

NTIRE 2024 Challenge on Short-form UGC Video Quality Assessment: Methods and Results

Xin Li[†] Kun Yuan[†] Yajing Pei[†] Yiting Lu[†] Ming Sun[†] Chao Zhou[†]
 Zhibo Chen[†] Radu Timofte[†] Wei Sun Haoning Wu Zicheng Zhang
 Jun Jia Zhichao Zhang Linhan Cao Qiubo Chen Xionghuo Min Weisi Lin
 Guangtao Zhai Jianhui Sun Tianyi Wang Lei Li Han Kong Wenxuan Wang
 Bing Li Cheng Luo Haiqiang Wang Xiangguang Chen Wenhui Meng
 Xiang Pan Huiying Shi Han Zhu Xiaozhong Xu Lei Sun Zhenzhong Chen
 Shan Liu Fangyuan Kong Haotian Fan Yifang Xu Haoran Xu Mengduo Yang
 Jie Zhou Jiase Li Shijie Wen Mai Xu Da Li Shunyu Yao Jiazhi Du
 Wangmeng Zuo Zhibo Li Shuai He Anlong Ming Huiyuan Fu Huadong Ma
 Yong Wu Fie Xue Guozhi Zhao Lina Du Jie Guo Yu Zhang
 Huimin Zheng Junhao Chen Yue Liu Dulan Zhou Kele Xu
 Qisheng Xu Tao Sun Zhixiang Ding Yuhang Hu

Abstract

This paper reviews the NTIRE 2024 Challenge on Short-form UGC Video Quality Assessment (S-UGC VQA), where various excellent solutions are submitted and evaluated on the collected dataset KVQ from popular short-form video platform, i.e., Kuaishou/Kwai Platform. The KVQ database is divided into three parts, including 2926 videos for training, 420 videos for validation, and 854 videos for testing. The purpose is to build new benchmarks and advance the development of S-UGC VQA. The competition had 200 participants and 13 teams submitted valid solutions for the final testing phase. The proposed solutions achieved state-of-the-art performances for S-UGC VQA. The project can be found at <https://github.com/lixinustc/KVQ-Challenge-CVPR-NTIRE2024>.

1. Introduction

Short-form UGC video platforms, e.g., Kwai, and Tiktok have attained significant success and widespread popularity, attracting billions of users globally. In contrast to long-form UGC videos, short-form UGC (S-UGC) videos present sev-

eral advantages, including mobile-friendly broadcasting, user-friendly engagement, kaleidoscope content creation, and snackable content, etc. [45]. However, S-UGC videos inevitably suffer from inconsistent and even poor subjective quality due to unprofessional creation, improper collection environments, or limited processing workflow. Building the quality assessment benchmarks tailed for S-UGC videos is crucial and urgent to advance the fast development of associated video processing techniques and quality control mechanisms for user-uploaded S-UGC videos.

Recently, amounts of studies [5, 24–26, 31, 32, 35, 43, 44, 66, 67, 72, 77, 78, 83, 84] have been proposed to assess the perceptual quality of long UGC videos or images by excavating the capability of pre-trained backbones, i.e., ResNet [15, 25, 33], and Transformer [34, 41, 42, 54, 64, 65, 76]. Since the scarcity of the VQA database, some well-designed sampling strategies [65, 66], e.g., fragment sampling, are introduced to reduce the computational complexity of the VQA method and enhance the diversity of VQA database. Apart from that, there are also some works exploring the potential of large cross-modality foundation models [31, 50, 68, 70] for VQA. However, the above works are only designed and validated on general long UGC VQA databases, e.g., LSVQ [75], YouTube-UGC [61], and KoNViD-1k [16], lacking the perception capability for short-form UGC VQA.

This NTIRE 2024 Short-form UGC Video Quality Assessment Challenge is held to promote related research on S-UGC VQA, and develop the powerful S-UGC VQA

[†] X. Li (lixin666@mail.ustc.edu.cn), K. Yuan (yuankun03@kuaishou.com), Y. Pei, Y. Lu, M. Sun, C. Zhou, Z. Chen and R. Timofte are the challenge organizers.

The other authors are participants of the NTIRE 2024 Short-form UGC Video Quality Assessment Challenge.

The NTIRE2024 website: <https://cvlai.net/ntire/2024/>

The KVQ database: <https://lixinustc.github.io/projects/KVQ/>

benchmark to assist the technique evolution of S-UGC video creation, compression, or processing, etc. In this competition, we utilize the KVQ dataset [45] collected from the Kwai platform to evaluate the submitted methods. The KVQ dataset contains 4200 S-UGC videos with distinct creation modes, content scenarios, and video processing workflows, satisfying the data distribution of practical short-form video platforms. For evaluation, the KVQ dataset is divided into three parts with a proportion of 7:1:2 for training, validation, and testing. Instead of only quality scores, we also provide the ranking labels for indistinguishable samples in the validation and testing parts, ensuring the fine-grained perception evaluation for S-UGC VQA methods. The evaluation metric is composed of four parts, including PLCC, SROCC, the ranking accuracies between homogeneous video pairs (*i.e.*, with the same content) and non-homogeneous video pairs. *Notably, we have released the quality labels and ranking scores of the validation data, and reopened the test submission on Codalab [49] after this competition, intending to facilitate the development of algorithms.*

The competition has two phases, *i.e.*, the development and testing phases, attracting 200 registered participants in total. There are 49 teams, and 56 teams submitting their predictions in the development and testing phases, respectively. Finally, 13 teams submitted their fact sheets and validation codes for competition ranking. Their algorithms will be summarized in the Section 4.

This challenge is one of the NTIRE 2024 Workshop¹ associated challenges on: dense and non-homogeneous dehazing [2], night photography rendering [3], blind compressed image enhancement [74], shadow removal [58], efficient super resolution [51], image super resolution ($\times 4$) [6], light field image super-resolution [62], stereo image super-resolution [60], HR depth from images of specular and transparent surfaces [79], bracketing image restoration and enhancement [82], portrait quality assessment [4], quality assessment for AI-generated content [36], restore any image model (RAIM) in the wild [28], RAW image super-resolution [8], short-form UGC video quality assessment [27], low light enhancement [37], and RAW burst alignment and ISP challenge.

2. Challenge

The NTIRE 2024 Short-form UGC Video Quality Assessment Challenge is organized to advance the development of VQA techniques for short-form UGC videos. Moreover, this is the first challenge focusing on S-UGC VQA, aiming to build a new benchmark to guide the perceptual quality improvement of S-UGC videos. The details of the whole challenge are described in the following parts, including

¹<https://cvlai.net/ntire/2024/>

datasets, evaluation protocol, and challenge phases.

2.1. Dataset

To ensure the reliable and comprehensive evaluation for each method, we utilize our KVQ dataset [45] collected from the practical S-UGC video platform, *i.e.*, Kwai/Kuaishou platform. The KVQ dataset contains 4200 S-UGC videos, consisting of 600 user-uploaded videos and 3600 processed videos via various practical processing workflows in Kuaishou platform. In the collection process, we select video samples from 9 content scenarios, *i.e.*, food, stage, computer graphic, night, caption, person, crowd, landscape, and portrait, covering several typical creation modes, including three-stage, live, subtitled, and special effect videos, etc.

In practical S-UGC video platforms, video processing tools are diverse and complicated. Consequently, we apply three typical processing techniques, including enhancement, pre-processing, and transcoding. Among them, enhancement tools consist of de-artifact, de-noise, and deblur, and pre-processing tools can be divided into global-level and ROI-level. The transcoding tools aim to reduce the bit-stream, where the quantization parameters (QPs) from 16 to 47 are divided into six groups. The above three processing strategies are serially applied to user-uploaded S-UGC videos, resulting in various degradations.

These S-UGC videos are annotated by 15 professional researchers specializing in image processing. The Mean Opinion Score (MOS) value is [1-5] and the scoring interval is 0.5, which makes it easy for humans to annotate the quality score. Notably, the fine-grained quality differences between different S-UGC video samples are hard to identify. In the validation and testing datasets of this challenge, we also provide the ranking score for the video samples with similar quality scores. In the test stage, we annotated the ranking scores for 169 homologous video pairs (*i.e.*, with the same contents but different degradations), and 231 non-homologous video pairs, respectively. For the validation dataset, we provide ranking scores for 62 homologous video pairs and 38 non-homologous video pairs, respectively. The above hybrid annotations enable a more thorough evaluation of submitted VQA methods. More details of the KVQ dataset can be found in work [45].

2.2. Evaluation protocol

In this challenge, we aim to evaluate each method from four perspectives: (i) the prediction monotonicity with Spearman rank-order correlation coefficient (SROCC) and (ii) the prediction accuracy with Person Linear Correlation coefficient (PLCC); (iii) fine-grained ranking accuracy between homologous video pairs, *i.e.*, Rank1; (iv) the challenging ranking accuracy between non-homologous video pairs, *i.e.*, Rank2. The final score used for ranking is computed by

Table 1. Quantitative results of the NTIRE 2024 Short-form UGC Video Quality Assessment Challenge.

