SkipSR: Faster Super Resolution with Token Skipping

Rohan Choudhury[*1,2] **Shanchuan Lin**[2] **Jianyi Wang**[2] **Hao Chen**[2] **Qi Zhao**[2]
Feng Cheng[2] **Lu Jiang**[2] **Kris M. Kitani**[1] **László A. Jeni**[1]
[1]Carnegie Mellon University [2]ByteDance Seed
rchoudhu@andrew.cmu.edu

Abstract

Diffusion-based super-resolution (SR) is a key component in video generation and video restoration, but is slow and expensive, limiting scalability to higher resolutions and longer videos. Our key insight is that many regions in video are inherently low-detail and gain little from refinement, yet current methods process all pixels uniformly. To take advantage of this, we propose SkipSR, a simple framework for accelerating video SR by identifying low-detail regions directly from low-resolution input, then skipping computation on them entirely, only super-resolving the areas that require refinement. This simple yet effective strategy preserves perceptual quality in both standard and one-step diffusion SR models while significantly reducing computation. In standard SR benchmarks, our method achieves up to 60% faster end-to-end latency than prior models on 720p videos with no perceptible loss in quality. Video demos are available at our project page.

1 Introduction

Diffusion transformers are the dominant paradigm in image and video generation, but due to the quadratic cost of self-attention, computational time grows steeply with resolution and sequence length. Furthermore, diffusion models typically require tens of steps to produce high-quality outputs. To handle this, a common design choice is *cascaded* diffusion [Ho et al., 2022, Saharia et al., 2022a, Gao et al., 2025], which first generates low-resolution images or videos, then uses a conditional diffusion model to upscale and refine them into high-resolution results with fewer steps. However, because of the larger input size at higher resolutions, diffusion-based super-resolution (SR) steps can be very slow, often dominating the computation time. Speeding these up is essential for high-resolution video generation and restoration.

Most prior work on addressing this issue focuses on reducing the number of diffusion steps. While successful, this leaves gains on the table since even a single diffusion step is still expensive on high-resolution video. Another recent line of work experiment with alternative attention mechanisms, such as windowed, sliding tile, or spatially sparse attention. These works typically restrict the attention mechanism to focus on nearby tokens, resulting in significant speedups.

A common assumption behind all these works is that they spend equal computation on every input patch. However, not every patch is created equal. Our key insight is that many videos contain *visually simple regions*: for example, a blue sky or a blurry background. This phenomenon is particularly pronounced in super-resolution tasks, where the low-resolution conditional input often includes extensive visually simple areas. These regions can be simply resized and placed directly in the output, since they do not need refinement like more detailed regions of the input do.

We present Skip Super Resolution (SkipSR), a method that uses this idea to accelerate video super-resolution. SkipSR uses a lightweight mask predictor to route only patches that require refinement

[*]Work done during an internship at ByteDance Seed.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).



Original Image | Standard Diffusion Transformer | SkipSR (Ours)

Figure 1: **Patch Skipping.**Standard diffusion SR models refine the entire input, while SkipSR identifies and upscales only the patches that need refinement. This significantly reduces computation with no loss in perceptual quality.

through the transformer, while the remaining patches skip the model entirely. The transformer model maintain positional awareness with mask-aware rotary positional encodings, and the two patch groups are composed together at the output stage, resulting in perceptually indistinguishable outputs at a significantly lower computational cost.

We primarily evaluate on video SR tasks, given their considerable inference cost, and validate SkipSR on top of a state-of-the-art video SR model [Wang et al., 2025b,a]. Our experimental results demonstrate its versatility in both multi-step and one-step diffusion paradigms. SkipSR significantly accelerates super-resolution, yielding outputs that are perceptually indistinguishable from dense attention equivalents. In particular, we achieve up to 60% faster super-resolution on 720p videos with no loss in quality, supported both by video-quality metrics and user studies, and reduce the diffusion time for 1080p videos by 70%.

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

- We demonstrate simple regions are common in videos, and that identifying and skipping these simple regions can match perceptual quality with a fraction of the compute.
- We propose SkipSR, consisting of a lightweight mechanism to identify complex regions and apply sparse attention only to them, leading to accurate and efficient super-resolution.
- We validate our hypotheses and design choices with extensive experiments, demonstrating a consistent speed-up while maintaining quality.

2 Related Work

Efficient Video Diffusion. Video generation is prohibitively slow, with full-attention models like HunyuanVideo [Kong et al., 2024] taking tens of minutes to generate a 5-second video [Xi et al., 2025]. Since most models require many steps [Song and Ermon, 2019, Ho et al., 2020, Meng et al., 2021], most work has focused on reducing the number of steps, through improving flow [Liu et al., 2022, Geng et al., 2025], consistency losses [Song et al., 2023], or distillation [Salimans and Ho, 2022, Yin et al., 2024, Lin et al., 2025]. We build upon these works by demonstrating speed-ups on both multi-step and single-step models.

Several methods have proposed speed-ups by increasing sparsity into the attention mechanism, thus reducing the number of tokens considered. Sliding window attention [Hassani et al., 2023, Liu et al., 2024, Zhang et al., 2025b] and other variants such as Radial Attention Li et al. [2025] restrict attention to focus on spatiotemporally local tokens. Other works builds on this by enforcing sparsity on attention heads in an online manner, such as SVG and Sparge [Xi et al., 2025, Zhang et al., 2025a]. Our method differs from these by skipping the entire Transformer, rather than just the attention mechanism.

Diffusion Models for Video Super-Resolution. Alternatively, diffusion can be made more efficient by first diffusing at low resolution, then upscaling at the end. [Ho et al., 2022]. Cascaded models like this are commonplace, such as SR3 [Saharia et al., 2022b], Stable Diffusion [Rombach et al., 2022], and Seedance [Gao et al., 2025]. Diffusion is also considered standard for stand-alone

