

# VQ-Insight: Teaching VLMs for AI-Generated Video Quality Understanding via Progressive Visual Reinforcement Learning

Xuanyu Zhang<sup>1,2,†</sup>, Weiqi Li<sup>1,†</sup>, Shijie Zhao<sup>2,◆,✉</sup>, Junlin Li<sup>2</sup>, Li Zhang<sup>2</sup>, Jian Zhang<sup>1,✉</sup>

<sup>1</sup>School of Electronic and Computer Engineering, Peking University,

<sup>2</sup>ByteDance Inc.

zhaoshijie.0526@bytedance.com, zhangjian.sz@pku.edu.cn.

## Abstract

Recent advances in AI-generated content (AIGC) have led to the emergence of powerful text-to-video generation models. Despite these successes, evaluating the quality of AI-generated videos remains challenging due to limited generalization, lack of temporal awareness, heavy reliance on large-scale annotated datasets, and the lack of effective interaction with generation models. Most current approaches rely on supervised finetuning of vision-language models (VLMs), which often require large-scale annotated datasets and tend to decouple understanding and generation. To address these shortcomings, we propose VQ-Insight, a novel reasoning-style VLM framework for AIGC video quality assessment. Our approach features: (1) a progressive video quality learning scheme that combines image quality warm-up, general task-specific temporal learning, and joint optimization with the video generation model; (2) the design of multi-dimension scoring rewards, preference comparison rewards, and temporal modeling rewards to enhance both generalization and specialization in video quality evaluation. Extensive experiments demonstrate that VQ-Insight consistently outperforms state-of-the-art baselines in preference comparison, multi-dimension scoring, and natural video scoring, bringing significant improvements for video generation tasks.

**Code** — <https://github.com/bytedance/Q-Insight>

## Introduction

In recent years, AIGC technology has demonstrated remarkable progress in video generation, giving rise to a variety of powerful text-to-video generative models (Hong et al. 2022; Yang et al. 2024; Zheng et al. 2024; Chen et al. 2024; Li et al. 2024b), such as Sora, Runway Gen-2, and Pika. These models have shown significant potential in producing longer-duration videos with higher quality and improved naturalness. Despite substantial advances, generated videos from these models still frequently suffer from issues including unnaturalness, consistency errors, and poor alignment

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

† Equal contributions, ✉ Corresponding authors, ◆: Project lead. This work was supported in part by National Natural Science Foundation of China (No. 62372016) and Shenzhen Science and Technology Program (No. JCYJ20241202125904007, No. SYSPG20241211173440004).

with human preferences, significantly hindering their practical application. Consequently, establishing reliable evaluation approaches for AIGC-generated videos is of crucial importance. Such evaluation methods can not only facilitate fine-grained manipulation of generated contents, but also serve as a robust basis for Reinforcement Learning with Human Feedback (RLHF), guiding models to more closely match user expectations.

A key challenge in applying RLHF to video generation is designing effective evaluation models. Existing VLM-based AIGC video evaluation methods (He et al. 2024; Xu et al. 2024) primarily rely on supervised finetuning (SFT), forcefully training the large models to regress video quality scores or directly judge human preferences. This approach suffers from three main drawbacks. **First**, it significantly diminishes the visual perception and general reasoning abilities of a general agent, reducing it to merely a scoring specialist. However, since different human annotators often exhibit biases when scoring the same video, simply regressing scores can in some sense be meaningless. We would prefer to inspire the model’s intrinsic potential for better understanding of AIGC video quality by teaching it scoring and preference comparison tasks. **Second**, existing methods (Wang et al. 2025b,a) typically require massive amounts of training data and continual construction of new benchmarks to keep pace with the rapidly evolving AIGC video generation methods. For instance, VisionReward (Xu et al. 2024) employed 80k visual question-answering annotations along with 2k preference comparisons to simulate human preferences, and VideoAlign (Liu et al. 2025b) even constructed a 182k human-labeled preference training samples, consuming enormous human and material resources. However, despite the variety of video generation models available, the produced videos often share common visual characteristics. This strongly motivates the need for an AIGC video evaluation method that achieves sufficient generalization capability with minimal training data. **Third**, there is no effective interaction between existing visual quality and generation models, as generation and understanding are mutually decoupled. This leads to the understanding model cannot obtain dynamic enhancement during the optimization process of the generation model, nor can it achieve a balance between the generalization and targeted capability.

To achieve this demand, we resort to the Group Rel-

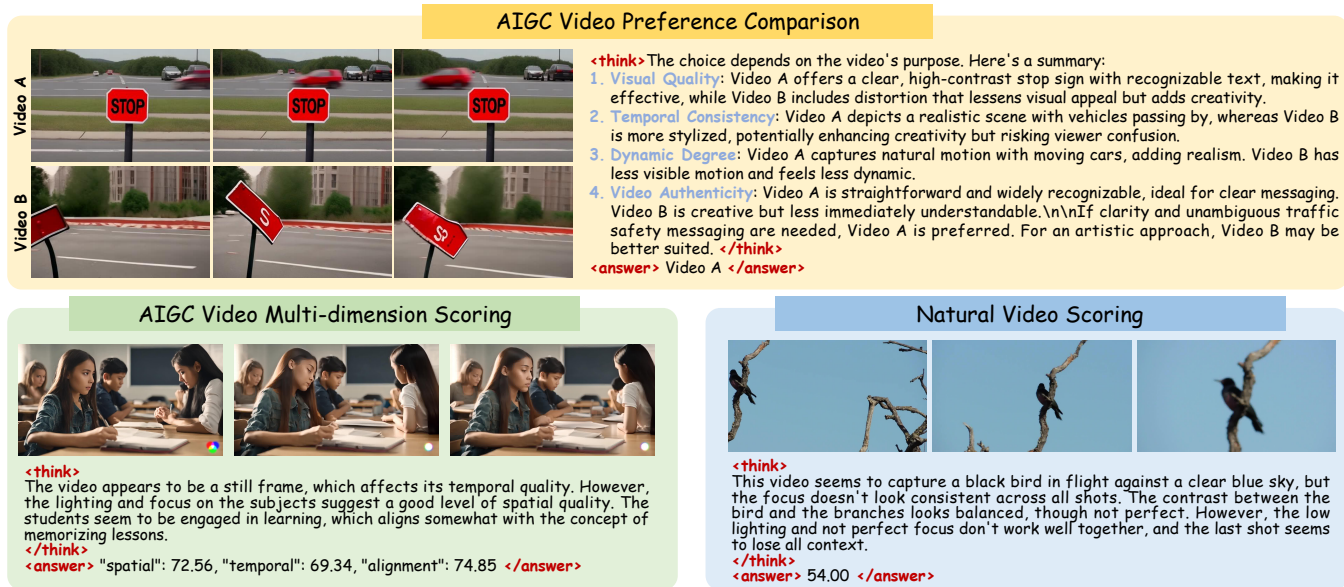


Figure 1: We propose a reasoning-style vision-language model VQ-Insight, which accurately performs AIGC video preference comparison, AIGC video multi-dimension scoring, and natural video scoring, accompanied by detailed and reasonable reasoning processes. Our VQ-Insight can be applied to post-training of video generation models and zero-shot content repairing.

