

PTM-VQA: Efficient Video Quality Assessment Leveraging Diverse PreTrained Models from the Wild

Kun Yuan^{1†}, Hongbo Liu^{1,2†}, Mading Li^{1†}, Muyi Sun^{3,4},
Ming Sun^{1(✉)}, Jiachao Gong¹, Jinhua Hao¹, Chao Zhou¹, Yansong Tang^{2(✉)}
¹Kuaishou Technology, ²Tsinghua University, ³School of AI, BUPT, ⁴CASIA

{yuankun03, limading, sunming03}@kuaishou.com, liuhbleon@gmail.com, tang.yansong@sz.tsinghua.edu.cn

Abstract

Video quality assessment (VQA) is a challenging problem due to the numerous factors that can affect the perceptual quality of a video, e.g., content attractiveness, distortion type, motion pattern, and level. However, annotating the Mean opinion score (MOS) for videos is expensive and time-consuming, which limits the scale of VQA datasets, and poses a significant obstacle for deep learning-based methods. In this paper, we propose a VQA method named PTM-VQA, which leverages PreTrained Models to transfer knowledge from models pretrained on various pre-tasks, enabling benefits for VQA from different aspects.

Specifically, we extract features of videos from different pretrained models with frozen weights and integrate them to generate representation. Since these models possess various fields of knowledge and are often trained with labels irrelevant to quality, we propose an Intra-Consistency and Inter-Divisibility (ICID) loss to impose constraints on features extracted by multiple pretrained models. The intra-consistency constraint ensures that features extracted by different pretrained models are in the same unified quality-aware latent space, while the inter-divisibility introduces pseudo clusters based on the annotation of samples and tries to separate features of samples from different clusters. Furthermore, with a constantly growing number of pretrained models, it is crucial to determine which models to use and how to use them. To address this problem, we propose an efficient scheme to select suitable candidates. Models with better clustering performance on VQA datasets are chosen to be our candidates. Extensive experiments demonstrate the effectiveness of the proposed method.

1. Introduction

In recent years, social network platforms that focus on videos have gained immense popularity. According to

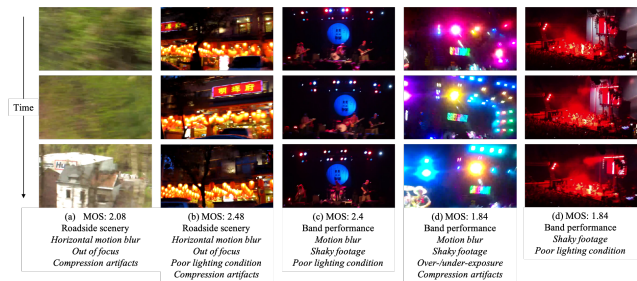


Figure 1. The KoNViD-1k dataset provides video frames that demonstrate a correlation between content/motion patterns and video quality. To identify potential reasons for poor perceptual video quality, we have highlighted specific factors in italics that correspond to the labeled MOS.

Cisco’s Visual Networking Index (VNI), global IP video traffic is predicted to account for 82% of all IP traffic by 2022, both in business and consumer sectors [2]. The significant increase in video content consumption poses significant challenges for video providers to deliver better services. Since the perceptual quality of videos has a significant impact on the Quality of Experience (QoE), identifying the quality of videos has become one of the most important issues [10, 11, 18, 26, 47]. Video quality assessment (VQA) aims to assess the perceptual quality of input videos automatically, imitating the subjective feedback of humans when viewing a video. It has been extensively studied in the context of assessing compression artifacts, transmission errors, and overall quality [27, 30, 34, 40, 45]. Data-driven deep learning-based methods have been drawing more and more attention compared to conventional methods based on hand-crafted features, as they possess better performance [7–9, 28, 31, 42, 57, 65, 68].

Compared with other high-level computer vision tasks, datasets for VQA are much smaller. One of the most popular datasets for human action classification Kinetics [5] has 650,000 clips, while the popular VQA dataset KoNViD-1k [22] has only 1,200 videos. One of the reasons is because VQA is a highly subjective task [58, 61]. To obtain an un-

[†] Equal contribution. [✉] Corresponding authors.

biased label, it is recommended by annotation guidelines [44] that the subjective quality of a single video should be measured in a laboratory test by calculating the arithmetic mean value of multiple subjective judgments, *i.e.*, Mean Opinion Score (MOS). Take *KoNViD-1k* as an example, it has 114 votes for each video on average. This significantly raises the cost of labeling and limits the size of the VQA dataset. Such a small amount of data limits the power of data-driven VQA methods. To deal with the problem, most existing VQA methods [7, 28, 65, 68] choose to finetune using weights pretrained on common larger datasets (*e.g.*, ImageNet [15]). However, existing works [29, 30, 57] show that the perceptual quality of a video is related to many factors, *e.g.*, content attractiveness, aesthetic quality, distortion type, motion pattern, and level. Only considering content-based pretrained models may not be sufficient for VQA. Thus, in this work, we focus on how to better utilize a large amount of available pretrained models to benefit VQA.

The present study initially observes a correlation between VQA tasks and other computer vision tasks. To illustrate, Fig. 1 displays several examples from the *KoNViD-1k* dataset. It is reasonable to assume that models pretrained on datasets for various pre-tasks have the ability to capture distinct characteristics regarding video quality. Consequently, we conduct a simple clustering experiment utilizing Large Margin Nearest Neighbor (LMNN) [59] to investigate the correlation between typical pretrained models and the VQA task. Based on the findings, we propose a practical approach, named *PTM-VQA (PreTrained Models-VQA)*, which leverages pretrained models as feature extractors and predicts the quality of input videos based on integrated features. As the parameters of pretrained models remain fixed, we can introduce more pretrained models without exhausting computational resources.

Moreover, we notice that labels in common datasets for pretraining are quite quality-irrelevant. For instance, **a clear photo of a puppy** with high quality and **a blurred photo of a puppy** may have the same object-wise label, whereas their quality-wise label may be significantly different. This will confuse the learning process for the VQA task. Therefore, we propose an *Intra-Consistency and Inter-Divisibility (ICID) loss*, which applies constraints on features extracted by multiple pretrained models from different samples. Specifically, model-wise intra-consistency requires features extracted by different pretrained models to be in the same unified quality-aware latent space. Meanwhile, sample-wise inter-divisibility introduces pseudo clusters based on the MOS of samples and aims to separate features of samples from different clusters.

Furthermore, as the number of pretrained models continues to grow (*e.g.*, PyTorch image models library (Timm) [60] supports over 700 pretrained models), finding models suitable for the VQA task through trial-and-error be-

comes unfeasible. Therefore, we propose to use the Davies-Bouldin Index (DBI) [14] to evaluate the clustering results and adopt it as the basis for model selection and weighting for feature integration. To summarize, the **main contributions** are specified below:

- We explore and confirm the association between pretrained models utilizing various pre-text tasks and their effectiveness in performing VQA. Moreover, we present a practical non-reference VQA method named PTM-VQA, which exploits cutting-edge pretrained models with diversity to benefit VQA effectively.
- To constrain features with diversity into a unified quality-aware space and eliminate the mismatch between objective and perceptual annotations, we propose an ICID loss. To avoid looking for a needle in a haystack, we propose an effective way to select candidate models based on DBI, which also determines the contributions of different pretrained models.
- PTM-VQA achieves SOTA performance with a rather small amount of learnable weights on three NR-VQA datasets, including *KoNViD-1k*, *Live-VQC*, and *YouTube-UGC*. Extensive ablations also prove the effectiveness of our method.

