ATTRIBUTE-CALIBRATED AND QUALITY-AWARE VISION-LANGUAGE LEARNING FOR IMAGE QUALITY ASSESSMENT

Anonymous authors
Paper under double-blind review

ABSTRACT

Existing image quality assessment (IQA) methods primarily rely on score regression paradigms while neglecting semantic analysis, making score-only supervision unreliable across diverse scenarios. Additionally, the relationship between quality-aware attributes and cross-scenario quality fluctuations remains underexplored. To address these issues, an image quality assessment method based on attribute-calibrated and quality-aware vision-language learning (ACQA) is proposed. Firstly, a dual-branch vision-language model (VLM) is employed to extract content and attribute features for subsequent quality analysis. Secondly, an attribute calibration learning (ACAL) module is proposed to calibrate the VLM in attribute analysis by jointly incorporating distortion ranking, stability, and disentanglement, thereby enabling adaptive attribute information across scenarios. Thirdly, a quality-aware analysis learning (QAL) module is proposed to integrate attribute contributions with content alignment to optimize the VLM. Finally, the optimized VLM provides reliable content and attribute information for quality prediction. Experimental results demonstrate that the proposed ACAL and QAL can effectively refine content and attribute information to enhance the reliability of quality prediction. Moreover, ACQA outperforms existing IQA methods in both accuracy and generalization. The codes and models are publicly available at https://github.com/dart-into/ACQA.

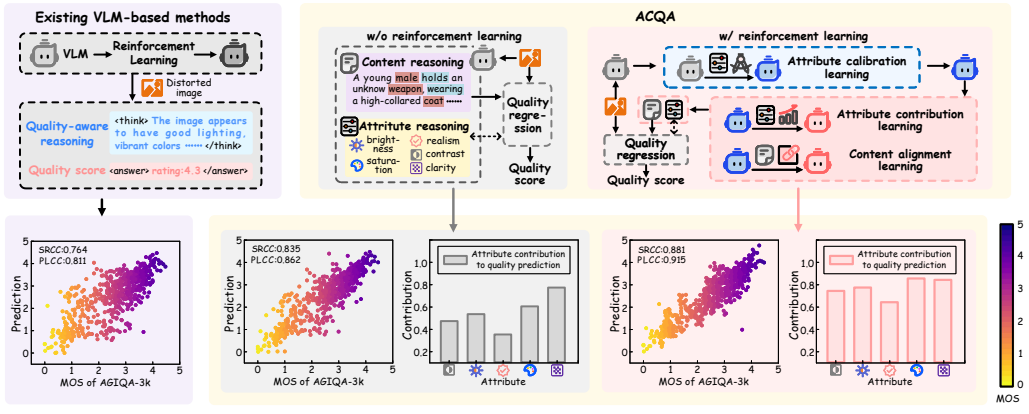


Figure 1: Compared to existing methods, our method focus on content alignment and attribute contribution to predict quality. Experimental results demonstrate the the effectiveness of ACQA.

1 INTRODUCTION

Image degradations objectively compromise visual fidelity, making image quality assessment (IQA) (Mittal et al. (2012a); Zhang et al. (2023)) important for ensuring the usability in computer vision tasks. Most IQA methods (Mittal et al. (2012b); Lu et al. (2024)) leverage score-oriented paradigm

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

based on handcrafted features or deep features for quality regression, which lacks further semantic analysis and struggles to generalize across scenarios.

Vision-language models (VLMs), e.g. Qwen, Deepseek, have recently been employed in several IQA methods (Huang et al. (2024); Wang et al. (2024)) to bridge the gap between semantic analysis and quality prediction. Some of these methods, e.g (Zhang et al. (2023); You et al. (2024)) have successfully introduced frozen VLMs to describe semantic content for quality prediction. Conversely, effectively adjusting semantic information across scenarios is still a challenge. Some recent methods, e.g. Q-Insight (Li et al. (2025)) and EvoQuality (Wen et al. (2026)), optimize VLMs for semantic information adjustment by employing reinforcement learning (RL). These RL-based methods primarily focus on the normativity, accuracy, and stability of generated outputs. However, the analysis of the relationship between the generative process and quality prediction is insufficient.

Furthermore, some methods (You et al. (2024); Jin et al. (2023); Yue et al. (2024)) utilize fixed attribute evaluation strategies to improve the performance of quality perception. However, human subjective perception of attributes varies across scenarios (Cao et al. (2026); Liu et al. (2025)), which causes different quality fluctuation. Therefore, analyzing the relationships between attributes and quality under changing scenarios is essential.

Based on the above analysis, we propose an attribute-calibrated and quality-aware vision-language learning for image quality assessment (ACQA), as illustrated in Figure 1. Firstly, a dual-branch VLM is employed to extract content descriptions and content scores. Secondly, an attribute calibration learning (ACAL) module is designed to calibrate the attribute perception of VLM across scenarios by jointly incorporating distortion ranking, stability, and disentanglement. Thirdly, a quality-aware analysis learning (QAL) module is proposed to combine attribute contributions with content alignment to optimize the VLM. Experimental results demonstrate that ACQA outperforms existing IQA methods in both accuracy and generalization. The contributions of our work are as follows:

- A dual-branch VLM is employed to extract content description and content scores for subsequent quality analysis. In addition, an ACAL module is designed to calibrate the VLM for attribute adjustment by integrating distortion ranking, stability, and disentanglement.
- A QAL module combining attribute contributions with content alignment is proposed to optimize the VLM for quality-aware adaptation, which enables quality-aware semantic information analysis.
- Experimental results demonstrate that proposed ACAL and QAL can provide a reference for semantic-based IQA. Furthermore, ACQA outperforms existing IQA methods in both accuracy and generalization.

2 Related Work

2.1 IQA

Early IQA methods (Mittal et al. (2012a;b)) relied on handcrafted features to mimic human subjective evaluation. With the advent of deep learning, deep networks have been extensively applied to IQA (Jiang et al. (2022); Fang et al. (2025)). These methods primarily focus on score-oriented paradigm, but lack semantic analysis.

Recent methods (Huang et al. (2024); Wang et al. (2024)) adopt VLMs to offer semantic information for quality prediction. By leveraging a VLM to generate quality descriptions, LIQE (Zhang et al. (2023)) further exploits CLIP to capture multimodal semantics. CLIP-AGIQA (Tang et al. (2024)) utilizes CLIP to effectively analyze keywords in text to extract features of subjects. These methods rely on frozen VLMs to obtain fixed semantic outputs, making it difficult to dynamically adapt to diverse scenarios. To optimize the VLM for different scenarios, Q-Insight (Li et al. (2025)) employs RL based on group relative policy optimization (GRPO). Based on this work, VQ-R1 (Wu et al. (2025)) utilizes RL to fine-tune the VLM by analyzing accuracy and stability of quality-aware information. These methods can achieve semantic information optimization by employing RL when scenario changes.