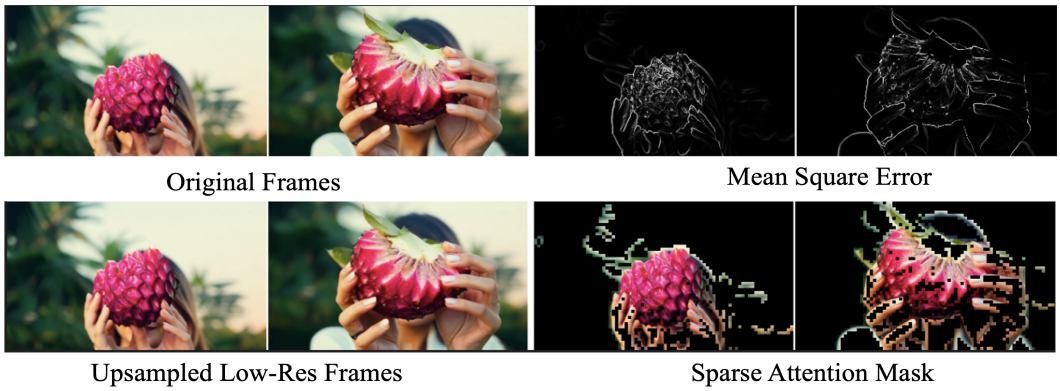


Figure 2: **Oracle Mask Computation.** We identify the low-resolution video regions by comparing the original high-resolution input to a spatial downsampled, then upsampled version. This results in a attention mask, shown in dark, that excludes these regions from refinement.

video restoration. Recent works adapt image diffusers or condition video diffusers on LR inputs via ControlNet [Zhou et al., 2024, Yang et al., 2024]. While improving fidelity, these methods inherit the high cost of multi-step sampling and heavy conditioning. Parallel efforts aim to accelerate diffusion through fewer or single sampling steps, using rectified flows [Liu et al., 2022], distillation [Wang et al., 2025a, Dong et al., 2025, Chen et al., 2025], or deterministic one-step models. However, these methods still process every region of the input equally, so that simple videos take the same time as more complex ones.

Efficient Super-Resolution. Prior works have also attempted to accelerate super-resolution, mainly on images, by identifying which regions need less refinement. Sparse Mask SR [Wang et al., 2021] enforces sparsity to this end on CNN filters, while other works [Kong et al., 2021, Hu et al., 2022, Fan et al., 2024] apply different sub-networks to sub-regions of an input image based on their predicted complexity. Although these methods reduce FLOPS, they result in significantly slower wall clock time, as stated in the papers. YODA [Moser et al., 2025] uses attention masks to guide and improve diffusion for image SR, but does not accelerate the model. We provide a simple solution to these issues while significantly speeding up wall-clock inference time.

3 Motivation and Analysis

Selective Patch Processing. High-resolution videos often contain regions that do not contain high-frequency details, e.g., a blue sky, a plain wall, or an out-of-focus background. We hypothesize that these regions can skip the expensive transformer computation and be upsampled by cheaper methods, such as bilinear interpolation, instead.

Concretely, we divide a high-resolution video I of size $T \times H \times W \times 3$ into non-overlapping patches P of size $4 \times 16 \times 16 \times 3$. We define a patch as *skippable* if, after area downsampling D and bilinear upsampling U, the reconstruction mean squared error (MSE) is smaller than a threshold τ:

$$\text{MSE}\left(P, U(D(P))\right) \;\leq\; \tau \tag{1}$$

Note that while there are many ways to classify the skippable regions, this is not the focus of our work. Empirically, we find this straightforward measurement suffices.

To measure the commonality of such skippable regions, we measure the percentage of patches satisfying the criterion in Eq. (1) with $\tau = 0.0002$ and a spatial downsampling factor of $4\times$ on a diverse set of videos. Furthermore, to validate skipping the patches does not degrade quality in the latent diffusion setting, we encode the original video I through a VAE and swap the latent at the corresponding skippable region with the encoding of $U(D(I))$ before decoding back to the pixel space. Then, PSNR, SSIM, and LPIPS [Zhang et al., 2018] metrics are used to measure degradation against the original video I.

Table 1 shows the results of our analysis on videos from VBench [Huang et al., 2024], AIGC-30 [Wang et al., 2025b], VideoLQ [Chan et al., 2022], and YouHQ-40 [Zhou et al., 2024]. In particular, VBench and AIGC-30 are AI generated videos, which yield the highest amount of skippable

Table 1: **Oracle Mask Analysis.** Using the oracle mask, we measure the reconstruction error of upsampling and swapping simple patches, as well as the resulting expected speedup. We spatially downsample by $4\times$ in all cases.

Dataset	Resolution	Skippable %	PSNR↑	SSIM↑	LPIPS↓	Speedup↑
VBench	720p	44.8 %	42.24	0.988	0.0400	1.8×
AIGC-30	720p	41.9%	43.11	0.989	0.038	1.7×
VideoLQ	720p	31.2%	46.61	0.994	0.0239	1.5×
YouHQ-40	720p	14.7%	46.64	0.994	0.0149	1.2×
YouHQ-40 (corrupted)	720p	0.9%	66.33	0.999	0.0002	1.0×
AIGC-30	1080p	44.9%	44.39	0.991	0.031	1.8×
VideoLQ	1080p	36.7%	48.32	0.993	0.013	1.6×

patches, up to 45%. For real-world videos, VideoLQ is a common SR evaluation dataset offered in low resolution. We use the upsampled version by SeedVR [Wang et al., 2025b] and yield 30% skippable patches. YouHQ has large camera movements and yields a lower 15% skippable patches. Overall, our approach can expect an average reduction of 30% tokens for regular super-resolution tasks while retaining a VAE reconstruction PSNR of 40+ and LPIPS of near 0, which are considered imperceptible to human eyes. However, as a limitation, we also show that the synthetically corrupted YouHQ contains heavy white noise everywhere. Such cases do not yield much patch savings and are less applicable to our approach.

Estimating the Theoretical Speedup. Next, we estimate the expected speed-up by completely skipping the patches from the transformer and profiling the diffusion model. We use the state-of-the-art SR method SeedVR2 [Wang et al., 2025a] as the baseline architecture, and the relative speed-up is provided in the last column of Table 1. This preliminary experiment shows that skipping patches can achieve $1.2\times$ to $1.8\times$ speed-up on the diffusion model. Inspired by these observations, our method is designed to (1) identify the simple regions from low-resolution input as accurately as possible and (2) produce high-quality super-resolution outputs from applying attention only on the complex regions.

4 Method

This section first provides a brief review of the video super-resolution model on which our method is built. We then introduce our approach in Section 4.2.

