

# QuantVSR: Low-Bit Post-Training Quantization for Real-World Video Super-Resolution

Bowen Chai<sup>1\*</sup>, Zheng Chen<sup>1\*</sup>, Libo Zhu<sup>1</sup>,  
Wenbo Li<sup>2</sup>, Yong Guo<sup>3</sup>, Yulun Zhang<sup>1†</sup>

<sup>1</sup>Shanghai Jiao Tong University,

<sup>2</sup>Joy Future Academy,

<sup>3</sup>Max Planck Institute for Informatics

## Abstract

Diffusion models have shown superior performance in real-world video super-resolution (VSR). However, the slow processing speeds and heavy resource consumption of diffusion models hinder their practical application and deployment. Quantization offers a potential solution for compressing the VSR model. Nevertheless, quantizing VSR models is challenging due to their temporal characteristics and high fidelity requirements. To address these issues, we propose QuantVSR, a low-bit quantization model for real-world VSR. We propose a spatio-temporal complexity aware (STCA) mechanism, where we first utilize the calibration dataset to measure both spatial and temporal complexities for each layer. Based on these statistics, we allocate layer-specific ranks to the low-rank full-precision (FP) auxiliary branch. Subsequently, we jointly refine the FP and low-bit branches to achieve simultaneous optimization. In addition, we propose a learnable bias alignment (LBA) module to reduce the biased quantization errors. Extensive experiments on synthetic and real-world datasets demonstrate that our method obtains comparable performance with the FP model and significantly outperforms recent leading low-bit quantization methods.

**Code** — <https://github.com/bowenchai/QuantVSR>

## Introduction

Video super-resolution (VSR) is a crucial task that aims to reconstruct high-resolution (HR) videos from low-resolution (LR) inputs. Early VSR methods (Jo et al. 2018; Wang et al. 2019; Nah et al. 2019; Chan et al. 2021; Liang et al. 2024) assume synthetic degradations (*e.g.*, bicubic). However, models trained on such degraded data perform poorly in restoring practical videos. In real-world scenarios, videos often suffer from various complex and unknown degradations, which increase the difficulty of video restoration. Numerous methods are proposed (Yang et al. 2021; Pan et al. 2021; Xie et al. 2023) to address this challenge. Among them, generative adversarial networks (GANs) gather attention for their ability to restore fine details (Goodfellow et al. 2014; Lucas et al. 2019). Nevertheless, they still face challenges such as over-smoothing and training instability.

\*These authors contributed equally.

†Corresponding author: Yulun Zhang, yulun100@gmail.com  
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 1: Visual comparison among the full-precision VSR model, SVDQuant (Li et al. 2025a), and our QuantVSR.

A new generative model, the diffusion model (Ho, Jain, and Abbeel 2020), has recently shown promising results across various tasks, including VSR. Compared to GANs, diffusion models demonstrate superior generative performance with robust training stability. Diffusion-based VSR models (Blattmann et al. 2023; Yang et al. 2024; Zhou et al. 2024; Xie et al. 2025) leverage the generative capability while incorporating temporal modules, achieving remarkable visual quality. However, the high performance of diffusion models often comes with slow processing speeds and substantial resource consumption, which poses challenges for practical application and deployment on edge devices.

To tackle these issues, quantization stands out as an effective compression method (Nagel et al. 2021; Liu et al. 2025). Model quantization maps floating-point weights and/or activations to low-precision representations, typically as integers, thereby decreasing computational and memory costs. Prior studies have demonstrated the impressive results of quantization in image/video generation (Li et al. 2023a; Zhao et al. 2024; Li et al. 2025a) and image restoration (Zhu et al. 2025) tasks. However, in VSR, model quantization has rarely been explored, especially for the diffusion-based models. There are two critical difficulties in low-bit quantization for VSR models: **(1) Temporal consistency loss.** Model quantization introduces inconsistent errors across frames, resulting in a loss of temporal coherence in the final generated video. **(2) Complex data distribution.** VSR models embed temporal dynamics into latent features, resulting in more intricate activation distributions. This increases quantization challenges, as both spatial and temporal features must be considered to narrow the performance gap between full-precision (FP) and its quantized counterpart.

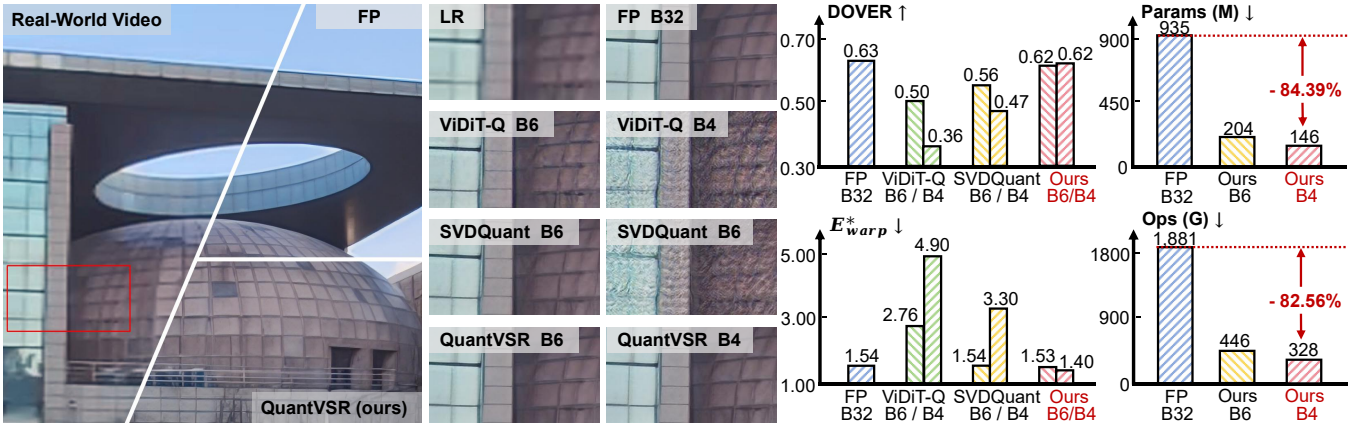


Figure 2: Performance comparisons on the real-world benchmark (*i.e.*, MVSR4x (Wang et al. 2023)) and compression ratio of our method. Bn represents n-bit quantization. Our QuantVSR compresses the model but retains performance comparable to the FP model, which surpasses existing quantization methods (*e.g.*, ViDiT-Q (Zhao et al. 2024) and SVDQuant (Li et al. 2025a)).

