

DCVQE: A Hierarchical Transformer for Video Quality Assessment

Zutong Li* and Lei Yang†

Weibo R&D Limited, USA
{zutongli0805, trilithy}@gmail.com

Abstract. The explosion of user-generated videos stimulates a great demand for no-reference video quality assessment (NR-VQA). Inspired by our observation on the actions of human annotation, we put forward a Divide and Conquer Video Quality Estimator (DCVQE) for NR-VQA. Starting from extracting the frame-level quality embeddings (QE), our proposal splits the whole sequence into a number of clips and applies Transformers to learn the clip-level QE and update the frame-level QE simultaneously; another Transformer is introduced to combine the clip-level QE to generate the video-level QE. We call this hierarchical combination of Transformers as a Divide and Conquer Transformer (DCTr) layer. An accurate video quality feature extraction can be achieved by repeating the process of this DCTr layer several times. Taking the order relationship among the annotated data into account, we also propose a novel correlation loss term for model training. Experiments on various datasets confirm the effectiveness and robustness of our DCVQE model.

1 Introduction

Recent years have witnessed a significant increase in user-generated content (UGC) on social media platforms like Youtube, Tiktok, and Weibo. Watching the UGC videos on computers or smartphones has even become part of our daily life. This trend stimulates a great demand for automatic video quality assessment (VQA), especially in popular video sharing/recommendation services.

UGC-VQA, also known as blind or No-Reference video quality assessment (NR-VQA), aims to evaluate in-the-wild videos without the corresponding pristine reference videos. Usually, UGC videos may suffer from complex distortions due to the diversity of capturing devices, uncertain shooting skills, compression, and poor editing process. Although many excellent algorithms have been proposed to evaluate video quality, it remains a challenging task to assess the quality of UGC videos accurately and consistently.

Besides the frame-level image information, temporal information is regarded as a critically important factor for video analysis tasks. Although many image quality assessment (IQA) models [35, 67, 33, 61, 22, 14, 62, 60, 64] can be applied

* Work done when Z. Li was at Weibo. Z. Li is currently with Microsoft.

† Corresponding author.



Fig. 1. Three consecutive frames from Video B304 in LIVE-VQA dataset. An overall high mean opinion score (MOS) of 91.73 was annotated.

to VQA base on a simple temporal pooling process [52], these models may not work very robustly because of the absence of proper time sequence aggregation. For instance, Fig. 1 shows three consecutive frames extracted from Video B304 in LIVE-VQA [48] dataset. As seen, the motion blur distortion appears on the actress’s hand area. These frame images are most likely to be recognized as middle, even low quality when applying a sophisticated IQA method to them individually. However, the quality of this video was labeled as high by human annotators, because a very smooth movement of the actress can be observed when playing the video stream. RNNs [1, 2, 23] and 3D-CNNs [20, 6, 13, 12, 25] are potential models to integrate spatial and temporal information for NR-VQA. Though these algorithms perform well on many datasets, the difficulty of parallelizing RNN models and the non-negligible computational cost of 3D-CNNs make them infeasible to be applied to many internet applications that require quick responses, just like the online video sharing/recommendation services.

On the other hand, research on NR-VQA is often limited by the lack of sufficient training data, due to the tedious work to label the mean opinion score (MOS) for each video. Among the publicly available datasets, MCL-JCV [56], VideoSet [57], UGC-VIDEO [27], CVD-2014 [38], LIVE-Qualcomm [16] are generated in lab environments, while KoNViD-1k [18], LIVE-VQC [48], YouTube-UGC [58] and LSVQ [63] are collected in-the-wild. As mentioned above, one could consider VQA as a temporal extension of IQA task, thus can apply some frame-level distortion and ranking processes [28] to augment the small video datasets for algorithm development. However, in-the-wild videos are usually hard to synthesize, since they may suffer from compound distortions which cannot be exactly parameterized as a combination of certain distortion cases. Recently, a large-scale LSVQ dataset [63], which contains 39,075 annotated videos with authentic distortions is released for public research. We know that not many studies have been conducted so far based on this new dataset.

Additionally, most previous works take L1 or L2 loss [64, 63, 70, 23, 54] as the optimization criterion for model training. Since these criteria to some extent ignore the order relationship of quality scores of the training samples, the trained model may be not stable to quantify the perceptual differences between the

videos with similar quality scores. For example, in our research we find that many existing NR-VQA models work well to identify both high and low quality videos, but struggle to distinguish the videos with middle quality scores. How to effectively quantify the difference between samples with similar perceptual scores therefore becomes the key to the success of a NR-VQA model.

To address the above problems, in this paper we put forward a new Divide and Conquer Video Quality Estimator (DCVQE) model for NR-VQA. We summarize our contributions as follows: (1) Inspired by our observation on the actions of human annotation, we propose a Divide and Conquer Transformer (DCTr) architecture to extract video quality features for NR-VQA. Our algorithm starts from extracting the frame-level quality representations. Regarded as a divide process, we split the input sequence into a number of clips and apply Transformers to learn the clip-level quality embeddings (QE) and update the frame-level QE simultaneously. Subsequently, a conquer process is conducted by using another Transformer to combine the clip-level QE to generate a video-level QE. After stacking several DCTr layers and topping with a linear regressor, our DCVQE model can be constructed to predict the quality value of the input video. (2) By taking the order relationship of the training samples into account, we propose a novel correlation loss to bring an additional order constraint of video quality to guide the training. Experiments indicate that the introduction of this correlation loss can consistently help to improve the performance of our DCVQE model. (3) We conduct plenty of experiments on different datasets and confirm that our DCVQE outperforms most other algorithms.

2 Related works

Traditional NR-VQA solutions: Many prior NR-VQA works are “distortion specific” because they are designed to identify different distortion types like blur [31], blockiness [39], or noise [49] in compressed videos or image frames. More recent and popular used models are deployed on natural video statistics (NVS) features, which are created by extending the highly regular parametric band-pass models of natural scene statistics (NSS) from IQA to VQA tasks. Among them, successful applications have explored in both frequency domain (BIQI [36], DIIVINE [37], BLINDS [44], BLINDS-II [45]) and spatial domain (NIQE [35], BRISQUE [33]). V-BLIINDS [46] combined spatio-temporal NSS with motion coherency models to estimate perceptual video quality. Inter-subband correlations, modeled by spatial domain statistical features in frame differences, were used to quantify the degree of distortion in VIIDEO [34]. 3D-DCT was applied on local space-time regions, to establish quality aware features in [26]. Based on hand-crafted features selection and combination, recent algorithms VIDEVAL [53] and TLVQM [21] demonstrated outstanding performance on many UGC datasets. The designation of these hand-crafted features is deliberate, though they are hard to be deployed in an end-to-end fashion for NR-VQA tasks.