2. Related Work

VQA. Based on whether the pristine reference video is required, VQA methods can be classified as Full Reference (FR), Reduced Reference (RR), and No Reference (NR). Our work will be focused on the NR-VQA method. Traditional NR-VQA methods either measure video quality by rule-based metric [66], or predict MOS by an estimator (*e.g.*, Multi-Layer Perceptron, Support Vector Machine) based on hand-crafted features [13]. In recent years, deep learning-based VQA methods have been studied and surpassed traditional methods. STDAM [65] introduced a graph convolution to extract features and a bidirectional long short-term memory network to handle motion information. StarVQA [64] proposed encode spatiotemporal information of each patch on video frames and feed them into a Transformer. RAPIQUE [53] proposed to combine texture features and deep convolutional features. These works, however, neglected the correlation between VQA and other tasks and did not explore other datasets. BVQA [29] took one step further and proposed to transfer knowledge from IQA and action recognition to VQA. Our work further investigates the possibility of using more kinds of tasks.

Pretrained models. Pretrained models reveal the great potential in deep learning. In Natural Language Processing (NLP), BERT [16] and GPT-3 [4] demonstrated substantial gains on many NLP tasks and benchmarks by pre-training on a large corpus of text followed by finetuning on a specific task. The advent of ViT [17] had migrated this

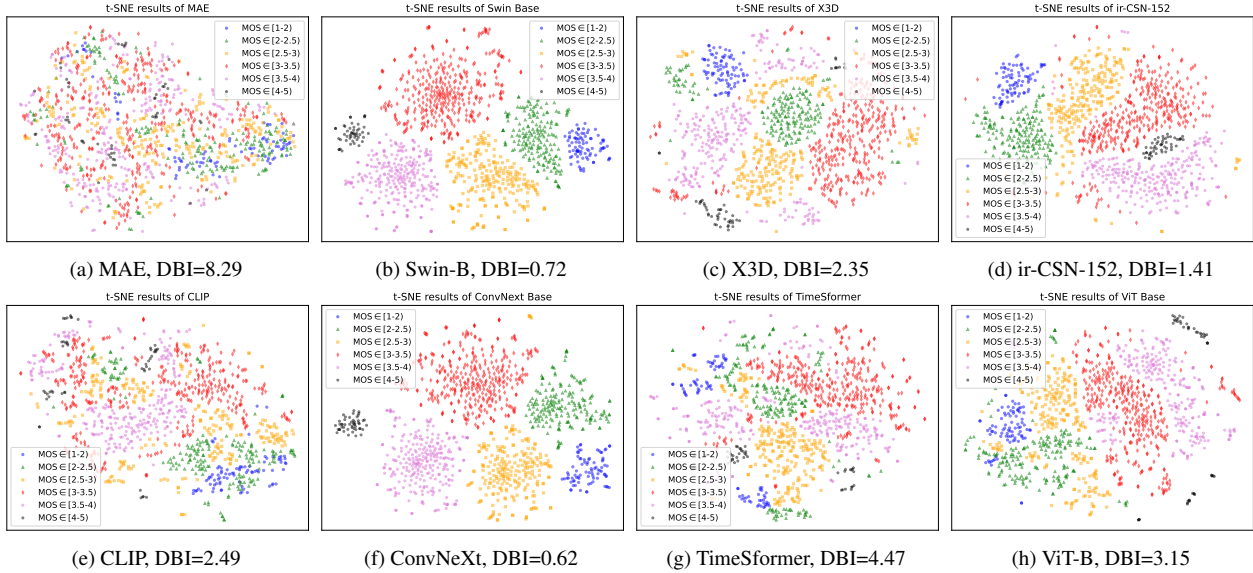


Figure 2. Visualization of clustering results of features extracted by different pretrained models using t-SNE [54]. Videos in KoNViD-1k [22] are used. The number of cluster centers is set to be 6 according to the range of MOS values. And DBI scores, which will be introduced in detail in Sec. 3.4, measure the divergence of clustering results (**the smaller, the better**).

capability into the visual realm. Some subsequent literature [21, 32, 43] had shown that the same benefits can be achieved. For example, CLIP [43] trained on the WebImageText matched the performance of the original ResNet-50 on ImageNet zero-shot, without using any of the original labeled data. In the field of quality assessment (QA), there are also efforts [8, 29, 30, 38] to introduce pretrained models to improve performance. Among them, VSFA [30] extracted features from a pretrained image classification neural network for its inherent content-aware property. And BVQA [29] proposed transferring knowledge from IQA and action recognition datasets with motion patterns. Recently, Ada-DQA [33] utilized diverse pretrained models to distill quality-related knowledge. However, its training cost is relatively high. We hope to tap the potential of the pretrained model itself and reduce the tuning process, in this work.

Metric learning. Metric learning can learn distance metrics from data to measure the difference between samples. It has been used in many research, including QA. RankIQA [35] trained a siamese network to rank synthesized images with different levels of distortions constrained by pairwise ranking hinge loss and then finetune the model on the target IQA dataset. UNIQUE [70] sampled ranked image pairs from individual IQA datasets and used a fidelity loss [51] and a hinge constraint to supervise the training process. FPR [6] extracted distortion/reference feature from the input/reference, hallucinated pseudo reference feature from the input alone, and used a triplet loss [46] to pull the pristine and hallucinated reference features closer while pushing the distortion feature away.

3. Method

3.1. Observations

In recent years, there has been a surge of research attention towards pretraining, as evidenced by a number of notable works [4, 16, 21, 43], that demonstrate the effectiveness of applying pretrained models to downstream tasks. This meets the main obstacle of VQA tasks, where the cost of annotation poses a significant challenge in scaling up datasets. In addition to such efforts, the field of VQA has also witnessed endeavors [8, 29, 30] towards leveraging pretraining to capture intrinsic content-aware or motion-related patterns, with a view to enhancing the representation of perceptual qualities. However, the impact of various factors inherent to pretrained models (*e.g.*, neural network architectures, pre-text tasks, and pretrained databases) on the performance of model transfer remains a subject of inquiry. To the best of our knowledge, there has been limited exploration and exploitation of these factors, as well as newly-appeared cutting-edge pretrained models, in VQA. Therefore, our objective is to investigate ways to fully leverage these models in VQA applications.

