

Zoom-VQA: Patches, Frames and Clips Integration for Video Quality Assessment

Kai Zhao, Kun Yuan, Ming Sun and Xing Wen
Kuaishou Technology

{zhaokai05, yuankun03, sunming03, wenxing}@kuaishou.com

Abstract

Video quality assessment (VQA) aims to simulate the human perception of video quality, which is influenced by factors ranging from low-level color and texture details to high-level semantic content. To effectively model these complicated quality-related factors, in this paper, we decompose video into three levels (i.e., patch level, frame level, and clip level), and propose a novel Zoom-VQA architecture to perceive spatio-temporal features at different levels. It integrates three components: patch attention module, frame pyramid alignment, and clip ensemble strategy, respectively for capturing region-of-interest in the spatial dimension, multi-level information at different feature levels, and distortions distributed over the temporal dimension. Owing to the comprehensive design, Zoom-VQA obtains state-of-the-art results on four VQA benchmarks and achieves 2nd place in the NTIRE 2023 VQA challenge. Notably, Zoom-VQA has outperformed the previous best results on two subsets of LSVQ, achieving 0.8860 (+1.0%) and 0.7985 (+1.9%) of SRCC on the respective subsets. Adequate ablation studies further verify the effectiveness of each component. Codes and models are released in <https://github.com/k-zha14/Zoom-VQA>.

1. Introduction

To achieve the ultimate viewing experience, streaming platforms often employ diverse enhancement techniques to ameliorate the quality of uploaded content [15, 52, 53]. The primary phase in this endeavor involves precise evaluation of the perceptual quality of the content, which can be achieved by subjective and objective quality assessment measures. Given that subjective quality assessment experiments can be time-consuming, expensive, and complex [5, 14], both industry and academia are engaged in the pursuit of establishing more scalable and feasible objective quality assessment methods, capable of managing ever-growing volumes of user-generated content [11].

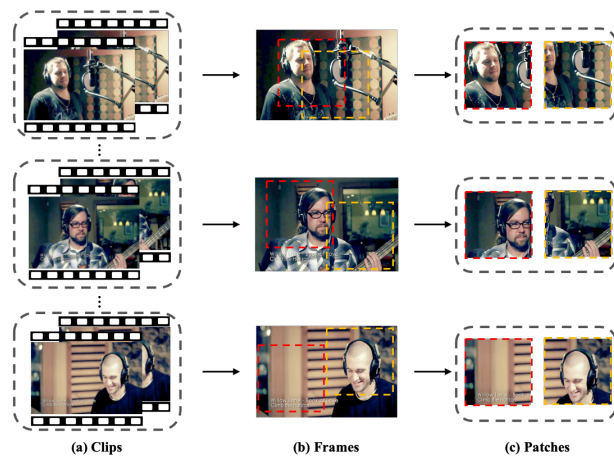


Figure 1. The decomposition of a typical video into three levels, (a) clips, (b) frames, and (c) patches. Videos captured at different levels contain various factors that affect video quality. For instance, motion blur is reflected in the clip level, semantic information is contained in the frame level, and local block artifacts and noise are manifested in the patch level.

Generally, VQA can be classified into three categories: full-reference, reduced reference, and no-reference, based on the availability of reference information [29]. Despite evident differences in the reference feature fusion, almost all VQA methods typically follow a conventional paradigm [2], that is, extracting visual features related to perceived quality and designing regression or classification head for quality prediction from the extracted features.

Established on this paradigm, the success of VQA methods hinges on two aspects: 1) extracting more powerful and representative features, including advanced backbones and modules [21, 49], quality-aware pre-training [30, 61] and multiple hand-crafted features ensemble [12, 45], and 2) designing quality score prediction head, such as normal regression heads [40, 45], patches fusion [59] and IP-NLR [49]. Additionally, some studies explore patches [59], frames [40] and clip integration [49] separately, inspired by locality and sensitivity of human visual system (HVS). Along with state-of-the-art deep neural networks [27, 28],

we take a step further by integrating patches, frames and clips together. We claim that the performance of VQA methods could be further improved through this method.

Specifically, we treat a video from a *zooming in* perspective. As shown in Fig. 1, a typical video can be decomposed into clips, frames and patches progressively. However, due to the absence of fine-grained labels (*e.g.* clip-level, frame-level and patch-level) in existing VQA datasets [17, 48], the multi-level supervised training and integration is practically difficult. Besides, with the popularization of deep learning-based video enhancement methods [24, 47], nonexistent artifacts are generated to improve perceptual aspect of processed videos. Consequently, the precise prediction of perceived quality has become even more challenging.

To overcome the above difficulties, we propose Zoom-VQA, a multi-level integration framework for VQA. The contributions of this framework are summarized as follows:

- To capture region-of-interest in the spatial dimension that affects quality, a *patch attention module* is proposed that generates weights and scores for patches.
- To obtain distortions that are reflected in different feature levels, a *frame pyramid alignment* module is proposed to integrate multi-scale information.
- To model spatio-temporal information effectively, a *clip ensemble strategy* is utilized to aggregate temporal information adaptively.
- Zoom-VQA obtains state-of-the-art results on four VQA benchmarks and achieves **2nd place** in the NTIRE 2023 VQA challenge. Notably, Zoom-VQA has outperformed the previous best results on two subsets of LSVQ, achieving **0.8860** (+1.0%) and **0.7985** (+1.9%) of SRCC on the respective subsets.