In conclusion, leveraging VLMs for IQA has advantages in semantic analysis over traditional methods, and RL can enhance the adaptation of VLMs in IQA tasks.

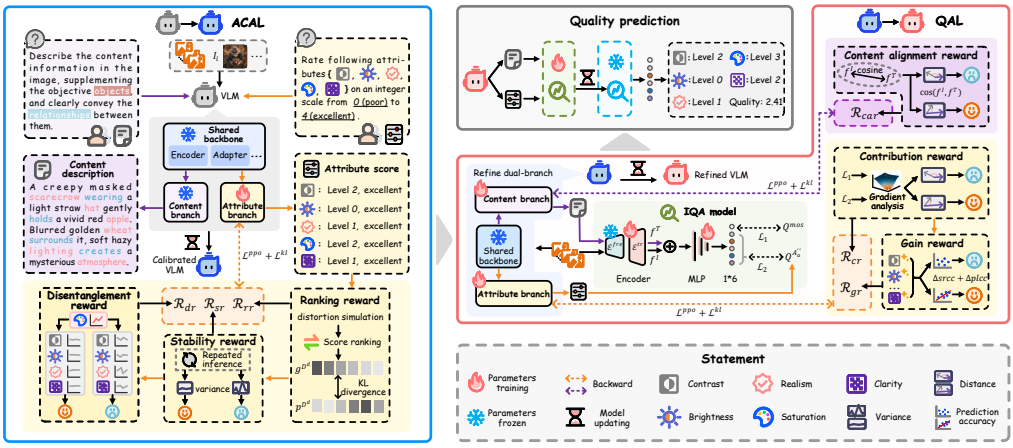


Figure 2: The illustration of ACQA framework. ACQA consists of a dual-branch VLM, ACAL module, and QAL module. The dual-branch VLM aims to extract content description and attribute scores. ACAL module can calibrate the VLM for attribute adjustment across scenarios. And, QAL module is designed to guide the VLM to align both content and attribute for quality prediction.

2.2 RL

With the rapid advancement of VLMs, recent studies have increasingly explored RL to enhance the reasoning capabilities in computer vision tasks (Guo et al. (2025a); Zheng et al. (2026); Mao et al. (2026)). For example, iRe-VLA (Guo et al. (2025b)) combines vision models with VLMs by adopting RL to improve the performance in vision tasks. RePrompt (Wu et al. (2026)) introduces explicit reasoning into the prompt enhancement and image generation via reinforcement learning. The success of reinforcement learning in visual tasks provides references for RL-based IQA.

Further analysis reveals that the core benefit of RL in IQA stems from a reasoning representation mechanism (Chopparapu et al. (2024); Li et al. (2025)). RL can guide models to explore quality-aware reasoning, thereby improving generalization. For instance, Refine-IQA (Jia et al. (2026)) employs RL based on GRPO to enhance the regularity and accuracy of semantic content. RALI (Zhao et al. (2026)) focuses on utilizing RL to align semantic information with images. EvoQuality (Wen et al. (2026)) and VQ-R1 (Wu et al. (2025)) utilizes RL based on accuracy and stability to fine-tune the VLM for quality-aware reasoning. However, these methods primarily design RL for the final generation outcomes of VLMs, lacking optimization of the generation process.

In conclusion, RL can effectively optimize VLMs for IQA tasks, and considering both generation outcomes and process for RL is crucial to obtain sufficient quality-aware semantics.

3 METHOD

ACQA consists of a dual-branch VLM, ACAL module, and QAL module. The framework of ACQA is illustrated in Figure 2.

3.1 DUAL-BRANCH VLM

A dual-branch VLM $\mathcal{M}^{vlm}$ based on Qwen 2.5-7B is proposed to analyze the content and quality-aware attributes. Unlike standard VLMs that focus solely on semantic analysis, $\mathcal{M}^{vlm}$ decouples the generation of content description and attribute scores. Specifically, the early-stage visual encoder and feature fusion modules of $\mathcal{M}^{vlm}$ are retained as shared backbone $\mathcal{M}^{s}$. And, the subsequent semantic decoding layers are bifurcated into a **content branch** $\mathcal{B}^{con}$ and an **attribute branch** $\mathcal{B}^{att}$.

Content Branch. $\mathcal{M}^{s}$ and $\mathcal{B}^{con}$ are employed to obtain content description of images $\{\mathcal{I}_i^{tr} | i \in [0, N]\}$. To capture the content information of $\{\mathcal{I}_i^{tr}\}$, the prompt $\mathcal{P}^{con}$ is designed to emphasize a clear description of the objective objects and the logical relationships between them. As shown in Figure 2, $\mathcal{M}^{vlm}$ focuses on objects such as 'male', 'bowl', and 'table', and the logical relationships

like 'holding', 'eating', and "lying." Consequently, both $\{\mathcal{I}_i^{tr}\}$ and $\mathcal{P}^{con}$ are input into $\mathcal{M}^s$ and $\mathcal{B}^{con}$ to obtain the content description $\{T_i^{con}\}$:

$$\{T_i^{con}\} = \mathcal{B}^{con}(\mathcal{M}^s(\{\mathcal{I}_i^{tr}\}, \mathcal{P}^{con})). \tag{1}$$

Attribute Branch. $\mathcal{M}^s$ and $\mathcal{B}^{att}$ are employed to obtain attribute scores of $\{\mathcal{I}_i^{tr}\}$. The prompt $\mathcal{P}^{att}$ instructs five attributes $\{\mathcal{A}_a | a \in [1, 5]\}$, namely contrast, brightness, realism, saturation, and blur, of $\{\mathcal{I}_i^{tr}\}$. The integer of attribute score is the scale of 0 to 4, corresponding to performance levels of 'poor', 'fair', 'good', 'very good', and 'excellent', respectively. Consequently, both $\{\mathcal{I}_i^{tr}\}$ and $\mathcal{P}^{att}$ are input into $\mathcal{M}^s$ and $\mathcal{B}^{att}$ to obtain the attribute scores $\{Q_i^a\}$:

$$\{Q_i^{\mathcal{A}_a}\} = \mathcal{B}^{att}(\mathcal{M}^s(\{\mathcal{I}_i^{tr}\}, \mathcal{P}^{att})). \tag{2}$$

3.2 ACAL

The ACAL module is proposed to calibrate the $\mathcal{B}^{att}$ performance of attribute scoring across diverse scenarios. Specifically, **ranking reward**, **stability reward**, and **disentanglement reward** in ACAL are designed to optimize $\mathcal{B}^{att}$ from the perspectives of distortion ranking, stability, and disentanglement, respectively. And, $\mathcal{M}^s$ and $\mathcal{B}^{con}$ are frozen in this stage.

