to cover different degradation and visual appearances: 1) deraining images from IVIPC-DQA (Wu et al. 2019); 2) deblurring images from MDD13 (Liu et al. 2013); 3) dehazing images from IVC-Dehazing (Ma, Liu, and Wang 2015) and exBeDDE (Zhao et al. 2020); 4) denoising images from PIPAL (Jinjin et al. 2020) and KADID-10k (Lin, Hosu, and Saupé 2019); 5) super-resolution images from QADS (Zhou et al. 2019), RealSRQ (Jiang et al. 2022), SISRSset (Shi et al. 2019), and PIPAL (Jinjin et al. 2020); 6) mixture restoration images from PIPAL (Jinjin et al. 2020). After image collection, we generate image pairs by randomly pairing all images with identical content and restoration tasks, resulting in a total of 1,275,297 image pairs.

### Data Filtration

The randomly generated image pairs can be categorized into three distinct types based on perceptual quality differences:

(1) *Coarse-grained pairs* with highly noticeable quality differences, (2) *Unnoticeable pairs* with negligible quality differences, and (3) *Fine-grained pairs* with subtle but perceptible quality differences. We employ a two-step filtration process to remove coarse-grained and unnoticeable pairs.

**Score-based Coarse-grained Pairs Filtering.** To eliminate coarse-grained pairs, we establish score difference thresholds for each source dataset through expert discussion. The coarse-grained filtering criterion is defined as:

$$\mathcal{F}_{\text{coarse}}(p_i) = \begin{cases} 1, & \text{if } |s_i^A - s_i^B| \leq \tau_d \\ 0, & \text{otherwise} \end{cases}, \quad (1)$$

where  $p_i = (I_i^A, I_i^B)$  represents the  $i$ -th image pair,  $s_i^A$  and  $s_i^B$  are the quality scores of images  $I_i^A$  and  $I_i^B$  respectively, and  $\tau_d$  is the dataset-specific threshold. Pairs with  $\mathcal{F}_{\text{coarse}}(p_i) = 1$  are retained for further processing.

**JND-based Unnoticeable Pairs Filtering.** To remove unnoticeable pairs, we propose a JND-based filtering approach. The JND represents the minimum perceptual threshold for noticeable quality differences, making it ideal for identifying imperceptible quality pairs. Our JND-based filtering process operates as follows: (1) randomly select 20,000 images from the collected datasets, (2) compute the JND map for each image using established JND estimation model (Wu et al. 2017), and (3) generate JND noise images by adding these JND maps as noise to the original images:

$$I_{\text{JND}} = I + \text{JND}(I), \quad (2)$$

where  $I$  is the original image and  $\text{JND}(I)$  is the corresponding JND noise map. We calculate SSIM (Wang et al. 2004) between each  $I$  and  $I_{\text{JND}}$ , and use the median SSIM value  $\text{SSIM}_{\text{med}}$  as our filtering threshold. The unnoticeable pair filtering criterion is then defined as:

$$\mathcal{F}_{\text{unnotice}}(p_j) = \begin{cases} 1, & \text{if } \text{SSIM}(I_j^A, I_j^B) \leq \text{SSIM}_{\text{med}} \\ 0, & \text{otherwise} \end{cases}, \quad (3)$$

where  $p_j = (I_j^A, I_j^B)$  represents an image pair that passed the coarse-grained filtering. Pairs with  $\mathcal{F}_{\text{unnotice}}(p_j) = 1$  are retained as fine-grained pairs.

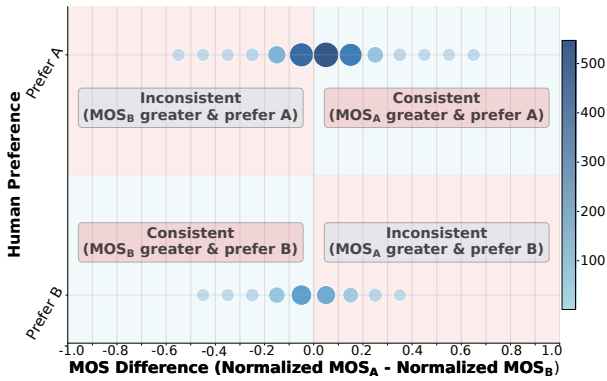


Figure 4: Consistency analysis between MOS scores and human preference rankings. Point sizes represent frequency of image pairs. Red regions indicate inconsistent cases where MOS scores and human preferences disagree.

