

Modular Blind Video Quality Assessment

Wen Wen¹, Mu Li², Yabin Zhang³, Yiting Liao³, Junlin Li³, Li Zhang³, and Kede Ma^{1*}

¹ City University of Hong Kong, ² The Chinese University of Hong Kong, Shenzhen

³ ByteDance Inc.

wwen29-c@my.cityu.edu.hk, limuhit@gmail.com

{zhangtao.ceb, liaoyiting, lijunlin.li, lizhang.idm}@bytedance.com

kede.ma@cityu.edu.hk

<https://github.com/winwinwenwen77/ModularBVQA>

Abstract

Blind video quality assessment (BVQA) plays a pivotal role in evaluating and improving the viewing experience of end-users across a wide range of video-based platforms and services. Contemporary deep learning-based models primarily analyze video content in its aggressively subsampled format, while being blind to the impact of the actual spatial resolution and frame rate on video quality. In this paper, we propose a modular BVQA model and a method of training it to improve its modularity. Our model comprises a base quality predictor, a spatial rectifier, and a temporal rectifier, responding to the visual content and distortion, spatial resolution, and frame rate changes on video quality, respectively. During training, spatial and temporal rectifiers are dropped out with some probabilities to render the base quality predictor a standalone BVQA model, which should work better with the rectifiers. Extensive experiments on both professionally-generated content and user-generated content video databases show that our quality model achieves superior or comparable performance to current methods. Additionally, the modularity of our model offers an opportunity to analyze existing video quality databases in terms of their spatial and temporal complexity.

1. Introduction

We undoubtedly reside in an era that exposes us to a diverse array of exponentially growing video data on a daily basis. Such growth is accompanied by advancements in video recording equipment and display devices, leading to high spatial resolution, high frame rate, high dynamic range, and wide color gamut video content. As a result, understanding how multifaceted video attributes together affect video quality is crucial to measure and improve the viewing experience of end-users.

Over the years, researchers have collected numerous supporting evidence from psychophysical and perceptual studies [11, 18, 22, 23, 27, 47, 49] that a higher spatial resolution and higher frame rate contribute positively to video quality, with exact perceptual gains depending on the video content.

In response to these subjective findings, early knowledge-driven BVQA models directly take spatial resolution and frame rate parameters as part of input for quality prediction of compressed videos [30]. Despite the simplicity, these video attribute parameters, being independent of content and distortion, are less relevant to perceived video quality. Consequently, simple content complexity features like spatial information¹ and temporal information², along with simple natural scene statistics derived from the pixel domain [26], the 3D discrete cosine transform domain [34], and the wavelet domain [3] are computed at the actual spatial resolution and frame rate [11]. Nevertheless, operating on the full-size video sequence is computationally daunting. Korhonen [13] proposed a two-level approach, extracting low-complexity features from each spatially downsampled frame and high-complexity features from key frames at the actual spatial resolution. An alternative approach is to work with spatiotemporal chips [6], which are spatial-time localized cuts of the full-size video in local motion flow directions. In general, knowledge-driven BVQA models perform marginally due to the limited expressiveness of manually crafted features.

The computational issues faced by deep learning-based data-driven BVQA methods are more pronounced. There are few attempts [16, 17] to assess full-size videos, which come with significant computational demands, especially

¹Spatial information is computed as the root mean squares of the gradient magnitudes of each frame.

²Temporal information is computed as the root mean squares of the differences of consecutive frames.

*Corresponding author.

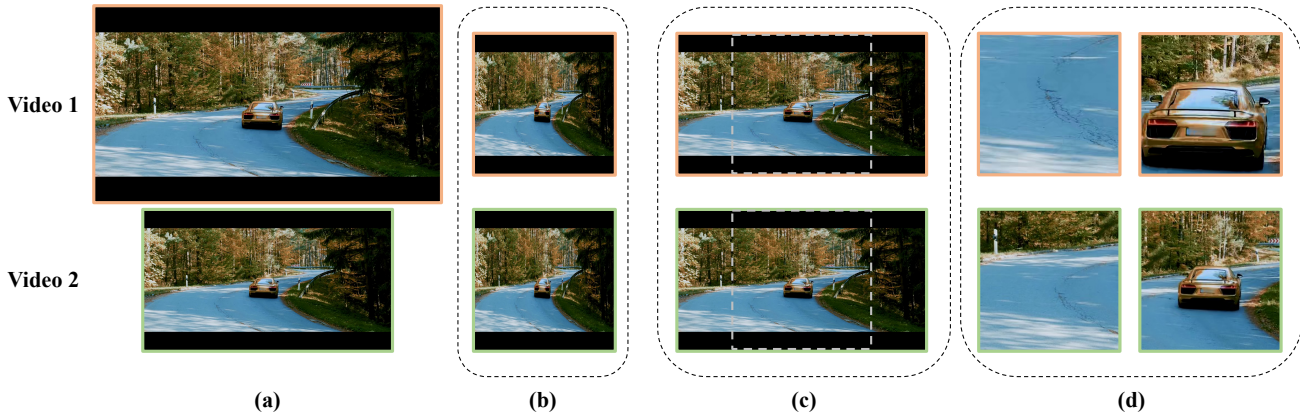


Figure 1. Conventional ways of pre-processing videos in spatial view. (a) Videos with the same content but different spatial resolutions taken from the Waterloo-IVC-4K dataset [18]. (b) Resizing without maintaining the aspect ratio leads to geometric distortions of structures and textures. (c) Aspect ratio-preserving resizing and cropping results in almost identical inputs of fixed size. (d) Cropping from videos at the actual spatial resolution reduces the field of view with limited content coverage.

for videos of high resolutions and frame rates. Moreover, due to the small scale of video quality datasets, many BVQA methods that utilize convolutional neural networks (CNNs) commonly depend on pretrained models from object recognition tasks, which expect small and fixed-size inputs. Consequently, videos need to be spatially resized [37, 43, 46] and/or cropped [43–45, 47], and temporally subsampled [37, 43, 46]. When two videos of the same scene but different spatial resolutions are resized to the same lower resolution, their quality variations, which are discernable by the human eye, diminish (see Figure 1 (b) and (c)). Cropping, on the other hand, significantly reduces the field of view and content coverage, which can impede the precision of BVQA. When videos that share the same content and length but differ in frame rates undergo temporal subsampling at a rate proportional to their frame rates, the frames that remain would be identical, as shown in Figure 2. Therefore, these BVQA methods are insensitive to changes in the spatial resolution and frame rate and their impact on video quality.