Ranking Reward. Based on prior work in distortion simulation, four types of synthetic distortions $\{\mathcal{D}_d | d \in [1, 4]\}$, namely contrast distortion, overexposure, color shift, and blur, are applied in $\mathcal{I}_i^{tr}$ with five distortion level $l \in [1, 5]$. In this way, distorted images $\{\mathcal{I}_{i,l}^{\mathcal{D}_d}\}$ can provide knowledge regarding contrast, brightness, saturation, and clarity for the reasoning process, respectively. The ground-truth ranking distribution for $\{\mathcal{I}_{i,l}^{\mathcal{D}_d}\}$ across the five distortion levels under the same distortion type is $g_i^{\mathcal{D}_d}$, which aligns with the exponential decay distribution of human visual quality perception:

$$g_i^{\mathcal{D}_d} = \frac{\exp(-\alpha \cdot l)}{\sum_{l=1}^{5} \exp(-\alpha \cdot l)}, \tag{3}$$

where α denotes the attenuation coefficient. And, $\mathcal{M}^s$ and $\mathcal{B}^{att}$ are employed to obtain scores $\{Q_{i,l}^{\mathcal{A}_a}\}$ of $\{\mathcal{I}_{i,l}^{\mathcal{D}_d}\}$ in attribute $\mathcal{A}_a$. In this way, the ranking probability distribution $p_i^{\mathcal{D}_d}$ predicted by $\mathcal{M}^s$ and $\mathcal{B}^{att}$ is:

$$p_i^{\mathcal{D}_d} = softmax(\{Q_{i,l}^{\mathcal{A}_a}\}). \tag{4}$$

The kullback-leibler (KL) divergence between the ranking distributions of $g_i^{\mathcal{D}_d}$ and $p_i^{\mathcal{D}_d}$ is used as the ranking reward value $\mathcal{R}_{rr}$ in a batch N^b:

$$\mathcal{R}_{rr} = \frac{1}{4N^b} \sum_{i=1}^{N^b} \sum_{d=1}^{4} g_i^{\mathcal{D}_d} \log(\frac{g_i^{\mathcal{D}_d}}{p_i^{\mathcal{D}_d}}). \tag{5}$$

Stability Reward. The $t \in [1, \mathcal{T}]$ times reasoning is employed to obtain attribute score $\{Q_{i,t}^{\mathcal{A}_a}\}$. Thus, the variance of the multiple inference results for $\{\mathcal{I}_{i,l}^{tr}\}$ is used as the stability reward value $\mathcal{R}_{sr}$:

$$\mathcal{R}_{sr} = \frac{1}{5N^b} \sum_{i=1}^{N^b} \sum_{a=1}^{5} [\frac{1}{\mathcal{T}} \sum_{t=1}^{\mathcal{T}} (Q_{i,t}^{\mathcal{A}_a} - \frac{1}{\mathcal{T}} \sum_{t=1}^{\mathcal{T}} Q_{i,t}^{\mathcal{A}_a})^2]. \tag{6}$$

By combining $\mathcal{R}_{rr}$ and $\mathcal{R}_{sr}$, a relative advantage reward $\hat{\mathcal{R}}_{(rr,sr)}$ can be obtained through group-wise normalization to eliminate reward scale discrepancies. The proximal policy optimization with clipped objective is employed to limit the single-step update magnitude and prevent training collapse. The clipped policy loss is denoted as $\mathcal{L}_{rr,sr}^{ppo}$:

$$\mathcal{L}_{rr,sr}^{ppo} = \frac{1}{N^b} \sum_{i=1}^{N^b} [min\{r_i \cdot \hat{\mathcal{R}}_{(rr,sr)}, f_{clip}(r_i, 1-\epsilon, 1+\epsilon) \cdot \hat{\mathcal{R}}_{(rr,sr)}\}], \tag{7}$$

where r_i denotes the ratio of the output probabilities of $\mathcal{M}^s$ and $\mathcal{B}^{att}$ before and after the update, f_{clip} denotes the clipping function that constrains r_i to the range $[1-\epsilon, 1+\epsilon]$. In addition, the regularization result of the KL divergence of $\mathcal{M}^s$ and $\mathcal{B}^{att}$ before and after the update is used as the

loss $\mathcal{L}^{kl}_{rr,sr}$, which prevents the policy from deviating significantly from the original model. Thus, parameters θ^{att} of $\mathcal{B}^{att}$ are updated to $\theta^{att'}$ with learning rate η:

$$\theta^{att'} \leftarrow \theta^{att} - \eta \nabla_{\theta^{att}} (\mathcal{L}^{ppo}_{rr,sr} + \mathcal{L}^{kl}_{rr,sr})(\theta^{att}). \tag{8}$$

Disentanglement Reward. In the stage of optimizing θ^{att}, when the distortion $\mathcal{D}_w, w \in [1,4]$ exhibits wrong ranking and others are correct, a disentanglement analysis is performed for $\mathcal{D}_w$. The $\mathcal{D}_w$ corresponds to the attribute $\mathcal{A}_w$. And, the updated attribute scores $\{Q_i^{\mathcal{A}'_a} | a \neq w\}$ are predicted according $\theta^{att'}$. The discrepancy between $\{Q_i^{\mathcal{A}'_a}\}$ and $\{Q_i^{\mathcal{A}_a}\}$ is calculated as $\mathcal{R}_{dr}$:

$$\mathcal{R}_{dr} = \frac{1}{4N^b} \sum_{i=1}^{N^b} \sum_{a=1}^{5} ||Q_i^{\mathcal{A}'_a} - Q_i^{\mathcal{A}_a}||^2, a \neq w. \tag{9}$$

Based on $\mathcal{R}_{dr}$, $\hat{\mathcal{R}}_{dr}$ can be obtained to optimize $\theta^{att'}$ to $\theta^{att''}$ via Eq. 7 and Eq. 8, and the process is described in Algorithm 1.

Algorithm 1 The process of ACAL.