After applying both filtering steps, we retain 30,886 fine-grained image pairs and their corresponding 18,408 images, forming the core of our FGRestore dataset. Due to space limitations, high-resolution fine-grained restoration examples are provided in the *supplementary material*.

### In-lab Subjective Study

FGRestore preserves the original score annotations from source datasets while introducing comprehensive fine-grained pairwise ranking annotations. For datasets originally annotated through pairwise ranking methodologies, fine-grained rankings can be directly derived from existing score relationships without additional human annotation. We design a resource-efficient two-round annotation protocol that incorporates quality control mechanisms to ensure reliable annotations. The annotation process is structured as follows:

**Annotation Protocol.** (1) Image pairs are divided into two groups, with each group assigned to a team of five trained annotators. For each image pair, annotators select the image with superior perceptual quality. Recognizing that some pairs may contain imperceptible differences despite filtering, annotators can also indicate equal quality when no clear preference exists. (2) Due to the subjective characteristics of fine-grained quality assessment, disagreements inevitably arise between annotators when quality differences are particularly subtle. Inconsistent annotations are resolved through expert review. This systematic annotation approach yields 45,318 pairwise preference annotations, providing a robust foundation for fine-grained quality assessment model training and evaluation.

### Dataset Analysis

To validate the necessity of fine-grained pairwise ranking, we analyze the consistency between MOS-based scoring and human preferences. Figure 4 compares normalized MOS differences ( $x$ -axis) with human preferences ( $y$ -axis), revealing significant inconsistencies. Points in the red regions represent cases where MOS scores and preferences disagree. Critically, image pairs with smaller MOS differences exhibit larger scatter, indicating that inconsistencies worsen as qual-

ity differences become subtler. This supports the necessity of fine-grained pairwise ranking, as MOS scoring becomes unreliable for discriminating subtle quality differences.

## FGResQ

Based on FGRestore, we propose FGResQ, a new fine-grained image quality assessment model for perceptual image restoration evaluation. The overall framework pipeline is illustrated in Figure 5, which consists of two main components: (a) Degradation-aware Feature Learning that incorporates restoration task knowledge to enable unified evaluation across multiple IR tasks, and (b) Dual-branch Quality Prediction that simultaneously handles both coarse-grained score regression and fine-grained pairwise ranking.

### Degradation-aware Feature Learning

To achieve unified quality assessment across diverse IR tasks, we propose a degradation-aware feature learning approach that enables our model to perceive and incorporate task-specific degradation characteristics. This design allows FGResQ to handle multiple IR scenarios within a unified framework, eliminating the need for task-specific model training. Specifically, we leverage a pretrained CLIP model to establish semantic alignment between visual content and degradation type. Specifically, we freeze the text encoder while fine-tuning a degradation encoder to learn degradation-aware representations. The degradation encoder maps images to a feature space that distinguishes different restoration scenarios. The learning objective employs bidirectional contrastive alignment to ensure robust degradation type perception:

$$\mathcal{L}_{cont} = \frac{1}{2} \left[ \mathcal{L}_{CE}(\mathbf{F}_I \mathbf{F}_T^T, \mathbf{y}) + \mathcal{L}_{CE}(\mathbf{F}_T \mathbf{F}_I^T, \mathbf{y}) \right], \quad (4)$$

where  $\mathbf{F}_I$  and  $\mathbf{F}_T$  represent image and text features respectively,  $\mathcal{L}_{CE}$  denotes the cross-entropy loss, and  $\mathbf{y} = [0, 1, 2, \dots, N-1]$  represents the ground-truth matching labels for  $N$  samples in a batch.

