

Multiview Contrastive Learning for Completely Blind Video Quality Assessment of User Generated Content

Shankhanil Mitra
Indian Institute of Science
Bangalore, Karnataka, India

Rajiv Soundararajan
Indian Institute of Science
Bangalore, Karnataka, India

ABSTRACT

Completely blind video quality assessment (VQA) refers to a class of quality assessment methods that do not use any reference videos, human opinion scores or training videos from the target database to learn a quality model. The design of this class of methods is particularly important since it can allow for superior generalization in performance across various datasets. We consider the design of completely blind VQA for user generated content. While several deep feature extraction methods have been considered in supervised and weakly supervised settings, such approaches have not been studied in the context of completely blind VQA. We bridge this gap by presenting a self-supervised multiview contrastive learning framework to learn spatio-temporal quality representations. In particular, we capture the common information between frame differences and frames by treating them as a pair of views and similarly obtain the shared representations between frame differences and optical flow. The resulting features are then compared with a corpus of pristine natural video patches to predict the quality of the distorted video. Detailed experiments on multiple camera captured VQA datasets reveal the superior performance of our method over other features when evaluated without training on human scores. Code will be made available at <https://github.com/Shankhanil006/VISION>.

CCS CONCEPTS

• **Computing methodologies** → **Computer vision; Image processing.**

KEYWORDS

Multiview contrastive learning, blind video quality assessment, user generated content

1 INTRODUCTION

The ubiquity of mobile devices and video sharing platforms has led to an explosion in the number of videos captured, processed, and shared. Thus, the perceptual quality assessment of such videos is of paramount importance in enabling a better user experience. The video quality assessment (VQA) of such user generated content obtained through camera captures is challenging on multiple counts. The authentic distortions that arise in user generated content are more complex and often suffer from multiple sources of video degradation at the same time. Further, unavailability of reference video motivates the study of no reference (NR) VQA.

NR VQA has been studied quite extensively in the supervised setting [3, 16–18, 37] for user generated content (UGC), where video features are regressed against human scores to learn a video quality model. However, such an approach suffers from two main drawbacks. Firstly, it requires a large number of human ratings to train a model. Secondly, the generalization performance of such

supervised approaches has also been found to be a limitation [19]. This motivates the need for completely blind NR VQA approaches where neither the human ratings nor the videos from the target database are used in any step of the algorithm design.

The problem of designing completely blind NR VQA where neither human scores nor the videos in the target database are used has been researched to some extent. The VIIDEO [26] model represents one such example where the models based on the statistics of natural videos are designed to measure video quality. Recently, the perceptual straightening hypothesis of the temporal information of natural videos has been used to design completely blind NR VQA [14]. The natural image quality evaluator (NIQE) [27] which is a completely blind image quality metric has also had limited success in NR VQA. Nevertheless, the role of deep features for completely blind NR VQA has not been explored to the best of our knowledge.

In this work, we explore self-supervised contrastive learning of video quality features and deploy them to predict video quality without training on human labels. Self-supervised learning for video quality representations has been explored recently using predicted frames as augmentations for contrastive learning [5]. However, the inaccuracy of the video prediction algorithm can limit the performance of the quality representation learning. Further, the learnt features have only been evaluated in a supervised setup as opposed to an unsupervised prediction of quality we explore in our work. HEKE [22] learns spatio-temporal representations from synthetically distorted videos suffering from artifacts due to compression and packet losses. But such a model may not perform well on authentic distortions which arise during camera captures.

Our main contribution in this work is the design of a multiview contrastive learning framework for quality representation learning. We refer to our framework and the resulting quality metric as Video quality Index uSing multiview cOntrastive learNing (**VISION**). We observe that frame differences are often interpreted as moving edges and contain information about the spatial and temporal quality. Thus, their joint distribution with frames and optical flow can be used to extract interesting quality related features. In particular, we capture the joint distribution of quality features in frames and frame differences for predicting video quality. Similarly, we also extract features that correspond to the joint distribution of frame differences and optical flow for quality prediction. The two sets of features are then compared against a corpus of such features from pristine videos to predict video quality.

We conduct detailed experiments on multiple datasets to show the effective performance of our model. In particular, we also design novel benchmarks for comparison, where we evaluate several features learnt using a variety of approaches for unsupervised quality prediction without training them on human scores. This analysis is

interesting in understanding the performance of the features for predicting video quality without human training.

In summary, the main contributions of this work are as follows:

- We design a multiview contrastive learning approach for learning quality representations for user generated content. In particular, our features capture the joint distributions of frame differences with frames and optical flow.
- We show through detailed experiments on multiple datasets that the learnt features can be used to effectively predict video quality by comparing with a corpus of such features from pristine videos.
- We evaluate the performance of several interesting features that can be learnt without human scores and evaluated for opinion unaware quality assessment by comparing against a corpus of pristine videos.

2 RELATED WORK

Supervised NR VQA: One of the most successful NR VQA approaches models natural scene statistics to generate quality aware features. Handcrafted methods involve designing a statistical model based upon discrete cosine transform (DCT) coefficients of frame differences [31], 3D DCT coefficients [20], and 3D mean subtracted contrast normalized coefficients [29] of video clips, and optical flow based [24]. Several other NR VQA methods have been developed by taking into account blockiness, sharpness, noise, and temporal correlation in videos [2, 8, 43]. TLVQM [16] deploys low complexity and high complexity features from video frames based on spatial, and temporal statistics of videos. VIDEVAL [37] combines handcrafted features from several existing blind VQA algorithms.

In recent years, convolutional neural networks (CNN) have been used to extract quality aware features from videos. One of the approaches involves combining CNN based features with other heuristics based features to achieve state-of-the-art NR VQA performance [1, 17, 38]. On the other hand, an end-to-end deep learning based method was designed to predict compressed video quality for specific codecs [21]. Motion representation based models [3] have also been studied for NR VQA. Features extracted by learning a 3D CNN model on video clips [47], and pretrained ResNet50 [12] features learnt on ImageNet database have been fed to a recurrent model to predict quality [18]. Quality aware features have also been extracted from image quality models such a PaQ-2-PiQ [46] and combined with pretrained 3D Resnet18 features to predict global video quality [45]. UCDA [4] adopts a domain adaptation approach to adapt a model learnt on synthetic video distortions to authentic video distortions. While UCDA does not use any human quality scores for the target database, it trains on a large set of human annotated labels for the source database.

Pseudo-label Training for NR VQA: A broad class of opinion unaware quality models is designed by learning to predict full reference metrics available in the training data. These models are relevant for synthetic distortions due to the need for a reference for generating the target quality index. Video CORNIA [42] was designed by first learning quality aware features through visual codebooks in an unsupervised manner and then regressing these features against frame level full reference quality measures. A weakly supervised approach was adopted for pretraining, where a deep network was

first trained to predict a full reference quality measure. The features were then fine-tuned on human opinion scores [49]. Pseudo-labels from multiple full reference video quality measures are used to obtain a richer quality prediction method on account of the heterogeneity of the measures [22]. A similar approach is adopted based on the spatio-temporal entropic differences index exploiting the complementarity of the spatial and temporal streams [25].

Completely Blind NR VQA: The class of completely blind VQA algorithms is very challenging and very few algorithms have been designed. One of the first such methods was the VIIDEO method [26], which models the intrinsic statistical regularities in natural videos to measure the disturbances in the presence of distortions. STEM [14] is a recent completely blind VQA model designed for user generated content that exploits the human perception of straighter temporal trajectories for natural videos in a transformed space. The loss in straightness or the resulting curvature in the presence of distortions is used to measure quality. This approach is used in conjunction with the completely blind natural image quality evaluator to obtain impressive performance.