In order to examine the relationship between pretrained models and VQA tasks, we designed a simple clustering experiment. Specifically, we selected a pretrained model and utilized its frozen weights as a feature extractor to obtain corresponding video features. We then clustered these features into multiple centers using LMNN [59], based on their range of MOS values. To this end, we selected eight models, including MAE [21] trained on ImageNet-1k [15],

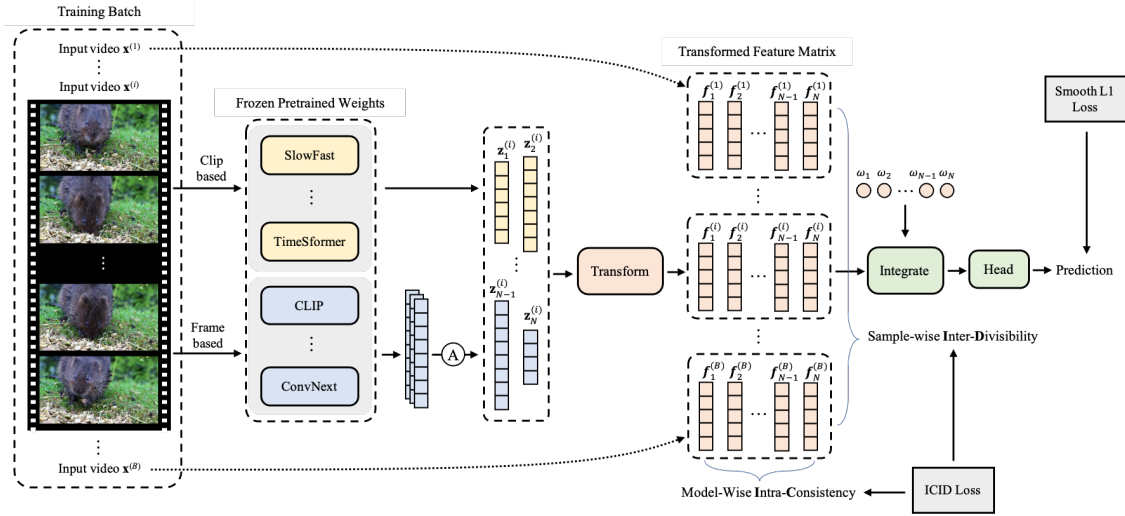


Figure 3. The pipeline of the proposed PTM-VQA. Features of input videos are extracted by pretrained models with frozen weights, transformed to the same dimension, and integrated to generate the final representation. Except for the ordinary smooth \mathcal{L}_1 loss for regression, we add an ICID loss to ensure model-wise consistency and sample-wise divisibility.

Swin-Base [36] trained on ImageNet-22k [15], X3D [19] trained on Kinetics-400 [24], ir-CSN-152 [50] trained on Sports-1M [23], CLIP [43] trained on WebImageText [43], ConvNeXt [37] trained on ImageNet-22k, TimeSformer [3] trained on Kinetics-400, and ViT-Base [17] trained on ImageNet-22k. In Fig. 2, we show that some of these models display surprisingly discriminatory results, despite not having been exposed to quality-related labels during pre-text task training. We hypothesize that some quality-aware representations were learned concurrently during the pre-text task training. For instance, CLIP, which learns visual concepts through natural language supervision, may include emotional descriptions relating to image quality in some texts. Similarly, other models trained on action tasks (e.g., ir-CSN-152) may be sensitive to motion-related distortions [29] (e.g., camera shaking or motion blurriness). As such, these broader pretrained models may be useful in improving VQA task performance.

3.2. Pipeline of the proposed PTM-VQA

Assuming the availability of multiple pre-trained models, the conventional approach for employing them in VQA tasks is through fine-tuning on target datasets while integrating extracted features for quality prediction. Nonetheless, this approach is computationally resource-intensive, making it less feasible as the number of pre-trained models increases and their sizes become larger. For instance, the ViT model requires a TPUv3 with eight cores and 30 days of training [17], while the MAE model consumes 128 TPUv3 cores and 800 epochs of training [21]. This would be unaffordable in a VQA task. However, the findings illustrated in Fig. 2 suggest that pretrained models have the

potential to be applied to VQA tasks with their weights frozen. In this paper, we propose an effective framework, named PTM-VQA, to utilize the knowledge from diverse pretrained models efficiently.

As shown in Fig. 3, given an input video $\mathbf{x}^{(i)}$, N pretrained models, whose weights are frozen, are utilized to extract features, resulting in representations from different perspectives. Specifically, for video clip-based models, we uniformly sample frames in the temporal dimension to form the input clip. Corresponding representations are then generated by these models. For frame-based models, they are fed with sampled frames and the output features are averaged to perform the spatiotemporal representation. Features extracted by models can be noted as $\mathbf{z}_n^{(i)}$, where $n \in \{1, \dots, N\}$. To further distill quality-aware features and perform dimension alignment, we apply a learnable transformation module following each feature extractor. Structurally, the transformation module consists of two fully connected layers, each followed by a normalization layer and an activation layer of GELU. The transformed features are defined as $\mathbf{f}_n^{(i)} \in \mathbb{R}^D$, where D represents the aligned dimension. Then features are integrated to obtain a unified representation through:

$$\mathbf{h}^{(i)} = \frac{\sum_{n=1}^N \omega_n \mathbf{f}_n^{(i)}}{\sum_{n=1}^N \omega_n}, \quad (1)$$

where ω_n is the coefficient for each model. When ω_n equals $1/N$, it means calculating an average, with each model contributing equally to the final representation. Last, $\mathbf{h}^{(i)}$ is used to get the quality prediction through a regression head, which is a single fully-connected layer.

Drawing on the proposed design, the training methodol-

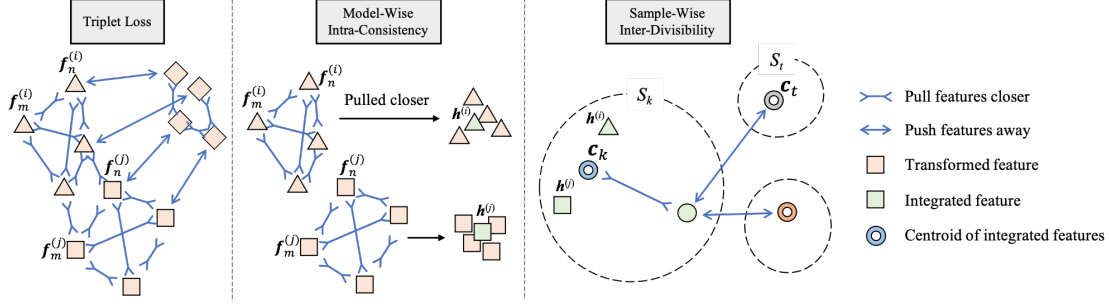


Figure 4. Illustration of ICID loss. The figure shows examples of several triplets of triplet loss; two sets of intra-consistency between features extracted by four pretrained models; and one sample (with two triplets) for inter-divisibility.

ogy exhibits remarkable efficiency, thereby circumventing the computational overhead associated with the finetuning approach aforementioned. As attested by the results presented in Tab. 1, the entire training regimen can be accomplished within a span of approximately **two hours, leveraging a single GPU**. This retains the information of the pretrained models well, but it also increases the difficulty of obtaining preferable performance due to the reduction of learnable parameters. Some concerns are as follows:

1. Due to various pre-texts of pretrained models, features generated by different models are of **large diversity**, which may distribute over inconsistent feature spaces [62]. How to constrain these abundant features into a unified quality-aware space is important.
2. Different from the objective category in common classification tasks, the perceptual quality of a video is more implicit and related to various factors (*e.g.*, content attractiveness, distortion type, and level, motion pattern and level), whereas videos of the same quality often render completely different content and vice versa. Therefore, it is difficult for the models trained based on objective annotations to distinguish these samples of **the same category** but with a large perceptual **quality difference**. A more comprehensive contrast approach beyond sample-wise comparison needs to be proposed to deal with these outliers.
3. There exist **hundreds of** pretrained models available in public libraries. How to select the desired models efficiently and how to determine the contribution of these models to represent the perceptual quality effectively is an urgent problem to be solved.