To reliably assess the perceptual quality of digital videos with great content and distortion diversities, and variable spatial resolutions and frame rates, we propose a modular BVQA model. Our model consists of three modules: a base quality predictor, a spatial rectifier, and a temporal rectifier, responding to the visual content and distortion, spatial resolution, and frame rate changes, respectively. The base quality predictor takes a sparse set of spatially downsampled key frames as input, and produces a scalar as the quality estimate. The spatial rectifier relies on a shallow CNN to process the Laplacian pyramids of the key frames at the actual spatial resolution, and computes the scaling and shift parameters to rectify the base quality score. Similarly, the temporal rectifier relies on a lightweight CNN to process the spatially downsampled video chunks centered around

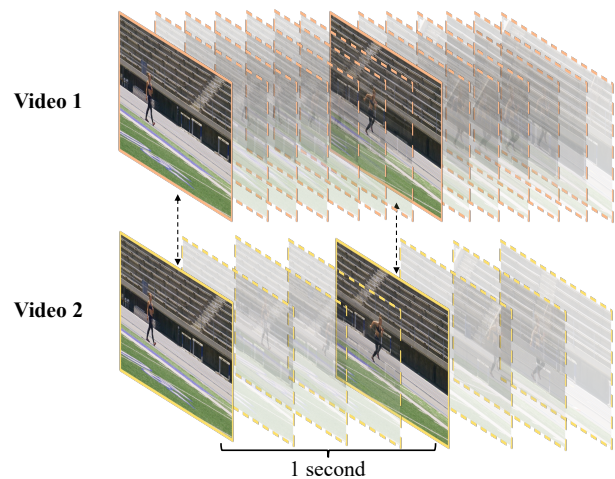


Figure 2. Two videos from the LIVE-YT-HFR dataset [23] with identical content and duration but different frame rates. When the subsampling rate is proportional to the frame rate, the remaining frames are identical.

the key frames at the actual frame rate, and computes another scaling and shift parameters for quality rectification. To enhance the modularity of our model, we introduce a dropout strategy during training. At each iteration, we randomly drop out the spatial and/or temporal rectifiers with pre-specified probabilities. In summary, the current work presents several novel elements:

- A modular BVQA model that is sensitive to visual distortions and changes in the spatial resolution and frame rate, and is readily extensible to add other video attributes (*e.g.*, dynamic range and color gamut) as additional rectifiers;
- A training strategy for our BVQA model to encourage the base quality predictor to function independently as a BVQA model, which would perform better when combined with the rectifiers;

- A experimental demonstration of the superiority of our BVQA model over existing methods on a total of *fourteen* video quality datasets [1, 8, 10, 15, 18, 21–23, 28, 35, 36, 42, 47, 48], covering professionally-generated content (PGC) and user-generated content (UGC), synthetic and authentic distortions, and variable spatial resolutions and frame rates;
- A computational analysis of eight UGC datasets in terms of their spatial and temporal complexity thanks to the modularity design of our BVQA model [38].

2. Related Work

In this section, we provide a concise overview of BVQA models for PGC and UGC content.

2.1. BVQA Models for PGC Content

Several publicly available PGC datasets [15, 18, 21–23, 33] have been specifically constructed to investigate the effects of spatial resolution and frame rate variations on video quality. However, there is a lack of research dedicated to BVQA methods for various spatial resolution and frame rate configurations. Ou *et al.* [29] designed an exponential function to compute video quality, taking into account the combined effect of a spatial quality factor (*e.g.*, the peak signal-to-noise ratio (PSNR)) and a temporal quality factor (*e.g.*, the frame rate). Janowski and Romaniak [11] proposed a generalized linear function to model spatial and temporal information independently. Ou *et al.* [30] developed Q-STAR, a model that measures the joint influence of the spatial resolution, temporal resolution, and quantization step size through the product of three one-parameter functions. FAVER [51] was designed to evaluate videos with diverse and high frame rates. It extracts natural scene statistics from the spatiotemporal wavelet domain, and employs a support vector regressor to estimate video quality at varying frame rates.

2.2. BVQA Models for UGC Content

UGC videos cover a wide range of natural scenes, spatial resolutions, and frame rates, leading to a diverse array of spatiotemporal distortions. Current BVQA models have achieved considerable progress in assessing the quality of UGC videos. In particular, knowledge-driven BVQA models primarily rely on manually crafted features from both spatial and temporal domains to determine video quality. Notable examples of such models are V-BLIINDS [34], VIIDEO [26], TLVQM [13], and VIDEVAL [39]. Recent studies have been focusing on enhancing BVQA models by integrating handcrafted features with those computed by pretrained CNNs. Representative models in this regard include CNN-TLVQM [14] and RAPIQUE [40]. Meanwhile, there has been a shift towards purely learning-based approaches. VSFA [17], Li22 [16], PVQ [47], and CoINVQ [43] leverage various pretrained networks as fixed

feature extractors, and train the regression module independently. Yi21 [46] and SimpleVQA [37] reduce the video resolution in space and time, followed by end-to-end fine-tuning. FastVQA [44] extracts spatially localized and temporally aligned cubes (termed as fragments) to facilitate end-to-end training. DOVER [45] further enhances FastVQA by integrating a pretrained ConvNeXt-T [19] designed for image aesthetics assessment.

3. Proposed Method

In this section, we present in detail our modular BVQA model, comprising a base quality predictor, a spatial rectifier, and a temporal rectifier, as illustrated in Figure 3.

3.1. Base Quality Predictor

We denote a video sequence as $\mathbf{x} = \{\mathbf{x}_i\}_{i=0}^{N-1}$, where $\mathbf{x}_i \in \mathbb{R}^{H \times W \times 3}$ represents the i -th frame, H and W are the frame height and width, and N is the total number of frames. The base quality predictor $f(\cdot) : \mathbb{R}^{H \times W \times 3 \times N} \mapsto \mathbb{R}$ takes \mathbf{x} as input, and computes a quality score, q_b . As part of the modular design, $f(\cdot)$ responds to semantic visual content and spatial distortions that remain unaffected by spatial resizing. Specifically, we uniformly sample a sparse set of M key frames, $\mathbf{y} = \{\mathbf{y}_i\}_{i=0}^{M-1}$, which are subject to bicubic resizing and cropping to $H_b \times W_b$. Other parametric temporal subsampling techniques [12] are also applicable with added computational complexity. H_b and W_b are determined by the input specifications of the pretrained network, which in our paper is a vision Transformer (ViT) [5] in the CLIP model [32]. We use the image representation corresponding to the `[class]` token, and add a multilayer perceptron (MLP) with two layers (*i.e.*, two fully connected layers with ReLU nonlinearity in between) as the quality regressor. We compute per-key-frame quality scores, and average them to obtain a base quality estimate of \mathbf{x} .

