

Sketch3DVE: Sketch-based 3D-Aware Scene Video Editing

FENG-LIN LIU, Institute of Computing Technology, CAS and University of Chinese Academy of Sciences, China

SHI-YANG LI, Institute of Computing Technology, CAS, China

YAN-PEI CAO, VAST, China

HONGBO FU, Hong Kong University of Science and Technology, China

LIN GAO*, Institute of Computing Technology, CAS and University of Chinese Academy of Sciences , China

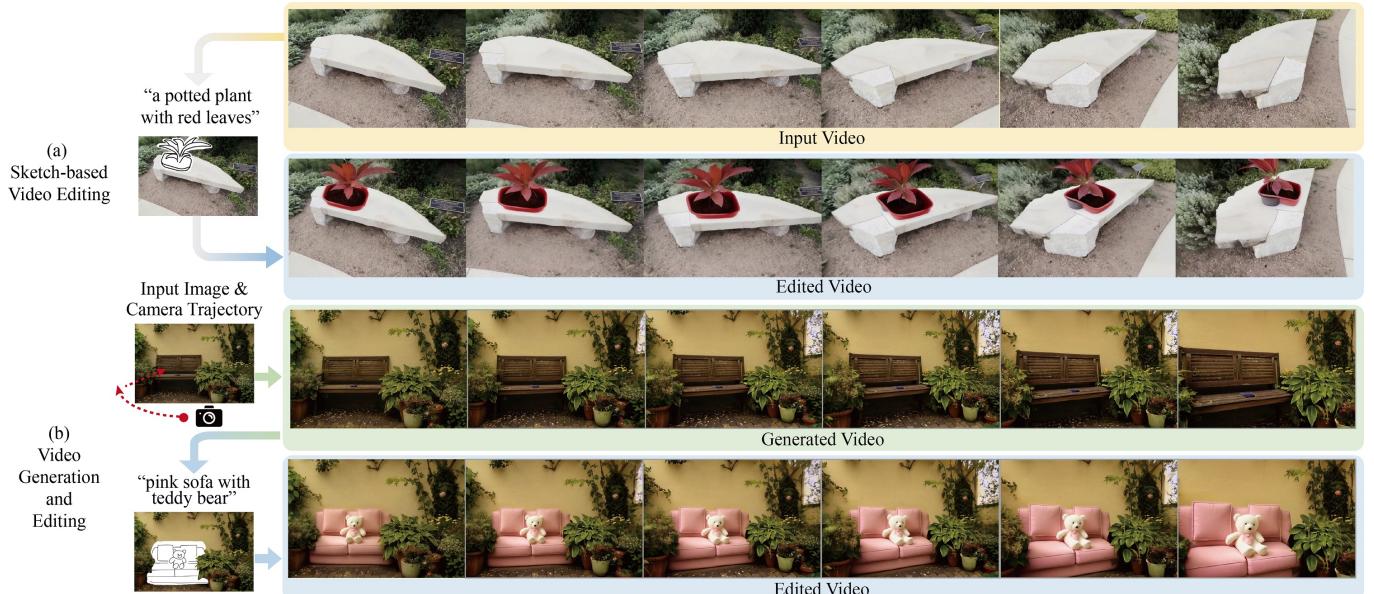


Fig. 1. Given an input video (yellow box in (a)) with significant viewpoint changes, Sketch3DVE generates realistic editing results (blue box) with the inputs of a text prompt, hand-drawn sketch, and a mask (in white). Our method can also edit a video generated given an input image and a camera trajectory (b). Input Video ©DL3DV-10K, Input image ©Pegah Sharifi.

Recent video editing methods achieve attractive results in style transfer or appearance modification. However, editing the structural content of 3D scenes in videos remains challenging, particularly when dealing with significant viewpoint changes, such as large camera rotations or zooms. Key challenges include generating novel view content that remains consistent with the original video, preserving unedited regions, and translating sparse 2D inputs into realistic 3D video outputs. To address these issues, we propose Sketch3DVE, a sketch-based 3D-aware video editing method to enable detailed local manipulation of videos with significant viewpoint changes.

*Corresponding author is Lin Gao (gaolin@ict.ac.cn).

Authors' addresses: Feng-Lin Liu, Institute of Computing Technology, CAS and University of Chinese Academy of Sciences, China, liufenglin21s@ict.ac.cn; Shi-Yang Li, Institute of Computing Technology, CAS, China, shiyangli179@gmail.com; Yan-Pei Cao, VAST, China, caoyanpei@gmail.com; Hongbo Fu, Hong Kong University of Science and Technology, China, hongbofu@ust.hk; Lin Gao, Institute of Computing Technology, CAS and University of Chinese Academy of Sciences , China, gaolin@ict.ac.cn.



This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

SIGGRAPH Conference Papers '25, August 10–14, 2025, Vancouver, BC, Canada

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1540-2/2025/08

<https://doi.org/10.1145/3721238.3730623>

To solve the challenge posed by sparse inputs, we employ image editing methods to generate edited results for the first frame, which are then propagated to the remaining frames of the video. We utilize sketching as an interaction tool for precise geometry control, while other mask-based image editing methods are also supported. To handle viewpoint changes, we perform a detailed analysis and manipulation of the 3D information in the video. Specifically, we utilize a dense stereo method to estimate a point cloud and the camera parameters of the input video. We then propose a point cloud editing approach that uses depth maps to represent the 3D geometry of newly edited components, aligning them effectively with the original 3D scene. To seamlessly merge the newly edited content with the original video while preserving the features of unedited regions, we introduce a 3D-aware mask propagation strategy and employ a video diffusion model to produce realistic edited videos. Extensive experiments demonstrate the superiority of Sketch3DVE in video editing. We will release our code upon acceptance.

CCS Concepts: • Human-centered computing → Graphical user interfaces; • Computer systems organization → Neural networks; • Computing methodologies → Image processing.

Additional Key Words and Phrases: Sketch-based interaction, video generation, video editing, video diffusion models

ACM Reference Format:

Feng-Lin Liu, Shi-Yang Li, Yan-Pei Cao, Hongbo Fu, and Lin Gao. 2025. Sketch3DVE: Sketch-based 3D-Aware Scene Video Editing. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25), August 10–14, 2025, Vancouver, BC, Canada*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3721238.3730623>

1 INTRODUCTION

Video generation and editing are popular research areas with broad applications in movie production, education, robotics, and AR/VR. Unlike creating videos from scratch, editing involves manipulating the geometry and appearance of objects in existing real or synthetic videos, enabling detailed user customization. While pioneering video editing methods [Jamriska et al. 2019; Kasten et al. 2021; Liu et al. 2022; Ruder et al. 2018; Tzaban et al. 2022] achieve attractive results, they are limited to specific domains or restricted editing types. Image editing methods, such as sketch-based works [Jo and Park 2019; Liu et al. 2024a; Mikaeili et al. 2023; Portenier et al. 2018; Zeng et al. 2022] and mask-based inpainting works [Alibaba 2024; Yang et al. 2023a; Zhang et al. 2025], have good flexibility and attractive results. However, applying similar ideas to video editing presents several challenges, including domain gaps between single-image 2D inputs and videos, the absence of novel view information for videos with significant view changes, and the requirement to identify and preserve unedited regions.

Recent advancements in diffusion models have led to impressive results in text- and image-based video generation, with both commercial [Kuaishou 2024; OpenAI 2024; runway 2024; StabilityAI 2024] and open source works [Lab and etc. 2024; Team and Laboratory 2024; Yang et al. 2024c; Zheng et al. 2024b]. Similarly, diffusion-based video editing methods [Ku et al. 2024; Mou et al. 2024a; Ouyang et al. 2024a; Zhang et al. 2024c] have made significant progress in diverse editing scenarios. However, as shown in Fig. 4, most of these methods transfer motion features from the input video to the edited video, which is effective for appearance modification but struggles with structural editing, such as component insertion and replacement. Furthermore, these methods use UNet or DiT networks to handle temporal motion implicitly. However, they lack extensive 3D information analysis, limiting their ability to process videos with substantial viewpoint changes.