3.3. Intra-Consistency and Inter-Divisibility Loss

To solve the above concerns and better satisfy VQA tasks, we intend to constrain the features between different pretrained models and different samples using metric learning. Triplet loss, which is one of the most widely adopted metric learning measures, can be formed as follows:

$$\mathcal{L}_{\text{triplet}}(\mathbf{f}_{\hat{a}}, \mathbf{f}_{\hat{p}}, \mathbf{f}_{\hat{n}}) = \max(\|\mathbf{f}_{\hat{a}} - \mathbf{f}_{\hat{p}}\|^2 - \|\mathbf{f}_{\hat{a}} - \mathbf{f}_{\hat{n}}\|^2 + \alpha, 0), \quad (2)$$

where $\mathbf{f}_{\hat{a}}, \mathbf{f}_{\hat{p}}, \mathbf{f}_{\hat{n}}$ are features of an anchor sample \hat{a} , a positive sample \hat{p} of the same class as \hat{a} , and a negative sample \hat{n} which has a different class of \hat{a} . And α is a margin between anchor-positive and anchor-negative pairs. Some previous studies [6, 20] in QA also applied triplet loss to measure the distance between the distorted feature and the reference feature of the same sample. Since the MOS values are continuous, the original triplet loss cannot be directly used to constrain the distance between arbitrary samples. We make some modifications to constrain features generated by different pretrained models and samples, as given in Fig. 4.

Intra-consistency constraint. To solve the first concern, and unify features generated by different pretrained models into a unified quality-aware latent space, we propose a model-wise intra-consistency constraint. Formally, it is defined to minimize the distance between arbitrary two of the transformed features through computing a cosine similarity, which is widely used in deep metric learning [55]:

$$\mathcal{L}_{\text{intra}} = \frac{2}{N \cdot (N - 1)} \sum_{n=1}^N \sum_{m, m \neq n}^N \left(1 - \frac{\mathbf{f}_n^{(i)} \cdot \mathbf{f}_m^{(i)}}{\|\mathbf{f}_n^{(i)}\|_2 \|\mathbf{f}_m^{(i)}\|_2}\right). \quad (3)$$

Inter-divisibility constraint. To solve the second concern, we split videos into distinct pseudo clusters under different numerical intervals, according to the annotated MOS values (on a scale of 1.0 to 5.0). For example, videos with MOS in the range of 1.0 to 2.0 are generally considered to be of poor quality, and whose content cannot be normally recognized due to the existence of various distortions. And videos with MOS in the range of 4.0 to 5.0 are of high quality, whose content is unambiguous, without noise, shaking, and blurring. We identify the videos within the same range as the same category, thus dividing them into K clusters. Each cluster can be noted as $\mathcal{S}_k = \{\mathbf{x}^{(i)} | y^{(i)} \in (p_k, q_k], q_k > p_k \in [1.0, 5.0]\}$, where $y^{(i)}$ is the labeled MOS for the i -th input video, p_k and q_k are the endpoints of the interval. Through this pseudo cluster, triplet loss can be utilized for samples belonging to the

Table 1. Details of PTM-VQA for different datasets. Time¹ is the training time. Time² represents the inference time using a 1080P/30FPS/20s video. The overall computational cost is relatively small compared with existing SOTA methods.

Dataset	Combination of Pretrained Models	Frames	Interval	LR	Param(M)	FLOPs(T)	Mem(G)	Time ¹ (h)	Time ² (s)
KoNViD-1k	ConvNeXt, ir-CSN-152, CLIP	16	2	1e-3	0.66	0.43	4.94	2.00	0.14
LIVE-VQC	CLIP, TimeSformer	16	4	5e-3	0.30	0.53	4.32	1.97	0.17
YouTube-UGC	ConvNeXt, ir-CSN-152, CLIP, Video Swin-B	32	8	1e-3	0.86	1.35	5.32	2.34	1.20

same cluster to be closer and samples of the different clusters to be farther away. Then Equ. 2 can be rewritten as:

$$\mathcal{L}_{\text{triplet}}(\mathbf{h}^{(i)}, \mathbf{h}^{(j)}, \mathbf{h}^{(l)}), \text{ where } \mathbf{x}^{(i)}, \mathbf{x}^{(j)} \in \mathcal{S}_k, \mathbf{x}^{(l)} \notin \mathcal{S}_k. \quad (4)$$

Besides, the original feature \mathbf{f} extracted by individual models is replaced by the integrated feature \mathbf{h} . As shown in Fig. 4, the original triplet loss performs a sample-to-sample form, which is highly affected by the sampling of triples. When facing outliers that are of the same quality but render different contents or vice versa, it may lead to bad local minima and prevent the model from achieving top performance. Thus we propose using the centroid of the cluster to represent the positive and negative points as:

$$\mathcal{L}_{\text{inter}} = \mathcal{L}_{\text{triplet}}(\mathbf{h}^{(i)}, \mathbf{c}_k, \mathbf{c}_t), \text{ where} \\ \mathbf{c}_k = \frac{1}{|\mathcal{S}_k|} \sum_{\{i|\mathbf{x}^{(i)} \in \mathcal{S}_k\}} \mathbf{h}^{(i)}, \mathbf{c}_t = \frac{1}{|\mathcal{S}_t|} \sum_{\{j|\mathbf{x}^{(j)} \in \mathcal{S}_t\}} \mathbf{h}^{(j)}. \quad (5)$$

Given a batch consisting B inputs, during training, the optimization objective can be summarized as:

$$\min \mathcal{L}_1 + \beta \left(\sum_{i=1}^B \mathcal{L}_{\text{intra}} + \mathcal{L}_{\text{inter}} \right), \quad (6)$$

where β is the coefficient balancing smooth \mathcal{L}_1 regression loss and the proposed ICID loss.

