

VIRES: Video Instance Repainting via Sketch and Text Guided Generation

Shuchen Weng^{1,2†‡} Haojie Zheng^{3,4†} Peixuan Zhang⁵ Yuchen Hong^{1,2}
Han Jiang³ Si Li⁵ Boxin Shi^{1,2*}

¹State Key Laboratory for Multimedia Information Processing, School of Computer Science, Peking University

²National Engineering Research Center of Visual Technology, School of Computer Science, Peking University

³OpenBayes Information Technology Co., Ltd. ⁴School of Software and Microelectronics, Peking University

⁵School of Artificial Intelligence, Beijing University of Posts and Telecommunications

scweng@baai.ac.cn, suimu@stu.pku.edu.cn, {pxzhang, lisi}@bupt.edu.cn
yuchenhong.cn@gmail.com, hahn@openbayes.com, shiboxin@pku.edu.cn



Figure 1. Our VIRES model demonstrates powerful video editing capabilities with sketch and text guidance, as shown in four typical scenarios: (a) Repainting the color and style of the man’s shirt. (b) Replacing the pickup truck with the dark-colored SUV. (c) Generating a running corgi within a video clip. (d) Removing a specified football from a video clip.

Abstract

We introduce VIRES, a video instance repainting method with sketch and text guidance, enabling video instance repainting, replacement, generation, and removal. Existing

approaches struggle with temporal consistency and accurate alignment with the provided sketch sequence. VIRES leverages the generative priors of text-to-video models to maintain temporal consistency and produce visually pleasing results. We propose the Sequential ControlNet with the standardized self-scaling, which effectively extracts structure layouts and adaptively captures high-contrast sketch details. We further augment the diffusion transformer back-

[†] Equal contribution.

[‡] Now at Beijing Academy of Artificial Intelligence.

* Corresponding author.

bone with the sketch attention to interpret and inject fine-grained sketch semantics. A sketch-aware encoder ensures that repainted results are aligned with the provided sketch sequence. Additionally, we contribute the VIRESET, a dataset with detailed annotations tailored for training and evaluating video instance editing methods. Experimental results demonstrate the effectiveness of VIREs, which outperforms state-of-the-art methods in visual quality, temporal consistency, condition alignment, and human ratings. The code, dataset and pretrained models are available at: <https://hjzheng.net/projects/VIREs>.

1. Introduction

Video is a crucial medium for people to capture and communicate their experiences and ideas. Traditionally, maintaining temporal consistency when editing videos requires specialized skills and considerable effort. With the advancements in diffusion models [6, 34], recent text-guided video editing methods [27, 43] attract considerable interest and offer the potential to revolutionize filmmaking and entertainment through a controllable approach.

While text can convey general video editing goals, it often leaves considerable room for interpretation of the user’s intent (e.g., logo details on clothing). As the improvement, recent works [14, 18, 47] introduce additional guidance (e.g., sketch) for more precise control. However, relying on pre-trained text-to-image models (e.g., Stable Diffusion [31]), these zero-shot methods struggle with temporal consistency, inevitably resulting in flickering. As an alternative, VideoComposer [38] fine-tunes the text-to-image model and incorporates temporal modeling layers. However, its focus on compositionality limits the accurate alignment with the provided sketch sequence, resulting in sub-optimal fine-grained instance-level edits.

In this paper, we propose the **VIREs**, a **Video Instance RE**painting method with **S**ketch and text guidance. As illustrated in Fig. 1, VIREs facilitates four application scenarios to repaint instances within the provided video (i.e., video instance repainting, replacement, generation, and removal). Leveraging the generative priors of text-to-video models (i.e., Open-Sora [52]), our approach maintains temporal consistency and produces visually pleasing results. To provide precise user control over instance properties, we design tailored modules that process sketch for the structure (e.g., pose and movement) and text for the appearance (e.g., color and style). Specifically, we present the Sequential ControlNet to effectively extract structure layouts and the standardized self-scaling to adaptively capture high-contrast structure details. Additionally, to interpret and inject fine-grained sketch semantics into the latent space, we augment the Diffusion Transformer (DiT) backbone with the sketch attention. Finally, the sketch-aware encoder aligns repainted results with multi-level texture fea-

tures during the latent code decoding.

We further contribute the VIRESET, a new dataset for training and evaluating video instance editing methods. VIRESET provides 85K training videos and 1K evaluation videos, sourced from the SA-V [30]. For each video, we generate text annotations from the large language model (LLM) [46] and extract sketch sequences using HED edge detection [45]. Our contributions are summaries as follows:

- We propose the first DiT-based framework for controllable video instance editing, along with a tailored dataset for training and evaluating relevant editing methods.
- We present the Sequential ControlNet and the standardized self-scaling to effectively extract structure layouts and adaptively capture high-contrast sketch details.
- We introduce the sketch attention to interpret and inject fine-grained sketch semantics and sketch-aware encoder to align repainted results with multi-level texture features.

2. Related work

2.1. Diffusion-based image generation

Diffusion models [6, 34] have demonstrated remarkable advancements in image generation [31, 33]. These models learn to reverse the forward diffusion process, generating images by iterative denoising a sample from a standard Gaussian distribution. To control this generation process, researchers use adapter mechanisms to incorporate additional conditions [22, 50]. Leveraging the generative priors encapsulated in the pre-trained generation model, users can customize content via fine-tuning [17, 32] and edit arbitrary images via training-free strategies [21, 24]. These advancements expand the application of diffusion models to various image-based tasks, e.g., super resolution [4, 39], image colorization [41, 42], and reflection removal [8, 53]. Given the similar spatial representations between images and videos, researchers are inspired to effectively lift these pre-trained image priors to the video domain.

2.2. Diffusion-based video generation

Early text-to-video generation works attempt to extend the UNet-based denoising network with space-only 3D convolutions, training them from scratch [7]. However, this approach proves computationally expensive, leading researchers [5, 36] to explore fine-tuning pre-trained text-to-image models [31]. This fine-tuning typically introduces additional temporal modules (e.g., temporal convolutions [35] or temporal attention [2]) to improve the temporal consistency, while freezing most of the pre-trained denoising network weights to preserve spatial generative priors. To further reduce generation costs, researchers explore zero-shot text-guided video editing that solely performs during the inference phase [14, 51]. Recently, impressive text-to-video generation results from OpenAI [20] have re-

newed interest in training models from scratch, leveraging the diffusion transformer architecture [48, 52, 54]. Given the significant improvements demonstrated on video generation benchmarks [10], we build upon this foundation and develop modules for controllable video generation.

3. Dataset

Existing datasets [1, 25] face challenges in providing precise instance masks due to temporal inconsistencies (*e.g.*, occlusions and reappearances) and variations in visual appearance (*e.g.*, complex motion, lighting, and scale). Therefore, we present the VIRESET, a dataset including high-quality instance masks tailored for training and evaluating video instance editing methods.

Data acquisition. Our initial video samples are sourced from SA-V [30], a dataset includes indoor (54%) and outdoor (46%) scenes recorded across 47 countries by diverse participants. The original videos have an average duration of 14 seconds and a resolution of 1401×1037 pixels, with multiple manual and automatic instance mask annotations provided at 6 FPS. SA-V adopts a data engine with a verification step to ensure high-quality instance mask annotations. To further improve temporal consistency for video instance editing, we leverage the pre-trained SAM-2 model [30] to annotate the intermediate frames, increasing the mask annotation rate to 24 FPS.