Recent video generation methods have employed 3D information to enable controllable viewpoint changes. Given a single image, novel view videos can be synthesized by inserting camera parameters into pretrained video diffusion models [He et al. 2024; Wang et al. 2024b; Yang et al. 2024a; Zheng et al. 2024a], or using point cloud rendering for explicit view control [Liu et al. 2024c; Yu et al. 2024]. These methods support editing by modifying the input image and generating videos with the original camera motions. However, for generated novel view frames, the obscured areas and components outside the viewport of the edited image are entirely imagined by the models, often resulting in unintended changes in these regions even if they remain unedited. Instead of video generation, NeRF [Mildenhall et al. 2022] and 3DGS [Kerbl et al. 2023] have been utilized for 4D generation [Wu et al. 2024; Xu et al. 2025; Zhang et al. 2024a], image-to-object [Liu et al. 2023b; Luo et al. 2023], and

image-to-scene [Seo et al. 2024], while focusing on generation from scratch rather than local modification. 3D scene editing methods [Li et al. 2024; Liu et al. 2024b; Mirzaei et al. 2023] modify 3D attributes to render edited videos, but struggle with structural modifications or produce rendering-style outputs.

To address these issues, we propose Sketch3DVE, a sketch-based 3D-aware method to produce realistic edited videos (see Fig. 1). We assume both the input and output are scene videos with significant viewpoint changes but slight object movements. To bridge the domain gap between 2D sketches and videos, we first generate the edited results of the first frame by MagicQuill [Liu et al. 2024d] and then propagate the editing effects across the video. Image editing on the first frame can also be achieved with other image editing methods (e.g., [Alibaba 2024; Yang et al. 2023a]) that use masks to label the local modification regions. To generate high-quality details in novel views of the edited content while ensuring seamless camera motion consistency with the input video, we perform explicit 3D geometry analysis and manipulation. Specifically, we first obtain a point cloud and the camera parameters from the input video using DUST3R [Wang et al. 2024a]. Then, we propose a point cloud editing approach to map 2D image edits into 3D space, ensuring that the edited regions align with the rest of the scene. Depth maps, which provide relative 3D geometry information and explicit pixel matching before and after editing, are utilized to achieve this alignment. To solve the issue of unedited region identification and preservation, we propose a 3D-aware mask propagation strategy that tracks the input 2D mask across other frames. Finally, a video diffusion model generates the novel view results of the edited components that are merged with unedited regions to synthesize a realistic edited video. Our method supports a wide range of editing operations, including object insertion, removal, replacement, and modification to shape and texture, as shown in Fig. 3 and 8.

Extensive experiments demonstrate that our method produces higher-quality video editing results than existing image-based video generation and editing approaches. Our main contributions can be summarized as follows:

- We propose a novel sketch-based 3D-aware video editing method that generates realistic and high-quality editing results for videos with significant viewpoint changes.
- We propose a point cloud editing approach that utilizes depth maps to represent and align the edited regions with the original 3D scene.
- We develop a 3D-aware mask propagation strategy and a precise region modification video diffusion model that synthesizes novel view results of edited components while accurately preserving unedited regions.

2 RELATED WORK

Our method is closely related to camera-controllable video generation, deep video editing, and sketch-based content editing.

Camera Controllable Video Generation. One category of methods implicitly inserts camera parameters into video generation models. The camera features influence the temporal attention mechanism [Wang et al. 2024b; Yang et al. 2024a], or transform to Plücker Embedding [Sitzmann et al. 2021] with additional condition networks

for insertion [Bahmani et al. 2024; He et al. 2024; Wang et al. 2024c]. Epipolar attention [Xu et al. 2024; Zheng et al. 2024a], LoRa fine-tuning [Sun et al. 2024a], and multi-view datasets [Bai et al. 2024] further enhance camera control performance. Explicit methods, on the other hand, construct 3D representations from input images and render them into novel views, which are then used as conditions for video generation. Training-free denoising resampling [Hou et al. 2024; Liu et al. 2024c] and video diffusion model finetuning [Yu et al. 2024] are utilized to translate point cloud renderings into videos. 3D tracking videos [Gu et al. 2025] and 3D box rendering depth maps [Wang et al. 2025] further serve as conditions for dynamic content synthesis. While these methods effectively control the camera for video generation, they are inherently limited for editing due to the absence of relationship analysis between the original video content and the newly introduced modifications. Our method generates high-quality editing contents that seamlessly merge into original videos with good view consistency.

Deep Video Editing. Pioneering works achieve effective video editing in facial videos [Liu et al. 2022; Tzaban et al. 2022], style transfer [Jamriska et al. 2019; Ruder et al. 2018], and layered editing propagation [Kasten et al. 2021]. With the success of diffusion models, many works [Ceylan et al. 2023; Liu et al. 2024e; Ouyang et al. 2024b; Qi et al. 2023; Wang et al. 2023; Yang et al. 2023b] extend text-to-image generation models [Rombach et al. 2022] to video editing by enforcing temporal coherence. Video generation models are further used to mitigate flickering and enhance robustness. AVID [Zhang et al. 2024c] treats video editing as a sketch-guided inpainting task and enables text-based editing. Similar to our method, many works propagate the editing from the first frame of a video to the rest. AnyV2V [Ku et al. 2024], a training-free method, merges edited results with attention features of the input video but sometimes generates structural distortion and temporal inconsistency. I2VEdit [Ouyang et al. 2024a] fine-tunes models using LoRa on temporal layers to extract motion information, producing consistent outputs but struggles with object insertion due to missing motion information in the original video. ReVideo [Mou et al. 2024a] addresses this by adding missing motion using 2D trajectory inputs. In contrast, our method explicitly extracts and analyzes 3D information, enabling accurate shape editing, object insertion, and replacement even when the input videos have significant view transformations.

Sketch-Based Content Editing. Sketch-based interfaces have been widely used in image generation [Chen et al. 2020; Gao et al. 2020; Isola et al. 2017; Ma et al. 2024; Mou et al. 2024b; Xie et al. 2024b; Zhang et al. 2023; Zhu et al. 2017], video generation [Huang et al. 2024; Jiang et al. 2025; Li et al. 2022; Xing et al. 2024], and 3D generation [Brodt and Bessmeltsev 2022; Gao et al. 2024; Han et al. 2017; Sun et al. 2024b; Wang et al. 2014; Zheng et al. 2024c, 2023]. Compared to generation, editing requires preserving original features while producing reasonable modifications. This challenge can be addressed as a sketch-guided inpainting problem in image editing [Jo and Park 2019; Liu et al. 2024d; Portenier et al. 2018; Yu et al. 2019; Zeng et al. 2022]. Similar ideas have been utilized for video editing. For example, VIRES [Weng et al. 2024] designs a sequence ControlNet to enable sketch-based video edits but requires sketches and masks of edited regions for all frames. FramePainter [Zhang

et al. 2025] further utilizes video models for sketch/dragging-based image editing. Instead of fast inference, sketch-based 3D editing methods [Liu et al. 2024a; Mikaeili et al. 2023; Xie et al. 2024a] rely on optimization techniques to modify 3D representations but only focus on objects instead of 3D scenes. Our method designs a 3D-aware approach to generate masks for each frame, using a 3D point cloud representation to handle videos with significant view changes. Additionally, our method generates realistic results across diverse scenarios, including both indoor and outdoor scenes, addressing the limitations of previous sketch-based editing techniques.