Deep learning based NR-VQA solutions: Deep neural networks have shown their superior abilities in many computer vision tasks. With the availabil-

ity of perceptual image quality datasets [15, 19, 41, 64], many successful applications have been reported in the past decade [28, 64, 29, 50, 70]. Combining with a convolutional neural aggregation network, DeepVQA [59] utilized the advantages of CNN to learn spatio-temporal visual sensitivity maps for VQA. Based on a weakly supervised learning and resampling strategy, Zhang et al. [69] proposed a general purpose NR-VQA framework which inherited the knowledge learned from full-reference VQA and can effectively alleviate the curse of inadequate training data. VSFA [23] used a pretrained CNN to abstract frame features, and introduced the gated recurrent unit (GRU) to learn the temporal dependencies. Following VSFA, MDTVSFA [24] proposed a mixed datasets training method to further improve VQA performance. Although the above methods performed well on synthetic distortion datasets, they may be unstable to analyze UGC videos with complex and diverse distortions. PVQ [63] reported a leading performance on the large-scale dataset LSVQ. For a careful study on the local and global spatio-temporal quality, spatial patch, temporal patch and spatio-temporal patch were introduced in [63]. RAPIQUE [54], by leveraging a set of NSS features concatenated with learned CNN features, shown the top performance on several public datasets, including KoNViD-1k, YouTube-UGC, and their combination.

Transformer techniques in computer vision: Self-attention mechanism-based Transformer architecture shows its exceptional performance in natural language processing [55, 9]. Recently, many researchers introduced Transformer to solve computer vision problems. ViT [10] directly run attention among image patches with positional embeddings for image classification. Detection Transformer (DETR) [5] reached a comparable performance with Faster-RCNN [43] by designing a new object detection systems based on Transformers and bipartite matching loss for direct set prediction. Through conducting contrastive learning on 400 million image-text pairs, CLIP [42] shown impressive performance to solve different zero-shot transfer learning problems. For IQA tasks, Transformer also shows its powerful strength. Inspired by ViT, TRIQ [66] connected Transformer with MLP head to predict perceptual image quality, where sufficient lengths of positional embeddings were set to analyze the images with different resolutions. IQT [7] achieved outstanding performance by applying Transformer encoder and decoder on the features of reference images and distorted images. Through introducing 1D CNN and Transformer to integrate short-term and long-term temporal information, a recent work LSCT [65] demonstrated excellent performance on VQA. Our proposal is also derived from Transformer, inspired by our observation on the actions of human annotation. Experiments on various datasets confirm the effectiveness and robustness of our method.

3 Divide and Conquer Video Quality Estimator (DCVQE)

Human judgements of video quality are usually content-dependent and affected by their temporal memory [3, 11, 32, 47, 51, 68, 57]. In our investigation, we notice

that many human annotators like to give their opinions on the quality of a video after the following two actions: first, watch the video quickly (usually in the fast forward mode) to get an overall impression of its quality, then they may scroll mouse forward and backward to review some specific parts of the video for their final decisions. Inspired by this observation, we propose a hierarchical architecture, dubbed Divide and Conquer Video Quality Estimator (DCVQE) for NR-VQA. Our model is worked by extracting three levels of video quality representations from frames, video clips to whole video sequence progressively and repeatedly, somewhat similar to the reverse processes of human annotation. An additional correlation loss term is also presented to bring an additional order constraint of video quality to guide the training. We find that our method can effectively improve the performance of NR-VQA. We will describe our work in detail in the following paragraphs.

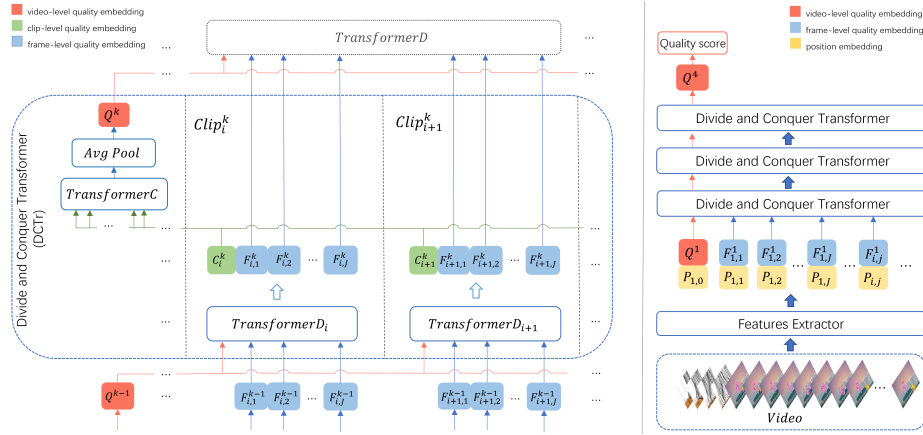


Fig. 2. The architectures of the proposed Divide and Conquer Transformer (DCTr) layer (left) and Divide and Conquer Video Quality Estimator (DCVQE) (right).

3.1 Overall Architecture

The left side of Fig. 2 represents the architecture of a key video quality analysis layer, Divide and Conquer Transformer (DCTr) in our proposal. In order to simulate the second action of human annotation mentioned above, we split the input sequence into a number of clips, and introduce a Transformer module $TransformerD$ to learn quality representations for each clip. As shown in Fig. 2, for the k^{th} DCTr layer, we split the whole sequence into I clips, and each clip covers J frame-level quality representations generated by the previous layer. For the i^{th} ($1 \leq i \leq I$) clip C_i^k , a module $TransformerD_i$ is applied to combine the two levels of quality embeddings (QE) generated by the previous layer, that

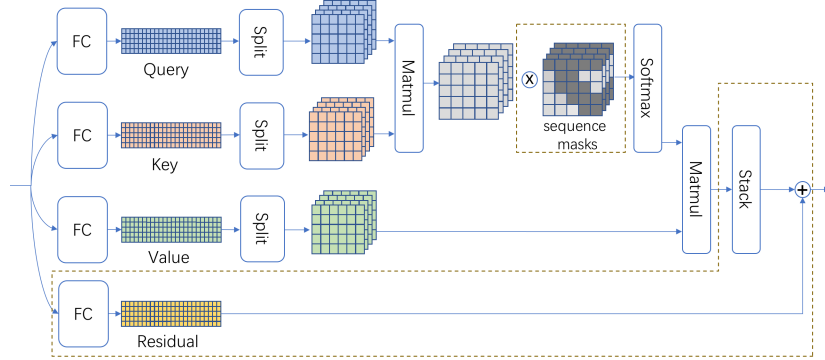


Fig. 3. The architecture of *TransformerD* (include dashed box) and *TransformerC* (exclude dashed box).