Self-supervised learning for image and video quality: The role of self-supervised feature learning in VQA is still nascent. Contrastive self-supervised pretraining [5] has been explored for learning deep video quality features for supervised NR VQA. In particular, frame prediction algorithms are deployed to create augmented videos which are used to learn quality distortion and content features for NR VQA. CONTRIQUE [23] uses information about distortion levels and types of synthetically distorted images, along with authentically distorted images for supervised NR image quality assessment. Nevertheless, self-supervised methods are only being deployed in a supervised setup, while such approaches offer the possibility of completely blind VQA, which we explore in our work.

3 METHOD

3.1 Overview

An overview of our quality aware representation learning from videos using multiple views of video frames without any human subjective score is shown in Figure 1. Our approach is inspired by the success of contrastive multiview coding (CMC) [35] in image classification, where the goal is to elicit features that maximize the shared information contained in both the views. This can be achieved by maximizing the mutual information between the resulting features. However, we adapt the CMC method to create positive and negative view pairs that vary only in quality to learn rich quality representations through contrastive learning. The key question that arises in CMC is the choice of views. We note that frame differences in videos contain information about variations in quality across frames and thus some information about the spatial quality of the frames. Further, frame differences also contain temporal information and the common information with optical flow may be used to capture temporal quality. We capture the common information between frame differences and frames and similarly between frame differences and optical flow to predict video quality. We note that the choice of views in our framework for quality representation learning is different from the choice of views in the original CMC work [35]. In the following sections, we first discuss how we generate a set of videos, consisting of variations of the

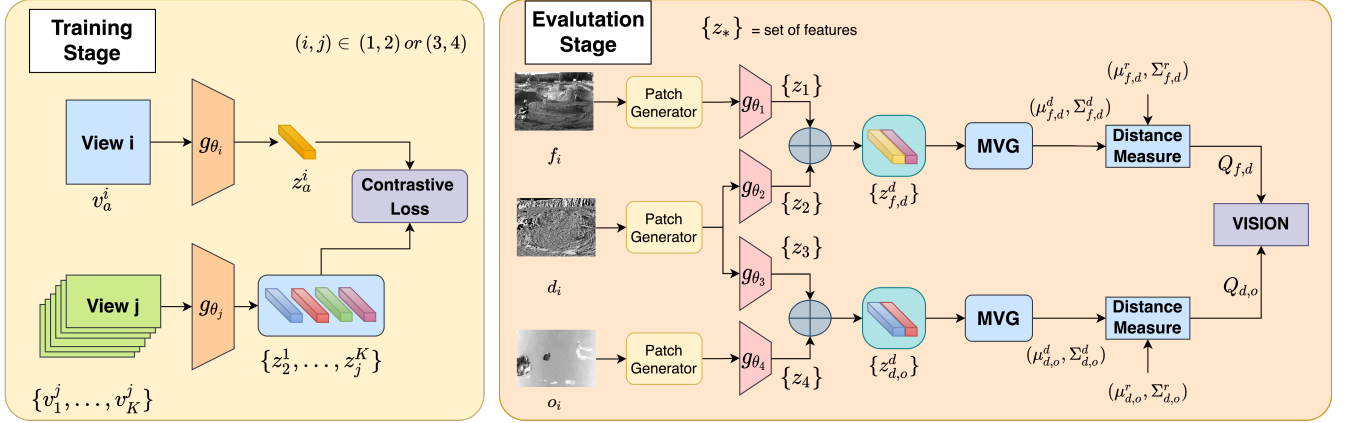


Figure 1: Illustration of VISION algorithm. Training Stage: For every sample of the first view, one positive and $(K - 1)$ negative sample are chosen from the second view. Encoders g_{θ_i} and g_{θ_j} are optimized with contrastive loss using the positive and negative set of features. **Evaluation Stage:** Feature embeddings are extracted from frames, frame differences, and optical flow using the learned encoder in the training stage. Quality is predicted by measuring the distance in Equation (9) using the MVG model parameters of pristine and test video patches.

video that differ only in quality. We then discuss our multiview contrastive learning framework on this dataset.

3.2 Data Generation

Two broad sets of videos are employed for learning our quality representations. We first consider databases of videos where pristine videos are corrupted with synthetic distortions such as compression, and transmission errors. We employ existing synthetically distorted datasets like the LIVE Mobile [28], LIVE VQA [32], EPFL-PoliMI [7], CSIQ VQD [39], and ECVQ-EVVQ [30] which contain pristine videos corrupted with rich synthetic distortions like MPEG-2, MPEG-4, H.264, noise, gaussian blur, IP loss, wavelet compression based snow codec, wireless transmission loss, and so on. Although our focus is on authentic distortions, we still believe that the synthetic distortions help learn good quality representations. Since our study mainly focuses on quality prediction for authentically distorted videos, we also generate different synthetically distorted versions of a camera captured video.

We hypothesize that by creating variations of the authentically distorted video in terms of further quality degradations, and then learning features that contrast these videos, we can learn quality representations of the underlying authentically distorted video. In particular, we randomly sample a camera captured video from an authentically distorted video database namely LIVE Large-Scale Social Video Quality [45]. We generate augmentations of this video with MPEG-2, and H.264 compression at different quality levels, and downsampling and upsampling at different scales. In addition, we use frame interpolation techniques [36] to create augmentations by first reducing the frame rate of the video and then generating the video at the original frame rate through interpolation. We create a combined database of synthetically distorted videos with compression and transmission errors, and also a database of authentically distorted videos with its various distorted versions.

3.3 Multiview Contrastive Feature Learning

As we remarked earlier, we choose two pairs of views, one between frame differences and frames and another between frame differences and optical flow to learn our quality representations. We note that there does not appear to exist any shared information between frames and optical flow and do not consider that pair. Thus, we constitute a two stream approach to learn quality aware representation. In the first stream, frame f and corresponding frame difference d with respect to its neighbour is used to form the two views of a multiview contrastive framework. In the other stream, frame difference d and optical flow o are considered as two views of a CMC [35] framework. Our quality representation learning framework is temporally localized.

We first describe the multiview representation learning with frames and frame differences as the pair of views. The multiview contrastive learning framework requires a pair of congruent and several incongruent pairs to learn quality representations. While the congruent pairs are chosen as the feature representations from frames and frame differences of the same video, the incongruent pairs are chosen as frames and frame differences from different videos having the same content but different distortions. This is a key aspect of our model where the positives and negatives are always drawn from the same scene. In a mini-batch, we sample S scenes from the training database, and for each scene, we have a set of K videos with different distortions. Let $\{V_1^s, V_2^s, \dots, V_K^s\}$ be a set of K videos with different distortion types and levels but having the same content, where, $s = \{1, 2, \dots, S\}$. Let, f_j and d_j denote the frame and frame difference (with frame f_j as one of the frames) of the video V_j^s respectively at a certain time instance where $j \in \{1, 2, \dots, K\}$. We describe the training method for V_j^s below and the same can be extended to other videos. Let frame f_j be chosen as the anchor view. The positive or congruent pair of views corresponding to this anchor is $\{f_j, d_j\}$, while the negative pairs

are given by $\{f_j, d_k\}_{k \neq j}$ by enumerating over the frame differences of all incongruent videos.