3.2. Spatial Rectifier

We incorporate a spatial rectifier, $\mathbf{f}_s(\cdot) : \mathbb{R}^{H \times W \times 3 \times N} \mapsto \mathbb{R}^{2 \times 1}$, that takes the same \mathbf{x} as input and computes a scale parameter α_s and a shift parameter β_s to rectify the base quality score:

$$q_s = \alpha_s q_b + \beta_s. \quad (1)$$

As part of the modular design, $\mathbf{f}_s(\cdot)$ responds to spatial distortions that arise from or are affected by spatial resizing. Similar to the base quality predictor, we work with the sparse set of M key frames, $\mathbf{y} = \{\mathbf{y}_i\}_{i=0}^{M-1}$. The difference lies in that we do not further perform spatial resizing, but build a Laplacian pyramid of $K + 1$ levels for each key frame at the actual spatial resolution. For the i -th key frame

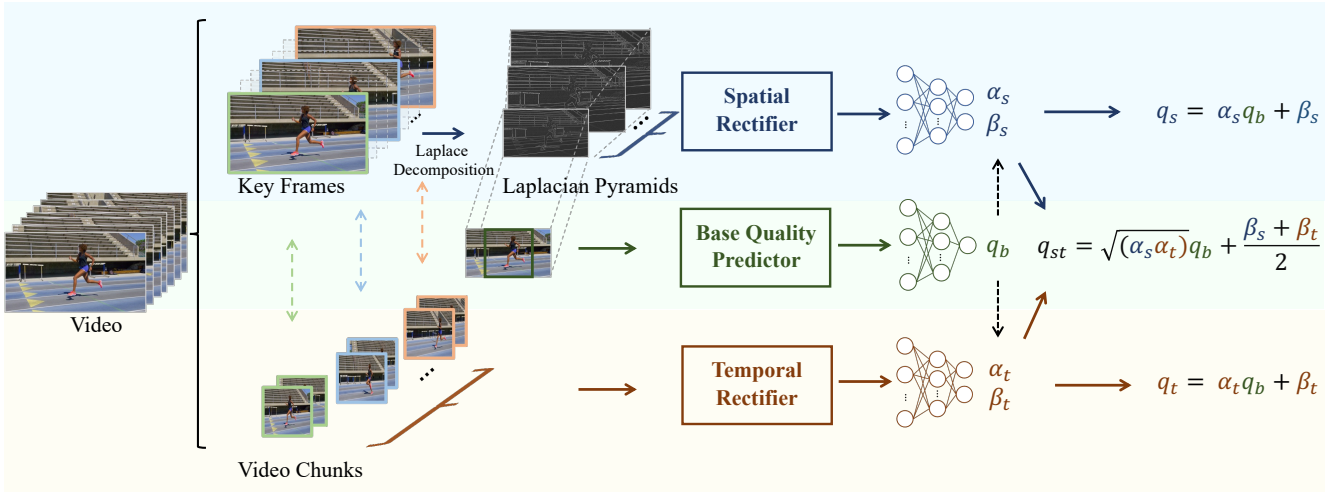


Figure 3. System diagram of our modular BVQA model. The base quality predictor takes a sparse set of spatially downsampled key frames as input, and generates a base quality score denoted by q_b . The spatial rectifier employs Laplacian pyramids derived from the key frames at their actual spatial resolution, and computes a scaling parameter α_s and a shift parameter β_s to rectify the base quality score. The temporal rectifier leverages features from the video chunks centered around the key frames at the actual frame rate to compute another scaling parameter α_t and shift parameter β_t for quality rectification.

and at the k -th level, we have

$$\mathbf{y}_i^{(k+1)} = \mathbf{D}\mathbf{L}\mathbf{y}_i^{(k)}, \quad k \in \{0, \dots, K-1\}, \quad (2)$$

$$\mathbf{z}_i^{(k)} = \mathbf{y}_i^{(k)} - \mathbf{L}\mathbf{U}\mathbf{y}_i^{(k+1)}, \quad (3)$$

$$\mathbf{z}_i^{(K)} = \mathbf{y}_i^{(K)}, \quad (4)$$

where $\mathbf{y}_i^{(0)} = \mathbf{y}_i$ and $\mathbf{z}_i^{(k)}$ is the k -th bandpass subband. \mathbf{D} and \mathbf{U} denote linear sub/upsampling operation, respectively, and \mathbf{L} denotes the lowpass filter in the matrix form. We set the subsampling factor ρ at each level to $\frac{\min\{H, W\}}{\min\{H_b, W_b\} \times K}$. In case ρ is a non-integer, we implement sub/upsampling together with the lowpass filtering using bicubic interpolation. The K -th level of the Laplacian pyramids $\mathbf{z}_i^{(K)}$ represents a low-frequency residual, which is related to the input of the base quality predictor by a cropping operator, and can thus be discarded.

We upsample all bandpass subbands to the actual resolution using bicubic interpolation, and feed each subband independently to a pretrained lightweight CNN for low-level feature extraction, followed by concatenation of $M \times K$ sets of spatial features. In our paper, we utilize the first two stages of ResNet-18 [9] pretrained on ImageNet-1K [4]. We last aggregate spatial information using global average and standard deviation pooling, and append a two-layer MLP to generate the spatial scaling parameter α_s and shift parameter β_s in Eq. (1).

3.3. Temporal Rectifier

We include a temporal rectifier $\mathbf{f}_t(\cdot) : \mathbb{R}^{H \times W \times 3 \times N} \rightarrow \mathbb{R}^{2 \times 1}$, that takes the same \mathbf{x} as input and computes another scaling parameter α_t and shift parameter β_t to rectify the base

quality score:

$$q_t = \alpha_t q_b + \beta_t. \quad (5)$$

As part of the modular design, $\mathbf{f}_t(\cdot)$ responds to temporal distortions resulting from motion anomaly and frame rate variations. We sample a set of M video chunks, denoted by $\mathbf{v} = \{\mathbf{v}_i\}_{i=0}^{M-1}$ from \mathbf{x} . Each \mathbf{v}_i is centered around the i -th key frame \mathbf{y}_i , and consists of L consecutive frames, resized to $H_v \times W_v$ without respecting the aspect ratio. We utilize the fast pathway of a pretrained SlowFast network [7], excluding the original classification head. Temporal features for different video chunks are concatenated, spatially pooled, and fed to a two-layer MLP to produce the temporal scaling parameter α_t and shift parameter β_t in Eq. (5).