Rank	Team	Leader	Final Score	SROCC	PLCC	Rank1	Rank1	Pretrained	Ensemble	Extra Data	
1	SJTU MMLab	Wei Sun	0.9228	0.9361	0.9359	0.7792	0.8284	✓	✓	✓ (public)	
2	IH-VQA	Jianhui Sun	0.9145	0.9298	0.9325	0.7013	0.8284	✓	✓	✓ (private)	
3	TVQE	Haiqiang Wang	0.9120	0.9268	0.9312	0.6883	0.8284	✓	✓	✓ (private)	
4	BDVQAGroup	Fangyuan Kong	0.9116	0.9275	0.9211	0.7489	0.8462	✓	✓	✓ (public)	
5	VideoFusion	Haoran Xu	0.8932	0.9026	0.9071	0.7186	0.8580	✓	✓	✓ (public)	
6	MC ² Lab	Shijie Wen	0.8855	0.8966	0.8977	0.7100	0.8521	✓	✓	×	
7	Padding	Da Li	0.8690	0.8841	0.8839	0.6623	0.8047	×	×	×	
8	ysy0129	Shunyu Yao	0.8655	0.8759	0.8777	0.6883	0.8402	✓	✓	×	
9	lizhibo	Zhibo Li	0.8641	0.8778	0.8822	0.6494	0.7929	✓	×	×	
10	YongWu	Yong Wu	0.8555	0.8629	0.8668	0.6970	0.8462	×	×	×	
11	we are a team	Lina Du, Jie Guo	0.8243	0.8387	0.8324	0.6234	0.8225	×	×	×	
12	dulan	Dulan Zhou	0.8098	0.8164	0.8297	0.5758	0.8047	✓	✓	×	
13	D-H	Zhixiang Ding	0.7677	0.7774	0.7832	0.5931	0.7160	✓	×	×	
Baseline			VSFA [25]	0.7869	0.7974	0.7950	0.6190	0.7870			
			SimpleVQA [54]	0.8159	0.8306	0.8202	0.6147	0.8461	×	×	×
			FastVQA [65]	0.8356	0.8473	0.8467	0.6494	0.8166			

combining the above metrics:

$$\text{Final_Score} = 0.45 * \text{SROCC} + 0.45 * \text{PLCC} + 0.05 * \text{Rank1} + 0.05 * \text{Rank2} \quad (1)$$

We expect the above equation can estimate both coarse- and fine-grained quality assessment capability for submitted S-UGC VQA methods.

2.3. Challenge phases

There are two phases in this challenge, *i.e.*, the development phase, and testing phase.

Development phase: In this phase, we divided our collected KVQ dataset into training, validation, and testing data based on their contents. We release 2926 S-UGC videos and their corresponding quality scores to participants, which aims to support them develop their S-UGC VQA algorithm. Besides, we provide 420 S-UGC videos without quality labels for validation. Participants can upload their submissions to the challenge platform and obtain their final score, SROCC, PLCC, Rank1, and Rank2 accuracies. We also release the benchmark code based on SimpleVQA [54] to help participants quickly familiarize the process of S-UGC VQA. In the development phase, 881 submissions from 49 teams are received.

Testing phase In the test phase, we release 854 S-UGC videos for testing on the challenge platform. The testing leaderboard is hidden and each team can observe their testing results individually. The ranking for this challenge is based on the final score as Eq. 1 in the testing stage. After testing, the participants are requested to submit the fact sheet and source code/executable to reproduce their results. There are 56 teams submitting their prediction results on the challenge platform. Finally, we received the fact sheets

and source codes from 13 teams, which are used for final ranking.

3. Challenge Results

The challenge results are shown in Table 1. We only report the performances of the teams submitting their fact sheets. From this Table, we can find the top five teams, including SJTU MMLab, IH-VQA, TVQE, BDVQAGroup, and VideoFusion have achieved excellent results on both PLCC and SROCC, which exceed 0.9. Among all teams, Padding, lizhibo, YongWu, we are a team and D-H do not use ensemble strategies and extra data. MC²Lab, ysy0129 and dulan do not utilize extra data. However, the above methods still achieve great performances in this competition.

We also report the performances for existing baselines, including VSFA [25], SimpleVQA [54], and FastVQA [65], without utilizing the ensemble and extra data. We can observe that 10 teams have achieved better performances than these baselines. The methods of each team are described in Sec. 4, which greatly promote the development of short-form UGC video quality assessment.

4. Teams and Methods

4.1. SJTU MMLab

This team proposes the BVQA method, which is based on SimpleVQA+ [54, 55], comprising a Swin Transformer-B [41] for spatial feature extraction and a temporal pathway of SlowFast for temporal feature extraction. The whole framework is shown in Fig. 1. Given the diverse visual content and complexity distortions in the KVQ dataset, they incorporate three BI/VQA models: LIQE [81], Q-Align [70], and FAST-VQA [65], to extract comprehensive quality-

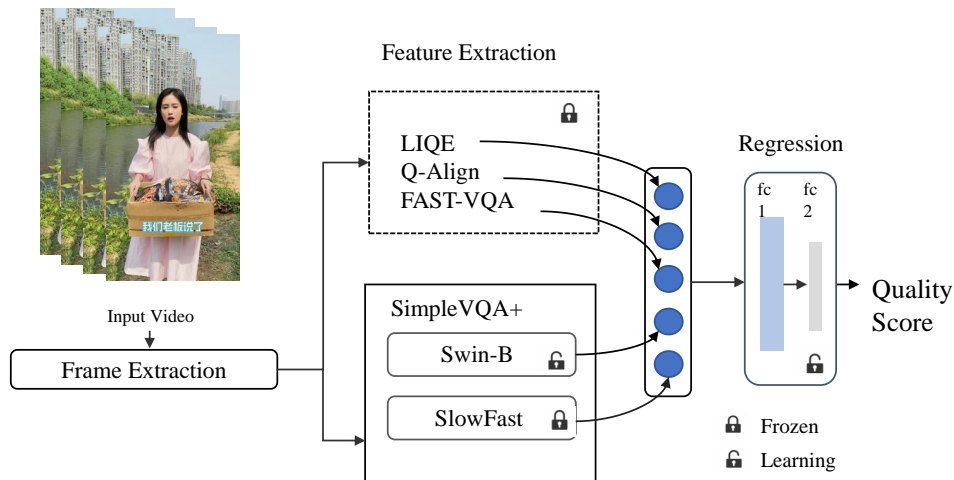


Figure 1. The network architecture [56] of the solution proposed by team SJTU MMLab.

aware features to aid their BVQA model. Specifically, LIQE extracts eleven types of distortions, nine scene categories, and five quality-level probabilities as quality-related and scene-specific features. Q-Align extracts quality-level features for video frames by the large multi-modality models, while FAST-VQA extracts spatial-temporal quality-aware features from local video patches. They concatenate these features with those from SimpleVQA+ and utilize a two-layer MLP to derive the video quality scores.

Training Details. In the above BVQA method, the weights of SimpleVQA+ are initialized by training it on the LSVQ [75] dataset. LIQE is trained on LIVE [53], CSIQ [23], KADID-10k [29], BID [7], CLIVE [13], and KonIQ-10k [17]. Q-Align is trained on SPAQ [11], KonIQ-10k [17], KADID-10k [29], LSVQ [75], and AVA [46]. FAST-VQA is also trained on LSVQ [75]. They sample one key frame from one-second video chunks (*i.e.*, 1 fps) for the spatial feature extraction module of SimpleVQA+ as well as for extracting LIQE and Q-Align features. The resolution of key frames is further resized to 384×384 for training. For the temporal feature extraction module, the resolution of the videos is resized to 224×224 . They then split the whole video into several one-second length video chunks to extract the corresponding temporal features. Fast-VQA features are extracted from the entire video using the fragment sampling method [65]. They train the proposed model on 2 Nvidia RTX 3090 GPUs with a batch size 6 for 30 epochs. The learning rate is set as $1e - 5$.

Testing Details. They randomly split the publicly available data of the KVQ dataset into ten different training-validation sets and train ten BVQA model. The video qual-

ity score is computed by averaging the quality scores obtained from these models.

4.2. IH-VQA

This team proposes an Ensemble-based Video Quality Assessment System, which consists of seven different expert models, *i.e.*, four regression expert models, including SigLIP-ViT-SO400M [80], SigLIP-ViT-B [80], ConvNeXt V2-H [63], ConvNeXt V2-L [63], and three classification expert models, consisting of ConvNeXt-B [40], MobileNetV3 [18] and SqueezeNet [19]. The system framework is outlined in Fig. 2. As shown in this figure, four video frames are inputted into these four regression models, respectively, and their features are fused in the last pooling layer. After that, two additional feed-forward layers are used to produce the final prediction of the input video. They also adopt three classification models to predict the probability of being rated in different score intervals. These classification expert models rate the given video frames, and their votes are averaged as the final prediction. To address the quality variance issues within different frames, they designed a novel loss function, *i.e.*, the mean absolute error between the target quality label and the mean of the predictions from multiple frames, jointly with the cross-entropy loss.