We deploy two CNNs, $g_{\theta_1}(\cdot)$ and $g_{\theta_2}(\cdot)$ with the same architecture but different parameters θ_1 and θ_2 to learn feature representations of the frames and frame differences, respectively. A discriminative function $h_{\theta}(\cdot)$ is trained to give high similarity between $g_{\theta_1}(f_j)$, and $g_{\theta_2}(d_j)$ and low similarity between $g_{\theta_1}(f_j)$ and $g_{\theta_2}(d_k)$, where $k \neq j$ and $k \in \{1, 2, \dots, K\}$. Let, the feature representation for frame f and frame difference d be given by,

$$z_1 = g_{\theta_1}(f), z_2 = g_{\theta_2}(d).$$

The cosine similarity between the embedding z_1 and z_2 is given by:

$$h_{\theta}(f, d) = \exp\left(\frac{z_1^T z_2}{\|z_1\| \|z_2\|} \cdot \frac{1}{\tau}\right), \quad (1)$$

where τ is the dynamic range adjuster. Therefore, the contrastive loss can be written as:

$$l(f_j, d_j) = -\log \frac{h_{\theta}(f_j, d_j)}{\sum_{k=1}^K h_{\theta}(f_j, d_k)} \quad (2)$$

Enumerating the anchor view over all the videos in $\{V_1^s, V_2^s, \dots, V_K^s\}_{s=1}^S$, we have the overall loss term with frames taken as anchor view as,

$$l(f, d) = \frac{1}{SK} \sum_{s=1}^S \sum_{j=1}^K l(f_j, d_j) \quad (3)$$

Similarly, taking the anchor view as frame differences d_j , we have, $l(d, f) = \frac{1}{SK} \sum_{s=1}^S \sum_{j=1}^K l(d_j, f_j)$. We obtain the overall two view loss function as

$$\mathcal{L}_{f,d} = l(f, d) + l(d, f). \quad (4)$$

For the second stream in our contrastive learning setup, we choose frame difference d and optical flow map o at a certain time instance of a video as the two views. Similar to the above approach, we choose two CNN based encoder networks $g_{\theta_3}(\cdot)$ and $g_{\theta_4}(\cdot)$ to learn the representations of d and o , respectively. Let z_3 , and z_4 represent the feature output of $g_{\theta_3}(\cdot)$ and $g_{\theta_4}(\cdot)$. Then the similarity between these embeddings is given as,

$$h_{\theta}(d, o) = \exp\left(\frac{z_3^T z_4}{\|z_3\| \|z_4\|} \cdot \frac{1}{\tau}\right). \quad (5)$$

Taking frame differences as the anchor view and enumerating them over optical flow for positive and negative pairs like Equation (3), we have the contrastive loss as

$$l(d_j, o_j) = -\log \frac{h_{\theta}(d_j, o_j)}{\sum_{k=1}^K h_{\theta}(d_j, o_k)} \quad (6)$$

$$l(d, o) = \frac{1}{SK} \sum_{s=1}^S \sum_{j=1}^K l(d_j, o_j). \quad (7)$$

The overall objective function with either of the views taken as anchor view is given as,

$$\mathcal{L}_{d,o} = l(d, o) + l(o, d). \quad (8)$$

We train all the four networks in the two streams with the respective loss functions in Equation (4), and (8). Ideally, either of

the CNNs from a given pair of views may be used to extract features during the prediction stage. However, we note that using both CNNs may be beneficial in overcoming any residual errors in perfectly contrasting the positive and negative pairs. Thus, during the prediction stage, we obtain the embedding for the first stream with frame and frame difference as views as $z_{f,d} = (z_1 + z_2)/2$, and for the second stream as $z_{d,o} = (z_3 + z_4)/2$. In Section 4.4, we provide a detailed study of using the average of the feature representation over using individual features.

3.4 Distance Measure and Quality Prediction

The goal of our work is to estimate the quality of videos in a completely blind setup. We choose a distance based approach similar to NIQE [27] to predict quality given the feature representation of a distorted video and a corpus of pristine videos. Since we extract two sets of features $z_{f,d}$ and $z_{d,o}$ from the video, we compute the distance with respect to each of these features and combine them. We first describe the distance computation using $z_{f,d}$. A similar approach is adopted for $z_{d,o}$.

We generate patches of size $R \times R$ from the frames of pristine videos corresponding to each of the views, i.e. frame and frame difference. Similar to NIQE [27], we select those patches in each frame that have a sharpness greater than τ_s times the sharpness of the sharpest patch of the frame. The frame differences are also drawn at the corresponding same locations as the sharp patches. A multivariate Gaussian (MVG) model with parameters $(\mu_r^{fd}, \Sigma_r^{fd})$ is learnt on the feature representation $z_{f,d}$ of these sharp patches drawn from the set of pristine videos. We now predict the quality of the i^{th} frame in a distorted video as follows. We extract $R \times R$ patches with no overlap from i^{th} frame and obtain the feature embeddings for all the patches similar to NIQE [27]. An MVG model with parameters $(\mu_d^{fd}, \Sigma_d^{fd})$ is learnt on these distorted patches. The quality estimate of the i^{th} frame in a distorted video is given as,

$$q_{f,d}^i = \sqrt{(\mu_r^{fd} - \mu_d^{fd})^T \left(\frac{\Sigma_r^{fd} + \Sigma_d^{fd}}{2} \right)^{-1} (\mu_r^{fd} - \mu_d^{fd})}. \quad (9)$$

We compute the above quality index for frames sampled at 1 frame per second in the distorted video. For a video of duration N seconds, $i \in \{1, 2, \dots, N\}$, the video level quality is estimated by average pooling the frame level quality predictions as,

$$Q_{f,d} = \frac{1}{N} \sum_{i=1}^N q_{f,d}^i. \quad (10)$$

Similarly, we extract the embedding $z_{d,o}$ using the second pair of networks corresponding to frame differences and optical flow maps from the same locations of sharp patches determined above. Let $q_{d,o}^i$ denote the quality of a video at the i^{th} time instance using $z_{d,o}$. The video level quality from the second stream is given by,

$$Q_{d,o} = \frac{1}{N} \sum_{i=1}^N q_{d,o}^i. \quad (11)$$

VISION. The overall video quality for the test video is given as the product of the predicted quality using the two streams as,

$$VISION = Q_{f,d} * Q_{d,o}. \quad (12)$$

Our overall quality prediction using the product of individual components is similar to that of ST-RRED [33], and NR-STED [25].

4 EXPERIMENTS AND RESULTS

4.1 Experimental Settings

Training Data. We train our model in a self-supervised fashion on both synthetically and authentically distorted datasets.

Synthetic Database. We use a combination of 5 databases comprising 850 distorted videos produced from 60 pristine videos at different resolutions and frame rates. The details of distortions present in the synthetic databases are as follows:

LIVE VQA [32]: The videos suffer from compression artifacts due to MPEG-2 and H.264 as well as transmission distortions obtained by sending H.264 compressed videos through error-prone IP and wireless networks.

LIVE Mobile [28]: A set of videos comprising compression, wireless packet-loss, rate-adaptation, and temporal dynamics are included in this database.

EPFL-PoliMI [7]: This database is composed of videos at two different resolutions encoded with H.264/AVC and passed through an error prone channel.