### Dual-branch Quality Prediction

Given degradation-aware features from the previous module, we first enhance them with learnable prompt embeddings to obtain quality-aware representations, then employ specialized prediction heads for different assessment tasks. Given an input image  $I$ , we extract general features  $\mathbf{F}_g \in \mathbb{R}^d$  using the image encoder and degradation features  $\mathbf{F}_d \in \mathbb{R}^d$  from the frozen degradation encoder. To effectively utilize degradation information, we employ learnable prompt embeddings  $\mathbf{p} \in \mathbb{R}^d$  to transform degradation features into quality-aware features:

$$\begin{aligned} \mathbf{F}_p &= \text{MLP}_1(\text{softmax}(\text{MLP}_2(\mathbf{F}_d)) \cdot \mathbf{p}) \\ \mathbf{F}_q &= \mathbf{F}_g \oplus \mathbf{F}_p \oplus (\mathbf{F}_g + \mathbf{F}_p) \end{aligned}, \quad (5)$$

where  $\text{MLP}_1$  and  $\text{MLP}_2$  are multi-layer perceptrons,  $\oplus$  denotes concatenation, and  $\mathbf{F}_q$  represents the final quality-aware features.

FGResQ employs two specialized prediction heads: *Regression Head* processes the quality features  $\mathbf{F}_q$  to predict absolute quality scores  $y_{pred}$  for individual images, enabling

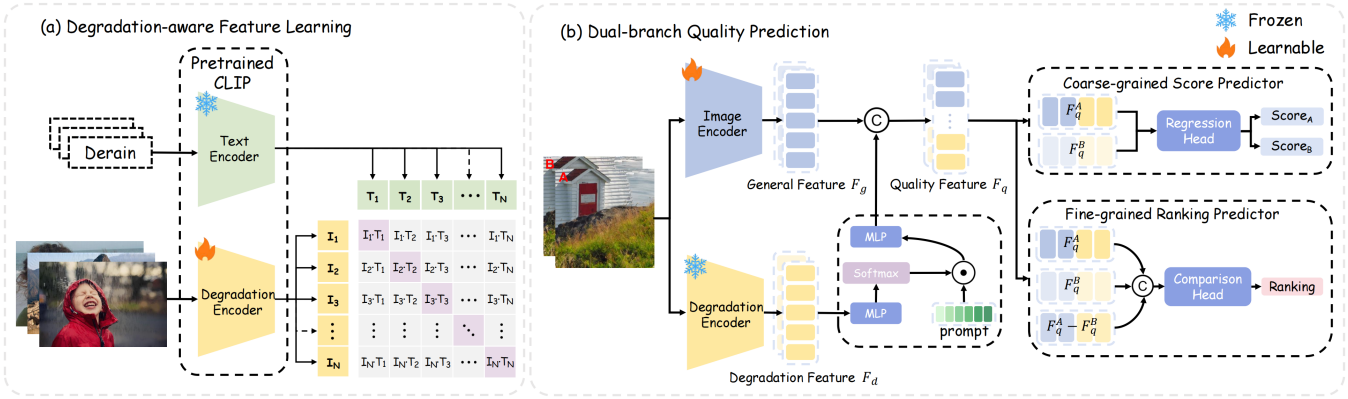


Figure 5: Overview of the proposed FGResQ framework.

coarse-grained quality assessment compatible with traditional evaluation protocols. *Comparison Head* handles fine-grained quality ranking by processing image pairs  $(I_A, I_B)$  to obtain their respective quality features  $(F_q^A, F_q^B)$  and predicting the ranking probability  $p_{AB}$ , directly addressing the fine-grained discrimination challenges identified in our preliminary analysis.

**Training Objectives** We design a comprehensive training strategy through carefully designed loss functions:

**Scene-aware Fidelity Loss.** Since FGRestore preserves original score annotations from different source datasets with varying quality scales, we employ a scene-aware fidelity loss that focuses on ranking relationships within each dataset (scene) rather than absolute score values:

$$\mathcal{L}_{fid}^{(s)} = \frac{1}{\binom{N_s}{2}} \sum_{i < j} \left( 1 - \sqrt{p_{ij}g_{ij} + \epsilon} - \sqrt{(1 - p_{ij})(1 - g_{ij}) + \epsilon} \right), \quad (6)$$

where  $p_{ij} = \sigma(y_{pred,i} - y_{pred,j})$ ,  $g_{ij} = \frac{1}{2}(\text{sign}(y_{gt,i} - y_{gt,j}) + 1)$ , and  $N_s$  is the number of samples in scene  $s$ .