Training Details. Four regression models are pre-trained on an in-house dataset with 20K data and then fine-tuned on the official dataset KVQ. During training, techniques like the Fast Gradient Method and Exponential Moving Average (EMA) [20] are also exploited for better performance. For three classification models, the weights of their spatial at-




Figure 2. The framework of the Ensemble-based Video Quality Assessment System proposed by IH-VQA.

tention blocks are initialized to zero. The remaining model blocks keep the pre-trained weights on the ImageNet-1K dataset [52]. They also use EMA to boost model performance. Once they get the well-trained ConvNeXt-B model, it will serve as a teacher model for distillation purposes, aiding in the training of two additional student models (*i.e.*, MobileNetV3 and SqueezeNet). Their models are trained for nine GPU hours with 20 epochs. An Adam optimizer with a learning rate of $3e^{-5}$ is leveraged to train large models like SigLIP and ConvNeXt, while an Adam optimizer with a learning rate of $1e^{-3}$ is used to train smaller models such as MobileNetV3 and SqueezeNet.

Testing Details. During inference, the final prediction is determined by averaging the predictions of the above seven models.

4.3. TVQE

This team proposes the TVQE, which is a hybrid model trained for VQA tasks. The overall framework is shown in Fig. 3. It combines two multi-modality models to extract visual information and semantic information and a classical convolution neural network to capture technical and aesthetic quality.

First, they introduce a feature pyramid aggregation mechanism on the backbone, *i.e.*, the ConvNeXt [40], of the aesthetics branch of Dover [69] to extract better aesthetic representations. This is motivated by the fact that the aesthetics of short-form video is vital to the final quality.

They adopt two large multi-modality models, *i.e.*, Q-align [70], with different levels of quality descriptions. The aim is to alleviate the quantization effect around 5-level

quality descriptors, which is introduced in the process to convert predicted rating levels to the final predicted score. Furthermore, as a dedicated video quality metric, they integrate the optical flow module to capture motion information into the assessment process. The extracted motion feature is combined with a visual feature to seamlessly represent the quality of the entire video. These three models were trained independently on the official KVQ dataset [45] and private datasets in an end-to-end manner. During the inference stage, the final predicted score could be obtained by heuristically fusing the prediction results of these models.

Experimental Analysis. They have tested several top-performing VQA methods in the validation stage. In general, two SOTA metrics, *i.e.* Dover [69] and Q-align [70] are selected due to their superior performance. Meanwhile, they have found the motion feature is absent in the original Q-align. This might lead to sub-optimal representation for the video quality assessment task. Furthermore, they find a quantization effect in Fig. 4 when the original Q-align is evaluated on a split of the training set, even though the obtained PLCC and SRCC are not bad.

Training Details. They train three models, *i.e.*, Dover [69], Q-Align [70], and IFRNet [21] on the above dataset independently. The training datasets consist of three parts, including the KVQ training part, a private UGC dataset, and a private PGC dataset. Among them, The private UGC dataset also contains short-form videos. The processing pipeline includes enhancement filters before transcoding. The private PGC dataset contains multiple distortion types, such as transcoding with different codecs, enhancement followed by transcoding. These three datasets are simply con-

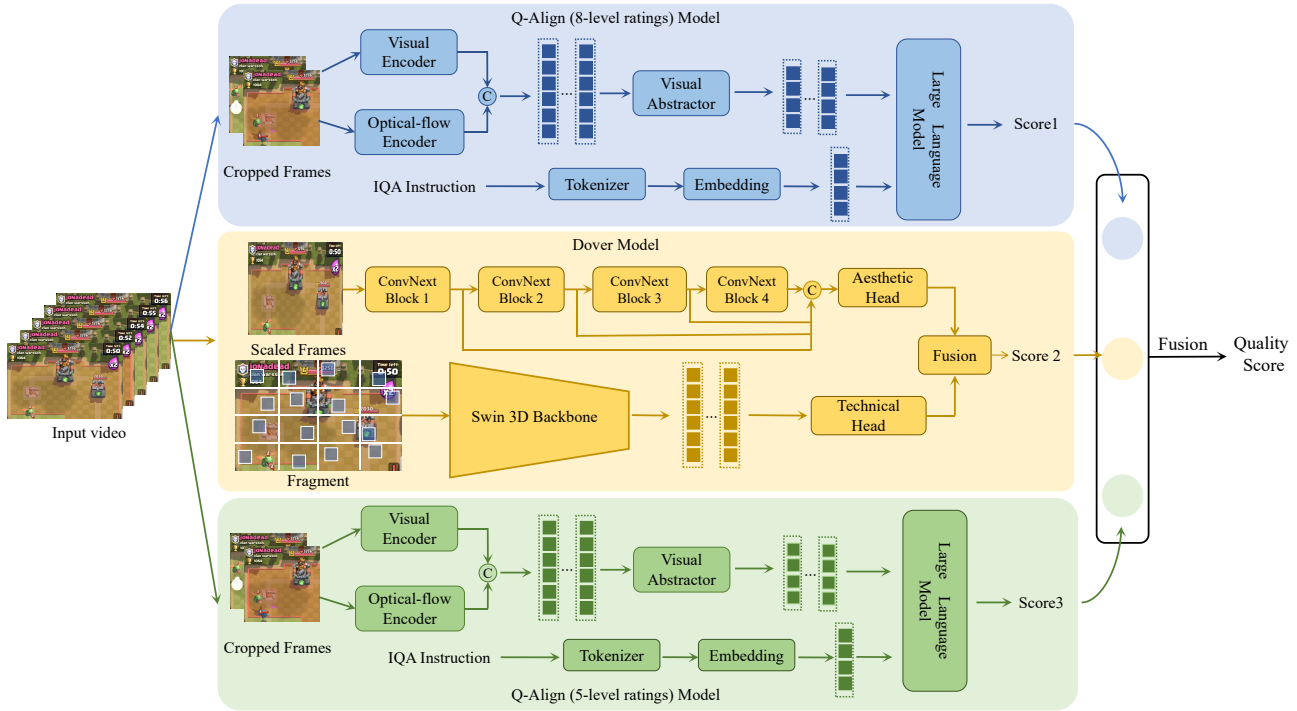


Figure 3. The network architecture of the solution proposed by team TVQE.




Figure 4. The quantization effect when the original Q-Align is tested on the split of the training dataset. The experimental results are from team TVQE.

concatenated together.

Testing Details. A test video is processed by the three models independently. The predicted scores are heuristically fused together as the final prediction.

4.4. BDVQAGroup

This team chooses a method called Q-Align to assess the quality of UGC videos. It is based on large multi-modality models (LMMs). Q-Align converts MOS scores into rating levels, and uses classification to teach LMMs with text-defined rating levels instead of scores. During inference, it extracts the close-set probabilities of rating levels and

performs a weighted average to obtain the LMM-predicted score. They raise several training and inference tricks to increase the performance of this method. The whole framework is depicted in Fig. 5.

Training Details. This team uses several data augmentation methods to increase the training dataset and enhance the robustness of the model, which are listed as follows:

- data resample: calculate the maximum and minimum mos score of training dataset, divide them into 5 intervals for rating level mapping and then resample offline until each rating level has the same amount of training data.
- random frame sample: during training, 8 frames are randomly selected from a video as input.
- random reverse: video frames are randomly arranged in chronological or reverse order.
- random crop: randomly crop frames at a ratio of 0.9, and the frames of the same video crop at the same position.

This team uses a Q-Align [70] model pre-trained on KonIQ [17], SPAQ [11], KADID [29], AVA [46], and LSVQ [75], and fine-tune this model through two strategies. The first model is fine-tuned for 3 epochs on the NTIRE 2024 KVQ dataset. The second model is fine-tuned for 1 epoch on the ICME 2021 UGCVQA [59], and then fine-tuned for another 3 epochs on the NTIRE 2024 KVQ dataset.

Testing Details. The following ensemble methods are implemented to increase model performance.

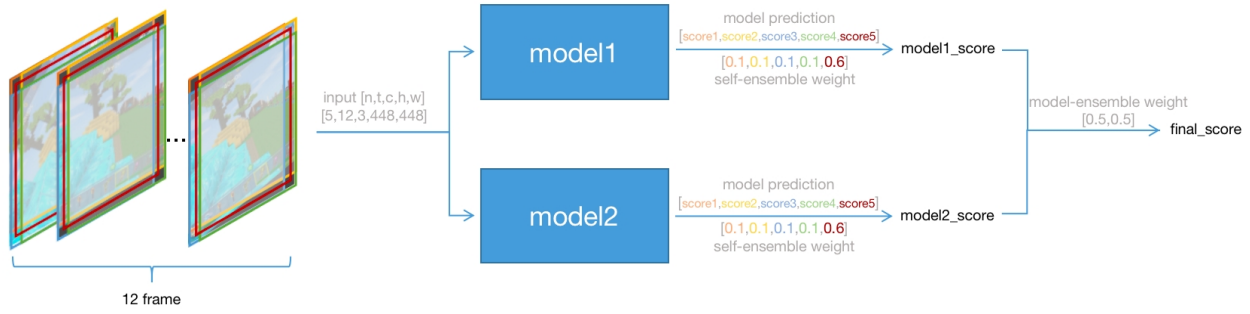


Figure 5. The network architecture of the solution proposed by team BDVQAGroup.

- self-ensemble: during inference, 12 frames are taken from each video at equal intervals, and FiveCrop is performed on each frame. Because of the sandwich video, they give higher weight to the prediction of the central area.
- model ensembles: they use two models to predict and give each prediction the same weight to get the final result.