3.4. Selection scheme through DBI

For the third concern, we observe an obvious difference in the clustering results of different pretrained models in Fig. 2. Since the weights of models are frozen both in the clustering test and subsequent training process, the divergence of clustering results can reflect the relevance of VQA tasks. We propose using the Davies-Bouldin Index (DBI) [14] as a metric for model selection, which is commonly employed for evaluating clustering results. In our particular setting, the DBI can be expressed as follows:

$$\psi = \frac{1}{K} \sum_{k=1}^K \max_{t \neq k} \frac{\mathbf{d}_k + \mathbf{d}_t}{\|\mathbf{c}_k - \mathbf{c}_t\|_2}, \text{ where} \\ \mathbf{c}_k = \frac{1}{|\mathcal{S}_k|} \sum_{\mathcal{S}_k} \mathbf{z}^{(i)}, \mathbf{d}_k = \frac{1}{|\mathcal{S}_k|} \sum_{\mathcal{S}_k} \|\mathbf{z}^{(i)} - \mathbf{c}_k\|_2, \quad (7)$$

where \mathbf{c}_k is the centroid of cluster \mathcal{S}_k for the set of extracted feature $\mathbf{z}^{(i)}$, \mathbf{d}_k represents the average distance between each sample and its corresponding centroid. For the

n -th model, its DBI score can be noted as ψ_n . A lower DBI indicates better clustering performance, which means that the pretrained model (e.g., ConvNeXt, Swin-Base, ir-CSN-152, CLIP in Fig. 2) is more relevant to downstream VQA tasks. During training, the DBI scores computed offline can be used in the aggregation procedure as given in Equ. 1, where ω_n can be replaced by $1/\psi_n$. It means the models that are more relevant to the VQA task contribute more to the feature representation.

4. Experiments

4.1. Datasets and evaluation criteria

Datasets. Our method is evaluated on 4 public NR-VQA datasets, including KoNViD-1k [22], LIVE-VQC [48], YouTube-UGC [56] and LSVQ [67]. In detail, KoNViD-1k contains 1,200 videos that are fairly filtered from a large public video dataset YFCC100M. The videos are 8 seconds long with 24/25/30 FPS and a resolution of 960×540 . The MOS ranges from 1.22 to 4.64. Each video owns 114 annotations to get a reliable MOS. LIVE-VQC consists of 585 videos with complex authentic distortions captured by 80 different users using 101 different devices, with 240 annotations for each video. YouTube-UGC has 1,380 UGC videos sampled from YouTube with a duration of 20 seconds and resolutions from 360P to 4K, with 123 annotations for each video. And LSVQ is the largest VQA dataset currently (proposed in 2021) with 39,076 videos. All the datasets contain no pristine videos, thus only NR methods can be evaluated on them. Following [49, 65], we split the dataset into a 80% training set and a 20% testing set randomly for the first three datasets. For LSVQ, we follow the official split setting. We perform 10 repeat runs in each dataset using different splittings to get the mean values of PLCC and SRCC.

Evaluation criteria. Pearson’s Linear Correlation Coefficient (PLCC) and Spearman’s Rank-Order Correlation Coefficient (SRCC) are selected as criteria to measure the accuracy and monotonicity. They are in the range of $[0, 1]$. A larger PLCC means a more accurate numerical fit with MOS scores. A larger SRCC shows a more accurate ranking between samples. Besides, the mean average of PLCC and SRCC is also reported as a comprehensive criterion.

4.2. Implementation details

Our experiments are performed using PyTorch [41] and MMAction2 [12], and are all conducted on **one** Nvidia

Table 2. Comparisons with existing methods. The “-” is an unreported result. The “**” is using extra labeled data for training. The best and second best results are **bolded** and underlined. Average scores weighted by the size of datasets are also reported.

Method	KoNViD-1k			LIVE-VQC			YouTube-UGC			Weighted Avg		
	PLCC	SRCC	Mean	PLCC	SRCC	Mean	PLCC	SRCC	Mean	PLCC	SRCC	Mean
VIIDEO [40]	0.3030	0.2980	0.3005	0.2164	0.0332	0.1248	0.1534	0.0580	0.1057	0.2218	0.1444	0.1832
NIQE [39]	0.5530	0.5417	0.5473	0.6286	0.5957	0.6121	0.2776	0.2379	0.2577	0.4469	0.4365	0.4417
BRISQUE [38]	0.626	0.654	0.640	0.638	0.592	0.615	0.395	0.382	0.388	0.5275	0.5240	0.5257
VSFA [30]	0.744	0.755	0.749	-	-	-	-	-	-	-	-	-
TLVQM [27]	0.7688	0.7729	0.7708	0.8025	0.7988	0.8006	0.6590	0.6693	0.6641	0.7272	0.7325	0.7298
RIRNet [8]	0.7812	0.7755	0.7783	0.7982	0.7713	0.7847	-	-	-	-	-	-
UGC-VQA [52]	0.7803	0.7832	0.7817	0.7514	0.7522	0.7518	0.7733	0.7787	0.7760	0.7663	0.7732	0.7697
CSPT [9]	0.8062	0.8008	0.8035	0.8194	0.7989	0.8091	-	-	-	-	-	-
RAPIQUE [53]	0.8175	0.8031	0.8103	0.7863	0.7548	0.7705	0.7684	0.7591	0.7637	0.7925	0.7789	0.7857
StarVQA [64]	0.796	0.812	0.804	0.808	0.732	0.770	-	-	-	-	-	-
BVQA* [29]	0.8335	0.8362	0.8348	0.8415	0.8412	0.8413	0.8194	0.8312	0.8253	0.8352	0.8349	0.8351
STDAM* [65]	0.8415	0.8448	0.8431	0.8204	0.7931	0.8067	0.8297	0.8341	0.8319	0.8325	0.8337	0.8331
Fast-VQA [63]	0.855	0.859	0.857	0.844	<u>0.823</u>	<u>0.834</u>	-	-	-	-	-	-
VQT [69]	0.8684	<u>0.8582</u>	<u>0.8633</u>	<u>0.8357</u>	0.8238	0.8298	0.8514	<u>0.8357</u>	<u>0.8436</u>	<u>0.8529</u>	<u>0.8421</u>	<u>0.8475</u>
PTM-VQA	0.8718	0.8568	0.8643	0.8198	0.8110	0.8154	0.8570	0.8578	0.8574	0.8591	0.8454	0.8523

V100 GPU by training for 60 epochs. For KoNViD-1k, we select ConvNeXt, ir-CSN-152, and CLIP as feature extractors. For LIVE-VQC, we use CLIP and TimeSformer. For YouTube-UGC, an extra Video Swin-Base is used together with those selected on KoNViD-1k. For KoNViD-1k, we sample 16 frames with a frame interval of 2. As videos in LIVE-VQC and YouTube-UGC has a longer time duration, we use larger intervals for these two datasets. Since most augmentations will introduce extra interference to the quality of videos [25], we only choose the center crop to produce an input with a size of 224×224 . During training, we use AdamW optimizer with a weight decay of 0.02. Cosine annealing with a warmup of 2 epochs is adopted to control the learning rate. The dimension D of transformed features is set to 128. The margin α is set to be 0.05. β is set to be 0.2. By default, we select the checkpoint generated by the last iteration for evaluation. During inference, we follow a similar procedure as given in [1] by using 4×5 views. To be specific, 4 clips are uniformly sampled from a video in the temporal domain. For each clip, we take 5 crops in the four corners and the center. The final score is computed as the average score. More details are given in Tab. 1.

