

DR.Experts: Differential Refinement of Distortion-Aware Experts for Blind Image Quality Assessment

Bohan Fu^{1*}, Guanyi Qin^{2*}, Fazhan Zhang¹, Zihao Huang¹, Mingxuan Li¹, Runze Hu^{1†},

¹ Beijing Institute of Technology, Beijing 100086, China

² National University of Singapore, Singapore 119276, Singapore

{fubohan809, zfz63622, huangzihhhh, limx1630, hrzlpk2015}@gmail.com, guanyi.qin@u.nus.edu

Abstract

Blind Image Quality Assessment, aiming to replicate human perception of visual quality without reference, plays a key role in vision tasks, yet existing models often fail to effectively capture subtle distortion cues, leading to a misalignment with human subjective judgments. We identify that the root cause of this limitation lies in the lack of reliable distortion priors, as methods typically learn shallow relationships between unified image features and quality scores, resulting in their insensitive nature to distortions and thus limiting their performance. To address this, we introduce DR.Experts, a novel prior-driven BIQA framework designed to explicitly incorporate distortion priors, enabling a reliable quality assessment. DR.Experts begins by leveraging a degradation-aware vision-language model to obtain distortion-specific priors, which are further refined and enhanced by the proposed Distortion-Saliency Differential Module through distinguishing them from semantic attentions, thereby ensuring the genuine representations of distortions. The refined priors, along with semantics and bridging representation, are then fused by a proposed mixture-of-experts style module named the Dynamic Distortion Weighting Module. This mechanism weights each distortion-specific feature as per its perceptual impact, ensuring that the final quality prediction aligns with human perception. Extensive experiments conducted on five challenging BIQA benchmarks demonstrate the superiority of DR.Experts over current methods and showcase its excellence in terms of generalization and data efficiency.

Code — <https://github.com/FuBohan01/DR.Experts>

Introduction

High-quality images are essential for reliable performance in vision tasks. Being instructional and serving as a crucial and valuable part in enabling quality control in real-world scenarios, **Blind Image Quality Assessment (BIQA)** aims to evaluate image visual quality without reference images, easing the burden of references that are hard to acquire in in-the-wild cases (Huang et al. 2024; Li et al. 2025; Mittal, Moorthy, and Bovik 2012). Albeit numerous efforts have been made to push the boundaries of BIQA, the advantage

*Equal contribution.

† Corresponding author.

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

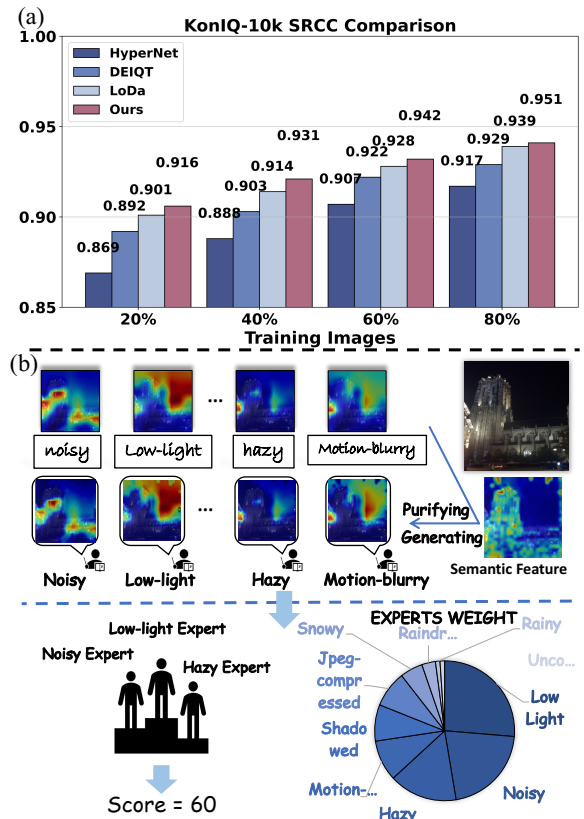


Figure 1: (a): DR.Experts maintains the advantage even with limited training data. (b): Given a distorted image, the proposed framework DR.Experts first leverages a vision-language model specialized on visual distortions to obtain attention corresponding to various distortions. By differentiating these cues from semantic attention, DR.Experts effectively further purifies distortion-aware representations. These refined features are then adaptively weighted according to their perceptual importance and integrated to yield a precise and perceptually consistent quality assessment.

of being reference-free still introduces significant performance challenges, especially in terms of distortion modeling (Moorthy and Bovik 2011). Existing BIQA techniques

still struggle to effectively characterize quality-aware representations, undermining their abilities to align with human evaluations of perceptual quality.

A primary ground of such limitations arises from their reduced sensitivity to subtle distortions and conceptual ambiguity under complicated and diverse distortions due to the lack of guidance from distortion priors. While BIQA has made significant progress (Chen et al. 2024) through the development of intricate learning architectures and advanced training strategies recently, these methods (Yang et al. 2024) still focus on learning a plain relationship between consolidated and unified image features and the final quality scores. However, in-the-wild images often suffer from diverse types and varying levels of distortions (Hosu et al. 2020), such as under-exposure, noise, and motion blur, affecting image visual quality in different ways and extents. Yet, the limited volume of existing BIQA datasets, coupled with a lack of distortion specifications and annotations, constrains the capabilities of models to learn and accurately capture these subtle yet important distortion-related features.

As the challenge in capturing distortion-specific characteristics highlights a fundamental gap in existing BIQA approaches, thus, naturally, a key motivation to advance BIQA is to effectively incorporate prior knowledge of distortions into current learning frameworks (Li et al. 2024a), enabling fine-grained and perceptually aligned modeling of quality-aware representations. Driven by this insight, we propose a **Differential Refinement of Distortion-Aware Experts BIQA** framework, hereby named **DR.Experts**, that incorporates distortion priors to guide the extraction of perceptual visual quality-relevant features, and distill and purify these features further by differentiating them with semantic features. These refined features are then aggregated with adaptive weights as per their respective contributions to the image quality, acting as multiple experts for different distortion perspectives, thus yielding a comprehensive and trustworthy quality prediction, as shown in Fig. 1. Such a design not only enhances the alignment between prediction and human perception but also offers transparency in the decision-making process, as the generated score can be traced back to certain distortion factors, thereby increasing the trustworthiness of the assessment.