ative Policy Optimization (GRPO) (Guo et al. 2025). As an outcome-driven reinforcement learning method, GRPO eliminates the need for an extra critic model and explicit reasoning processes during training, reducing dependence on human-labeled data and enhancing generalization. Although widely used in various vision tasks (Shen et al. 2025; Feng et al. 2025), GRPO has two key issues in AIGC video evaluation: limited multi-dimension analysis and poor temporal information handling. Specifically, we propose a GRPO-based AIGC video quality understanding model. Through image scoring warm-up, VLMs gain preliminary understanding of image quality. By incorporating temporal modeling rewards and task-specific rewards, VLM is encouraged to acquire general video scoring and preference comparison capabilities while capturing temporal cues. Furthermore, we conduct alternating optimization between specific video generation models and understanding models to foster mutual promotion. Fig. 1 shows the application scenarios of our VQ-Insight. Our contributions are summarized as follows.

- (1) We propose VQ-Insight, a reasoning-style VLM for understanding AIGC video quality. With limited data, VQ-Insight can effectively simulate human preferences and fine-grained scoring, providing feedback for video generation.
- (2) We propose a progressive video quality learning framework, which integrates image scoring warm-up, general task-specific temporal learning, and untied finetuning of the video generation and understanding. It enables the model to progressively move from image quality understanding to temporal perception, ultimately enhancing preference accuracy for specific video generation models.
- (3) We design a multi-dimension scoring reward and preference comparison reward, complemented by a temporal modeling and length control reward to effectively enhance

the model's capability in temporal perception.

- (4) Extensive experiments show that our approach outperforms state-of-the-art methods across AIGC video preference comparison, multi-dimension scoring and even natural video scoring. Additionally, our method can be applied to alignment and editing tasks of video generation models.

## Related Works

### VLM-based Video Quality Understanding

VLM-based quality assessment approaches (You et al. 2024, 2025; Wu et al. 2024b; Li et al. 2025; Zhao et al. 2025) combine both reasoning capabilities of LLMs and its powerful score regression abilities, achieving great success. In the field of video quality assessment, VQA-Scorer (Jia et al. 2024) introduced the SlowFast-R50 (Feichtenhofer et al. 2019) encoder to enhance motion capturing capabilities, and applied instruction tuning to the MLLM to guide the model to focus more on the description of low-level visual cues. In particular, evaluating AI-generated videos generally requires more complex fine-grained analyses. For instance, VideoScore (He et al. 2024) enabled automatic video quality assessment by training a VLM on the large-scale, multi-aspect human-annotated dataset VideoFeedback. VisionReward (Xu et al. 2024) employed a hierarchical visual assessment framework and multi-dimension consistent preference learning to capture fine-grained human preferences for both image and video generation. UnifiedReward (Wang et al. 2025b) introduced a unified preference learning framework to enable joint pairwise ranking and pointwise scoring for multimodal generation and understanding.

## Preference Learning in Generative Models

Recently, an increasing number of studies (Wang et al. 2024) have explored aligning generative models with human preferences using methods such as DPO (Rafailov et al. 2023) and GRPO (Guo et al. 2025). For example, VideoDPO (Liu et al. 2025c) proposed the OmniScore video evaluation pipeline, which is used to generate win-lose pairs for subsequent direct preference optimization. VADER (Prabhudesai et al. 2024) utilized an online reward model to finetune the video generation model. VideoAlign (Liu et al. 2025b) constructed a large-scale, multi-dimension human preference annotation dataset and proposed the VideoReward model, extending existing DiffusionDPO (Wallace et al. 2024) to flow-based models for more fine-grained alignment. Flow-GRPO (Liu et al. 2025a) integrated online RL method into flow matching by leveraging ODE-to-SDE conversion and a denoising reduction strategy. However, these methods are still limited by the accuracy and generalization issues of the reward model, making them inaccurate in assisting diffusion models to learn human preferences.

## Methodology

### Preliminaries

Group Relative Policy Optimization (GRPO) is a recent advanced reinforcement learning framework for LLMs and VLMs. Distinct from proximal policy optimization like PPO which rely on a dedicated value-critic to estimate policy quality, GRPO eliminates an explicit critic by leveraging relative comparison among grouped responses. Concretely, for each query  $q$ , GRPO samples a set of  $N$  candidate outputs  $\{o_1, o_2, \dots, o_N\}$  from the current or previous policy  $\pi_{\theta_{\text{old}}}$ . Each output receives rewards  $\{r_1, r_2, \dots, r_N\}$  based on task-specific functions, and GRPO computes the normalized advantage for each response as their reward’s deviation from the group mean, scaled by the standard deviation:

$$\hat{A}_i = \frac{r_i - \text{mean}(\{r_1, r_2, \dots, r_N\})}{\text{std}(\{r_1, r_2, \dots, r_N\})}. \quad (1)$$

After obtaining the relative advantage  $\hat{A}_i$ , GRPO computes the likelihood ratio of each response under the new policy  $\pi_{\theta_{\text{new}}}$  and the old policy  $\pi_{\theta_{\text{old}}}$ , and clips this ratio into the interval  $[1 - \delta, 1 + \delta]$  to prevent overly large updates and unstable training. The policy is then updated to increase the likelihood of responses with higher relative advantage, while penalizing large deviations from a given reference policy via a KL divergence term. The objective can be expressed as:

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}_{[q \sim Q, o_i \sim \pi_{\theta_{\text{old}}}(o|q)]} \left\{ \min \left[ \rho_i \hat{A}_i, \right. \right. \quad (2)$$
$$\left. \left. \text{clip}(\rho_i, 1 - \delta, 1 + \delta) \hat{A}_i \right] - \beta \cdot \mathbb{D}_{\text{KL}}[\pi_{\theta} \parallel \pi_{\text{ref}}] \right\},$$

where  $\rho_i = \pi_{\theta_{\text{new}}}(o_i | q) / \pi_{\theta_{\text{old}}}(o_i | q)$  denotes the update ratio between new and old policy for response  $o_i$ ,  $\delta$  controls update stability, and  $\beta$  weights the KL-regularization relative to the reference model.  $Q$  denotes the question set.