4.3. Comparison with SOTA methods

We select existing VQA methods for comparison in three datasets. As shown in Tab. 2, our method obtains competitive results on all three datasets. Compared with traditional methods that rely on statistical regularities (e.g., VIIDEO [40], NIQE [39], and BRISQUE [38]), PTM-VQA models outperform by large margins. Compared with some deep learning-based methods that apply well-designed networks (e.g., TLVQM [27], StarVQA [64]), PTM-VQA still obtains higher performances. Especially, VSFA [30] and RIRNet [8] also adopt pretrained mod-

Table 3. Comparisons on the largest LSVQ dataset.

Method	LSVQ-Test		LSVQ-1080P	
	PLCC	SRCC	PLCC	SRCC
BRISQUE [38]	0.576	0.569	0.531	0.497
VSFA [30]	0.796	0.801	0.704	0.675
TLVQM [27]	0.774	0.772	0.616	0.589
PVQ (w/o patch) [67]	0.816	0.814	0.708	0.686
PVQ (w patch) [67]	0.828	0.827	0.739	0.711
PTM-VQA-1k	0.8536	0.8530	0.7784	0.7279
PTM-VQA-VQC	0.8637	0.8545	0.7817	0.7359
PTM-VQA-UGC	0.8443	0.8429	0.7769	0.7300

els that contain content-dependency or motion information to finetune in VQA tasks. PTM-VQA demonstrates that features extracted directly from pretrained models can also achieve better results. As the best two SOTA methods BVQA [29] and STDAM [65] who utilize extra IQA datasets, **PTM-VQA proves that transferring knowledge from pretrained models can achieve competitive results compared with a model trained with additional data.**

To assess the generalizability of the selected combinations, we evaluate on the largest LSVQ dataset using the three combinations utilized in KoNViD-1k, LIVE-VQC, and YouTube-UGC. As given in Tab. 3, PTM-VQA models demonstrate a significant performance advantage over existing methods, **indicating the benefits of leveraging pretrained models with a larger amount of data.**

We also compare the cost of inference time with open-source methods on a 1080P video (100 repeat runs). And the inference time cost is 75s (BRISQUE), 248s (TLVQM), 117s (VSFA), 0.12s (StarVQA), and 2.45s (BVQA) respectively. Thanks to the **reduced dimensions (e.g., frame sampling, center cropping) and model selection using**

Table 4. Cross-database comparison using models trained on LSVQ, then evaluated on the entire KoNViD-1k and LIVQ-VQC datasets without fine-tuning.

Method	KoNViD-1k		LIVQ-VQC	
	PLCC	SRCC	PLCC	SRCC
BRISQUE [38]	0.647	0.646	0.536	0.524
VSFA [30]	0.794	0.784	0.772	0.734
TLVQM [27]	0.724	0.732	0.691	0.670
PVQ (w/o patch) [67]	0.781	0.781	0.776	0.747
PVQ (w patch) [67]	0.795	0.791	0.807	0.770
PTM-VQA-1k	0.8249	0.8164	0.7772	0.7215
PTM-VQA-VQC	0.8292	0.8243	0.7772	0.7231
PTM-VQA-UGC	0.8303	0.8223	0.7852	0.7367

Table 5. Ablation on the ICID loss in KoNViD-1k.

\mathcal{L}_1	\mathcal{L}_{intra}	\mathcal{L}_{inter}	\mathcal{L}_{tri}	PLCC	SRCC
✓	✓	✓		0.8718	0.8568
✓			✓	0.7968	0.7850
✓				0.7867	0.7655
✓	✓			0.8545	0.8299
✓		✓		0.8172	0.7707

Table 6. Ablation on cluster settings in KoNViD-1k.

K	intervals	PLCC	SRCC
2	$\mathcal{S}_1=[1,3), \mathcal{S}_2=[3, 5]$	0.8277	0.8066
4	$\mathcal{S}_1=[1,2), \mathcal{S}_2=[2,3), \mathcal{S}_3=[3,4), \mathcal{S}_4=[4,5]$	0.8431	0.8012
6	$\mathcal{S}_1=[1,2), \mathcal{S}_2=[2, 2.5), \mathcal{S}_3=[2.5, 3), \mathcal{S}_4=[3, 3.5), \mathcal{S}_5=[3.5, 4), \mathcal{S}_6=[4, 5]$	0.8718	0.8568

DBI, PTM-VQA models do not significantly increase inference time over StarVQA and BVQA, as given in Tab. 1. Meanwhile, due to the different number and composition of pretrained models, the calculation cost of PTM-VQA models vary. Even so, the largest PTM-VQA can process a high-resolution video in about 1s, and the smaller models can process nearly 6/7 videos per second.

Cross-database comparison. To emphasize the validity and generalizability, we perform the cross-database evaluation in Tab. 4. Models trained on LSVQ are tested on much smaller datasets of KoNViD-1k and LIVE-VQA directly. It can be seen that PTM-VQA transferred very well to both datasets, highlighting the general efficacy.

4.4. Ablation studies

We conduct experimental analysis to evaluate the effectiveness of each component. By default, experiments are performed following the best configurations in Sec. 4.2.

Ablation on different constraints. As given in Tab. 5, direct usage of triplet loss cannot obtain satisfying results.

Table 7. Ablation on the effectiveness of DBI.

Datasets	ω_n	PLCC	SRCC
KoNViD-1k	$1/N$	0.8631	0.8521
	$1/\psi$	0.8718	0.8568
LIVE-VQC	$1/N$	0.8205	0.8107
	$1/\psi$	0.8198	0.8110
YouTube-UGC	$1/N$	0.8427	0.8446
	$1/\psi$	0.8570	0.8578

When either or both constraints are absent, performance degrades significantly. These prove the effectiveness of intra-consistency constraints in transferring knowledge from different pretrained models and inter-divisibility constraints in generating stable predictions.

Ablation on the clustering settings. Tab. 6 gives the results with different numbers of clusters. When K is 2, videos are simply classified as low-quality and high-quality ones. When K is 4, videos are evenly divided into four parts on a scale of 1.0 to 5.0. Due to the relatively small amount of data at both endpoints, a 6-split setting can be obtained by using fine-grained division in the middle fraction segment. Since the need to ensure the number of samples per cluster within the batch, a larger number of clusters are not attempted. The best result can be acquired when K is 6.

Ablation on the DBI strategy. The effectiveness of DBI can be evaluated in two aspects: (1) Model selection strategy. We performed 10 experiments based on randomly selected pretrained models in KoNViD-1k, resulting in PLCC (0.7917 ± 0.0578 , SRCC (0.7583 ± 0.0492). Compared with the DBI-based strategy, the performance is poor and the randomness is high. (2) Feature integration. Tab. 7 shows the effectiveness of DBI in guiding the integration of different models, allowing more relevant models to contribute more.

5. Conclusion

In this paper, we proposed PTM-VQA that utilizes in-the-wild pretrained models as feature extractors for NR-VQA tasks, transferring quality-related knowledge from diverse pre-text domains. The DBI scores are used to select candidates from a large amount of available pretrained models. To constrain features with large diversity into a unified latent space of quality and tackle outliers, we propose a new ICID loss. Under small computational cost, PTM-VQA models obtain SOTA results in widely-used benchmarks. Experiments in larger datasets and cross-database evaluation further prove generalizability.