3 METHODOLOGY

Given an input video, users incorporate hand-drawn sketches, masks, and text prompts to indicate desired modification effects. Our objective is to ensure the newly generated content remains consistent with the original video’s view changes while accurately preserving unedited regions. We assume the video captures static primarily scenes, mainly with significant viewpoint changes. During editing, we leverage MagicQuill [Liu et al. 2024d], a sketch-based image editing method (detailed in Sec. 3.1), to edit the first frame. As shown in Fig. 2, to propagate these edits across the video, our method extracts 3D information from the input video, including a point cloud and camera poses. We introduce a point cloud editing approach, which utilizes depth maps to represent and align the 3D geometry of the edited regions, as discussed in Sec. 3.2. Additionally, Sec. 3.3 presents our video mask propagation strategy, designed to track the edited regions across all frames. Finally, in Sec. 3.4, we develop a video diffusion model that synthesizes realistic edited videos with input conditions of point cloud renderings, the original video, and propagated masks. This pipeline ensures coherent and realistic editing across diverse view transformations, bridging 2D edits and 3D-aware video generation.

3.1 Preliminary

Image Editing Model. MagicQuill [Liu et al. 2024d] generates realistic edited images from an input image, text prompt, mask, and a sketch with optional color strokes. It builds on Stable Diffusion [Rombach et al. 2022] with two additional networks: ControlNet [Zhang et al. 2023] for effective editing and an inpainting U-Net [Ju et al. 2024] for preserving unedited regions. The model is trained to reconstruct images under conditions that imitate the editing process.

Video Diffusion Model. CogVideoX [Yang et al. 2024c], a state-of-the-art text-to-video generation model, is utilized for video synthesis. It employs a 3D causal VAE for video compression into a latent space, followed by a denoising process to generate latent codes. The DiT architecture, which contains multiple transformer blocks, is utilized to generate high-quality videos with good temporal consistency. Our video generation model adds an additional control block based on the CogVideoX model.

3.2 Depth-based Point Cloud Editing

Given a video with frames $I = \{I^1, I^2, \dots, I^N\}$, where N is the number of frames, the first frame is edited by MagicQuill to generate image I_{edit} . Directly applying an image-to-video generation model

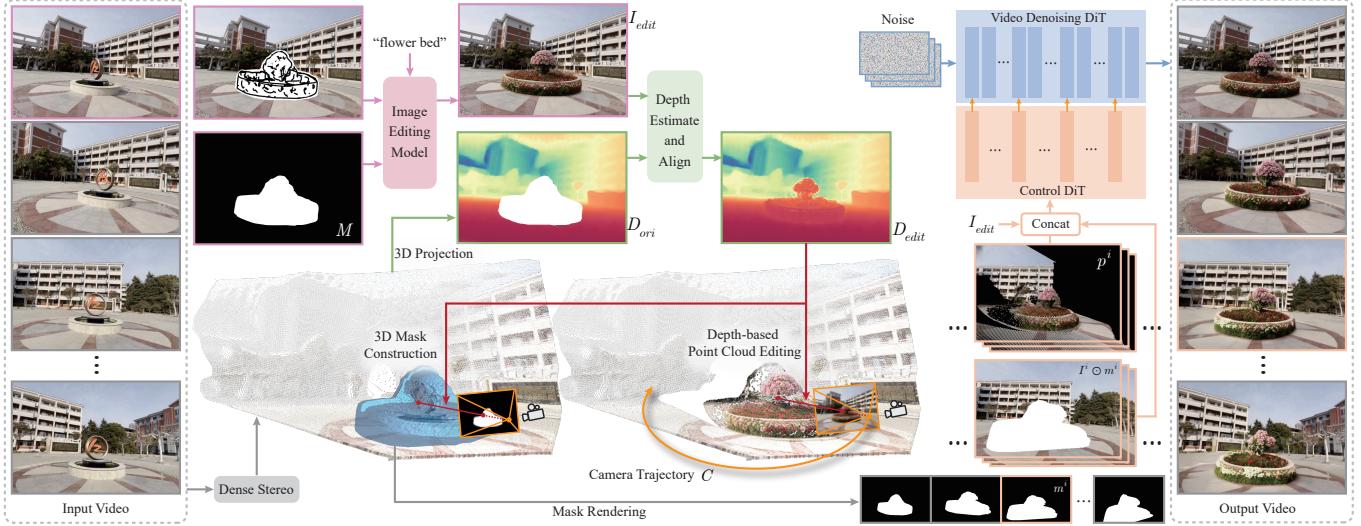


Fig. 2. The pipeline of our sketch-based 3D-ware video editing method. Given an input video, the edited result of the first frame, I_{edit} is generated by the image editing model [Liu et al. 2024d] using the sketch, mask, and text prompt input. We derive a 3D point cloud and camera parameters from the input video to propagate the editing. Then, we utilize a depth map to represent and align the newly edited content with the original scene to generate the edited point cloud. To identify edited regions, we construct a 3D mask that is rendered for all the frames to generate a mask sequence. The point cloud rendering results, the original video containing unedited regions, and the edited images serve as inputs for the conditional video diffusion model to synthesize the final edited video. Input Video ©DL3DV-10K.

on I_{edit} cannot control camera movement, leading to totally different viewpoint changes from the input video. To address this, we argue that explicit 3D geometry analysis and manipulation is a promising solution. Following [Yu et al. 2024], we use point clouds as a 3D representation to guide video generation, ensuring that the viewpoint changes of new content align with the original video.

We employ DUST3R [Wang et al. 2024a], a dense stereo model, to obtain a 3D point cloud and camera parameters from the input frames. Given two images I^1, I^2 from different views, DUST3R predicts point maps $O^{1,1}, O^{2,1}$ in the same camera coordinate system of view I^1 . This process is repeated between all pairs of frames in the video, followed by global point map alignment to obtain the output 3D information. Please refer to [Wang et al. 2024a] for more details. We denote the point cloud of the first frame as P and the camera parameters for all frames as $C = \{C^1, C^2, \dots, C^N\}$.

To effectively propagate the edits in I_{edit} across novel view frames, we translate the edited image I_{edit} into an edited point cloud P_{edit} , which is rendered with camera C to provide the pixel level guidance for novel view frame generation. To achieve this, we find that depth maps effectively encode the relative geometry relationship among objects within a 3D scene and explicitly have correspondence between the pixels in unedited regions, which can be used for point cloud alignment and editing.

Specifically, with the point cloud P and camera parameters C^1 , the depth map of the first frame D_{ori} before editing is generated by 3D projection. Similarly, the depth map of I_{edit} can also be obtained with a new point cloud and camera estimated by DUST3R (with I_{edit} duplicated to imitate paired input, as in [Yu et al. 2024]). Depth estimation models like [Yang et al. 2024b] can also be optionally

used to get depth maps. We normalize the values of the edited depth map to get \hat{D}_{edit} . Then, two alignment coefficients, translation t and scale s , are required to transform the relative depth maps \hat{D}_{edit} into the absolute one like D_{ori} . Notably, the geometry in unedited regions \bar{M} before and after editing naturally has correspondence, which can be used to obtain these coefficients. We utilize a simple least squares algorithm to achieve it. The optimization objective is to minimize the pixel distances:

$$E(s, t) = \sum_{i=1}^m \left\| (s\hat{d}_i + t) - d_i \right\|^2, \quad (1)$$

where d_i and \hat{d}_i are depth values in unedited regions of D_{ori} and \hat{D}_{edit} , respectively, and m is the number of pixels in unedited regions \bar{M} . Users can also adjust t, s to obtain the best performance. We merge the edited regions into the original depth maps to get the final edited depth maps:

$$D_{edit} = D_{ori} \odot \bar{M} + (\hat{D}_{edit} \cdot s + t) \odot M. \quad (2)$$

The final edited point cloud P_{edit} is obtained by back-projection:

$$P_{edit} = o^1 + r^1 \cdot D_{edit}, \quad (3)$$

where o^1 is the ray origin and r^1 is the ray direction corresponding to camera C^1 .