3.4. Module Aggregation

During training, we implement a dropout strategy by randomly disabling the spatial and temporal rectifiers with some probabilities p_s and p_t , respectively. We derive the combined scaling and shift parameters as the geometric and arithmetic mean of those from the activated rectifiers:

$$\alpha_{st} = \left(\alpha_s^{\mathbb{I}[u_s \geq p_s]} \alpha_t^{\mathbb{I}[u_t \geq p_t]} \right)^{\frac{1}{\max\{1, \mathbb{I}[u_s \geq p_s] + \mathbb{I}[u_t \geq p_t]\}}}, \quad (6)$$

and

$$\beta_{st} = \frac{\mathbb{I}[u_s \geq p_s] \beta_s + \mathbb{I}[u_t \geq p_t] \beta_t}{\max\{1, \mathbb{I}[u_s \geq p_s] + \mathbb{I}[u_t \geq p_t]\}}, \quad (7)$$

where $\mathbb{I}[\cdot]$ represents the indicator function, u_s and u_t are two samples drawn from the uniform distribution $\mathcal{U}[0, 1]$. The final rectified quality score becomes

$$q_{st} = \alpha_{st} q_b + \beta_{st}. \quad (8)$$

Dataset	# Videos	# Scenes	Duration (Sec)	Spatial Resolution	Frame Rate	Distortion Type
BVI-SR [21]	240	24	5	4K	60	Subsampling
Waterloo-IVC-4K [18]	1,200	20	9-10	540p, 1080p, 4K	24, 25, 30	Compression
BVI-HFR [22]	88	22	10	1080p	15, 30, 60, 120	Frame Averaging
LIVE-YT-HFR [23]	480	16	6-10	1080p	24, 30, 60, 82, 98, 120	Compression
LIVE-Livestream [35]	315	45	7	1080p, 4K	25, 30	Compression, Streaming
ETRI [15]	437	15	5-7	540p, 720p, 1080p, 4K	30, 60, 120	Compression
CVD2014 [28]	234	5	10-25	480p, 720p	10-32	Authentic
LIVE-Qualcomm [8]	208	54	15	1080p	30	Authentic
KoNViD-1K [10]	1,200	1,200	8	540p	24, 25, 30	Authentic
LIVE-VQC [36]	585	585	10	240p-1080p	30	Authentic
YouTube-UGC [42]	1,142	1,142	20	360p-4K	30	Authentic
LBVD [1]	1,013	1,013	10	240p-540p	< 30	Authentic, Streaming
LSVQ [47]	38,811	38,811	5-12	99p-4K	< 60	Authentic
LIVE-YT-Gaming [48]	600	600	8-9	360p-1080p	30, 60	Authentic, Rendering

Table 1. Summary of PGC and UGC VQA datasets.

As specific instances, if both rectifiers are dropped, $q_{st} = q_b$, and if both rectifiers are enabled:

$$q_{st} = \sqrt{(\alpha_s \alpha_t)} q_b + \frac{\beta_s + \beta_t}{2}. \quad (9)$$

Such a modular combination takes the contributions of the rectifiers into separate account, and is readily extensible to incorporate future rectifiers into the model design.

4. Experiment

In this section, we first describe the experimental setups. We then present the experimental results on the six PGC and eight UGC datasets, with emphasis on the computational analysis of the spatial and temporal complexity of UGC datasets. Last, we conduct a set of ablation experiments to complete our model analysis.

4.1. Experimental Setups

Benchmark Datasets. To assess the spatial rectifier, we employ BVI-SR [21] and Waterloo-IVC-4K [18], which are dedicated to examining how spatial resolutions affect video quality. To evaluate the temporal rectifier, we utilize BVI-HFR [22] and LIVE-YT-HFR [23], which are tailored to study the influence of frame rates on video quality. LIVE-Livestream and ETRI are selected to encompass video streaming-related artifacts. We further validate the proposed BVQA model on eight UGC datasets [1, 8, 10, 28, 36, 42, 47, 48]. A comprehensive summary of these datasets is provided in Table 1.

Competing Methods. For PGC datasets, we choose six knowledge-driven BVQA models for comparison: 1) NIQE [25], 2) BRISQUE [24], 3) VIDEVAL [39], 4) TLVQM [13], 5) RAPIQUE [40] and 6) FAVER [51], which extract spatiotemporal features from full-size videos. FAVER [51] is the only model capable of responding to different frame rates. We also include VSFA [17], a

learning-based model that takes full-size videos as input, and BTURA [20], which takes cropped frame patches as input. For UGC datasets, we select 1) VSFA, 2) Li22 [16], 3) SimpleVQA [37], 4) FastVQA [44], and 5) DOVER [45] as competing models, among which FastVQA and DOVER deliver the state-of-the-art performance.

Performance Criteria. We employ two evaluation criteria: the Spearman’s rank correlation coefficient (SRCC) as a measure of prediction monotonicity and the Pearson linear correlation coefficient (PLCC) as a measure of prediction linearity. For the six PGC datasets, we adhere to the dataset splitting strategy described in the original papers, ensuring the content independence between the training and test sets. For the eight UGC datasets expect for LSVQ [47], we randomly split each into three non-overlapping subsets: 60% for training, 20% for validation, and 20% for testing. All learning-based models are retrained in accordance with their original optimization procedures. For LSVQ, we conform to the training and testing protocol in the original paper. We repeat all random splitting processes 10 times, and report the median results.

Implementation Details. For the BVI-SR, Waterloo-IVC-4K, BVI-HFR, LIVE-YT-HFR, LIVE-Livestream, and ETRI datasets, we set the number of key frames M to 5, 9, 10, 6, 7, and 5, respectively. For the eight UGC datasets, $M = 8$. All key frames are first resized such that the shorter side has 224 pixels and then cropped to 224×224 . For the spatial rectifier, the number of Laplacian pyramid levels K is set to 5, spanning from the actual spatial resolution down to 224. Similarly, for the temporal rectifier, we extract an equal number of video chunks to that of the key frames. During training, the spatial and temporal rectifiers are dropped out with a probability $p_s = p_t = 0.2$. The optimization of all learnable parameters is carried out by the Adam method for 30 epochs, with an initial learning rate of 1×10^{-5} , a decay ratio of 0.9 per two epochs, and a

Method	BVI-SR	Waterloo-IVC-4K
NIQE [25]	0.494 / 0.590	0.099 / 0.160
BRISQUE [24]	0.167 / 0.222	0.154 / 0.248
VSFA [17]	0.831 / 0.780	0.122 / 0.148
BTURA [20]	0.737 / 0.851	0.562 / 0.625
q_b (Ours)	0.448 / 0.554	0.785 / 0.834
q_s (Ours)	0.837 / 0.891	0.805 / 0.844
q_t (Ours)	0.304 / 0.411	0.770 / 0.832
q_{st} (Ours)	0.787 / 0.831	0.807 / 0.843

Table 2. Performance comparison of our models against competing methods on BVI-SR and Waterloo-IVC-4K, with emphasis on spatial resolution changes. All numbers are presented in the SRCC / PLCC format. The top-2 results on each dataset are highlighted in **bold**.

Method	BVI-HFR	LIVE-YT-HFR
NIQE [25]	0.225 / 0.419	0.137 / 0.418
BRISQUE [24]	0.260 / 0.445	0.319 / 0.420
VIDEVAL [39]	0.345 / 0.474	0.475 / 0.567
TLVQM [13]	0.373 / 0.491	0.430 / 0.505
RAPIQUE [40]	0.304 / 0.463	0.457 / 0.567
FAVER [51]	0.556 / 0.639	0.635 / 0.692
q_b (Ours)	0.035 / 0.206	0.588 / 0.696
q_s (Ours)	0.194 / 0.460	0.584 / 0.690
q_t (Ours)	0.620 / 0.782	0.789 / 0.799
q_{st} (Ours)	0.541 / 0.685	0.798 / 0.828

Table 3. Performance comparison of our models against competing methods on BVI-HFR and LIVE-YT-HFR, with emphasis on frame rate changes.

minibatch size of 16. We employ PLCC as the optimization goal as it combines the advantages of both the regression and learning-to-rank formulations of BVQA.