4.5. VideoFusion

This team proposes a three-level (frame-segment-video) integration framework [73] for short-form UGC VQA in Fig. 6. The contributions of this framework are summarized as follows:

- They propose a multi-level framework. Globally, a three-level architecture is proposed to capture features at each level, and locally, features on backbones from low level to high level are fused.
- Based on the view of data augmentation, data augmentation in spatial and temporal domains is employed on the three-level architecture respectively to improve the robustness of the model.
- In order to distinguish between two kinds of hard samples and relative rank information, they design an adaptive relative rank loss.
- By using the redistributed model integration strategy, the distribution of labels is aligned and the training of the model is more stable.

Training Details. They set the batch size to 4 and assigned a weight of 0.3 to the rank loss. They employ the AdamW optimizer with an initial learning rate of 1×10^{-3} . Following a warm-up period spanning 3 epochs, the learning rate is modulated using a cosine decay schedule. The weight decay for the optimizer is configured at 1×10^{-2} , and the model underwent training for a total of 30 epochs.

Adaptive Rank-Aware Loss. In their method, they propose an Adaptive Rank-Aware Loss function to effectively handle the challenges posed by the coexistence of homogenous and heterogeneous data within Kwai UGC video dataset.

This loss function is designed to differentiate between hard samples and enhance the model’s ability to make fine-grained distinctions in video quality.

The loss function is formulated as follows:

$$e(y_i^{gt}, y_j^{gt}) = \begin{cases} 1, & y_i^{gt} \geq y_j^{gt} \\ -1, & y_i^{gt} < y_j^{gt} \end{cases} \quad (2)$$

$$\mathcal{L}_r = \frac{1}{m^2} \sum_{i=1}^m \sum_{j=1}^m [\max(0, -e(y_i^{gt}, y_j^{gt})(y_i - y_j))]^2 \quad (3)$$

$$\mathcal{L}_{\text{Rank}} = M\mathcal{L}_r + \lambda(1 - M)M_\alpha\mathcal{L}_r \quad (4)$$

In this equation, M_α denotes a distance indicator function that utilizes the threshold of the ground truth score of video pair (y_i^{gt}, y_j^{gt}) , where M_α is 1 if $|y_i^{gt} - y_j^{gt}| < c$ else is 0 if $|y_i^{gt} - y_j^{gt}| \geq c$. m is the number of all pairs in the same batch, and $M = [y_i^{\text{class}} = y_j^{\text{class}}]$ is an indicator that is 1 if the videos in the pair have sample class label (homogeneous videos), and 0 otherwise. The margin parameter λ is introduced to control the trade-off between homogeneous loss and non-homogeneous pairs loss.

4.6. MC²Lab

This team utilizes the Swin Transformer V2 [39] to extract features from salient regions within videos, employs ConvNext [40] for capturing the overall distortion information of the videos, and leverages the SlowFast [54] network for extracting information related to motion distortions. These features are then concatenated. Ultimately, a three-layer Multilayer Perceptron (MLP) is applied to yield the final quality score.

Training Details. This team utilizes the optimizer and schedule of FastVQA [65]. The total epochs of the training procedure are 20. The batch size is 8, and the video clip length is 5. The input resolution of the video clip is

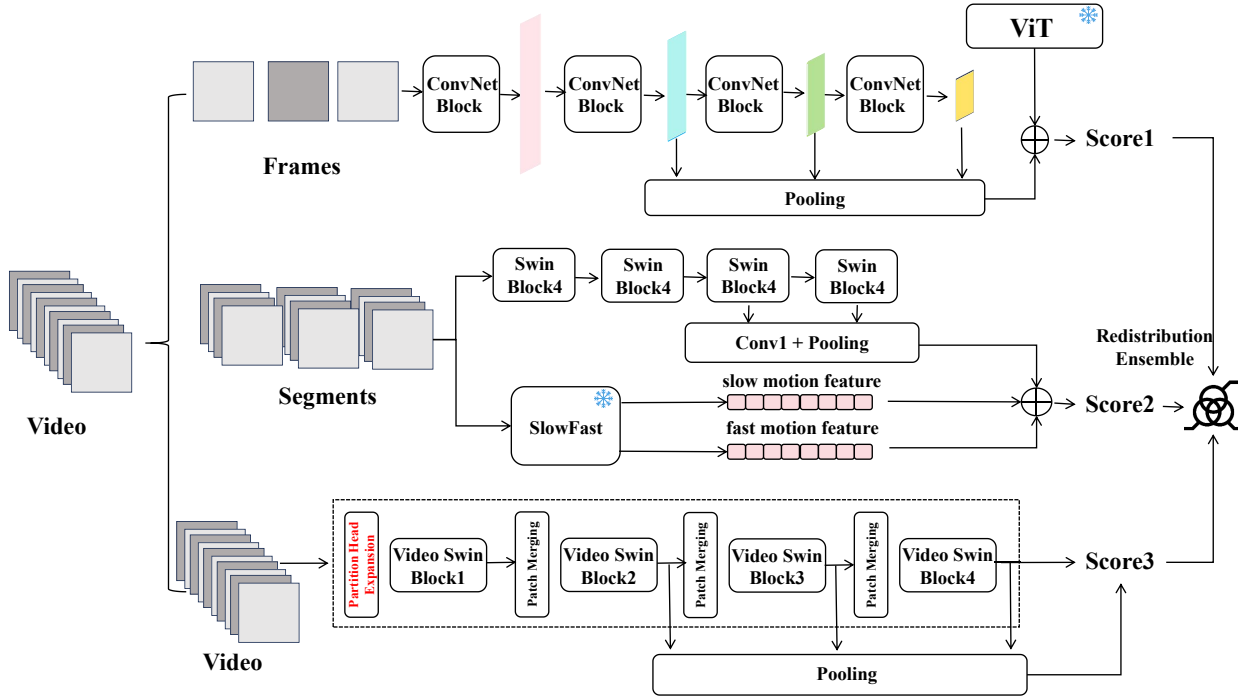


Figure 6. An overview of VideoFusion VQA framework, proposed by team VideoFusion.

768×448 for ConvNext [40] and 256×256 for Swin Transformer [39].

Testing Details. For each video, They predict the quality score 5 times and average them as the final score.

4.7. Padding

This team proposes the zero-padding strategy [1, 10, 14, 30, 47] for S-UGC VQA, which is utilized in all scenarios requiring frame supplementation or duplication as:

- When the number of frames read is less than the length of the video, the remaining positions are now padded with tensors consisting entirely of zeros.
- During the processing of each video segment, if the length of the segment is inadequate, it is now padded with a tensor that is entirely zeros.
- If the actual count of video segments falls short of the minimum required number of video clips, tensors filled with zeros, equal in size to the first segment, are now appended until the minimum segment count is met.

There are two advantages of zero-padding for S-UGC VQA: (i) Zero-element padding does not introduce any extra information to the image, thus avoiding the introduction of potential biases during the learning process. In contrast, other padding methods, such as replicating edge pixels, could introduce information that does not belong to the original video content, potentially affecting the network’s learning. (ii) Zero-padding helps maintain the spatial loca-

tion of features within a certain degree. Replicating the final frame suggests its multiple occurrences, which is detrimental to the convolutional network’s capacity to learn features in relation to their positional context. This team has found that when they modified the approach to generate feature inputs by substituting the prior method of replicating the final frame with a strategy of zero-element padding, the resultant accuracy of the test set rose to 86.90% from 86.16%, thereby affirming their belief that padding with zero elements is indeed a highly effective strategy.

4.8. ysy0129

This team proposes a video quality assessment method that embeds powerful self-supervised model priors. The framework of the proposed model is shown in Fig. 7. They sample frames from input video in sparse and dense manners, respectively, for extracting motion distortion and spatial distortion.

As for motion distortion, a self-supervised pre-trained model VideoMAE [57] is utilized to extract motion features. When evaluating the quality of a video, there exists a certain continuity in the quality between frames. Therefore, there is a certain amount of redundant information. Since the MAE model learns the feature representation of the data through the reconstruction task of an autoencoder. In this process, the model might ignore or compress redundant information. Consequently, Masked Autoencoders (MAE) based model




Figure 7. The network architecture of the solution proposed by team ysy0129.

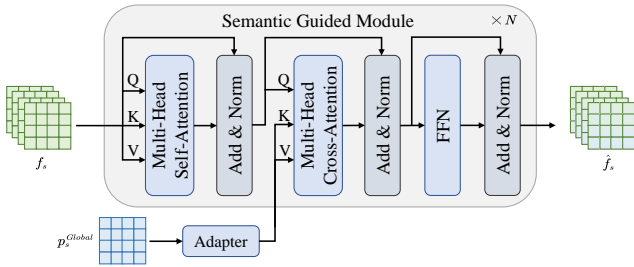


Figure 8. The Semantic Guided Module of the solution proposed by team ysy0129.




Figure 9. The framework of SimpleVQA-OPT proposed by team ysy0129.

VideoMAE [57] is chosen as the video sequence extraction algorithm, which is related to the information redundancy.

As for spatial distortion, the pre-trained semantic prior model is adopted to guide the spatial features. DINOv2 [48] has been revealed powerful semantic perception capability, attributed to self-supervised training on a dedicated, diverse, and curated image dataset, which can facilitate our model to focus on the semantic information which is related to the quality of a UGC video.

Specifically, as shown in the framework, a spatial feature extractor with the backbone ConvNeXt [40] is adopted to capture the multi-scale spatial distortion for video frames.