4.1 Preliminaries

Recent diffusion-based SR methods adopt the latent diffusion paradigm [Rombach et al., 2022], where a pre-trained variational autoencoder (VAE) [Kingma and Welling, 2013] compresses images or videos into a latent space, and a diffusion transformer generates the high-resolution latent conditioned on text and low-resolution latents. During training, low-resolution inputs are obtained by applying synthetic degradations (e.g. blurring, downsampling, noise injection, compression) to high-resolution data. This procedure works well for super-resolution as well as more complex restoration. For more details, we refer the reader to Wang et al. [2025b].

We adopt the architecture used in SeedVR [Wang et al., 2025b], which combines shifted windows [Liu et al., 2021] and the native-resolution trick of Dehghani et al. [2023] and use MMDiT [Esser et al., 2024] for text-conditioning in the DiT. We employ a causal video VAE that compresses the pixel-space input by $8\times$ spatially and $4\times$ temporally, and the DiT operates on $1 \times 2 \times 2$ patches in latent space.

4.2 Sparse Super-Resolution

Skip Prediction Model. Unlike our preliminary analysis in Section 3, at inference time, we do not have a high-resolution ground-truth video to identify skippable patches. We train a lightweight predictor network to predict skippable patches given low-resolution videos. Specifically, our predictor network operates in the VAE latent space and is composed of 4 layers of 3D convolution and

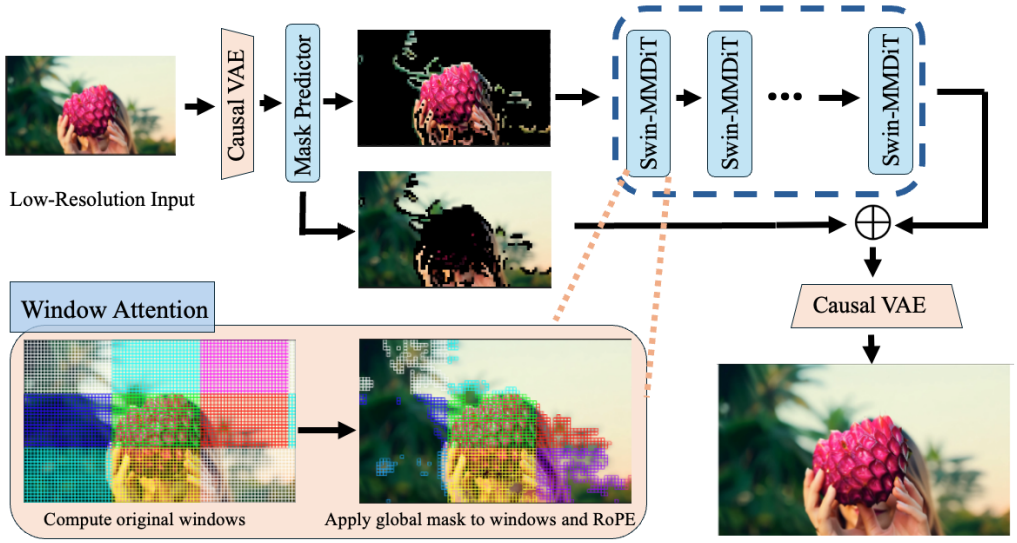


Figure 3: **SkipSR Overview.** We take as input low-resolution videos, project into latent space, then compute the complexity mask. Simple patches skip computation and are routed around the transformer, then composed with the refined output.

ReLU activations. The first convolution has a spatial stride of 2 to match the patch size. The network takes latent of shape $t \times h \times w \times 16$ as input, and outputs binary classification logits of shape $t \times h/2 \times w/2 \times 1$. This corresponds to a patch size of $1 \times 2 \times 2$ in the latent space and $4 \times 16 \times 16$ in the pixel space.

During training, given high-resolution videos I and their synthetic low-resolution pairs $D(I)$, we encode $D(I)$ through VAE and provide it as input to the predictor. The predictor network is trained using binary cross-entropy loss, where the ground-truth classification target is given by Eq. (1). During inference, low-resolution videos i are used instead of $D(I)$, and the predictor network classifies skippable patches with high accuracy.

Skip-Aware Diffusion Model. Prior works that accelerate diffusion models modify the attention algorithm, but leave the other components, such as the FFN and LayerNorm, unchanged. Our key finding is that the predicted skipped tokens *can skip the Transformer altogether*. Typically, these skipped tokens are in the background, while the un-skipped tokens come from the more detailed foreground. We speculate that the un-skipped tokens, along with the text conditioning, provide enough context to produce high-quality super-resolutions.

After the mask predictor produces a mask M, we apply M to the set of input latent patches, P. P has N patches, where $N = t \times h/2 \times w/2$, indexed as $P = \{p_1, p_2, \dots p_N\}$ We then apply M to P, partitioning it into two sets P_{skip} and P_{unskip}. P_{skip} receives no further computation.

P_{unskip} is next passed through the Transformer's patch projection layer, and proceeds through each Transformer block. However, as shown in Figure 3, the patches in P_{unskip} are no longer spatiotemporally contiguous. For example, P_{unskip} could consist of $\{p_3, p_4, p_1 3 \dots p_N\}$. We thus need to ensure that the Transformer is aware of their relative positions. We accomplish this by modifying the rotary positional encodings (RoPE) [Su et al., 2021] in the attention operation. In RoPE, the embeddings in each token are rotated by an angle θ that is scaled by the position index:

$$\text{RoPE}(p_i, i) = \mathbf{R}_i \cdot p_i \tag{2}$$

where p_i is the i-th patch in P, d is the feature dimension, and $\mathbf{R}_m \in \mathbb{R}^{d \times d}$ is a block-diagonal rotation matrix with angles $i \cdot \theta_k$ for each dimension pair k, where $\theta_k = 10000^{-2k/d}$.

Rather than use the index m of each patch in P_{unskip} , we use each the original index from P, encoding the relative position of each patch. With our previous example, if the first element of P_{unskip} was p_3, we would apply RoPE with $i = 3$ rather than $i = 1$, ensuring that patches in P_{unskip} are aware of their relative distances. After running the Transformer and the patch output projection layer on P, the output P'_{unskip} is composed with P_{skip} using M, resulting in a mixed sequence of processed and unchanged patches.

Table 2: **Main Quantitative Results.** We compare SkipSR's super-resolution output against SeedVR and SeedVR2 baselines on real-world (VideoLQ) and AI-generated (AIGC30) benchmarks. The best value is in **blue**, and the second-best in red. SkipSR matches the performance of SeedVR and SeedVR2 with significantly less compute.