3.3 Video Mask Propagation

For ease of interaction, we require users to draw a 2D mask M only in the first frame to label edited regions. This requires us to track the mask in novel view frames to merge new edited content with the original video while accurately preserving unedited regions. This

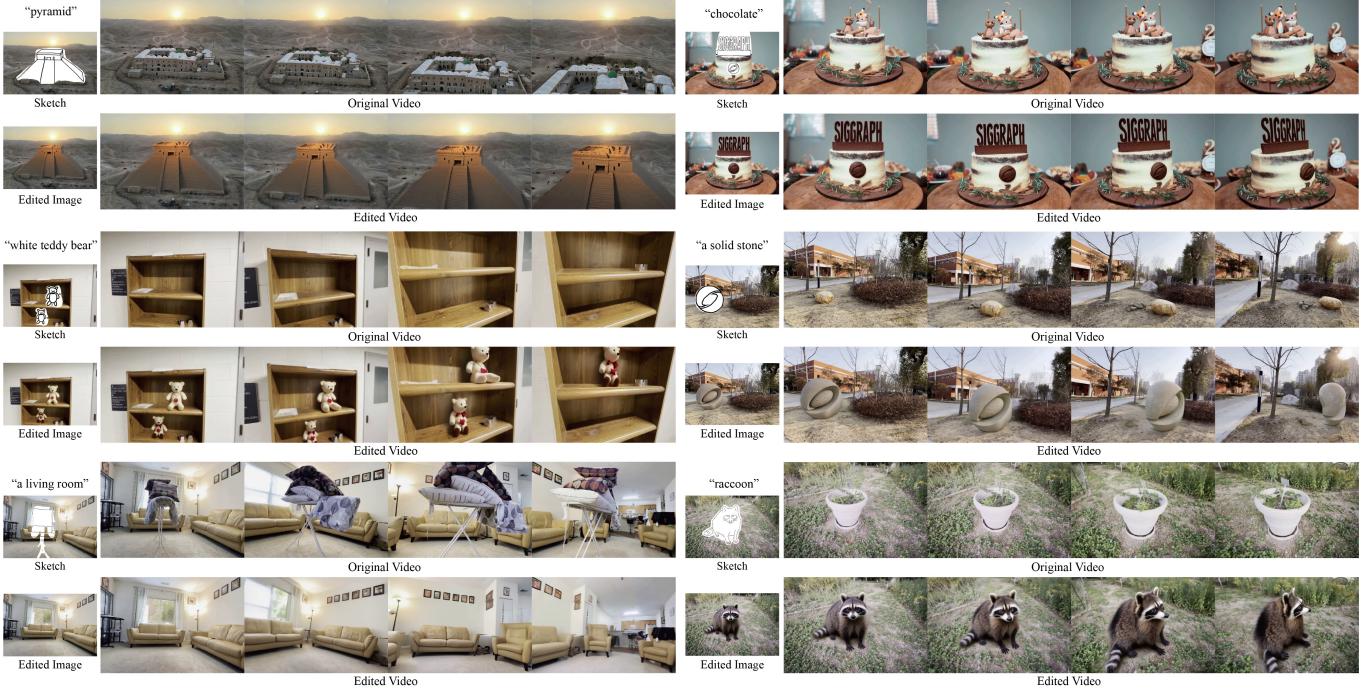


Fig. 3. Sketch-based video editing results. For each example, we show the input text prompt, sketch, and edited image on the left and the original and edited videos on the right. Our method generates realistic editing results for videos with significant viewpoint changes. Original Video ©DL3DV-10K, ©Taryn Elliott.

is nontrivial because the mask shapes are dynamically changed for each frame and must align with the original view transformation.

We address this by leveraging the depth maps before and after editing as priors. Specifically, we utilize a mesh model to build a 3D mask. We start with a cylindrical shape, whose top and bottom surfaces are defined by back-projecting 2D mask M along camera rays. This ensures the rendered 3D mask matches the 2D mask's shape in the edited view. To refine the 3D mask's geometry, for the top surface, we merge pre- and post-edit geometries by minimizing depth values D_{ori} and D_{edit} for each pixel. Then, vertex positions are adjusted based on this merged depth map, while edges are constructed according to pixel adjacency. For simplicity, we assume the bottom surface remains a plane structure, constructed with the uniform depth values obtained by the maximum of the edited regions in D_{ori} and D_{edit} . A side face further connects the frontal and back surfaces. With the camera C , we render the 3D mask to get a sequence of 2D masks, m^1, m^2, \dots, m^N , which are utilized for precise region modification in video editing. More details of the 3D mask construction are discussed in the supplemental material.

3.4 Diffusion-based Video Generation

A video diffusion model generates the final realistic edited videos. We render the edited point cloud P_{edit} with camera parameters C to obtain conditions $p = \{p^1, p^2, \dots, p^n\}$. The point cloud rendering results encode the camera transformation and enable explicit propagation of the editing manipulation. The input video (with the edited regions masked) also serves as a condition to preserve the original features. Additionally, we add the edited image I_{edit} as a

condition for each frame to provide the appearance reference and thus enhance the video details. We utilize a modified ControlNet [Zhang et al. 2023] branch to insert the camera control and original video information. The input of the control branch is formulated as,

$$c_{edit} = \{c_{edit}^i\}_{i=1}^N = \{concat(p^i, I^i \odot m^i, I_{edit})\}_{i=1}^N. \quad (4)$$

This video diffusion model removes the artifacts of point cloud rendering by inpainting the missing holes in rendering results, and corrects the slight geometry distortion caused by point cloud estimation. It also analyzes the relationship between the newly generated components and the original video to achieve seamless fusion.

Training Strategy. We use a self-supervised training strategy to train the video diffusion models. We create paired data, including input frames I , the first frame's point cloud rendering images p , and randomly generated 2D mask M imitating user-drawn cases. We utilize the mask propagation approach in Sec. 3.3 to obtain a mask sequence $\{m_i\}_{i=1}^N$. These conditions are concatenated as the input condition signals c similar to Eq. 4, with the first frame I^1 as the editing input of I_{edit} . The proposed training strategy is the same as that used for video inpainting but with explicit 3D guidance. The video denoising transformer network ϵ_θ is optimized by the following diffusion loss:

$$L(\theta) = \mathbb{E}_{t \sim u(0,1), \epsilon \sim N(0,1)} [\|\epsilon_\theta(z_t, t, y, c) - \epsilon\|_2^2], \quad (5)$$

where $z_t = \alpha_t z_0 + \sigma_t \epsilon$, with z_0 being the encoded ground-truth latent code of frames I . α_t and σ_t are hyperparameters of the diffusion process. The model takes an additional input of text prompt y detected by LLaVA [Liu et al. 2023a] during training and inference.

4 EVALUATION

In Sec. 4.1, we show the video editing results of our method. We have conducted extensive qualitative and quantitative experiments, including comparison (Sec. 4.2), ablation study (Sec.4.3), and user study (Sec. 4.4) to validate the superiority of our method compared with existing or baseline approaches.

Implementation Details. Our video diffusion model is built based on CogVideoX-2b [Yang et al. 2024c], with an additional conditional network [Zhang et al. 2023] composed of 10 control blocks. We fixed the weights of base video generation model while updating the condition network. Our model is trained on the DL3DV [Ling et al. 2024] and RealEstate [Zhou et al. 2018] datasets. The training pairs include videos, point cloud renderings, and random masks. We first train the control network with only point cloud rendering inputs with 20,000 steps and then with additional masked video inputs for an additional 15,000 steps. More details on the network, training, and data processing are in the supplemental materials.

4.1 Results

Our method enables detailed sketch-based editing for scene videos with significant viewpoint changes. As shown in Fig. 3, given user-drawn sketches, masks, and text prompts, our method automatically synthesizes realistic edited videos. The video diffusion model generates high-quality novel view results for editing regions while accurately preserving the unedited regions. Thanks to our 3D analysis and condition guidance, the newly generated parts exhibit reasonable view transformations and seamlessly merge with the original scene. Our method supports diverse editing operations, including object insertion, replacement, removal, and component modification. It works for input videos with various camera transformations, including rotation, translation, and zoom variation. Please refer to the supplemental video to examine edited video results.