2. Related work

2.1. Video Quality Assessment

The fundamental goal of objective VQA is to accurately predict the perceived quality of human viewers. To this end, a naive solution is to apply classic image quality assessment metrics [31, 33, 58] on each video frame, and then perform a feature/score pooling stage. Since the efficacy of temporal pooling is temporal-dependent [44], it is non-trivial to incorporate spatio-temporal information. V-BLINDS [36] proposes a spatio-temporal natural scene statistics (NSS) model by quantifying the frame difference of NSS features. VIIDEO [32] exploits the inherent statistical characteristics of natural videos to deal with distortion-specific disturbances. TLVQM [19] and VIDEVAL [45] combine abundant spatio-temporal features and typical quality assessment methods to regress them into the video quality score. However, these hand-crafted features are not sufficient to capture

complicated factors affecting video quality, such as semantic information. Subsequent works [20, 46] attempt to integrate hand-crafted features with semantic features extracted by shallow pre-trained CNN models.

Recently, deep learning-based VQA methods have become dominant in the VQA field, with the rise of large-scale VQA datasets [17, 48, 59]. Among them, VSFA [22] utilizes the pre-trained ResNet-50 [16] from ImageNet-1k [7] to extract spatial features, and then models the temporal relationship by GRU [6]. Based on this, MDVSFA [23] explores a mixed dataset training strategy with VSFA to mitigate the cross-dataset evaluation challenge. MLSP-FF [13] adopts deeper and heavier Inception-ResNet-V2 for feature extraction. Bi-LSTM [37] and graph convolution network [18] are also introduced into the VQA field, to enhance temporal feature extraction [56]. In addition, some methods [21, 50, 59] use video models (*e.g.* 3D-CNN [4, 10]) pre-trained on action recognition datasets to extract spatio-temporal features. However, limited by high computational cost on high-resolution videos, the above methods usually regress the subjective quality scores with fixed features instead of end-to-end training [21]. Aimed at this issue, FAST-VQA [49] and DOVER [51] propose grid mini-patch sampling (GMS) strategy, which samples patches at their raw resolution to maintain the global quality with the original video. In order to accommodate high-resolution inputs of the challenge dataset, implementation and modification of GMS will be further discussed in Sec. 3.4.

2.2. Advancement in Visual Networks

The advancement of visual backbone networks prospers various downstream tasks tremendously [3, 25, 28, 55, 60]. Based on the modality of input data, visual networks can be classified into two types, image networks (*i.e.* 2D networks) and video networks (*i.e.* 3D networks). C3D [42] pioneeringly devises an 11-layer CNN with 3D-CNN to adapt to video inputs. Subsequent P3D [35], S3D [54] and R(2+1)D [43] observe that disentangled spatial and temporal convolutions results in a more favorable speed-accuracy trade-off than the pure 3D convolution. Remarkably, I3D [4] reveals that successful 2D networks could be seamlessly *inflated* to corresponding 3D networks and even their parameter.

In recent days, a shift in backbone architecture, from CNNs to Transformers (ViT) [8, 41], has begun. Especially, Swin Transformer [26] reintroduces the inductive bias of convolutions (*i.e.*, locality, translation invariance and hierarchy), which enables it to serve as a general-purpose backbone. The success of image Transformer leads to further investigation of Transformer-based video networks (*e.g.*, ViViT [1], MViT [9], Video Swin Transformer [28]). Among all characteristics of Transformers, the patch-wise operations inherently differentiate the edges of patches, thus making them ideal for handling input sampled by GMS.

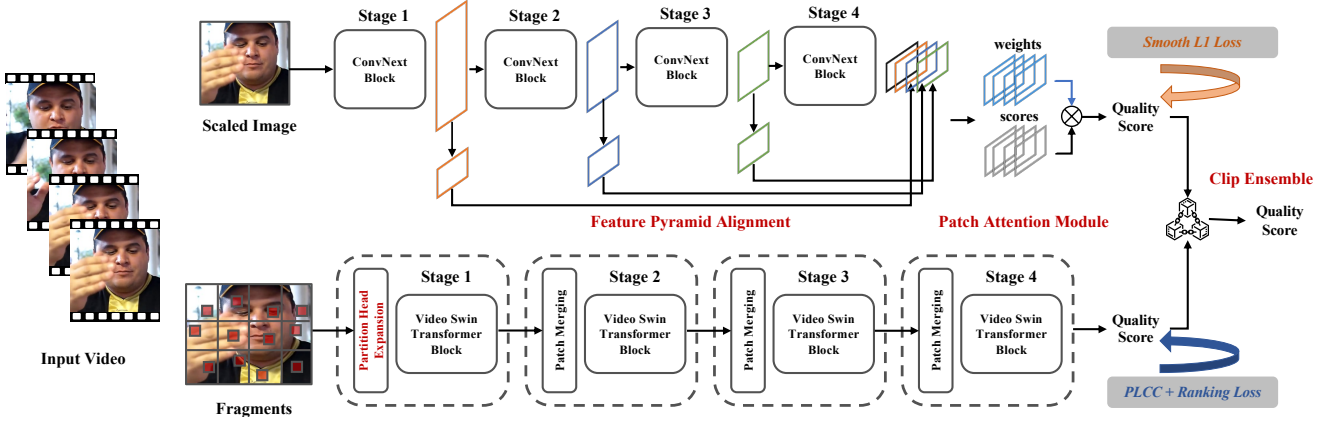


Figure 2. Illustration of the framework of Zoom-VQA. The overall architecture consists of two parts. One part is a perception network based on image input that obtains global information. The other part is a perception network based on video clip input, obtaining texture information by sampling fragment regions. During training, the two parts are trained separately using different optimization objectives. During inference, we use an ensemble to combine the prediction results of the two parts to enhance the model’s generalization ability.