In this work, we propose QuantVSR, an effective low-bit quantization model for real-world VSR. We select MGLD-VSR (Yang et al. 2024) as our FP backbone due to its outstanding performance. We adopt several common techniques from low-bit quantization, such as the high-precision branch (Li et al. 2025a) and rotation (Ashkboos et al. 2024). Two novel designs are proposed to address the aforementioned issues. **First**, leveraging the unique temporal characteristics of VSR, we propose a spatio-temporal complexity aware (STCA) mechanism. A special low-rank branch is employed to skip the quantization process and is directly connected to the output. We allocate layer-specific ranks to the FP branch matrices based on the temporal and spatial complexities of the calibration dataset. Thus we achieve a trade-off between performance and computational efficiency. Subsequently, joint refinement of the FP and low-bit branches is applied to achieve simultaneous optimization. **Second**, we propose a learnable bias alignment (LBA) module to the quantization process, aiming to mitigate the severe biased quantization errors in low-bit quantization.

Comprehensive experiments (*i.e.*, Figs. 1 and 2) demonstrate that QuantVSR experiences minimal performance degradation even with 4-bit quantization. It outperforms recent leading quantization methods significantly. Compared to the FP model, 4-bit QuantVSR reduces parameter (Params) and operation (Ops) numbers by **84.39%** and **82.56%**, respectively (Fig. 2). Our contributions are:

- We propose QuantVSR, a low-bit quantized model for real-world VSR. To the best of our knowledge, this is the first work to explore low-bit quantization (*e.g.*, 4-bit and 6-bit) for diffusion-based VSR models.
- Our spatio-temporal complexity aware (STCA) mechanism effectively preserves the performance of the FP model in both the temporal and spatial domains. Learnable bias alignment (LBA) module is proposed to reduce the severe biased error in low-bit quantization.
- Extensive experiments on the synthetic and real-world datasets demonstrate the superior performance of our method in comparison to other quantization methods.

## Related Work

### Video Super-Resolution

Video super-resolution (VSR) aims to recover high-resolution (HR) videos from low-resolution (LR) inputs. With the advancement of deep learning, various methods (Jo et al. 2018; Chan et al. 2021, 2022a; Liang et al. 2024) have shown promising results, roughly categorized into recurrent-based (Liang et al. 2022; Shi et al. 2022) and sliding-window-based (Yi et al. 2019; Li et al. 2020) approaches. However, these methods often perform poorly on real-world videos due to their assumption of a fixed degradation process (Liu and Sun 2013; Xue et al. 2019). A growing focus has emerged on real-world VSR recently, aiming to address complex and unknown degradations. RealVSR (Yang et al. 2021) and MVSR4x (Wang et al. 2023) propose leveraging HR-LR paired data from real environments. While some approaches (Chan et al. 2022b; Xie et al. 2023) introduce diverse degradations for data augmentation to avoid labor-intensive data collection. Moreover, some works modify the model’s structure to enhance its adaptability in real-world VSR, such as kernel estimation based on the image formation model (Pan et al. 2021) and selective cross-attention modules (Xie et al. 2023). Although significant progress has been made, these methods still struggle to generate realistic details and fine textures.

### Diffusion Model

Diffusion models (Ho, Jain, and Abbeel 2020) demonstrate remarkable performance initially in image generation (Ramesh et al. 2022; Rombach et al. 2022) and have been extended to diverse vision tasks, including image restoration (Wang, Yu, and Zhang 2023; Fei et al. 2023), video generation (Blattmann et al. 2023; Chen et al. 2024; Zhang et al. 2023), and video super-resolution (VSR) (Zhou et al. 2024; Yang et al. 2024; He et al. 2024; Li et al. 2025b; Xie et al. 2025; Wang et al. 2025). With rich generative priors encoded in pretrained diffusion models, diffusion-based approaches can generate highly realistic details, driving a new wave of VSR methods. Among them, some approaches

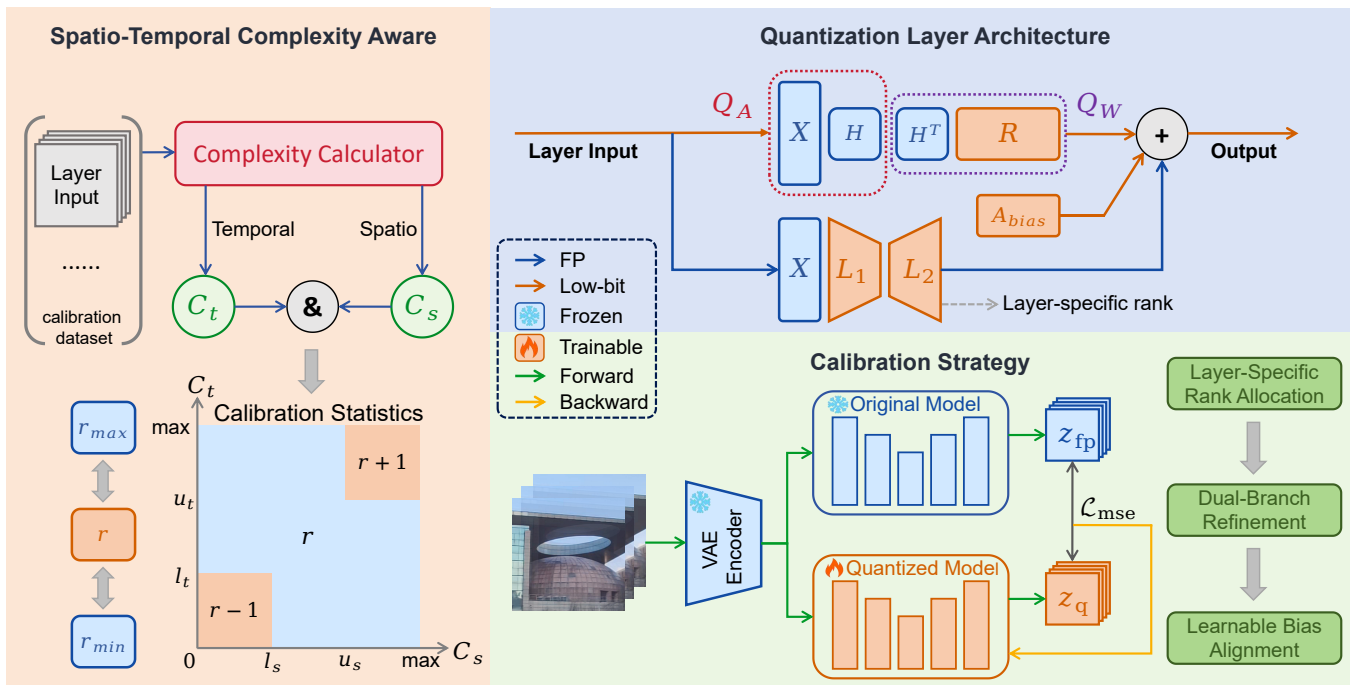


Figure 3: Overview of our QuantVSR. **First**, we analyze the temporal and spatial complexity distribution of the calibration dataset and leverage these statistics to allocate layer-specific ranks. **Next**, we jointly refine the two branches in spatio-temporal complexity aware mechanism. **Finally**, we train the learnable bias alignment module.

extend pretrained image generation models (Rombach et al. 2022) with temporal mechanisms, such as adding temporal layers (Zhou et al. 2024) or guiding the denoising process with optical flows of LR videos (Yang et al. 2024). Other methods (He et al. 2024; Xie et al. 2025) directly leverage video generation models (Zhang et al. 2023). Despite their success, diffusion-based VSR models are often associated with high computational and memory costs.