4.2. Results on PGC Datasets

The proposed BVQA model is modular, and is thus capable of outputting four quality scores: q_b , q_s , q_t , and q_{st} . Table 2 displays the results on BVI-SR and Waterloo-IVC-4K, which are predominantly influenced by the distortions arising from spatial resolution changes. Table 3 shows the results on BVI-HFR and LIVE-YT-HFR, emphasizing the distortions arising from frame rate conversions. Table 4 includes the results on LIVE-Livestream and ETRI, which encompass video streaming-based artifacts. By examining these results, we arrive at several interesting observations.

Spatial Rectifier Matters in Responding to Spatial Resolution Changes. This is evident in Table 2, especially on the BVI-SR dataset [21]. By incorporating the spatial rectifier, our model shows substantial improvements in performance compared to all other methods. Furthermore, it is worth noting that the inclusion of the temporal rectifier does not contribute significantly and can even harm the accuracy

Method	LIVE-Livestream	ETRI
NIQE [25]	0.323 / 0.496	0.346 / 0.342
BRISQUE [24]	0.638 / 0.670	0.248 / 0.207
TLVQM [13]	0.750 / 0.751	0.270 / 0.251
q_b (Ours)	0.641 / 0.678	0.924 / 0.943
q_s (Ours)	0.669 / 0.720	0.926 / 0.944
q_t (Ours)	0.730 / 0.728	0.929 / 0.947
q_{st} (Ours)	0.802 / 0.811	0.933 / 0.950

Table 4. Performance comparison of our models against competing methods on LIVE-Livestream and ETRI, with emphasis on video streaming-based artifacts.

of predictions concerning distortions related to spatial resolution changes. This underscores the value of a modular approach in the construction of BVQA models.

Nevertheless, we notice that the base quality predictor, despite working with low-spatial-resolution input, performs far better on Waterloo-IVC-4K [18] than BVI-SR. This performance discrepancy is largely attributed to the dominant compression artifacts present in Waterloo-IVC-4K, which are relatively unaffected by changes in spatial resolution. Conversely, BVI-SR is specifically designed to address distortions resulting from spatial subsampling.

Additionally, even though NIQE, BRISQUE, and VSFA extract global frame features at the actual spatial resolution, they show subpar performance on Waterloo-IVC-4K. Similarly, relying on cropped patches (*i.e.*, local frame features), BTURA also shows limited effectiveness on this dataset. These outcomes indicate that simple global or complex local feature extraction is inadequate when confronted with videos compressed at varying spatial resolutions.

Temporal Rectifier Matters in Responding to Frame Rate Changes. This is evident in Table 3, especially on BVI-HFR [22]. The inclusion of the temporal rectifier significantly enhances the performance of our model, surpassing all alternative methods by a clear margin. We find that the spatial rectifier tends to be effective primarily when combined with the temporal rectifier, as in the complete model, q_{st} . Similarly, the base quality predictor delivers noticeably better performance on LIVE-YT-HFR [23] compared to BVI-HFR. This difference can be ascribed to the fact that LIVE-YT-HFR contains a certain portion of spatial compression distortions, which the base quality predictor is adept at quantifying.

Out of the six competing models in Table 3, only FAVER [51] is specifically engineered to assess quality fluctuations due to frame rate changes, and it performs the best among them. Nevertheless, its subpar performance relative to our model suggests that the handcrafted features in FAVER have limited capability in addressing distortions related to frame rate variations.

Full Model Performs the Best in Complex Video Streaming Applications. LIVE-Livestream [35] is subject to both

Method	CVD2014	LIVE-Qualcomm	KoNViD-1K	LIVE-VQC	YouTube-UGC	LBVD	LIVE-YT-Gaming	Weighted Average
VSFA [17]	0.850 / 0.869	0.708 / 0.774	0.794 / 0.799	0.718 / 0.771	0.787 / 0.789	0.834 / 0.830	0.784 / 0.819	0.789 / 0.805
Li22 [16]	0.863 / 0.883	0.833 / 0.837	0.839 / 0.830	0.841 / 0.839	0.825 / 0.818	0.891 / 0.887	0.852 / 0.868	0.849 / 0.847
SimpleVQA [37]	0.834 / 0.864	0.722 / 0.774	0.792 / 0.798	0.740 / 0.775	0.819 / 0.817	0.872 / 0.878	0.814 / 0.836	0.810 / 0.823
FastVQA [44]	0.883 / 0.901	0.807 / 0.814	0.893 / 0.887	0.853 / 0.873	0.863 / 0.859	0.804 / 0.809	0.869 / 0.880	0.856 / 0.860
DOVER [45]	0.858 / 0.881	0.736 / 0.789	0.892 / 0.900	0.853 / 0.872	0.875 / 0.874	0.824 / 0.824	0.882 / 0.906	0.860 / 0.870
q_b (Ours)	0.870 / 0.892	0.759 / 0.775	0.864 / 0.875	0.737 / 0.786	0.841 / 0.847	0.701 / 0.700	0.859 / 0.895	0.806 / 0.822
q_s (Ours)	0.873 / 0.892	0.804 / 0.806	0.868 / 0.878	0.714 / 0.776	0.857 / 0.859	0.678 / 0.683	0.857 / 0.898	0.805 / 0.822
q_t (Ours)	0.879 / 0.899	0.825 / 0.822	0.892 / 0.891	0.833 / 0.851	0.854 / 0.858	0.887 / 0.885	0.857 / 0.894	0.868 / 0.875
q_{st} (Ours)	0.883 / 0.901	0.832 / 0.842	0.901 / 0.905	0.860 / 0.880	0.876 / 0.877	0.898 / 0.892	0.867 / 0.902	0.882 / 0.890

Table 5. Performance comparison of our models against five competing methods on seven small-scale UGC VQA datasets. The weighted average represents the average results across different datasets, weighted by the size of each dataset.