is, all the J frame-level QE $F_{i,j}^{k-1}$ ($1 \leq j \leq J$) and the video-level QE Q^{k-1} , to simultaneously learn the clip-level QE C_i^k and update the frame-level QE $F_{i,j}^k$ for the current layer. We further simulate the first action of human annotation by integrating the learned clip-level QE to generate a video-level QE. The other Transformer module *TransformerC* with a topped average pooling layer are proposed to merge all clip-level quality representations C_i^k to predict the video-level QE Q^k for the current layer. In general, our proposed video quality analysis module is constructed in a divide and conquer format, so we coin it as a Divide and Conquer Transformer (DCTr) layer.

The overall architecture of our DCVQE model is shown in the right side of Fig. 2. As seen, we stack several DCTr layers (3 layers for our proposal, refer to supplementary material for details about this setting) to extract the final video-level QE Q^k for the input video. The final quality score is predicted by topping this embedding Q^k with a regressor. In practice, to improve the temporal sensitivity of our model, we progressively expand the coverage of video clips in the DCTr as the layers deepen. For example, supposing that each clip in the k^{th} DCTr layer covers J frame-level QE from the previous layer, the coverage of each clip in the $(k+1)^{th}$ DCTr layer will be $2J$, which means two neighbor clips of the k^{th} DCTr layer are combined to extract the clip-level QE in the $(k+1)^{th}$ DCTr layer. For the first DCTr layer, the input frame-level QE $F_{*,*}^1$ is generated by our feature extractor described in subsection 3.2, and the input video-level QE Q^1 is initialized randomly. Both of them are integrated with positional embeddings $P_{*,*}$ [9] for the following processes.

TransformerD and *TransformerC* are the two most important modules in our DCVQE. As shown in Fig. 3, they have similar architecture to the classic Transformer [55]. Three fully connected layers are used to convert the input features into Query, Key, and Value tensors. Split operations are then adopted to group these three tensors into H heads, respectively. For the input frame-level

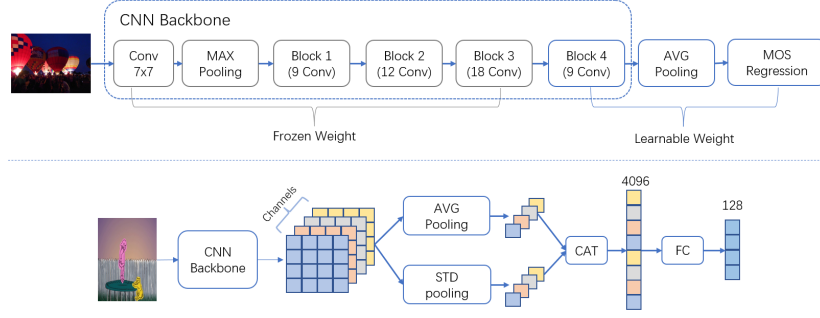


Fig. 4. Details of CNN backbone fine-tuning (top) and feature extraction (bottom).

feature with the shape of $B \times S \times D$, where B , S and D denote the batch size, sequence length and dimension of the frame embedding respectively, the shapes of its corresponding grouped Query, Key and Value tensors will be $H \times B \times S / H \times D$. The attention weights can be computed by the dot-product operations between the grouped Query and grouped Key tensors. To make *TransformerD* module more concerned with the clip-level quality information, except the first position is reserved for the input video-level QE, we mask out the attention weights outside a predefined temporal range by using sequence masks (the dashed square plotted in Fig. 3) and adopt a Softmax layer to normalize the rests. The normalized weights are then applied to the grouped Value tensor to update the quality features. A stack operation is conducted on the head axis to recover the shape of the tensor. Here we also introduce the residual technique [17] to improve the robustness of *TransformerD* module, so the final output of this module is the addition of the recovered features and residual. *TransformerC* module is designed to update the video-level QE for each DCTr layer. The difference between *TransformerC* and *TransformerD* modules is that *TransformerC* module does not contain sequence masks and residual block, as shown in Fig. 3.

3.2 Feature Extraction

To avoid the additional distortions that may be caused by applying some preprocessing steps to the input data, such as image resizing and filtering, we take the original full-size image frames as the input of our model. The pretrained Resnet-50 [17] is adopted as the CNN backbone of our feature extractor. To increase the sensitivity of the feature extractor to capture frame-level distortions, we first conduct a typical IQA task to fine-tune the CNN backbone. As shown in the top figure of Fig. 4, all CNN layers except the last group of CNN blocks are frozen for fine-tuning. We find this process can help to transfer the learned knowledge from ImageNet [8] to our application. It will be verified by our experiments.

Based on the fine-tuned backbone, as the bottom figure of Fig. 4 shows, the quality feature of each input frame is then extracted by adopting a concatenation

operator to the outputs of global average (AVG) pooling and global standard deviation (STD) pooling processes [23, 24]. Since both the dimensions of the outputs of AVG pooling and STD pooling are 2048, the frame-level QE for each frame will be a vector with a dimension of 4096. In practice, refer to [23], we further add a fully connected (FC) layer to reduce the dimension of the frame-level QE to 128. It is an efficient way to balance the accuracy of frame-level feature representation and overall computational cost.