## Model Quantization

Quantization is an effective method for model compression, as evidenced by its widespread and successful applications in the field of large language models (Hu et al. 2022; Liu et al. 2023; Xiao et al. 2023; Shao et al. 2023; Huang et al. 2024). With the rapid evolution of diffusion models, researchers have concentrated on enhancing their efficiency by applying quantization techniques. PTQ4DM (Shang et al. 2023) pioneers the study of quantized diffusion models, achieving 8-bit quantization. Subsequent works have made breakthroughs with methods such as specialized calibration strategies (Li et al. 2023a; He et al. 2023b), sensitivity analysis (Yang et al. 2023), quantization-aware training (Li et al. 2023b), weight-only quantization (Sui et al. 2024), and outlier removal (Zhao et al. 2024). Recently, low-bit quantization has made significant progress. EfficientDM (He et al. 2023a) proposes a low-rank quantization fine-tuning strategy. While SVDQuant (Li et al. 2025a) employs a 16-bit parallel low-rank branch. Both EfficientDM and SVDQuant achieve 4-bit quantization with minimal performance loss. Furthermore, there are methods designed specifically for particular tasks and models, such as PassionSR (Zhu et al.

2025) for one-step diffusion-based super-resolution models, ViDiT-Q (Zhao et al. 2024) for image and video generation. However, studies on the quantization of VSR models remain limited. Existing quantization methods cause significant performance loss when applied to VSR models.

## Method

In this section, we introduce our low-bit quantized real-world VSR model, QuantVSR. First, we explain some necessary quantization preliminaries. Then, we present the overall framework of QuantVSR and the calibration process. Finally, we focus on two key designs proposed to mitigate the accuracy loss of quantizing VSR model: the spatio-temporal complexity aware (STCA) mechanism and the learnable bias alignment (LBA) module.

### Quantization Preliminary

Model quantization reduces storage and computational costs by mapping weights and/or activations to low-bit integers. The most common quantization process can be defined as:

$$x_{\text{int}} = \text{Clip} \left( \left\lfloor \frac{x}{s} \right\rfloor - z, l, u \right), \hat{x} = s \cdot x_{\text{int}} + z, \quad (1)$$

where  $x_{\text{int}}$  refers to the quantized integer and  $\hat{x}$  is the de-quantized floating-point value.  $s$  and  $z$  denote the scale factor and zero point, which are related to data distribution and quantization bit-width.  $\lfloor \cdot \rfloor$  is the round-to-nearest operator.  $\text{Clip}(\cdot, l, u)$  ensures values remain within the range  $[l, u]$ .

Quantization-aware training (QAT) and several recent post-training quantization (PTQ) methods rely on gradient back-propagation, while the rounding operation introduced by quantization is non-differentiable. To address this issue,

the straight-through estimator (STE) is widely adopted to approximate the gradients:

$$\frac{\partial Q(x)}{\partial x} \approx \begin{cases} 1 & \text{if } x \in [l, u], \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

## Overall Framework

QuantVSR is built on MGLD-VSR (Yang et al. 2024) for its superior performance in real-world video super-resolution (VSR). In the whole restoration process, the multi-step denoising of the diffusion model (*i.e.*, UNet) dominates the computational cost compared to the VAE decoding stage. Therefore, we focus primarily on quantizing the UNet structure, by replacing the original layers (*i.e.*, Linear, Conv2d, and Conv3d) with custom-designed counterparts.

The overall architecture of the quantization layer is illustrated in Fig. 3. Spatio-temporal complexity aware (STCA) mechanism proposes an enhanced skip connection (SC) structure, where a full-precision (FP) branch allows part of the computation to bypass quantization. For the low-bit branch, following previous work (Ashkboos et al. 2024), layer inputs and weights are first smoothed using a randomized Hadamard transform to avoid outliers dominating limited quantization levels. Subsequently, the main computations are performed using low-bit integers. Learnable bias alignment (LBA) is added to the result finally. The entire quantization layer is formulated as follows:

$$\mathbf{XW} = \underbrace{\mathbf{XL}_1\mathbf{L}_2}_{\text{FP, STCA}} + \underbrace{Q_A(\mathbf{XH})Q_W(\mathbf{H}^\top\mathbf{R})}_{\text{Low-Bit}} + \underbrace{\mathbf{A}_{\text{bias}}}_{\text{LBA}}, \quad (3)$$

where  $L_1 \in \mathbb{R}^{m \times r}$  and  $L_2 \in \mathbb{R}^{r \times n}$  are two low-rank matrices.  $\mathbf{R} = \mathbf{W} - \mathbf{L}_1\mathbf{L}_2$  is the residual.  $\mathbf{H}$  is the random Hadamard matrix.  $Q_W$  and  $Q_A$  refer to the weight and activation quantizer.  $\mathbf{A}_{\text{bias}}$  is the LBA module with the same dimensionality as the layer’s bias.

The quantization calibration process consists of three main stages. First, we analyze the spatio-temporal complexity distribution of the calibration dataset. Such calibration statistics allocate higher ranks to the FP branches in layers with high complexity. Second, we refine two low-rank matrices and residual, allowing the STCA structure to attain optimal performance across both FP and low-bit branches in each layer. Finally, as changes in quantized weights introduced by previous modules affect the biased quantization errors, the LBA module is trained in the last stage with all other parameters frozen. For fine-tuning and training in the final two stages, we adopt the mean squared error (MSE) between FP and quantized model outputs (*i.e.*,  $z_{fp}$  and  $z_q$ ).

## Spatio-Temporal Complexity Aware

Under extremely low-bit quantization settings, the broad dynamic range of floating-point weights and activations is only represented by a limited set of integer values. This severely compromises the restoration ability of the quantized model. A commonly adopted strategy to mitigate this issue incorporates a low-rank full-precision (FP) branch to preserve the original model’s performance (Li et al. 2025a; He et al. 2023a). However, this approach still has two main limitations: (1) The rank allocation strategy is suboptimal, incur-

ring unnecessary overhead. (2) The FP branch may degrade the performance of the low-bit branch. Based on these analyses, we propose STCA that leverages the characteristics of video inputs in VSR model, consisting of two key steps: layer-specific rank allocation and dual-branch refinement.

**Layer-Specific Rank Allocation.** The FP branch in Eq. (3) maintains baseline accuracy by constraining the product  $L_1L_2$  to approximate the original weight. A larger rank  $r$  theoretically preserves more information, but it introduces a computational cost of  $r \frac{m+n}{mn}$  that grows linearly with  $r$ . Consequently, gains in accuracy must be balanced against the accompanying increase in complexity.