3. Method

We propose a novel Zoom-VQA network architecture that evaluates video quality by comprehensively analyzing the “patch-frame-clip” dimensions of the video. Through this approach, Zoom-VQA can effectively capture both local and global information that impacts video quality, making it more suitable for assessing enhanced videos. In this section, we first give the overall framework in Sec. 3.1 and adopt a perspective that gradually “zooms out” from the micro to the macro level to introduce Zoom-VQA, covering a *patch attention module* which considers qualities of different spatial regions in Sec. 3.2, a *frame pyramid alignment module* which tackles multi-level feature information in Sec. 3.3 and a *clip ensemble strategy* which utilizes the spatio-temporal information fully in Sec. 3.4. Last, the optimization objective will be given in Sec. 3.5.

3.1. Overall Framework

As given in Fig. 2, the overall framework of Zoom-VQA consists of two parts. Given an input video \mathcal{V} , one part is a perception network based on image input (named IQA branch) \mathbf{x}_t , which denotes the t -th frame sampled from \mathcal{V} , that obtains global information using ConvNext [27] as the backbone, analyzes the video frame by frame, and outputs the average results as the final quality prediction score. The other part is a perception network based on video segment input (named VQA branch) $\mathcal{G} = \{\mathbf{g}_t, \dots, \mathbf{g}_T\}$ (will be described below), using Video Swin Transformer [28] as the backbone, and obtaining spatio-temporal information by sampling fragment regions as FAST-VQA [49], outputting the overall quality prediction score. The final prediction of the quality score y' is obtained through an ensemble strategy, enhancing the model’s generalization ability.

3.2. Patch Attention Module

To comprehensively consider the quality of different regions in the spatial dimension, we added a patch-based attention module to the head of the IQA branch. Specifically, as given in Fig. 2, this module receives extracted features as input and generates *weights* and *scores* through two separate branches consisting of lightweight networks. The quality score of each patch in the feature map can be obtained by multiplication of corresponding score and weight in a self-attention manner. Then the final prediction of the whole image is generated by summarizing all patches. Given the input features $\mathbf{f}_t \in \mathbb{R}^{c \times h \times w}$, this can be noted as follows:

$$\begin{aligned} \mathbf{w}_{\text{patch}} &= \text{ReLU}(\text{FC}(\text{ReLU}(\text{FC}(\mathbf{f}_t))))), \\ \mathbf{s}_{\text{patch}} &= \text{Sigmoid}(\text{FC}(\text{ReLU}(\text{FC}(\mathbf{f}_t))))), \\ y_t^{\text{IQA}} &= \sum \mathbf{w}_{\text{patch}} \odot \mathbf{s}_{\text{patch}}, \end{aligned} \quad (1)$$

where $\mathbf{w}_{\text{patch}} \in \mathbb{R}^{c \times h \times w}$ is the learned weights for each patch, and $\mathbf{s}_{\text{patch}} \in \mathbb{R}^{c \times h \times w}$ is corresponding scores. The \odot indicates the Hadamard product and y_t^{IQA} is the prediction score for the t -th frame. By this explicit constraint, the network can pay more attention to region-of-interest in the frame, capturing features about distortions that are non-uniformly distributed across spatial dimensions.

3.3. Frame Pyramid Alignment

Video quality is determined by multiple aspects, including content, distortions, and compression artifacts [48]. The factors that may affect quality have different responses to features at different scales. For example, content-related semantic information is usually reflected in high-level features, while distortion-related information such as texture and color is more reflected in low-level features. To ob-

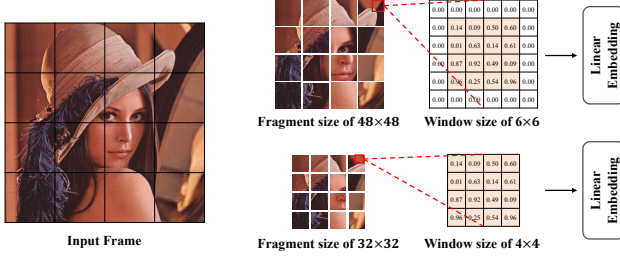


Figure 3. Patch head expansion for the larger size of fragments. The window size is expanded from 4 to 6 with zero padding during the weight initialization phase. In this type, the change in fragment size can be well-compatible with the pre-trained weights.

tain these multi-level features, we design a frame pyramid alignment module in the IQA branch, performing a bottom-up pathway. As given in Fig. 2, features at different stages are extracted, resulting in $\mathbf{f}_l \in \mathbb{R}^{c_l \times h_l \times w_l}$ for the l -th stage. Then, an adaptive average pooling operation is performed on the spatial dimension to align the spatial scale with the last stage of features. Lastly, the aligned features from different stages are concatenated and sent into the patch attention module. Given the ConvNext-Tiny with 4 stages, this procedure can be noted as follows:

$$\begin{aligned} \mathbf{f}'_l &= \text{AdaptiveAvgPool}(\mathbf{f}_l, h_l/h_4)|_{i=1}^3, \\ &\text{where } \mathbf{f}'_l \in \mathbb{R}^{c_l \times h_4 \times w_4}, \\ \mathbf{f}'_t &= \text{Concatenate}\{\mathbf{f}'_1, \mathbf{f}'_2, \mathbf{f}'_3, \mathbf{f}'_4\}, \\ &\text{where } \mathbf{f}'_t \in \mathbb{R}^{\sum_{i=1}^4 c_i \times h_4 \times w_4}, \end{aligned} \quad (2)$$

where h_l/h_4 is the pooling ratio. Then the input \mathbf{f}_t for the patch attention module in Eq. (1) can be replaced by \mathbf{f}'_t with more abundant features.