Dataset	Method	NIQE↓	MUSIQ↑	CLIP-IQA↑	DOVER↑	BRISQUE↓	TFLOPs↓	s/step↓	Skipped%
VideoLQ	SeedVR-3B	**4.069**	**57.41**	**0.318**	**8.009**	35.45	2001	6.41	0
	SeedVR-7B	4.933	48.35	0.258	7.416	**28.120**	3139	7.74	0
	Ours (3B)	4.112	54.175	0.288	7.992	38.094	**1343**	**4.47 (1.4×)**	**27.7%**
	SeedVR2-3B	4.687	51.09	0.295	8.176	38.1	2001	6.41	0 %
	SeedVR2-7B	4.948	45.76	0.257	7.236	41.332	3139	7.74	0%
	Ours (3B, one-step)	**4.3449**	**57.783**	**0.324**	**8.199**	**39.220**	**1343**	**4.47 (1.4×)**	**27.7%**
AIGC30	SeedVR-3B	**3.655**	64.4	**0.589**	12.9	**25.3**	2001	6.41	0%
	SeedVR-7B	4.053	64.26	0.564	**16.47**	35.7	3139	7.74	0%
	Ours (3B)	3.99	**65.14**	0.575	11.9	35.89	**1135**	**3.85 (1.6×)**	**39.6%**
	SeedVR2-3B	3.801	62.99	0.561	15.77	**28.021**	2001	6.41	0%
	SeedVR2-7B	4.138	65.09	0.574	13.4	39.14	3139	7.74	0%
	Ours (3B, one-step)	**3.6628**	**67.218**	**0.581**	**16.021**	29.534	**1135**	**3.85 (1.6×)**	**39.6%**

Handling Window Attention. Our DiT uses shifted window attention, which partitions the input into non-overlapping windows at each layer. In this case, the rotary positional encodings are not defined with respect to the entire input sequence length, but instead are based on the window size. To handle this with the SkipSR mask, we first assign each patch in P to a window as normal, then apply the mask, resulting in P_{unskip}. Each patch simply P keeps its original window assignment from P, as shown in the inset of Figure 3. Although each window now has an imbalanced number of tokens, this is handled natively by FlashAttention and NaViT, while preserving rotary positional embeddings dependent on a fixed window size.

Training. To train the mask predictor, we freeze the VAE encoder and train on high-resolution images and videos. For the main super-resolution model, we follow standard practice and train on a mixture of images and videos [Chen et al., 2025, Wang et al., 2025b], using standard diffusion training. Furthermore, while standard diffusion transformers are multi-step, several works accelerate these models by reducing the number of steps required to just one step. We train a one-step version of SkipSR to measure against these approaches, based on the procedure introduced in APT [Lin et al., 2025]. We apply progressive distillation [Salimans and Ho, 2022] to distill our model to reduce one-step, then adversarially post-train it using a transformer-based discriminator to restore sharpness, resulting in a one-step model that matches the quality of the original.

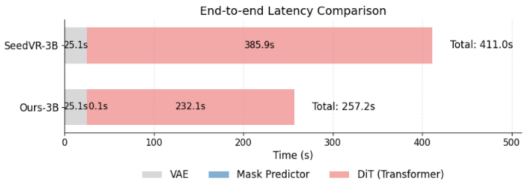
5 Experiments

Implementation Details. Models are initialized from public SeedVR/SeedVR2 (3B/7B) checkpoints and fine-tuned jointly with the same data and objectives as the baselines. Training uses 40 NVIDIA A100-80G GPUs with batches of about 100 frames at 720p, using sequence [Korthikanti et al., 2023] and data parallelism [Li et al., 2020]. We use the same data for training as Wang et al. [2025b]. The mask predictor is trained on the same data as the super-resolution model with 8 GPUs for 10k iterations, as the mask predictor converges significantly faster due to its smaller size.

Experimental Settings. Our goal is to simply compare the generation quality of the base model (SeedVR) on real-world videos while reducing wall-clock time. We conduct our main evaluation on the VideoLQ [Chan et al., 2022] and our AIGC30 datasets, reporting commonly used reference-free quality metrics such as NIQE [Mittal et al., 2013], CLIP-IQA [Wang et al., 2023], MUSIQ [Ke et al., 2021] and DOVER [Wu et al., 2023]. Efficiency is measured by runtime and skipped token fraction (token%). All comparisons use unmodified SeedVR / SeedVR2 with our masked variants initialized from the same checkpoints and trained using the same procedure. Though not the main focus of our work, we also measure restoration ability on heavily degraded synthetic benchmarks, namely SPMCS [Yi et al., 2019], UDM10 [Yi et al., 2020], REDS30 [Nah et al., 2019], and YouHQ40, where we compute PSNR, SSIM, LPIPS and DISTS.

Methods-{Steps}	Speed (seconds)	Overall Quality
SeedVR-3B-50	320	+4%
DOVE-1	14.1	-18%
SeedVR2-3B-1	6.41	0%
Ours-3B-50	**3.5**	**+7%**

(a) **User Study.** Our method is preferred by users relative to SeedVR2 while also being significantly faster.



(b) **Latency.** The end-to-end generation latency for the transformer is significantly improved by SkipSR, reducing the total time by 60% while the mask predictor adds negligible overhead.

Table 3: **Results on Synthetic VSR Benchmarks.** SkipSR matches the performance of other methods on heavily degraded benchmarks. Best numbers are **blue**; second best are red.

Datasets	Metrics	RealViformer	MGLD-VSR	STAR	DOVE	SeedVR-7B	SeedVR2-3B	SeedVR2-7B	SkipSR-3B *(Ours)*
SPMCS	PSNR ↑	**24.185**	23.41	22.58	23.11	20.78	22.97	22.90	23.9892
	SSIM ↑	0.663	0.633	0.609	0.621	0.575	0.646	0.638	**0.6807**
	LPIPS ↓	0.378	0.369	0.420	0.288	0.395	0.306	0.322	**0.2865**
	DISTS ↓	0.186	0.166	0.229	0.171	0.166	0.131	0.134	**0.1234**
UDM10	PSNR ↑	**26.70**	26.11	24.66	26.48	24.29	25.61	26.26	26.5243
	SSIM ↑	0.796	0.772	0.747	0.783	0.731	0.784	**0.798**	0.788
	LPIPS ↓	0.285	0.273	0.359	0.270	0.264	0.218	**0.203**	0.2429
	DISTS ↓	0.166	0.144	0.195	0.149	0.124	0.106	**0.101**	0.117
REDS30	PSNR ↑	**23.34**	22.74	22.04	22.11	21.74	21.90	22.27	22.4298
	SSIM ↑	**0.615**	0.578	0.593	22.67	0.596	0.598	0.606	0.6093
	LPIPS ↓	0.328	0.271	0.487	0.277	0.340	0.350	0.337	**0.2682**
	DISTS ↓	0.154	**0.097**	0.229	0.106	0.122	0.135	0.127	0.1081
YouHQ40	PSNR ↑	**23.26**	22.62	22.15	24.3	20.60	22.10	22.46	23.1495
	SSIM ↑	0.606	0.576	0.575	0.674	0.546	0.595	0.600	**0.6277**
	LPIPS ↓	0.362	0.356	0.451	0.299	0.323	0.284	0.274	**0.2361**
	DISTS ↓	0.193	0.166	0.213	0.148	0.134	0.122	0.110	**0.1007**