A straightforward approach is to allocate layer-specific ranks. Layers that process more complex inputs encode richer information and therefore have a greater impact on the final output. Accordingly, we assign them higher ranks to preserve more of the original performance. For VSR models, this complexity spans both temporal and spatial dimensions. Specifically, for the layer input  $\mathbf{X} \in \mathbb{R}^{T \times C \times H \times W}$  ( $T$ : frames,  $C$ : channels,  $H$ : height, and  $W$ : width), we define its temporal complexity  $C_t$  as:

$$C_t = \frac{1}{T-1} \sum_{t=1}^{T-1} \frac{1}{CHW} \|\mathbf{X}_{t+1} - \mathbf{X}_t\|_2^2, \quad (4)$$

where  $\|\cdot\|_2$  denotes the Frobenius norm. This measures the inter-frame difference energy. Higher values reflect more intense motion between frames, which increases difficulty in restoration. Meanwhile, the spatial complexity is defined as:

$$C_s = \frac{1}{TC} \sum_{t=1}^T \sum_{c=1}^C \sigma_{h,w}(\mathbf{X}_{t,c}), \quad (5)$$

where  $\sigma(\cdot)$  denotes the standard deviation calculated over the spatial dimensions (*i.e.*, height and width). Since each channel represents a latent feature, its spatial variance reflects the complexity of that feature within each frame (*e.g.*, textures or edges). Features with higher spatial variance are generally more informative and, therefore, should be assigned higher ranks to preserve overall performance.

As shown in Fig. 3, the rank allocation process involves the following steps. First, we compute the distributions of temporal and spatial complexity for each layer using the calibration set, and define upper and lower thresholds for both (*i.e.*,  $u_s$ ,  $u_t$ ,  $l_s$ , and  $l_t$ ). Then, for each layer input in the calibration set, we assess its complexity: (1) If both temporal and spatial complexities exceed the upper threshold, the rank is incremented by one. (2) If both fall below the lower threshold, the rank is decremented by one. (3) Otherwise, the rank remains unchanged. Finally, the rank is constrained to the interval  $[r_{min}, r_{max}]$ . To ensure compatibility with GPU-parallel computation, the rank is then rounded to the nearest integer that is a multiple of 8.

**Dual-Branch Refinement.** The FP branch preserves the original model’s capabilities but alters the data distribution of the low-bit branch (*i.e.*,  $\mathbf{R}$  in Eq. (3)), affecting quantization difficulty. However, both branches contribute to the final output. This implies that enhancing the FP branch (*e.g.*, by allocating a higher rank) may degrade the performance of the low-bit branch, leading to suboptimal overall results.

Dataset	Bits	SC	LBA	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	DISTS $\downarrow$	CLIP-IQA $\uparrow$	MUSIQ $\uparrow$	NIQE $\downarrow$	MANIQA $\uparrow$	DOVER $\uparrow$	$E_{warp}^*$ $\downarrow$
SPMCS	W4A4	-	-	17.13	0.2595	0.6480	0.4111	0.2553	44.14	7.26	0.2957	0.0802	10.00
		-	$\checkmark$	21.38	0.4996	0.3666	0.2352	<u>0.5097</u>	64.74	<b>3.37</b>	<u>0.3183</u>	0.6515	3.27
		SVDSC	-	18.94	0.2820	0.5921	0.3756	<b>0.5378</b>	58.41	5.34	<b>0.3471</b>	0.4028	6.40
		SVDSC	$\checkmark$	22.58	0.5783	0.3296	0.2089	0.4647	63.32	3.55	0.3144	0.6673	1.90
		STCA	-	<u>22.75</u>	<u>0.6071</u>	<u>0.2914</u>	<u>0.1816</u>	0.4586	<u>65.49</u>	3.45	0.3155	<u>0.6886</u>	<b>1.74</b>
		STCA	$\checkmark$	<b>22.76</b>	<b>0.6075</b>	<b>0.2857</b>	<b>0.1747</b>	0.4553	<b>65.75</b>	<u>3.41</u>	0.3168	<b>0.6969</b>	<u>1.76</u>

Table 1: Ablation study on our key designs: STCA and LBA. Experiments are conducted on the synthetic dataset SPMCS (Yi et al. 2019) with W4A4 quantization. SC stands for the skip connection structure. SVDSC means that the skip connection structure follows the same settings as SVDQuant (Li et al. 2025a).

To ensure the effectiveness of both branches, we train two low-rank matrices  $L_1$  and  $L_2$  after layer-specific rank allocation, subject to the constraint  $R = W - L_1 L_2$ . However, training them from scratch converges slowly, which undermines the goal of post-training quantization (PTQ). Inspired by SVDQuant (Li et al. 2025a), we initialize the two matrices using singular value decomposition (SVD), providing a solid starting point for fine-tuning. With minimal training on the calibration set, the two branches achieve an optimal state, resulting in significantly improved performance.

### Learnable Bias Alignment

The quantization error is biased, which means the average output of the FP model and its quantized counterpart is different. This issue is a major source of the overall quantization error, especially under low-bit settings. Previous works have acknowledged this issue (Nagel et al. 2019, 2021), but have focused primarily on weight-only quantization. They use calibration statistics or batch normalization (BN) parameters to reduce the discrepancy. For low-bit quantization of both weights and activations, the error becomes considerably more complex and can be formalized as follows:

$$\mathbb{E}(\widehat{W}\widehat{X}) - \mathbb{E}(WX) = \Delta W \mathbb{E}(\widehat{X}) + W \mathbb{E}(\Delta X), \quad (6)$$

where  $\widehat{W} = W + \Delta W$  and  $\widehat{X} = X + \Delta X$  are quantized weights and activations. The weights and their quantization loss (*i.e.*,  $W$  and  $\Delta W$ ) are fixed. So the error term is a function of the activation distribution and is complicated by activation quantization error (*i.e.*,  $\Delta X$ ).

We propose a learnable bias alignment module to capture this error. As indicated by Eqs. (3) and (6), it is added to the output of the quantized layer and has the same dimensionality as the layer bias. The trainable parameters are relatively small compared with the whole model, allowing rapid convergence and preserving quantization efficiency. Moreover, during inference, this module can be fused into the layer bias, incurring negligible computational overhead.