Method	Inference Time (Sec)	In-dataset Testing		Cross-dataset Testing			
		LSVQ-test	LSVQ-1080p	CVD2014	KoNViD-1K	LIVE-VQC	YouTube-UGC
VSFA [17]	11.109	0.801 / 0.796	0.675 / 0.704	0.756 / 0.760	0.810 / 0.811	0.753 / 0.795	0.718 / 0.721
Li22 [16]	27.632	0.852 / 0.854	0.772 / 0.788	0.817 / 0.811	0.843 / 0.835	0.793 / 0.811	0.802 / 0.792
SimpleVQA [37]	0.714	0.866 / 0.863	0.750 / 0.793	0.780 / 0.802	0.826 / 0.820	0.749 / 0.789	0.802 / 0.806
FastVQA [44]	0.045	0.876 / 0.877	0.779 / 0.814	0.805 / 0.814	0.859 / 0.855	0.823 / 0.844	0.730 / 0.747
DOVER [45]	0.047	0.888 / 0.889	0.795 / 0.830	0.829 / 0.832	0.884 / 0.883	0.832 / 0.855	0.777 / 0.792
q_b (Ours)	0.159	0.849 / 0.843	0.754 / 0.802	0.740 / 0.769	0.822 / 0.836	0.731 / 0.793	0.782 / 0.801
q_s (Ours)	0.159	0.838 / 0.842	0.764 / 0.808	0.775 / 0.796	0.845 / 0.856	0.773 / 0.823	0.723 / 0.743
q_t (Ours)	0.159	0.886 / 0.883	0.796 / 0.831	0.816 / 0.837	0.851 / 0.853	0.803 / 0.837	0.774 / 0.791
q_{st} (Ours)	0.159	0.895 / 0.895	0.809 / 0.844	0.823 / 0.839	0.878 / 0.884	0.806 / 0.844	0.788 / 0.804

Table 6. In-dataset and cross-dataset testing of our model against five competing models, all retrained on the official training split of the large-scale LSVQ dataset [47] and tested on other VQA datasets without fine-tuning. Inference time is computed using an NVIDIA A100 GPU, on an 8-second, 1080p, and 30 fps video.

streaming and compression distortions, whereas ETRI [15] is affected by spatiotemporal subsampling due to different levels of compression, primarily for streaming applications. The outcomes in Table 4 demonstrate that incorporating both spatial and temporal rectifiers is crucial for optimal performance on both datasets.

4.3. Results on UGC Datasets

Our method is pretrained on the official training split of the large-scale LSVQ dataset [47], followed by fine-tuning on seven small-scale UGC datasets [1, 8, 10, 28, 36, 42, 48]. All competing models are retrained following the procedures outlined in their respective papers. Results are shown in Table 5. Furthermore, a cross-dataset testing is performed, with results shown in Table 6.

Base Quality Predictor Works Well on CVD2014 [28], KoNViD-1K [10], and LIVE-YT-Gaming [48]. The base quality predictor achieves commendable performance on CVD2014. Akin to a blind image quality model, it outperforms most alternatives, except for FastVQA [44]. Incorporating the spatial and temporal rectifiers yields marginal performance gains. These results provide a strong indication that CVD2014 is primarily influenced by spatial distortions that are insensitive to spatial resizing with limited content diversity (*i.e.*, composition of merely five natural scenes). Similar observations can be made on KoNViD-1K

and LIVE-YT-Gaming, which are dominated by spatial distortions with some spatial complexity.

Spatial Rectifier Matters on YouTube-UGC [42] and LIVE-Qualcomm [8]. The spatial rectifier offers favorable performance improvements on YouTube-UGC and LIVE-Qualcomm. This can be explained by the presence of ultra-high-spatial-resolution videos, such as 4K in YouTube-UGC and 1080p in LIVE-Qualcomm. Similar conclusions can be drawn by comparing q_b with q_s on LSVQ-1080p.

Temporal Rectifier Matters on LIVE-VQC [36], LIVE-Qualcomm [8], and LBVD [1]. The temporal rectifier proves beneficial across multiple datasets, particularly on LIVE-VQC, LIVE-Qualcomm, and LBVD. This suggests the prominence of temporal distortions in these datasets. Interestingly, the spatially rectified model q_s is inferior to the base model q_b on LIVE-VQC and LBVD, providing additional evidence of the dominance of temporal distortions.

Full Model Performs the Best with Strong Cross-Dataset Generalization. The combination of both rectifiers consistently improves the quality prediction performance across all eight UGC datasets, yielding results that rival those of the leading methods. More importantly, the fully rectified model exhibits strong cross-dataset generalization, affirming the effectiveness of our modular BVQA method in addressing spatial and temporal aspects of video quality distinctly and thoroughly.

	LSVQ-test	LSVQ-1080p
RN50-ImageNet	0.858 / 0.857	0.741 / 0.782
RN50-CLIP	0.870 / 0.870	0.780 / 0.812
ViT-ImageNet	0.866 / 0.867	0.758 / 0.801
ViT-CLIP	0.895 / 0.895	0.809 / 0.844

(a) **Backbone initialization.** ViT-B outperforms ResNet-50, and the initialization using contrastive learning in CLIP is favored over the supervised learning on ImageNet-1K.

	LSVQ-test	LSVQ-1080p
ℓ_1	0.882 / 0.883	0.779 / 0.823
ℓ_2	0.887 / 0.886	0.797 / 0.828
PLCC	0.895 / 0.895	0.809 / 0.844

(d) **Loss function.** PLCC proves to be significantly better than the ℓ_1 -norm and ℓ_2 -norm induced metrics as loss functions.

	LSVQ-test	LSVQ-1080p
NN	0.893 / 0.893	0.809 / 0.844
Bilinear	0.891 / 0.891	0.809 / 0.843
Bicubic	0.895 / 0.895	0.809 / 0.844

(b) **Spatial resizing.** NN stands for the nearest neighbor interpolation. Bicubic interpolation marginally outperforms bilinear and nearest neighbor interpolation.

	BVI-SR	Waterloo-IVC-4K
RN18	0.859 / 0.904	0.850 / 0.870
RN34	0.866 / 0.907	0.810 / 0.855
RN50	0.861 / 0.911	0.809 / 0.846

(e) **Spatial rectifier.** As a shallower architecture with fewer parameters, ResNet-18 is sufficient to achieve satisfactory results.

	LSVQ-test	LSVQ-1080p
MLP	0.895 / 0.895	0.809 / 0.844
GRU	0.894 / 0.895	0.808 / 0.844
Transformer	0.892 / 0.891	0.802 / 0.841

(c) **Regressor.** GRU indicates the gated recurrent unit [2]. Even in the absence of complex temporal modeling, MLP exhibits a slight advantage over GRU and Transformer models.

	BVI-HFR	LIVE-YT-HFR
S	0.423 / 0.478	0.515 / 0.629
T	0.629 / 0.733	0.791 / 0.801
S+T	0.529 / 0.659	0.765 / 0.789

(f) **Temporal rectifier.** S and T stand for the slow and fast pathways in SlowFast [7]. The fast pathway solely yields excellent performance with reduced complexity.

Table 7. Ablation experiments. Default settings are marked in gray .

Our computational analysis of the eight UGC datasets resonates with a recent computational investigation of VQA datasets via design of minimalistic VQA models [38]. Specifically, we also expose the *easy dataset problem* in the construction of current VQA datasets, which allows oversimplistic models, such as those close to blind image quality models, to achieve state-of-the-art performance. The root cause of this problem is the superficial treatment of sample selection and excessive reliance on absolute category rating as the de facto subjective testing methodology to gather single-valued quality indicators.