CSIQ VQD [39]: Here, we have videos comprising H.264, HEVC, Motion JPEG, wavelet based compression using Snow codec and H.264 were subjected to wireless transmission channel.

ECVQ and EVVQ [30]: The artifacts include H.264 and MPEG-4 visual compression.

Authentic Database. Since our objective in this work is completely blind estimation of video quality of user generated content, which is typically authentically distorted, we use videos from the LIVE-FB Large-Scale Social Video Quality (LSVQ) [45] database for our training. Since our model works with frame level data, we randomly sample 200 videos from the 39K videos in the LSVQ video database. LSVQ contains only a single distorted video for each content, so we augment the video with further distortions for contrastive learning. For distortion augmentation, we use MPEG-2, H.264, video downsampling followed by upsampling, and video frame interpolation. For every distortion type, we generate videos at three different distortion levels. We use ffmpeg [36] to generate MPEG-2 distorted videos with *qscale* index varying between 1 to 20. Similarly, to corrupt the authentic videos with H.264, we use ffmpeg [36] at a *crf* value between 10 to 50. Sampling is done by downscaling the video at a rate of 2, 4, and 8 and upscaling them back to original resolution. For frame interpolation, we read each video at 0.25, 0.33, and 0.5 times the original frame rate. Then we fill the frame by interpolation using *minterpolate* filter in ffmpeg [36] to get the distorted video at the original frame rate.

Details of Views The input data to the encoder $g_{\theta_1}(\cdot)$ consists of grayscale frames, while the input to the encoders $g_{\theta_2}(\cdot)$, and $g_{\theta_3}(\cdot)$ consists of the difference of grayscale frames. The input to $g_{\theta_4}(\cdot)$ is a two channel optical flow map, each channel representing horizontal and vertical displacement.

Training Details The encoder chosen to extract features from frames, frame differences, and optical flow viz. $g_{\theta_1}(\cdot)$, $g_{\theta_2}(\cdot)$, $g_{\theta_3}(\cdot)$, and $g_{\theta_4}(\cdot)$ have the same architecture and each of them outputs a 256 dimensional feature vector. The encoder architecture is given in the supplementary material. Each of the encoders takes a frame level

input at a certain time instance. The optical flow map is estimated using TV-L1 algorithm [48]. The training was done using a batch size $S = 8$ with Adam optimizer [15] at a learning rate of $1e - 4$ for 5000 iterations. For each batch of input, we have 1 positive and 10 negative pairs. Due to the computational complexity, we centre crop the views, taking the input at a resolution of 224×224 . The dynamic range adjuster is chosen to be 0.1 similar to [6, 23].

Evaluation Databases Since our goal is to design a completely blind VQA algorithm for user generated content, we omit the LSVQ [45] dataset for quality prediction of authentically distorted videos since we trained on that dataset. We conduct experiments on four user generated datasets as follows:

KoNViD-1K [13]: This dataset contains 1200 videos filtered from the YFCC100m database consisting of 793436 sequences. The videos in this database contain a wide variety of content, distortion types, and subjective quality variations. The videos are of 720×540 resolution, corresponding to a frame rate of 24, 25, or 30 frames per second, and 8 seconds in duration.

LIVE Video Quality Challenge (LVQC) Database[9]: The LIVE VQC database consists of 585 videos of unique content captured from 101 different devices leading to a widespread of complex authentic distortions. LVQC has 10 second long videos available at 18 different spatial resolutions ranging between 1980×1080 to 320×240 across landscape and portrait modes.

LIVE Qualcomm Database (LQCOMM)[10]: This database consists of 208 videos accounting for distortions generated during the camera capture process using eight mobile devices. The videos are of spatial resolution 1920×1080 , 15 seconds long when played at 30 fps.

YouTube-UGC [40]: This database contains 1380 user generated videos at resolutions varying between 360p to 4k. The videos are 20 seconds long in duration. This database contains videos belonging to 15 categories (e.g. gaming, sports, and music videos).

Evaluation Details. Similar to NIQE [27], we choose patches of size 96×96 ($R = 96$). To generate pristine patches, we choose reference videos from [7, 28, 30, 32, 39]. We choose a sharpness threshold $\tau_s = 0.85$ for generating the pristine set of patches for all our experiments and also for the other benchmarking algorithms. To reduce the computational capacity in estimating the optical flow map using the TV-L1 algorithm at test time, we estimate the optical flow at 1/8th of the spatial resolution of the video and upsample the flow to the original resolution with appropriate scaling. Since the computational time required to predict video quality is large if the resolution is greater than 720p as shown in prior work [16, 37, 38], we see that quality prediction at 1 frame per second followed by averaging pooling of the scores gives a similar performance to taking all the frames for prediction. Thus, we compute the quality of videos at 1 frame per second for VISION.

We evaluate the performance of the completely blind VQA methods using the conventional measures such as Spearman’s rank order correlation coefficient (SROCC), and linear correlation coefficient (LCC) between the predicted quality scores and the ground truth quality scores. Since no training is required on the test dataset, we do not need any training or validation splits of the test dataset. All the videos in the test datasets were used for evaluation. We pass our predicted quality through a four parameter monotonic

Table 1: Performance evaluation of VISION against other completely blind benchmarking algorithms on four user generated content datasets. Methods marked with (*) are modified to predict quality in a blind fashion using the distance metric in Section 9. The *Emphasised*, and **Boldfaced entries indicate the best and second-best performance in each database.**

Methods	KoNVid-1K		LIVE VQC		Youtube-UGC		LIVE Qualcomm	
	SROCC	PLCC	SROCC	PLCC	SROCC	PLCC	SROCC	PLCC
VIIDEO	0.013	-0.015	0.029	0.137	0.130	146	-0.141	0.098
VCORNIA* (1fr/sec)	0.112	0.132	0.166	0.133	0.461	0.455	0.186	0.267
NIQE (1 fr/sec)	0.542	0.544	0.563	0.610	0.236	0.105	0.467	0.504
ResNet50* (1 fr/sec)	0.273	0.288	0.240	0.275	0.466	0.465	0.313	0.362
STEM	0.629	0.629	0.656	0.670	0.284	0.318	0.483	0.537
HEKE* (1 fr/sec)	0.487	0.508	0.444	0.525	0.462	0.501	0.236	0.327
$Q_{f,d}$ (1 fr/sec)	0.545	0.558	0.549	0.592	0.492	0.501	0.441	0.479
$Q_{d,o}$ (1 fr/sec)	0.496	0.497	0.647	0.664	0.466	0.479	0.502	0.534
VISION (1 fr/sec)	0.598	0.597	0.676	0.701	0.503	0.510	0.547	0.576

logistic function given by, $Q_{trans} = \beta_2 + \frac{\beta_1 - \beta_2}{1 + \exp(-((Q - \beta_2)/\beta_4))}$ as in [14, 25, 29], where Q is the predicted quality before computing the performance measures.