Since the extracted features have different sizes, the extracted features are adaptive pooling to the same size and concatenated together as f_s . Then, f_s is enhanced by the semantic feature p_s^{Global} generated by DINOv2 [48] in the Semantic Guided Module (SGM). As shown in Fig. 8, the SGM consists of N cross attention structure and an adapter composed of Linear-Relu-Linear architecture, which takes f_s and p_s^{Global} as input and output the guided feature \hat{f}_s . Next, they apply global average and stand deviation pooling operations on \hat{f}_s and concatenate them together with the motion features p_t^{Local} extracted by VideoMAE [57] after pooling. Lastly, the above spatial and motion features are fused together with a simple regression head, integrating both spatial and temporal aspects of the video.

Training Details. In the training phase, their end-to-end network is trained with AdamW optimizer for 50 epochs. The initial learning rate and weight decay are set to $3e^{-5}$ and 0.05 respectively. And batch size is set as 4.

They uniformly partition each input video into 8 chunks. Within the spatial feature extractor branch, they resize the original video to dimensions of 520×520 and randomly sample 1 frame from each chunk. Subsequently, they crop the obtained 8 frames to the dimension of 448×448 . For the temporal feature branch, they uniformly selected 32 frames from each chunk and resized them to the dimension of 224×224 .

Testing Details. The video processing procedure is the same as the training process, ensuring consistency in the handling of data across both stages. They conduct tests on two models, *i.e.*, their VQA framework and SimpleVQA-OPT as shown in Fig. 7 and Fig. 9, respectively. The SimpleVQA-OPT is achieved by incorporating optical flow information into SimpleVQA [54]. Inspired by StableVQA [22], they utilize the RAFT to extract the optical flow. Then, they utilize the pre-trained 3D ResNet of StableVQA [22], which is pre-trained on a video quality as-

assessment dataset with diversely-shaky UGC videos, to extract optical flow features. Subsequently, these features are combined with spatial and motion features. They perform 8 tests for each model, as the spatial feature branch samples different frames during each test process. All the scores are averaged to obtain the final results.

4.9. lizhibo

This team first segments the video into continuous chunks, then employs SlowFast R50 [12] as the backbone to extract motion features from a key frame, and utilizes ResNet50 [15] as the backbone to extract spatial features from all frames. Subsequently, a Multi-Layer Perceptron (MLP) network maps quality-aware features to chunk-level quality scores, followed by a temporal average pooling strategy to obtain the final video quality.

Training Details. This team first removes distorted videos from the training set and uses the first 500 video data as the validation set. The total epochs of the training process is 50, the initial learning rate of the Adam optimizer is $3e-5$, and the batch size is 8. During the training process, the Exponential Moving Average (EMA) [20] strategy is employed to enhance the model’s stability and improve the generalization ability of the model.

Testing Details. During the testing phase, they use the Exponential Moving Average (EMA) model for inference, reducing the risk of overfitting.

4.10. YongWu

This team utilizes various data augmentation for S-UGC VQA. At the onset of the competition, the team delves into various Video Quality Assessment (VQA) methodologies, including SimpleVQA [54], FastVQA [65], and DOVER [66]. Following a thorough comparison, it is found that DOVER outperforms the others, aligning with the findings presented in KVQ [45]. Consequently, DOVER is chosen as the baseline model for the competition.

To improve the performance of the baseline model, they aim to acquire additional videos without compromising the quality assessment score of the training data. To achieve this, they employ various methods as follow:

- **Padding.** To prevent distortion caused by resizing, they implement zero-padding to preprocess the training and validation videos while maintaining the integrity of the quality assessment scores. Upon evaluation, this method does not yield favorable results for the baseline model. The reason for this is that the aesthetic branch of the baseline model relies on proportional scaling, and the technical branch samples fragments that are unaffected by resizing. However, this approach proves to be effective for FastVQA.
- **Flip.** To maintain the quality assessment score, flipping videos with minimal text proved to be effective. Their

analysis of the training data revealed that videos with little text remained unaffected in terms of quality scores post-flipping. They select 747 such videos and applied flipping to them.

- **Average.** The dataset comprises an original video and five derivative versions, all classified under a single category. They organize the training data into separate folders, each containing videos of identical content. Subsequently, they randomly select two videos from each folder and calculate their average. Empirically, the quality score of the averaged video typically lies between that of the lower and higher-scored videos. For simplicity, they calculate the average quality score for these two videos, yielding a total of 418 new videos. Additionally, one could employ a predictive model to estimate the score and introduce slight variations to the averaged video’s quality score. However, due to time constraints, no further exploration was undertaken. It is important to recognize that the video quality assessment model essentially functions to fit the quality scores, and having a greater number of sampling points enhances the accuracy of this fit.

Rank Loss. To boost the rank score, they have experimented with increasing the weight of rank loss from 0.3 to 0.5, which proved beneficial for the model. they also consider replacing the rank loss function:

$$L_{rank}^{ij} = \max(0, |\hat{Q}_i - \hat{Q}_j| - \text{sign}(\hat{Q}_i, \hat{Q}_j) \cdot (Q_i - Q_j)) \quad (5)$$

where i and j are two video indexes in a mini-batch, and sign is formulated as:

$$\text{sign}(\hat{Q}_i, \hat{Q}_j) = \begin{cases} 1, & \hat{Q}_i \geq \hat{Q}_j \\ -1, & \hat{Q}_i < \hat{Q}_j \end{cases} \quad (6)$$

but as the rank score had a negligible impact, this change does not significantly improve the model.

Training Details. During the training phase, they train the end-to-end model using the AdamW optimizer for 50 epochs, with an initial learning rate of 0.001 and a weight decay of 0.05. The batch size was set at 12.

For training the model, they utilize 99 percent of the data instead of using cross-validation and normalized the original scores to the range of $[0, 1]$, aiding in model convergence. They sample 32 frames at 2-frame intervals from each video, with both the aesthetic and technical branches receiving inputs of size 224×224 .

Testing Details. In the testing phase, they use the last saved epoch model with Exponential Moving Average (EMA) [20] to enhance the model’s stability and prevent overfitting.

4.11. We are a team

This team proposes a multi-task based short video quality assessment method [9], which is shown in Fig. 10. Video

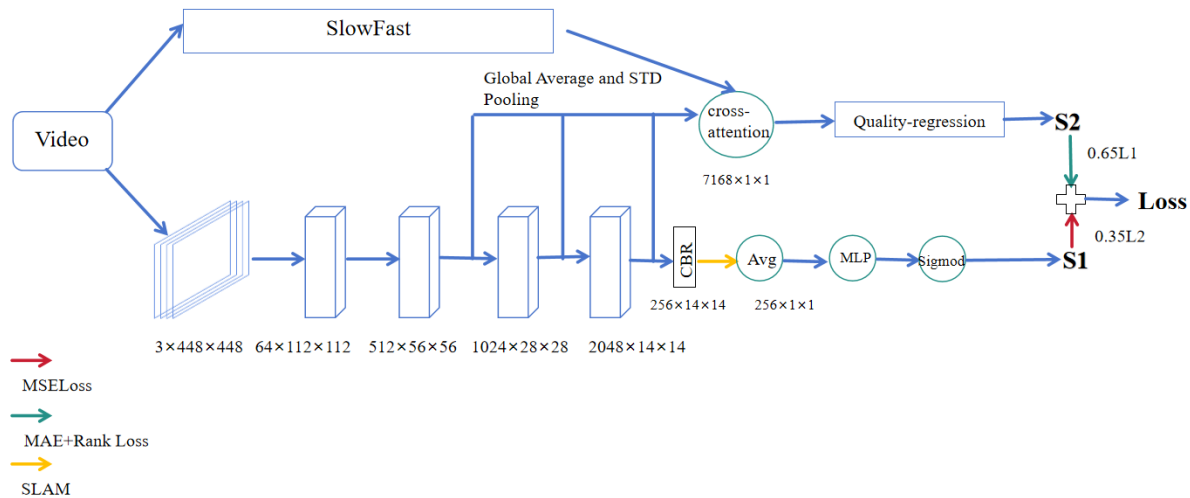


Figure 10. The network architecture of the solution proposed by team We are a team

quality assessment is closely related to the task of video complexity estimation. For example, when high-frequency noise is introduced into the video, the complexity increases, while when low-frequency blur is introduced, the complexity decreases. Therefore, this method introduces complexity pseudo-labels into the KVQ data set, uses a multi-task training method to evaluate the quality of short videos, and forces the backbone network to extract more general and robust features to meet the needs of two tasks at the same time. Their experimental results demonstrate that the performance of the video quality assessment task can also be improved when the video complexity estimation task is introduced.

Similar to most video tasks, their method also has two branches: spatial feature extraction branch and temporal feature extraction branch. For each video, they sample 1 frame per second and adjust its size to 520×520 , crop it to 448×448 and use ResNet50 [15] for spatial feature extraction; sample 32 frames per second and adjust the size to 224×224 and use SlowFast [12] for temporal feature extraction. Use cross-attention on the extracted spatial and temporal features to obtain spatio-temporal fusion features. Finally, the quality assessment head performs video quality score regression based on spatio-temporal fusion features, and the complexity estimation head performs video complexity score regression based on spatial features.