Input: $\{\mathcal{I}_i^{tr}\}$, $\mathcal{B}^{att}$, θ^{att}, $\mathcal{A}_a$, l, $\mathcal{D}_d$, and N^b.
Output: $\theta^{att''}$.
1: $\{\mathcal{I}_{i,l}^{\mathcal{D}_d}\} \leftarrow \{\mathcal{I}_i^{tr}\}$. // Distortion simulation
2: $g_i^{\mathcal{D}_d} = \exp(-\alpha \cdot l) / \sum_{l=1}^{5} \exp(-\alpha \cdot l)$.
3: $\{Q_{i,l}^{\mathcal{A}_a}\} = \mathcal{B}^{att}(\{\mathcal{I}_{i,l}^{\mathcal{D}_d}\}, \theta^{att})$.
4: $p_i^{\mathcal{D}_d} = softmax(\{Q_{i,l}^{\mathcal{A}_a}\})$.
5: $\mathcal{R}_{rr} = \frac{1}{4N^b} \sum_{i=1}^{N^b} \sum_{d=1}^{4} g_i^{\mathcal{D}_d} \log(\frac{g_i^{\mathcal{D}_d}}{p_i^{\mathcal{D}_d}})$.
6: t times reasoning $\{\mathcal{I}_{i,t}^{tr}\}$ for $\{Q_{i,t}^{\mathcal{A}_a}\}$.
7: $\mathcal{R}_{sr} = \frac{1}{5N^b} \sum_{i=1}^{N^b} \sum_{a=1}^{5} [variance(Q_{i,t}^{\mathcal{A}_a})]$.
8: Update θ^{att} to $\theta^{att'}$ with $\mathcal{R}_{rr}, \mathcal{R}_{sr}$.
9: Select $\mathcal{A}_w$ with wrong ranking $\mathcal{D}_w$.
// Attribute disentanglement reward.
10: $\{Q_i^{\mathcal{A}'_a}\} = \mathcal{B}^{att}(\mathcal{I}_i^{tr}, \theta^{att'})$.
11: $\mathcal{R}_{dr} = \frac{1}{4N^b} \sum_{i=1}^{N^b} \sum_{1}^{4} ||Q_i^{\mathcal{A}'_a} - Q_i^{\mathcal{A}_a}||^2, a \neq w$.
12: Update $\theta^{att'}$ to $\theta^{att''}$ with $\mathcal{R}_{dr}$.
Return $\theta^{att''}$

3.3 QAL

The QAL module is proposed to guide $\mathcal{M}^{vlm}$ to learn quality-aware knowledge. Specifically, **content alignment reward** is designed to optimize $\mathcal{B}^{con}$ for better semantic analysis. Meanwhile, **contribution reward** and **gain reward** are designed to adjust $\mathcal{B}^{att}$ in attribute analysis to suit for quality prediction. Moreover, $\mathcal{M}^s$ is frozen in this stage.

Content Alignment Reward. The correlation between $\{\mathcal{I}_i^{tr}\}$ and $\{T_i^{con}\}$ is analyzed to optimize $\mathcal{B}^{con}$. Specifically, a CLIP-based IQA model $\mathcal{M}^{IQA}$ is proposed to obtain the image feature $\{f_i^I\}$ and text feature $\{f_i^T\}$. Only the last four layers $\mathcal{E}^{tr}$ of the image and text encoder are trained, while the remaining layers $\mathcal{E}^{fre}$ are frozen. $\{f_i^I\}$ and $\{f_i^T\}$ are concatenated to obtain five attribute scores $\{Q_i^{\mathcal{A}_a,pre}\}$ and quality scores $\{Q_i^{qua}\}$. The cosine similarity between $\{f_i^I\}$ and $\{f_i^T\}$ is used as the content alignment reward $\mathcal{R}_{car}$:

$$\mathcal{R}_{car} = \frac{1}{N^b} \sum_{i=1}^{N^b} \frac{f_i^I \cdot f_i^T}{||f_i^I||^2 ||f_i^T||^2}. \tag{10}$$

Based on $\mathcal{R}_{car}$, $\hat{\mathcal{R}}_{car}$ can be obtained to optimize parameters θ^{con} of $\mathcal{B}^{con}$ to $\theta^{con'}$ via Eq. 7 and Eq. 8.

Contribution Reward. Through quantitative analysis of $\{T_i^{con}\}$ and $\{Q_i^{\mathcal{A}''_a}\}$ obtaining by $\theta^{att''}$, the learning of $\mathcal{M}^{IQA}$ for quality fluctuations is optimized. Specifically, quality-oriented training is conducted on $\mathcal{M}^{IQA}$. The euclidean norm between $\{Q_i^{qua}\}$ and the quality label $\{Q_i^{mos}\}$ is employed as the loss $\mathcal{L}_1$:

$$\mathcal{L}_1 = \frac{1}{N^b} \sum_{i=1}^{N^b} ||Q_i^{mos} - Q_i^{qua}||^2. \tag{11}$$

Table 1: Performance comparisons of existing IQA methods in synthetic, authentic, and AIGC tasks.

	Methods (year)	LIVEC		TID2013		AGIQA-3k		AGIQA-20k		Average	
		SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC
Hand-craft IQA	BRISQUE (2012)	0.629	0.629	0.626	0.571	0.473	0.556	0.431	0.459	0.539	0.553
	NIQE (2012)	0.594	0.589	0.521	0.648	0.405	0.462	0.421	0.439	0.485	0.534
	ILNIQE (2019)	0.508	0.508	0.521	0.648	0.490	0.553	0.526	0.541	0.511	0.562
Classic DL IQA	MUSIQ (2021)	0.702	0.746	0.773	0.815	0.764	0.812	0.793	0.805	0.758	0.794
	DBCNN (2023)	0.851	0.869	0.816	0.865	0.820	0.875	0.836	0.849	0.830	0.864
	Re-IQA (2023)	0.840	0.854	0.804	0.861	0.832	0.853	0.842	0.869	0.829	0.859
	MCOLE (2024)	0.812	0.831	0.895	0.905	0.845	0.871	0.819	0.805	0.842	0.853
	DGQA (2024)	0.782	0.807	0.857	0.852	0.763	0.811	0.774	0.780	0.794	0.812
	GDCIQA (2026)	0.829	0.841	0.860	0.869	0.847	0.872	0.813	0.832	0.837	0.853
VLM-based IQA	Q-Align (2024)	0.865	0.873	0.859	0.845	0.723	0.786	0.749	0.757	0.799	0.815
	VQ-R1 (2025)	0.750	0.794	0.829	0.841	0.783	0.849	0.726	0.720	0.772	0.801
	EvoQuality (2026)	0.813	0.847	0.611	0.674	0.831	0.771	0.806	0.827	0.765	0.779
	Refine-IQA-S2 (2026)	0.870	0.892	0.739	0.760	0.798	0.841	0.817	0.826	0.806	0.829
	RALI (2026)	0.876	0.896	0.817	0.844	0.715	0.779	0.726	0.742	0.783	0.815
	Q-Insight (2026)	0.865	0.893	0.822	0.831	0.764	0.811	0.718	0.729	0.792	0.816
	(Ours)	0.890	0.914	0.903	0.928	0.881	0.915	0.876	0.862	0.887	0.907

Furthermore, attribute-guided training is performed on $\mathcal{M}^{IQA}$, which can assist quality analysis. The euclidean norm between $\{Q_i^{\mathcal{A}_a''}\}$ and $\{Q_i^{\mathcal{A}_a,pre}\}$ is employed as the loss $\mathcal{L}_2$:

$$\mathcal{L}_2 = \frac{1}{5N^b}\sum_{i=1}^{N^b}\sum_{a=1}^{5}||Q_i^{\mathcal{A}_a,pre} - Q_i^{\mathcal{A}_a''}||^2. \tag{12}$$