4.2 Comparison

To validate the superiority of our method, we compare it with existing techniques. Specifically, we compare with video editing methods, including AnyV2V [Ku et al. 2024] and I2VEdit [Ouyang et al. 2024a], which propagate the editing of the first frame to the entire video. We also compare with camera-controllable video generation methods, including ViewExtra [Liu et al. 2024c] and ViewCrafter [Yu et al. 2024], which take point cloud rendering as inputs (same as our method) and generate video results.

AnyV2V [Ku et al. 2024] is a training-free method that merges the features of the original video with edited images. It is effective for appearance editing, but encounters information conflict for geometry editing, resulting in shape blending of leaves and flowers in Fig. 4 (b). I2VEdit [Ouyang et al. 2024a] utilizes trainable LoRa to distill motion from input video, generating results (Fig. 4 (c)) with clear shapes but exhibiting fuzzy details in frames distant from the edited image. ViewExtrapolator [Liu et al. 2024c] (shortened as ViewExtra) is a training-free method that refines the point cloud rendering with diffusion prior but struggles with large viewpoint changes, leading to frame quality degradation, as shown in Fig. 4 (d). ViewCrafter [Yu et al. 2024] fine-tunes the video diffusion model to accept additional input of point cloud rendering, making it more

Table 1. The quantitative comparisons with existing methods. Our method outperforms existing approaches in all metrics.

Method	Auto Evaluations		Human Preference			
	CLIP↑	PSNR↑	VC	UP	TC	EQ
AnyV2V	92.28	20.09	4.21	4.18	4.25	4.40
I2VEdit	94.92	22.31	3.23	3.05	3.18	3.25
ViewExtra	94.57	17.76	3.75	3.83	3.75	3.74
ViewCrafter	94.95	17.94	2.68	2.80	2.66	2.48
Ours	95.47	31.20	1.14	1.15	1.15	1.12

robust for view changes. However, the generated components (e) exhibit strange geometry, and the unedited regions are completely changed. Merging the results into the input video (f) preserves the unedited regions, but the editing regions are inconsistent and display obvious seam artifacts. Our method (g) generates high-quality video editing results with clear details and accurate preservation of unedited regions.

Quantitative Comparison. We further employ automatic metrics to measure the performance of the compared approaches. Following [Mou et al. 2024a], we utilize the CLIP [Radford et al. 2021] score between two consecutive frames to assess temporal consistency. The PSNR metric in unedited regions is used to measure the preservation of input videos. We randomly select 20 examples from all our editing cases to calculate the above-mentioned metrics. As shown in Table 1, our method significantly outperforms other methods in PSNR, validating that our method effectively preserves the features of unedited regions. Our method also outperforms existing methods in CLIP similarity, supporting the best temporal consistency.

4.3 Ablation Study

We conducted an ablation study to validate the necessity of our key components. In our method, depth maps are utilized to represent the geometry of edited regions, followed by scale and shift transformations to align with the original 3D scene. Since we utilize DUSt3R [Wang et al. 2024a] to estimate the point cloud of the edited images, an intuitive baseline is to directly use traditional alignment approaches, such as ICP [Besl and McKay 1992], to align the new point cloud with the original point cloud in 3D space. During editing, since only the unedited regions have correspondences, we calculate the rotation and translation based on the edited point cloud in unedited regions and then apply transformations to the whole edited point cloud. However, for this baseline, the edited point cloud has a slightly different scale from the original one, and the correspondences of unedited regions are not considered. As shown in Fig. 5 (b), without the depth-guided alignment, the point cloud rendering results have holes in the first frame, and the new edited content is in the wrong position, resulting in a strange geometry of the new content. Our method generates reasonable point cloud rendering (d) with realistic video editing results (e).

During video generation, the masked original video and point cloud rendering serve as conditions to control the unedited and edited regions, respectively. As shown in Fig. 6, if we remove the conditions of 3D-aware masks and the original video, the results (c) cannot preserve the original features, such as the different shapes

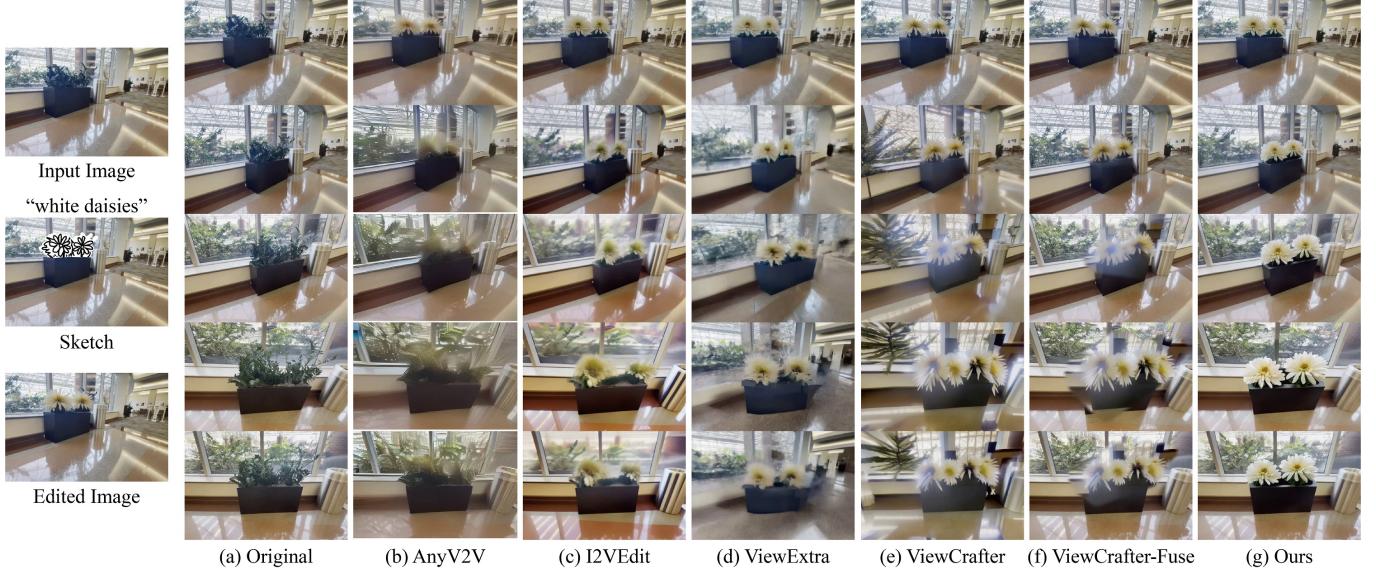


Fig. 4. Comparison with existing approaches. Given original video (a) and sketch-based editing on the first frame (Left), existing video editing methods, including AnyV2V [Ku et al. 2024] and I2VEdit [Ouyang et al. 2024a], generate fuzzy details. 3D-aware video generation methods, including ViewExtrapolation [Liu et al. 2024c] and ViewCrafter [Yu et al. 2024] cannot handle significant viewpoint changes. Directly fusing the generation results of Viewcrafter with the original video in pixel space improves the results but with obvious seams (f). In contrast, our method (g) generates the most realistic video editing results. Original Video ©DL3DV-10K.

Table 2. The quantitative results of the ablation study. Our method outperforms all baseline approaches. LPIPS values are scaled by a factor of 100.

Metric	w/o depth	w/o mask	w/o point cloud	ours
PSNR \uparrow	25.87	18.29	23.62	26.88
LPIPS \downarrow	9.12	20.35	9.65	8.28

of stones and green leaves, because these occlusion components are imagined by the models. If we remove the point cloud rendering (d), the explicit 3D guidance is missing, and thus, the complicated structure (shape of red leaves) cannot be maintained during viewpoint changes. Our method (e) generates view-consistent new content and accurately preserves the unedited regions.