4.2 Benchmarking Algorithms

We compare VISION with other completely blind VQA algorithms like VIIDEO [26] and STEM [14]. Since our quality estimation using the distance metric in Equation (9) is applied at a frame level, we extend our benchmarking analysis to completely blind IQA methods like NIQE [27] to compute frame level quality. We estimate the video quality using NIQE [27] as in Equation (10) at 1 frame/sec. In addition to the above, we present several interesting benchmarks where we evaluate methods that learn features without human supervision and use them to evaluate quality by comparing with a corpus of pristine videos. We note that VCORNIA [42] learns frame level quality features without supervision. So, we apply our distance measurement criteria to the set of pristine and test video features generated using VCORNIA [42]. HEKE [22] learns quality aware features using synthetically distorted databases by regressing against pseudo-quality scores generated by various full reference metrics. Since the training of HEKE [22] is blind with respect to the authentic data, we also extract pristine and test video features using this model. Thereafter, we apply our distance measure to estimate the quality at a 1 frame per second and average pool the scores like our method. Pretrained ResNet50 [12] features extracted at frame level perform well for NR VQA [37]. In our experiment, we extract ResNet50 features at patch level like VISION from the pristine corpus and the test video to measure quality of the test video using our distance measure. We term VCORNIA [42], HEKE [22], and pretrained ResNet50 [12] formulated in a blind setup as VCORNIA*, HEKE*, and ResNet50* respectively.

4.3 Performance Analysis

From Table 1, we see that VIIDEO [26] performs very poorly on the authentically distorted videos across all databases. While HEKE* (1 fr/sec) and NIQE (1fr/sec) [27] perform well on KoNVid-1K and LIVE VQC, NIQE (1fr/sec) performs poorly on Youtube-UGC [40], and HEKE* (1fr/sec) performs poorly on LIVE Qualcomm [10]. Though VCORNIA* (1fr/sec), and pretrained ResNet50* (1 fr/sec)

Table 2: Ablation analysis on different choice of views measured through SROCC. In the first column "View" denotes the views considered for learning in the multiview setup.

Views	Feature	KoNVid-1K	LVQC	LQCOMM
Frame -	z_1	0.498	0.485	0.434
Frame	z_2	0.507	0.473	0.382
Diff.	$z_{f,d}$	0.545	0.549	0.441
Optical Fl.	z_3	0.371	0.568	0.480
- Frame	z_4	0.472	0.575	0.372
Diff.	$z_{d,o}$	0.496	0.647	0.502
VISION	$(z_{f,d} + z_{d,o})/2$	0.598	0.676	0.547
CMC	-	0.401	0.492	0.443

performs well on the Youtube-UGC [40] dataset, its performance on other datasets is poor. We see that VISION gives comparable performance against the current state-of-the-art completely blind algorithm STEM [14] on KoNVid-1K [13], and LIVE VQC [9] databases. On YouTube-UGC [40], and LIVE Qualcomm [10], VISION outperforms STEM [14]. Overall, VISION achieves a good performance consistently across all datasets.

4.4 Ablation Studies

Impact of Averaging the features from Multiple Views. In Table 2, we analyze whether the average of the features extracted from $g_{\theta_1}(\cdot)$, and $g_{\theta_2}(\cdot)$ gives better performance than the individual features extracted. A similar experiment is also conducted for $g_{\theta_3}(\cdot)$, and $g_{\theta_4}(\cdot)$. Averaging tends to often give a better performance than the individual performance as seen in Table 2. Thus we choose the averaged features from each stream to get $Q_{f,d}$, and $Q_{d,o}$. We also learned representations based on the views in CMC [35] and presented the results in Table 2. We infer that the views with our approach perform significantly better in learning representations for the blind VQA task.

Performance on Synthetically Distorted Databases. We evaluate the performance of VISION on synthetically distorted databases such as LIVE VQA [32], LIVE Mobile [28], CSIQ VQD [39], and EPFL-PoliMI [7]. Since, VISION was trained using all these synthetic databases, for this analysis, we leave out the test database during

Table 3: SROCC performance analysis on four synthetically distorted video databases. VCORNIA*, NIQE [27], ResNet50*, HEKE*, and VISION are applied at 1 frame/second for quality prediction.

Methods	CSIQ VQD	EPFL-Po	LIVE VQA	LIVE Mobile
VIIDEO	0.02	0.205	0.624	0.216
VCORNIA*	-0.02	0.05	0.15	0.289
NIQE	0.440	0.185	0.174	0.427
ResNet50*	0.164	0.005	0.070	0.176
STEM	0.380	0.199	0.205	0.361
HEKE*	0.291	0.347	-0.139	0.370
VISION	0.463	0.225	0.273	0.433

Table 4: SROCC performance analysis of VISION on three UGC datasets when trained on synthetically distorted, authentically distorted and combined database.

Data Type	KoNVid-1K	LIVE VQC	LIVE QCOMM
Synthetic	0.592	0.630	0.524
Authentic	0.567	0.667	0.542
Combined	0.598	0.676	0.547

the training and the computation of the model for the pristine corpus. Since there are very few unique scenes in the synthetic datasets and many videos are distorted versions of these, we find that all the methods are sensitive to the content causing a drop in performance. Nevertheless, VISION gives a more stable performance than other benchmarking algorithms on most of the databases.

Learning on Synthetic vs Authentic Datasets. VISION is learnt on a mix of synthetically distorted videos [7, 28, 30, 32, 39], and authentically distorted videos generated from LSVQ [45]. We study the impact of learning with synthetic distortions and authentic distortions individually. In this experiment, VISION is trained with either synthetically distorted or authentically distorted videos. We show a comparative study of our model on three UGC test databases [9, 10, 13] in Table 4. Though, learning on synthetic or authentic distortions alone gives a similar performance, the combined learning on mixed data gives superior performance on most databases. Combining synthetically distorted videos with the authentically distorted videos benefits the performance since the synthetically distorted databases contain richer sets of distortions generated by various study groups as mentioned in Section 4.1.

Impact of Distortion Augmentation. To generate distorted views of the authentically distorted videos, we corrupt each UGC video with synthetic distortions such as MPEG-2, H.264, sampling, and interpolations as described in Section 4.1. We study the impact of synthetic augmentations when our model is trained on LSVQ [45] videos corrupted with synthetic distortions, without one of the above four distortion types. In Figure 2, a comparative study on the performance of VISION is given when the encoders are trained with videos corrupted with all the above four distortion types vs when trained on videos corrupted with all but one distortion type. Overall, we see that no single augmentation is crucial and we get roughly similar performances even if we remove one of the augmentations.

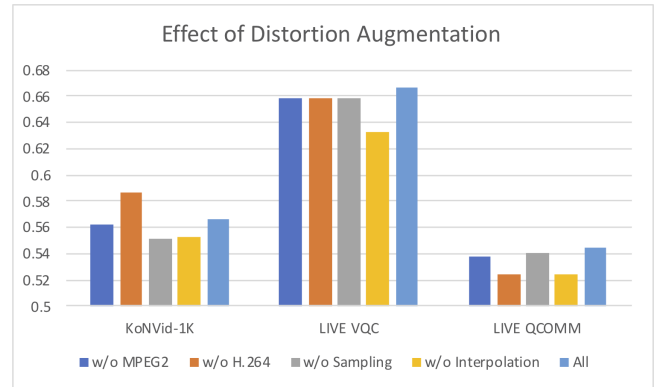


Figure 2: Effect of distortion augmentation on VISION performance. First four bars denote the performance measured by SROCC without one of the four distortion types mentioned in Section 4.1, while the last bar denotes the performance when all distortions are considered during training.

Table 5: SROCC performance comparison with other self-supervised video representation learning methods on three UGC datasets. Entry marked '-' denotes data is unavailable.