## Experiments

### Experimental Settings

**Data Construction.** We sample intermediate input-output pairs from the full-precision (FP) UNet model on REDS30 (Nah et al. 2019) at fixed intervals during the denoising process (*i.e.*, 5 samples across 50 timesteps) (Nagel et al. 2021). This results in a calibration set of 1,800 pairs, where each input has a shape of  $5 \times 4 \times 64 \times 64$ . For eval-

Method	Bits	Params / M ( $\downarrow$ Ratio)	Ops / G ( $\downarrow$ Ratio)
MGLD-VSR	W32A32	935 ( $\downarrow$ 0%)	1,881 ( $\downarrow$ 0%)
QuantVSR	W8A8	263 ( $\downarrow$ 71.87%)	563 ( $\downarrow$ 70.07%)
	W6A6	204 ( $\downarrow$ 78.18%)	446 ( $\downarrow$ 76.29%)
	W4A4	146 ( $\downarrow$ 84.39%)	328 ( $\downarrow$ 82.56%)

Table 2: Params, Ops, and compression ratio (UNet only) of different quantization settings. Ops are computed with latent input size  $5 \times 4 \times 64 \times 64$ , corresponding to a 5-frame video.

uation, we apply both synthetic and real-world datasets. The synthetic datasets include REDS4 (Nah et al. 2019) and SPMCS (Yi et al. 2019), using multiple degradations (*i.e.*, random blur, resizing, noise, JPEG compression, and video compression). For real-world datasets, we apply MVS4x (Wang et al. 2023).

**Evaluation Metrics.** We apply both image quality assessment (IQA) and video quality assessment (VQA). The IQA include reference-based metrics: PSNR, SSIM (Wang et al. 2004), LPIPS (Zhang et al. 2018), and DISTS (Ding et al. 2020), and no-reference metrics: CLIP-IQA (Wang, Chan, and Loy 2023), MUSIQ (Ke et al. 2021), NIQE (Zhang, Zhang, and Bovik 2015), and MANIQA (Yang et al. 2022). For VQA, we adopt two metrics: DOVER (Wu et al. 2023) and  $E_{warp}^*$  (Lai et al. 2018).

**Implementation Details.** We quantize the weights and activations in the main components (*i.e.*, UNet) with low bit-widths (*e.g.*, 6 and 4 bits).  $r_{min}$  and  $r_{max}$  in STCA are set to 16 and 64 respectively. For both temporal and spatial complexity, the lower and upper thresholds (*i.e.*,  $l_s$ ,  $l_t$ ,  $u_s$ , and  $u_t$ ) are set to the 25th and 75th percentiles of the calibration dataset’s complexity distribution, respectively. The calibration training is performed on NVIDIA RTX A6000 GPU for 2 epochs. The learning rate is set as  $1 \times 10^{-3}$  and  $2 \times 10^{-4}$  during the first and second epoch, respectively.

**Compared Methods.** We select several representative and leading quantization methods: MaxMin (Jacob et al. 2018), Q-Diffusion (Li et al. 2023a), QuaRot (Ashkboos et al. 2024), ViDiT-Q (Zhao et al. 2024), and SVDQuant (Li et al. 2025a). We implement these methods on MGLD-VSR (Yang et al. 2024) based on their released code.

### Ablation Study

**Spatio-Temporal Complexity Aware (STCA).** To verify the effectiveness of STCA, we conduct experiments under three configurations: (1) without the skip connection (SC)

Datasets	Bits	Methods	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	DISTS $\downarrow$	CLIP-IQA $\uparrow$	MUSIQ $\uparrow$	NIQE $\downarrow$	MANIQA $\uparrow$	DOVER $\uparrow$	$E_{warp}^*$ $\downarrow$
REDS4	W32A32	MGLD-VSR	23.27	0.6180	0.2117	0.0890	0.3641	65.81	2.65	0.3254	0.6761	7.24
		MaxMin	22.53	0.5059	0.4150	0.2444	<b>0.3905</b>	55.68	2.77	0.2512	0.6044	8.84
	W6A6	Q-Diffusion	22.99	0.5589	0.3534	0.2036	0.3402	56.43	<u>2.62</u>	0.2389	0.6244	7.40
		QuaRot	<b>23.42</b>	<b>0.6213</b>	0.2388	<u>0.1179</u>	0.2731	60.66	2.80	0.2639	0.6562	<b>6.66</b>
		ViDiT-Q	22.68	0.5670	0.2759	0.1461	0.3346	62.86	2.66	0.2918	0.6267	8.75
		SVDQuant	23.26	0.6035	<u>0.2379</u>	0.1211	0.3289	<u>63.89</u>	<b>2.54</b>	<u>0.3018</u>	<u>0.6717</u>	<u>7.02</u>
		QuantVSR (ours)	<u>23.30</u>	<u>0.6167</u>	<b>0.2138</b>	<b>0.0921</b>	<u>0.3614</u>	<b>65.46</b>	2.63	<b>0.3230</b>	<b>0.6755</b>	7.12
	W4A4	MaxMin	16.18	0.1995	0.6720	0.4026	0.2821	46.52	5.90	0.2931	0.1451	52.27
		Q-Diffusion	19.99	0.3176	0.5279	0.3096	<b>0.5887</b>	52.81	3.69	<b>0.3415</b>	0.4936	19.63
		QuaRot	20.21	0.3182	0.5346	0.3179	<u>0.5561</u>	51.86	3.77	<u>0.3109</u>	0.5195	17.82
		ViDiT-Q	20.98	0.4118	<u>0.4436</u>	<u>0.2523</u>	0.4503	52.85	<u>3.03</u>	0.2986	0.5404	14.93
		SVDQuant	<u>21.19</u>	<u>0.4138</u>	0.4718	0.2908	0.5116	<u>54.09</u>	3.18	0.2699	<u>0.5865</u>	<u>12.46</u>
		QuantVSR (ours)	<b>23.31</b>	<b>0.6143</b>	<b>0.2286</b>	<b>0.1055</b>	0.3364	<b>64.19</b>	<b>2.61</b>	0.3046	<b>0.6822</b>	<b>6.88</b>
	SPMCS	W32A32	MGLD-VSR	22.81	0.6157	0.2807	0.1643	0.4500	65.89	3.56	0.3211	0.6957
MaxMin			21.79	0.4891	0.4395	0.2914	<b>0.5163</b>	61.08	3.44	0.2997	0.5978	2.45
W6A6		Q-Diffusion	22.37	0.5504	0.3937	0.2618	<u>0.4788</u>	60.19	<b>3.41</b>	0.2844	0.5912	1.77
		QuaRot	<b>22.95</b>	<b>0.6173</b>	<u>0.2991</u>	<u>0.1892</u>	0.3716	60.64	3.75	0.2839	0.6251	<b>1.66</b>
		ViDiT-Q	22.08	0.5652	0.3204	0.2158	0.4590	63.36	<u>3.41</u>	<u>0.3106</u>	0.6630	2.43
		SVDQuant	22.74	0.5959	0.3118	0.2013	0.4541	<u>63.48</u>	3.51	0.3066	<u>0.6724</u>	1.76
		QuantVSR (ours)	<u>22.80</u>	<u>0.6126</u>	<b>0.2833</b>	<b>0.1697</b>	0.4569	<b>65.73</b>	3.54	<b>0.3215</b>	<b>0.6848</b>	1.79
W4A4		MaxMin	17.51	0.2819	0.6239	0.3886	0.2823	47.07	6.55	0.3040	0.1331	8.92
		Q-Diffusion	19.85	0.3370	0.5275	0.3381	<b>0.5765</b>	<u>60.48</u>	4.67	<b>0.3568</b>	0.4996	4.82
		QuaRot	20.16	0.3505	0.5277	0.3397	<u>0.5691</u>	59.44	4.52	<u>0.3246</u>	0.5554	4.40
		ViDiT-Q	20.60	0.4276	<u>0.4600</u>	<u>0.2911</u>	0.4984	59.55	3.78	0.3075	0.5294	3.98
		SVDQuant	<u>21.17</u>	<u>0.4459</u>	0.4739	0.3126	0.5417	57.39	<u>3.74</u>	0.2897	<u>0.5627</u>	<u>2.77</u>
		QuantVSR (ours)	<b>22.76</b>	<b>0.6075</b>	<b>0.2857</b>	<b>0.1747</b>	0.4553	<b>65.75</b>	<b>3.41</b>	0.3168	<b>0.6969</b>	<b>1.76</b>
MVS4x		W32A32	MGLD-VSR	22.77	0.7422	0.3571	0.2248	0.3699	53.65	5.06	0.2880	0.6321
	MaxMin		22.29	0.5664	0.5478	0.3335	<b>0.5734</b>	51.38	<b>3.92</b>	<b>0.3166</b>	0.4538	2.47
	W6A6	Q-Diffusion	22.59	0.6377	0.5011	0.2997	<u>0.5486</u>	50.41	3.93	0.3041	0.4353	1.80
		QuaRot	<b>22.85</b>	<b>0.7494</b>	<b>0.3540</b>	<u>0.2333</u>	0.3194	47.08	5.51	0.2531	0.5463	<b>1.40</b>
		ViDiT-Q	22.16	0.6396	0.4487	0.2861	0.5131	51.88	<u>3.92</u>	0.3022	0.5043	2.76
		SVDQuant	<u>22.84</u>	0.7107	0.3842	0.2528	0.4720	<u>53.24</u>	4.12	<u>0.3075</u>	<u>0.5557</u>	1.54
		QuantVSR (ours)	22.80	<u>0.7319</u>	<u>0.3559</u>	<b>0.2276</b>	0.4201	<b>54.68</b>	4.62	0.3009	<b>0.6156</b>	<u>1.53</u>
	W4A4	MaxMin	18.70	0.2813	0.7245	0.4330	0.2818	33.29	5.90	0.3146	0.1627	8.97
		Q-Diffusion	21.10	0.3971	0.6263	0.3855	0.5973	47.53	4.52	<b>0.3660</b>	0.3362	5.33
		QuaRot	21.00	0.3880	0.6253	0.3919	<u>0.6197</u>	49.09	4.88	<u>0.3416</u>	0.3795	5.31
		ViDiT-Q	21.18	0.4471	0.5956	<u>0.3600</u>	0.5833	48.09	<b>4.09</b>	0.3253	0.3618	4.90
		SVDQuant	<u>21.70</u>	<u>0.5021</u>	<u>0.5780</u>	0.3659	<b>0.6394</b>	<u>53.68</u>	<u>4.24</u>	0.3328	<u>0.4727</u>	<u>3.30</u>
		QuantVSR (ours)	<b>22.90</b>	<b>0.7367</b>	<b>0.3590</b>	<b>0.2309</b>	0.4339	<b>55.35</b>	4.57	0.3017	<b>0.6219</b>	<b>1.40</b>