### Overview of Our VQ-Insight

**Motivation:** Previous approaches using VLMs for video quality assessment either relied on scoring labels without reasoning processes or required explicitly constructed

Chain-of-Thought (CoT) data using powerful foundational models (GPT-4o), thus consuming substantial resources. However, forceful SFT or cold-start tend to impair the general understanding capability of these models. In contrast, we hypothesize that as a heuristic and self-discovery training approach, reinforcement learning can be employed throughout the entire optimization process. Furthermore, we observe that inevitable annotator biases across diverse data sources and scoring tasks. Blindly mixing all biased data together during training can significantly harm the model’s performance on each individual data domain. Thus, designing a training pipeline that progressively transfers biased knowledge from simple to complicated, general to specific scenarios emerges as the key challenge to address for training robust VLMs for video quality assessment.

We propose a curriculum-style progressive visual reinforcement learning strategy consisting of three stages: image scoring warm-up, general task-specific temporal learning, and united finetuning of generation and understanding. At each stage, we flexibly handle different tasks and data by employing tailored reward functions and training strategies, guiding the model to progressively focus on spatial relationships, temporal modeling, and text-video alignment.

### Image Scoring Warm-up

Image quality understanding forms the foundational basis for video quality comprehension. At this stage, our main goal is to help the model learn the reasoning and response formats while improving its spatial understanding of images. As shown in Fig. 2, we warm up a general pretrained VLM using an image scoring task to obtain the initial policy model. Specifically, we use two distinct reward functions: a format reward and an image scoring reward. The format reward encourages the model to explicitly provide the reasoning between the `<think>` and `</think>` tags, and the numerical quality score between the `<answer>` and `</answer>` tags. Meanwhile, the image scoring reward is implemented as a continuous absolute norm ( $\ell_1$ -norm) to guide accurate score prediction. Given predicted score of the  $i$ -th response  $s_i^{\text{pred}}$  and its ground truth  $s^{\text{gt}}$ , the reward value  $r_i^{\text{score}}$  is calculated as follows.

$$r_i^{\text{score}} = 1 - \|s_i^{\text{pred}} - s^{\text{gt}}\|_1. \quad (3)$$

After warming up, the policy model can better understand image structures and visual quality, shifting its descriptive focus from high-level semantic cues towards low-level details, thus facilitating subsequent task-specific optimization.

### General Task-Specific Temporal Learning

After gaining preliminary image understanding capability, we further require our VQ-Insight to move towards temporal modeling and task-specific learning. As plotted in Fig. 2, in the following sections, we mainly focus on three tasks: natural video quality assessment, AIGC multi-dimension scoring, and AIGC preference comparison.

**Temporal Modeling Reward:** In temporal learning, to encourage the model to assess video quality based on temporal cues, we consider using random shuffling operations

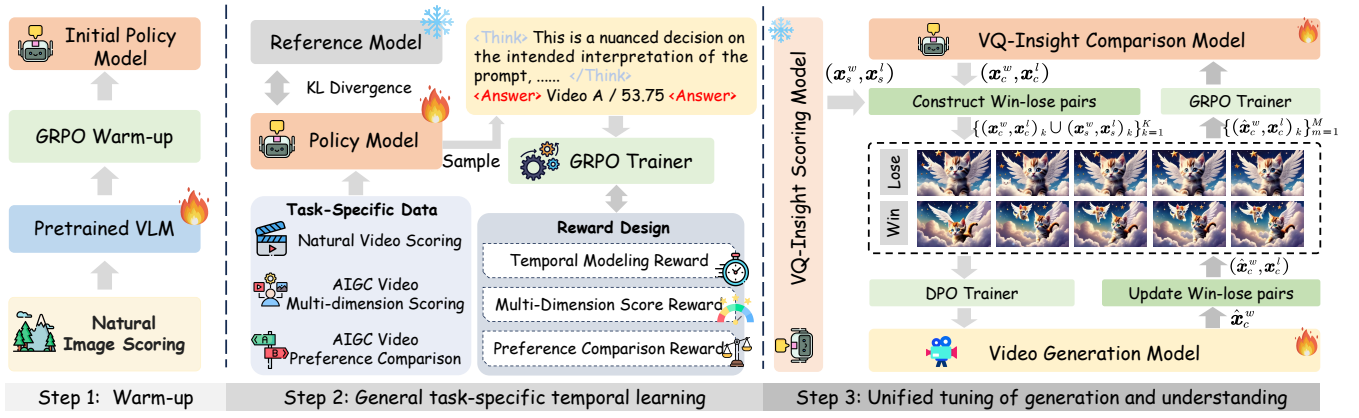


Figure 2: Illustration of the proposed VQ-Insight and our progressive visual reinforcement learning framework. In stage 1, we use the image scoring task and GRPO to warm up the pre-trained VLM; in stage 2, we employ temporal modeling rewards and task-specific rewards to enable the policy model to learn general tasks and temporal patterns; in stage 3, we jointly and alternately finetune the VQ-Insight comparison model and the video generation model, achieving a mutually beneficial effect.

to evaluate whether the model possesses sufficient temporal awareness. Specifically, given a question  $q$ , we first convert the instruction into text tokens and the video into visual tokens, then concat them as input to the VLM, obtaining a set of answers  $o^{seq}$ . Meanwhile, we randomly shuffle the tokens derived from video frames, feed these shuffled tokens into the policy model, and obtain another set of candidate answers  $o^{rand}$ . Assuming that the model’s predictions after shuffling should significantly differ from the ground truth, we compute the probability  $w^{seq}$  of giving the correct answer with sequentially ordered tokens and the probability  $w^{rand}$  with randomly shuffled tokens. If  $w_i^{seq}$  of the  $i$ -th response is significantly greater than  $w_i^{rand}$ , we can conclude the model has successfully captured temporal information and can assign it a reward value  $r_i^{temp}$  as compensation. Formally, the process is as follows.

$$r_i^{temp} = \alpha \text{ if } w_i^{seq} > \mu \cdot w_i^{rand}, \text{ else } 0, \quad (4)$$

where  $\alpha$  and  $\mu$  are set to 0.3 and 0.8, respectively.