We further conduct quantitative experiments to validate the effectiveness of each component. We randomly select 30 test examples from DL3DV and RealEState and estimate the point cloud and camera parameters. We detect object masks using SAM [Kirillov et al. 2023] and then remove a random object in the point cloud and original video to imitate the editing process. We treat the first frame of the original videos as edited images and then utilize our method to reconstruct the original videos. As shown in Table 2, our method achieves the best reconstruction values, validating that our method outperforms the baseline approaches.

4.4 User Study

We conducted a user study to visually compare our method with the existing approaches. We randomly selected 10 editing examples and invited 24 people to participate in our questionnaires. They ranked

different methods in four aspects: viewpoint change consistency with the original video (VC), unedited preservation (UP), temporal consistency (TC), and overall editing quality (EQ). The rankings provided by the participants were used as scores. As shown in Table 1, our method outperforms others in all criteria, validating the superior performance of our method.

5 APPLICATIONS

5.1 Camera-Controllable Video Generation

As shown in Fig. 7, given a single image, similar to ViewCrafter [Yu et al. 2024], users can define a camera trajectory to produce point cloud rendering results, which serve as conditions to achieve camera-controllable video generation. Users can further edit the input image to update the generated videos while preserving unedited regions, showcasing detailed controllable generation ability.

5.2 More Editing interaction tools

Stroke-based Video Editing. As shown in Fig. 8, users can draw color strokes to change the appearance of objects in images (supported by MagicQuill [Liu et al. 2024d]), and the editing is reasonably propagated to the whole video. Additionally, as shown in the island insertion examples in Fig. 8, the sketches and color strokes can be simultaneously input to control the geometry and appearance.

Inpainting-based Video Editing. Our method extends versatile image inpainting methods to 3D scene video editing. As shown in Fig. 9, text-based [Alibaba 2024] or reference-based [Yang et al. 2023a] image inpainting methods can be utilized to edit the first

frame. These edits are then well propagated throughout the other frames to generate realistic video results.

6 CONCLUSION

In this paper, we have presented a novel sketch-based 3D-aware video editing method to change the objects and components in videos involving significant viewpoint changes. Our method supports diverse editing operations, including insertion, removal, replacement, and appearance modification of components and objects. The editing operations are applied to the first frame and propagated throughout the video. We utilize an explicit 3D point cloud representation to handle viewpoint changes and a depth-guided point cloud editing approach to represent the geometry of new content and align it with the original scene. To identify and preserve the unedited regions, we propose a 3D-aware mask propagation strategy, followed by a video diffusion model to generate the final realistic video editing results. Extensive experiments support the superior performance of our method compared with existing approaches.

Limitation and Discussion. Our method utilizes DUST3R [Wang et al. 2024a] to estimate the camera parameters and point cloud during 3D analysis. However, for challenging cases far from the training data of DUST3R, as shown in Fig. 10, it predicts incorrect 3D information, which is inherited during the editing and causes obvious artifacts. More robust 3D analysis approaches [Tang et al. 2024; Zhang et al. 2024b] or model fine-tuning with data augmentation might be adopted to address this issue. Additionally, due to the restriction of the training dataset, our method currently cannot handle cases with 360° viewpoint changes. Employing a large-scale dataset with progressive generation could possibly solve this problem. Moreover, one limitation of our method is the restriction on mostly static 3D scenes. Processing videos of objects simultaneously with large dynamic motion and significant viewpoint changes is interesting research, but it is challenging because disentangling this information is required for reasonable propagation. Furthermore, our method might introduce slight blurriness in background details, such as the sofa in Fig. 3 and the trees in Fig. 5. A potential solution is to leverage more powerful video diffusion models.

ACKNOWLEDGMENTS

This work was sponsored by Beijing Municipal Science and Technology Commission (No. Z231100005923031), Innovation Funding of ICT, CAS (No. E461020), and the National Natural Science Foundation of China (No. 62322210). The authors would like to acknowledge the Nanjing Institute of InforSuperBahn OneAiNexus for providing the training and evaluation platform.

REFERENCES

Alibaba. 2024. <https://github.com/alimama-creative/FLUX-Controlnet-Inpainting>

Sherwin Bahmani, Ivan Skorokhodov, Aliaksandr Siarohin, Willi Menapace, Guocheng Qian, Michael Vasilkovsky, Hsin-Ying Lee, Chaoyang Wang, Jiaxu Zou, Andrea Tagliasacchi, David B. Lindell, and Sergey Tulyakov. 2024. VD3D: Taming Large Video Diffusion Transformers for 3D Camera Control. *CoRR* abs/2407.12781 (2024).

Jianhong Bai, Menghan Xia, Xintao Wang, Ziyang Yuan, Xiao Fu, Zuozhu Liu, Haoji Hu, Pengfei Wan, and Di Zhang. 2024. SynCamMaster: Synchronizing Multi-Camera Video Generation from Diverse Viewpoints. *arXiv*:2412.07760

Paul J. Besl and Neil D. McKay. 1992. A Method for Registration of 3-D Shapes. *IEEE Trans. Pattern Anal. Mach. Intell.* 14, 2 (1992), 239–256.

Kirill Brodt and Mikhail Bessmeltsev. 2022. Sketch2Pose: estimating a 3D character pose from a bitmap sketch. *ACM Trans. Graph.* 41, 4 (2022), 85:1–85:15.

Duygu Ceylan, Chun-Hao Paul Huang, and Niloy J. Mitra. 2023. Pix2Video: Video Editing using Image Diffusion. In *Int. Conf. Comput. Vis.* 23149–23160.

Shu-Yu Chen, Wanchoo Su, Lin Gao, Shihong Xia, and Hongbo Fu. 2020. DeepFace-Drawing: deep generation of face images from sketches. *ACM Trans. Graph.* 39, 4 (2020), 72.

Chengying Gao, Qi Liu, Qi Xu, Limin Wang, Jianzhuang Liu, and Changqing Zou. 2020. SketchyCOCO: Image Generation From Freehand Scene Sketches. In *IEEE Conf. Comput. Vis. Pattern Recog.* 5173–5182.

Chenjian Gao, Xilin Wang, Qian Yu, Lu Sheng, Jing Zhang, Xiaoguang Han, Yi-Zhe Song, and Dong Xu. 2024. 3D Reconstruction From a Single Sketch via View-Dependent Depth Sampling. *IEEE Trans. Pattern Anal. Mach. Intell.* 46, 12 (2024), 9661–9676.

Zekai Gu, Rui Yan, Jiahao Lu, Peng Li, Zhiyuan Dou, Chenyang Si, Zhen Dong, Qifeng Liu, Cheng Lin, Ziwei Liu, Wenping Wang, and Yuan Liu. 2025. Diffusion as Shader: 3D-aware Video Diffusion for Versatile Video Generation Control. *CoRR* abs/2501.03847 (2025).

Xiaoguang Han, Chang Gao, and Yizhou Yu. 2017. DeepSketch2Face: a deep learning based sketching system for 3D face and caricature modeling. *ACM Trans. Graph.* 36, 4 (2017), 126:1–126:12.

Hao He, Yinghao Xu, Yuwei Guo, Gordon Wetzstein, Bo Dai, Hongsheng Li, and Ceyuan Yang. 2024. CameraCtrl: Enabling Camera Control for Text-to-Video Generation. *CoRR* abs/2404.02101 (2024).

Chen Hou, Guoqiang Wei, Yan Zeng, and Zhibo Chen. 2024. Training-free Camera Control for Video Generation. *arXiv*:2406.10126

Zhitong Huang, Mohan Zhang, and Jing Liao. 2024. LVCD: Reference-based Lineart Video Colorization with Diffusion Models. *ACM Trans. Graph.* 43, 6 (2024), 177:1–177:11.

Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. 2017. Image-to-Image Translation with Conditional Adversarial Networks. In *IEEE Conf. Comput. Vis. Pattern Recog.* 5967–5976.