Instance selection. We filter instances to include only those covering at least 10% of the frame area. Additionally, We only use instances present for at least 51 consecutive frames, from which we randomly sample 51-frame clips. Each selected instance is then cropped from the original video with a small margin around its bounding box and resized to 512×512 resolution.

Additional annotations. We adopt the pre-trained PLLaVa [46] to generate text descriptions for each cropped video clip, capturing their visual appearance. To ensure the quality, we recruit 10 volunteers to review 1% samples, achieving a 91% acceptance rate. Following VideoComposer [38], we extract sketch sequences using HED edge detection [45] to provide the structure guidance.

Dataset statistics. The final VIRESET includes 85K video clips for training and 1K video clips for evaluation. Each clip consists of 51 frames at 24 FPS with a resolution of 512×512 , accompanied by a sketch sequence for the structure and a text for the appearance.

4. Methodology

This section begins with an overview of the framework (Sec. 4.1). We then delve into the proposed modules: the Sequential ControlNet with the standardized self-scaling (Sec. 4.2), the sketch-based DiT backbone with the sketch attention (Sec. 4.3), and the sketch-aware encoder to im-

prove the decoding process (Sec. 4.4).

4.1. Overview

Given an N -frame original video clip $x = \{x^i\}_{i=1}^N$, a corresponding sketch sequence $s = \{s^i\}_{i=1}^N$ presenting the structure, a text description y describing the appearance, and an instance mask sequence $m = \{m^i\}_{i=1}^N$ indicating the repainted regions, VIRESET is proposed to repaint the specific video instance. The overall pipeline is illustrated in Fig. 2.

Input encoding. The original video clip x is encoded into the latent code z using the pre-trained spatial-temporal VAE encoder \mathcal{E}^x [52] as $z = \mathcal{E}^x(x)$. The sketch sequence s is processed by our proposed Sequential ControlNet \mathcal{F}^s to effectively extract structure layouts as $f^s = \mathcal{F}^s(s)$. The text y is encoded with the pre-trained text encoder \mathcal{F}^y [29] to obtain the word embeddings as $f^y = \mathcal{F}^y(y)$.

Condition injection. To adaptively capture the high-contrast structure details of the sketch sequence, we introduce the standardized self-scaling. We further augment the DiT backbone with the sketch attention, which interprets and injects fine-grained sketch semantics into the latent space. Word embeddings are injected into the DiT backbone with the pre-trained cross-attention modules.

Forward process. We adopt the Flow Matching formulation [19] for robust and stable training. For simplicity, we consider a linear path between the sampled latent code z and the Gaussian noise $\epsilon \sim \mathcal{N}(0, 1)$:

$$\hat{z}_t = tz + (1 - t)\epsilon, \quad (1)$$

where $t \in [0, 1]$ is the timestep, with $\hat{z}_0 = \epsilon$ and $\hat{z}_1 = z$.

Latent masking. After concatenating the mask with the latent code, we further selectively add noise according to the instance mask, allowing the model to repaint specific instances within the video:

$$z_t = \hat{z}_t \odot \hat{m} + z \odot (1 - \hat{m}), \quad (2)$$

where \hat{m} is the downsampled mask sequence indicating where the repainted instance places.

Backward process. The reverse diffusion process can be defined by the Ordinary Differential Equation (ODE):

$$dz/dt = v_\theta(z_t, t, s, m, y), \quad (3)$$

where $v_\theta(z_t, t, s, m, y)$ is the estimated vector field that guides the generation process. By solving the ODE from $t = 0$ to $t = 1$, we transform the noise into the latent code.

Denoising learning. We adopt the flow matching objective [19] to train the denoising network $v_\theta(z_t, t, s, m, y)$. In the case of the linear path, the target velocity v_t simplifies to $z - \epsilon$. We optimize the estimated velocity $v_\theta(z_t, t, s, m, y)$ towards the target velocity v_t by minimizing the loss:

$$\mathcal{L}_{\text{fm}} = \mathbb{E}_{t, z, \epsilon} [\|v_t(z_t) - v_\theta(z_t, t, s, m, y)\|^2]. \quad (4)$$

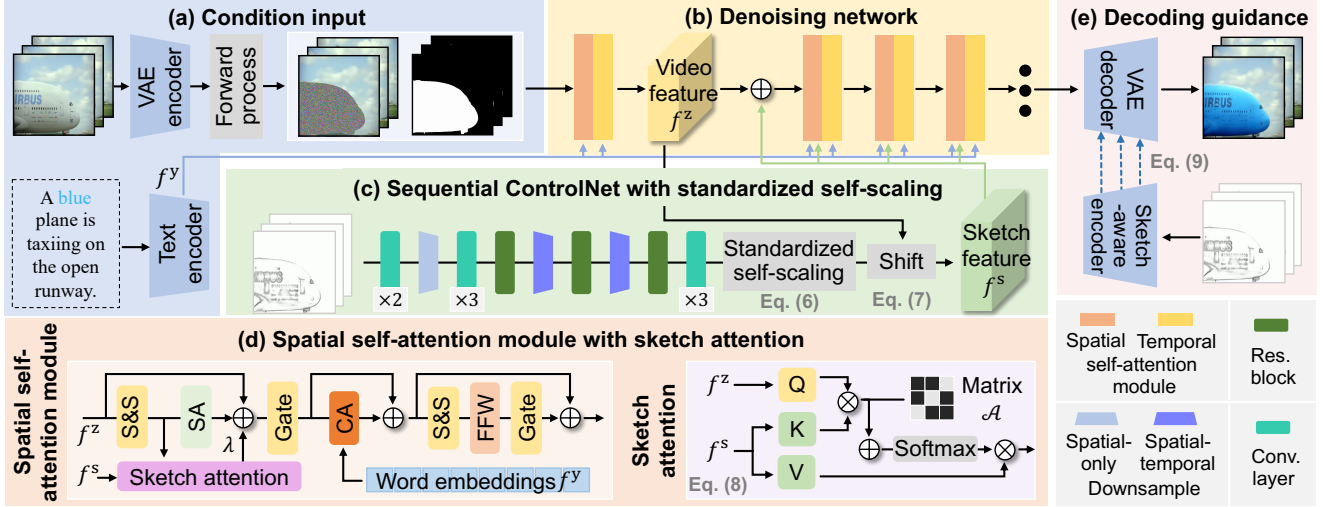


Figure 2. The pipeline of VIRES model. (a) The input video clip and corresponding text descriptions are encoded using independent encoders. Noise is selectively added to the latent code according to the instance mask. (b) This noised latent code is fed into the denoising network, composed of stacked transformer blocks with spatial and temporal self-attention modules, trained to reverse the diffusion process. (c) The Sequential ControlNet extracts structure layouts and injects them into the denoising network with the standardized self-scaling. (d) The sketch attention is implemented as a parallel branch within the spatial self-attention module, injecting structure semantics into subsequent transformer blocks. (e) A sketch-aware encoder additionally provides multi-level texture features during decoding, generating the repainted video clip from the denoised latent code.

Latent decoding. The denoised latent code z' is then decoded into a video clip \bar{x} using the pre-trained spatial-temporal VAE decoder \mathcal{D} [52]. To further align the repainted results with the sketch, we introduce a sketch-aware encoder \mathcal{E}^s , which adopts the same VAE encoder architecture to provide multi-level texture features as the guidance during decoding:

$$\bar{x} = \mathcal{D}(z', \mathcal{E}^s(s)). \quad (5)$$

4.2. Sequential sketch feature extraction

As presented in Fig. 2 (c), we propose the Sequential ControlNet to extract structure layouts from the sketch sequence, followed by the standardized self-scaling to adaptively capture the high-contrast structure details.