**Length Control Reward:** To control the completion length and avoid overthinking or underthinking, we introduce a length control reward. If the length of the model’s answer  $o_i$  falls within a predefined interval  $[l_{min}, l_{max}]$ , we grant an additional reward  $r_i^{len}$  for the  $i$ -th response.

$$r_i^{len} = \gamma \text{ if } l_{min} < \text{len}(o_i) < l_{max}, \text{ else } 0, \quad (5)$$

where  $\gamma$  is set to 0.1.  $l_{min}$  and  $l_{max}$  are empirically set to 320 and 512. We observe that introducing the length control reward lead to an “aha moment” in the large model’s understanding of video quality, and the model paid greater attention to temporal modeling during reasoning.

**Multi-Dimension Scoring Reward:** Unlike scoring images, video quality assessment often requires consideration from multiple aspects. Following UGVQ (Zhang et al. 2024), we mainly focus on three aspects: spatial quality, temporal quality, and text-video alignment, each represented

by a Mean Opinion Score (MOS). Given a query  $q$ , as shown in Fig. 2, we prompt the VLM to directly output a set of scores with  $M$  dimensions  $\{v_{i,j}^{pred}\}_{j=1}^M$  in the  $i$ -th response. Similar to image scoring warm-up stage, we also adopt the  $\ell_1$  norm to fit scores in each dimension separately.

$$r_i^{multd} = 1 - \sum_{j=1}^M \lambda_j \|v_{i,j}^{pred} - v_j^{gt}\|_1, \quad (6)$$

where  $v_j^{gt}$  denotes the ground truth score of the  $j$ -th dimension.  $\lambda_j$  is used to balance the weights of different dimensions. Note that, for natural VQA, we still adopt a single-dimension reward due to the limitations of existing datasets.

**Preference Comparison Reward:** Rather than obtaining an absolute score, it is often more meaningful to directly provide the relative ranking between two generated results. To this end, we introduce the preference comparison reward. Specifically, given two input videos, we first convert them into visual tokens separately and feed them into the VLM to produce a group of preference choice  $c_i^{pred}$ . If the chosen answer  $c_i^{pred}$  matches the ground truth  $c_i^{gt}$ , we set  $r_i^{comp}$  to 1. To enhance the model’s multimodal scoring capability while preserving its general understanding ability, we include additional AIGC visual question answering data as an auxiliary task. Given an input video, the VLM is sequentially asked questions such as “Is the motion pattern in this video reasonable?” The large model is required only to answer “yes” or “no”. If the model answers correctly, we assign a reward value  $r_i^{comp}$  of 1. The rewards for these two tasks can be expressed uniformly as follows.

$$r_i^{comp} = 1 \text{ if } c_i^{pred} = c_i^{gt}, \text{ else } 0, \quad (7)$$

where  $c_i^{gt}$  is “video A” or “video B” in the preference comparison task, and “yes” or “no” in the VQA task. Ultimately, we combine the temporal-aware rewards and task-specific rule-based rewards to jointly optimize our VQ-Insight with robust and general video understanding capabilities.

## United Tuning of Generation and Understanding

**Motivation:** To apply VQ-Inspight to the video generation model, a common practice is to employ DPO to align generated outputs with human preferences. However, since DPO is an offline RL method, its preference dataset cannot dynamically evolve alongside the generation model during optimization, causing it to quickly reach a performance ceiling. Furthermore, video generation models finetuned by DPO typically possess stronger generative capabilities, further widening the gap between the newly generated positive samples and the original samples, thus enabling updates to our preference dataset. This updated preference dataset has the potential to enhance VQ-Inspight’s understanding capabilities, allowing it to better focus on preference comparisons specific to certain generative models.

**Constructing win-lose pairs:** Specifically, we first use the video generation model  $\mathcal{G}_\theta$  to produce  $N$  candidate videos  $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$ . Then, we form pairs of these videos and employ the VQ-Inspight comparison model  $\mathcal{D}_c$  in the Stage 2 to conduct  $\binom{N}{2}$  preference estimations. Finally, by counting the number of times each candidate is preferred, we identify the most- and least-chosen videos within the set, thus obtaining win-lose pairs.

$$\mathbf{x}_c^w, \mathbf{x}_c^l = \underset{\mathbf{x}_i, \mathbf{x}_j \in \mathcal{X}}{\operatorname{argmax}} \mathcal{D}_c(\mathbf{x}_i, \mathbf{x}_j), \underset{\mathbf{x}_i, \mathbf{x}_j \in \mathcal{X}}{\operatorname{argmin}} \mathcal{D}_c(\mathbf{x}_i, \mathbf{x}_j), \quad (8)$$

where  $\mathbf{x}_c^w, \mathbf{x}_c^l$  respectively denote the win and lose sample generated by the comparison model. Furthermore, to enhance the generative model’s performance in terms of spatial quality, temporal quality, and video–text consistency, we average the fine-grained scores from VQ-Inspight’s AIGC scoring model as a key criterion for constructing win–lose pairs. Considering that natural video evaluation models tend to focus more on visual quality and exhibit stronger degradation-awareness, we additionally incorporate scores from the VQ-Inspight’s natural video QA model together with those from the AIGC scoring model to compute the final score  $S_{total}$ . Thus, we also leverage the scoring model of VQ-Inspight to generate the win–lose pairs  $(\mathbf{x}_s^w, \mathbf{x}_s^l)$ .

$$S_{total} = \mathcal{D}_a(\mathbf{x}_i) + \beta \cdot \mathcal{D}_n(\mathbf{x}_i), \quad (9)$$

$$\mathbf{x}_s^w, \mathbf{x}_s^l = \underset{\mathbf{x}_i \in \mathcal{X}}{\operatorname{argmax}}(S_{total}), \underset{\mathbf{x}_i \in \mathcal{X}}{\operatorname{argmin}}(S_{total}), \quad (10)$$

where  $\beta$  is set to 1. Finally, we combine the two sets of win–lose pairs to form the complete preference dataset  $\mathcal{C}$ .