3.4. Clip Ensemble Strategy

To reflect the impact caused by temporal jitters, we employed a clip ensemble strategy to attain the final prediction quality score of a video. For the IQA branch, the video-level prediction is obtained by averaging the predicted scores for all video frames, which is formulated as:

$$y^{\text{IQA}} = \frac{1}{T} \sum_{t=1}^T y_t^{\text{IQA}}. \quad (3)$$

For more effective spatio-temporal information extraction, we introduce a Video Swin Transformer based on video clip input in the VQA branch. To preserve the original video quality and obtain local texture information which benefits quality assessment, we utilize the sampling strategy of fragments used in FAST-VQA. The fragments are obtained through uniform grid mini-patch sampling. This method vastly reduces the computational cost by 97.6%

compared with computing attention on the whole resolution. It allows the network to pay more attention to low-quality texture information (e.g., noise, blur, blocky artifacts, etc.) and reduces its focus on higher-level semantics.

In practice, we found that appropriately increasing the size of the fragment helps to better model local information (e.g., expanding the size from 32 to 48), but changes in size can prevent the subsequent tokenization from utilizing existing pre-trained weights (i.e., mismatching of window size). To address this issue, we propose a patch head expansion module that performs zero-padding around existing convolutional kernels of tokenization as shown in Fig. 3. This approach enables the pre-trained weights to be effectively utilized while increasing the receptive field for extracting low-level information. Then the expanded fragments are sent to the subsequent Video Swin Transformer, generating the quality prediction of y^{VQA} .

Finally, the prediction results of these two branches are aggregated to obtain the final video quality. To ensure consistency in the distribution of the predicted range, we use the sigmoid function for normalization. The clip ensemble strategy for the final quality prediction can be noted as:

$$y' = \frac{1}{2} (\text{Sigmoid}(y^{\text{IQA}}) + \text{Sigmoid}(y^{\text{VQA}})). \quad (4)$$

3.5. Optimization Objective

During training, to ensure convergence, the two branches of Zoom-VQA are **trained separately using different optimization objectives**. For the IQA branch, the network is optimized using a smooth \mathcal{L}_1 regression loss. Given the labeled MOS y , the optimization objective can be noted as:

$$\min \mathcal{L}_{reg} = \begin{cases} 0.5(y - y^{\text{IQA}})^2, & \text{if } |y - y^{\text{IQA}}| < 1 \\ |y - y^{\text{IQA}}| - 0.5, & \text{otherwise.} \end{cases} \quad (5)$$

For the VQA branch, a PLCC-induced loss and a ranking-based loss are utilized. Assume we have m videos on the training batch. Given the predicted quality scores $\mathbf{Y}^{\text{VQA}} = \{y_1^{\text{VQA}}, y_2^{\text{VQA}}, \dots, y_m^{\text{VQA}}\}$ and corresponding MOS values $\mathbf{Y} = \{y_1, y_2, \dots, y_m\}$, the PLCC-induced loss is defined as:

$$\mathcal{L}_{plcc} = \left(1 - \frac{\sum_{i=1}^m (y_i^{\text{VQA}} - a^{\text{VQA}})(y_i - a)}{\sqrt{\sum_{i=1}^m (y_i^{\text{VQA}} - a^{\text{VQA}})^2 \sum_{i=1}^m (y_i - a)^2}} \right) / 2, \quad (6)$$

where a' and a are the mean values of \mathbf{Y}^{VQA} and \mathbf{Y} respectively. And the ranking-based loss can be denoted as:

$$\mathcal{L}_{rank} = \frac{1}{m^2} \sum_{i=1}^m \sum_{j=1}^m \max(0, |y_i - y_j| - e(y_i, y_j) \cdot (y_i^{\text{VQA}} - y_j^{\text{VQA}})), \quad (7)$$

where $e(y_i, y_j)$ is 1 if $y_i \geq y_j$, else is -1 if $y_i < y_j$. And the optimization objective can be written as:

$$\min \mathcal{L}_{plcc} + \beta \cdot \mathcal{L}_{rank}, \quad (8)$$

where β is the balancing coefficient which is set to 0.3.

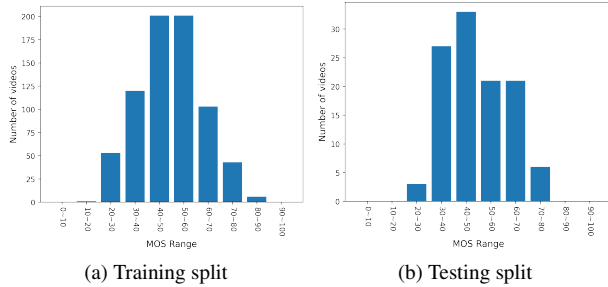


Figure 4. Distribution of labeled MOS in the custom divided train-test splits of VDPVE.




Figure 5. Distribution of adopted video enhancement methods in the custom divided train-test splits of VDPVE.

4. Experiments

4.1. Datasets and Evaluation Criteria

Datasets. Four NR-VQA datasets are selected for evaluation, including VDPVE [11], LSVQ [38], KoNViD-1k [17] and LIVE-VQC [39]. In detail, VDPVE has 1,211 videos with different enhancements, enabling the quality assessment of enhanced videos. Specifically, they can be divided into three sub-datasets. The first sub-dataset consists of 600 videos with color, brightness, and contrast enhancements. The second sub-dataset consists of 310 videos with deblurring. The third sub-dataset consists of 301 deshaked videos. LSVQ, which is the largest VQA dataset by far, samples hundreds of thousands of open-source Internet UGC digital videos to match the feature distributions of social media UGC videos, collecting 39,000 real-world videos of diverse sizes, contents, and distortions. LSVQ also provides two official test subsets, including $LSVQ_{test}$ and $LSVQ_{1080p}$. The $LSVQ_{test}$ consists of 7,400 various resolution videos from 240P to 720P, and $LSVQ_{1080p}$ consists of 3,600 1080P high resolution videos. KoNViD-1k contains 1,200 videos that are fairly filtered from a large public video dataset YFCC100M. The videos are 8 seconds long with 24/25/30 FPS and a resolution of 960×540 . Besides, LIVE-VQC consists of 585 videos with complex authentic distortions captured by 80 different users using 101 different devices.