5.1 Main Results

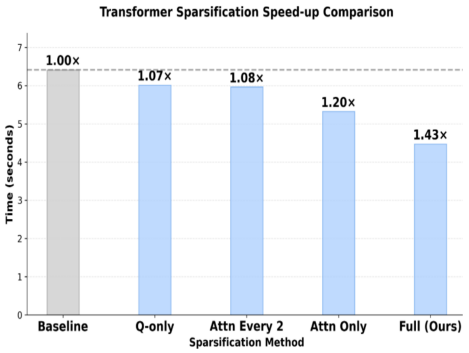
Since SkipSR is a general way to accelerate video SR, our goal is to match the performance of existing diffusion-based methods with minimal loss in quality while achieving a significant speed-up. We compare against the state-of-the-art diffusion methods, including other prior methods for completeness. Our main results on super-resolution benchmarks are in Table 2, where we evaluated SkipSR on real-world and AI-generated video SR test sets. In general, SkipSR matches or exceeds SeedVR's performance in video quality metrics, while significantly reducing the total time required, reducing it by 40% in VideoLQ and 60% on AIGC-30. We also match the performance of SeedVR2 in one-step generation, demonstrating the versatility of our framework and its potential applicability to cascaded diffusion. We also measure the end-to-end latency of the system in Figure 4b. The computation time is dominated by the diffusion transformer, with the VAE taking around 30s total and the mask predictor adding 60ms of overhead. The mask predictor also adds 23M parameters and 22MB of memory to the base model's 3B parameters and 13GB memory, incurring little cost overall. We find that on average, SkipSR reduces the overall end-to-end generation time by 60%, more than 2 minutes, strongly supporting its general applicability.

User Study. User studies are crucial for evaluating generative outputs. We conducted a GSB study with three domain experts, asking each to decide whether the samples are better, same, or worse. The preference score is calculated as $\frac{G-B}{G+S+B}$, which ranges from -100% to 100%. We randomly select 25 samples from VideoLQ and AIGC-30 and ask subjects to evaluate the overall quality of our generations compared to the baseline, SeedVR2. As shown in Table 4a, we match SeedVR-3B and SeedVR2, and significantly outperform DOVE, supporting our claim that we can accelerate the model without losing visual quality as perceived by users.

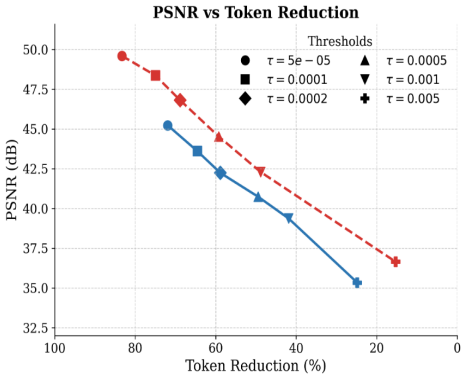
Synthetic Benchmarks. Synthetic video restoration benchmarks are created by applying small degradations to the entire frame, making them unsuitable for our core hypothesis, that upsampled

Table 4: **Predicted Mask Analysis.** We repeat the analysis in Table 4 with the learned predictor and find that we generally maintain PSNR, but retain slightly more tokens.

Dataset	Resolution	PSNR↑	SSIM↑	Skipped (%)↓
VBench	720p	43.67 (+1.4)	0.980 (–0.008)	58.1 (-2.9)
VideoLQ	720p	43.87 (–2.7)	0.991 (–0.001)	27.7 (-3.5)
YouHQ-40 (clean)	720p	44.52 (-2.1)	0.988 (-0.006)	8.6 (-6.1)
AIGC-30	720p	42.03 (–1.1)	0.991 (-0.003)	39.6 (-2.3)
VideoLQ	1080p	43.87 (–2.7)	0.991 (–0.001)	33.4 (-3.3)
AIGC-30	1080p	43.24 (–1.2)	0.993 (+0.002)	42.9 (-2.0)



(a) **Sparsity mechanism.** We measure the speed of routing around the model entirely compared to other design choices, such as only masking the attention, using interspersed global layers, or only masking the attention queries. Ours is significantly faster, while maintainng quality.



(b) **Threshold effect.** As we increase the threshold τ, the reduction increases at the cost of maximum PSNR. $\tau = 0.0002$ is a reasonable choice for token reduction while maintaining perceptual quality.

regions from low-resolution videos map well onto their high-resolution outputs. Achieving state-of-the-art performance on these is not the main focus of this work; we are emphatically focused on general acceleration without significant loss of quality. However, we include the results of these benchmarks in Table 3 for completeness. Our model performs surprisingly well, matching the performance of comparable models, such as SeedVR, DOVE, and SeedVR2.

5.2 Analysis

Mask Predictor. We repeat the analysis of the oracle experiment in Section 3 using our learned mask predictor, comparing the theoretical maximum performance using our predicted mask in absence of the ground-truth high-resolution video with that of the oracle using the ground-truth. The results in Table 4 demonstrate that our mask performs well: it does not quite reach the theoretical optimum but still skips a significant proportion of patches while maintaining visual fidelity. Notably, the difference between swapping with our predicted mask and the ground-truth high-resolution is still consistently above 40 PSNR, which is an essentially imperceptible difference.

Profiling and Speed Design Choices. We measure the efficiency of other potential sparsity mechanisms in Figure 5a. Most works address the attention layer; we demonstrate that this loses 23% speed. Similarly, strategies like only masking out the attention queries or interspersing sparse and dense layers are not much faster than the baseline. On the other hand, our method achieves the largest speed-up while maintaining quality.