4.4. Ablation Studies

We expand our computational investigation by ablating 1) backbones of the based quality predictor with different initializations in Table 7a, 2) spatial resizing methods in Table 7b, 3) quality regressors in Table 7c, 4) loss functions in Table 7d, 5) spatial rectifiers in Table 7e, and 6) temporal rectifiers in Table 7f. For the former four studies, we ablate the full model q_{st} on LSVQ [47], for the spatial rectifier, we ablate q_s on BVI-SR [21] and Waterloo-IVC-4K [18], and for the temporal rectifier, we ablate q_t on BVI-HFR [22] and LIVE-YT-HFR [23].

The ablation studies show that our base quality predictor benefits from a stronger ViT-B [5] backbone, pretrained through contrastive learning. The impact of employing different spatial resizing techniques for preparing inputs to the base predictor and the two rectifiers appears to be negligible, so do the quality regressors. Even more sophisticated models like the gated recurrent unit (GRU) [2] and the Transformer model [41] do not show expected performance improvements over the MLP, when considering the increased computational complexity. In contrast, PLCC as the loss function yields noticeable performance gains compared to the ℓ_1 -norm and ℓ_2 -norm induced metrics. More-

over, the simple ResNet-18 as the spatial rectifier and the Fast pathway as the temporal rectifier provide a good trade-off between performance and computational complexity.

5. Conclusion and Discussion

We have described a modular BVQA model, accompanied by a training strategy that improves its modularity. Our method addresses the BVQA challenges posed by intricacies of natural scenes and spatiotemporal distortions, and variations in spatial resolutions and frame rates. Extensive experiments on fourteen VQA datasets verify the promise of the proposed method. Additionally, the modular design of our model gives us an opportunity to analyze current UGC datasets in terms of their spatial and temporal complexity.

The current work serves as an initial step towards modular BVQA, yet numerous challenges remain unaddressed. For example, our method assumes a fixed routing function [31], where activation of the rectifiers is determined by expert insight into the VQA datasets. A routing function that adapts to specific datasets or individual samples is compelling for future research. Meanwhile, the hand-crafted aggregation function in Eq. (9) might seem rudimentary. It is interesting to learn a neural network to aggregate active modules, potentially at the feature level. Finally, adding a new rectifier currently necessitates the retraining or fine-tuning of all parameters. Exploring parameter-efficient training methodologies, such as continual learning [50] and transfer learning, could offer substantial benefits.

6. Acknowledgement

This work was supported in part by the National Natural Science Foundation of China under Grants 62071407 and 62102339, and the Hong Kong RGC Early Career Scheme (2121382).

References

- [1] Pengfei Chen, Leida Li, Yipo Huang, Fengfeng Tan, and Wenjun Chen. QoE evaluation for live broadcasting video. In *IEEE International Conference on Image Processing*, pages 454–458, 2019. **3, 5, 7**
- [2] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Conference on Empirical Methods in Natural Language Processing*, pages 1724–1734, 2014. **8**
- [3] Sathya V. R. Dendi and Sumohana S. Channappayya. No-reference video quality assessment using natural spatiotemporal scene statistics. *IEEE Transactions on Image Processing*, 29:5612–5624, 2020. **1**
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. **4**
- [5] 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 *International Conference on Learning Representations*, pages 1–22, 2021. **3, 8**
- [6] Joshua P. Ebenezer, Zaixi Shang, Yongjun Wu, Hai Wei, Sri-ran Sethuraman, and Alan C. Bovik. ChipQA: No-reference video quality prediction via space-time chips. *IEEE Transactions on Image Processing*, 30:8059–8074, 2021. **1**
- [7] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. SlowFast networks for video recognition. In *IEEE International Conference on Computer Vision*, pages 6202–6211, 2019. **4, 8**
- [8] Deepti Ghadiyaram, Janice Pan, Alan C. Bovik, Anush K. Moorthy, Prasanjit Panda, and Kai-Chieh Yang. In-capture mobile video distortions: A study of subjective behavior and objective algorithms. *IEEE Transactions on Circuits and Systems for Video Technology*, 28(9):2061–2077, 2018. **3, 5, 7**
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. **4**
- [10] 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 *IEEE International Conference on Quality of Multimedia Experience*, pages 1–6, 2017. **3, 5, 7**
- [11] Lucjan Janowski and Piotr Romaniak. QoE as a function of frame rate and resolution changes. In *International Workshop on Future Multimedia Networking*, pages 34–45. Springer, 2010. **1, 3**
- [12] Bruno Korbar, Du Tran, and Lorenzo Torresani. SCSampler: Sampling salient clips from video for efficient action recognition. In *IEEE International Conference on Computer Vision*, pages 6232–6242, 2019. **3**
- [13] Jari Korhonen. Two-level approach for no-reference consumer video quality assessment. *IEEE Transactions on Image Processing*, 28(12):5923–5938, 2019. **1, 3, 5, 6**
- [14] Jari Korhonen, Yicheng Su, and Junyong You. Blind natural video quality prediction via statistical temporal features and deep spatial features. In *ACM International Conference on Multimedia*, pages 3311–3319, 2020. **3**
- [15] Dae Yeol Lee, Somdyuti Paul, Christos G. Bampis, Hyun-suk Ko, Jongho Kim, Se Yoon Jeong, Blake Homan, and Alan C. Bovik. A subjective and objective study of space-time subsampled video quality. *IEEE Transactions on Image Processing*, 31:934–948, 2021. **3, 5, 7**
- [16] 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 Transactions on Circuits and Systems for Video Technology*, 32(9):5944–5958, 2022. **1, 3, 5, 7**
- [17] Dingquan Li, Tingting Jiang, and Ming Jiang. Quality assessment of in-the-wild videos. In *ACM International Conference on Multimedia*, pages 2351–2359, 2019. **1, 3, 5, 6, 7**
- [18] Zhuoran Li, Zhengfang Duanmu, Wentao Liu, and Zhou Wang. AVC, HEVC, VP9, AVS2 or AV1?—A comparative study of state-of-the-art video encoders on 4K videos. In *International Conference on Image Analysis and Recognition*, pages 162–173, 2019. **1, 2, 3, 5, 6, 8**
- [19] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A ConvNet for the 2020s. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 11976–11986, 2022. **3**
- [20] Wei Lu, Wei Sun, Xiongkuo Min, Wenhan Zhu, Quan Zhou, Jun He, Qiyuan Wang, Zicheng Zhang, Tao Wang, and Guangtao Zhai. Deep neural network for blind visual quality assessment of 4K content. *IEEE Transactions on Broadcasting*, 69(2):406–421, 2023. **5, 6**
- [21] Alex Mackin, Mariana Afonso, Fan Zhang, and David R. Bull. A study of subjective video quality at various spatial resolutions. In *IEEE International Conference on Image Processing*, pages 2830–2834, 2018. **3, 5, 6, 8**
- [22] Alex Mackin, Fan Zhang, and David R. Bull. A study of high frame rate video formats. *IEEE Transactions on Multimedia*, 21(6):1499–1512, 2018. **1, 5, 6, 8**
- [23] Pavan C. Madhusudana, Xiangxu Yu, Neil Birkbeck, Yilin Wang, Balu Adsumilli, and Alan C. Bovik. Subjective and objective quality assessment of high frame rate videos. *IEEE Access*, 9:108069–108082, 2021. **1, 2, 3, 5, 6, 8**
- [24] Anish Mittal, Anush K. Moorthy, and Alan C. Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on Image Processing*, 21(12):4695–4708, 2012. **5, 6**
- [25] Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal Processing Letters*, 20(3):209–212, 2013. **5, 6**
- [26] Anish Mittal, Michele A. Saad, and Alan C. Bovik. A completely blind video integrity oracle. *IEEE Transactions on Image Processing*, 25(1):289–300, 2016. **1, 3**
- [27] Rasoul M. Nasiri, Jiheng Wang, Abdul Rehman, Shiqi Wang, and Zhou Wang. Perceptual quality assessment of high frame