To this end, CLIP (Radford et al. 2021) is incorporated into the proposed framework for distortion priors. Serving as a foundation between visuals and text, Contrastive Language-Image Pre-Training (CLIP) enables transferable priors by mapping both modalities into a shared embedding space, where text can prompt distortion-related visual features. Specifically, we employ DA-CLIP (Luo et al. 2024), a derivative optimized for low-level vision tasks and fine-tuned to increase sensitivity to distortion features. DA-CLIP demonstrates strong capabilities for fine-grained distortion perception, it achieves an impressive average accuracy of 99.2% on tasks involving ten distinct distortion types, including blur, low light, and compression errors. Based on this, we obtain the distortion-aware visual attentions activated by DA-CLIP under various distortion-specific prompts as priors by the dot product between the prompt representations from the text encoder and the visual features

from the image encoder. To further strengthen these representations and suppress the attention noise of semantic redundancy originating from upstream classification pre-training of CLIP and Vision Transformer (ViT) (Dosovitskiy et al. 2020), we introduce the **Distortion-Saliency Differential Module (DSDM)**, with its **differential refinement attention mechanism** to isolate distortion-related features. DSDM refines the priors by contrasting DA-CLIP’s distortion-aware attention with the semantic attention extracted by ViT, thereby ensuring a precise quality assessment.

Furthermore, to effectively integrate and leverage features of different types, along with considerations that various distortions impact perceptual quality in different ways and extents, we design the **Dynamic Distortion Weighting Module (DDWM)**, a Mixture-of-Experts (MoE) architecture that adaptively aggregates these feature groups: distortion priors refined by DSDM, semantic features extracted by ViT, and bridging features as intermediate supplementaries derived from the difference between the two aforementioned feature groups. This module not only model complementary cues across diverse representations, and allows to assign adaptive weights to different distortion types based on their respective perceptual impacts. By doing so, our proposed framework aligns more closely with the fine-grained assessment criteria of the human vision system, further enhancing both the accuracy and trustworthiness of the generated quality score. To summarize, our contributions can be regarded as follows:

- We propose a novel BIQA framework that leverages distortion priors as guidance, enabling fine-grained, distortion-aware quality assessment. By leveraging DA-CLIP’s text-prompted transferable visual attention as priors, our framework learns towards trustworthy and explainable results. The proposed framework enhances the perceptual alignment between human-perceived distortion representations and quality metrics.
- A novel module called Distortion-Saliency Differential Module is introduced to refine distortion priors by differentiating DA-CLIP’s distortion-aware attention with ViT-derived semantic attention. DSDM effectively suppresses suppressing redundant semantic noise from pre-training, while enhances the saliency of distortion features.
- Dynamic Distortion Weighting Module is further proposed to dynamically assign significance weights to score token regarding different types and levels of distortions, aiming at simulating a comprehensive analysis as per subject quality assessment of distortion characteristics by specialists.
- We verify the proposed framework on five diverse and challenging BIQA benchmarks, where it consistently outperforms other competitors and showcases strong generalization and data efficiency capabilities.

Related Work

Blind Image Quality Assessment

Conventional BIQA methods mainly use hand-crafted features based on natural scene statistics (Zhang, Zhang, and

Bovik 2015; Wang et al. 2021; Saad, Bovik, and Charrier 2012). However, due to the constrained representation resulting from manual feature selection, these methods exhibit limited generalization in in-the-wild scenes. With the advancement of deep learning, BIQA has seen notable improvements. Current learning-based approaches can be categorized into CNN-based and ViT-based methods (Chu et al. 2025b), extracting image features for straightforward end-to-end regression of quality scores (Zhang et al. 2020; Kang et al. 2014; Bosse et al. 2017). By increasing capacity and depth, models, *e.g.*, ResNet (He et al. 2016) and ViT (Dosovitskiy et al. 2020), pre-trained on large datasets (Deng et al. 2009) are used to further perceive image distortion (Shin, Lee, and Kim 2024; Xu et al. 2023). To overcome the restricted volume of IQA datasets, some self-supervised learning methods conduct contrastive learning on degraded images formed by distortion models (Agnolucci et al. 2024a; Saha, Mishra, and Bovik 2023), to fully perceive the differences and similarities of distortions. At the same time, domain adaptation, multi-scale feature adaptation, and other techniques (Chu et al. 2025a; Hu et al. 2025; Qin et al. 2025) have also been explored to advance BIQA (Su et al. 2020). Recently, multimodal large language models have been widely applied in downstream vision tasks (Wu et al. 2024b; Li et al. 2023), with researchers also exploring their powerful visual understanding and perception capabilities on BIQA by fine-tuning, and thus to promote the interpretability and robustness of BIQA research (Zhou et al. 2025; Wu et al. 2024a).

CLIP for Low-level Vision

Built upon the foundations of large-scale (image, text) pairs and contrastive pre-training, CLIPs (Radford et al. 2021), a series of models, demonstrate strong zero-shot capabilities for transfer learning and have thus been widely applied in downstream vision tasks (Agnolucci et al. 2024b; Liang et al. 2023). CLIP-IQA, serving as the first piece of work, explores the possibility of using CLIP for BIQA through a sequence of carefully designed prompts (Wang, Chan, and Loy 2023). Meanwhile, methods such as multitask learning and self-supervised strategies have been further employed to leverage the knowledge of CLIP in the field of BIQA. (Agnolucci, Galteri, and Bertini 2024; Kwon et al. 2024). DA-CLIP adapted CLIP for low-level vision tasks by introducing a controller to accurately predict the distortion type of the input image, yielding excellent results of 99.2 accuracy (Luo et al. 2024). In this paper, we utilize the prior of specific distortion types obtained by DA-CLIP to guide our proposed model to assess image quality in a fine-grained manner.