The overall scene-aware loss combines fidelity losses from different scenes with inverse sample count weighting:

$$\mathcal{L}_{scene} = \sum_{s \in \mathcal{S}} w_s \mathcal{L}_{fid}^{(s)}, \quad (7)$$

where  $w_s$  is the inverse sample count weight for scene  $s$ .

**Pairwise Ranking Loss.** For fine-grained pairwise preferences, we employ binary cross-entropy loss:

$$\mathcal{L}_{rank} = -\frac{1}{M} \sum_{i=1}^M [r_i \log(p_i) + (1 - r_i) \log(1 - p_i)], \quad (8)$$

where  $M$  is the number of image pairs,  $r_i$  is the ground-truth ranking label, and  $p_i$  is the predicted ranking probability. The overall training objective combines both components:

$$\mathcal{L}_{total} = \lambda_1 \mathcal{L}_{scene} + \lambda_2 \mathcal{L}_{rank}, \quad (9)$$

where  $\lambda_1$  and  $\lambda_2$  are balancing hyperparameters.

## Experiments

### Evaluation Protocol

We compare the proposed FGResQ model against state-of-the-art FR and NR IQA methods. For fair comparison, all

learning-based methods with public source codes are re-trained using the same training-testing protocol with our scene-aware loss replacing their original regression loss, except for Compare2Score (Zhu et al. 2024) which operates in a zero-shot manner without training. We evaluate performance using SRCC and PLCC for quality score prediction, and accuracy (ACC) for fine-grained pairwise ranking. We split the dataset at an 8:2 ratio by image pairs. Individual image division following the corresponding pair.

### Implementation Details

We use CLIP ViT-B/16 as the backbone encoder for degradation-aware feature learning and dual-branch quality prediction. The model is trained on NVIDIA RTX4090 GPUs using the Adam optimizer with cosine annealing. The maximum learning rate is  $5 \times 10^{-6}$  and the batch size is 64. To ensure training stability, we employ scene-aware sampling that groups samples from the same source dataset within each batch.  $\lambda_1$  and  $\lambda_2$  are set to 5 and 1, respectively.

### Performance Evaluation

**Overall Performance Analysis.** Table 3 presents comprehensive performance comparisons across all six IR tasks in our FGRestore dataset. FGResQ achieves state-of-the-art performance across most evaluation metrics, with our method demonstrating particularly significant advantages in fine-grained ranking accuracy (ACC), which serves as the most critical metric for fine-grained quality assessment. The substantial improvement in ACC demonstrates FGResQ’s superior capability in distinguishing subtle quality differences. When compared to the strongest baseline methods, FGResQ shows consistent improvements across different evaluation paradigms. Among existing approaches, DeQA-Score exhibits the best coarse-grained regression performance, yet FGResQ surpasses it by significant margins with 2.5% SRCC improvement and 7.2% ACC improvement. More remarkably, our proposed method outperforms FR methods like DISTS despite these methods having access to additional reference images for quality assessment. This superior performance highlights the critical importance of explicitly modeling fine-grained quality relationships rather than relying solely on reference-based similarity measures,