3.3 Correlation loss

Prior works usually use the following L1 loss to optimize the model:

$$L_1 = \frac{1}{N} \sum_{n=1}^N |p_n - g_n| \quad (1)$$

where $p_n, n \in [1, N]$ represents the n^{th} predicted MOS, g_n denotes its corresponding ground truth, N is the total number of videos. Mathematically, this criterion to some extent ignores the relative order relationship of quality scores of the training samples in a batch, it sometimes may lead to limited ability to quantify the perceptual differences between the videos with similar quality scores. For example, considering the training processes of two NR-VQA models based on L1 loss criterion, for two video samples A and B with the MOS ground truths of 7.0 and 7.1, if their quality scores are predicted as $P_{A_1} = 7.1$, $P_{B_1} = 7.0$ and $P_{A_2} = 6.9$, $P_{B_2} = 7.0$ by the two models respectively (where P_{A_1} denotes the quality score of video A predicted by model 1, and so forth), the same strengths coming from the L1 losses (0.1) of these two training samples will be contributed to optimize the models correspondingly. Though it is hard to say which model is better than the other, in our opinion, it may be easier to optimize model 2 because of a positive correlation between the ground truths and its loss. Based on this idea, a new correlation loss, aiming to bring an additional order constraint of video quality to guide the training procedure, is proposed as follows:

$$L_c = \frac{1}{N} \sum_{n=1}^N \max \left(0, - \left(\sum_{m=1}^N (p_n - p_m) \right) \left(\sum_{m=1}^N (g_n - g_m) \right) \right) \quad (2)$$

where p_m represents the prediction of a training sample in a batch with size N , and p_n denotes the n^{th} prediction which is picked to compare the ranking order with each element in the prediction set. g_m and g_n stand for the corresponding ground truth scores of p_m and p_n , respectively. This loss equation can be considered as that, we take the n^{th} video in the batch as an anchor and compute the positive or negative correlations between predictions and ground truths for this anchor and any other one in the batch. Eq. 2 can be further simplified as:

$$L_c = \frac{1}{N} \sum_{n=1}^N \max \left(0, -N^2 \left(p_n - \sum_{m=1}^N \frac{p_m}{N} \right) \left(g_n - \sum_{m=1}^N \frac{g_m}{N} \right) \right) \quad (3)$$

$$L_c = N \sum_{n=1}^N \max (0, - (p_n - \bar{p}) (g_n - \bar{g})) \quad (4)$$

where \bar{p} and \bar{g} represent the mean values of predictions and ground truths, respectively. We find that this simplified L_c term is actually equivalent to the Spearman correlation coefficient without normalization. A total loss L , which combines the above two loss equations (1) and (4), is finally constructed to optimize our DCVQE model as

$$L = \alpha \times L_1 + \beta \times L_c \quad (5)$$

where α and β represent the weights of L_1 loss and L_c loss. We will discuss the optimal settings for these two weights in subsection 4.4.

4 Experiments

4.1 Datasets

We conduct experiments on four in-the-wild VQA datasets: KoNViD-1k [18], LIVE-VQC [48], YouTube-UGC [58], and LSVQ [63]. In these datasets, KoNViD-1k contains 1,200 unique contents with a duration of 8 seconds; LIVE-VQE, the smallest among them, consists of 585 video contents; YouTube-UGC contains 1,500 UGC video clips sampled from millions of YouTube videos; A recently released dataset LSVQ, containing 39,075 video samples with diverse durations and resolutions, is the largest publicly available UGC datasets for research right now. We also follow the suggestion given in [53, 40] to select YouTube-UGC dataset as the anchor and map MOS values of KoNViD-1k and LIVE-VQC data onto a common scale to generate an All-Combined dataset for performance evaluation. The typical random 60-20-20 strategy [23] is applied to split data into three sets, i.e. 60% for training, 20% for evaluation, and the remaining 20% for testing. Especially, for LSVQ we follow the setting in [63] to first generate a Test-1080p set and then conduct the random 80-20 splitting on the rests to generate training and testing sets for experiments.

4.2 Evaluation metrics

Here we adopt four commonly used criteria, Spearman Rank-Order Correlation Coefficient (SRCC), Kendall Rank-Order Correlation Coefficient (KRCC), Pearson Linear Correlation Coefficient (PLCC) and Root Mean Square Error (RMSE) to measure prediction monotonicity and prediction accuracy. For fair comparisons, we conduct each evaluation 100 times individually and report the median values of these four metrics as their final results.

4.3 Implementation Details

We first fine-tune ImageNet-pretrained Resnet-50 backbone on KonIQ-10k dataset [19] with random 80-20 splitting and maximum 20 epochs. The CNN backbone with the best performance of SRCC on the test set is selected to construct the feature extractor for our VQA task. The frame-level features are extracted using

this feature extractor, where all frames of each video are considered without sampling. The temporal range of the sequence mask in the proposed *TransformerD* (see Fig. 3) is empirically set to 15. For a tradeoff between GPU memory, running speed, and performance, we set the maximum length of each input video to 600 (i.e. only the first 600 frames of a video are considered if the video is longer than this maximum length) and the maximum training epoch to 75. An evaluation process will be conducted after each epoch, and the model with the lowest loss on the validation set will be stored.

4.4 Ablation Studies

In this subsection, we perform ablation studies to better understand the different components of the proposed DCVQE model.

CNN Backbone: Training of the CNN feature extraction backbone can also be regarded as a typical IQA task. Two series of architectures based on ResNet and ViT backbones are fine-tuned fully or partially in our study. Here we note the partial fine-tuning means that only the last blocks of these models, topped with average pooling and fully connected layers (see Fig. 4) are trainable. We determine that fully fine-tuned backbones act negatively with regard to the complexities of the networks. i.e. deeper structure results in worse performance, and partial fine-tuned backbones usually work better than fully fine-tuned ones for both two series of architectures. We also incorporate a new attention-based IQA architecture PHIQNet [65] to our study, and similar performance as partial fine-tuning on ResNet is observed (refer to supplementary material for details). In our application, for a tradeoff between model complexity and performance, also a fair comparison with previous work [23], we select the partially fine-tuned ResNet-50 as the feature extraction backbone to construct our DCVQE model.

Temporal Range of Sequence Mask: We introduce a group of experiments on KoNViD-1k dataset to study how the temporal range affects the DCVQE model. By directly replacing the proposed DCTr layers with the traditional Transformer layers [55], a baseline Transformer model is constructed for performance comparisons. Test results are plotted in Fig. 5. From these figures, we can see that the best temporal range for the baseline Transformer model is about 9, and the performance of this baseline model slightly decreases as the temporal range increases. Especially if we set the range to “all”, which means that all the frame-level representations are involved in the self-attention processes [55] of the Transformer, the performance of this baseline model drops dramatically. Switching to our DCVQE, we can see that the overall performance stays at a high level. The best performance can be observed when the temporal range reaches 15. The results of the “all” setting for DCVQE, for which one frame only conducts self-attention with all frames in its own clip, also demonstrate the robust performance of our DCVQE model.