Methodology

Overview

In this work, We propose a novel BIQA network that evaluates images precisely under the guidance of the prior of specific distortion types within the image to form a mixture-of-Experts. The overall framework is illustrated in the Fig. 2. Specifically, an RGB image is used as input, and a Vision Transformer is employed as the Image Encoder. Meanwhile,

DA-CLIP is utilized to identify the distortion types present in the image. Subsequently, we introduce a Distortion-Saliency Differential Module to refine different image distortion information by removing noise unrelated to distortion features and enhances the saliency of distortion features. To avoid the influence of non-dominant or absent distortion types in the image on the final quality assessment, we designed the Dynamic Distortion Weighting module to integrate the semantic feature from Image Encoder, distortion priors refined by DSDM, and supplementary interval features derived from the difference between them. The comprehensive feature group formed by the above three features assigns dynamic weights to experts of different distortion types and obtains the final image quality score.

Distortion-specific Prior

In most previous works, data-driven approaches have been heavily relied upon. However, we introduce prior knowledge of distortion types, enabling the model to extract features in a more targeted manner and thereby reducing excessive dependence on data-driven methods. In this work, we adopt DA-CLIP as the Distortion-specific Prior Module. Built upon the frozen CLIP text and Image Encoder, DA-CLIP trains an additional image controller through contrastive learning to predict high-quality image distortion feature embeddings. CLIP is a multimodal model that aligns images and text in a shared embedding space, widely used in cross-modal retrieval and understanding tasks. DA-CLIP is designed to generate features that match the actual distortion type of the input image, making it better suited to the needs of degraded scenarios. Specifically, we use DA-CLIP’s feature embeddings for image distortion types as prior knowledge for image quality assessment tasks. DA-CLIP maps the input RGB image $I \in \mathbb{R}^{H \times W \times C}$ and ten predefined common image distortion types $T = [\text{motion-blurry, hazy, jpeg-compressed, low-light, noisy, raindrop, rainy, shadowed, snowy, uncompleted}]$, into the same feature space through the image controller \mathcal{E}_D and text encoder \mathcal{E}_T , respectively. Specifically, the degraded image representation is expressed as $E_{dis} = \mathcal{E}_D(I)$, while the text representation for the i -th distortion type is given by $E_T^i = \mathcal{E}_T(T^i)$. We use the Hadamard product similar to CLIP-IQA (Wang, Chan, and Loy 2023) to obtain the representation of the image under different distortion categories. This process can be formulated as:

$$F_D^i = \text{Linear}^i (E_{dis} \odot E_T^i). \quad (1)$$

Distortion-Saliency Differential Module

Inspired by the Differential Transformer (Ye et al. 2025), which effectively suppresses noise in homogenous attention through a differential mechanism, we extend this approach to heterogeneous attention and propose the Differential Refinement Attention Mechanism. This mechanism is capable of suppressing redundant or overlapping information, including attention noise and certain high-level semantic information, between distortion prior knowledge and the semantic information of the Image Encoder. Furthermore, it

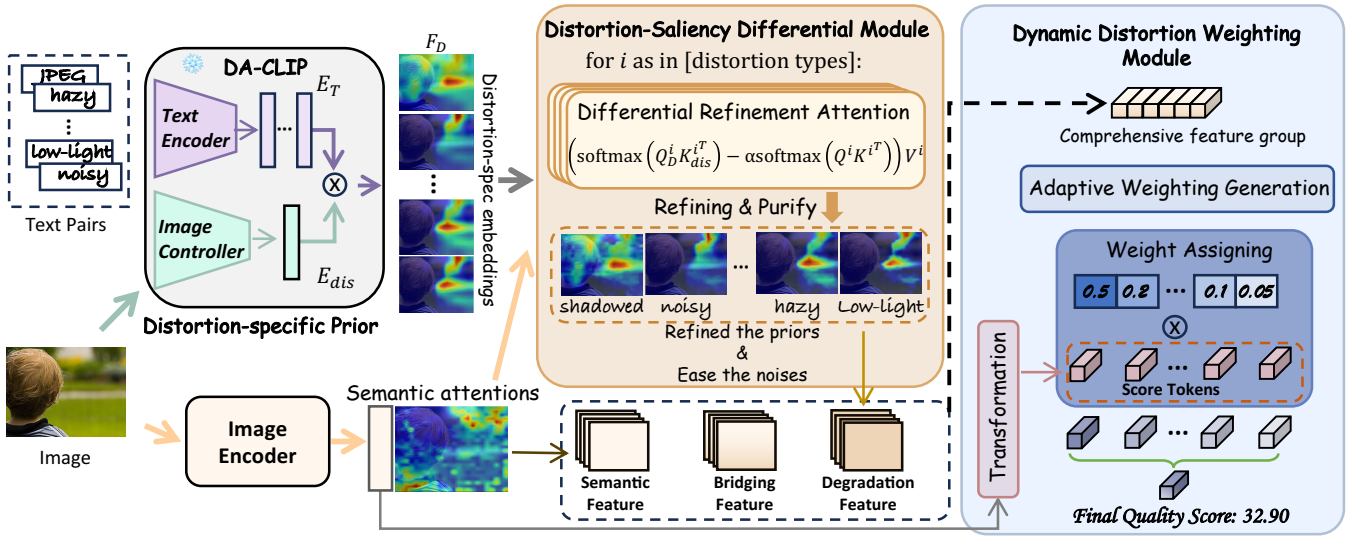


Figure 2: Overall architecture of the proposed DR.Experts. We leverage DA-CLIP to obtain priors and use DSDM to refine the attentions. DDWM then, serving as experts, to weigh the importance of distortions and give final predictions.