Table 3: Quantitative results on synthetic and real-world datasets. The full-precision backbone is MGLD-VSR (Yang et al. 2024). **Bold** and underline represent the best and second best scores, respectively.

structure, (2) using the SVDQuant (Li et al. 2025a) settings (*i.e.*, fixed rank and SVD, denoted as SVDSC), and (3) our proposed STCA method. As shown in Tab. 1, STCA demonstrates a significant performance improvement compared to SVDSC and the method without SC.

**Learnable Bias Alignment (LBA).** We adopt LBA under various settings, as shown in Tab. 1. In all cases, LBA leads to substantial improvements, particularly when accuracy degradation is severe. This strongly supports our previous conclusion regarding the importance of biased quantization errors in low-bit quantization scenarios. Although modest, integrating LBA after STCA still yields slight enhancements across multiple metrics.

## Main Results

**Quantitative Results.** Quantitative comparisons are shown in Tab. 3. For VQA and reference-based IQA, our method

outperforms prior methods on most datasets, especially in 4-bit settings. Some methods achieve high scores on no-reference IQA metrics (*e.g.*, CLIP-IQA), even surpassing FP models by a large margin, yet their performance on structural metrics (*e.g.*, PSNR and SSIM) remains poor. Prior studies (Zhu et al. 2025) note that images with significant noise may still yield high no-reference scores. These approaches do not surpass ours, as further illustrated in Fig. 4.

**Qualitative Results.** Qualitative comparisons are shown in Figs. 4 and 5. QuantVSR produces sharper details and more faithful textures than prior methods, with minimal difference from the FP model. Leading quantization methods tend to introduce unrealistic artifacts when applied to MGLD-VSR (*e.g.*, ViDiT-Q on REDS4 011), whereas our QuantVSR suppresses such distortions. Moreover, our QuantVSR achieves higher temporal consistency than competing techniques, owing to our special consideration of temporal features.




Figure 4: Visual comparison on synthetic (SPMCS (Yi et al. 2019), REDS4 (Nah et al. 2019)) and real-world (MVSR4x (Wang et al. 2023)) datasets at 6 / 4-bit quantization. Our approach outperforms existing methods especially in the 4-bit setting.

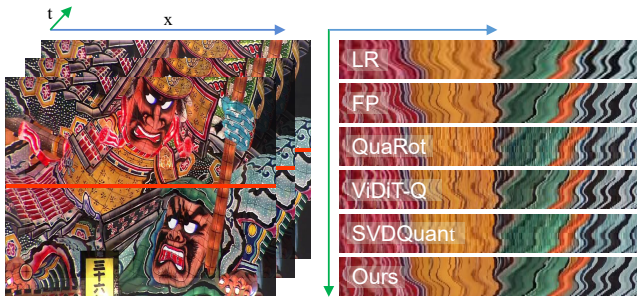


Figure 5: Comparison of temporal consistency (stacking the red line across frames).

**Rank in Full-Precision Branch.** The rank of two matrices in the skip connection structure determines the computational burden. Using our proposed STCA, most layers preserve the minimal rank  $r_{min}$  (i.e., 16), and only a few complex layers require higher ranks (up to  $r_{max}$ , 64). Consequently, the average rank across all layers of QuantVSR is 24, which is lower than the fixed rank of 32 in SVDQuant (Li et al. 2025a) at 4-bit quantization.

**Compression Ratio.** We calculate the model size (Params / M) and the number of operations (Ops / G) with the same method as previous quantization studies (Qin et al. 2023). Table 2 shows that 6-bit QuantVSR shrinks model size and computation by 78.18% and 76.29%, while 4-bit cuts them by 84.39% and 82.56% versus the FP MGLD-VSR.