Sequential ControlNet. Previous ControlNet [50] is designed for image editing, which lacks temporal consistency and can produce flickering artifacts when applied to video-based tasks. This motivates us to introduce Sequential ControlNet for the video instance repainting. Our Sequential ControlNet includes convolutional layers, residual blocks, and downsampling layers. Each convolutional layer consists of a 3D causal convolution [49], followed by a Group Normalization [44] and a SiLU activation function [3], to effectively capture spatial-temporal dependencies between frames. Residual blocks consist of two such convolutional layers with a residual connection. Downsampling layers are convolutional layers with a stride of 2, applied to compress the spatial or spatial-temporal dimensions. Since the

Further variations are discussed in the Supp.

spatial-temporal VAE encoder [52] downsamples the original video clip by $4 \times$ temporally and $64 \times$ spatially, we adopt one spatial downsampling layer and two spatial-temporal downsampling layers to match the feature map dimensions of the DiT backbone. To align the embedding channels, we progressively increase the number of channels. Specifically, we increase the embedding channels in the shallow convolutional layers and double the channels at each downsampling layer. The final three layers maintain a high channel dimension to effectively extract structure layouts from the sketch sequence.

Standardized self-scaling. Feature modulation has proven effective in conditional image editing (e.g., AdaIN [9], FiLM [26], and SPADE [23]). Observing that the sketch has high-contrast transitions between black lines and the white background, we introduce the standardized self-scaling to adaptively capture sketch details, instead of performing simply addition. Specifically, we use sketch features f^s extracted by the Sequential ControlNet and standardize them to scale the features themselves, effectively highlighting the high-contrast regions:

$$\hat{f}^s = ((f^s - \mu(f^s)) / \sigma(f^s)) \odot f^s, \quad (6)$$

where $\mu(\cdot)$ and $\sigma(\cdot)$ represent the function of mean and standard deviation, respectively. We then shift the feature domain from sketch to video by aligning their means:

$$\bar{f}^s = \hat{f}^s - \mu(\hat{f}^s) + \mu(f^z), \quad (7)$$

where f^z represents the video features. To reduce computational cost, standardized self-scaling is applied only once

to the first transformer block of the DiT backbone.

4.3. Latent sketch semantic interpretation

As illustrated in Fig. 2 (d), following the standardized self-scaling that provides the high-contrast structure details, we further augment the DiT backbone with the sketch attention to interpret and inject fine-grained sketch semantics within the latent space.

DiT backbone. The VIRES framework builds upon the pre-trained text-to-video generation model [52], leveraging its DiT architecture that consists of stacked transformer blocks. Each block incorporates separate spatial and temporal self-attention modules to capture intra-frame and inter-frame dependencies, respectively. In each module, Self-Attention (SA) and Feed-Forward Network (FFN) extract contextual features, Cross-Attention (CA) injects semantics from word embeddings, and the scale and shift (S&S) and the gate modulate the latent code with timestep embeddings. We concatenate the latent code with the instance mask before feeding it to the transformer, and patchify both the latent code and the sketch sequence into non-overlapping $1 \times 2 \times 2$ tokens, resulting in a $4 \times$ spatial downsampling.

Sketch attention. To interpret and inject sketch semantics into the latent space, we augment the DiT backbone with the sketch attention within each spatial self-attention module except for the first. The sketch attention incorporates a predefined binary matrix \mathcal{A} to indicate correspondences between the latent code and the sketch sequence:

$$\bar{f}^z = \text{Softmax}((QK^\top + \mathcal{A})/\sqrt{C})V, \quad (8)$$

where Q and K, V are transformed features from the video features f^z and the extracted structure layouts f^s , respectively. C is the number of embedding channels. Sketch attention is implemented as a parallel branch, and its outputs are added with a learnable scaling parameter λ , allowing adaptive weighting of injected sketch semantics.

4.4. Visual sketch texture guidance

As illustrated in Fig. 2 (e), during decoding the latent code into a video clip, we additionally adopt a sketch-aware encoder, which provides multi-level texture features to the spatial-temporal VAE decoder and aligns the structure of repainted results with the provided sketch sequence.

Spatial-temporal VAE. We use the pre-trained spatial-temporal VAE [52] to encode the original video clip into a latent code and subsequently decode it back into a video clip. Specifically, the spatial-temporal encoder firstly uses a stack of 2D vanilla convolution blocks with three downsampling layers, achieving $64 \times$ spatially compression. Following this, 3D causal convolution blocks capture spatial-temporal dependencies with two downsampling layers applied only to the temporal dimension, resulting in $4 \times$ temporal compression. The spatial-temporal VAE decoder mir-

rors the architecture of the encoder, using 3D causal convolutions for spatial-temporal modeling and upsampling the temporal dimension before the spatial dimension.

Sketch-aware encoder. During decoding the latent code into a video clip, the spatial-temporal VAE decoder can produce numerous visually pleasing results. However, lacking guidance from the sketch, many results may not align with expected visual representations. Therefore, we introduce the sketch-aware encoder to guide this process. Specifically, the sketch-aware encoder adopts the same architecture as the VAE encoder, extracting multi-level texture features before each downsampling layer. Then, these features are added to the decoder at each corresponding level:

$$\mathcal{D}(z', \mathcal{E}^s(s))_i = \mathcal{E}^s(s)_i + \mathcal{D}(z')_i, \quad (9)$$

where subscript i denotes the feature level. We train the sketch-aware encoder by minimizing a combination of an SSIM loss [40] \mathcal{L}_{sm} for structure alignment, L1 loss \mathcal{L}_1 for reconstruction fidelity, perceptual loss [12] \mathcal{L}_{pc} for sharp details, and KL loss [16] \mathcal{L}_{kl} for regularization. The combined loss is calculated for each frame and then summed over all frames:

$$\mathcal{L}_{\text{vae}} = \sum_i^N \mathcal{L}_{\text{sm}} + \lambda_1 \mathcal{L}_1 + \lambda_2 \mathcal{L}_{\text{pc}} + \lambda_3 \mathcal{L}_{\text{kl}}, \quad (10)$$

where i indexes the frames. $\lambda_1 = 1$, $\lambda_2 = 0.1$, and $\lambda_3 = 1 \times 10^{-6}$ are hyper-parameters that are not sensitive to variations in a certain range.

5. Experiment

Training details. We initialize our spatial-temporal VAE and DiT backbone with pre-trained weights from OpenSora v1.2 and train the VIRES model to repaint videos at a 512×512 resolution. The training process consists of three steps: (i) With the pre-trained spatial-temporal VAE encoder and decoder frozen, we train the sketch-aware encoder for 22K steps; (ii) With both VAEs and the pre-trained DiT backbone frozen, we train the Sequential ControlNet, the standardized self-scaling, and output matrices of the self-attention layers for 35K steps; and (iii) After augmenting the DiT backbone with the sketch attention, we train the model for an additional 45K steps. All experiments are conducted on 8 H100 GPUs using the Adam optimizer [15] with a learning rate of 1×10^{-5} .

Evaluation Datasets. We evaluate our approach on two datasets: our collected VIRESET and the widely-used DAVIS [25]. 50 videos are randomly selected from each dataset for evaluation. The condition annotation process of the DAVIS is consistent with that of VIRESET, and details are presented in Sec. 3.



Figure 3. Visual quality comparisons of video editing methods with text and sketch guidance.