Acknowledgments

This work was partly supported by the Natural Science Foundation of China (NSFC) under Grant No. 62306309.

References

- [1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lucic, and Cordelia Schmid. Vivit: A video vision transformer. *CoRR*, abs/2103.15691, 2021. 7
- [2] Thomas Barnett, Shruti Jain, Usha Andra, and Taru Khurana. Cisco visual networking index (vni) complete forecast update, 2017–2022. *Americas/EMEAR Cisco Knowledge Network (CKN) Presentation*, 2018. 1
- [3] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, pages 813–824. PMLR, 2021. 4
- [4] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *NeurIPS*, 2020. 2, 3
- [5] João Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *CoRR*, abs/1907.06987, 2019. 1
- [6] Baoliang Chen, Lingyu Zhu, Chenqi Kong, Hanwei Zhu, Shiqi Wang, and Zhu Li. No-reference image quality assessment by hallucinating pristine features. *IEEE Trans. Image Process.*, 2022. 3, 5
- [7] Baoliang Chen, Lingyu Zhu, Guo Li, Fangbo Lu, Hongfei Fan, and Shiqi Wang. Learning generalized spatial-temporal deep feature representation for no-reference video quality assessment. *IEEE Trans. Circuits Syst. Video Technol.*, 32(4): 1903–1916, 2022. 1, 2
- [8] Pengfei Chen, Leida Li, Lei Ma, Jinjian Wu, and Guangming Shi. Rirnet: Recurrent-in-recurrent network for video quality assessment. In *ACM Multimedia*, pages 834–842. ACM, 2020. 3, 7
- [9] Pengfei Chen, Leida Li, Jinjian Wu, Weisheng Dong, and Guangming Shi. Contrastive self-supervised pre-training for video quality assessment. *IEEE Trans. Image Process.*, 31: 458–471, 2022. 1, 7
- [10] Yanjiao Chen, Kaishun Wu, and Qian Zhang. From qos to qoe: A tutorial on video quality assessment. *IEEE Commun. Surv. Tutorials*, 17(2):1126–1165, 2015. 1
- [11] Shyamprasad Chikkerur, Vijay Sundaram, Martin Reisslein, and Lina J. Karam. Objective video quality assessment methods: A classification, review, and performance comparison. *IEEE Trans. Broadcast.*, 57(2):165–182, 2011. 1
- [12] MMAAction2 Contributors. Openmmlab’s next generation video understanding toolbox and benchmark. <https://github.com/open-mmlab/mmaaction2>, 2020. 6
- [13] Dubravko Culibrk, Dragan Kukulj, Petar Vasiljevic, Maja Pokric, and Vladimir Zlokolic. Feature selection for neural-network based no-reference video quality assessment. In *ICANN (2)*, pages 633–642. Springer, 2009. 2
- [14] David L. Davies and Donald W. Bouldin. A cluster separation measure. *IEEE Trans. Pattern Anal. Mach. Intell.*, 1(2): 224–227, 1979. 2, 6
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. IEEE Computer Society, 2009. 2, 3, 4
- [16] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*, pages 4171–4186. Association for Computational Linguistics, 2019. 2, 3
- [17] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*. OpenReview.net, 2021. 2, 4
- [18] Qiang Fan, Wang Luo, Yuan Xia, Guozhi Li, and Daojing He. metrics and methods of video quality assessment: a brief review. *Multim. Tools Appl.*, 78(22):31019–31033, 2019. 1
- [19] Christoph Feichtenhofer. X3D: expanding architectures for efficient video recognition. In *CVPR*, pages 200–210. Computer Vision Foundation / IEEE, 2020. 4
- [20] S. Alireza Golestaneh, Saba Dadsetan, and Kris M. Kitani. No-reference image quality assessment via transformers, relative ranking, and self-consistency. In *WACV*, pages 1220–1230, 2022. 5
- [21] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. *CoRR*, abs/2111.06377, 2021. 3, 4
- [22] Vlad Hosu, Franz Hahn, Mohsen Jenadeleh, Hanhe Lin, Hui Men, Tamás Szirányi, Shujun Li, and Dietmar Saupe. The konstanz natural video database (konvid-1k). In *QoMEX*, pages 1–6. IEEE, 2017. 1, 3, 6
- [23] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014. 4
- [24] Will Kay, João Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. *CoRR*, abs/1705.06950, 2017. 4
- [25] Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. MUSIQ: multi-scale image quality transformer. In *ICCV*, pages 5128–5137. IEEE, 2021. 7
- [26] Janusz Klink and Tadeus Uhl. Video quality assessment: Some remarks on selected objective metrics. In *SoftCOM*, pages 1–6. IEEE, 2020. 1
- [27] Jari Korhonen. Two-level approach for no-reference consumer video quality assessment. *IEEE Trans. Image Process.*, 28(12):5923–5938, 2019. 1, 7, 8
- [28] Koffi Kossi, Stéphane Coulombe, Christian Desrosiers, and Ghyslain Gagnon. No-reference video quality assessment using distortion learning and temporal attention. *IEEE Access*, 10:41010–41022, 2022. 1, 2