Correlation Loss: To show how the proposed Correlation Loss (CL) contributes to our VQA task, the performances of DCVQE with different weight settings of CL and L1 loss terms are reported in Table 1. The best performance can be identified when L1 Loss weight α reaches 0.7 and CL weight β reaches

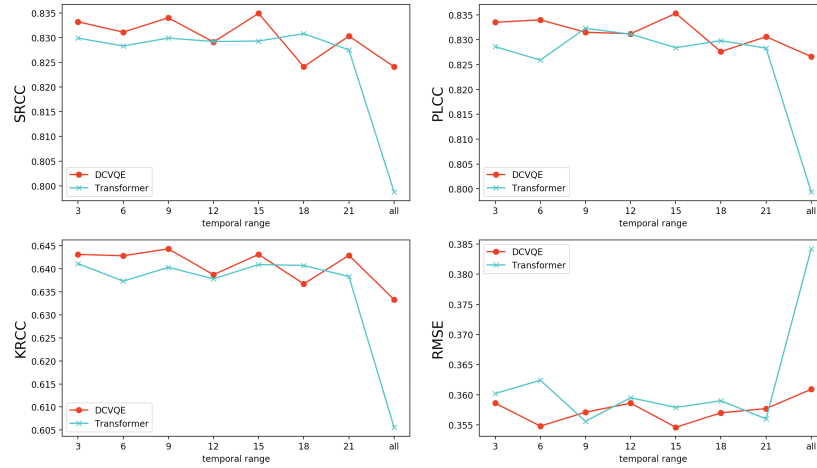


Fig. 5. Performance comparisons between a baseline Transformer model and our DCVQE with different settings of the temporal range.

Table 1. Performance comparisons of different weight combinations of α (for L1 loss term) and β (for proposed CL term) on KoNViD-1k dataset.

α	β	SRCC	PLCC	KRCC	RMSE
0.0	1.0	0.6980	0.6956	0.5149	3.0966
0.3	0.7	0.8259	0.8264	0.6360	0.3924
0.5	0.5	0.8349	0.8353	0.6431	0.3546
0.7	0.3	0.8382	0.8375	0.6500	0.3515
1.0	0.0	0.8278	0.8306	0.6396	0.3623

0.3. Here we note that our CL does not intend to play a critical role in model optimization since it tries to describe the relative order relationships between videos, particularly for those videos with similar quality scores, so the overall improvement by adding this CL term may not look very significant. Overall, we confirm that our model can achieve better performance with the introduction of CL (Eq. 4). The rest experiments are conducted with the setting of $\alpha = 0.7$ and $\beta = 0.3$ if not specified.

Component by Component Ablation: We further stack the key parts in our DCVQE component-by-component to analyze their impacts individually and gradually. To additionally study the effectiveness of our new CL term, a well-known pairwise ranking loss (PW-RL) [4], which is similar to our proposal, is also involved in our experiments. As shown in Table 2, compared with Vanilla-Transformer (row 1), only introduces the Divide operation (row 2) does not impact the system’s performance. In row 3 we employ the AVG pooling operator to integrate clip-level embeddings to generate video-level embeddings,

Table 2. Component-by-component comparisons of several key parts in our DCVQE on KoNViD-1k dataset. Here we note that L1 loss is a default setting for all tests. Rows 3 and 4 apply AVG pooling (AVG) and *Transformer_C* (*Trans_C*) respectively to conduct the “conquer” operation. “TR 15”s in rows 5, 6, 7, and 8 mean the temporal ranges of sequence masks are set to 15 for the tests.

Row #	Components	SRCC	PLCC	RMSE
1	Vanilla-Transformer	0.7988	0.7994	0.3842
2	Divide Only	0.7987	0.8032	0.3723
3	Divide + AVG	0.8025	0.8140	0.3644
4	Divide + <i>Trans_C</i>	0.8031	0.8121	0.3608
5	Divide + AVG + TR 15	0.8181	0.8267	0.3474
6	Divide + <i>Trans_C</i> + AVG + TR 15	0.8278	0.8289	0.3453
7	Divide + <i>Trans_C</i> + AVG + TR 15 + PW-RL	0.8307	0.8326	0.3542
8	Divide + <i>Trans_C</i> + AVG + TR 15 + CL	0.8382	0.8375	0.3515

and the results show that this simple “conquer” operation can slightly improve the performance. In row 4 the conquer operation is switched to the proposed *Transformer_C*, where similar correlation scores with a lower RMSE are presented (compared with row 3). The temporal range of sequence mask 15 is introduced to tests in rows 5 and 6, where we can see that this change apparently benefits the system. Furthermore, the losses CL and PW-RL are compared directly in rows 7 and 8. The results confirm that the proposal CL provides the best performance for our VQA task.

4.5 Comparison with the State-of-the-art Methods

We compare our method with 13 state-of-the-art VQA methods, which can be roughly divided into NR-IQA based methods: BRISQUE [33], GM-LOG [61], HIGRADE [22], FRIQUEE [14], CORNIA [62], HOSA [60], Koncept-512 [19], PaQ-2-PiQ [64], and NR-VQA based methods: V-BLIINDS [46], TLVQM [21], VSFA [23], VIDEVAL [53] and RAPIQUE [54]. For NR-IQA based methods, we abstract one frame feature per second and regard the average value of the features as the video representation to train a support vector regressor (SVR) for prediction [54, 53, 63].

Test results are reported in Table 3 and the best solution for each evaluation metric is marked in bold. Since it is hard to reproduce VIDEVAL and RAPIQUE methods, in Table 3 we directly cite the results given in [53, 54] (marked as “*”). Different from our default 60-20-20 data splitting setting, these two methods used the random 80-20 strategy to split data for training and testing. For a fair comparison, we train an additional version of the DCVQE model under the same 80-20 data splitting setting. Test results of this additional DCVQE model are marked as “†”. As seen from this table, our model (DCVQE†) significantly outperforms the current top method RAPIQUE on KoNViD-1K dataset by 4.37% SRCC and 2.45% PLCC, and greatly exceeds the second place method VIDEVAL on YouTube-UGC dataset by 5.62% SRCC and 5.67% PLCC. TLVQM

method performs best on LIVE-VQC dataset. Looking into this dataset, we can see that most videos are captured using mobile devices, thus camera motion blurring is one of the common distortions affecting the video quality. TLVQM introduces several hand-crafted motion-related features to handle this issue, so it achieves the best performance, especially compared with all the deep learning based solutions. Nevertheless, besides TLVQM, our proposal performs the best among the remaining. The evaluation on the All-Combined dataset of KoNViD-1k, LIVE-VQC, YouTube-UGC also confirms the robustness of our model. As shown in the column “All Combined”, our method (DCVQE†) surpasses the second place method (RAPIQUE) by 2.84% SRCC and 1.12% PLCC.