5.1. Quantitative evaluation metrics

We quantitatively evaluate performance using five metrics: (i) the **Peak Signal-to-Noise Ratio (PSNR)** [11] for the visual perceptual quality; (ii) the **Structural Similarity Index Measure (SSIM)** [40] for spatial structure consistency; (iii) the **Warp Error (WE)**, calculated by warping repainted video clips with the estimated optical flow, following VidToMe [18] to evaluate motion accuracy; (iv) the **Frame Consistency (FC)**, measuring cosine similarity between consecutive frames for temporal consistency using the CLIP image encoder [28]; and (v) the **Text Consistency (TC)**, measuring cosine similarity between repainted video clips and text descriptions for text-video alignment using the VideoCLIP-XL [37].

5.2. Comparison with state-of-the-art methods

We compare our approach with state-of-the-art video editing methods for generating 51-frame video clips, including VideoComposer [38], Text2Video-Zero [14], Rerender [47], VidToMe [18], and RAVE [13] using their publicly released models with default configurations.

Qualitative comparisons. We present visual quality comparison results with aforementioned methods [13, 14, 18, 38, 47]. As shown in Fig. 3, the left column edits instance appearance using text descriptions while preserving structure with extracted sketches. Conversely, the right column maintains its appearance using annotated text descriptions and edits structure via sketch modifications. As a result, Rerender [47] produces excessively smooth textures (e.g.,

Table 1. Quantitative experiment results of comparison and ablation. \uparrow (\downarrow) means higher (lower) is better. All scores except PSNR are percentages. Throughout the paper, best performances are highlighted in **bold**.

Method	VIRESET							DAVIS						
	PSNR \uparrow	SSIM \uparrow	WE \downarrow	FC \uparrow	TC \uparrow	VQE \uparrow	TAE \uparrow	PSNR \uparrow	SSIM \uparrow	WE \downarrow	FC \uparrow	TC \uparrow	VQE \uparrow	TAE \uparrow
Comparison with state-of-the-art methods														
Rerender	11.85	35.21	5.77	90.40	15.78	6.40	5.60	11.79	29.67	8.01	89.75	15.56	7.20	8.00
VidToMe	16.40	57.35	5.25	89.14	16.13	4.00	6.40	15.82	52.63	6.51	90.00	15.76	6.00	7.60
Text2VideoZero	10.04	27.74	9.03	92.13	15.80	5.20	8.40	11.11	26.63	10.38	92.15	15.79	3.20	5.60
RAVE	18.99	70.81	6.18	92.17	15.89	18.00	23.60	19.08	68.35	8.74	91.99	15.81	23.20	25.60
VideoComposer	16.57	53.99	5.15	89.82	15.49	12.00	12.80	17.15	49.70	6.52	90.24	15.70	9.60	10.80
Ours (VIRES)	23.87	77.87	5.09	92.23	16.19	54.40	43.20	25.36	76.74	6.49	92.27	16.08	50.80	42.40
Ablation study														
W/o SSS	22.72	75.81	5.17	91.94	16.13	N/A	N/A	24.59	75.66	6.55	91.88	16.02	N/A	N/A
W/o SS	23.09	76.31	5.13	91.99	16.16	N/A	N/A	24.95	75.84	6.52	91.96	16.04	N/A	N/A
W/o SA	22.89	76.45	5.12	92.01	16.08	N/A	N/A	24.86	76.06	6.54	91.92	15.99	N/A	N/A
W/o SE	22.94	76.13	5.17	91.80	16.11	N/A	N/A	24.72	75.78	6.58	91.89	16.02	N/A	N/A

Fig. 3 left, the unnatural smoothness of the horse and tree). VidToMe [18] generates a cartoonish effect (e.g., Fig. 3 left, the stylized appearance of the entire frame). Text2VideoZero [14] presents inaccurate color tone (e.g., Fig. 3 right, the white jersey of the player). RAVE [13] suffers from inter-frame flickering (e.g., Fig. 3 right, the player’s neck flashing with frame brightness fluctuations). VideoComposer [38] struggles to follow the target structure (e.g., Fig. 3 right column, the missing sleeves of the player). In contrast, our VIRES faithfully repaint the video instance.

Quantitative comparisons. We show comprehensive quantitative comparison results in Tab. 1, demonstrating that VIRES outperforms relevant state-of-the-art methods across all five quantitative metrics on both datasets.

User study. In addition to qualitative and quantitative comparisons, we conduct two user studies to evaluate human preference for our results: (i) **Visual Quality Evaluation (VQE):** Participants are shown the edited results produced by VIRES and relevant video editing methods. They are asked to select the most visually pleasing video clip. (ii) **Text Alignment Evaluation (TAE):** Given a corresponding text description, participants are instructed to select the video clip from the same set of edited results that best matches the description. For each experiment, we randomly select 10 samples from each dataset, and recruit 25 volunteers from Amazon Mechanical Turk (AMT) to provide independent evaluations. As shown in Tab. 1, our model achieves the highest preference scores in both experiments.

5.3. Ablation study

We discard several modules and establish four baselines to study the impact of the corresponding modules. The evaluation scores and repainted results of the ablation study are presented in Tab. 1 and Fig. 4, respectively.

W/o Standardized Self-Scaling (SSS). We replace the standardization self-scaling with simple addition when ini-

tially injecting the structure layout. This leads to distorted textures in the repainted results (Fig. 4 third column, an unusual texture appears around the fish’s eye).

W/o Self-Scaling (SS). We discard the self-scaling, but preserve the standardization for the extracted structure layout before adding them to the DiT backbone. As a result, the model struggles to capture structure details (Fig. 4 second column, the fish head appears less detailed).

W/o Sketch Attention (SA). We remove the sketch attention in the DiT backbone, thereby removing the third stage of training. This prevents the model from accurately interpreting fine-grained sketch semantics (Fig. 4 fourth column, abrupt yellow patches appear on the fish’s body).

W/o Sketch-aware Encoder (SE). We disable the sketch-aware encoder that provides multi-level features to guide the decoding process, resulting in repainted results that are not well aligned with the expected visual representations (Fig. 4 first column, the fish body appearing blurry and flat).

5.4. Application

In addition to four application scenarios illustrated in Fig. 1 (i.e., video instance repainting, video instance replacement, custom instance generation, and specified instance removal), we present additional three applications.

Sketch-to-video generation. VIRES is capable of generating entire video clips from sketch sequences with text guidance, not limited to repainting specific instances. This is achieved by discarding the original video content and expanding the instance mask to the entire frame (e.g., Fig. 5 first row, generating a realistic lotus flower from a provided sketch sequence).

Sparse sketch guidance. To reduce user effort when repainting video instances, we incorporate sparse sketch guidance into the VIRES using two fine-tuning strategies: (i) random dropout of sketch frames with 20% probability; and splitting the video into 2^d intervals with 80% probability,




Figure 4. Ablation study results of different VIRES variants.

where $d \in \{0, \dots, 4\}$ is a randomly selected interval index. For each interval, only the first sketch frame is provided, with missing frames replaced by black images. As a result, VIRES can repaint video instances with even a single sketch frame (e.g., Fig. 5 second row, repainting the pattern on the person’s clothes with the first sketch frame).

Long-duration video repainting. Due to the computational resource limitation, processing the full long-duration video at once is challenging. However, VIRES can effectively repaint long-duration videos by randomly unmasking the frames during training (e.g., the first k frames). Specifically, we initially repaint the first 51-frame video clip, leveraging the last 17 frames of this clip as a hint to iteratively repaint the subsequent 34-frame clips until the entire video is processed. We present the results in the Supp.