## Conclusion

In this paper, we propose QuantVSR, a low-bit quantization framework for real-world video super-resolution (VSR). To maintain full-precision (FP) model performance across both spatial and temporal dimensions, we propose a spatio-temporal complexity aware (STCA) mechanism. This also enables the joint optimization of both branches in the skip connection (SC) structure. In addition, a learnable bias alignment (LBA) module mitigates the severe biased quantization error without incurring any additional overhead during inference. Comprehensive experiments on synthetic and real-world benchmarks demonstrate that QuantVSR achieves perceptual quality comparable to its FP counterpart at 6-bit and even 4-bit precision. It significantly outperforms the latest state-of-the-art quantization methods overall.

## Acknowledgments

This work was supported by Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102) and the Fundamental Research Funds for the Central Universities.

## References

- Ashkboos, S.; Mohtashami, A.; Croci, M. L.; Li, B.; Cameron, P.; Jaggi, M.; Alistarh, D.; Hoefler, T.; and Hensman, J. 2024. QuaRot: Outlier-free 4-bit inference in rotated llms. In *NeurIPS*.
- Blattmann, A.; Dockhorn, T.; Kulal, S.; Mendeleevitch, D.; Kilian, M.; Lorenz, D.; Levi, Y.; English, Z.; Voleti, V.; Letts, A.; et al. 2023. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*.
- Chan, K. C.; Wang, X.; Yu, K.; Dong, C.; and Loy, C. C. 2021. Basicvsr: The search for essential components in video super-resolution and beyond. In *CVPR*.
- Chan, K. C.; Zhou, S.; Xu, X.; and Loy, C. C. 2022a. Basicvsr++: Improving video super-resolution with enhanced propagation and alignment. In *CVPR*.
- Chan, K. C.; Zhou, S.; Xu, X.; and Loy, C. C. 2022b. Investigating tradeoffs in real-world video super-resolution. In *CVPR*.
- 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 *CVPR*.
- Ding, K.; Ma, K.; Wang, S.; and Simoncelli, E. P. 2020. Image quality assessment: Unifying structure and texture similarity. *TPAMI*.
- Fei, B.; Lyu, Z.; Pan, L.; Zhang, J.; Yang, W.; Luo, T.; Zhang, B.; and Dai, B. 2023. Generative diffusion prior for unified image restoration and enhancement. In *CVPR*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *NeurIPS*.
- He, J.; Xue, T.; Liu, D.; Lin, X.; Gao, P.; Lin, D.; Qiao, Y.; Ouyang, W.; and Liu, Z. 2024. Venhancer: Generative space-time enhancement for video generation. *arXiv preprint arXiv:2407.07667*.
- He, Y.; Liu, J.; Wu, W.; Zhou, H.; and Zhuang, B. 2023a. Efficientdm: Efficient quantization-aware fine-tuning of low-bit diffusion models. *arXiv preprint arXiv:2310.03270*.
- He, Y.; Liu, L.; Liu, J.; Wu, W.; Zhou, H.; and Zhuang, B. 2023b. PTQD: Accurate Post-Training Quantization for Diffusion Models. In *NeurIPS*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. In *NeurIPS*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*.
- Huang, W.; Zheng, X.; Ma, X.; Qin, H.; Lv, C.; Chen, H.; Luo, J.; Qi, X.; Liu, X.; and Magno, M. 2024. An empirical study of llama3 quantization: From llms to mllms. *Visual Intelligence*.
- Jacob, B.; Kligys, S.; Chen, B.; Zhu, M.; Tang, M.; Howard, A.; Adam, H.; and Kalenichenko, D. 2018. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *CVPR*.
- Jo, Y.; Oh, S. W.; Kang, J.; and Kim, S. J. 2018. Deep video super-resolution network using dynamic upsampling filters without explicit motion compensation. In *CVPR*.
- Ke, J.; Wang, Q.; Wang, Y.; Milanfar, P.; and Yang, F. 2021. Musiq: Multi-scale image quality transformer. In *ICCV*.
- Lai, W.-S.; Huang, J.-B.; Wang, O.; Shechtman, E.; Yumer, E.; and Yang, M.-H. 2018. Learning blind video temporal consistency. In *ECCV*.
- Li, M.; Lin, Y.; Zhang\*, Z.; Cai, T.; Li, X.; Guo, J.; Xie, E.; Meng, C.; Zhu, J.-Y.; and Han, S. 2025a. SVDQuant: Absorbing Outliers by Low-Rank Components for 4-Bit Diffusion Models. In *ICLR*.
- Li, W.; Tao, X.; Guo, T.; Qi, L.; Lu, J.; and Jia, J. 2020. Mucan: Multi-correspondence aggregation network for video super-resolution. In *ECCV*.
- Li, X.; Liu, Y.; Cao, S.; Chen, Z.; Zhuang, S.; Chen, X.; He, Y.; Wang, Y.; and Qiao, Y. 2025b. DiffVSR: Enhancing Real-World Video Super-Resolution with Diffusion Models for Advanced Visual Quality and Temporal Consistency. *arXiv preprint arXiv:2501.10110*.
- Li, X.; Liu, Y.; Lian, L.; Yang, H.; Dong, Z.; Kang, D.; Zhang, S.; and Keutzer, K. 2023a. Q-diffusion: Quantizing diffusion models. In *ICCV*.
- Li, Y.; Xu, S.; Cao, X.; Sun, X.; and Zhang, B. 2023b. Q-dm: An efficient low-bit quantized diffusion model. In *NeurIPS*.
- Liang, J.; Cao, J.; Fan, Y.; Zhang, K.; Ranjan, R.; Li, Y.; Timofte, R.; and Van Gool, L. 2024. Vrt: A video restoration transformer. *TIP*.
- Liang, J.; Fan, Y.; Xiang, X.; Ranjan, R.; Ilg, E.; Green, S.; Cao, J.; Zhang, K.; Timofte, R.; and Gool, L. V. 2022. Recurrent video restoration transformer with guided deformable attention. In *NeurIPS*.
- Liu, C.; and Sun, D. 2013. On Bayesian adaptive video super resolution. *IEEE transactions on pattern analysis and machine intelligence*.
- Liu, K.; Zheng, Q.; Tao, K.; Li, Z.; Qin, H.; Li, W.; Guo, Y.; Liu, X.; Kong, L.; Chen, G.; et al. 2025. Low-bit model quantization for deep neural networks: A survey. *arXiv preprint arXiv:2505.05530*.
- Liu, Z.; Oguz, B.; Zhao, C.; Chang, E.; Stock, P.; Mehdad, Y.; Shi, Y.; Krishnamoorthi, R.; and Chandra, V. 2023. LLM-QAT: Data-Free Quantization Aware Training for Large Language Models. *arXiv preprint arXiv:2305.17888*.
- Lucas, A.; Lopez-Tapia, S.; Molina, R.; and Katsaggelos, A. K. 2019. Generative adversarial networks and perceptual losses for video super-resolution. *TIP*.
- Nagel, M.; Baalen, M. v.; Blankevoort, T.; and Welling, M. 2019. Data-free quantization through weight equalization and bias correction. In *ICCV*.
- Nagel, M.; Fournarakis, M.; Amjad, R. A.; Bondarenko, Y.; Van Baalen, M.; and Blankevoort, T. 2021. A white