enhances the prior knowledge related to distortions. Specifically, an RGB image is processed by a Vision Transformer to extract high-level semantic features $F \in \mathbb{R}^{N \times E}$, where N is the sequence length and E is the embedding dimension. Simultaneously, the RGB image I is processed by DA-CLIP to obtain features $F_D^i \in \mathbb{R}^{N \times E}$, where $i \in [1, 2, \dots, 10]$, corresponding to a specific distortion type. F_D^i , proposed as the degraded category information query in DA-CLIP, aims to extract features related to specific types of distortion from the image features E_{dis} generated by the image controller of DA-CLIP. At this stage, the distortion features contain a small amount of semantic information and attention noise. The query and key for the semantic features extracted by Image Encoder are defined as $Q^i = W_Q^i F$ and $K^i = W_K^i F$, while the query and key for the features of i -th distortion type are defined as $Q_D^i = W_{DQ}^i F_D^i$ and $K_{dis}^i = W_{disK}^i E_{dis}$. Moreover, the value vector is computed as $V^i = W_V^i \text{Cat}([F, F_D^i])$. Using these, the differential refinement attention mechanism extracts the image features under the i -th distortion type as follows:

$$F_{distortion}^i = (\text{softmax}(Q_D^i K_{dis}^{iT}) - \alpha \text{softmax}(Q^i K^{iT})) V^i, \quad (2)$$

where, $F_{distortion}^i$ is the feature of the image under the i -th distortion type. α is a learnable parameter. As a result, the module isolates and refine the features that correspond exclusively to a specific distortion type.

Subsequently, the features of all distortion types are fed into a Feed-Forward Network (FFN) composed of two linear layers with GELU as the activation function, thereby generating expert knowledge for identifying different distortion features.

$$F_{multi-distortion} = \text{FFN}(F_{distortion}^1, \dots, F_{distortion}^i), \quad (3)$$

where, $\text{FFN}(\cdot)$ is the Feed-Forward Network. Benefiting from the Differential Refinement Attention Mechanism, this module can effectively suppress redundant or overlapping information, accurately extract features related to different distortion types, and further integrate and optimize multi-distortion features through the FFN, thereby generating expert-level representations for image distortions.

Dynamic Distortion Weighting Module

To assign distortion weights based on the impact of image distortions on perceived image quality and thereby achieve more reliable quality assessment scores, we propose a hybrid expert system named the Dynamic Distortion Weighting Module (DDWM). This module first utilizes the semantic features F extracted by the Image Encoder, and the distortion priors $F_{Group} = (1 - \lambda) \cdot F_{multi-distortion}$, as refined by the DSDM with a learnable parameter λ to balance the contributions to the following modeling. Supplementary interval features, named by bridging feature $F_{bridging} = F - F_{Group}$ are further incorporated into this module to form a comprehensive feature group $F_{com} = \text{Cat}([F_{distortion}^i, F, F_{bridging}])$ for multidimensional quality evaluation. The supplementary interval features serve to bridge distributional differences between dimensions or levels of features, ensuring a more cohesive representation.

The supplementary interval features serve to bridge distributional differences between dimensions or levels of features, ensuring a more cohesive representation.

Based on this feature group, the Adaptive Weighting Generation (WG) module dynamically calculates the importance of each distortion type in terms of its impact on quality perception, *i.e.*, the dynamic expert weight for each distortion type. This process could be described as:

$$W_{distortion}^1, \dots, W_{distortion}^{10} = \text{WG}(F_{com}). \quad (4)$$

Here, Weighting Generation refers to a multi-layer percep-

tron (MLP) with PReLU as the activation function. Finally, the computed weights are combined with the token T_{score} derived from the ViT’s class token through a linear transformation, representing the current image quality. The score weight is employed to optimize and adjust the final image quality score, which can be written as:

$$\text{Score} = \sum_{i=1}^{10} W_{distortion}^i T_{score} \quad (5)$$

Experiments

Benchmark Datasets

We assess the performance of the proposed DR.Experts on five in-the-wild BIQA datasets collected from real-world scenarios. The BID (Ciancio et al. 2010) and LIVEC (Ghadiyaram and Bovik 2015) datasets comprise 590 and 1,162 images taken by photographers using mobile devices, respectively. KonIQ-10k (Hosu et al. 2020) contains 10,073 images selected from publicly available multimedia sources. SPAQ (Fang et al. 2020) includes 11,125 images captured with smartphone cameras, covering diverse content and resolutions. LIVEFB (Ying et al. 2020), the largest authentic dataset to date, consists of 39,810 images.

Two widely used BIQA metrics are adopted as performance evaluation metrics, including Spearman’s Rank-order Correlation Coefficient (SRCC) and Pearson’s Linear Correlation Coefficient (PLCC), to measure the effectiveness of DR.Experts. SRCC evaluates the monotonic relationship between predicted and ground-truth scores, while PLCC reflects the prediction accuracy of models. Both metrics range from 0 to 1, with higher values indicating better performance. Ideally, achieving values closer to 1 demonstrates superior prediction monotonicity and accuracy.

Implementation Details

For our approach, we apply a standard data augmentation technique commonly used in BIQA, where each image is randomly cropped into smaller patches. In addition, we follow the same experimental settings as QPT (Zhao et al. 2023) and QCN (Shin, Lee, and Kim 2024). The number of patches is adjusted based on the size of each dataset. The Image Encoder in our method is the small version of DeiT-III (Touvron, Cord, and Jégou 2022), pre-trained on ImageNet with a weight decay of 0.05, a batch size of 1024, and trained for 400 epochs. As for the prior knowledge module, we utilize the pre-trained DA-CLIP image controller and text encoder (Luo et al. 2024). This module is frozen during training.