6. Conclusion

In this paper, we present the VIRES, a Video Instance REpainting method with Sketch and text guidance, built upon a pre-trained text-to-video generation model. We propose a Sequential ControlNet with a standardized self-

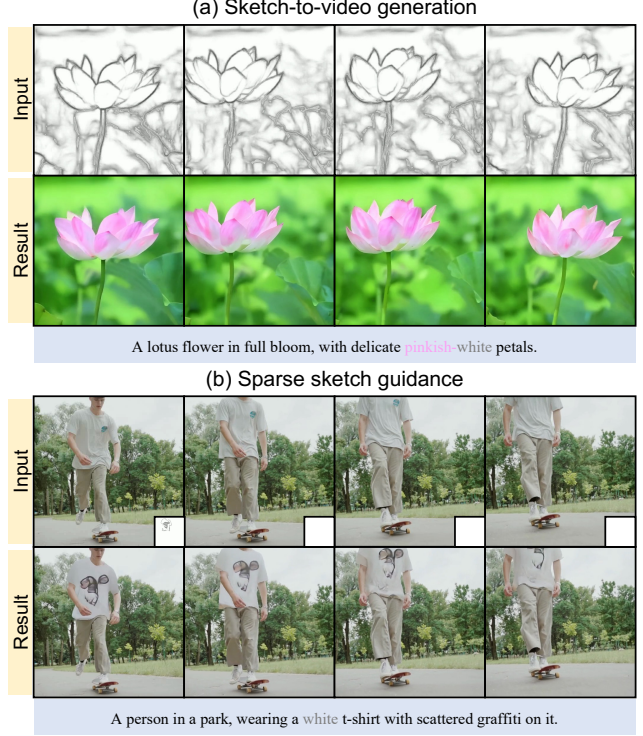


Figure 5. Examples of extended applications: (a) Generating an entire video clip with the sketch sequence. (b) Repainting the video instance with a single sketch frame.

scaling to effectively extract structure layouts and adaptively capture high-contrast sketch details. We further augment the DiT backbone with the sketch attention to interpret and inject fine-grained sketch semantics, and introduce a sketch-aware encoder for structure alignment during decoding. For training and evaluating video editing methods, we contribute VIRESET. In addition to video instance repainting, replacement, generation, and removal, VIRES is applicable to sketch-to-video generation, sparse sketch guidance, and long-duration video repainting. Extensive experiments demonstrate that VIRES outperforms existing video editing methods and receives higher ratings from human evaluators.

Limitation. Our VIRES is based on the DiT architecture [52], which leads to higher computational requirements. For example, editing a 51-frame video at 512×512 resolution takes approximately 323 seconds on a single A6000 GPU, which is about 1.6 times slower than other diffusion-based methods [14, 38]. We believe faster sampling methods will reduce the inference time in the future.

Acknowledgements

This work is supported by National Natural Science Foundation of China under Grant No. 62136001. PKU-affiliated authors thank *openbayes.com* for providing computing resources.

References

- [1] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, 2021. 3
- [2] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, 2021. 2
- [3] Stefan Elfving, Eiji Uchibe, and Kenji Doya. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural networks*, 2018. 4
- [4] Sicheng Gao, Xuhui Liu, Bohan Zeng, Sheng Xu, Yanjing Li, Xiaoyan Luo, Jianzhuang Liu, Xiantong Zhen, and Baochang Zhang. Implicit diffusion models for continuous super-resolution. In *CVPR*, 2023. 2
- [5] Yingqing He, Tianyu Yang, Yong Zhang, Ying Shan, and Qifeng Chen. Latent video diffusion models for high-fidelity long video generation. *arXiv preprint arXiv:2211.13221*, 2022. 2
- [6] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 2
- [7] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. In *NeurIPS*, 2022. 2
- [8] Yuchen Hong, Haofeng Zhong, Shuchen Weng, Jinxiu Liang, and Boxin Shi. L-Differ: Single image reflection removal with language-based diffusion model. *ECCV*, 2024. 2
- [9] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, 2017. 4
- [10] 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 *CVPR*, 2024. 3
- [11] Quan Huynh-Thu and Mohammed Ghanbari. Scope of validity of psnr in image/video quality assessment. *Electronics letters*, 2008. 6
- [12] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016. 5
- [13] Ozgur Kara, Bariscan Kurtkaya, Hidir Yesiltepe, James M Rehg, and Pinar Yanardag. RAVE: Randomized noise shuffling for fast and consistent video editing with diffusion models. In *CVPR*, 2024. 6, 7
- [14] Levon Khachatryan, Andranik Movsisyan, Vahram Tadevosyan, Roberto Henschel, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Text2Video-Zero: Text-to-image diffusion models are zero-shot video generators. In *ICCV*, 2023. 2, 6, 7, 8
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [16] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 1951. 5
- [17] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *CVPR*, 2023. 2
- [18] Xirui Li, Chao Ma, Xiaokang Yang, and Ming-Hsuan Yang. VidToMe: Video token merging for zero-shot video editing. In *CVPR*, 2024. 2, 6, 7
- [19] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *ICLR*, 2023. 3
- [20] Yixin Liu, Kai Zhang, Yuan Li, Zhiling Yan, Chujie Gao, Ruoxi Chen, Zhengqing Yuan, Yue Huang, Hanchi Sun, Jianfeng Gao, et al. Sora: A review on background, technology, limitations, and opportunities of large vision models. *arXiv preprint arXiv:2402.17177*, 2024. 2
- [21] 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. In *ICLR*, 2021. 2
- [22] Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, and Ying Shan. T2I-Adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In *AAAI*, 2024. 2
- [23] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, 2019. 4
- [24] Gaurav Parmar, Krishna Kumar Singh, Richard Zhang, Yijun Li, Jingwan Lu, and Jun-Yan Zhu. Zero-shot image-to-image translation. In *SIGGRAPH*, 2023. 2
- [25] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, 2016. 3, 5
- [26] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. FiLM: Visual reasoning with a general conditioning layer. In *AAAI*, 2018. 4
- [27] Chenyang Qi, Xiaodong Cun, Yong Zhang, Chenyang Lei, Xintao Wang, Ying Shan, and Qifeng Chen. FateZero: Fusing attentions for zero-shot text-based video editing. In *ICCV*, 2023. 2
- [28] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 6
- [29] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 2020. 3
- [30] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. SAM 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 2, 3
- [31] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2