Data Type	KoNVid-1K	LIVE VQC	LIVE QCOMM
DPC	0.426	0.485	0.308
VCOP	0.452	0.494	-
PRP	0.349	0.415	0.281
IIC	0.527	0.534	0.361
CSPT	0.702	0.623	-
VISION	0.724	0.698	0.623

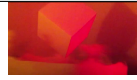





Self-Supervision based Video Representation Learning. We also provide a comparative study of various self-supervision based video representation learning methods on three UGC databases in Table 5. DPC [11], VCOP [41], PRP [44], and IIC [34] learn spatio-temporal video features by learning an action recognition task. CSPT [5] is a contrastive learning based VQA method modelled on learning quality based features by pretraining on a future video frame prediction task. We train a linear regressor (ridge regression) on top of the video level features extracted from all the above mentioned self-supervised algorithms and VISION (average pooled across frames) except CSPT [5], and VCOP [41]. We conduct this test on 100 splits by dividing each of the 3 UGC datasets in the ratio of 80 : 20 for the train-test splits. The numbers reported for CSPT [5], and VCOP [41] are taken from literature. We see that VISION outperforms these self-supervised algorithms in a linear evaluation testing protocol. Thus, our multiview contrastive learning based approach using sets of distorted video samples with similar content learns good features for quality prediction.

Linear Evaluation of VQA Methods. We benchmark the quality aware feature representation following the linear evaluation protocol for SOTA feature based VQA methods. We divide three different authentically distorted datasets, each in the ratio of 80 : 20 for the train-test splits and train a linear regressor for 100 splits. In Table 6, we see VISION gives a competitive performance with other

Table 6: SROCC performance comparison of VISION against other state-of-the-art VQA methods under linear evaluation protocol. Entry marked ‘-’ denotes data is unavailable.

Data Type	KoNVid-1K	LIVE VQC	LIVE QCOMM
ResNet50	0.670	0.653	0.749
TLVQM	0.755	0.758	0.727
VIDEVAL	0.738	0.731	0.647
VCORNIA	0.475	0.526	0.462
HEKE	0.678	0.648	0.609
CSPT	0.702	0.623	-
VISION	0.724	0.698	0.623

Table 7: Analysis of complementarity between frame-frame difference ($z_{f,d}$), and frame difference-optical flow-based features ($z_{d,o}$). Three consecutive frames from two videos in KoNVid-1K [13] dataset are given in each row.

frame 1	frame 2	frame 3	$E_{f,d}$	$E_{d,o}$
			0.52	4.02
			1.82	0.61

benchmarking methods on KoNVid-1K [13], and LIVE VQC [9], while on LIVE Qualcomm [10] it is slightly lower than some other methods. We note that although the pretrained ResNet50 model achieves better performance than VISION in a linear supervised setting, in a completely blind setup, VISION features completely outperform the pretrained ResNet50 features. On the other hand, TLVQM [16] and VIDEVAL [37] provide video level feature representations. We also experimented with video level features from numerous video clips using the features from the TLVQM [16] model to evaluate it in a completely blind setting. However, the quality of the videos predicted using distance measure in Equation (9) from these features of the clips tends to perform poorly in terms of correlation with human scores.

Feature Complementarity Analysis. We study the complementarity among the features extracted from each stream by predicting the absolute error in quality prediction. For a given test video with human opinion score Q , let the error in quality estimation with features ($z_{f,d}$), and ($z_{d,o}$) be given by $E_{f,d} = |Q - Q_{f,d}|$, and $E_{d,o} = |Q - Q_{d,o}|$, respectively. In Table 7, we show consecutive frames from two videos from KoNVid-1K [13] dataset mainly corrupted with illumination change and stabilization error respectively. We see that $E_{f,d}$ is higher than $E_{d,o}$ for the video mainly corrupted by illumination variation, while the trend is opposite when the distortion is due to motion only. We infer that when spatial distortion is predominant, using ($z_{f,d}$) is more useful due to the presence of frame level spatial information. While for the other video where temporal shakiness is predominant, ($z_{d,o}$) is more useful as optical flow information helps in understanding the distortion due to large motion. Thus, a combination of both sets of features for predicting quality will provide a better estimate for all types of distortion.




Figure 3: Runtime Comparison

4.5 Runtime Analysis

We present an analysis of the computational complexity of different methods through a runtime analysis. The runtime analysis is done on a Ubuntu 18.04.4 LTS system with an Intel® Core™ i7-8700 CPU @ 3.20GHz × 12. We also show the computational time for CNN-based methods using a 11GB GeForce RTX 2080 Ti graphics card. Since learning with 1 frame per second is an important aspect in our algorithm, we show runtime for VISION, HEKE*, VCORNIA*, NIQE [27], and pretrained ResNet50* at 1 frame/second. The average computational time to predict quality for 10 videos from the LIVE Qualcomm [10] dataset consisting of 450 frames at 1920×1080 resolution is measured. In Figure 3, we see that on a CPU device VISION takes a similar time as that of other benchmarking methods. On a GPU platform, the CNN based methods are 10–15 times faster than that on CPU device. As the time taken to compute the optical flow is around 6 seconds, VISION’s runtime is slightly higher than other CNN based methods on a GPU.

5 CONCLUSION

We designed a framework for completely blind NR VQA of user generated videos using a multiview contrastive setup. We learn a quality aware feature representation by leveraging the common quality information in multiple views of a video. A combination of quality aware features from frame and frame difference, as well as frame difference and optical flow is used in a blind fashion to predict video quality. We show that these features can be used to compute an effective distance measure between the test video and a corpus of pristine videos to predict video quality. VISION gives a stable and consistent performance across all UGC datasets compared to other completely blind algorithms.