Type	Method	Pub.	Deblurring			Denoising			Deraining		
			SRCC	PLCC	ACC	SRCC	PLCC	ACC	SRCC	PLCC	ACC
FR	PSNR	-	0.187	0.167	0.634	0.487	0.482	0.775	-	-	-
	SSIM	TIP'04	0.441	0.348	0.695	0.642	0.652	<b>0.789</b>	-	-	-
	LPIPS	CVPR'18	0.776	0.700	0.755	0.673	0.680	0.765	-	-	-
	DISTS	TPAMI'20	0.907	0.901	0.845	0.679	0.672	0.739	-	-	-
	A-FINE	CVPR'25	<u>0.907</u>	<u>0.796</u>	<u>0.865</u>	0.624	0.620	0.747	-	-	-
NR	NIQE	SPL'12	0.382	0.410	0.529	0.240	0.151	0.564	0.030	0.057	0.694
	BRISQUE	TIP'12	0.354	0.354	0.555	0.254	0.114	0.569	0.112	0.102	0.603
	DB-CNN	TCSVT'22	0.788	0.786	0.688	0.478	0.431	0.611	0.243	0.259	0.437
	HyperIQA	CVPR'25	0.871	0.887	0.402	0.605	0.625	0.675	0.264	0.294	0.484
	CLIP-IQA	AAAI'23	0.867	0.785	0.765	0.474	0.440	0.625	0.241	0.221	0.349
	Q-Align	ICML'24	0.767	0.804	0.795	0.676	0.687	0.731	0.433	0.421	0.455
	DeQA-Score	CVPR'25	0.815	0.843	0.819	0.754	0.771	0.778	<b>0.507</b>	<b>0.576</b>	0.426
	Compare2Score	NeurIPS'24	0.769	0.813	0.757	0.661	0.679	0.687	0.074	0.108	<b>0.790</b>
FGResQ	-	<b>0.926</b>	<b>0.910</b>	<b>0.873</b>	<b>0.759</b>	<b>0.777</b>	0.760	<u>0.496</u>	<u>0.518</u>	<u>0.778</u>	

Method	Dehazing			MixtureRestoration			SuperResolution			Average		
	SRCC	PLCC	ACC	SRCC	PLCC	ACC	SRCC	PLCC	ACC	SRCC	PLCC	ACC
PSNR	-	-	-	0.280	0.248	0.643	0.296	0.303	0.624	0.313	0.300	0.669
SSIM	-	-	-	0.421	0.379	0.684	0.351	0.361	0.641	0.464	0.435	0.702
LPIPS	-	-	-	0.466	0.475	0.658	0.448	0.460	0.666	0.591	0.579	0.711
DISTS	-	-	-	0.495	0.464	0.644	0.482	0.488	0.658	0.640	0.631	0.721
A-FINE	-	-	-	0.536	0.528	0.645	0.421	0.434	0.688	0.622	0.594	0.736
NIQE	0.036	0.063	0.446	0.041	0.027	0.541	0.176	0.072	0.503	0.151	0.130	0.546
BRISQUE	0.132	0.088	0.436	0.201	0.064	0.575	0.142	0.040	0.514	0.199	0.127	0.542
DB-CNN	0.643	0.645	0.524	0.415	0.460	0.614	0.459	0.454	0.643	0.504	0.506	0.586
HyperIQA	0.643	0.674	0.409	0.523	0.535	0.499	0.437	0.433	0.538	0.557	0.574	0.501
CLIP-IQA	0.547	0.499	0.546	0.302	0.300	0.580	0.244	0.186	0.571	0.446	0.405	0.573
Q-Align	0.715	0.765	0.584	0.569	0.571	0.658	0.376	0.366	0.662	0.589	0.603	0.648
DeQA-Score	0.762	0.803	0.644	0.669	0.679	0.697	<b>0.561</b>	<b>0.573</b>	0.718	0.678	0.707	0.680
Compare2Score	0.334	0.381	0.436	0.494	0.519	0.635	0.317	0.315	0.607	0.441	0.469	0.652
FGResQ	<b>0.821</b>	<b>0.854</b>	<b>0.669</b>	<b>0.698</b>	<b>0.706</b>	<b>0.713</b>	<u>0.521</u>	<u>0.536</u>	<b>0.721</b>	<b>0.703</b>	<b>0.717</b>	<b>0.752</b>

Table 3: Performance comparison on FGRestore dataset across different IR tasks. ”-” indicates no reference images available.

demonstrating that specialized fine-grained assessment capabilities are crucial for accurate IR evaluation.

**Fine-grained vs. Coarse-grained Assessment.** A critical observation is the performance gap between regression metrics and ranking accuracy across different methods. While some methods achieve reasonable regression performance, their fine-grained ranking capabilities remain limited.