- [29] Bowen Li, Weixia Zhang, Meng Tian, Guangtao Zhai, and Xianpei Wang. Blindly assess quality of in-the-wild videos via quality-aware pre-training and motion perception. *IEEE Trans. Circuits Syst. Video Technol.*, 32(9):5944–5958, 2022. [2](#), [3](#), [4](#), [7](#)
- [30] Dingquan Li, Tingting Jiang, and Ming Jiang. Quality assessment of in-the-wild videos. In *ACM Multimedia*, pages 2351–2359. ACM, 2019. [1](#), [2](#), [3](#), [7](#), [8](#)
- [31] Dingquan Li, Tingting Jiang, and Ming Jiang. Unified quality assessment of in-the-wild videos with mixed datasets training. *Int. J. Comput. Vis.*, 129(4):1238–1257, 2021. [1](#)
- [32] Yanghao Li, Saining Xie, Xinlei Chen, Piotr Dollár, Kaiming He, and Ross B. Girshick. Benchmarking detection transfer learning with vision transformers. *CoRR*, abs/2111.11429, 2021. [3](#)
- [33] Hongbo Liu, Mingda Wu, Kun Yuan, Ming Sun, Yansong Tang, Chuanchuan Zheng, Xing Wen, and Xiu Li. Ada-dqa: Adaptive diverse quality-aware feature acquisition for video quality assessment. In *ACM Multimedia*, pages 6695–6704. ACM, 2023. [3](#)
- [34] Wentao Liu, Zhengfang Duanmu, and Zhou Wang. End-to-end blind quality assessment of compressed videos using deep neural networks. In *ACM Multimedia*, pages 546–554. ACM, 2018. [1](#)
- [35] Xialei Liu, Joost van de Weijer, and Andrew D. Bagdanov. Rankiqa: Learning from rankings for no-reference image quality assessment. In *ICCV*, pages 1040–1049. IEEE Computer Society, 2017. [3](#)
- [36] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 9992–10002. IEEE, 2021. [4](#)
- [37] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *CoRR*, abs/2201.03545, 2022. [4](#)
- [38] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE Trans. Image Process.*, 21(12):4695–4708, 2012. [3](#), [7](#), [8](#)
- [39] Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal Process. Lett.*, 20(3):209–212, 2013. [7](#)
- [40] Anish Mittal, Michele A. Saad, and Alan C. Bovik. A completely blind video integrity oracle. *IEEE Trans. Image Process.*, 25(1):289–300, 2016. [1](#), [7](#)
- [41] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, pages 8024–8035, 2019. [6](#)
- [42] Lihui Qian, Tianxiang Pan, Yunfei Zheng, Jiajie Zhang, Mading Li, Bing Yu, and Bin Wang. No-reference nonuniform distorted video quality assessment based on deep multiple instance learning. *IEEE Multim.*, 28(1):28–37, 2021. [1](#)
- [43] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. [3](#), [4](#)
- [44] ITUT Rec. P. 800.1, mean opinion score (mos) terminology. *International Telecommunication Union, Geneva*, 2006. [2](#)
- [45] Michele A. Saad, Alan C. Bovik, and Christophe Charrier. Blind prediction of natural video quality. *IEEE Trans. Image Process.*, 23(3):1352–1365, 2014. [1](#)
- [46] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, pages 815–823. IEEE Computer Society, 2015. [3](#)
- [47] Muhammad Shahid, Andreas Rossholm, Benny Löfvström, and Hans-Jürgen Zepernick. No-reference image and video quality assessment: a classification and review of recent approaches. *EURASIP J. Image Video Process.*, 2014:40, 2014. [1](#)
- [48] Zeina Sinno and Alan Conrad Bovik. Large-scale study of perceptual video quality. *IEEE Trans. Image Process.*, 28(2):612–627, 2019. [6](#)
- [49] Shaolin Su, Qingsen Yan, Yu Zhu, Cheng Zhang, Xin Ge, Jinqiu Sun, and Yanning Zhang. Blindly assess image quality in the wild guided by a self-adaptive hyper network. In *CVPR*, pages 3664–3673. Computer Vision Foundation / IEEE, 2020. [6](#)
- [50] Du Tran, Heng Wang, Matt Feiszli, and Lorenzo Torresani. Video classification with channel-separated convolutional networks. In *ICCV*, pages 5551–5560. IEEE, 2019. [4](#)
- [51] Ming-Feng Tsai, Tie-Yan Liu, Tao Qin, Hsin-Hsi Chen, and Wei-Ying Ma. Frank: a ranking method with fidelity loss. In *SIGIR*, pages 383–390. ACM, 2007. [3](#)
- [52] Zhengzhong Tu, Yilin Wang, Neil Birkbeck, Balu Adsumilli, and Alan C. Bovik. UGC-VQA: benchmarking blind video quality assessment for user generated content. *IEEE Trans. Image Process.*, 30:4449–4464, 2021. [7](#)
- [53] Zhengzhong Tu, Xiangxu Yu, Yilin Wang, Neil Birkbeck, Balu Adsumilli, and Alan C. Bovik. RAPIQUE: rapid and accurate video quality prediction of user generated content. *CoRR*, abs/2101.10955, 2021. [2](#), [7](#)
- [54] Laurens van der Maaten. Learning a parametric embedding by preserving local structure. In *AISTATS*, pages 384–391. JMLR.org, 2009. [3](#)
- [55] Xun Wang, Xintong Han, Weilin Huang, Dengke Dong, and Matthew R. Scott. Multi-similarity loss with general pair weighting for deep metric learning. In *CVPR*, pages 5022–5030. Computer Vision Foundation / IEEE, 2019. [5](#)
- [56] Yilin Wang, Sasi Inguva, and Balu Adsumilli. Youtube UGC dataset for video compression research. In *MMSP*, pages 1–5. IEEE, 2019. [6](#)
- [57] Yilin Wang, Junjie Ke, Hossein Talebi, Joong Gon Yim, Neil Birkbeck, Balu Adsumilli, Peyman Milanfar, and Feng Yang. Rich features for perceptual quality assessment of UGC videos. In *CVPR*, pages 13435–13444. Computer Vision Foundation / IEEE, 2021. [1](#), [2](#)

- [58] Zhou Wang and Qiang Li. Video quality assessment using a statistical model of human visual speed perception. *JOSA A*, 24(12):B61–B69, 2007. [1](#)
- [59] Kilian Q. Weinberger and Lawrence K. Saul. Distance metric learning for large margin nearest neighbor classification. *J. Mach. Learn. Res.*, 10:207–244, 2009. [2](#), [3](#)
- [60] Ross Wightman. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019. [2](#)
- [61] Stefan Winkler. Issues in vision modeling for perceptual video quality assessment. *Signal Process.*, 78(2):231–252, 1999. [1](#)
- [62] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo Lopes, Ari S. Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *ICML*, pages 23965–23998. PMLR, 2022. [5](#)
- [63] Haoning Wu, Chaofeng Chen, Jingwen Hou, Liang Liao, Annan Wang, Wenxiu Sun, Qiong Yan, and Weisi Lin. FAST-VQA: efficient end-to-end video quality assessment with fragment sampling. In *ECCV (6)*, pages 538–554. Springer, 2022. [7](#)
- [64] Fengchuang Xing, Yuan-Gen Wang, Hanpin Wang, Leida Li, and Guopu Zhu. Starvqa: Space-time attention for video quality assessment. *CoRR*, abs/2108.09635, 2021. [2](#), [7](#)
- [65] Jiahua Xu, Jing Li, Xingguang Zhou, Wei Zhou, Baichao Wang, and Zhibo Chen. Perceptual quality assessment of internet videos. In *ACM Multimedia*, pages 1248–1257. ACM, 2021. [1](#), [2](#), [6](#), [7](#)
- [66] Fuzheng Yang, Shuai Wan, Yilin Chang, and Hong Ren Wu. A novel objective no-reference metric for digital video quality assessment. *IEEE Signal Process. Lett.*, 12(10):685–688, 2005. [2](#)
- [67] Zhenqiang Ying, Maniratnam Mandal, Deepti Ghadiyaram, and Alan C. Bovik. Patch-vq: ‘patching up’ the video quality problem. In *CVPR*, pages 14019–14029. Computer Vision Foundation / IEEE, 2021. [6](#), [7](#), [8](#)
- [68] Junyong You. Long short-term convolutional transformer for no-reference video quality assessment. In *ACM Multimedia*, pages 2112–2120. ACM, 2021. [1](#), [2](#)
- [69] Kun Yuan, Zishang Kong, Chuanchuan Zheng, Ming Sun, and Xing Wen. Capturing co-existing distortions in user-generated content for no-reference video quality assessment. In *ACM Multimedia*, pages 1098–1107. ACM, 2023. [7](#)
- [70] Weixia Zhang, Kede Ma, Guangtao Zhai, and Xiaokang Yang. Uncertainty-aware blind image quality assessment in the laboratory and wild. *IEEE Trans. Image Process.*, 30:3474–3486, 2021. [3](#)