Table 3. Performance comparisons with the state-of-the-art methods.

DATASET MODEL	KoNViD-1k			LIVE-VQC			YouTube-UGC			All Combined		
	SRCC	PLCC	RMSE	SRCC	PLCC	RMSE	SRCC	PLCC	RMSE	SRCC	PLCC	RMSE
BRISQUE	0.6567	0.6576	0.4813	0.5925	0.6380	13.100	0.3820	0.3952	0.5919	0.5695	0.5861	0.5617
GM-LOG	0.6578	0.6636	0.4818	0.5881	0.6212	13.223	0.3678	0.3920	0.5896	0.5650	0.5942	0.5588
HIGRADE	0.7206	0.7269	0.4391	0.6103	0.6332	13.027	0.7376	0.7216	0.4471	0.7398	0.7368	0.4674
FRIQUEE	0.7472	0.7482	0.4252	0.6579	0.7000	12.198	0.7652	0.7571	0.4169	0.7568	0.7550	0.4549
CORNIA	0.7169	0.7135	0.4486	0.6719	0.7183	11.832	0.5972	0.6057	0.5136	0.6764	0.6974	0.4946
HOSA	0.7654	0.7664	0.4142	0.6873	0.7414	11.353	0.6025	0.6047	0.5132	0.6957	0.7082	0.4893
KonCept-512	0.7349	0.7489	0.4260	0.6645	0.7278	11.626	0.5872	0.5940	0.5135	0.6608	0.6763	0.5091
PaQ-2-PiQ	0.6130	0.6014	0.5148	0.6436	0.6683	12.619	0.2658	0.2935	0.6153	0.4727	0.4828	0.6081
V-BLINDS	0.7101	0.7037	0.4595	0.6939	0.7178	11.765	0.5590	0.5551	0.5356	0.6545	0.6599	0.5200
TLVQM	0.7729	0.7668	0.4102	0.7988	0.8025	10.145	0.6693	0.6590	0.4849	0.7271	0.7342	0.4705
VSFA	0.7728	0.7754	0.4205	0.6978	0.7426	11.649	0.7611	0.7500	0.4269	0.7690	0.7862	0.4253
VIDEVAL*	0.7832*	0.7803*	0.4026*	0.7522*	0.7514*	11.100*	0.7787*	0.7733*	0.4049*	0.7960*	0.7939*	0.4268*
RAPIQUE*	0.8031*	0.8175*	0.3623*	0.7548*	0.7863*	10.518*	0.7591*	0.7684*	0.4060*	0.8070*	0.8279*	0.3968*
DCVQE†	0.8382†	0.8375†	0.3515†	0.7620†	0.7858†	10.549†	0.8225†	0.8172†	0.3770†	0.8299†	0.8372†	0.3824†
DCVQE	0.8206	0.8224	0.3671	0.7479	0.7648	11.599	0.8069	0.8050	0.3974	0.8239	0.8350	0.3914

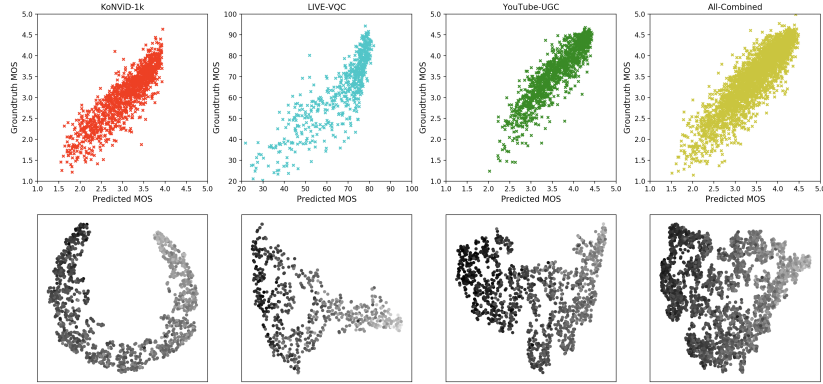


Fig. 6. Visualize test results. Top figures: Video quality predictions versus their ground truth MOS scores; Bottom figures: dimension-reduced video-level QE colored by ground truth MOS scores, where we apply t-SNE algorithm [30] to reduce the dimension of the embeddings to 2 for display.

To demonstrate how well our model is learned from a given dataset, in Fig. 6 we also visualize the correlation between the video quality predictions and their ground truths (top row), as well as the video-level QE learned by the “conquer” operation of the last DCTr layer (bottom row). From the top row, we can see a strong correlation between our predictions and their ground truths. The figures in the bottom row plot the dimension-reduced video-level QE, where the gray value of each point is assigned with the normalized ground truth MOS score of the corresponding video sample, that is, the darker a point, the greater its corresponding MOS score. From these figures, we can observe that the points with higher MOS scores scatter on one side while the points with lower MOS on the other. It indicates that the separability of quality for those video samples has been significantly enhanced by using our DCVQE model.

Table 4. Performance comparisons on the large-scale LSVQ dataset.

DATASET	Test		Test-1080p	
MODEL	SRCC	PLCC	SRCC	PLCC
BRISQUE	0.579	0.576	0.497	0.531
TLVQM	0.772	0.774	0.589	0.616
VIDEVAL	0.794	0.783	0.545	0.554
VSFA	0.801	0.796	0.675	0.704
PVQ	0.827	0.828	0.711	0.739
DCVQE	0.836	0.834	0.727	0.758

Additionally, we compare our model with several others on the recently released large-scale LSVQ dataset. The PVQ model [63], which was proposed along with this dataset, is also involved in our evaluation. Test results listed in Table 4 confirm the excellent performance of our model again.

5 Conclusions

Inspired by our observation on the actions of human annotation, in this paper, we propose a new Divide and Conquer Transformer (DCTr) architecture to extract the video quality features for NR-VQA. Starting from extracting the frame-level quality embeddings (QE) of an input video, two types of Transformers are introduced to extract the clip-level QE and video-level QE progressively in our DCTr layer. Through stacking several DCTr layers and topping with a regressor, a hierarchical model, named Divide and Conquer Video Quality Estimator (DCVQE) is constructed to predict the quality score of the input video. We also put forward an additional correlation loss regarding the order relationship among the training data to guide the training. Experiments confirm that our proposal outperforms most other methods. What is more, our model is purely deep learning based, compared with other top methods where NSS/NVS features are more or less needed, so we believe our proposal is more practical.