- [32] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. DreamBooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *CVPR*, 2023. 2
- [33] 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. In *NeurIPS*, 2022. 2
- [34] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021. 2
- [35] Tae Soo Kim and Austin Reiter. Interpretable 3D human action analysis with temporal convolutional networks. In *CVPR workshops*, 2017. 2
- [36] Fu-Yun Wang, Wenshuo Chen, Guanglu Song, Han-Jia Ye, Yu Liu, and Hongsheng Li. Gen-L-Video: Multi-text to long video generation via temporal co-denoising. *arXiv preprint arXiv:2305.18264*, 2023. 2
- [37] Jiapeng Wang, Chengyu Wang, Kunzhe Huang, Jun Huang, and Lianwen Jin. VideoCLIP-XL: Advancing long description understanding for video clip models. *arXiv preprint arXiv:2410.00741*, 2024. 6
- [38] Xiang Wang, Hangjie Yuan, Shiwei Zhang, Dayou Chen, Jiniun Wang, Yingya Zhang, Yujun Shen, Deli Zhao, and Jingren Zhou. VideoComposer: Compositional video synthesis with motion controllability. In *NeurIPS*, 2024. 2, 3, 6, 7, 8
- [39] Yufei Wang, Wenhan Yang, Xinyuan Chen, Yaohui Wang, Lanqing Guo, Lap-Pui Chau, Ziwei Liu, Yu Qiao, Alex C Kot, and Bihan Wen. SinSR: diffusion-based image super-resolution in a single step. In *CVPR*, 2024. 2
- [40] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 2004. 5, 6
- [41] Shuchen Weng, Jimeng Sun, Yu Li, Si Li, and Boxin Shi. CT²: Colorization transformer via color tokens. In *ECCV*, 2022. 2
- [42] Shuchen Weng, Peixuan Zhang, Yu Li, Si Li, Boxin Shi, et al. L-CAD: Language-based colorization with any-level descriptions using diffusion priors. In *NeurIPS*, 2024. 2
- [43] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-A-Video: One-shot tuning of image diffusion models for text-to-video generation. In *ICCV*, 2023. 2
- [44] Yuxin Wu and Kaiming He. Group normalization. In *ECCV*, 2018. 4
- [45] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. In *ICCV*, 2015. 2, 3
- [46] Lin Xu, Yilin Zhao, Daquan Zhou, Zhijie Lin, See Kiong Ng, and Jiashi Feng. PLLaVA: Parameter-free llava extension from images to videos for video dense captioning. *arXiv preprint arXiv:2404.16994*, 2024. 2, 3
- [47] Shuai Yang, Yifan Zhou, Ziwei Liu, and Chen Change Loy. Rerender a video: Zero-shot text-guided video-to-video translation. In *SIGGRAPH Asia*, 2023. 2, 6
- [48] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. CogVideoX: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. 3
- [49] Lijun Yu, Jose Lezama, Nitesh Bharadwaj Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G Hauptmann, et al. Language model beats diffusion-tokenizer is key to visual generation. In *ICLR*. 4
- [50] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, 2023. 2, 4
- [51] Yabo Zhang, Yuxiang Wei, Dongsheng Jiang, XIAOPENG ZHANG, Wangmeng Zuo, and Qi Tian. ControlVideo: Training-free controllable text-to-video generation. In *ICLR*. 2
- [52] Zangwei Zheng, Xiangyu Peng, Tianji Yang, Chenhui Shen, Shenggui Li, Hongxin Liu, Yukun Zhou, Tianyi Li, and Yang You. Open-Sora: Democratizing efficient video production for all, 2024. 2, 3, 4, 5, 8
- [53] Haofeng Zhong, Yuchen Hong, Shuchen Weng, Jinxiu Liang, and Boxin Shi. Language-guided image reflection separation. In *CVPR*, 2024. 2
- [54] Yuan Zhou, Qiuyue Wang, Yuxuan Cai, and Huan Yang. Allegro: Open the black box of commercial-level video generation model. *arXiv preprint arXiv:2410.15458*, 2024. 3

VIRES: Video Instance Repainting via Sketch and Text Guided Generation

Supplementary Material

Shuchen Weng^{1,2†‡} Haojie Zheng^{3,4†} Peixuan Zhang⁵ Yuchen Hong^{1,2}
 Han Jiang³ Si Li⁵ Boxin Shi^{1,2*}

¹State Key Laboratory for Multimedia Information Processing, School of Computer Science, Peking University

²National Engineering Research Center of Visual Technology, School of Computer Science, Peking University

³OpenBayes Information Technology Co., Ltd. ⁴School of Software and Microelectronics, Peking University

⁵School of Artificial Intelligence, Beijing University of Posts and Telecommunications

scweng@baai.ac.cn, suimu@stu.pku.edu.cn, {pxzhang, lisi}@bupt.edu.cn
 yuchenhong.cn@gmail.com, hahn@openbayes.com, shiboxin@pku.edu.cn

A. Appendix

A.1. Variations of Sequential ControlNet

We present four typical architectures of Sequential ControlNet in Tab. S1, where "Conv" denotes a convolutional layer, "Block" means a residual block, and "Down" is a down-sampling layer. Numbers in brackets indicate the input and output channel dimensions, respectively. To determine the optimal architecture under constrained computational resources, we train these variations for 10K steps, excluding the sketch attention and sketch-aware encoder. Quantitative results on VIRESET are shown in Tab. S2, and the best-performing architecture is selected for our model.

Table S1. Architecture of Sequential ControlNet variations.

	Ours (VIRES)	Variation_1	Variation_2	Variation_3
block_1	Conv (3, 72)		Conv (3, 36)	Conv (3, 72)
	Conv (72, 72)		Conv (36, 36)	Conv (72, 72)
	Down (72, 144)	Down (3, 64)	Down (36, 72)	Down (72, 144)
block_2	Conv (144, 144)	Conv (64, 64)		Conv (144, 144)
	Conv (144, 288)	Conv (64, 128)		Conv (144, 288)
	Conv (288, 288)			
	Block (288, 288)	Block (128, 128)	Block (72, 72)	Block (288, 288)
	Down (288, 576)	Down (128, 128)	Down (72, 288)	Down (288, 576)
block_3	Block (576, 576)	Block (128, 256)	Block (288, 288)	Block (576, 576)
	Down (576, 1152)	Down (256, 256)	Down (288, 1152)	Down (576, 1152)
block_4	Block (1152, 1152)			
	Conv (1152, 1152)	Block (256, 256)	Conv (1152, 1152)	Conv (1152, 1152)
	Conv (1152, 1152)		Conv (1152, 1152)	Conv (1152, 1152)
	Conv (1152, 1152)	Conv (256, 1152)	Conv (1152, 1152)	Conv (1152, 1152)

A.2. Validation of additional conditions

VIRES is designed specifically for sketch-based video instance repainting, adopting the Standardized Self-Scaling (SSS) to extract condition features. To explore its generalization to other conditioning signals, we augment our dataset with edge maps [1] and depth maps [6] to pro-

[†] Equal contribution.

[‡] Now at Beijing Academy of Artificial Intelligence.

* Corresponding author.

Table S2. Quantitative experiment results of Sequential ControlNet variations. All scores except PSNR are percentages.

Method	PSNR ↑	SSIM ↑	WE ↓	FC ↑	TC ↑
Variation_1	18.54	69.12	5.36	90.79	15.85
Variation_2	18.49	68.63	5.47	89.14	15.83
Variation_3	18.64	69.89	5.40	90.94	15.85
Ours (VIRES)	19.93	72.85	5.33	91.07	15.87

vide the structure guidance. We then retrain VIRES from scratch on these conditions. Due to limited computational resources, we train for 10K steps with edge maps as a preliminary investigation and 30K steps with depth maps due to its slower convergence. Finally, we compare two VIRES variants: one using SSS and the other using simple addition for feature extraction. As shown in Tab. S3, SSS provides considerable advantages for edge maps, but only marginal improvements for depth maps. We suggest that this difference arises because both sketch and edge maps have high-contrast transitions between black lines and the white background, allowing self-scaling to effectively capture structure details. In contrast, depth maps are smoother and lack such sharp transitions, limiting the benefits of the self-scaling operation. We further present the qualitative results in Fig. S1.

Table S3. Quantitative experiment results of VIRES variants using edge and depth maps. All scores except PSNR are percentages.