Ondrej Jamriska, Sárka Sochorová, Ondrej Texler, Michal Lukáč, Jakub Fiser, Jingwan Lu, Eli Shechtman, and Daniel Sykora. 2019. Styling video by example. *ACM Trans. Graph.* 38, 4 (2019), 107:1–107:11.

Lifan Jiang, Shuang Chen, Boxi Wu, Xiaotong Guan, and Jiahui Zhang. 2025. VidSketch: Hand-drawn Sketch-Driven Video Generation with Diffusion Control. *CoRR* abs/2502.01101 (2025).

Youngjoo Ju and Jongyool Park. 2019. SC-FEGAN: Face Editing Generative Adversarial Network With User’s Sketch and Color. In *Int. Conf. Comput. Vis.* 1745–1753.

Xuan Ju, Xian Liu, Xintao Wang, Yuxuan Bian, Ying Shan, and Qiang Xu. 2024. BrushNet: A Plug-and-Play Image Inpainting Model with Decomposed Dual-Branch Diffusion. In *Eur. Conf. Comput. Vis.*, Vol. 15078. 150–168.

Yoni Kasten, Dolev Ofri, Oliver Wang, and Tali Dekel. 2021. Layered Neural Atlases for Consistent Video Editing. *ACM Trans. Graph.* 40, 6 (2021), 210:1–210:12.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.* 42, 4 (2023), 139:1–139:14.

Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloé Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross B. Girshick. 2023. Segment Anything. In *Int. Conf. Comput. Vis.* 3992–4003.

Max Ku, Cong Wei, Weiming Ren, Harry Yang, and Wenhui Chen. 2024. AnyV2V: A Tuning-Free Framework For Any Video-to-Video Editing Tasks. *arXiv preprint arXiv:2403.14468* (2024).

Kuaishou. 2024. <https://kling.kuaishou.com>

PKU- Yuan Lab and Tuzhan AI etc. 2024. Open-Sora-Plan. <https://doi.org/10.5281/zenodo.10948109>

Jinyu Li, Xiaokun Pan, Gan Huang, Ziyang Zhang, Nan Wang, Hujun Bao, and Guofeng Zhang. 2024. RD-VIO: Robust Visual-Inertial Odometry for Mobile Augmented Reality in Dynamic Environments. *IEEE Trans. Vis. Comput. Graph.* 30, 10 (2024), 6941–6955.

Xiaoyu Li, Bo Zhang, Jing Liao, and Pedro V. Sander. 2022. Deep Sketch-Guided Cartoon Video Inbetweening. *IEEE Trans. Vis. Comput. Graph.* 28, 8 (2022), 2938–2952.

Lu Ling, Yichen Sheng, Zhi Tu, Wentian Zhao, Cheng Xin, Kun Wan, Lantao Yu, Qianyu Guo, Zixun Yu, Yawen Lu, Xuanmao Li, Xingpeng Sun, Rohan Ashok, Aniruddha Mukherjee, Hao Kang, Xiangru Kong, Gang Hua, Tianyi Zhang, Bedrich Benes, and Aniket Bera. 2024. DL3DV-10K: A Large-Scale Scene Dataset for Deep Learning-based 3D Vision. In *IEEE Conf. Comput. Vis. Pattern Recog.* 22160–22169.

Feng-Lin Liu, Shu-Yu Chen, Yu-Kun Lai, Chunpeng Li, Yue-Ren Jiang, Hongbo Fu, and Lin Gao. 2022. DeepFaceVideoEditing: sketch-based deep editing of face videos. *ACM Trans. Graph.* 41, 4 (2022), 167:1–167:16.

Feng-Lin Liu, Hongbo Fu, Yu-Kun Lai, and Lin Gao. 2024a. SketchDream: Sketch-based Text-To-3D Generation and Editing. *ACM Trans. Graph.* 43, 4 (2024), 44:1–44:13.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023a. Visual Instruction Tuning. In *Adv. Neural Inform. Process. Syst.*

Kunhao Liu, Ling Shao, and Shijian Lu. 2024c. Novel View Extrapolation with Video Diffusion Priors. *CoRR* abs/2411.14208 (2024).

Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. 2023b. Zero-1-to-3: Zero-shot One Image to 3D Object. In *Int. Conf. Comput. Vis.* 9264–9275.

Shaoteng Liu, Yuechen Zhang, Wenbo Li, Zhe Lin, and Jiaya Jia. 2024e. Video-P2P: Video Editing with Cross-Attention Control. In *IEEE Conf. Comput. Vis. Pattern Recog.* 8599–8608.

Zhiheng Liu, Hao Ouyang, Qiuyu Wang, Ka Leong Cheng, Jie Xiao, Kai Zhu, Nan Xue, Yu Liu, Yujun Shen, and Yang Cao. 2024b. InFusion: Inpainting 3D Gaussians via Learning Depth Completion from Diffusion Prior. *CoRR* abs/2404.11613 (2024).

Zichen Liu, Yue Yu, Hao Ouyang, Qiuyu Wang, Ka Leong Cheng, Wen Wang, Zhiheng Liu, Qifeng Chen, and Yujun Shen. 2024d. MagicQuill: An Intelligent Interactive Image Editing System. *CoRR* abs/2411.09703 (2024).

Fei Luo, Yongqiang Zhu, Yanping Fu, Huajian Zhou, Zexheng Chen, and Chunxia Xiao. 2023. Sparse RGB-D images create a real thing: A flexible voxel based 3D reconstruction pipeline for single object. *Vis. Informatics* 7, 1 (2023), 66–76.

Hao Ma, Jingyuan Yang, and Hui Huang. 2024. Taming diffusion model for exemplar-based image translation. *Computational Visual Media* 10, 6 (2024), 1031–1043.

Aryan Mikaeili, Or Perel, Mehdi Safaei, Daniel Cohen-Or, and Ali Mahdavi-Amiri. 2023. SKED: Sketch-guided Text-based 3D Editing. In *Int. Conf. Comput. Vis.* 14561–14573.

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2022. NeRF: representing scenes as neural radiance fields for view synthesis. *Commun. ACM* 65, 1 (2022), 99–106.

Ashkan Mirzaei, Tristan Amentado-Armstrong, Konstantinos G. Derpanis, Jonathan Kelly, Marcus A. Brubaker, Igor Gilitschenski, and Alex Levinshtein. 2023. SPI-NeRF: Multiview Segmentation and Perceptual Inpainting with Neural Radiance Fields. In *IEEE Conf. Comput. Vis. Pattern Recog.* 20669–20679.

Chong Mou, Mingdeng Cao, Xintao Wang, Zhaoyang Zhang, Ying Shan, and Jian Zhang. 2024a. ReVideo: Remake a Video with Motion and Content Control. *CoRR* abs/2405.13865 (2024).

Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, and Ying Shan. 2024b. T2I-Adapter: Learning Adapters to Dig Out More Controllable Ability for Text-to-Image Diffusion Models. In *AAAI* 4296–4304.

OpenAI. 2024. Sora Overview: <https://openai.com/index/sora/>. <https://openai.com/index/sora/>

Hao Ouyang, Qiuyu Wang, Yuxi Xiao, Qingyan Bai, Juntao Zhang, Kecheng Zheng, Xiaowei Zhou, Qifeng Chen, and Qifeng Chen. 2024b. CoDef: Content Deformation Fields for Temporally Consistent Video Processing. In *IEEE Conf. Comput. Vis. Pattern Recog.* 8089–8099.

Wenqi Ouyang, Yi Dong, Lei Yang, Jianlou Si, and Xingang Pan. 2024a. I2VEDit: First-Frame-Guided Video Editing via Image-to-Video Diffusion Models. In *ACM SIGGRAPH Asia*, Takeo Igarashi, Ariel Shamir, and Hao (Richard) Zhang (Eds.). 95:1–95:11.