References

1. Baccouche, M., Mamalet, F., Wolf, C., Garcia, C., Baskurt, A.: Sequential deep learning for human action recognition. In: International workshop on human behavior understanding. pp. 29–39. Springer (2011)
2. Ballas, N., Yao, L., Pal, C., Courville, A.: Delving deeper into convolutional networks for learning video representations. CoRR (2016)
3. Bampis, C.G., Li, Z., Moorthy, A.K., Katsavounidis, I., Aaron, A., Bovik, A.C.: Study of temporal effects on subjective video quality of experience. *IEEE Transactions on Image Processing* **26**(11), 5217–5231 (2017)
4. Cao, Z., Qin, T., Liu, T.Y., Tsai, M.F., Li, H.: Learning to rank: From pairwise approach to listwise approach. vol. 227, pp. 129–136 (01 2007)
5. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European conference on computer vision. pp. 213–229. Springer (2020)
6. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4724–4733 (2017)
7. Cheon, M., Yoon, S.J., Kang, B., Lee, J.: Perceptual image quality assessment with transformers. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 433–442 (2021)
8. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. pp. 248–255 (2009)
9. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: NAACL-HLT (2019)
10. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
11. Duanmu, Z., Ma, K., Wang, Z.: Quality-of-experience of adaptive video streaming: Exploring the space of adaptations. In: Proceedings of the 25th ACM international conference on Multimedia. pp. 1752–1760 (2017)
12. Feichtenhofer, C.: X3d: Expanding architectures for efficient video recognition. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 200–210 (2020)
13. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: 2019 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 6201–6210 (2019)
14. Ghadiyaram, D., Bovik, A.: Perceptual quality prediction on authentically distorted images using a bag of features approach. *Journal of Vision* **17** (09 2016)
15. Ghadiyaram, D., Bovik, A.C.: Massive online crowdsourced study of subjective and objective picture quality. *IEEE Transactions on Image Processing* **25**(1), 372–387 (2016)
16. Ghadiyaram, D., Pan, J., Bovik, A.C., Moorthy, A.K., Panda, P., Yang, K.C.: 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)
17. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016)

18. Hosu, V., Hahn, F., Jenadeleh, M., Lin, H., Men, H., Szirányi, T., Li, S., Saupe, D.: The konstanz natural video database (konvid-1k). In: 2017 Ninth International Conference on Quality of Multimedia Experience (QoMEX). pp. 1–6 (2017)
19. Hosu, V., Lin, H., Szirányi, T., Saupe, D.: Koniq-10k: An ecologically valid database for deep learning of blind image quality assessment. *IEEE Transactions on Image Processing* **29**, 1–1 (01 2020)
20. Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., Fei-Fei, L.: Large-scale video classification with convolutional neural networks. In: 2014 IEEE Conference on Computer Vision and Pattern Recognition. pp. 1725–1732 (2014)
21. Korhonen, J.: Two-level approach for no-reference consumer video quality assessment. *IEEE Transactions on Image Processing* **28**(12), 5923–5938 (2019)
22. Kundu, D., Ghadiyaram, D., Bovik, A.C., Evans, B.L.: No-reference quality assessment of tone-mapped hdr pictures. *IEEE Transactions on Image Processing* **26**(6), 2957–2971 (2017)
23. Li, D., Jiang, T., Jiang, M.: Quality assessment of in-the-wild videos. *Proceedings of the 27th ACM International Conference on Multimedia* (Oct 2019)
24. Li, D., Jiang, T., Jiang, M.: Unified quality assessment of in-the-wild videos with mixed datasets training. *International Journal of Computer Vision* **129**(4), 1238–1257 (Jan 2021)
25. Li, X., Wang, Y., Zhou, Z., Qiao, Y.: Smallbignet: Integrating core and contextual views for video classification. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1089–1098 (2020)
26. Li, X., Guo, Q., Lu, X.: Spatiotemporal statistics for video quality assessment. *IEEE Transactions on Image Processing* **25**(7), 3329–3342 (2016)
27. Li, Y., Meng, S., Zhang, X., Wang, S., Wang, Y., Ma, S.: Ugc-video: Perceptual quality assessment of user-generated videos. In: 2020 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR). pp. 35–38 (2020)
28. Liu, X., van de Weijer, J., Bagdanov, A.D.: Rankiq: Learning from rankings for no-reference image quality assessment. In: 2017 IEEE International Conference on Computer Vision (ICCV). pp. 1040–1049 (2017)
29. Ma, K., Liu, W., Zhang, K., Duanmu, Z., Wang, Z., Zuo, W.: End-to-end blind image quality assessment using deep neural networks. *IEEE Transactions on Image Processing* **27**(3), 1202–1213 (2018)
30. van der Maaten, L., Hinton, G.: Visualizing data using t-sne. *Journal of Machine Learning Research* **9**, 2579–2605 (11 2008)
31. Marziliano, P., Dufaux, F., Winkler, S., Ebrahimi, T.: A no-reference perceptual blur metric. In: *Proceedings. International Conference on Image Processing*. vol. 3 (2002)
32. Mirkovic, M., Vrgović, P., Stefanović, D., Anderla, A.: Evaluating the role of content in subjective video quality assessment. *The Scientific World Journal* **2014**, 625219 (01 2014)
33. Mittal, A., Moorthy, A.K., Bovik, A.C.: No-reference image quality assessment in the spatial domain. *IEEE Transactions on Image Processing* **21**(12), 4695–4708 (2012)
34. Mittal, A., Saad, M.A., Bovik, A.C.: A completely blind video integrity oracle. *IEEE Transactions on Image Processing* **25**(1), 289–300 (2016)
35. Mittal, A., Soundararajan, R., Bovik, A.C.: Making a “completely blind” image quality analyzer. *IEEE Signal Processing Letters* **20**(3), 209–212 (2013)
36. Moorthy, A.K., Bovik, A.C.: A two-step framework for constructing blind image quality indices. *IEEE Signal Processing Letters* **17**(5), 513–516 (2010)