DR.Experts is then fine-tuned on BIQA datasets for 9 epochs with the prior module frozen. The learning rate starts at 2×10^{-4} and is reduced by a factor of 10 after every 3 epochs. The Smooth L1 loss is utilized as the loss function for model training, and the batch size is determined by the scale of the datasets, *e.g.*, 64 for the LIVEC dataset and 156 for the KonIQ dataset. For all datasets, 80% of the images are splitted and used for training, while the remaining 20% are reserved for testing. We repeat this train-test setting 10 times to mitigate the splitting bias, and record the medians

of the results. We use PyTorch to implement DR.Experts, and all experiments are conducted using 4 RTX 4090 GPUs.

Comparison with SOTAs

In this experiment, 14 representative BIQA methods are involved and compared with DR.Experts in terms of prediction performance on the in-the-wild datasets. Among them, DIIVINE (Zhang, Zhang, and Bovik 2015) and BRISQUE (Mittal, Moorthy, and Bovik 2012) are based on handcrafted perceptual features, while the remaining methods are based on learning architectures, *e.g.*, DB-CNN (Zhang et al. 2020) and HyperIQA (Su et al. 2020) utilize CNNs, while MUSIQ (Ke et al. 2021) and DEIQT (Qin et al. 2023) leverage vision transformers. Additionally, we included TRoS (Golestaneh, Dadsetan, and Kitani 2022) and LODA (Xu et al. 2024), methods that combine the strengths of both CNNs and ViT, as well as the LQMamba (Guan et al. 2025) based on the vision mamba architecture.

As shown in Table 1, DR.Experts exhibits significant improvements across five diverse BIQA datasets involving a range of complex scenes and levels of distortion. Given that achieving robust performance on these datasets is inherently challenging due to their varied nature, and compared to the SOTA approaches such as QFM-IQM (Li et al. 2024b), LODA, LQMamba and QCN (Shin, Lee, and Kim 2024), DR.Experts achieves at least a 0.75- and a 0.95-point improvement on the KonIQ-10k dataset in terms of the SRCC and PLCC metrics, respectively. Our design of leverage distortion priors are further confirm. Additionally, DR.Experts outperforms methods that leverage quality-pretraining techniques, such as QPT (Zhao et al. 2023) and CONRTIQUE (Madhusudana et al. 2022), further demonstrating its robustness and adaptability. These results firmly emphasize the capabilities of our DR.Experts to effectively utilize and refine distortion priors and characteristics, particularly in the context of challenging in-the-wild scenarios.

Generalization Validation

We conducted a cross-dataset experiment to verify the generalization ability of the model, which is an essential indicator of BIQA practices, where a model is trained and tested on different datasets. The results are shown in Table 2. In short, across all different training and test pairs, DR.Experts has achieved better results in terms of the SRCC metric compared to previous SOTA methods. This result also verifies the core idea of leveraging prior knowledge of distortion to help enhance the robustness and generalization.

Data Efficiency Validation

Benefiting from priors, together with the quality-aware attention further refined through DSDM, our model could quickly attend to distortions and thus alleviate its performance reliance on large-scale data. To verify this, we conducted data efficiency validation by reducing the training set data to 60%, 40%, and 20% of the original size, while maintaining the same model architecture and training settings, and compared it with existing models. As shown in Table 3, under the 60% mode, DR.Experts are 2.3- and 3-point higher than LoDa on the LIVEC dataset in terms of

Methods	KonIQ		LIVEC		SPAQ		LIVEFB		BID	
	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC
ILNIQE (Zhang, Zhang, and Bovik 2015)	0.503	0.496	0.453	0.511	0.719	0.654	0.219	0.255	0.495	0.454
BRISQUE (Mittal, Moorthy, and Bovik 2012)	0.715	0.702	0.601	0.621	0.802	0.806	0.320	0.356	0.574	0.540
WaDIQaM-NR (Bosse et al. 2017)	0.729	0.754	0.692	0.730	0.840	0.845	0.435	0.430	0.653	0.636
DB-CNN (Zhang et al. 2020)	0.878	0.887	0.844	0.862	0.910	0.913	0.554	0.652	0.845	0.850
HyperIQA (Su et al. 2020)	0.906	0.917	0.859	0.882	0.916	0.919	0.535	0.623	0.869	0.878
MUSIQ (Ke et al. 2021)	0.916	0.928	0.702	0.746	0.917	0.921	–	–	0.646	0.739
TReS (Golestaneh, Dadsetan, and Kitani 2022)	0.915	0.928	0.846	0.877	–	–	0.554	0.625	–	–
DEIQT (Qin et al. 2023)	0.921	0.934	0.875	0.894	0.919	0.923	0.571	0.663	–	–
CONRTIQUE* (Madhusudana et al. 2022)	0.894	0.906	0.845	0.857	0.914	0.919	0.580	0.641	–	–
QPT* (Zhao et al. 2023)	0.927	0.941	0.895	0.914	0.925	0.928	0.578	0.675	0.888	0.911
QCN (Shin, Lee, and Kim 2024)	0.934	0.945	0.875	0.893	0.923	0.928	–	–	0.892	0.890
QFM-IQM (Li et al. 2024b)	0.922	0.936	0.891	0.913	0.920	0.924	0.567	0.667	–	–
LODA (Xu et al. 2024)	0.932	0.944	0.876	0.899	0.925	0.928	0.578	0.679	–	–
LQMamba (Guan et al. 2025)	0.928	0.943	0.863	0.903	0.927	0.933	0.574	0.672	–	–
DR.Experts (Ours)	0.941	0.954	0.914	0.926	0.928	0.933	0.585	0.690	0.896	0.919

Table 1: Performance comparison measured by medians of SRCC and PLCC, where bold entries indicate the top two results. Pre-training methods are marked with the *.