Method	PSNR ↑	SSIM ↑	WE ↓	FC ↑	TC ↑
Repainting with edge maps					
W/o SSS	19.43	72.01	6.31	90.04	15.83
W/ SSS	20.48	74.28	6.00	90.49	16.16
Repainting with depth maps					
W/o SSS	18.04	66.16	5.86	91.28	15.98
W/ SSS	18.39	67.28	5.80	91.32	16.14

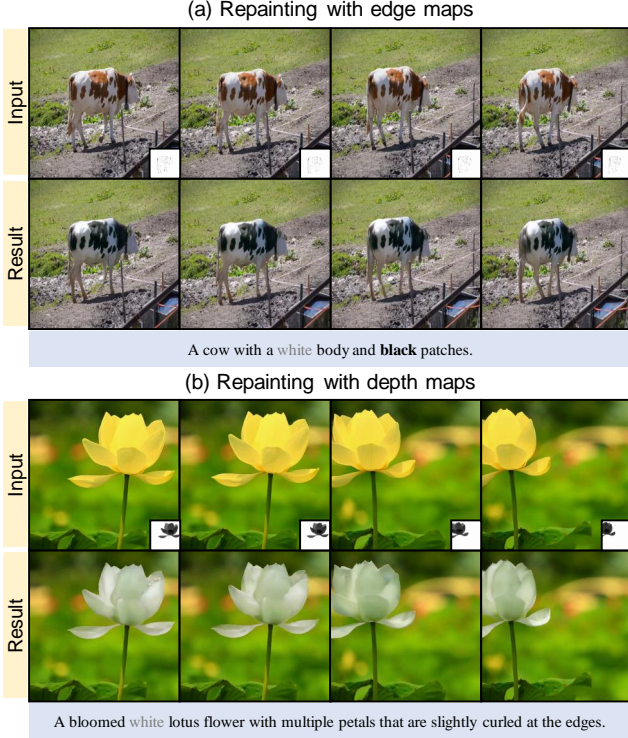


Figure S1. Examples of repainting with additional conditions.

Table S4. Quantitative experiment results of VIREs using sketch guidance of varying sparsity. All scores except PSNR are percentages.

Method	PSNR \uparrow	SSIM \uparrow	WE \downarrow	FC \uparrow	TC \uparrow
$D = 1$	21.87	73.41	5.12	92.21	16.15
$D = 2$	22.64	74.69	5.14	92.21	16.17
$D = 4$	23.24	76.00	5.13	92.18	16.16
$D = 8$	23.36	76.51	5.14	92.19	16.18
$D = 16$	23.67	77.15	5.12	92.19	16.19
$D = 51$	23.87	77.87	5.09	92.23	16.19

A.3. Robustness of sparse conditions

VIREs allows users to provide sparse sketch guidance for video instance repainting, minimizing user effort. To investigate the impact of sketch guidance sparsity on repainting performance, we evaluate VIREs on the VIREsET, using varying interval indices $d \in \{0, \dots, 4\}$, corresponding to $D = 2^d$ sketch frames, and a full set of $D = 51$. As shown in Tab. S4, even with a single sketch frame, VIREs can produce high-quality results (PSNR and SSIM) with robust temporal consistency (WE and FC).

A.4. Compatibility with DiT backbone

In this paper, we build VIREs upon the pre-trained OpenSora v1.2 [9]. Given that many text-to-video models [8, 10] share a similar DiT backbone architecture, our proposed modules offer potential compatibility with these approaches, including the Sequential ControlNet for layout initialization (Sec. 4.2), the standardized self-scaling for details capture (Sec. 4.2), the sketch attention (Sec. 4.3) for semantic injection, and the sketch-aware encoder for structure alignment (Sec. 4.4). We believe our work will inspire further research on guiding pre-trained text-to-video models and open new avenues for conditional video repainting.

A.5. Organization of supplementary video

We provide a supplementary video to dynamically showcase our repainting results. The video is structured as follows: (i) **Typical application scenarios**: We demonstrate four typical repainting scenarios and compare our results with relevant methods [2–5, 7]. Instance repainting/replacement results are shown in Fig. S2, and instance generation/removal results are in Fig. S3. (ii) **Sketch-to-video generation and inpainting**: We demonstrate sketch-to-video generation and conditional video inpainting, comparing our results with VideoComposer [5], as it is the only relevant method that supports this functionality. Results are shown in Fig. S4. (iii) **Sparse sketch guidance**: We showcase sparse sketch guidance, repainting two distinct variations of the same video using only two different first sketch frames. This functionality is not supported by existing methods. Results are shown in Fig. S5. (iv) **Long-duration video repainting**: We demonstrate repainting on a long-duration (13-second) video, with representative frames presented in Fig. S6. (v) **Comparison and ablation study**: Finally, the video includes comparisons with other methods (Sec. 5.1) and additional ablation studies (Sec. 5.2). To improve visual clarity and facilitate detailed comparison, the video playback speed is halved ($2\times$ slower).

References

- [1] Rafael C Gonzales and Paul Wintz. *Digital image processing*. Addison-Wesley Longman Publishing Co., Inc., 1987. 1
- [2] Ozgur Kara, Bariscan Kurtkaya, Hidir Yesiltepe, James M Rehg, and Pinar Yanardag. RAVE: Randomized noise shuffling for fast and consistent video editing with diffusion models. In *CVPR*, 2024. 2
- [3] Levon Khachatryan, Andranik Movsisyan, Vahram Tadevosyan, Roberto Henschel, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Text2Video-Zero: Text-to-image diffusion models are zero-shot video generators. In *ICCV*, 2023.
- [4] Xirui Li, Chao Ma, Xiaokang Yang, and Ming-Hsuan Yang.

VidToMe: Video token merging for zero-shot video editing. In *CVPR*, 2024.

- [5] Xiang Wang, Hangjie Yuan, Shiwei Zhang, Dayou Chen, Jiniu Wang, Yingya Zhang, Yujun Shen, Deli Zhao, and Jingren Zhou. VideoComposer: Compositional video synthesis with motion controllability. In *NeurIPS*, 2024. 2
- [6] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *arXiv preprint arXiv:2406.09414*, 2024. 1
- [7] Shuai Yang, Yifan Zhou, Ziwei Liu, and Chen Change Loy. Rerender a video: Zero-shot text-guided video-to-video translation. In *SIGGRAPH Asia*, 2023. 2
- [8] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. CogVideoX: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. 2
- [9] Zangwei Zheng, Xiangyu Peng, Tianji Yang, Chenhui Shen, Shenggui Li, Hongxin Liu, Yukun Zhou, Tianyi Li, and Yang You. Open-Sora: Democratizing efficient video production for all, 2024. 2
- [10] Yuan Zhou, Qiuyue Wang, Yuxuan Cai, and Huan Yang. Allegro: Open the black box of commercial-level video generation model. *arXiv preprint arXiv:2410.15458*, 2024. 2




Figure S2. Typical application scenarios. **Left:** Video instance repainting. **Right:** Video instance replacement.