**Qualitative Analysis.** Figure 6 presents qualitative comparisons on representative fine-grained image pairs across different restoration tasks. The results demonstrate systematic limitations of existing IQA methods in fine-grained scenarios. Traditional metrics like PSNR and SSIM often produce nearly identical scores, failing to provide meaningful quality discrimination. Advanced learning-based methods such as CLIP-IQA and DeQA-Score also struggle with subtle quality differences, frequently producing incorrect rankings. In contrast, FGResQ consistently identifies the superior image across tested cases, demonstrating the effectiveness of our pairwise comparison approach. Additional qualitative examples are provided in the *supplementary material*.

## Conclusion

In this work, we have addressed a critical yet underexplored challenge in image restoration evaluation: the inadequacy of existing IQA methods for fine-grained quality assessment. To address this fundamental limitation, we introduced FGRestore, the first fine-grained IQA dataset specifically designed for IR evaluation, comprising 18,408 restored images across six restoration tasks with 30,886 fine-grained pair-

	Image A	Image B	Metric	Score <sub>A</sub>	RANK	Score <sub>B</sub>
Deblurring			SSIM ↑	0.681	✗	0.703
			CLIP-IQA ↑	0.980	✗	0.999
			DeQA-Score ↑	3.226	✗	3.270
			FGResQ	Image A is better ✓		
Denoising			SSIM ↑	0.929	✗	0.933
			CLIP-IQA ↑	0.999	✓	0.996
			DeQA-Score ↑	3.256	✗	3.425
			FGResQ	Image A is better ✓		
Deraining			SSIM ↑	0.868	✗	0.837
			CLIP-IQA ↑	0.979	✓	0.993
			DeQA-Score ↑	1.764	✓	2.213
			FGResQ	Image B is better ✓		
Mixture Restoration			SSIM ↑	0.702	✗	0.617
			CLIP-IQA ↑	0.021	✓	0.830
			DeQA-Score ↑	3.329	✓	3.479
			FGResQ	Image B is better ✓		
Super Resolution			SSIM ↑	0.602	✗	0.653
			CLIP-IQA ↑	0.987	✗	0.993
			DeQA-Score ↑	2.667	✗	3.021
			FGResQ	Image A is better ✓		
Dehazing			SSIM ↑	0.792	✗	0.775
			CLIP-IQA ↑	0.261	✓	0.273
			DeQA-Score ↑	2.457	✗	2.264
			FGResQ	Image B is better ✓		

Figure 6: Qualitative comparison. (Best viewed zoomed in.)

wise preference annotations. Based on this, our proposed FGResQ model achieves state-of-the-art with significant improvements in fine-grained ranking accuracy. This work establishes a new evaluation framework for image restoration and provides valuable insights for developing more perceptually-aligned quality assessment methods in the era of advanced generative restoration models.

## Acknowledgments

This work is supported by National Natural Science Foundation of China under Grants 62471349, 62171340 and 62301378, Fundamental Research Funds for the Central Universities under Grant QTZX25076, and partly supported by the China Postdoctoral Science Foundation under Grant 2024M762553.