TRAINING	LIVEFB		LIVEC	KonIQ
TESTING	KonIQ	LIVEC	KonIQ	LIVEC
DBCNN	0.716	0.724	0.754	0.755
P2P-BM	0.755	0.738	0.740	0.770
HperIQA	0.758	0.735	0.772	0.785
TReS	0.713	0.740	0.733	0.786
DEIQT	0.733	0.781	0.744	0.794
LODA	0.763	0.805	0.745	0.811
DR.Experts	0.783	0.807	0.785	0.841
Gains	0.020(↑)	0.002(↑)	0.040(↑)	0.030(↑)

Table 2: SRCC on the generalization validation. The best performance is highlighted in bold. Gains are calculated versus the second-best performance.

the SRCC and PLCC metrics, respectively. Notably, under the 20% mode, DR.Experts surpasses with the largest performance improvement against other methods. Those results further confirm the good robustness and generalization ability of DR.Experts and its core idea of priors and refinement.

Ablation Study

Module Study We conducted ablation experiments to evaluate the effectiveness of the proposed Distortion-Saliency Differential Module and Dynamic Distortion Weighting Module. As shown in Table 4, the *Image Encoder* represents the performance when using only ViT as the BIQA backbone, with the results indicating that high-level information for quality assessment is insufficient and thus comes with data scarcity, as that the model performance on the smaller dataset LIVEC is less comparable against large dataset. *DA-CLIP* refers to the approach where only the prior module is involved with the vision and language encoder frozen, followed by a dot-product operation identical to that of CLIPIQA to obtain distortion-specific features.

Mode	Methods	KonIQ		LIVEC	
		SRCC	PLCC	SRCC	PLCC
20%	HyperNet	0.869	0.873	0.776	0.809
	DEIQT	0.888	0.908	0.792	0.822
	LoDa	0.907	0.923	0.815	0.854
	DR.Experts	0.917	0.931	0.837	0.861
40%	HyperNet	0.892	0.908	0.832	0.849
	DEIQT	0.903	0.922	0.838	0.855
	LoDa	0.922	0.935	0.849	0.879
	DR.Experts	0.929	0.942	0.874	0.896
60%	HyperNet	0.901	0.914	0.843	0.862
	DEIQT	0.914	0.931	0.848	0.877
	LoDa	0.928	0.940	0.869	0.891
	DR.Experts	0.939	0.950	0.899	0.914

Table 3: Data-efficient learning validation with the training set containing 20%, 40% and 60% images. Bold entries indicate the best performance.

We only trained the prediction head with these features to regress scores. The results demonstrate that the features extracted by DA-CLIP are not directly applicable to BIQA, as certain distortion types, *e.g.*, rain-drop and uncompleted, are either irrelevant or detrimental to visual quality. These results are also broadly consistent with the observations in CLIPIQA. Based on these features, *DSDM* is introduced to improve the priors and purify the attentions, thereby mitigating noise interference in BIQA and enhancing the final performance, especially on smaller datasets. Furthermore, *DR.Experts* means the inclusion of the proposed DDWM module, which leverages a comprehensive feature group to differentiate and optimize the contributions of various distortions, *i.e.*, levels and types, to quality scores, further enhancing the metrics on both the KonIQ and LIVEC datasets.

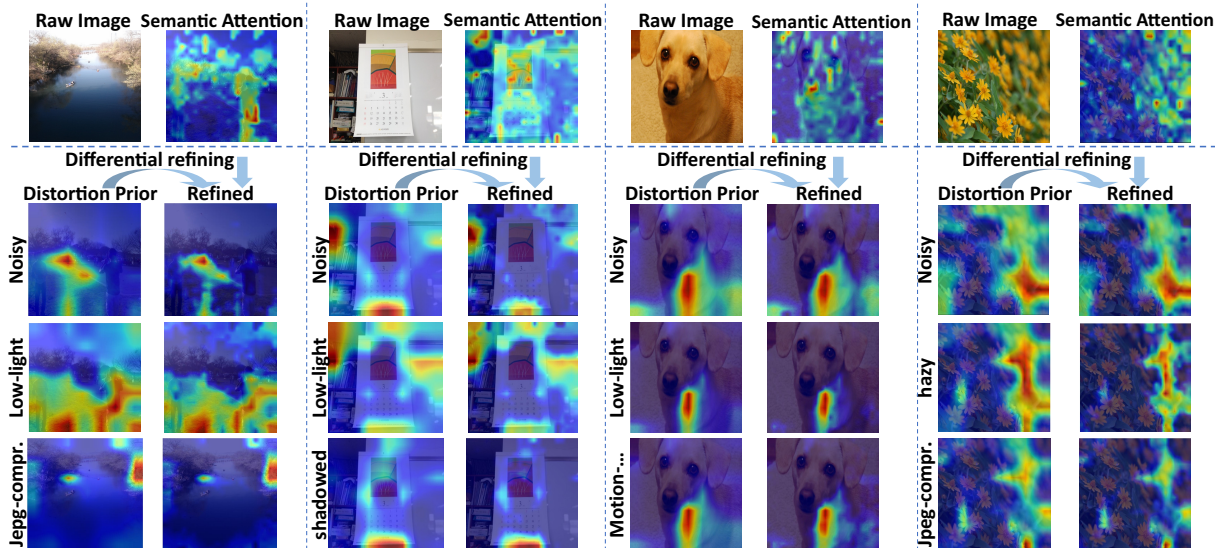


Figure 3: Comparison of attention maps from DA-CLIP, the image encoder, and DSDM outputs.

Methods	KonIQ		LIVEC	
	SRCC	PLCC	SRCC	PLCC
Image Encoder	0.916	0.929	0.857	0.884
DA-CLIP	0.720	0.754	0.587	0.635
DSDM	0.930	0.941	0.885	0.904
DR.Experts	0.941	0.954	0.914	0.926

Table 4: Ablation experiments on LIVEC and KonIQ datasets. Bold entries indicate the best performance.