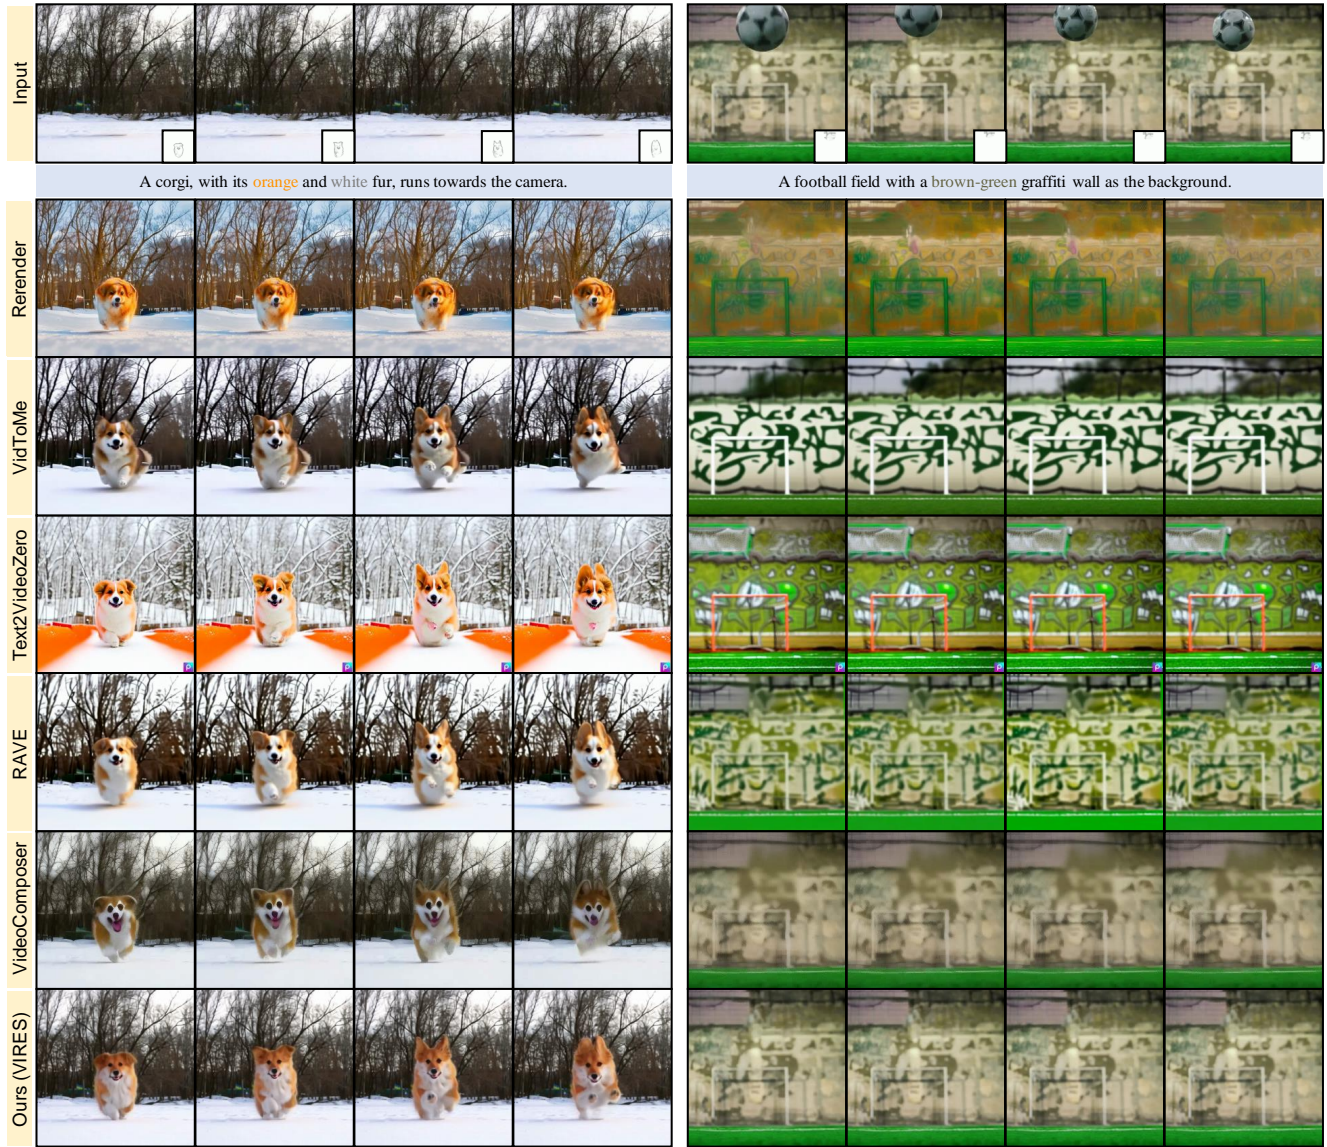


Figure S3. Typical application scenarios. **Left:** Custom instance generation. **Right:** Specified instance removal.




Figure S4. Sketch-to-video generation and inpainting. **Topleft:** Sketch-to-video generation. **Others:** Sketch-to-video inpainting.



Figure S5. Sparse sketch guidance. Repainting the video using different first sketch frames. **Left:** First variation. **Right:** Second variation.

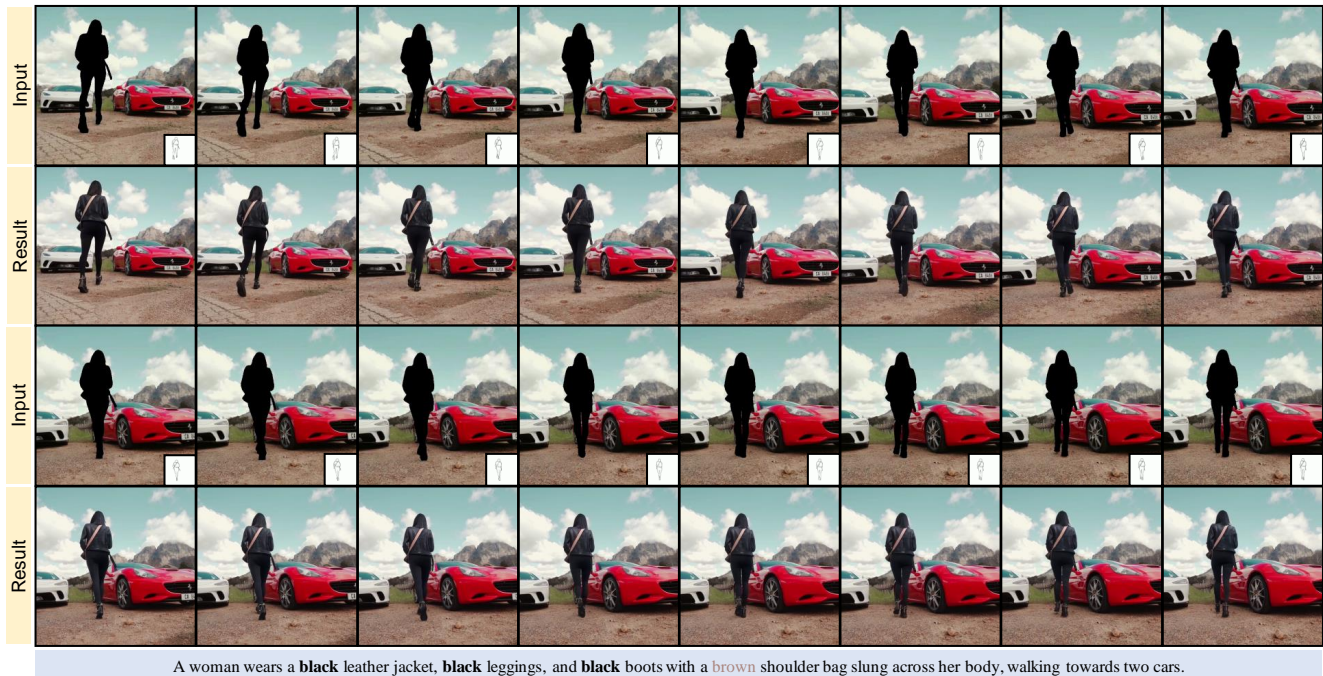


Figure S6. Long-duration video repainting. Restoring the damaged video to clearly depict a realistic female character.