The parameters θ^{IQA} of $\mathcal{M}^{IQA}$ can be updated to $\theta^{IQA'}$ by combining $\mathcal{L}_1$ and $\mathcal{L}_2$. When $\mathcal{M}^{IQA}$ approaches convergence, the cosine distance $cos(\cdot)$ between the gradient descent directions corresponding to five attributes and quality is computed to quantify the attribute contribution. Thus, the $cos(\cdot)$ is used as the contribution reward value $\mathcal{R}_{cr}$ of attributes:

$$\mathcal{R}_{cr} = \frac{1}{5}\sum_{a=1}^{5} cos(vec(\nabla\theta^{IQA}\mathcal{L}_1), vec(\nabla\theta^{IQA}(\frac{1}{N}\sum_{i=1}^{N}||Q_i^{\mathcal{A}_a,pre} - Q_i^{\mathcal{A}_a''}||^2))), \tag{13}$$

where $vec(\cdot)$ flattens gradients into a single vector. Then, $\hat{\mathcal{R}}_{cr}$ can be obtained to optimize $\theta^{att'''}$ to $\theta^{att'''}$ via Eq. 7 and Eq. 8.

Gain Reward. Based on $\theta^{att'''}$, the optimized attribute scores $\{Q_i^{\mathcal{A}_a'''}\}$ are derived. Quantify the gain of the adjustment from $\{Q_i^{\mathcal{A}_a''}\}$ to $\{Q_i^{\mathcal{A}_a'''}\}$ to optimize attribute analysis. Specifically, $\theta^{IQA'}$ is fine-tuned to $\theta^{IQA''}$ via $\mathcal{L}_1$ and $\mathcal{L}_2$. The quality scores predicted by $\theta^{IQA'}$ and $\theta^{IQA''}$ are defined as $\{Q_i^{qua'}\}$ and $\{Q_i^{qua''}\}$, respectively. The spearman rank correlation coefficient (SRCC) and pearson linear correlation coefficient (PLCC) between $\{Q_i^{qua'}\}$ and $\{Q_i^{mos}\}$ are $w^{srcc'}$ and $w^{plcc'}$, while those of $\{Q_i^{qua''}\}$ are $w^{srcc''}$ and $w^{plcc''}$. The gain reward value $\mathcal{R}_{gr}$ of quality can be obtained by:

$$\mathcal{R}_{gr} = \sigma(\gamma \cdot \frac{1}{2}[(w^{srcc''} - w^{srcc'}) + (w^{plcc''} - w^{plcc'})]), \tag{14}$$

where $\sigma(\cdot)$ and γ are a sigmoid function and temperature, respectively. Then, $\mathcal{R}_{gr}$ is employed to optimize $\theta^{att'''}$ to $\theta^{att''''}$. Finally, $\theta^{con'}$ and $\theta^{att''''}$ are used to obtain refined content description $\{T_i^{con*}\}$ and attribute scores $\{Q_i^{a*}\}$ to fine-tune $\theta^{IQA''}$ to $\theta^{IQA'''}$.

4 Experiments

4.1 Experimental Settings

This section includes the expression of **datasets**, **metrics**, and **method details**.

Table 2: Effectiveness of dual-branch VLM.

#	Content branch	Attribute branch	Distorted datasets		Specific distortions in KADID		
			LIVEC	AGIQA-3k	Guassian blur	Darken	Color shift
1	✗	✗	0.802/0.839	0.814/0.856	0.869/0.872	0.886/0.728	0.845/0.862
2	✓	✗	0.849/0.852↑	0.845/0.878↑	0.882/0.895↑	0.904/0.913↑	0.852/0.875↑
3	✗	✓	0.843/0.849↑	0.850/0.895↑	0.892/0.898↑	0.895/0.901↑	0.871/0.880↑
4	✓	✓	0.890/0.914↑	0.881/0.915↑	0.897/0.916↑	0.910/0.925↑	0.908/0.892↑

Table 3: Effectiveness of ACAL and QAL.

#	$\mathcal{R}_{rr}$	$\mathcal{R}_{sr}$	$\mathcal{R}_{dr}$	$\mathcal{R}_{cr}$	$\mathcal{R}_{gr}$	$\mathcal{R}_{car}$	Distorted datasets		Specific distortions in KADID		
							LIVEC	AGIQA-3k	Gaussian blur	Darken	Color shift
1	✗	✗	✗	✗	✗	✗	0.792/0.785	0.803/0.844	0.804/0.826	0.829/0.841	0.816/0.813
2	✓	✗	✗	✗	✗	✗	0.853/0.853↑	0.849/0.893↑	0.863/0.875↑	0.861/0.869↑	0.872/0.869↑
3	✗	✓	✗	✗	✗	✗	0.819/0.825↑	0.809/0.868↑	0.819/0.828↑	0.837/0.844↑	0.832/0.839↑
4	✗	✗	✓	✗	✗	✗	0.810/0.812↑	0.811/0.839↑	0.837/0.829↑	0.831/0.841↑	0.820/0.840↑
5	✗	✗	✗	✓	✗	✗	0.847/0.859↑	0.850/0.896↑	0.808/0.832↑	0.840/0.843↑	0.818/0.829↑
6	✗	✗	✗	✗	✓	✗	0.832/0.847↑	0.857/0.892↑	0.842/0.838↑	0.859/0.852↑	0.849/0.848↑
7	✗	✗	✗	✗	✗	✓	0.826/0.819↑	0.819/0.853↑	0.812/0.818↓	0.830/0.842↑	0.820/0.815↑
8	✗	✗	✗	✓	✓	✗	0.863/0.871↑	0.864/0.904↑	0.859/0.865↑	0.869/0.878↑	0.860/0.881↑
9	✗	✗	✗	✓	✓	✓	0.869/0.862↑	0.872/0.906↑	0.864/0.867↑	0.875/0.883↑	0.862/0.881↑
10	✓	✓	✓	✗	✗	✗	0.874/0.881↑	0.870/0.902↑	0.872/0.876↑	0.882/0.906↑	0.890/0.887↑
11	✓	✓	✓	✓	✓	✓	0.890/0.914↑	0.881/0.915↑	0.897/0.916↑	0.910/0.925↑	0.908/0.892↑

Datasets and Metrics. Experiments are conducted on 4 natural scenario datasets, namely 2 synthetic distortion datasets, KADID (Lin et al. (2019)) and TID2013 (Ponomarenko et al. (2013)), and 2 authentic distortion datasets, LIVEC (Ghadiyaram & Bovik (2015)) and BID (Ciancio et al. (2010)), and 2 AI-generated (AIGC) datasets, namely AGIQA-3k (Li et al. (2023)) and AGIQA-20k (Li et al. (2024b)). The SRCC and PLCC were employed to evaluate the performance of methods.