Threshold Effect. Finally, we analyze the effect of varying the threshold τ that defines the boundary between simple and complex patches in Figure 5b. We measure the PSNR compared to the original image for both AIGC and VideoLQ using the oracle mask, as well as the fraction of tokens removed by this procedure. Since the mask is dependent on the input data, the curves are offset from each other, but follow the same general trend: as we increase the threshold, the PSNR consistently

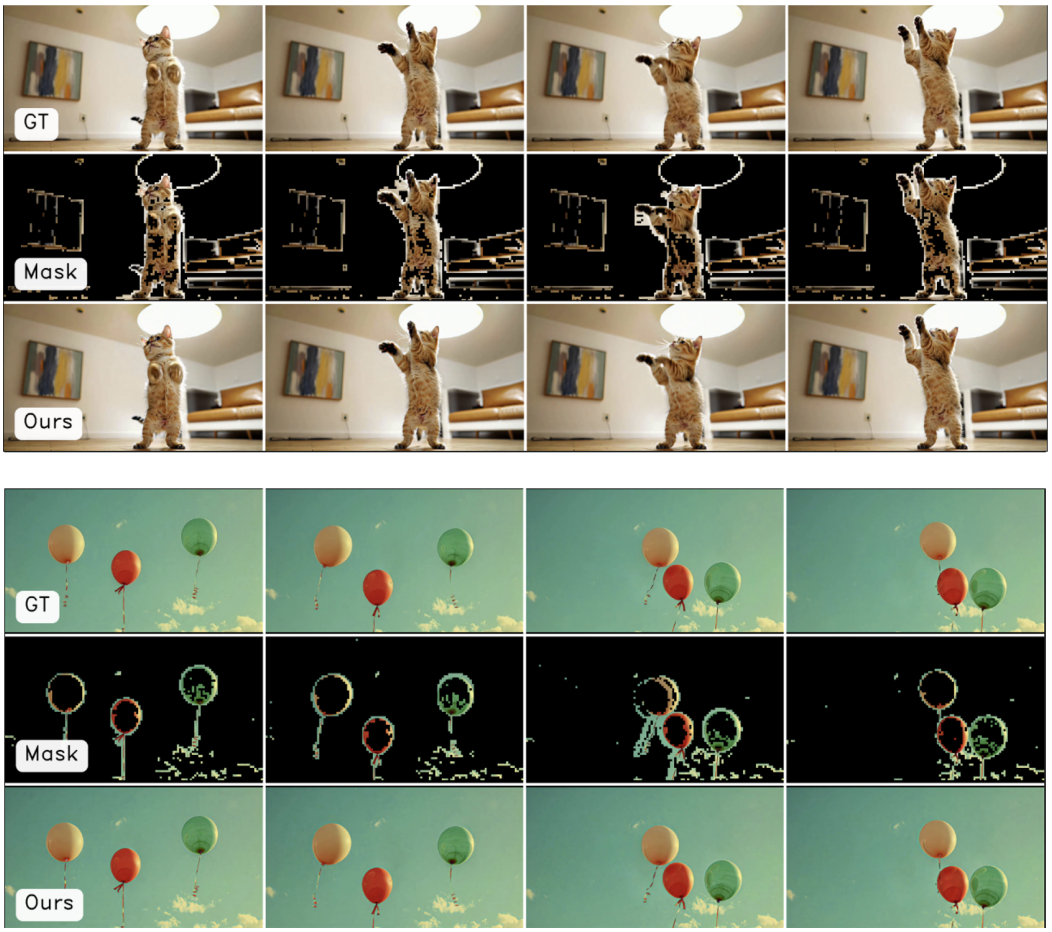


Figure 6: **Visual Comparison.** In the above two examples, the ground truth is on top, followed by the predicted mask, and our output. We produce perceptually indistinguishable results while only refining a small subset of the input video. Additional results are on our project page.

decreases as more and more complex patches are skipped rather than refined. From the plot, we can see that $\tau = 0.0002$ represents a reasonable balance between fidelity and token reduction, as it is paramount that any acceleration does not come at a cost to visual quality.

6 Conclusion

This paper tackles the acceleration of video super-resolution by focusing refinement only on complex regions that actually require it, rather than uniformly upscaling the entire input. We find that such regions make up a substantial fraction of video, and that using a cheap upsampling operation on them rather than expensive super-resolution can maintain quality. We propose SkipSR, which identifies these regions from the low-resolution input and routes them entirely around the transformer, composing them with the upscaled output afterward. Implementing our simple strategy enables significant acceleration, reducing the diffusion time by $1.8\times$ without visible loss in quality and the end-to-end generation time by $1.6\times$.

Limitations. A limitation of SkipSR is its lack of acceleration on video restoration tasks with severe degradations, where we are unable to skip patches due to widespread corruptions. However, our method focuses more on videos without such synthetic corruptions, where it achieves much larger speed-ups. SkipSR also does not provide significant speed-ups on videos with extremely crowded scenes or camera jitter, but we believe its strong performance on more typical videos is well worth it; SkipSR is no slower than the existing methods in these cases while performing much faster in more common, simpler videos.

Acknowledgments and Disclosure of Funding

This work was done while RC was an intern at ByteDance Seed, and supported by the NSF Graduate Research Fellowship Program. The authors declare no competing interests.

References

Kelvin C.K. Chan, Shangchen Zhou, Xiangyu Xu, and Chen Change Loy. Investigating tradeoffs in real-world video super-resolution. In *CVPR*, 2022.

Zheng Chen, Zichen Zou, Kewei Zhang, Xiongfei Su, Xin Yuan, Yong Guo, and Yulun Zhang. Dove: Efficient one-step diffusion model for real-world video super-resolution. *arXiv preprint arXiv:2505.16239*, 2025.

Mostafa Dehghani, Basil Mustafa, Josip Djolonga, Jonathan Heek, Matthias Minderer, Mathilde Caron, Andreas Steiner, Joan Puigcerver, Robert Geirhos, Ibrahim M Alabdulmohsin, et al. Patch n'pack: Navit, a vision transformer for any aspect ratio and resolution. *Advances in Neural Information Processing Systems*, 36:2252–2274, 2023.

Linwei Dong, Qingnan Fan, Yihong Guo, Zhonghao Wang, Qi Zhang, Jinwei Chen, Yawei Luo, and Changqing Zou. Tsd-sr: One-step diffusion with target score distillation for real-world image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.

Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024.

Yuanting Fan, Chengxu Liu, Nengzhong Yin, Changlong Gao, and Xueming Qian. Adadiffsr: Adaptive region-aware dynamic acceleration diffusion model for real-world image super-resolution. In *European Conference on Computer Vision*, pages 396–413. Springer, 2024.

Yu Gao, Haoyuan Guo, Tuyen Hoang, Weilin Huang, Lu Jiang, Fangyuan Kong, Huixia Li, Jiashi Li, Liang Li, Xiaojie Li, et al. Seedance 1.0: Exploring the boundaries of video generation models. *arXiv preprint arXiv:2506.09113*, 2025.

Zhengyang Geng, Mingyang Deng, Xingjian Bai, J Zico Kolter, and Kaiming He. Mean flows for one-step generative modeling. *arXiv preprint arXiv:2505.13447*, 2025.

Ali Hassani, Steven Walton, Jiachen Li, Shen Li, and Humphrey Shi. Neighborhood attention transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6185–6194, 2023.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research*, 23(47):1–33, 2022.

Xiaotao Hu, Jun Xu, Shuhang Gu, Ming-Ming Cheng, and Li Liu. Restore globally, refine locally: A mask-guided scheme to accelerate super-resolution networks. In *European Conference on Computer Vision*, pages 74–91. Springer, 2022.

Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qingyang Jin, Nattapol Chanpaisit, et al. Vbench: Comprehensive benchmark suite for video generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21807–21818, 2024.

Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5148–5157, 2021.

Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*, 2024.

Xiangtao Kong, Hengyuan Zhao, Yu Qiao, and Chao Dong. Classsr: A general framework to accelerate super-resolution networks by data characteristic. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12016–12025, 2021.

Vijay Anand Korthikanti, Jared Casper, Sangkug Lym, Lawrence McAfee, Michael Andersch, Mohammad Shoeybi, and Bryan Catanzaro. Reducing activation recomputation in large transformer models. *Proceedings of Machine Learning and Systems*, 5:341–353, 2023.

Shen Li, Yanli Zhao, Rohan Varma, Omkar Salpekar, Pieter Noordhuis, Teng Li, Adam Paszke, Jeff Smith, Brian Vaughan, Pritam Damania, et al. Pytorch distributed: Experiences on accelerating data parallel training. *arXiv preprint arXiv:2006.15704*, 2020.

Xingyang Li, Muyang Li, Tianle Cai, Haocheng Xi, Shuo Yang, Yujun Lin, Lvmin Zhang, Songlin Yang, Jinbo Hu, Kelly Peng, et al. Radial attention: O (nlog n)sparse attention with energy decay for long video generation. *arXiv preprint arXiv:2506.19852*, 2025.

Shanchuan Lin, Xin Xia, Yuxi Ren, Ceyuan Yang, Xuefeng Xiao, and Lu Jiang. Diffusion adversarial post-training for one-step video generation. *arXiv preprint arXiv:2501.08316*, 2025.

Songhua Liu, Zhenxiong Tan, and Xinchao Wang. Clear: Conv-like linearization revs pre-trained diffusion transformers up. *arXiv preprint arXiv:2412.16112*, 2024.

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.

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 *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.

Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021.

Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. Making a "completely blind" image quality analyzer. *IEEE Signal Processing Letters*, 20(3):209–212, 2013. doi: 10.1109/LSP.2012.2227726.

Brian B Moser, Stanislav Frolov, Federico Raue, Sebastian Palacio, and Andreas Dengel. Dynamic attention-guided diffusion for image super-resolution. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 451–460. IEEE, 2025.

Seungjun Nah, Sungyong Baik, Seokil Hong, Gyeongsik Moon, Sanghyun Son, Radu Timofte, and Kyoung Mu Lee. Ntire 2019 challenge on video deblurring and super-resolution: Dataset and study. In *CVPR Workshops*, June 2019.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022.

Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494, 2022a.

Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J. Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. In *TPAMI*, 2022b.

Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022.

Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.

Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, 2023.

Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding, 2021.

Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *AAAI*, 2023.

Jianyi Wang, Shanchuan Lin, Zhijie Lin, Yuxi Ren, Meng Wei, Zongsheng Yue, Shangchen Zhou, Hao Chen, Yang Zhao, Ceyuan Yang, et al. Seedvr2: One-step video restoration via diffusion adversarial post-training. *arXiv preprint arXiv:2506.05301*, 2025a.

Jianyi Wang, Zhijie Lin, Meng Wei, Yang Zhao, Ceyuan Yang, Chen Change Loy, and Lu Jiang. Seedvr: Seeding infinity in diffusion transformer towards generic video restoration. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 2161–2172, 2025b.

Longguang Wang, Xiaoyu Dong, Yingqian Wang, Xinyi Ying, Zaiping Lin, Wei An, and Yulan Guo. Exploring sparsity in image super-resolution for efficient inference. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4917–4926, 2021.

Haoning Wu, Erli Zhang, Liang Liao, Chaofeng Chen, Jingwen Hou, Annan Wang, Wenxiu Sun, Qiong Yan, and Weisi Lin. Exploring video quality assessment on user generated contents from aesthetic and technical perspectives. In *International Conference on Computer Vision (ICCV)*, 2023.

Haocheng Xi, Shuo Yang, Yilong Zhao, Chenfeng Xu, Muyang Li, Xiuyu Li, Yujun Lin, Han Cai, Jintao Zhang, Dacheng Li, et al. Sparse videogen: Accelerating video diffusion transformers with spatial-temporal sparsity. *arXiv preprint arXiv:2502.01776*, 2025.

Xi Yang, Chenhang He, Jianqi Ma, and Lei Zhang. Motion-guided latent diffusion for temporally consistent real-world video super-resolution. In *European Conference on Computer Vision (ECCV)*, 2024.

Peng Yi, Zhongyuan Wang, Kui Jiang, Junjun Jiang, and Jiayi Ma. Progressive fusion video super-resolution network via exploiting non-local spatio-temporal correlations. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3106–3115, 2019.

Peng Yi, Zhongyuan Wang, Kui Jiang, Zhenfeng Shao, and Jiayi Ma. Multi-temporal ultra dense memory network for video super-resolution. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(8):2503–2516, 2020. doi: 10.1109/TCSVT.2019.2925844.

Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and Bill Freeman. Improved distribution matching distillation for fast image synthesis. *Advances in neural information processing systems*, 37:47455–47487, 2024.