In our experiments, VDPVE and LSVQ are used for training and evaluation. While KoNViD-1k and LIVE-VQC are only used for cross-dataset evaluation. The train-test

split of LSVQ is aligned with previous studies [49, 59]. As for VDPVE, we conducted ablation studies based on a custom train-test split, as the official validation and testing splits for the NTIRE 2023 VQA Challenge have not yet been released. Particularly, we carefully divide the annotated 839 enhanced videos into two parts, 728 videos for training and 111 videos for testing. As shown in Figs. 4 and 5, the distributions of our train-test splits are consistent in MOS and video enhancement methods.

Evaluation Criteria. Spearman’s Rank-Order Correlation Coefficient (SRCC) and Pearson’s Linear Correlation Coefficient (PLCC) are selected as criteria to measure the accuracy and monotonicity, respectively. They are in the range of $[0.0, 1.0]$. A larger SRCC shows a more accurate ranking between different samples. A larger PLCC means a more accurate numerical fit with MOS scores. Besides, the mean average of PLCC and SRCC (named main score) is also reported as a comprehensive criterion.

4.2. Implementation Details

Our experiments are conducted on an NVIDIA V100 GPU with PyTorch 1.12 [34] and CUDA 11.3 for training and testing. All ablation studies are performed on the custom-divided train-test splits of VDPVE. As stated in Sec. 3, the two branches of Zoom-VQA are trained independently and evaluated in a late-fusion manner. Below, we describe relevant settings in turn:

IQA branch. We choose the pre-trained ConvNext-Tiny from ImageNet-1k as the backbone. Following the standard training strategy [57, 61], we sample frames at 2 fps and resize the smaller edge of images to 512 while maintaining the aspect ratio. During training, images are cropped at a random size of 320×320 and subject to horizontal flipping with a probability of 0.5. We utilize the AdamW optimizer with a learning rate of 0.002, a weight decay of 0.01 and a cosine annealing scheduler. The mini-batch size is set to 32. During testing, the center-cropped patches of all extracted frames are used for evaluation.

VQA branch. We inherit most of the settings from FAST-VQA while enlarging the size of fragments and patch partition head. Specifically, we expand the path size from 4 to 6 and the corresponding fragment size from 32 to 48. During training, we choose the Video Swin Transformer-Tiny which is pre-trained on LSVQ. The AdamW optimizer is used for 30-epoch training, with a weight decay of 0.01 and a mini-batch size of 16. The initial global learning rate is 0.001 and decayed by a cosine annealing schedule. We sample a clip of 32 frames from each full-length video using a temporal stride of 2 and spatial size of 336×336 . During testing, we infer 4 views for each video, and the final score is computed as the average score of all views.

Table 1. Comparison with SOTA methods on public VQA benchmarks. LSVQ_{test} and LSVQ_{1080p} are official intra-dataset test subsets. While KoNViD-1k and LIVE-VQC are used for cross-dataset testing. The best and second best results are **bolded** and underlined. Note that our reproduced results of FAST-VQA are a little lower than the official ones. We speculate that the reason for the difference is the variation in the positions of the sampled fragments (both in the spatial and temporal dimensions).

Testing Type	Intra-dataset Test Sets				Cross-dataset Test Sets			
Testing Set	LSVQ _{test}		LSVQ _{1080p}		KoNViD-1k		LIVE-VQC	
Methods	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC
BRISQUE [31]	0.569	0.576	0.497	0.531	0.646	0.647	0.524	0.536
TLVQM [19]	0.772	0.774	0.589	0.616	0.732	0.724	0.670	0.691
VIDEVAL [45]	0.794	0.783	0.545	0.554	0.751	0.741	0.630	0.640
VSFA [22]	0.801	0.796	0.675	0.704	0.784	0.794	0.734	0.772
PVQ _{w/o/patch} [59]	0.814	0.816	0.686	0.708	0.781	0.781	0.747	0.776
PVQ _{w/patch} [59]	0.827	0.828	0.711	0.739	0.791	0.795	0.770	0.807
FAST-VQA-M [49]	0.852	0.854	0.739	0.773	0.841	0.832	0.788	0.810
FAST-VQA [49]	<u>0.876</u>	<u>0.877</u>	0.779	<u>0.814</u>	<u>0.859</u>	<u>0.855</u>	0.823	0.844
FAST-VQA* (our reproduced)	0.8686	0.8649	0.7707	0.7899	0.8418	0.8390	0.8124	0.8310
Zoom-VQA _{w/IQA}	0.8583	0.8557	0.7469	0.7907	0.8436	0.8373	0.7675	0.7970
Zoom-VQA _{w/VQA}	0.8735	0.8684	<u>0.7800</u>	0.7965	0.8448	0.8452	0.8137	0.8319
Zoom-VQA	0.8860	0.8791	0.7985	0.8189	0.8772	0.8745	<u>0.8141</u>	<u>0.8327</u>

4.3. Comparison with SOTA Results

In Tab. 1, we compare with current SOTA methods, including classical and deep learning-based VQA methods. And 4 testing sets are used for evaluation models trained on LSVQ. Among them, LSVQ_{test} and LSVQ_{1080p} are official intra-dataset test subsets. While KoNViD-1k and LIVE-VQC are used for the evaluation of generalization ability through cross-dataset testing. Compared with classical methods that rely on statistical regularities (e.g., BRISQUE), Zoom-VQA outperforms by large margins. Compared with some deep learning-based methods that apply well-designed networks (e.g., TLVQM), Zoom-VQA still obtains higher performances. Compared with PVQ which also considers local and global perceptual qualities using 2D/3D features, Zoom-VQA obtains higher results, improving the SRCC from 0.827 to 0.886 (+0.059) in LSVQ_{test}. Furthermore, compared with the current SOTA FAST-VQA, **Zoom-VQA improves the best results from 0.876 to 0.886 of SRCC in LSVQ_{test}, 0.779 to 0.798 in LSVQ_{1080p}**. Even without fine-tuning, Zoom-VQA obtains 0.877 of SRCC in KoNViD-1k, which is higher than the current best results of FAST-VQA (0.859). These prove the generalization ability of Zoom-VQA. Furthermore, we also report the results using separate IQA/VQA branches. The clip ensemble further boosts the final performance.