- rate video. In *IEEE International Workshop on Multimedia Signal Processing*, pages 1–6, 2015. [1](#)
- [28] Mikko Nuutinen, Toni Virtanen, Mikko Vaahteranoksa, Tero Vuori, Pirkko Oittinen, and Jukka Häkkinen. CVD2014—A database for evaluating no-reference video quality assessment algorithms. *IEEE Transactions on Image Processing*, 25(7):3073–3086, 2016. [3](#), [5](#), [7](#)
- [29] Yen-Fu Ou, Zhan Ma, and Yao Wang. Modeling the impact of frame rate and quantization stepsizes and their temporal variations on perceptual video quality: A review of recent works. In *Annual Conference on Information Sciences and Systems*, pages 1–6, 2010. [3](#)
- [30] Yen-Fu Ou, Yuanyi Xue, and Yao Wang. Q-STAR: A perceptual video quality model considering impact of spatial, temporal, and amplitude resolutions. *IEEE Transactions on Image Processing*, 23(6):2473–2486, 2014. [1](#), [3](#)
- [31] Jonas Pfeiffer, Sebastian Ruder, Ivan Vulić, and Edoardo Ponti. Modular deep learning. *Transactions on Machine Learning Research*, pages 1–76, 2023. [8](#)
- [32] 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 *International Conference on Machine Learning*, pages 8748–8763, 2021. [3](#)
- [33] Rakesh R. R. Rao, Steve Göring, Werner Robitzka, Bernhard Feiten, and Alexander Raake. AVT-VQDB-UHD-1: A large scale video quality database for UHD-1. In *IEEE International Symposium on Multimedia*, pages 17–177, 2019. [3](#)
- [34] Michele A. Saad, Alan C. Bovik, and Christophe Charrier. Blind prediction of natural video quality. *IEEE Transactions on Image Processing*, 23(3):1352–1365, 2014. [1](#), [3](#)
- [35] Zaixi Shang, Joshua P. Ebenezer, Alan C. Bovik, Yongjun Wu, Hai Wei, and Sriram Sethuraman. Assessment of subjective and objective quality of live streaming sports videos. In *IEEE Picture Coding Symposium*, pages 1–5, 2021. [3](#), [5](#), [6](#)
- [36] Zeina Sinno and Alan C. Bovik. Large-scale study of perceptual video quality. *IEEE Transactions on Image Processing*, 28(2):612–627, 2018. [3](#), [5](#), [7](#)
- [37] Wei Sun, Xiongkuo Min, Wei Lu, and Guangtao Zhai. A deep learning based no-reference quality assessment model for UGC videos. In *ACM International Conference on Multimedia*, pages 856–865, 2022. [2](#), [3](#), [5](#), [7](#)
- [38] 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. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, to appear, 2024. [3](#), [8](#)
- [39] Zhengzhong Tu, Yilin Wang, Neil Birkbeck, Balu Adsumilli, and Alan C. Bovik. UGC-VQA: Benchmarking blind video quality assessment for user generated content. *IEEE Transactions on Image Processing*, 30:4449–4464, 2021. [3](#), [5](#), [6](#)
- [40] 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. *IEEE Open Journal of Signal Processing*, 2:425–440, 2021. [3](#), [5](#), [6](#)
- [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 6000–6010, 2017. [8](#)
- [42] Yilin Wang, Sasi Inguva, and Balu Adsumilli. YouTube UGC dataset for video compression research. In *IEEE International Workshop on Multimedia Signal Processing*, pages 1–5, 2019. [3](#), [5](#), [7](#)
- [43] Yilin Wang, Junjie Ke, Hossein Talebi, Joong Gon Yim, Neil Birkbeck, Balu Adsumilli, Peyman Milanfar, and Feng Yang. Rich features for perceptual quality assessment of UGC videos. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 13435–13444, 2021. [2](#), [3](#)
- [44] 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 *European Conference on Computer Vision*, pages 538–554, 2022. [3](#), [5](#), [7](#)
- [45] 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 *IEEE International Conference on Computer Vision*, pages 20144–20154, 2023. [2](#), [3](#), [5](#), [7](#)
- [46] Fuwang Yi, Mianyi Chen, Wei Sun, Xiongkuo Min, Yuan Tian, and Guangtao Zhai. Attention based network for no-reference UGC video quality assessment. In *IEEE International Conference on Image Processing*, pages 1414–1418, 2021. [2](#), [3](#)
- [47] Zhenqiang Ying, Maniratnam Mandal, Deepti Ghadiyaram, and Alan C. Bovik. Patch-VQ: ‘Patching up’ the video quality problem. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 14019–14029, 2021. [1](#), [2](#), [3](#), [5](#), [7](#), [8](#)
- [48] Xiangxu Yu, Zhengzhong Tu, Zhenqiang Ying, Alan C. Bovik, Neil Birkbeck, Yilin Wang, and Balu Adsumilli. Subjective quality assessment of user-generated content gaming videos. In *IEEE Winter Conference on Applications of Computer Vision*, pages 74–83, 2022. [3](#), [5](#), [7](#)
- [49] Guangtao Zhai, Jianfei Cai, Weisi Lin, Xiaokang Yang, Wenjun Zhang, and Minoru Etoh. Cross-dimensional perceptual quality assessment for low bit-rate videos. *IEEE Transactions on Multimedia*, 10(7):1316–1324, 2008. [1](#)
- [50] Weixia Zhang, Dingquan Li, Chao Ma, Guangtao Zhai, Xiaokang Yang, and Kede Ma. Continual learning for blind image quality assessment. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):2864–2878, 2022. [8](#)
- [51] Qi Zheng, Zhengzhong Tu, Yibo Fan, Xiaoyang Zeng, and Alan C. Bovik. No-reference quality assessment of variable frame-rate videos using temporal bandpass statistics. In *IEEE International Conference on Acoustics, Speech & Signal Processing*, pages 1795–1799, 2022. [3](#), [5](#), [6](#)