Jintao Zhang, Chendong Xiang, Haofeng Huang, Jia Wei, Haocheng Xi, Jun Zhu, and Jianfei Chen. Spargeattn: Accurate sparse attention accelerating any model inference. In *International Conference on Machine Learning (ICML)*, 2025a.

Peiyuan Zhang, Yongqi Chen, Runlong Su, Hangliang Ding, Ion Stoica, Zhengzhong Liu, and Hao Zhang. Fast video generation with sliding tile attention. *arXiv preprint arXiv:2502.04507*, 2025b.

Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.

Shangchen Zhou, Peiqing Yang, Jianyi Wang, Yihang Luo, and Chen Change Loy. Upscale-a-video: Temporal-consistent diffusion model for real-world video super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2535–2545, 2024.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction (Section 1) clearly state the three contributions—identifying simple regions in videos, proposing SkipSR with mask prediction and skip-aware diffusion, and validating it experimentally—and these claims match the experimental results in Section 5 (up to $1.8\times$ diffusion speed-up and $\sim$60% end-to-end latency reduction with no perceptible quality loss).

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Section 3 discusses cases where SkipSR yields fewer skippable patches (e.g., the synthetically corrupted YouHQ benchmark with heavy white noise everywhere), and the dedicated **Limitations** paragraph at the end of Section 5/conclusion notes that SkipSR offers little speed-up on heavily-degraded restoration, extremely crowded scenes, or videos with strong camera jitter.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best

judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

   Justification: The paper presents an empirical method for accelerating video super-resolution and contains no theorems or formal proofs.

   Guidelines:

   - The answer NA means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: Section 4 describes the mask predictor architecture (4 layers of 3D convolution with ReLU activations, stride-2 first conv, latent-space input/output shapes and patch sizes), the skip-aware modifications to RoPE and shifted-window attention, and the training procedure. Section 5 reports the base checkpoints used (public SeedVR-3B and SeedVR2-7B), training data (the same as Wang et al. [2025b]), batch size, sequence/data parallelism, GPU configuration, and the one-step distillation procedure built on APT Lin et al. [2025]. The patch-skippability criterion (Eq. (1)), threshold $\tau = 0.0002$, and downsampling factor are stated explicitly.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example

(a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
(b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
(c) If the contribution is primarily a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [No]

   Justification: Code and trained checkpoints are not released with this submission. The base diffusion model SeedVR/SeedVR2 Wang et al. [2025b,a] is publicly available, and our mask-predictor architecture and training procedure are described in sufficient detail in Sections 4 and 5 for re-implementation. We will release code and weights upon publication.

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: Section 5 (Implementation Details and Experimental Settings) reports the base checkpoints, training data, batch size ($\sim$100 frames at 720p), parallelism strategy, GPU count (40 A100-80G), the mask predictor's separate training run (8 GPUs, 10k iterations), evaluation datasets (VideoLQ, AIGC-30, SPMCS, UDM10, REDS30, YouHQ-40), and metrics (NIQE, CLIP-IQA, MUSIQ, DOVER, PSNR, SSIM, LPIPS, DISTS). The skippability threshold τ and downsampling factor used in the analysis are also specified.

   Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: Quantitative super-resolution metrics in Tables 2 and 3 report point estimates per dataset. Re-running diffusion-based video SR over the full evaluation suites for multiple seeds is prohibitively expensive (each run requires tens of minutes per video on A100-80G GPUs across 40+ videos and several datasets), so error bars were not collected. The user-study Good–Same–Bad preference score in Table 4a is itself an aggregate measure intended to reduce single-rater variance.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
   - For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
   - If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Implementation Details in Section 5 state that the main super-resolution model is fine-tuned on 40 NVIDIA A100-80G GPUs at $\sim$100 frames per batch at 720p, and the lightweight mask predictor is trained for 10k iterations on 8 A100-80G GPUs. End-to-end latencies and per-component breakdowns (diffusion transformer, VAE, mask predictor) on A100-80G are reported in Figure 4b and the surrounding text.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.

- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: The work uses existing publicly-available video benchmarks and pre-trained models, performs a small internal preference study with no personally-identifiable subject data, and complies with the NeurIPS Code of Ethics. Anonymity is preserved in this submission.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: SkipSR is an efficiency improvement applied on top of existing diffusion-based super-resolution models. The primary positive impact is reducing the energy and compute cost of high-resolution video generation and restoration, broadening accessibility. As a possible negative impact, accelerated video SR could be used to upscale outputs of generative models (including potentially deceptive or non-consensual content) more efficiently; however, SkipSR does not by itself improve generative quality or capability, and downstream safety concerns are inherited from the underlying generative model rather than introduced by our acceleration technique.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [N/A]

Justification: This submission does not release any new pretrained generative model, image/video generator, or scraped dataset. SkipSR is an architectural acceleration technique applied on top of existing publicly-released super-resolution models, whose own release came with their authors' safeguards.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All used assets are credited via citation: the SeedVR / SeedVR2 base models Wang et al. [2025b,a], evaluation datasets VBench Huang et al. [2024], VideoLQ Chan et al. [2022], YouHQ-40 Zhou et al. [2024], SPMCS Yi et al. [2019], UDM10 Yi et al. [2020], and REDS30 Nah et al. [2019], as well as standard libraries and metrics (LPIPS Zhang et al. [2018], MUSIQ Ke et al. [2021], DOVER Wu et al. [2023], CLIP-IQA Wang et al. [2023], NIQE Mittal et al. [2013]). All are used for non-commercial academic research consistent with their respective licenses. Code and weights released alongside this work upon publication will be distributed under the CC BY 4.0 license.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [N/A]

Justification: No new datasets or pretrained models are released with this submission. The mask predictor architecture and training procedure are documented in Sections 4 and 5 for the purpose of reproducibility, but no asset is being distributed.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [N/A]

Justification: The Good–Same–Bad evaluation in Section 5 was an internal preference study with three domain experts who are employees of the authors' institutions (no crowdworkers, no external compensated participants, no personally identifiable data collected). It does not constitute crowdsourcing or formal human-subjects research as defined by the NeurIPS guidelines.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [N/A]

Justification: The paper does not involve crowdsourcing or formal human-subjects research; the small internal preference study described in Section 5 carries no risk to participants and does not require IRB approval.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: Large language models are not used as a component of the core methodology in this research.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.