## References

- Chen, D.; Wu, T.; Ma, K.; and Zhang, L. 2025. Toward generalized image quality assessment: Relaxing the perfect reference quality assumption. *arXiv preprint arXiv:2503.11221*.
- Ding, K.; Ma, K.; Wang, S.; and Simoncelli, E. P. 2020. Image quality assessment: Unifying structure and texture similarity. *IEEE transactions on pattern analysis and machine intelligence*, 44(5): 2567–2581.
- Jiang, Q.; Liu, Z.; Gu, K.; Shao, F.; Zhang, X.; Liu, H.; and Lin, W. 2022. Single image super-resolution quality assessment: a real-world dataset, subjective studies, and an objective metric. *IEEE Transactions on Image Processing*, 31: 2279–2294.
- Jinjin, G.; Haoming, C.; Haoyu, C.; Xiaoxing, Y.; Ren, J. S.; and Chao, D. 2020. Pipal: a large-scale image quality assessment dataset for perceptual image restoration. In *European conference on computer vision*, 633–651. Springer.
- Ke, J.; Wang, Q.; Wang, Y.; Milanfar, P.; and Yang, F. 2021. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5148–5157.
- Li, L.; Sheng, X.; Chen, P.; Wu, J.; and Dong, W. 2025. Towards Explainable Image Aesthetics Assessment With Attribute-Oriented Critiques Generation. *IEEE Transactions on Circuits and Systems for Video Technology*, 35(2): 1464–1477.
- Lin, H.; Hosu, V.; and Saupé, D. 2019. KADID-10k: A large-scale artificially distorted IQA database. In *2019 Eleventh International Conference on Quality of Multimedia Experience (QoMEX)*, 1–3. IEEE.
- Liu, Y.; Wang, J.; Cho, S.; Finkelstein, A.; and Rusinkiewicz, S. 2013. A no-reference metric for evaluating the quality of motion deblurring. *ACM Trans. Graph.*, 32(6): 175–1.
- Luo, Z.; Gustafsson, F. K.; Zhao, Z.; Sjölund, J.; and Schön, T. B. 2023. Controlling Vision-Language Models for Universal Image Restoration. *arXiv preprint arXiv:2310.01018*.
- Ma, K.; Liu, W.; and Wang, Z. 2015. Perceptual evaluation of single image dehazing algorithms. In *2015 IEEE International Conference on Image Processing (ICIP)*, 3600–3604. IEEE.
- Mittal, A.; Moorthy, A. K.; and Bovik, A. C. 2012. No-reference image quality assessment in the spatial domain. *IEEE Transactions on image processing*, 21(12): 4695–4708.
- Mittal, A.; Soundararajan, R.; and Bovik, A. C. 2012. Making a “completely blind” image quality analyzer. *IEEE Signal processing letters*, 20(3): 209–212.
- Potlapalli, V.; Zamir, S. W.; Khan, S. H.; and Shahbaz Khan, F. 2023. Promptir: Prompting for all-in-one image restoration. *Advances in Neural Information Processing Systems*, 36: 71275–71293.
- Sheng, X.; Li, L.; Chen, P.; Wu, J.; Dong, W.; Yang, Y.; Xu, L.; Li, Y.; and Shi, G. 2023. AesCLIP: Multi-attribute contrastive learning for image aesthetics assessment. In *Proceedings of the 31st ACM International Conference on Multimedia*, 1117–1126.
- Sheng, X.; Xie, P.; Zou, W.; Chen, P.; Zhu, T.; and Li, L. 2025a. InstructCrop: Teaching Multimodal Large Language Models to Crop Aesthetic Images. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 6830–6839.
- Sheng, X.; Zou, W.; Chen, P.; Cai, L.; He, C.; and Li, L. 2025b. Text-to-Image Diffusion Models are AI-Generated Image Quality Scorers. In *Proceedings of the IEEE International Conference on Multimedia and Expo (ICME)*, 1–6.
- Shi, G.; Wan, W.; Wu, J.; Xie, X.; Dong, W.; and Wu, H. R. 2019. SISRSSet: Single image super-resolution subjective evaluation test and objective quality assessment. *Neurocomputing*, 360: 37–51.
- Su, S.; Yan, Q.; Zhu, Y.; Zhang, C.; Ge, X.; Sun, J.; and Zhang, Y. 2020. Blindly assess image quality in the wild guided by a self-adaptive hyper network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3667–3676.
- 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*, volume 37, 2555–2563.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.
- Wu, H.; Zhang, Z.; Zhang, W.; Chen, C.; Liao, L.; Li, C.; Gao, Y.; Wang, A.; Zhang, E.; Sun, W.; et al. 2023. Q-align: Teaching llms for visual scoring via discrete text-defined levels. *arXiv preprint arXiv:2312.17090*.
- Wu, H.; Zhu, H.; Zhang, Z.; Zhang, E.; Chen, C.; Liao, L.; Li, C.; Wang, A.; Sun, W.; Yan, Q.; et al. 2024. Towards open-ended visual quality comparison. In *European Conference on Computer Vision*, 360–377. Springer.
- Wu, J.; Li, L.; Dong, W.; Shi, G.; Lin, W.; and Kuo, C.-C. J. 2017. Enhanced just noticeable difference model for images with pattern complexity. *IEEE Transactions on Image Processing*, 26(6): 2682–2693.
- Wu, Q.; Wang, L.; Ngan, K. N.; Li, H.; and Meng, F. 2019. Beyond synthetic data: A blind deraining quality assessment metric towards authentic rain image. In *2019 IEEE International Conference on Image Processing (ICIP)*, 2364–2368. IEEE.
- Wu, T.; Zou, J.; Liang, J.; Zhang, L.; and Ma, K. 2025. VisualQuality-R1: Reasoning-Induced Image Quality Assessment via Reinforcement Learning to Rank. *arXiv preprint arXiv:2505.14460*.