Method Details. During training the IQA model, a 224×224 patch from each image is randomly cropped and flipped as input, while whole images are input of VLM. The batch number of reinforcement learning in VLM is set to 4. The learning rate is set to 1×10^{-5} and multiplied by 0.8 every 10 epochs. Training is completed on 2 NVIDIA RTX 4090 GPUs with PyTorch 1.12.

4.2 Performance Comparisons

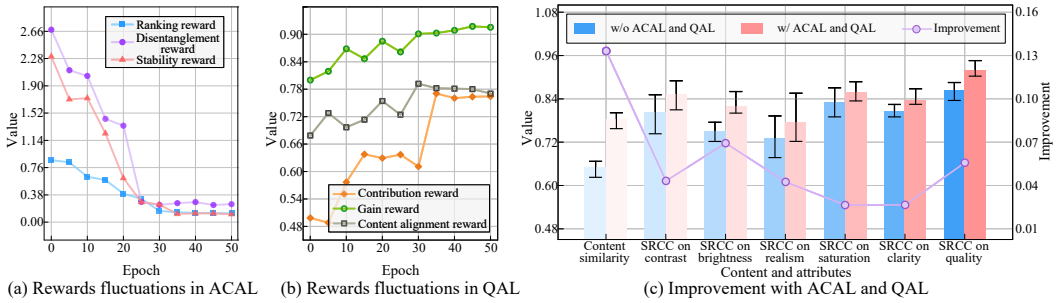
To demonstrate the performance in IQA tasks, ACQA was compared with 15 IQA methods, including 3 handcrafted methods, namely BRISQUE (Mittal et al. (2012a)), NIQE (Mittal et al. (2012b)) and ILNIQE (Zhang et al. (2015)), 6 classic deep learning-based (DL) IQA methods, namely MUSIQ (Ke et al. (2021)), DBCNN (Zhang et al. (2023)), Re-IQA (Saha et al. (2023)), MCOLE (Jiang et al. (2024)), DGQA (Li et al. (2024a)), and GDCIQA (Zhu et al. (2026)), and 6 VLM-based methods, namely Q-Align (Wu et al. (2024)), VQ-R1 (Wu et al. (2025)), EvoQuality (Wen et al. (2026)), Refine-IQA-S2 (Jia et al. (2026)), RALI (Zhao et al. (2026)), and Q-Insight (Li et al. (2025)). Given that methods such as Refine-IQA-S2 are proposed for special tasks like inference-free scenarios, we retain the experimental results reported in their corresponding original papers. The best and second-best results are shaded in dark pink and light pink. The experimental results demonstrate that ACQA achieves the best performance across all datasets. On LIVEC, ACQA improves by 1.5% and 2.0%, respectively. On AGIQA-20k, ACQA improves by 4.0% in SRCC over the runner-up. The largest improvements appear on AGIQA-3k, where ACQA improves by 4.0% in SRCC and 4.5% in PLCC. In terms of overall performance, ACQA also achieves the highest average scores. Therefore, ACQA has advantages over existing IQA methods.

4.3 Ablation Study

This section includes experiments for demonstrating the effectiveness of **dual-branch VLM, ACAL, and QAL**, respectively. The best and second-best results are shaded in dark and light pink.

Effectiveness of dual-branch VLM. To validate the effectiveness of dual-branch VLM, Table 2 is illustrated. In Table 2, scheme #1 removes dual-branch VLM and utilizes image encoder of $\mathcal{M}^{IQA}$ for quality prediction. Scheme #2 and scheme #3 remove one of content and attribute branch. Experimental results demonstrate that the performance improves by 8.23% and 6.89% over the baseline With dual-branch in datasets. Therefore, dual-branch VLM is effective to addressing IQA tasks.

Effectiveness of ACAL and QAL. To validate the effectiveness of ACAL and QAL, Table 3 is illustrated. In Table 3, scheme #1 removes all rewards of ACAL and QAL, which only employ the frozen $\mathcal{M}^{vlm}$ without reinforcement learning. The scheme #2 to #7 utilizes one of rewards in ACAL and QAL. The scheme #8 utilizes rewards related to attribute in QAL, while scheme # 9 and #10 utilize one of the whole ACAL or QAL. Experimental results demonstrate that the performance improves by 8.34% and 8.59% over the baseline by employing ACAL and QAL, respectively. Only Scheme 7 exhibits a slight decrease on AGIQA-3k. Therefore, ACAL and QAL can effectively optimize content and attribute information for IQA tasks.



(a) Rewards fluctuations in ACAL (b) Rewards fluctuations in QAL (c) Improvement with ACAL and QAL

Figure 3: The illustration of reward fluctuations and changes in semantic information.

4.4 Analysis of ACAL and QAL

This section includes experiments for analyzing the **reward fluctuations**, **content and attribute change** to demonstrate the optimization advantages of ACAL and QAL.

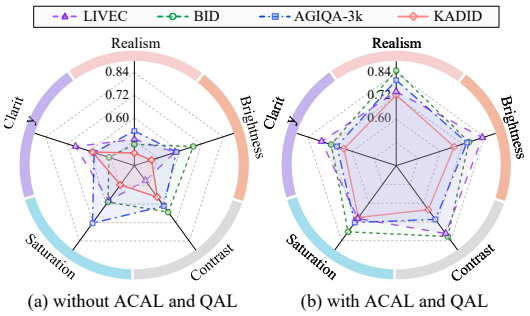
Analysis of Reward Fluctuations. To validate the rationality of the reward fluctuation, Figure 3 (a) and (b) is illustrated. Figure 3 (a) and (b) record the value fluctuations of different reward in ACAL and QAL over 50 training epochs in the TID2013 dataset. Experimental results demonstrate that while the initial epochs exhibit significant volatility across all rewards, the values gradually converge and stabilize as training progresses. The reward in the left chart decrease towards a steady baseline, whereas those in the right chart show an upward trend before plateauing. Therefore, the reward can effectively balances the VLM to obtain content and attribute information for IQA tasks.

Analysis of Content and Attributes Change. To validate the effectiveness of analysis description and attribute scores in IQA tasks, Figure 3 (c) is illustrated. Figure 3 (c) records a comparative analysis of the model performance with and without ACAL and QAL in TID2013 dataset, including content similarity and various attributes accuracy in SRCC. Experimental results demonstrate that after applying ACAL and QAL, significant improvements are achieved across all metrics. Therefore, the analysis description and attribute scores can better serve IQA tasks.

4.5 Visual Analysis

This section includes experiments for analyzing the **stability** of ACQA, and **VLM performance** in obtain content description and attribue scores.