Feature Group Study Table 5 compares different compositions of the feature group in DDWM. Specifically, *Dis*, *Sem*, and *Bri* represent the distortion features refined by DSDM, the semantic features extracted by the image encoder, and the supplementary bridging features derived from the difference between the two aforementioned feature groups, respectively. To note, *Only* indicates that only a specific feature group is included, while *w/o* suggests the inclusion of the other two feature groups yet excluding the current one. The results demonstrate that as more feature groups are introduced, the proposed method encompasses more information, resulting in more comprehensive weighing and better reflecting the quality score of the current image.

Qualitative Analysis Fig. 3 illustrates a comparative analysis involving three types of attention maps: attention maps on specific distortion obtained from DA-CLIP, semantic attention maps produced by the image encoder, and refined attention maps on distortions through DSDM. Technically, we visualize only the top three distortions, identified by the DDWM, of the most significance on image quality. From the figure, it can be observed that the refined attention maps effectively suppress irrelevant areas related to semantic information, while easing attention noise introduced by fictitious

Methods	KonIQ		LIVEC	
	SRCC	PLCC	SRCC	PLCC
Only Dis	0.933	0.944	0.897	0.912
Only Sem	0.936	0.944	0.899	0.913
Only Bri	0.936	0.947	0.897	0.914
w/o Dis	0.938	0.948	0.913	0.918
w/o Sem	0.940	0.951	0.911	0.919
w/o Bri	0.939	0.950	0.909	0.916
Full(Ours)	0.941	0.954	0.914	0.926

Table 5: Ablation study on different compositions of feature groups. Bold entries indicate the best performance.

distortions. These results indicate that the proposed DSDM significantly enhances the ability to capture and refine important distortion features in BIQA tasks.

Conclusion

We propose DR.Experts, a BIQA framework leveraging distortion priors refined via a distortion-aware vision-language model. Its Dynamic Distortion Weighting Module adaptively weights features using cues like distortion characteristics and semantics to simulate human vision. This design reduces over-reliance on unified features and enhances data efficiency. Extensive experiments confirm DR.Experts' superior performance and generalization across BIQA benchmarks.

Acknowledgments

This work was supported in part by the National Science Foundation of China (Grant No. 62301041), and in part by Beijing Institute of Technology Research Fund Program for Young Scholars.

References

- Agnolucci, L.; Galteri, L.; and Bertini, M. 2024. Quality-aware image-text alignment for opinion-unaware image quality assessment. *arXiv preprint arXiv:2403.11176*.
- Agnolucci, L.; Galteri, L.; Bertini, M.; and Del Bimbo, A. 2024a. Arniqa: Learning distortion manifold for image quality assessment. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 189–198.
- Agnolucci, L.; Galteri, L.; Bertini, M.; and Del Bimbo, A. 2024b. Reference-based Restoration of Digitized Analog Videotapes. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1659–1668.
- Bosse, S.; Maniry, D.; Müller, K.-R.; Wiegand, T.; and Samek, W. 2017. Deep neural networks for no-reference and full-reference image quality assessment. *IEEE Transactions on image processing*, 27(1): 206–219.
- Chen, C.; Mo, J.; Hou, J.; Wu, H.; Liao, L.; Sun, W.; Yan, Q.; and Lin, W. 2024. Topiq: A top-down approach from semantics to distortions for image quality assessment. *IEEE Transactions on Image Processing*, 33: 2404–2418.
- Chu, X.; Duan, H.; Wen, Z.; Xu, L.; Hu, R.; and Xiang, W. 2025a. Union-Domain Knowledge Distillation for Underwater Acoustic Target Recognition. *IEEE Transactions on Geoscience and Remote Sensing*.
- Chu, X.; Zhou, H.; Zhang, Y.; Zhang, Y.; Hu, R.; Duan, H.; Huang, Y.; Zheng, Y.; and Ji, R. 2025b. Attention-driven acoustic properties learning for underwater target ranging. *Pattern Recognition*, 164: 111560.
- Ciancio, A.; Da Silva, E. A.; Said, A.; Samadani, R.; Obrador, P.; et al. 2010. No-reference blur assessment of digital pictures based on multifeature classifiers. *IEEE Transactions on image processing*, 20(1): 64–75.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Fang, Y.; Zhu, H.; Zeng, Y.; Ma, K.; and Wang, Z. 2020. Perceptual quality assessment of smartphone photography. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3677–3686.
- Ghadiyaram, D.; and Bovik, A. C. 2015. Massive online crowdsourced study of subjective and objective picture quality. *IEEE transactions on image processing*, 25(1): 372–387.
- Golestaneh, S. A.; Dadsetan, S.; and Kitani, K. M. 2022. No-reference image quality assessment via transformers, relative ranking, and self-consistency. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 1220–1230.
- Guan, F.; Li, X.; Yu, Z.; Lu, Y.; and Chen, Z. 2025. QMamba: On First Exploration of Vision Mamba for Image Quality Assessment. In *Forty-second International Conference on Machine Learning*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hosu, V.; Lin, H.; Sziranyi, T.; and Saupe, D. 2020. KonIQ-10k: An ecologically valid database for deep learning of blind image quality assessment. *IEEE Transactions on Image Processing*, 29: 4041–4056.
- Hu, R.; Chu, X.; Dou, D.; Liu, X.; Liu, Y.; and Qi, B. 2025. Toward Real-World Applicability: Lightweight Underwater Acoustic Localization Model Through Knowledge Distillation. *IEEE Journal of Oceanic Engineering*.
- Huang, Z.; Li, X.; Fu, B.; Chu, X.; Li, K.; Shen, Y.; and Zhang, Y. 2024. Scale Contrastive Learning with Selective Attentions for Blind Image Quality Assessment. *arXiv preprint arXiv:2411.09007*.
- Kang, L.; Ye, P.; Li, Y.; and Doermann, D. 2014. Convolutional neural networks for no-reference image quality assessment. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1733–1740.
- 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.
- Kwon, D.; Kim, D.; Ki, S.; Jo, Y.; Lee, H.-E.; and Kim, S. J. 2024. ATTIQA: Generalizable image quality feature extractor using attribute-aware pretraining. In *Proceedings of the Asian Conference on Computer Vision*, 4526–4543.
- Li, A.; Wu, J.; Liu, Y.; and Li, L. 2024a. 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*, 28422–28431.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, 19730–19742. PMLR.
- Li, X.; Gao, T.; Hu, R.; Zhang, Y.; Zhang, S.; Zheng, X.; Zheng, J.; Shen, Y.; Li, K.; Liu, Y.; et al. 2024b. Adaptive Feature Selection for No-Reference Image Quality Assessment by Mitigating Semantic Noise Sensitivity. In *International Conference on Machine Learning*, 27808–27821. PMLR.
- Li, X.; Huang, Z.; Zhang, Y.; Shen, Y.; Li, K.; Zheng, X.; Cao, L.; and Ji, R. 2025. Few-Shot Image Quality Assessment via Adaptation of Vision-Language Models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10442–10452.
- Liang, Z.; Li, C.; Zhou, S.; Feng, R.; and Loy, C. C. 2023. Iterative prompt learning for unsupervised backlit image enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8094–8103.
- Luo, Z.; Gustafsson, F. K.; Zhao, Z.; Sjölund, J.; and Schön, T. B. 2024. Controlling Vision-Language Models for Multi-Task Image Restoration. In *International Conference on Learning Representations*.