- Xu, J.; Wu, M.; Hu, X.; Fu, C.-W.; Dou, Q.; and Heng, P.-A. 2024. Towards real-world adverse weather image restoration: Enhancing clearness and semantics with vision-language models. In *European Conference on Computer Vision*, 147–164. Springer.
- Yang, Z.; Li, L.; Chen, P.; Wu, J.; and Dong, W. 2024a. Semantics-Aware Image Aesthetics Assessment using Tag Matching and Contrastive Ranking. In *ACM International Conference on Multimedia*, 2632–2641.
- Yang, Z.; Li, L.; Chen, P.; Wu, J.; and Valenzise, G. 2025. Language-Guided Visual Perception Disentanglement for Image Quality Assessment and Conditional Image Generation. *arXiv preprint arXiv:2503.02206*.
- Yang, Z.; Li, L.; Yang, Y.; Li, Y.; and Lin, W. 2024b. Multi-Level Transitional Contrast Learning for Personalized Image Aesthetics Assessment. *IEEE Transactions on Multimedia*, 26: 1944–1956.
- You, Z.; Cai, X.; Gu, J.; Xue, T.; and Dong, C. 2025. Teaching large language models to regress accurate image quality scores using score distribution. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 14483–14494.
- Zhang, L.; Zhang, L.; and Bovik, A. C. 2015. A feature-enriched completely blind image quality evaluator. *IEEE Transactions on Image Processing*, 24(8): 2579–2591.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018a. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zhang, W.; Ma, K.; Yan, J.; Deng, D.; and Wang, Z. 2020. Blind Image Quality Assessment Using A Deep Bilinear Convolutional Neural Network. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(1): 36–47.
- Zhang, W.; Zhai, G.; Wei, Y.; Yang, X.; and Ma, K. 2023. Blind image quality assessment via vision-language correspondence: A multitask learning perspective. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14071–14081.
- Zhang, X.; Lin, W.; and Huang, Q. 2021. Fine-grained image quality assessment: A revisit and further thinking. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(5): 2746–2759.
- Zhang, X.; Lin, W.; Wang, S.; Liu, J.; Ma, S.; and Gao, W. 2018b. Fine-grained quality assessment for compressed images. *IEEE Transactions on Image Processing*, 28(3): 1163–1175.
- Zhao, S.; Zhang, L.; Huang, S.; Shen, Y.; and Zhao, S. 2020. Dehazing evaluation: Real-world benchmark datasets, criteria, and baselines. *IEEE Transactions on Image Processing*, 29: 6947–6962.
- Zhou, F.; Yao, R.; Liu, B.; and Qiu, G. 2019. Visual quality assessment for super-resolved images: Database and method. *IEEE Transactions on Image Processing*, 28(7): 3528–3541.
- Zhou, M.; Ye, K.; Delbracio, M.; Milanfar, P.; Patel, V. M.; and Talebi, H. 2025. UniRes: Universal Image Restoration for Complex Degradations. *arXiv preprint arXiv:2506.05599*.
- Zhu, H.; Li, L.; Wu, J.; Dong, W.; and Shi, G. 2020. MetalQA: Deep meta-learning for no-reference image quality assessment. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14143–14152.
- Zhu, H.; Wu, H.; Li, Y.; Zhang, Z.; Chen, B.; Zhu, L.; Fang, Y.; Zhai, G.; Lin, W.; and Wang, S. 2024. Adaptive image quality assessment via teaching large multimodal model to compare. *Advances in Neural Information Processing Systems*, 37: 32611–32629.