Analysis of Stability. To validate the stability of employing attribute contribution in ACQA, Figure 4 is illustrated. Figure 4 presents radar charts depicting the attribute contribution on four distinct datasets, namely LIVEC, BID, AGIQA-3k, and KADID. Experimental results demonstrate that the contributions of the five attributes are significantly improved by applying ACAL and QAL. Especially on the KADID dataset, the mean contribution of five attributes to quality saw a boost of 0.28 by employing ACAL and QAL. On other datasets, the average attribute contribution improved by 68.5% with ACAL and QAL. And, the contributions across different datasets are consistent.



(a) without ACAL and QAL (b) with ACAL and QAL

Figure 4: The illustration of attribute contribution.

Analysis of VLM Performance. To validate the stability of employing attribute contribution in ACQA, Figure 5 is illustrated. Figure 5 reports the effects of applying ACAL and QAL on two images from the VLM evaluation. Each row shows: the input image with its MOS, content description produced without and with ACAL/QAL, and predicted attributes and quality without and with ACAL/QAL. Across both samples, the corrected descriptions reduce some erroneous elements which are emphasized in red and generally adjust the quality scores for accuracy, e.g., from 4.19 to 3.41 and from 5.23 to 3.07, suggesting a trade-off between descriptiveness and perceptual quality. This highlights the need to balance factual correction with perceptual fidelity in IQA tasks.

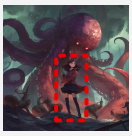
Images and MOS	Content description without ACAL and QAL	Content description with ACAL and QAL	Prediction without ACAL and QAL	Prediction with ACAL and QAL
 MOS: 3.52	A towering eldritch octopus beast looms over rugged dark rocks. Small horned creatures scatter nearby, intense dramatic lighting casts deep shadows across this haunting fantasy scene.	A calm young girl stands on dark rocks before a massive eldritch octopus monster with glowing red eyes. Misty eerie jungle backdrop, small horned beasts around, dramatic lighting.	Contrast: 1.00 Brightness: 2.00 Realism: 3.00 Saturation: 3.00 Blur: 4.00 Quality: 4.19	Contrast: 1.00 Brightness: 0.00 Realism: 1.00 Saturation: 1.00 Blur: 4.00 Quality: 3.41
 MOS: 2.73	Two mirrored views show an armored figure seemingly standing before a mirror. Shiny red-blue tokusatsu suit, grid helmet, reflecting identical costume against simple plain wall.	A full-body shot shows two identical retro hero suits side by side. Glossy red-and-blue armor, grid-pattern helmet with a small spike, standing upright against plain beige indoor wall.	Contrast: 4.00 Brightness: 4.00 Realism: 0.00 Saturation: 2.00 Blur: 4.00 Quality: 5.23	Contrast: 3.00 Brightness: 4.00 Realism: 0.00 Saturation: 1.00 Blur: 3.00 Quality: 3.07

Figure 5: The illustration of VLM Performance.

5 Conclusion and Discussion

Conclusion. In this paper, we presented ACQA, a novel image quality assessment framework that bridges the gap between semantic analysis and quality regression. By moving beyond traditional score-only supervision, we introduced the ACAL and QAL modules to explicitly model the relationships between attributes and quality fluctuations. Experiments demonstrate that VLM optimized by the ACAL and QAL can significantly improve both accuracy and generalization in IQA tasks, which can offer reference for IQA research related to VLMs.

Limitations and Future Directions. Despite the promising results, our work has limitations that suggest clear directions for future research. Firstly, to enhance the generalization of attribute analysis, we will explore automated strategies for quality-aware attribute selection across scenarios. Second, considering the computational cost incurred by the multiple optimization steps of the dual-branch VLM, we will further investigate cross-domain strategies to improve efficiency. Finally, an adaptive prompt optimization framework will be designed to refine VLM outputs by incorporating model performance feedback.

REFERENCES

Shuo Cao, Nan Ma, Jiayang Li, Xiaohui Li, Lihao Shao, Kaiwen Zhu, Yu Zhou, Yuandong Pu, Jiarui Wu, Jiaquan Wang, et al. Artimuse: Fine-grained image aesthetics assessment with joint scoring and expert-level understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15313–15322, 2026.

SaiTeja Chopparapu, Gowthami Chopparapu, and Divija Vasagiri. Enhancing visual perception in real-time: A deep reinforcement learning approach to image quality improvement. *Engineering, Technology & Applied Science Research*, 14(3):14725–14731, 2024.

Alexandre Ciancio, Eduardo AB da Silva, Amir Said, Ramin Samadani, Pere Obrador, et al. No-reference blur assessment of digital pictures based on multifeature classifiers. *IEEE Transactions on Image Processing*, 20(1):64–75, 2010.

Xiaohan Fang, Peilin Chen, Meng Wang, and Shiqi Wang. Examining the role of compression in influencing ai-generated image authenticity. *Scientific Reports*, 15(1):12192, 2025.

Deepti Ghadiyaram and Alan C Bovik. Massive online crowdsourced study of subjective and objective picture quality. *IEEE Transactions on Image Processing*, 25(1):372–387, 2015.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, et al. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081):633–638, 2025a.

Yanjiang Guo, Jianke Zhang, Xiaoyu Chen, Xiang Ji, Yen-Jen Wang, Yucheng Hu, and Jianyu Chen. Improving vision-language-action model with online reinforcement learning. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 15665–15672. IEEE, 2025b.

Yipo Huang, Leida Li, Pengfei Chen, Haoning Wu, Weisi Lin, and Guangming Shi. Multi-modality multi-attribute contrastive pre-training for image aesthetics computing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(2):1205–1218, 2024.

Ziheng Jia, Jiaying Qian, Zicheng Zhang, Zijian Chen, and Xiongkuo Min. Refine-iqa: Multi-stage reinforcement finetuning for perceptual image quality assessment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 22301–22309, 2026.

Qiuping Jiang, Yuese Gu, Chongyi Li, Runmin Cong, and Feng Shao. Underwater image enhancement quality evaluation: Benchmark dataset and objective metric. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(9):5959–5974, 2022.

Qiuping Jiang, Xiao Yi, Li Ouyang, Jingchun Zhou, and Zhihua Wang. Towards dimension-enriched underwater image quality assessment. *IEEE Transactions on Circuits and Systems for Video Technology*, 2024.

Xin Jin, Xinning Li, Hao Lou, Chenyu Fan, Qiang Deng, Chaoen Xiao, Shuai Cui, and Amit Kumar Singh. Aesthetic attribute assessment of images numerically on mixed multi-attribute datasets. *ACM Transactions on Multimedia Computing, Communications and Applications*, 18(3s):1–16, 2023.

Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5148–5157, 2021.

Aobo Li, Jinjian Wu, Yongxu Liu, and Leida Li. Bridging the synthetic-to-authentic gap: Distortion-guided unsupervised domain adaptation for blind image quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 28422–28431, 2024a.