Training Details. In the training phase, they use the Adam optimizer for 100 epochs of end-to-end training, with the initial learning rate and decay weight set to $3e-5$ and 0.9, and the batch size set to 32. For the quality assessment task, MAE loss and Rank loss are used, and the complexity estimation task uses MSE loss.

Testing Details. During the testing phase, for each video, the model only predicts the quality score once as the final

result.

4.12. dulan

This team has explored different VQA methods such as FastVQA [65] and Dover [69]. The empirical results demonstrate that these previously proposed methods can provide satisfactory performance. However, they usually struggle to obtain optimal performance owing to the existence of domain shifts caused by different data distributions. To solve the above issue, this team designed a simple yet effective UGC VQA model based on SimpleVQA [54], which trains an end-to-end spatial feature extractor to learn the quality-aware feature representation.

To enhance the performance of feature extraction, they adopt the SlowFast [12] approach, which is known for its effectiveness in capturing temporal and spatial features in videos. Following the approach in [71], they utilize the Tiny Swin-Transformer [38] as the backbone network, given its superior performance in various visual recognition tasks. Furthermore, they incorporate the Exponential Moving Average (EMA) [20] strategy to improve the model's stability during training, which in turn enhances its performance on the test set.

By leveraging the predictions of multiple models, the strengths of each model can be harnessed to generate a more resilient and precise overall prediction. During the inference stage, they implement an ensemble strategy that combines model parameters from different epochs. This ensemble technique not only diversifies the model's predictions but also significantly contributes to the overall improvement of its performance. The integration of these techniques effectively aggregates multi-view information, thereby enhancing the capabilities of their UGC VQA model, positioning it as a robust and efficient solution for this VQA

challenge.

Training Details. Their model is based on SimpleVQA [54] and they modify the spatial feature extraction module. Specifically, they utilize Swin Transformer [38] pre-trained on ImageNet and SlowFast network [12] pre-trained on Kinetics as the feature extraction model. During the training phase, they employ the EMA strategy to enhance the stability of the model’s convergence, which will be elaborated on the experimental results of the EMA strategy in Table 2. There are 26 millions of parameters in the model. Batch size is set to 4, the learning rate is set to 10^{-4} , and the AdamW optimizer is adopted with a weight decay of 5×10^{-4} . In addition, they use the cosine decay learning rate with the minimum learning rate of 10^{-7} and use linear preheating in the first 2 epochs with start learning rate 5×10^{-7} .

Testing Details. In the test stage, they adjust the model’s weights by assigning a bigger weight to the model with the lowest loss and smaller weights to the models from the subsequent epoch, thereby achieving a fusion of model parameters.

Analysis. (1) **EMA strategy.** They employ the EMA strategy and preserve models both with and without the application of EMA. Table 2 presents a comparison of the performance between these two sets of models. In this comparison, “Model-EMA” refers to the model that employs the EMA strategy, whereas “Model-None” does not utilize this strategy. It is evident from the table that the application of the EMA strategy significantly enhances the performance of the model.

Table 2. Validation Accuracy Comparison between Model-EMA and Model-Without EMA. ↓ donates the lower, the better and ↑ donates the higher, the better. The experimental results are from the team dulan.

Metric	Model-EMA	Model-None
SRCC↑	0.8554	0.3776
PLCC↑	0.8598	0.3979
KRCC↑	0.6646	0.2565
RMSE↓	0.3220	0.6672

(2) **Truncation.** In addition, since the optimal range for the output scores is between 1 and 5, they implemented data truncation for unreasonable outputs, setting negative values to 0. However, this measure inadvertently leads to a decrease in the output scores.

(3) **Ensembles and fusion strategies.** This team argues that no single model can capture all aspects of the data distribution. By combining the predictions of multiple models, the strengths of each model can be leveraged to produce a more robust and accurate overall prediction. They perform a weighted sum fusion based on models saved at differ-

ent epochs. In the training process, model-l represents the model weights with the lowest loss value achieved during the training stage. Additionally, model-s refers to the model weights saved from the previous epoch. The weights p_1 and p_2 are the proportions of the output results from model-l and model-s, respectively. The final score is calculated as follows:

$$Score = p_1 \times W_{model-s} + p_2 \times W_{model-l} \quad (7)$$

Through experiments, they identify the optimal ensemble weight ratio yielding the highest score. The model with the lowest loss is assigned a weight of 0.6, and the model from the subsequent epoch is assigned a weight of 0.4.

4.13. D-H

This team utilizes SimpleVQA [54] model. The S-UGC videos are first sampled and each frame is divided into patches, which are inputted to the spatio-temporal features models and regression module to obtain the predicted score. The loss function is the summation of PLCC loss and ranking loss.

Training Details. They randomly sampled 20% samples from each fraction segment of training samples as a validation set, and the remaining samples were used as the training set. The model weights with the highest evaluation index of the validation set are stored in each epoch. They also utilize the exponential moving average strategy to make the training to be stable and smooth.

Acknowledgements

This work was partially supported by the Humboldt Foundation. We thank the NTIRE 2024 sponsors: Meta Reality Labs, OPPO, Kuaishou, Huawei and University of Würzburg (Computer Vision Lab). We also thank Kuaishou for sponsoring this challenge.

A. Teams and Affiliations

NTIRE2024 Organizers

Title: NTIRE 2024 Challenge on Short-form UGC Video Quality Assessment

Members: Xin Li¹ (lixin666@mail.ustc.edu.cn), Kun Yuan², Yajing Pei¹, Yiting Lu¹, Ming Sun², Chao Zhou², Zhibo Chen¹, and Radu Timofte³

Affiliations:

¹ University of Science and Technology of China

² Kuaishou Technology

³ University of Würzburg

SJTU MMLab

Title: Enhancing Blind Video Quality Assessment with Rich Quality-aware Features

Members:

Wei Sun¹ (sunguwei@sjtu.edu.cn), Haoning Wu², Zicheng Zhang¹, Jun Jia¹, Zhichao Zhang¹, Linha Cao¹, Qiubo Chen², Xiongkuo Min¹, Weisi Lin², and Guangtao Zhai¹.

Affiliations:

¹ Shanghai Jiao Tong University

² Nanyang Technological University

³ Xiaohongshu

IH-VQA

Title: Ensemble-based Video Quality Assessment System

Members:

Jianhui Sun (55530544@qq.com), Tianyi Wang, Lei Li, Han Kong, Wenxuan Wang, Bing Li and Cheng Luo

Affiliations:

WeChat

TVQE

Title: Tencent Video Quality Evaluator

Members:

Haiqiang Wang¹ (walleewang@tencent.com), Xiangguang Chen¹, Wenhui Meng¹, Xiang Pan¹, Huiying Shi², Han Zhu², Xiaozhong Xu¹, Lei Sun¹, Zhenzhong Chen², and Shan Liu¹

Affiliations:

¹ Tencent

² Wuhan University

BDVQAGroup

Title: UGC Video Quality Assessment Method based on LMMs

Members:

Fangyuan Kong (kongfangyuan@bytedance.com), Haotian Fan, Yifang Xu

Affiliations:

ByteDance, Shenzhen

VideoFusion

Title: VideoFusion

Members:

Haoran Xu (xhr964691257@163.com), Mengduo Yang, Jie Zhou, and Jiase Li

Affiliations:

Zhejiang University

MC²Lab

Title: Short-form UGC Video Quality Assessment: A Combined Approach

Members:

Shijie Wen (wenshijie@buaa.edu.cn) and Mai Xu

Affiliations:

Department of Electronic and Information Engineering, Beihang University, Beijing, China

Padding

Title: How Should We Apply Padding?

Members:

Da Li (1048222617@qq.com)

Affiliations:

Southern University of Science and Technology

ysy0129

Title: Video Quality Assessment Based on Self-Supervised Model Priors

Members:

Shunyu Yao (shunyuyao19@gmail.com), Jiazhi Du and Wangmeng Zuo

Affiliations:

Harbin Institute of Technology

lizhibo

Title: Short-form Video Quality Evaluator

Members:

Zhibo Li (lizhibo_zyber@bupt.edu.cn), Shuai He, Anlong Ming, Huiyuan Fu, Huadong Ma

Affiliations:

Beijing University of Posts and Telecommunications, China

YongWu

Title: Data augmentation for the Short-form UGC Video Quality Assessment

Members:

Yong Wu (wuyong139@hotmail.com), Fei Xue and Guozhi Zhao

Affiliations:

China Merchants Bank

We are a team

Title: Short video quality assessment combined with video complexity

Members:

Lina Du (1340398672@qq.com), Jie Guo, Yu Zhang, Huimin Zheng, Junhao Chen and Yue Liu

Affiliations:

Shandong Jianzhu University

dulan

Title: No title

Members:

Dulan Zhou (zd1190301000@163.com), Kele Xu, Qisheng Xu and Tao Sun

Affiliations:

Key Laboratory for Parallel and Distributed Processing, Changsha, China

D-H

Title: No title

Members:

Zhixiang Ding (dingzhixiang23@mails.ucas.ac.cn) and Yuhang Hu

Affiliations:

¹ Institute of Automation, CAS

² University of Chinese Academy of Sciences

References

- [1] Fahad Alrasheedi, Xin Zhong, and Pei-Chi Huang. Padding module: Learning the padding in deep neural networks. *IEEE Access*, 11:7348–7357, 2023. [8](#)
- [2] Cosmin Ancuti, Codruta O Ancuti, Florin-Alexandru Vasluianu, Radu Timofte, et al. NTIRE 2024 dense and non-homogeneous dehazing challenge report. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. [2](#)
- [3] Nikola Banić, Egor Ershov, Artyom Panshin, Oleg Karasev, Sergey Korchagin, Shepelev Lev, Alexandr Startsev, Daniil Vladimirov, Ekaterina Zaychenkova, Dmitrii R Iarchuk, Maria Efimova, Radu Timofte, Arseniy Terekhin, et al. NTIRE 2024 challenge on night photography rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. [2](#)
- [4] Nicolas Chahine, Marcos V. Conde, Sira Ferradans, Radu Timofte, et al. Deep portrait quality assessment. a NTIRE 2024 challenge survey. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. [2](#)
- [5] Baoliang Chen, Lingyu Zhu, Guo Li, Fangbo Lu, Hongfei Fan, and Shiqi Wang. Learning generalized spatial-temporal deep feature representation for no-reference video quality assessment. *IEEE Trans. Circuits Syst. Video Technol.*, 32(4): 1903–1916, 2022. [1](#)
- [6] Zheng Chen, Zongwei WU, Eduard Sebastian Zamfir, Kai Zhang, Yulun Zhang, Radu Timofte, Xiaokang Yang, et al. NTIRE 2024 challenge on image super-resolution (x4): Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. [2](#)
- [7] 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. [4](#)
- [8] Marcos V. Conde, Florin-Alexandru Vasluianu, Radu Timofte, et al. Deep raw image super-resolution. a NTIRE 2024 challenge survey. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. [2](#)
- [9] Lina Du, Jie Guo, Yu Zhang, Huimin Zheng, Junhao Chen, and Yue Liu. Ugc-icvqa: Ugc video no-reference quality evaluation model based on complexity multi-task learning. 2024. [10](#)
- [10] Mahidhar Dwarampudi and NV Reddy. Effects of padding on lstms and cnns. *arXiv preprint arXiv:1903.07288*, 2019. [8](#)
- [11] Yuming Fang, Hanwei Zhu, Yan Zeng, Kede Ma, and Zhou Wang. Perceptual quality assessment of smartphone photography. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3677–3686, 2020. [4](#), [6](#)
- [12] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019. [10](#), [11](#), [12](#)
- [13] 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. [4](#)
- [14] Mahdi Hashemi. Enlarging smaller images before inputting into convolutional neural network: zero-padding vs. interpolation. *Journal of Big Data*, 6(1):1–13, 2019. [8](#)
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [1](#), [10](#), [11](#)
- [16] Vlad Hosu, Franz Hahn, Mohsen Jenadeleh, Hanhe Lin, Hui Men, Tamás Szirányi, Shujun Li, and Dietmar Saupe. The konstanz natural video database (konvid-1k). In *QoMEX*, pages 1–6. IEEE, 2017. [1](#)
- [17] Vlad Hosu, Hanhe Lin, Tamas Sziranyi, and Dietmar Saupe. Koniq-10k: An ecologically valid database for deep learning of blind image quality assessment. *IEEE Transactions on Image Processing*, 29:4041–4056, 2020. [4](#), [6](#)
- [18] Andrew G. Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, Quoc V. Le, and Hartwig Adam. Searching for mobilenetv3. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1314–1324, 2019. [4](#)
- [19] Forrest N. Iandola, Matthew W. Moskewicz, Khalid Ashraf, Song Han, William J. Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 1mb model size. *ArXiv*, abs/1602.07360, 2016. [4](#)
- [20] Frank Klinker. Exponential moving average versus moving exponential average. *Mathematische Semesterberichte*, 58: 97–107, 2011. [4](#), [10](#), [11](#)
- [21] Lingtong Kong, Boyuan Jiang, Donghao Luo, Wenqing Chu, Xiaoming Huang, Ying Tai, Chengjie Wang, and Jie Yang. Ifrnet: Intermediate feature refine network for efficient frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1969–1978, 2022. [5](#)
- [22] Tengchuan Kou, Xiaohong Liu, Wei Sun, Jun Jia, Xiongkuo Min, Guangtao Zhai, and Ning Liu. Stablevqa: A deep no-reference quality assessment model for video stability. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 1066–1076, 2023. [9](#)
- [23] Eric C Larson and Damon M Chandler. Most apparent distortion: full-reference image quality assessment and the role

- of strategy. *Journal of electronic imaging*, 19(1):011006–011006, 2010. 4
- [24] Bowen Li, Weixia Zhang, Meng Tian, Guangtao Zhai, and Xianpei Wang. Blindly assess quality of in-the-wild videos via quality-aware pre-training and motion perception. *IEEE Trans. Circuits Syst. Video Technol.*, 32(9):5944–5958, 2022. 1
- [25] Dingquan Li, Tingting Jiang, and Ming Jiang. Quality assessment of in-the-wild videos. In *ACM Multimedia*, pages 2351–2359. ACM, 2019. 1, 3
- [26] Xin Li, Yiting Lu, and Zhibo Chen. Freqalign: Excavating perception-oriented transferability for blind image quality assessment from a frequency perspective. *IEEE Transactions on Multimedia*, 2023. 1
- [27] Xin Li, Kun Yuan, Yajing Pei, Yiting Lu, Ming Sun, Chao Zhou, Zhibo Chen, Radu Timofte, et al. NTIRE 2024 challenge on short-form UGC video quality assessment: Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. 2
- [28] Jie Liang, Qiaosi Yi, Shuaizheng Liu, Lingchen Sun, Rongyuan Wu, Xindong Zhang, Hui Zeng, Radu Timofte, Lei Zhang, et al. NTIRE 2024 restore any image model (RAIM) in the wild challenge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. 2
- [29] 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)*, pages 1–3. IEEE, 2019. 4, 6
- [30] Guilin Liu, Aysegul Dundar, Kevin J Shih, Ting-Chun Wang, Fitsum A Reda, Karan Sapra, Zhiding Yu, Xiaodong Yang, Andrew Tao, and Bryan Catanzaro. Partial convolution for padding, inpainting, and image synthesis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5):6096–6110, 2022. 8
- [31] Hongbo Liu, Mingda Wu, Kun Yuan, Ming Sun, Yansong Tang, Chuanchuan Zheng, Xing Wen, and Xiu Li. Ada-dqa: Adaptive diverse quality-aware feature acquisition for video quality assessment. *CoRR*, abs/2308.00729, 2023. 1
- [32] Hongbo Liu, Mingda Wu, Kun Yuan, Ming Sun, Yansong Tang, Chuanchuan Zheng, Xing Wen, and Xiu Li. Ada-dqa: Adaptive diverse quality-aware feature acquisition for video quality assessment. In *ACM Multimedia*, pages 6695–6704. ACM, 2023. 1
- [33] Jianzhao Liu, Xin Li, Shukun An, and Zhibo Chen. Source-free unsupervised domain adaptation for blind image quality assessment. *arXiv preprint arXiv:2207.08124*, 2022. 1
- [34] Jianzhao Liu, Xin Li, Yanding Peng, Tao Yu, and Zhibo Chen. Swiniqa: Learned swin distance for compressed image quality assessment. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 1795–1799, 2022. 1
- [35] Jianzhao Liu, Wei Zhou, Xin Li, Jiahua Xu, and Zhibo Chen. Liqa: Lifelong blind image quality assessment. *IEEE Transactions on Multimedia*, 2022. 1
- [36] Xiaohong Liu, Xiongkuo Min, Guangtao Zhai, Chunyi Li, Tengchuan Kou, Wei Sun, Haoning Wu, Yixuan Gao, Yuqin Cao, Zicheng Zhang, Xiele Wu, Radu Timofte, et al. NTIRE 2024 quality assessment of AI-generated content challenge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. 2
- [37] Xiaoning Liu, Zongwei WU, Ao Li, Florin-Alexandru Vasluianu, Yulun Zhang, Shuhang Gu, Le Zhang, Ce Zhu, Radu Timofte, et al. NTIRE 2024 challenge on low light image enhancement: Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. 2
- [38] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 11, 12
- [39] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, et al. Swin transformer v2: Scaling up capacity and resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12009–12019, 2022. 7, 8
- [40] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11976–11986, 2022. 4, 5, 7, 8, 9
- [41] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *CVPR*, pages 3192–3201. IEEE, 2022. 1, 3
- [42] Yiting Lu, Jun Fu, Xin Li, Wei Zhou, Sen Liu, Xinxin Zhang, Wei Wu, Congfu Jia, Ying Liu, and Zhibo Chen. Rtn: Reinforced transformer network for coronary ct angiography vessel-level image quality assessment. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 644–653. Springer, 2022. 1
- [43] Yiting Lu, Xin Li, Jianzhao Liu, and Zhibo Chen. Styleam: Perception-oriented unsupervised domain adaption for non-reference image quality assessment. *arXiv preprint arXiv:2207.14489*, 2022. 1
- [44] Yiting Lu, Xin Li, Bingchen Li, Zihao Yu, Fengbin Guan, Xinrui Wang, Ruling Liao, Yan Ye, and Zhibo Chen. Aigc-vqa: A holistic perception metric for aigc video quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2024. 1
- [45] Yiting Lu, Xin Li, Yajing Pei, Kun Yuan, Qizhi Xie, Yunpeng Qu, Ming Sun, Chao Zhou, and Zhibo Chen. Kvq: Kwai video quality assessment for short-form videos. *CVPR*, 2024. 1, 2, 5, 10
- [46] Naila Murray, Luca Marchesotti, and Florent Perronnin. Ava: A large-scale database for aesthetic visual analysis. In *2012 IEEE conference on computer vision and pattern recognition*, pages 2408–2415. IEEE, 2012. 4, 6
- [47] HamidReza Naseri and Vahid Mehrdad. Novel cnn with investigation on accuracy by modifying stride, padding, kernel size and filter numbers. *Multimedia Tools and Applications*, 82(15):23673–23691, 2023. 8