4.4. Ablation Studies

To further analyze the effectiveness of each component, we conduct ablation studies on the custom-divided train-test splits of VDPVE, which are illustrated in Sec. 4.1.

Table 2. Ablation study on the clip ensemble in VDPVE. The proposed ensemble strategy further boosts final performance.

IQA	VQA	SRCC	PLCC	Main Score
✓	✓	0.7872	0.7897	0.7885
✓	×	0.7727	0.7867	0.7797
×	✓	0.7389	0.7449	0.7419

Table 3. Ablation study on the usage of patch attention module and feature pyramid alignment using the IQA branch in VDPVE.

PAM	FPA	SRCC	PLCC	Main Score
✓	✓	0.7727	0.7867	0.7797
✓	×	0.7576	0.7606	0.7591
×	✓	0.7424	0.7438	0.7431
×	×	0.7185	0.7438	0.7311

Ablation on the clip ensemble. Both the results in Tab. 1 and Tab. 2 prove the effectiveness of the ensemble of the proposed IQA and VQA branches. By the proposed dual-branch design, Zoom-VQA can effectively capture both local and global information that impacts video quality.

Ablation on the patch attention module and the feature pyramid alignment. In Tab. 3, removing the patch head attention module (PAM) leads to a drop of 0.0366 (from 0.7797 to 0.7431) in the main score. It proves the importance of considering qualities in different spatial regions. While removing the feature pyramid alignment (FPA) module decreases the main score from 0.7797 to 0.7591, which verifies the utility of using multi-scale information for capturing distortions in different feature levels.

Table 4. Ablation study on the partition head expansion in the VQA branch using different expansion types or window sizes.

Expansion	Size	SRCC	PLCC	Main Score
zero padding	6	0.7389	0.7449	0.7419
zero padding	8	0.7274	0.7352	0.7313
<i>w/o</i>	4	0.7244	0.7324	0.7284
reflect	6	0.6838	0.6709	0.6773
replicate	6	0.6973	0.7067	0.7020

Table 5. Comparisons on the NTIRE23 VQA challenge. The results are given based on the testing phase. Zoom-VQA achieved 2nd place according to the main score.

Team Name	Main Score	Ranking
TB-VQA	0.8576	1
QuoVadis (ours)	0.8396	2
OPDAI	0.8289	3
TIA Team	0.8199	4
VCCIP	0.7994	5
IVL	0.7859	6
HXHHXH	0.7850	7
fmgvtv	0.7727	8
KKARC	0.7635	9
DTVQA	0.7325	10
sqiyx	0.7302	11
402Lab	0.7136	12
one-for-all	0.6990	13
NTU-SLab	0.6972	14
HNU-LIMMC	0.6923	15
Drealitym	0.6863	16
LION-Vaader	0.6596	17
Caption Timor	0.6499	18
IVLab	0.5851	19

Ablation on the partition head expansion. In Tab. 4, we test different types of expansion for the initialization of kernels in the linear embedding. And the used zero padding achieves the best performance. We claim that zero padding can maintain consistency in the distribution of features during the initialization stage, which is beneficial for subsequent network training. And further enlarging the window size does not introduce continuous improvement.

4.5. NTIRE23 VQA Challenge Report

The final ranking of the competition in the testing phase is shown in Tab. 5. Our team QuoVadis, adopting Zoom-VQA, has achieved *2nd place* based on the main score. This demonstrates the good generalization ability of Zoom-VQA when facing unseen domains.

4.6. Visualization

Following the visualization approach in [49], we generate the spatial-temporal local quality map of our well-trained VQA branch. In this way, we can obtain an intuitive understanding of what is learned by Zoom-VQA.

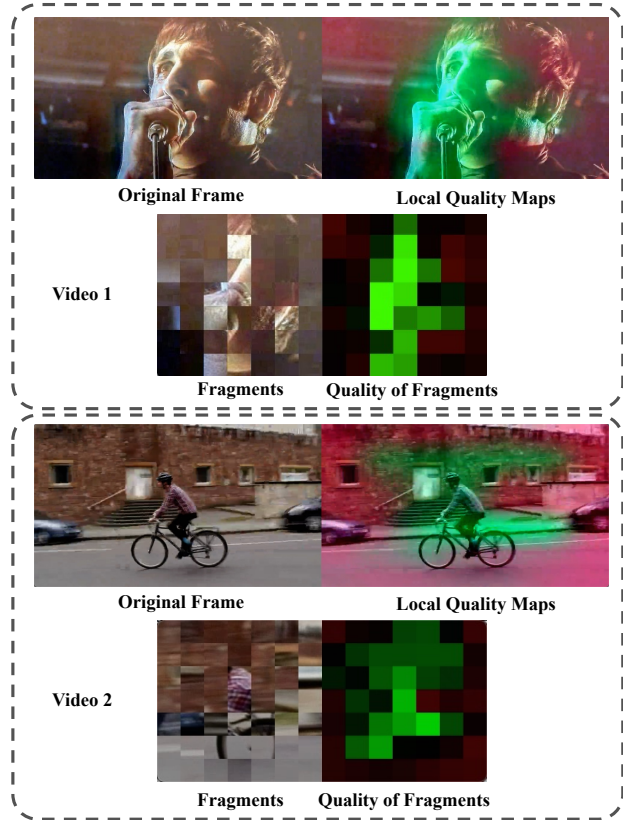


Figure 6. Spatio-temporal local quality maps, where red areas refer to relatively low quality scores and green areas refer to high scores. These demo videos are selected from VDPVE in 720P.