- paper on neural network quantization. *arXiv preprint arXiv:2106.08295*.
- Nah, S.; Baik, S.; Hong, S.; Moon, G.; Son, S.; Timofte, R.; and Mu Lee, K. 2019. Ntire 2019 challenge on video deblurring and super-resolution: Dataset and study. In *CVPRW*.
- Pan, J.; Bai, H.; Dong, J.; Zhang, J.; and Tang, J. 2021. Deep blind video super-resolution. In *CVPR*.
- Qin, H.; Zhang, Y.; Ding, Y.; Liu, X.; Danelljan, M.; Yu, F.; et al. 2023. QuantSR: accurate low-bit quantization for efficient image super-resolution. In *NeurIPS*.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *CVPR*.
- Shang, Y.; Yuan, Z.; Xie, B.; Wu, B.; and Yan, Y. 2023. Post-training Quantization on Diffusion Models. In *CVPR*.
- Shao, W.; Chen, M.; Zhang, Z.; Xu, P.; Zhao, L.; Li, Z.; Zhang, K.; Gao, P.; Qiao, Y.; and Luo, P. 2023. Omniquant: Omnidirectionally calibrated quantization for large language models. *arXiv preprint arXiv:2308.13137*.
- Shi, S.; Gu, J.; Xie, L.; Wang, X.; Yang, Y.; and Dong, C. 2022. Rethinking alignment in video super-resolution transformers. In *NeurIPS*.
- Sui, Y.; Li, Y.; Kag, A.; Idelbayev, Y.; Cao, J.; Hu, J.; Sagar, D.; Yuan, B.; Tulyakov, S.; and Ren, J. 2024. Bits-fusion: 1.99 bits weight quantization of diffusion model. In *NeurIPS*.
- Wang, J.; Chan, K. C.; and Loy, C. C. 2023. Exploring clip for assessing the look and feel of images. In *AAAI*.
- Wang, J.; Lin, Z.; Wei, M.; Zhao, Y.; Yang, C.; Xiao, F.; Loy, C. C.; and Jiang, L. 2025. Seedvr: Seeding infinity in diffusion transformer towards generic video restoration. In *CVPR*.
- Wang, R.; Liu, X.; Zhang, Z.; Wu, X.; Feng, C.-M.; Zhang, L.; and Zuo, W. 2023. Benchmark dataset and effective inter-frame alignment for real-world video super-resolution. In *CVPRW*.
- Wang, X.; Chan, K. C.; Yu, K.; Dong, C.; and Change Loy, C. 2019. Edvr: Video restoration with enhanced deformable convolutional networks. In *CVPRW*.
- Wang, Y.; Yu, J.; and Zhang, J. 2023. Zero-Shot Image Restoration Using Denoising Diffusion Null-Space Model. In *ICLR*.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *TIP*.
- Wu, H.; Zhang, E.; Liao, L.; Chen, C.; Hou, J.; Wang, A.; Sun, W.; Yan, Q.; and Lin, W. 2023. Exploring video quality assessment on user generated contents from aesthetic and technical perspectives. In *ICCV*.
- Xiao, G.; Lin, J.; Seznec, M.; Wu, H.; Demouth, J.; and Han, S. 2023. Smoothquant: Accurate and efficient post-training quantization for large language models. In *ICML*.
- Xie, L.; Wang, X.; Shi, S.; Gu, J.; Dong, C.; and Shan, Y. 2023. Mitigating artifacts in real-world video super-resolution models. In *AAAI*.
- Xie, R.; Liu, Y.; Zhou, P.; Zhao, C.; Zhou, J.; Zhang, K.; Zhang, Z.; Yang, J.; Yang, Z.; and Tai, Y. 2025. STAR: Spatial-Temporal Augmentation with Text-to-Video Models for Real-World Video Super-Resolution. *arXiv preprint arXiv:2501.02976*.
- Xue, T.; Chen, B.; Wu, J.; Wei, D.; and Freeman, W. T. 2019. Video enhancement with task-oriented flow. *IJCV*.
- Yang, S.; Wu, T.; Shi, S.; Lao, S.; Gong, Y.; Cao, M.; Wang, J.; and Yang, Y. 2022. Maniq: Multi-dimension attention network for no-reference image quality assessment. In *CVPRW*.
- Yang, X.; He, C.; Ma, J.; and Zhang, L. 2024. Motion-guided latent diffusion for temporally consistent real-world video super-resolution. In *ECCV*.
- Yang, X.; Xiang, W.; Zeng, H.; and Zhang, L. 2021. Real-world video super-resolution: A benchmark dataset and a decomposition based learning scheme. In *CVPR*.
- Yang, Y.; Dai, X.; Wang, J.; Zhang, P.; and Zhang, H. 2023. Efficient quantization strategies for latent diffusion models. *arXiv preprint arXiv:2312.05431*.
- Yi, P.; Wang, Z.; Jiang, K.; Jiang, J.; and Ma, J. 2019. Progressive fusion video super-resolution network via exploiting non-local spatio-temporal correlations. In *ICCV*.
- Zhang, L.; Zhang, L.; and Bovik, A. C. 2015. A feature-enriched completely blind image quality evaluator. *IEEE Transactions on Image Processing*.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*.
- Zhang, S.; Wang, J.; Zhang, Y.; Zhao, K.; Yuan, H.; Qin, Z.; Wang, X.; Zhao, D.; and Zhou, J. 2023. I2vgen-xl: High-quality image-to-video synthesis via cascaded diffusion models. *arXiv preprint arXiv:2311.04145*.
- Zhao, T.; Fang, T.; Huang, H.; Liu, E.; Wan, R.; Soedarmadji, W.; Li, S.; Lin, Z.; Dai, G.; Yan, S.; et al. 2024. Viditq: Efficient and accurate quantization of diffusion transformers for image and video generation. *arXiv preprint arXiv:2406.02540*.
- Zhou, S.; Yang, P.; Wang, J.; Luo, Y.; and Loy, C. C. 2024. Upscale-a-video: Temporal-consistent diffusion model for real-world video super-resolution. In *CVPR*.
- Zhu, L.; Li, J.; Qin, H.; Li, W.; Zhang, Y.; Guo, Y.; and Yang, X. 2025. PassionSR: Post-training quantization with adaptive scale in one-step diffusion based image super-resolution. In *CVPR*.