- [48] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. **9**
- [49] Adrien Pavao, Isabelle Guyon, Anne-Catherine Letournel, Dinh-Tuan Tran, Xavier Baro, Hugo Jair Escalante, Sergio Escalera, Tyler Thomas, and Zhen Xu. Codalab competitions: An open source platform to organize scientific challenges. *Journal of Machine Learning Research*, 24(198):1–6, 2023. **2**
- [50] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. **1**
- [51] Bin Ren, Yawei Li, Nancy Mehta, Radu Timofte, et al. The ninth NTIRE 2024 efficient super-resolution challenge report. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. **2**
- [52] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115:211–252, 2014. **5**
- [53] Hamid R Sheikh, Muhammad F Sabir, and Alan C Bovik. A statistical evaluation of recent full reference image quality assessment algorithms. *IEEE Transactions on image processing*, 15(11):3440–3451, 2006. **4**
- [54] Wei Sun, Xiongkuo Min, Wei Lu, and Guangtao Zhai. A deep learning based no-reference quality assessment model for UGC videos. In *ACM Multimedia*, pages 856–865. ACM, 2022. **1, 3, 7, 9, 10, 11, 12**
- [55] Wei Sun, Wen Wen, Xiongkuo Min, Long Lan, Guangtao Zhai, and Kede Ma. Analysis of video quality datasets via design of minimalistic video quality models. *arXiv preprint arXiv:2307.13981*, 2023. **3**
- [56] Wei Sun, Haoning Wu, Zicheng Zhang, Jun Jia, Zhichao Zhang, Linhan Cao, Qiubo Chen, Xiongkuo Min, Weisi Lin, and Zhai Guangtao. Enhancing blind video quality assessment with rich quality-aware features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2024. **4**
- [57] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems*, 35:10078–10093, 2022. **8, 9**
- [58] Florin-Alexandru Vasluianu, Tim Seizinger, Zhuyun Zhou, Zongwei WU, Cailian Chen, Radu Timofte, et al. NTIRE 2024 image shadow removal challenge report. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. **2**
- [59] haiqiang Wang, Gary Li, Shan Liu, and C.-C. Jay Kuo. Icme 2021 ugc-vqa challenge. In *Available: <http://ugcvqa.com/>*. **6**
- [60] Longguang Wang, Yulan Guo, Juncheng Li, Hongda Liu, Yang Zhao, Yingqian Wang, Zhi Jin, Shuhang Gu, Radu Timofte, et al. NTIRE 2024 challenge on stereo image super-resolution: Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. **2**
- [61] Yilin Wang, Sasi Inguva, and Balu Adsumilli. Youtube UGC dataset for video compression research. In *MMSP*, pages 1–5. IEEE, 2019. **1**
- [62] Yingqian Wang, Zhengyu Liang, Qianyu Chen, Longguang Wang, Jungang Yang, Radu Timofte, Yulan Guo, et al. NTIRE 2024 challenge on light field image super-resolution: Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. **2**
- [63] Sanghyun Woo, Shoubhik Debnath, Ronghang Hu, Xinlei Chen, Zhuang Liu, In-So Kweon, and Saining Xie. Convnext v2: Co-designing and scaling convnets with masked autoencoders. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16133–16142, 2023. **4**
- [64] Bichen Wu, Chenfeng Xu, Xiaoliang Dai, Alvin Wan, Peizhao Zhang, Zhicheng Yan, Masayoshi Tomizuka, Joseph E Gonzalez, Kurt Keutzer, and Peter Vajda. Visual transformers: Where do transformers really belong in vision models? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 599–609, 2021. **1**
- [65] Haoning Wu, Chaofeng Chen, Jingwen Hou, Liang Liao, Annan Wang, Wenxiu Sun, Qiong Yan, and Weisi Lin. Fastvqa: Efficient end-to-end video quality assessment with fragment sampling. In *European conference on computer vision*, pages 538–554. Springer, 2022. **1, 3, 4, 7, 10, 11**
- [66] Haoning Wu, Liang Liao, Chaofeng Chen, Jingwen Hou, Annan Wang, Wenxiu Sun, Qiong Yan, and Weisi Lin. Disentangling aesthetic and technical effects for video quality assessment of user generated content. *CoRR*, abs/2211.04894, 2022. **1, 10**
- [67] Haoning Wu, Chaofeng Chen, Liang Liao, Jingwen Hou, Wenxiu Sun, Qiong Yan, and Weisi Lin. Discovqa: Temporal distortion-content transformers for video quality assessment. *IEEE Trans. Circuits Syst. Video Technol.*, 33(9):4840–4854, 2023. **1**
- [68] Haoning Wu, Erli Zhang, Liang Liao, Chaofeng Chen, Jingwen Hou, Annan Wang, Wenxiu Sun, Qiong Yan, and Weisi Lin. Towards explainable in-the-wild video quality assessment: A database and a language-prompted approach. In *Proceedings of the 31st ACM International Conference on Multimedia, MM 2023, Ottawa, ON, Canada, 29 October 2023- 3 November 2023*, pages 1045–1054. ACM, 2023. **1**
- [69] Haoning Wu, Erli Zhang, Liang Liao, Chaofeng Chen, Jingwen Hou, Annan Wang, Wenxiu Sun, Qiong Yan, and Weisi Lin. Exploring video quality assessment on user generated contents from aesthetic and technical perspectives. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20144–20154, 2023. **5, 11**
- [70] Haoning Wu, Zicheng Zhang, Weixia Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Yixuan Gao, Annan Wang,

- Erli Zhang, Wenxiu Sun, et al. Q-align: Teaching lms for visual scoring via discrete text-defined levels. *arXiv preprint arXiv:2312.17090*, 2023. 1, 3, 5, 6
- [71] Wei Wu, Shuming Hu, Pengxiang Xiao, Sibin Deng, Yilin Li, Ying Chen, and Kai Li. Video quality assessment based on swin transformer with spatio-temporal feature fusion and data augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1846–1854, 2023. 11
- [72] Fengchuang Xing, Yuan-Gen Wang, Hanpin Wang, Leida Li, and Guopu Zhu. Starvqa: Space-time attention for video quality assessment. In *2022 IEEE International Conference on Image Processing, ICIP 2022, Bordeaux, France, 16-19 October 2022*, pages 2326–2330. IEEE, 2022. 1
- [73] Haoran Xu, Jie Zhou, Mengduo Yang, and Jiaye Li. Short-form ugc video quality assessment based on multi-level video fusion with rank-aware. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2024. 7
- [74] Ren Yang, Radu Timofte, et al. NTIRE 2024 challenge on blind enhancement of compressed image: Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. 2
- [75] Zhenqiang Ying, Maniratnam Mandal, Deepti Ghadiyaram, and Alan Bovik. Patch-vq: patching up the video quality problem. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14019–14029, 2021. 1, 4, 6
- [76] Zihao Yu, Fengbin Guan, Yiting Lu, Xin Li, , and Zhibo Chen. Sf-iqa: Quality and similarity integration for ai generated image quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2024. 1
- [77] Zihao Yu, Fengbin Guan, Yiting Lu, Xin Li, and Zhibo Chen. Video quality assessment based on swin transformerv2 and coarse to fine strategy. *arXiv preprint arXiv:2401.08522*, 2024. 1
- [78] Kun Yuan, Zishang Kong, Chuanchuan Zheng, Ming Sun, and Xing Wen. Capturing co-existing distortions in user-generated content for no-reference video quality assessment. In *ACM Multimedia*, pages 1098–1107. ACM, 2023. 1
- [79] Pierluigi Zama Ramirez, Fabio Tosi, Luigi Di Stefano, Radu Timofte, Alex Costanzino, Matteo Poggi, et al. NTIRE 2024 challenge on HR depth from images of specular and transparent surfaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. 2
- [80] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11941–11952, 2023. 4
- [81] 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*, pages 14071–14081, 2023. 3
- [82] Zhilu Zhang, Shuohao Zhang, Renlong Wu, Wangmeng Zuo, Radu Timofte, et al. NTIRE 2024 challenge on bracketing image restoration and enhancement: Datasets, methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2024. 2
- [83] Kai Zhao, Kun Yuan, Ming Sun, Mading Li, and Xing Wen. Quality-aware pretrained models for blind image quality assessment. In *CVPR*, pages 22302–22313. IEEE, 2023. 1
- [84] Kai Zhao, Kun Yuan, Ming Sun, and Xing Wen. Zoom-vqa: Patches, frames and clips integration for video quality assessment. In *CVPR Workshops*, pages 1302–1310. IEEE, 2023. 1