Chunyi Li, Zicheng Zhang, Haoning Wu, Wei Sun, Xiongkuo Min, Xiaohong Liu, Guangtao Zhai, and Weisi Lin. Agiqa-3k: An open database for ai-generated image quality assessment. *IEEE Transactions on Circuits and Systems for Video Technology*, 2023.

Chunyi Li, Tengchuan Kou, Yixuan Gao, Yuqin Cao, Wei Sun, Zicheng Zhang, Yingjie Zhou, Zhichao Zhang, Weixia Zhang, Haoning Wu, et al. Aigiqa-20k: A large database for ai-generated image quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6327–6336, 2024b.

Weiqi Li, Xuanyu Zhang, Shijie Zhao, Yabin Zhang, Junlin Li, Li Zhang, and Jian Zhang. Q-insight: Understanding image quality via visual reinforcement learning. In *Advances in Neural Information Processing Systems*, 2025.

Hanhe Lin, Vlad Hosu, and Dietmar Saupe. Kadid-10k: A large-scale artificially distorted iqa database. In *2019 Eleventh International Conference on Quality of Multimedia Experience (QoMEX)*, pp. 1–3. IEEE, 2019.

Mushui Liu, Yuhang Ma, Zhen Yang, Jun Dan, Yunlong Yu, Zeng Zhao, Zhipeng Hu, Bai Liu, and Changjie Fan. Llm4gen: Leveraging semantic representation of llms for text-to-image generation. In *Proceedings of the AAAI conference on Artificial Intelligence*, volume 39, pp. 5523–5531, 2025.

Yiting Lu, Xin Li, Jianzhao Liu, and Zhibo Chen. Styleam: Perception-oriented unsupervised domain adaption for no-reference image quality assessment. *IEEE Transactions on Multimedia*, 2024.

Weijia Mao, Hao Chen, Zhenheng Yang, and Mike Zheng Shou. The image as its own reward: Reinforcement learning with adversarial reward for image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5999–6009, 2026.

Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on Image Processing*, 21(12):4695–4708, 2012a.

Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a "completely blind" image quality analyzer. *IEEE Signal Processing Letters*, 20(3):209–212, 2012b.

Nikolay Ponomarenko, Oleg Ieremeiev, Vladimir Lukin, Karen Egiazarian, Lina Jin, Jaakko Astola, Benoit Vozel, Kacem Chehdi, Marco Carli, Federica Battisti, et al. Color image database tid2013: Peculiarities and preliminary results. In *European workshop on Visual Information Processing (EUVIP)*, pp. 106–111. IEEE, 2013.

Avinab Saha, Sandeep Mishra, and Alan C Bovik. Re-iqa: Unsupervised learning for image quality assessment in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5846–5855, 2023.

Zhenchen Tang, Zichuan Wang, Bo Peng, and Jing Dong. Clip-agiqa: boosting the performance of ai-generated image quality assessment with clip. In *International Conference on Pattern Recognition*, pp. 48–61. Springer, 2024.

Puyi Wang, Wei Sun, Zicheng Zhang, Jun Jia, Yanwei Jiang, Zhichao Zhang, Xiongkuo Min, and Guangtao Zhai. Large multi-modality model assisted ai-generated image quality assessment. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 7803–7812, 2024.

Wen Wen, Tianwu Zhi, Kanglong Fan, Yang Li, Xinge Peng, Yabin ZHANG, Yiting Liao, Junlin Li, and Li zhang. Self-evolving vision-language models for image quality assessment via voting and ranking. In *International Conference on Learning Representations*, pp. 39444–39457, 2026.

Haoning Wu, Zicheng Zhang, Weixia Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Yixuan Gao, Annan Wang, Erli Zhang, Wenxiu Sun, Qiong Yan, Xiongkuo Min, Guangtao Zhai, and Weisi Lin. Q-align: Teaching lmms for visual scoring via discrete text-defined levels. In *International Conference on Machine Learning*, pp. 54015–54029, 2024.

Mingrui Wu, Lu Wang, Pu Zhao, Fangkai Yang, Jianjin Zhang, Jianfeng Liu, Yuefeng Zhan, Weihao Han, Hao Sun, Jiayi Ji, et al. Reprompt: Reasoning-augmented reprompting for text-to-image generation via reinforcement learning. In *International Conference on Learning Representations*, volume 2026, pp. 14030–14057, 2026.

Tianhe Wu, Jian Zou, Jie Liang, Lei Zhang, and Kede Ma. Visualquality-r1: Reasoning-induced image quality assessment via reinforcement learning to rank. In *Advances in Neural Information Processing Systems*, volume 38, pp. 88167–88190, 2025.

Zhiyuan You, Zheyuan Li, Jinjin Gu, Zhenfei Yin, Tianfan Xue, and Chao Dong. Depicting beyond scores: Advancing image quality assessment through multi-modal language models. In *European Conference on Computer Vision*, pp. 259–276. Springer, 2024.

Guanghui Yue, Honglv Wu, Weiqing Yan, Tianwei Zhou, Hantao Liu, and Wei Zhou. Subjective and objective quality assessment of multi-attribute retouched face images. *IEEE Transactions on Broadcasting*, 70(2):570–583, 2024.

Lin Zhang, Lei Zhang, and Alan C Bovik. A feature-enriched completely blind image quality evaluator. *IEEE Transactions on Image Processing*, 24(8):2579–2591, 2015.

Weixia Zhang, Guangtao Zhai, Ying Wei, Xiaokang Yang, and Kede Ma. 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*, pp. 14071–14081, 2023.

Shijie Zhao, Xuanyu Zhang, Weiqi Li, Junlin Li, Tianfan Xue, Jian Zhang, et al. Reasoning as representation: Rethinking visual reinforcement learning in image quality assessment. In *International Conference on Learning Representations*, pp. 101666–101684, 2026.

Ziwei Zheng, Minghao Yang, Jack Hong, Chenxiao Zhao, Guohai Xu, Le Yang, and Chao Shen. Deepeyes: Incentivizing" thinking with images" via reinforcement learning. In *International Conference on Learning Representations*, volume 2026, pp. 126775–126798, 2026.

Shun Zhu, Xichen Yang, Yan Zhang, Tianshu Wang, Zhongyuan Mao, Tianyin Li, Zhuoyan Sun, and Xiaobo Shen. Graph-driven domain co-adaptation for cross-domain image quality assessment. AAAI Conference on Artificial Intelligence, pp. 13970–13978, 2026.

A Appendix

B AI Use Statement

During the preparation of this manuscript, the authors used AI for grammar checking and language refinement of author-written text. All scientific claims, experimental results, core methodologies, and conceptual contributions were independently developed by the human authors. The authors have carefully reviewed, verified, and edited all AI-assisted outputs and assume full responsibility for the accuracy and integrity of the final manuscript.