As shown in Fig. 6, we sample three typical videos from VDPVE in 720p and plot the original frame, local quality map, fragments, and quality map of fragments column by column. From video 1, we can find our VQA model tends to concentrate on region-of-interest (*i.e.*, face regions of the singer), which is consistent with HVS. Moreover, our VQA is assumed to be sensitive to textual quality information, since the qualities of clear areas are remarkably different from motion blur areas in video 2.

5. Conclusion

In this paper, we proposed a novel Zoom-VQA framework that assesses video quality by comprehensively analyzing the “patch-frame-clip” dimension of the video. Structurally, Zoom-VQA consists of two branches that receive image/clip input separately, considering both spatial and temporal information. Through designs of patch attention module, feature pyramid alignment, and clip ensemble, distortions in spatio-temporal and different feature levels can be effectively captured. Experimental results in VDPVE, LSVQ, KoNViD-1k, and LIVE-VQC prove the generalization ability of Zoom-VQA. Ablation studies further verify the effectiveness of each component.

References

- [1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lucic, and Cordelia Schmid. Vivit: A video vision transformer. In *ICCV*. IEEE, 2021. 2
- [2] Sebastian Bosse, Dominique Maniry, Klaus-Robert Müller, Thomas Wiegand, and Wojciech Samek. Deep neural networks for no-reference and full-reference image quality assessment. *IEEE TIP*, 27(1):206–219, 2018. 1
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229. Springer, 2020. 2
- [4] João Carreira and Andrew Zisserman. Quo vadis, action recognition? A new model and the kinetics dataset. In *CVPR*, pages 4724–4733. IEEE Computer Society, 2017. 2
- [5] Manri Cheon, Sung-Jun Yoon, Byungyeon Kang, and Junwoo Lee. Perceptual image quality assessment with transformers. In *CVPR Workshops*. IEEE, 2021. 1
- [6] Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *EMNLP*, pages 1724–1734. ACL, 2014. 2
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. IEEE, 2009. 2
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*. OpenReview.net, 2021. 2
- [9] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *ICCV*, pages 6804–6815. IEEE, 2021. 2
- [10] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. *CoRR*, abs/1812.03982, 2018. 2
- [11] Yixuan Gao, Yuqin Cao, Tengchuan Kou, Wei Sun, Yunlong Dong, Xiaohong Liu, Xiongkuo Min, and Guangtao Zhai. Vdpve: Vqa dataset for perceptual video enhancement. *CoRR*, abs/2303.09290, 2023. 1, 5
- [12] Deepti Ghadiyaram and Alan C. Bovik. Perceptual quality prediction on authentically distorted images using a bag of features approach. *CoRR*, abs/1609.04757, 2016. 1
- [13] Franz Götz-Hahn, Vlad Hosu, Hanhe Lin, and Dietmar Saupe. Konvid-150k: A dataset for no-reference video quality assessment of videos in-the-wild. *IEEE Access*, 9:72139–72160, 2021. 2
- [14] Haiyang Guo, Yi Bin, Yuqing Hou, Qing Zhang, and Hengliang Luo. IQMA network: Image quality multi-scale assessment network. In *CVPR Workshops*, pages 443–452. IEEE, 2021. 1
- [15] Jingwen He, Wu Shi, Kai Chen, Lean Fu, and Chao Dong. GCFSR: a generative and controllable face super resolution method without facial and GAN priors. In *CVPR*, pages 1879–1888. IEEE, 2022. 1
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778. IEEE, 2016. 2
- [17] 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. 2, 5
- [18] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR (Poster)*. OpenReview.net, 2017. 2
- [19] Jari Korhonen. Two-level approach for no-reference consumer video quality assessment. *IEEE Trans. Image Process.*, 28(12):5923–5938, 2019. 2, 6
- [20] Jari Korhonen, Yicheng Su, and Junyong You. Blind natural video quality prediction via statistical temporal features and deep spatial features. In *ACM MM*, 2020. 2
- [21] 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. Circ. Syst. Video Tech.*, 32(9):5944–5958, 2022. 1, 2
- [22] Dingquan Li, Tingting Jiang, and Ming Jiang. Quality assessment of in-the-wild videos. In *ACM Multimedia*, pages 2351–2359. ACM, 2019. 2, 6
- [23] Dingquan Li, Tingting Jiang, and Ming Jiang. Unified quality assessment of in-the-wild videos with mixed datasets training. *Int. J. Comput. Vis.*, 129(4):1238–1257, 2021. 2
- [24] Haoying Li, Yifan Yang, Meng Chang, Shiqi Chen, Huajun Feng, Zhihai Xu, Qi Li, and Yueting Chen. Srdiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing*, 479:47–59, 2022. 2
- [25] Jingyun Liang, Jie Zhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *ICCV Workshops*, pages 1833–1844. IEEE, 2021. 2
- [26] 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 *ICCV*, pages 9992–10002. IEEE, 2021. 2
- [27] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *CVPR*, pages 11966–11976. IEEE, 2022. 1, 3
- [28] 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, 2, 3
- [29] Kede Ma and Yuming Fang. Image quality assessment in the modern age. In *ACM MM*, pages 5664–5666, 2021. 1
- [30] Pavan C. Madhusudana, Neil Birkbeck, Yilin Wang, Balu Adsumilli, and Alan C. Bovik. Image quality assessment using contrastive learning. *IEEE Trans. Image Process.*, 31:4149–4161, 2022. 1
- [31] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE TIP*, 21(12):4695–4708, 2012. 2, 6
- [32] Anish Mittal, Michele A. Saad, and Alan C. Bovik. A completely blind video integrity oracle. *IEEE Trans. Image Process.*, 25(1):289–300, 2016. 2