$$\mathcal{C} = \{(\mathbf{x}_c^w, \mathbf{x}_c^l)_k \cup (\mathbf{x}_s^w, \mathbf{x}_s^l)_k\}_{k=1}^K. \quad (11)$$

**Unified Optimization:** Following the approach of DiffusionDPO (Wallace et al. 2024), we can optimize  $\mathcal{G}_\theta$  by comparing the noise prediction differences between the finetuned model and the reference model. After obtaining the finetuned video generation model  $\mathcal{G}'_\theta$ , we generate  $N$  additional samples via this updated model. Then, we utilize VQ-Inspight comparison model  $\mathcal{D}_c$  to select the best candidate  $\hat{\mathbf{x}}_c^w$  from these newly generated samples, pairing it with the previous lose sample  $\mathbf{x}_c^l$  to form a new preference

Dataset	GenAI		MonetBench	
	tau	diff	tau	diff
VQAScore (Lin et al. 2024)	46.96	69.14	54.00	59.39
VideoScore (He et al. 2024)	47.43	70.50	49.10	54.90
VisionReward (Xu et al. 2024)	46.68	68.86	<b>59.40</b>	<b>72.44</b>
VideoReward (Liu et al. 2025b)	45.84	69.00	53.60	59.88
Qwen-SFT (Bai et al. 2025)	40.69	59.43	59.20	72.07
Q-Inspight (Li et al. 2025)	47.52	70.43	49.60	60.37
UnifiedReward-Think (Wang et al. 2025a)	<b>49.67</b>	<b>74.42</b>	52.10	62.56
VQ-Inspight (Ours)	<b>51.26</b>	<b>76.57</b>	<b>62.10</b>	<b>75.61</b>

Table 1: Preference comparison between our VQ-Inspight and other competitive methods on GenAI and MonetBench.

set  $\{(\hat{\mathbf{x}}_c^w, \mathbf{x}_c^l)\}$ . As shown in Fig. 2, we combine the newly generated  $\{(\hat{\mathbf{x}}_c^w, \mathbf{x}_c^l)_k\}_{m=1}^M$  with the original preference data and continue to finetune  $\mathcal{D}_c$  via the training strategy in Stage 2, resulting in a preference model  $\mathcal{D}'_c$  specialized for  $\mathcal{G}_\theta$ . Finally, the updated model  $\mathcal{D}'_c$  and the original  $\mathcal{D}_a, \mathcal{D}_n$  are used to generate a new preference set  $\mathcal{C}'$  for next round of video DPO, resulting in a better generation model  $\mathcal{G}''_\theta$ .

## Experimental Results

### Experimental Setup

**Dataset and Metrics:** We use 7k images in KonIQ (Hosu et al. 2020) for image warm-up. In Stage 2, only 2k comparison videos (Xu et al. 2024) and 1k VQA data (Xu et al. 2024) are used to train for preference comparison. LGVQ (Zhang et al. 2024) and LSVQ (Ying et al. 2021) are used for AIGC multi-dimension scoring and natural video scoring. In Stage 3, we choose T2V-Turbo (Li et al. 2024a) as generation models and select 5k prompts from Vidprom (Wang and Yang 2024) for united finetuning. To evaluate the preference comparison capability, Gen-AI (Jiang et al. 2024) and MonetBench (Xu et al. 2024) are used, with preference selection accuracy used as the evaluation metric. Moreover, LGVQ is used to assess the performance in multi-dimension scoring. LSVQ-Test, LSVQ-1080p, LIVE-VQC (Sinno and Bovik 2018), and KonViD-1k (Hosu et al. 2017) are adopted to evaluate the model’s natural video quality scoring ability, using PLCC, SRCC, and KRCC as metrics. For video generation tasks, we use VBench (Huang et al. 2024) to evaluate the generation quality of the finetuned models.

**Implementation Details:** Qwen-2.5-VL-7B-Instruct (Bai et al. 2025) is used as our pretrained VLM. The generation number  $N$  and the weight of KL penalty  $\beta$  in the GRPO trainer are 8 and 0.001. The model is trained for 3 epochs on 8 NVIDIA A100 80G GPUs, with a learning rate of  $1 \times 10^{-6}$ .

### AIGC Video Preference Comparison

To evaluate the performance of our VQ-Inspight in preference comparison, we selected classic methods VQAScore (Jia et al. 2024), SFT-based VLM methods VideoScore (He et al. 2024), VisionReward (Xu et al. 2024), VideoReward (Liu et al. 2025b), and Qwen-SFT (Bai et al. 2025), and RL-based VLM methods like Q-Inspight (Li et al. 2025) and

Method	Spatial Quality			Temporal Quality			Text-Video Alignment		
	SRCC	KRCC	PLCC	SRCC	KRCC	PLCC	SRCC	KRCC	PLCC
CLIP-IQA (Wang, Chan, and Loy 2023)	0.684	0.502	0.709	-	-	-	-	-	-
FastVQA (Wu et al. 2022)	-	-	-	0.849	0.672	0.878	-	-	-
CLIPScore (Hessel et al. 2021)	-	-	-	-	-	-	0.446	0.301	0.453
UGVQ (Zhang et al. 2024)	<u>0.764</u>	<u>0.571</u>	<u>0.793</u>	<u>0.894</u>	<u>0.703</u>	<u>0.910</u>	0.545	0.391	0.569
UnifiedReward-Think (Wang et al. 2025a)	0.580	0.432	0.594	0.466	0.330	0.500	0.589	0.433	0.589
Qwen-SFT (Bai et al. 2025)	0.687	0.520	0.735	0.723	0.539	0.750	<u>0.605</u>	<u>0.462</u>	<u>0.660</u>
VQ-Insight (Ours)	<b>0.823</b>	<b>0.640</b>	<b>0.844</b>	<b>0.911</b>	<b>0.744</b>	<b>0.927</b>	<b>0.825</b>	<b>0.652</b>	<b>0.836</b>

Table 2: SRCC, KRCC, PLCC Comparison between VQ-Insight and competitive methods across three different dimensions.