37. Moorthy, A.K., Bovik, A.C.: Blind image quality assessment: From natural scene statistics to perceptual quality. *IEEE Transactions on Image Processing* **20**(12), 3350–3364 (2011)
38. Nuutinen, M., Virtanen, T., Vaahteranoksa, M., Vuori, T., Oittinen, P., Häkkinen, J.: Cvd2014—a database for evaluating no-reference video quality assessment algorithms. *IEEE Transactions on Image Processing* **25**(7), 3073–3086 (2016)
39. Pan, F., Lin, X., Rahardja, S., Ong, E.P., Lin, W.: Using edge direction information for measuring blocking artifacts of images. *Multidimensional Systems and Signal Processing* **18**, 297–308 (2007)
40. Pinson, M., Wolf, S.: Objective method for combining multiple subjective data sets. pp. 583–592 (06 2003)
41. Ponomarenko, N., Ieremeiev, O., Lukin, V., Egiazarian, K., Jin, L., Astola, J., Vozel, B., Chehdi, K., Carli, M., Battisti, F., Kuo, C.C.J.: Color image database tid2013: Peculiarities and preliminary results. In: *European Workshop on Visual Information Processing (EUVIP)*. pp. 106–111 (2013)
42. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International Conference on Machine Learning*. pp. 8748–8763. PMLR (2021)
43. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **39**(6), 1137–1149 (2017)
44. Saad, M.A., Bovik, A.C., Charrier, C.: A dct statistics-based blind image quality index. *IEEE Signal Processing Letters* **17**(6), 583–586 (2010)
45. Saad, M.A., Bovik, A.C., Charrier, C.: Blind image quality assessment: A natural scene statistics approach in the dct domain. *IEEE Transactions on Image Processing* **21**(8), 3339–3352 (2012)
46. Saad, M.A., Bovik, A.C., Charrier, C.: Blind prediction of natural video quality. *IEEE Transactions on Image Processing* **23**(3), 1352–1365 (2014)
47. Siahaan, E., Hanjalic, A., Redi, J.: Semantic-aware blind image quality assessment. *Signal Processing: Image Communication* **60** (10 2017)
48. Sinno, Z., Bovik, A.C.: Large-scale study of perceptual video quality. *IEEE Transactions on Image Processing* **28**(2), 612–627 (2019)
49. Tai, S.C., Yang, S.M.: A fast method for image noise estimation using laplacian operator and adaptive edge detection. In: *2008 3rd International Symposium on Communications, Control and Signal Processing*. pp. 1077–1081 (2008)
50. Talebi, H., Milanfar, P.: Nima: Neural image assessment. *IEEE Transactions on Image Processing* **27**(8), 3998–4011 (2018)
51. Triantaphillidou, S., Allen, E., Jacobson, R.: Image quality comparison between jpeg and jpeg2000. ii. scene dependency, scene analysis, and classification. *Journal of Imaging Science and Technology - J IMAGING SCI TECHNOL* **51** (05 2007)
52. Tu, Z., Chen, C.J., Chen, L.H., Birkbeck, N., Adsumilli, B., Bovik, A.C.: A comparative evaluation of temporal pooling methods for blind video quality assessment. In: *2020 IEEE International Conference on Image Processing (ICIP)*. pp. 141–145 (2020)
53. Tu, Z., Wang, Y., Birkbeck, N., Adsumilli, B., Bovik, A.C.: Ugc-vqa: Benchmarking blind video quality assessment for user generated content. *IEEE Transactions on Image Processing* **30**, 4449–4464 (2021)
54. Tu, Z., Yu, X., Wang, Y., Birkbeck, N., Adsumilli, B., Bovik, A.C.: Rapique: Rapid and accurate video quality prediction of user generated content. *IEEE Open Journal of Signal Processing* **2**, 425–440 (2021)

55. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
56. Wang, H., Gan, W., Hu, S., Lin, J.Y., Jin, L., Song, L., Wang, P., Katsavounidis, I., Aaron, A., Kuo, C.C.J.: Mcl-jcv: A jnd-based h.264/avc video quality assessment dataset. In: 2016 IEEE International Conference on Image Processing (ICIP). pp. 1509–1513 (2016)
57. Wang, H., Katsavounidis, I., Zhou, J., Park, J., Lei, S., Zhou, X., Pun, M.O., Jin, X., Wang, R., Wang, X., Zhang, Y., Huang, J., Kwong, S., Kuo, C.C.J.: Videoset: A large-scale compressed video quality dataset based on jnd measurement. *Journal of Visual Communication and Image Representation* **46** (01 2017)
58. Wang, Y., Inguva, S., Adsumilli, B.: Youtube ugc dataset for video compression research. In: 2019 IEEE 21st International Workshop on Multimedia Signal Processing (MMSP). pp. 1–5 (2019)
59. Woojae, K., Kim, J., Ahn, S., Kim, J., Lee, S.: Deep video quality assessor: From spatio-temporal visual sensitivity to a convolutional neural aggregation network. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *Computer Vision – ECCV 2018*. pp. 224–241. Springer International Publishing, Cham (2018)
60. Xu, J., Ye, P., Li, Q., Du, H., Liu, Y., Doermann, D.: Blind image quality assessment based on high order statistics aggregation. *IEEE Transactions on Image Processing* **25**(9), 4444–4457 (2016)
61. Xue, W., Mou, X., Zhang, L., Bovik, A.C., Feng, X.: Blind image quality assessment using joint statistics of gradient magnitude and laplacian features. *IEEE Transactions on Image Processing* **23**(11), 4850–4862 (2014)
62. Ye, P., Kumar, J., Kang, L., Doermann, D.: Unsupervised feature learning framework for no-reference image quality assessment. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition. pp. 1098–1105 (2012)
63. Ying, Z., Mandal, M., Ghadiyaram, D., Bovik, A.: Patch-vq: patching up the video quality problem. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 14019–14029 (2021)
64. Ying, Z., Niu, H., Gupta, P., Mahajan, D., Ghadiyaram, D., Bovik, A.: From patches to pictures (paq-2-piq): Mapping the perceptual space of picture quality. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3572–3582 (2020)
65. You, J.: Long short-term convolutional transformer for no-reference video quality assessment. pp. 2112–2120 (10 2021)
66. You, J., Korhonen, J.: Transformer for image quality assessment. In: 2021 IEEE International Conference on Image Processing (ICIP). pp. 1389–1393. IEEE (2021)
67. Zhang, L., Zhang, L., Bovik, A.C.: A feature-enriched completely blind image quality evaluator. *IEEE Transactions on Image Processing* **24**(8), 2579–2591 (2015)
68. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 586–595 (2018)
69. Zhang, Y., Gao, X., He, L., Lu, W., He, R.: Blind video quality assessment with weakly supervised learning and resampling strategy. *IEEE Transactions on Circuits and Systems for Video Technology* **29**(8), 2244–2255 (2019)
70. Zhu, H., Li, L., Wu, J., Dong, W., Shi, G.: Metaiqa: Deep meta-learning for no-reference image quality assessment. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 14131–14140 (2020)