REFERENCES

- [1] S. Ahn and S. Lee. 2018. Deep Blind Video Quality Assessment Based on Temporal Human Perception. In *2018 25th IEEE International Conference on Image Processing (ICIP)*. 619–623. <https://doi.org/10.1109/ICIP.2018.8451450>
- [2] Jorge E. Caviedes and Franco Oberti. 2003. No-reference quality metric for degraded and enhanced video. In *Visual Communications and Image Processing 2003*, Touradj Ebrahimi and Thomas Sikora (Eds.), Vol. 5150. International Society for Optics and Photonics, SPIE, 621 – 632. <https://doi.org/10.1117/12.510112>
- [3] Pengfei Chen, Leida Li, Lei Ma, Jinjian Wu, and Guangming Shi. 2020. *RIRNet: Recurrent-In-Recurrent Network for Video Quality Assessment*. Association for Computing Machinery, New York, NY, USA, 834–842. <https://doi.org/10.1145/3394171.3413717>
- [4] Pengfei Chen, Leida Li, Jinjian Wu, Weisheng Dong, and Guangming Shi. 2021. Unsupervised Curriculum Domain Adaptation for No-Reference Video Quality Assessment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 5178–5187.
- [5] Pengfei Chen, Leida Li, Jinjian Wu, Weisheng Dong, and Guangming Shi. 2022. Contrastive Self-Supervised Pre-Training for Video Quality Assessment. *IEEE Transactions on Image Processing* 31 (2022), 458–471. <https://doi.org/10.1109/TIP.2021.3130536>
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daum^Å III and Aarti Singh (Eds.). PMLR, 1597–1607.
- [7] F. De Simone, M. Tagliasacchi, M. Naccari, S. Tubaro, and T. Ebrahimi. 2010. A H.264/AVC video database for the evaluation of quality metrics. In *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*. 2430–2433.
- [8] M.C.Q. Farias and S.K. Mitra. 2005. No-reference video quality metric based on artifact measurements. In *IEEE International Conference on Image Processing 2005*, Vol. 3. III–141. <https://doi.org/10.1109/ICIP.2005.1530348>
- [9] Deepti Ghadiyaram and Alan C. Bovik. 2016. Massive Online Crowdsourced Study of Subjective and Objective Picture Quality. *IEEE Transactions on Image Processing* 25, 1 (2016), 372–387. <https://doi.org/10.1109/TIP.2015.2500021>
- [10] Deepti Ghadiyaram, Janice Pan, Alan C. Bovik, Anush Krishna Moorthy, Prasanjit Panda, and Kai-Chieh Yang. 2018. In-Capture Mobile Video Distortions: A Study of Subjective Behavior and Objective Algorithms. *IEEE Transactions on Circuits and Systems for Video Technology* 28, 9 (2018), 2061–2077. <https://doi.org/10.1109/TCSVT.2017.2707479>
- [11] Tengda Han, Weidi Xie, and Andrew Zisserman. 2019. Video Representation Learning by Dense Predictive Coding. In *Workshop on Large Scale Holistic Video Understanding, ICCV*.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*.
- [13] Vlad Hosu, Franz Hahn, Mohsen Jenadeleh, Hanhe Lin, Hui Men, Tamás Szirányi, Shujun Li, and Dietmar Saupe. 2017. The Konstanz natural video database (KoNViD-1k). In *2017 Ninth International Conference on Quality of Multimedia Experience (QoMEX)*. IEEE, 1–6.
- [14] Parimala Kancharla and Sumohana S. Channappayya. 2022. Completely Blind Quality Assessment of User Generated Video Content. *IEEE Transactions on Image Processing* 31 (2022), 263–274. <https://doi.org/10.1109/TIP.2021.3130541>
- [15] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [16] J. Korhonen. 2019. Two-Level Approach for No-Reference Consumer Video Quality Assessment. *IEEE Transactions on Image Processing* 28, 12 (2019), 5923–5938.
- [17] Jari Korhonen, Yicheng Su, and Junyong You. 2020. Blind Natural Video Quality Prediction via Statistical Temporal Features and Deep Spatial Features. In *Int. Conf. Multimedia (ACM-MM)*.
- [18] Dingquan Li, Tingting Jiang, and Ming Jiang. 2019. Quality Assessment of In-the-Wild Videos. In *Int. Conf. Multimedia (ACM-MM)*.
- [19] Dingquan Li, Tingting Jiang, and Ming Jiang. 2021. Unified Quality Assessment of in-the-Wild Videos with Mixed Datasets Training. *International Journal of Computer Vision* 129, 4 (01 Apr 2021), 1238–1257. <https://doi.org/10.1007/s11263-020-01408-w>
- [20] X. Li, Q. Guo, and X. Lu. 2016. Spatiotemporal Statistics for Video Quality Assessment. *IEEE Transactions on Image Processing* 25, 7 (2016), 3329–3342.
- [21] Wentao Liu, Zhengfang Duanmu, and Zhou Wang. 2018. End-to-End Blind Quality Assessment of Compressed Videos Using Deep Neural Networks. In *Proceedings of the 26th ACM International Conference on Multimedia (Seoul, Republic of Korea) (MM '18)*. Association for Computing Machinery, New York, NY, USA, 546–554. <https://doi.org/10.1145/3240508.3240643>
- [22] Yongxu Liu, Jinjian Wu, Leida Li, Weisheng Dong, Jimpeng Zhang, and Guangming Shi. 2021. Spatiotemporal Representation Learning for Blind Video Quality Assessment. *IEEE Transactions on Circuits and Systems for Video Technology* (2021). <https://doi.org/10.1109/TCSVT.2021.3114509>
- [23] Pavan C. Madhusudana, Neil Birkbeck, Yilin Wang, Balu Adsumilli, and Alan C. Bovik. 2021. Image Quality Assessment using Contrastive Learning. *CoRR abs/2110.13266* (2021). [arXiv:2110.13266](https://arxiv.org/abs/2110.13266)
- [24] K. Manasa and S. S. Channappayya. 2016. An optical flow-based no-reference video quality assessment algorithm. In *2016 IEEE International Conference on Image Processing (ICIP)*. 2400–2404.
- [25] Shankhanil Mitra, Rajiv Soundararajan, and Sumohana S. Channappayya. 2021. Predicting Spatio-Temporal Entropic Differences for Robust No Reference Video Quality Assessment. *IEEE Signal Processing Letters* 28 (2021), 170–174. <https://doi.org/10.1109/LSP.2021.3049682>
- [26] Anish Mittal, Michele A. Saad, and Alan C. Bovik. 2016. A Completely Blind Video Integrity Oracle. *IEEE Transactions on Image Processing* 25, 1 (2016), 289–300. <https://doi.org/10.1109/TIP.2015.2502725>
- [27] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. 2013. Making a “Completely Blind” Image Quality Analyzer. *IEEE Signal Process. Lett.* 20, 3 (2013), 209–212.
- [28] A. K. Moorthy, L. K. Choi, A. C. Bovik, and G. de Veciana. 2012. Video Quality Assessment on Mobile Devices: Subjective, Behavioral and Objective Studies. *IEEE Journal of Selected Topics in Signal Processing* 6, 6 (2012), 652–671.
- [29] S. V. Reddy Dendi and S. S. Channappayya. 2020. No-Reference Video Quality Assessment Using Natural Spatiotemporal Scene Statistics. *IEEE Transactions on Image Processing* 29 (2020), 5612–5624.
- [30] Snjezana Rimac-Drije, Mario Vranješ, and Drago Žagar. 2010. Foveated Mean Squared Error—a Novel Video Quality Metric. *Multimedia Tools Appl.* 49, 3 (Sept. 2010), 425–445. <https://doi.org/10.1007/s11042-009-0442-1>
- [31] Michele A Saad, Alan C Bovik, and Christophe Charrier. 2014. Blind Prediction of Natural Video Quality. *IEEE Transactions on Image Processing* 23, 3 (2014), 1352–1365.
- [32] Kalpana Seshadrinathan, Rajiv Soundararajan, Alan C Bovik, and Lawrence K Cormack. 2010. Study of subjective and objective quality assessment of video. *IEEE Transactions on Image Processing* 19, 6 (2010), 1427–1441.
- [33] Rajiv Soundararajan and Alan C Bovik. 2013. Video Quality Assessment by Reduced Reference Spatio-Temporal Entropic Differencing. *IEEE Transactions on Circuits and Systems for Video Technology* 23, 4 (2013), 684–694.
- [34] Li Tao, Xueting Wang, and Toshihiko Yamasaki. 2020. *Self-Supervised Video Representation Learning Using Inter-Intra Contrastive Framework*. Association for Computing Machinery, New York, NY, USA, 2193–2201. <https://doi.org/10.1145/3394171.3413694>
- [35] Yonglong Tian, Dilip Krishnan, and Phillip Isola. 2020. Contrastive Multiview Coding. In *Computer Vision – ECCV 2020*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer International Publishing, Cham, 776–794.
- [36] Suramya Tomar. 2006. Converting Video Formats with FFMpeg. *Linux J.* 2006, 146 (jun 2006), 10.
- [37] Zhengzhong Tu, Yilin Wang, Neil Birkbeck, Balu Adsumilli, and Alan C. Bovik. 2020. UGC-VQA: Benchmarking Blind Video Quality Assessment for User Generated Content. [arXiv:2005.14354](https://arxiv.org/abs/2005.14354) [cs.CV]
- [38] Zhengzhong Tu, Xiangxu Yu, Yilin Wang, Neil Birkbeck, Balu Adsumilli, and Alan C. Bovik. 2021. RAPIQUE: Rapid and Accurate Video Quality Prediction of User Generated Content. *CoRR abs/2101.10955* (2021). [arXiv:2101.10955](https://arxiv.org/abs/2101.10955)
- [39] Phong V. Vu and Damon M. Chandler. 2014. ViS3: an algorithm for video quality assessment via analysis of spatial and spatiotemporal slices. *Journal of Electronic Imaging* 23, 1 (2014), 1 – 25. <https://doi.org/10.1117/1.JEI.23.1.013016>
- [40] Yilin Wang, Sasi Inguva, and Balu Adsumilli. 2019. YouTube UGC Dataset for Video Compression Research. In *2019 IEEE 21st International Workshop on Multimedia Signal Processing (MMSP)*. 1–5. <https://doi.org/10.1109/MMSP.2019.8901772>
- [41] Dejing Xu, Jun Xiao, Zhou Zhao, Jian Shao, Di Xie, and Yueting Zhuang. 2019. Self-supervised Spatiotemporal Learning via Video Clip Order Prediction. In *Computer Vision and Pattern Recognition (CVPR)*.
- [42] J. Xu, P. Ye, Y. Liu, and D. Doermann. 2014. No-reference video quality assessment via feature learning. In *2014 IEEE International Conference on Image Processing (ICIP)*. 491–495. <https://doi.org/10.1109/ICIP.2014.7025098>
- [43] Fuzheng Yang, Shuai Wan, Yilin Chang, and Hong Ren Wu. 2005. A novel objective no-reference metric for digital video quality assessment. *IEEE Signal Processing Letters* 12, 10 (2005), 685–688. <https://doi.org/10.1109/LSP.2005.855553>
- [44] Yuan Yao, Chang Liu, Dezhao Luo, Yu Zhou, and Qixiang Ye. 2020. Video Playback Rate Perception for Self-Supervised Spatio-Temporal Representation Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [45] Zhenqiang Ying, Maniratnam Mandal, Deepti Ghadiyaram, and Alan Bovik. 2021. Patch-VQ: ‘Patching Up’ the Video Quality Problem. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 14019–14029.
- [46] Zhenqiang Ying, Haoran Niu, Praful Gupta, Dhruv Mahajan, Deepti Ghadiyaram, and Alan Bovik. 2020. From Patches to Pictures (PaQ-2-PiQ): Mapping the Perceptual Space of Picture Quality. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