- Madhusudana, P.; Birkbeck, N.; Wang, Y.; Adsumilli, B.; and Bovik, A. 2022. Image Quality Assessment Using Contrastive Learning. *IEEE Transactions on Image Processing: a Publication of the IEEE Signal Processing Society*, 31: 4149–4161.
- 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.
- Moorthy, A. K.; and Bovik, A. C. 2011. Blind image quality assessment: From natural scene statistics to perceptual quality. *IEEE transactions on Image Processing*, 20(12): 3350–3364.
- Qin, G.; Hu, R.; Liu, Y.; Zheng, X.; Liu, H.; Li, X.; and Zhang, Y. 2023. Data-efficient image quality assessment with attention-panel decoder. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 2091–2100.
- Qin, G.; Wang, Z.; Shen, D.; Liu, H.; Zhou, H.; Wu, J.; Hu, R.; and Jin, Y. 2025. Structure Matters: Revisiting Boundary Refinement in Video Object Segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 14431–14442.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Saad, M. A.; Bovik, A. C.; and Charrier, C. 2012. Blind Image Quality Assessment: A Natural Scene Statistics Approach in the DCT Domain. *IEEE Transactions on Image Processing*, 21(8): 3339–3352.
- Saha, A.; Mishra, S.; and Bovik, A. C. 2023. Re-iqa: Unsupervised learning for image quality assessment in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5846–5855.
- Shin, N.-H.; Lee, S.-H.; and Kim, C.-S. 2024. Blind image quality assessment based on geometric order learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12799–12808.
- 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.
- Touvron, H.; Cord, M.; and Jégou, H. 2022. Deit iii: Revenge of the vit. In *European conference on computer vision*, 516–533. Springer.
- 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, J.; Chen, P.; Zheng, N.; Chen, B.; Principe, J. C.; and Wang, F.-Y. 2021. Associations between MSE and SSIM as cost functions in linear decomposition with application to bit allocation for sparse coding. *Neurocomputing*, 422: 139–149.
- Wu, H.; Zhang, Z.; Zhang, W.; Chen, C.; Liao, L.; Li, C.; Gao, Y.; Wang, A.; Zhang, E.; Sun, W.; et al. 2024a. Q-Align: Teaching LMMs for Visual Scoring via Discrete Text-Defined Levels. In *International Conference on Machine Learning*, 54015–54029. PMLR.
- Wu, H.; Zhu, H.; Zhang, Z.; Zhang, E.; Chen, C.; Liao, L.; Li, C.; Wang, A.; Sun, W.; Yan, Q.; et al. 2024b. Towards open-ended visual quality comparison. In *European Conference on Computer Vision*, 360–377. Springer.
- Xu, K.; Liao, L.; Xiao, J.; Chen, C.; Wu, H.; Yan, Q.; and Lin, W. 2023. Local Distortion Aware Efficient Transformer Adaptation for Image Quality Assessment. *CoRR*.
- Xu, K.; Liao, L.; Xiao, J.; Chen, C.; Wu, H.; Yan, Q.; and Lin, W. 2024. Boosting image quality assessment through efficient transformer adaptation with local feature enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2662–2672.
- Yang, J.; Fu, J.; Zhang, Z.; Liu, L.; Li, Q.; Zhang, W.; and Cao, W. 2024. Align-IQA: aligning image quality assessment models with diverse human preferences via customizable guidance. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 10008–10017.
- Ye, T.; Dong, L.; Xia, Y.; Sun, Y.; Zhu, Y.; Huang, G.; and Wei, F. 2025. Differential Transformer. In *The Thirteenth International Conference on Learning Representations*.
- Ying, Z.; Niu, H.; Gupta, P.; Mahajan, D.; Ghadiyaram, D.; and Bovik, A. 2020. From patches to pictures (PaQ-2-PiQ): Mapping the perceptual space of picture quality. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3575–3585.
- 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, 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.
- Zhao, K.; Yuan, K.; Sun, M.; Li, M.; and Wen, X. 2023. Quality-aware pre-trained models for blind image quality assessment. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 22302–22313.
- Zhou, H.; Yang, R.; Tang, L.; Qin, G.; Hu, R.; and Li, X. 2025. Gamma: Toward Generic Image Assessment with Mixture of Assessment Experts. In *Proceedings of the 33rd ACM International Conference on Multimedia*, MM '25, 8815–8824. New York, NY, USA: Association for Computing Machinery. ISBN 9798400720352.