UnifiedReward-Think (Wang et al. 2025a). As reported on Tab. 1, our method surpasses existing SOTA approach VisionReward and UnifiedReward-Think on both “tau” and “diff” accuracy, which shows the effectiveness and strong generalization capability of our method. Meanwhile, although UnifiedReward-Think performs excellently on the GenAI dataset, its accuracy drops significantly on the out-of-domain MonetBench. Note that “tau” uses a tau-corrected result (Deutsch, Foster, and Freitag 2023) for preference accuracy, while “diff” excludes Tie cases. This is attributed to our refined reward design and progressive training.

### AIGC Video Multi-Dimension Scoring

To evaluate our VQ-Insight on fine-grained video quality assessment, we follow LGVQ setup and mainly consider three dimensions: spatial quality, temporal quality, and text-video alignment. For comparison, we select several metrics that are limited to a single dimension, such as CLIP-IQA (Wang, Chan, and Loy 2023), CLIPScore (Hessel et al. 2021), and FAST-VQA (Wu et al. 2022). Moreover, we include more comprehensive scorers such as UGVQ (Zhang et al. 2024), UnifiedReward-Think, and Qwen-SFT. As reported on Tab. 2, our method significantly outperforms the SOTA approaches UGVQ and Qwen-SFT across all dimensions. On the spatial quality dimension, our method surpasses UGVQ by 0.051 and 0.059 on PLCC and SRCC, respectively. Furthermore, on the text-video alignment dimension, our approach achieves an improvement of up to 0.2 over previous methods, demonstrating that our progressive reinforcement learning strategy effectively preserves the VLM’s general language understanding ability and world knowledge priors. Fig. 1 presents an example of VQ-Insight performing fine-grained scoring. Our VQ-Insight can comprehensively consider the spatial and temporal quality of the video while analyzing alignment based on the given prompt.

### Natural Video Scoring

In addition to AIGC video evaluation, our method can also be extended to natural video scoring. We conduct experiments on four datasets, namely LSVQ-Test, LSVQ-1080p, Live-VQC, and Konvid-1k (Hosu et al. 2017). The comparison baselines include the classic VQA model Fast-VQA (Wu et al. 2022), as well as VLM-based methods such as Q-Align (Wu et al. 2024b), Q-Instruct (Wu et al.

Model	Metric	LSVQ -Test	LSVQ -1080p	LIVE -VQC	KonViD -1k
Fast-VQA (Wu et al. 2022)	PLCC	0.878	0.810	0.815	0.857
Minimalist-VQA (Sun et al. 2024)	SRCC	0.874	0.765	0.769	0.859
Q-Align (Wu et al. 2024b)	PLCC	0.872	0.818	0.812	0.861
Q-Instruct (Wu et al. 2024a)	SRCC	0.880	0.769	0.765	0.859
VQA <sup>2</sup> (Jia et al. 2024)	PLCC	<b>0.882</b>	<b>0.833</b>	0.813	<u>0.876</u>
Q-Insight (Li et al. 2025)	SRCC	<b>0.883</b>	<u>0.758</u>	0.777	<u>0.865</u>
VQ-Insight (Ours)	PLCC	0.580	0.640	0.673	0.520
	SRCC	0.602	0.644	0.660	0.492
	PLCC	0.856	0.819	<u>0.823</u>	0.844
	SRCC	<u>0.882</u>	0.760	<u>0.776</u>	0.833
	PLCC	0.639	0.648	0.708	0.753
	SRCC	0.644	0.601	0.624	0.751
	PLCC	<u>0.876</u>	<u>0.823</u>	<b>0.835</b>	<b>0.884</b>
	SRCC	0.875	<b>0.786</b>	<b>0.790</b>	<b>0.875</b>

Table 3: PLCC and SRCC comparisons on the natural video scoring tasks between our VQ-Insight and other methods.

2024a), VQA<sup>2</sup> (Jia et al. 2024), and Q-Insight (Li et al. 2025). As reported on Tab. 3, our VQ-Insight achieves the best PLCC and SRCC on the LIVE-VQC and KonViD-1k datasets, demonstrating its strong generalization ability on out-of-domain data. On the in-domain dataset LSVQ-Test, the performance of VQ-Insight is comparable to that of the SOTA methods Q-Align and VQA<sup>2</sup>. This can be attributed to our progressive visual reinforcement learning strategy and temporal modeling reward used by our VQ-Insight. Furthermore, as shown in Fig. 1, the results produced by VQ-Insight effectively describe the content of the video and identify the imbalance between the theme and the environment.

### Ablation Studies

To validate the contributions of each component in our VQ-Insight, we design some variations and retrain them for the tasks of AIGC multi-dimension scoring and preference comparison by targeting image scoring warm-up (Warm-up), temporal modeling reward (TMR), length control reward (LCR), and unified finetuning (UF). Tab. 4 and Tab. 5 report our results. We observe that skipping the image scoring warm-up step and directly starting training from the pre-trained weights of Qwen2.5-VL results in noticeable degradation in both scoring and comparison performance. Re-



Figure 3: Video generation result comparisons between our method and other competitive methods.

Case	Warm-up	TMR	LCR	PLCC	KRCC	SRCC
(a)	✗	✓	✓	0.716	0.518	0.690
(b)	✓	✗	✓	0.787	0.590	0.761
(c)	✓	✓	✗	0.819	0.614	0.791
(d)	✓	✓	✓	0.869	0.679	0.853

Table 4: Ablation study on different components of our VQ-Insight for the AIGC multi-dimension scoring tasks.

Case	LCR	UF	GenAI		MonetBench	
			tau	diff	tau	diff
(e)	✗	✓	45.74	68.14	60.00	73.05
(f)	✓	✗	50.14	75.14	60.20	73.29
(g)	✓	✓	51.26	76.57	62.10	75.61

Table 5: Ablation study on different components of our VQ-Insight for the AIGC preference comparison tasks.

moving the TMR leads to a significant drop of 0.082 in PLCC for VQ-Insight’s performance on multi-dimension scoring tasks, demonstrating the critical role of TMR in helping the model capture motion patterns. Since performing temporal shuffle on two videos simultaneously can lead to preference label confusion, we do not validate the effectiveness of TWR on preference comparison. Additionally, we find that the length control reward (LCR) can better guide the model to produce detailed reasoning results while outputting more accurate scores or preference choices. Finally, when removing the unified finetuning strategy (case (f)), we observe that VQ-Insight’s tau accuracy on GenAI and MonetBench datasets decreased by 1.12 and 1.90, respectively. Moreover, as shown in Tab. 6, the performance of video generation models also experience certain degradation, which is caused by the inaccuracy of comparison models.