- [33] Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. Making a "completely blind" image quality analyzer. *IEEE Signal Process. Lett.*, 20(3):209–212, 2013. 2
- [34] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, pages 8024–8035, 2019. 5
- [35] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatiotemporal representation with pseudo-3d residual networks. In *ICCV*, pages 5534–5542. IEEE, 2017. 2
- [36] Michele A. Saad, Alan C. Bovik, and Christophe Charrier. Blind prediction of natural video quality. *IEEE Trans. Image Process.*, 23(3):1352–1365, 2014. 2
- [37] Mike Schuster and Kuldip K. Paliwal. Bidirectional recurrent neural networks. *IEEE Trans. Signal Process.*, 45(11):2673–2681, 1997. 2
- [38] Zeina Sinno and Alan Conrad Bovik. Large-scale study of perceptual video quality. *IEEE Trans. Image Process.*, 28(2):612–627, 2019. 5
- [39] Zeina Sinno and Alan Conrad Bovik. Large-scale study of perceptual video quality. *IEEE Trans. Image Process.*, 28(2):612–627, 2019. 5
- [40] Wei Sun, Xiongkuo Min, Wei Lu, and Guangtao Zhai. A deep learning based no-reference quality assessment model for UGC videos. In *ACM MM*, pages 856–865, 2022. 1
- [41] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, pages 10347–10357. PMLR, 2021. 2
- [42] Du Tran, Lubomir D. Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, pages 4489–4497. IEEE Computer Society, 2015. 2
- [43] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, pages 6450–6459. IEEE, 2018. 2
- [44] Zhengzhong Tu, Chia-Ju Chen, Li-Heng Chen, Neil Birkbeck, Balu Adsumilli, and Alan C. Bovik. A comparative evaluation of temporal pooling methods for blind video quality assessment. In *ICIP*, pages 141–145. IEEE, 2020. 2
- [45] Zhengzhong Tu, Yilin Wang, Neil Birkbeck, Balu Adsumilli, and Alan C. Bovik. UGC-VQA: benchmarking blind video quality assessment for user generated content. *IEEE Trans. Image Process.*, 30:4449–4464, 2021. 1, 2, 6
- [46] Zhengzhong Tu, Xiangxu Yu, Yilin Wang, Neil Birkbeck, Balu Adsumilli, and Alan C. Bovik. RAPIQUE: rapid and accurate video quality prediction of user generated content. *CoRR*, abs/2101.10955, 2021. 2
- [47] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *ICCV Workshops*, pages 1905–1914. IEEE, 2021. 2
- [48] Yilin Wang, Sasi Inguva, and Balu Adsumilli. Youtube UGC dataset for video compression research. In *MMSP*, pages 1–5. IEEE, 2019. 2, 3
- [49] Haoning Wu, Chaofeng Chen, Jingwen Hou, Liang Liao, Annan Wang, Wenxiu Sun, Qiong Yan, and Weisi Lin. FAST-VQA: efficient end-to-end video quality assessment with fragment sampling. In *ECCV*, pages 538–554. Springer, 2022. 1, 2, 3, 5, 6, 7
- [50] Haoning Wu, Chaofeng Chen, Liang Liao, Jingwen Hou, Wenxiu Sun, Qiong Yan, and Weisi Lin. Discovqa: Temporal distortion-content transformers for video quality assessment. *CoRR*, abs/2206.09853, 2022. 2
- [51] 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. 2
- [52] Lei Xiao, Anton Kaplanyan, Alexander Fix, Matthew Chapman, and Douglas Lanman. Deepfocus: learned image synthesis for computational displays. *ACM Trans. Graph.*, 37(6):200, 2018. 1
- [53] Lei Xiao, Salah Nouri, Matthew Chapman, Alexander Fix, Douglas Lanman, and Anton Kaplanyan. Neural super-sampling for real-time rendering. *ACM Trans. Graph.*, 39(4):142, 2020. 1
- [54] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*, pages 318–335. Springer, 2018. 2
- [55] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simsim: A simple framework for masked image modeling. *CoRR*, abs/2111.09886, 2021. 2
- [56] Jiahua Xu, Jing Li, Xingguang Zhou, Wei Zhou, Baichao Wang, and Zhibo Chen. Perceptual quality assessment of internet videos. In *ACM MM*, pages 1248–1257, 2021. 2
- [57] Sidi Yang, Tianhe Wu, Shuwei Shi, Shanshan Lao, Yuan Gong, Mingdeng Cao, Jiahao Wang, and Yujiu Yang. MANIQA: multi-dimension attention network for no-reference image quality assessment. In *CVPR Workshops*, pages 1190–1199. IEEE, 2022. 5
- [58] Peng Ye, Jayant Kumar, Le Kang, and David S. Doermann. Unsupervised feature learning framework for no-reference image quality assessment. In *CVPR*, pages 1098–1105. IEEE Computer Society, 2012. 2
- [59] Zhenqiang Ying, Maniratnam Mandal, Deepti Ghadiyaram, and Alan C. Bovik. Patch-vq: 'patching up' the video quality problem. In *CVPR*, pages 14019–14029. Computer Vision Foundation / IEEE, 2021. 1, 2, 5, 6
- [60] Kun Yuan, Shaopeng Guo, Ziwei Liu, Aojun Zhou, Fengwei Yu, and Wei Wu. Incorporating convolution designs into visual transformers. In *ICCV*, pages 559–568. IEEE, 2021. 2
- [61] Kai Zhao, Kun Yuan, Ming Sun, Mading Li, and Xing Wen. Quality-aware pre-trained models for blind image quality assessment. In *CVPR*. IEEE, 2023. 1, 5