- [47] J. You and J. Korhonen. 2019. Deep Neural Networks for No-Reference Video Quality Assessment. In *2019 IEEE International Conference on Image Processing (ICIP)*. 2349–2353. <https://doi.org/10.1109/ICIP.2019.8803395>
- [48] C. Zach, T. Pock, and H. Bischof. 2007. A Duality Based Approach for Realtime TV-L1 Optical Flow. In *Pattern Recognition*, Fred A. Hamprecht, Christoph Schnörr, and Bernd Jähne (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 214–223.
- [49] Y. Zhang, X. Gao, L. He, W. Lu, and R. He. 2019. Blind Video Quality Assessment With Weakly Supervised Learning and Resampling Strategy. *IEEE Transactions on Circuits and Systems for Video Technology* 29, 8 (2019), 2244–2255.

A FEATURE ENCODER ARCHITECTURE

As mentioned in Section 3.3 in the main paper, we deploy a pair of convolutional neural networks (CNN) for each stream of our two-stream feature extraction module. To learn feature embeddings from frames, and frame-differences, we deploy g_{θ_1} , and g_{θ_2} . Similarly, to learn representations from frame-differences, and optical flow, we deploy g_{θ_3} , and g_{θ_4} . All the four CNN based encoders have the same architecture as shown in Figure 4. Each encoder includes four convolutional blocks consisting of two convolutional layer followed by maxpooling layer, and a batch normalization layer. To extract feature representations, a global average pooling layer is applied at the output of the last convolutional block. The convolutional layers’ kernels and biases are initialized with random normal initializer. We used a rectified linear unit (ReLU) activation after each convolutional layer.

B IMPACT OF FRAME SAMPLING

In this section we study the impact of varying the frame sampling rate. VISION estimates video quality by averaging the predicted frame level quality at 1 frame/second instead of averaging over the whole duration of the video. We conduct a study to check the variation in performance of VISION on KoNVid-1K [13] dataset due to the different frame sampling rate. We also provide the average

computational time required to estimate the quality of videos. This study is done on an Ubuntu 18.04.4 LTS system with a Intel® Core™ i7-8700 CPU @ 3.20GHz × 12 with 11GB GeForce RTX 2080 Ti graphics card. In Table 8, we see that the performance is nearly similar as the sampling rate increases from 1 frame per second. There is a considerable dip in performance when the sampling rate is reduced below 1 frame per second. Thus, predicting the video quality by average pooling the estimated frame level quality at 1 frame per second gives similar performance while the average computational time is 9 times faster on a GPU machine.

Table 8: Performance of VISION with different frame sampling rate on KoNVid-1K [13] database.

Sampling Rate	SROCC	Computational Time (in seconds)
1 frame/4 sec	0.5336	1.56
1 frame/2 sec	0.578	1.96
1 frame/sec	0.598	2.86
2 frame/ sec	0.599	4.81
All frames	0.601	25.85

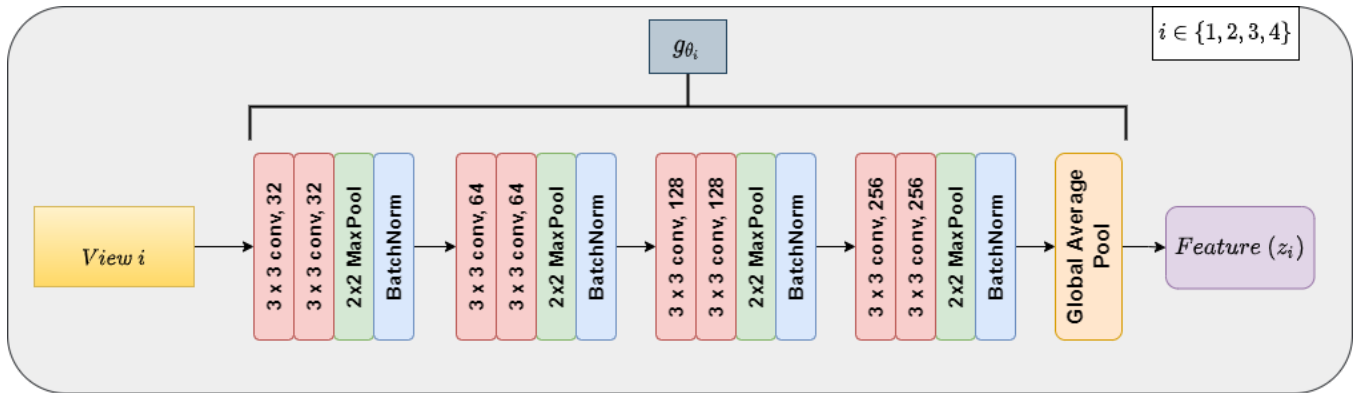


Figure 4: Architecture of our encoders g_{θ_1} , g_{θ_2} , g_{θ_3} , and g_{θ_4} .