## Application

To validate that our VQ-Insight can effectively support generation tasks, we conduct DPO on the video generation model T2V-Turbo (Li et al. 2024a). We select 1w prompts from Vidprom (Wang and Yang 2024) and use the

Method	Overall Score	D.D.	Scene	S.R.	M.O.
Original	0.8095	0.5667	0.5134	0.3668	0.5122
VideoDPO	0.8167	0.5889	0.5080	0.3342	0.5215
Ours-w/o UF	<u>0.8171</u>	<u>0.6128</u>	<u>0.5151</u>	<u>0.3697</u>	<u>0.5289</u>
Ours	<b>0.8202</b>	<b>0.6306</b>	<b>0.5225</b>	<b>0.3913</b>	<b>0.5357</b>

Table 6: VBench score comparison between our method, VideoDPO and T2V-Turbo. D.D., S.R., and M.O. denote dynamic degree, spatial relationship, and multiple objects.

T2V-Turbo to produce 10 results for each prompt. Subsequently, VQ-Insight is used to select the preference set for DPO training. The performance of our method evaluated on VBench (Huang et al. 2024) is shown in Tab. 6. It can be observed that, compared to VideoDPO and baseline results, our method achieves significant gains in overall score. Specifically, in multiple objects and spatial relationship, VQ-Insight shows strong general understanding capabilities. Meanwhile, since our method uses AIGC fine-grained scoring and natural scoring model for temporal cues, it surpasses T2V-Turbo in terms of dynamic degree, reaching 0.0639. Furthermore, Fig. 3 shows that our method achieves noticeable gains over baseline and VideoDPO in the subject consistency, background richness, and color vividness.

## Conclusion

In this paper, we propose VQ-Insight, a novel reasoning-style vision-language model framework for AIGC video evaluation. By introducing a progressive learning scheme that combines image warm-up, temporal learning, and joint optimization with generation models, as well as task-specific rewards, our method achieves superior accuracy and generalization with limited data. Extensive experiments demonstrate that VQ-Insight consistently outperforms SOTA baselines across multiple video scoring and comparison benchmarks and can be effectively applied to both generation alignment and content editing. Looking ahead, our unified approach sets the stage for more dynamic, human-aligned video evaluation and optimization, highlighting the potential for reinforcement learning in the field.

## References

- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Chen, H.; Zhang, Y.; Cun, X.; Xia, M.; Wang, X.; Weng, C.; and Shan, Y. 2024. Videocrafter2: Overcoming data limitations for high-quality video diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7310–7320.
- Deutsch, D.; Foster, G.; and Freitag, M. 2023. Ties Matter: Meta-Evaluating Modern Metrics with Pairwise Accuracy and Tie Calibration. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 12914–12929.
- Feichtenhofer, C.; Fan, H.; Malik, J.; and He, K. 2019. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6202–6211.
- Feng, K.; Gong, K.; Li, B.; Guo, Z.; Wang, Y.; Peng, T.; Wang, B.; and Yue, X. 2025. Video-r1: Reinforcing video reasoning in mllms. *arXiv preprint arXiv:2503.21776*.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- He, X.; Jiang, D.; Zhang, G.; Ku, M.; Soni, A.; Siu, S.; Chen, H.; Chandra, A.; Jiang, Z.; Arulraj, A.; et al. 2024. VideoScore: Building Automatic Metrics to Simulate Fine-grained Human Feedback for Video Generation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 2105–2123.
- Hessel, J.; Holtzman, A.; Forbes, M.; Bras, R. L.; and Choi, Y. 2021. Clipscore: A reference-free evaluation metric for image captioning. *arXiv preprint arXiv:2104.08718*.
- Hong, W.; Ding, M.; Zheng, W.; Liu, X.; and Tang, J. 2022. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *arXiv preprint arXiv:2205.15868*.
- Hosu, V.; Hahn, F.; Jenadeleh, M.; Lin, H.; Men, H.; Szirányi, T.; Li, S.; and Saupé, D. 2017. The Konstanz natural video database (KoNViD-1k). In *2017 Ninth international conference on quality of multimedia experience (QoMEX)*, 1–6. IEEE.
- Hosu, V.; Lin, H.; Sziranyi, T.; and Saupé, D. 2020. KonIQ-10k: An ecologically valid database for deep learning of blind image quality assessment. *IEEE Transactions on Image Processing*, 29: 4041–4056.
- Huang, Z.; He, Y.; Yu, J.; Zhang, F.; Si, C.; Jiang, Y.; Zhang, Y.; Wu, T.; Jin, Q.; Chanpaisit, N.; et al. 2024. Vbench: Comprehensive benchmark suite for video generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21807–21818.
- Jia, Z.; Zhang, Z.; Qian, J.; Wu, H.; Sun, W.; Li, C.; Liu, X.; Lin, W.; Zhai, G.; and Min, X. 2024. VQA2: Visual Question Answering for Video Quality Assessment. *arXiv preprint arXiv:2411.03795*.
- Jiang, D.; Ku, M.; Li, T.; Ni, Y.; Sun, S.; Fan, R.; and Chen, W. 2024. Genai arena: An open evaluation platform for generative models. *Advances in Neural Information Processing Systems*, 37: 79889–79908.
- Li, J.; Feng, W.; Fu, T.-J.; Wang, X.; Basu, S.; Chen, W.; and Wang, W. Y. 2024a. T2v-turbo: Breaking the quality bottleneck of video consistency model with mixed reward feedback. *arXiv preprint arXiv:2405.18750*.
- Li, W.; Zhang, X.; Zhao, S.; Zhang, Y.; Li, J.; Zhang, L.; and Zhang, J. 2025. Q-insight: Understanding image quality via visual reinforcement learning. *arXiv preprint arXiv:2503.22679*.
- Li, W.; Zhao, S.; Mou, C.; Sheng, X.; Zhang, Z.; Wang, Q.; Li, J.; Zhang, L.; and Zhang, J. 2024b. OmniDrag: Enabling Motion Control for Omnidirectional Image-to-Video Generation. *arXiv preprint arXiv:2412.09623*.
- Lin, Z.; Pathak, D.; Li, B.; Li, J.; Xia, X.; Neubig, G.; Zhang, P.; and Ramanan, D. 2024. Evaluating text-to-visual generation with image-to-text generation. In *European Conference on Computer Vision*, 366–384. Springer.
- Liu, J.; Liu, G.; Liang, J.; Li, Y.; Liu, J.; Wang, X.; Wan, P.; Zhang, D.; and Ouyang, W. 2025a. Flow-grpo: Training flow matching models via online rl. *arXiv preprint arXiv:2505.05470*.
- Liu, J.; Liu, G.; Liang, J.; Yuan, Z.; Liu, X.; Zheng, M.; Wu, X.; Wang, Q.; Qin, W.; Xia, M.; et al. 2025b. Improving Video Generation with Human Feedback. *arXiv preprint arXiv:2501.13918*.
- Liu, R.; Wu, H.; Zheng, Z.; Wei, C.; He, Y.; Pi, R.; and Chen, Q. 2025c. Videodpo: Omni-preference alignment for video diffusion generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 8009–8019.
- Prabhudesai, M.; Mendonca, R.; Qin, Z.; Fragkiadaki, K.; and Pathak, D. 2024. Video diffusion alignment via reward gradients. *arXiv preprint arXiv:2407.08737*.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36: 53728–53741.
- Shen, H.; Liu, P.; Li, J.; Fang, C.; Ma, Y.; Liao, J.; Shen, Q.; Zhang, Z.; Zhao, K.; Zhang, Q.; et al. 2025. Vlm-r1: A stable and generalizable r1-style large vision-language model. *arXiv preprint arXiv:2504.07615*.
- Sinno, Z.; and Bovik, A. C. 2018. Large-scale study of perceptual video quality. *IEEE Transactions on Image Processing*, 28(2): 612–627.
- Sun, W.; Wen, W.; Min, X.; Lan, L.; Zhai, G.; and Ma, K. 2024. Analysis of video quality datasets via design of minimalistic video quality models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Wallace, B.; Dang, M.; Rafailov, R.; Zhou, L.; Lou, A.; Purushwalkam, S.; Ermon, S.; Xiong, C.; Joty, S.; and Naik, N. 2024. Diffusion model alignment using direct preference optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8228–8238.

- Wang, J.; Chan, K. C.; and Loy, C. C. 2023. Exploring clip for assessing the look and feel of images. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 2555–2563.
- Wang, W.; and Yang, Y. 2024. Vidprom: A million-scale real prompt-gallery dataset for text-to-video diffusion models. *arXiv preprint arXiv:2403.06098*.
- Wang, Y.; Li, Z.; Zang, Y.; Wang, C.; Lu, Q.; Jin, C.; and Wang, J. 2025a. Unified multimodal chain-of-thought reward model through reinforcement fine-tuning. *arXiv preprint arXiv:2505.03318*.
- Wang, Y.; Tan, Z.; Wang, J.; Yang, X.; Jin, C.; and Li, H. 2024. Lift: Leveraging human feedback for text-to-video model alignment. *arXiv preprint arXiv:2412.04814*.
- Wang, Y.; Zang, Y.; Li, H.; Jin, C.; and Wang, J. 2025b. Unified reward model for multimodal understanding and generation. *arXiv preprint arXiv:2503.05236*.
- Wu, H.; Chen, C.; Hou, J.; Liao, L.; Wang, A.; Sun, W.; Yan, Q.; and Lin, W. 2022. Fast-vqa: Efficient end-to-end video quality assessment with fragment sampling. In *European conference on computer vision*, 538–554. Springer.
- Wu, H.; Zhang, Z.; Zhang, E.; Chen, C.; Liao, L.; Wang, A.; Xu, K.; Li, C.; Hou, J.; Zhai, G.; et al. 2024a. Q-instruct: Improving low-level visual abilities for multi-modality foundation models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 25490–25500.
- Wu, H.; Zhang, Z.; Zhang, W.; Chen, C.; Liao, L.; Li, C.; Gao, Y.; Wang, A.; Zhang, E.; Sun, W.; et al. 2024b. Q-ALIGN: teaching LMMs for visual scoring via discrete text-defined levels. In *Proceedings of the 41st International Conference on Machine Learning*, 54015–54029.
- Xu, J.; Huang, Y.; Cheng, J.; Yang, Y.; Xu, J.; Wang, Y.; Duan, W.; Yang, S.; Jin, Q.; Li, S.; et al. 2024. Vision-reward: Fine-grained multi-dimensional human preference learning for image and video generation. *arXiv preprint arXiv:2412.21059*.
- Yang, Z.; Teng, J.; Zheng, W.; Ding, M.; Huang, S.; Xu, J.; Yang, Y.; Hong, W.; Zhang, X.; Feng, G.; et al. 2024. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*.
- Ying, Z.; Mandal, M.; Ghadiyaram, D.; and Bovik, A. 2021. Patch-vq: 'patching up' the video quality problem. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14019–14029.
- You, Z.; Cai, X.; Gu, J.; Xue, T.; and Dong, C. 2025. Teaching Large Language Models to Regress Accurate Image Quality Scores using Score Distribution. *arXiv preprint arXiv:2501.11561*.
- You, Z.; Li, Z.; Gu, J.; Yin, Z.; Xue, T.; and Dong, C. 2024. Depicting beyond scores: Advancing image quality assessment through multi-modal language models. In *European Conference on Computer Vision*, 259–276.
- Zhang, Z.; Li, X.; Sun, W.; Jia, J.; Min, X.; Zhang, Z.; Li, C.; Chen, Z.; Wang, P.; Ji, Z.; et al. 2024. Benchmarking AIGC Video Quality Assessment: A Dataset and Unified Model. *arXiv preprint arXiv:2407.21408*.
- Zhao, S.; Zhang, X.; Li, W.; Li, J.; Zhang, L.; Xue, T.; and Zhang, J. 2025. Reasoning as Representation: Rethinking Visual Reinforcement Learning in Image Quality Assessment. *arXiv preprint arXiv:2510.11369*.
- Zheng, Z.; Peng, X.; Yang, T.; Shen, C.; Li, S.; Liu, H.; Zhou, Y.; Li, T.; and You, Y. 2024. Open-sora: Democratizing efficient video production for all. *arXiv preprint arXiv:2412.20404*.