Tiziano Portenier, Qiyang Hu, Attila Szabó, Siavash Arjomand Bigdeli, Paolo Favaro, and Matthias Zwicker. 2018. Faceshop: deep sketch-based face image editing. *ACM Trans. Graph.* 37, 4 (2018), 99.

Chenyang Qi, Xiaodong Cun, Yong Zhang, Chenyang Lei, Xintao Wang, Ying Shan, and Qifeng Chen. 2023. FateZero: Fusing Attentions for Zero-shot Text-based Video Editing. In *Int. Conf. Comput. Vis.* 15886–15896.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision (*Proceedings of Machine Learning Research*, Vol. 139). 8748–8763.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *IEEE Conf. Comput. Vis. Pattern Recog.* 10674–10685.

Manuel Ruder, Alexey Dosovitskiy, and Thomas Brox. 2018. Artistic Style Transfer for Videos and Spherical Images. *Int. J. Comput. Vis.* 126, 11 (2018), 1199–1219.

runway. 2024. <https://runwayml.com/research/introducing-gen-3-alpha>

Junyoung Seo, Kazumi Fukuda, Takashi Shibuya, Takuuya Narihira, Naoki Murata, Shoukang Hu, Chieh-Hsin Lai, Seungryong Kim, and Yuki Mitsuishi. 2024. GenWarp: Single Image to Novel Views with Semantic-Preserving Generative Warping. In *Adv. Neural. Inform. Process. Syst.*

Vincent Sitzmann, Semon Rezhikov, Bill Freeman, Josh Tenenbaum, and Frédéric Durand. 2021. Light Field Networks: Neural Scene Representations with Single-Evaluation Rendering. In *Adv. Neural. Inform. Process. Syst.* 19313–19325.

StabilityAI. 2024. <https://www.stablevideo.com/>

Jia-Mu Sun, Tong Wu, and Lin Gao. 2024b. Recent advances in implicit representation-based 3D shape generation. *Vis. Intell.* 2, 1 (2024). <https://doi.org/10.1007/S44267-024-00042-1>

Wenqiang Sun, Shuo Chen, Fangfu Liu, Zilong Chen, Yueqi Duan, Jun Zhang, and Yikai Wang. 2024a. DimensionX: Create Any 3D and 4D Scenes from a Single Image with Controllable Video Diffusion. *CoRR* abs/2411.04928 (2024).

Zhenggang Tang, Yuchen Fan, Dilin Wang, Hongyu Xu, Rakesh Ranjan, Alexander Schwing, and Zhicheng Yan. 2024. MV-DUS3R+: Single-Stage Scene Reconstruction from Sparse Views In 2 Seconds. *arXiv preprint arXiv:2412.06974* (2024).

Vchitect Team and Shanghai Artificial Intelligence Laboratory. 2024. *Vchitect-2.0: Parallel Transformer for Scaling Up Video Diffusion Models*. <https://github.com/Vchitect/Vchitect-2.0>

Rotem Tzaban, Ron Mokady, Rinon Gal, Amit Bermano, and Daniel Cohen-Or. 2022. Stitch it in Time: GAN-Based Facial Editing of Real Videos. In *SIGGRAPH Asia Conference Papers*. 29:1–29:9.

Lili Wang, Qinglin Qi, Yi Chen, Wei Ke, and Aimin Hao. 2014. Interactive texture design and synthesis from mesh sketches. *Frontiers Comput. Sci.* 8, 2 (2014), 330–341.

Qinghe Wang, Yawen Luo, Xiaoyu Shi, Xu Jia, Huchuan Lu, Tianfan Xue, Xintao Wang, Pengfei Wan, Di Zhang, and Kun Gai. 2025. CineMaster: A 3D-Aware and Controllable Framework for Cinematic Text-to-Video Generation. *CoRR* abs/2502.08639 (2025).

Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jérôme Revaud. 2024a. DUST3R: Geometric 3D Vision Made Easy. In *IEEE Conf. Comput. Vis. Pattern Recog.* 20697–20709.

Wen Wang, Kangyang Xie, Zide Liu, Hao Chen, Yue Cao, Xinlong Wang, and Chunhua Shen. 2023. Zero-Shot Video Editing Using Off-The-Shelf Image Diffusion Models. *CoRR* abs/2303.17599 (2023).

Yuelei Wang, Jian Zhang, Peng-Tao Jiang, Hao Zhang, Jinwei Chen, and Bo Li. 2024c. CPA: Camera-pose-awareness Diffusion Transformer for Video Generation. *CoRR* abs/2412.01429 (2024).

Zhouxiao Wang, Ziyang Yuan, Xintao Wang, Yaowei Li, Tianshui Chen, Menghan Xia, Ping Luo, and Ying Shan. 2024b. MotionCtrl: A Unified and Flexible Motion Controller for Video Generation. In *ACM SIGGRAPH*. 114.

Shuchen Weng, Haojie Zheng, Peixuan Zhang, Yuchen Hong, Han Jiang, Si Li, and Boxin Shi. 2024. VIRES: Video Instance Repainting with Sketch and Text Guidance. *CoRR* abs/2411.16199 (2024).

Rundi Wu, Ruiqi Gao, Ben Poole, Alex Trevithick, Changxi Zheng, Jonathan T. Barron, and Aleksander Holynski. 2024. CAT4D: Create Anything in 4D with Multi-View Video Diffusion Models. *CoRR* abs/2411.18613 (2024).

Haoran Xie, Keisuke Arihara, Syuhei Sato, and Kazunori Miyata. 2024b. DualSmoke: Sketch-based smoke illustration design with two-stage generative model. *Computational Visual Media* 10, 5 (2024), 965–979.

Tianhao Xie, Noam Aigerman, Eugene Belilovsky, and Tiberiu Popa. 2024a. Sketch-guided Cage-based 3D Gaussian Splatting Deformation. *arXiv:2411.12168 [cs.CV]* <https://arxiv.org/abs/2411.12168>

Jinbo Xing, Hanyuan Liu, Menghan Xia, Yong Zhang, Xintao Wang, Ying Shan, and Tien-Tsin Wong. 2024. ToonCrafter: Generative Cartoon Interpolation. *ACM Trans. Graph.* 43, 6 (2024), 245:1–245:11.

Dejia Xu, Weili Nie, Chao Liu, Sifei Liu, Jan Kautz, Zhangyang Wang, and Arash Vahdat. 2024. CamCo: Camera-Controllable 3D-Consistent Image-to-Video Generation. *CoRR* abs/2406.02509 (2024).

Yukun Xu, Keyang Ye, Tianjia Shao, and Yanlin Weng. 2025. Animatable 3D Gaussians for modeling dynamic humans. *Frontiers Comput. Sci.* 19, 9 (2025), 199704. <https://doi.org/10.1007/S11704-024-40497-5>

Binxin Yang, Shuyang Gu, Bo Zhang, Ting Zhang, Xuejin Chen, Xiaoyan Sun, Dong Chen, and Fang Wen. 2023a. Paint by Example: Exemplar-based Image Editing with Diffusion Models. In *IEEE Conf. Comput. Vis. Pattern Recog.* 18381–18391.

Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. 2024b. Depth Anything V2. *CoRR* abs/2406.09414 (2024).

Shiyuan Yang, Liang Hou, Haibin Huang, Chongyang Ma, Pengfei Wan, Di Zhang, Xiaodong Chen, and Jing Liao. 2024a. Direct-a-Video: Customized Video Generation with User-Directed Camera Movement and Object Motion. In *ACM SIGGRAPH*. 113.

Shuai Yang, Yifan Zhou, Ziwei Liu, and Chen Change Loy. 2023b. Rerender A Video: Zero-Shot Text-Guided Video-to-Video Translation. In *ACM SIGGRAPH*, June Kim, Ming C. Lin, and Bernd Bickel (Eds.). 95:1–95:11.

Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. 2024c. CogVideoX: Text-to-Video Diffusion Models with An Expert Transformer. *arXiv preprint arXiv:2408.06072* (2024).

Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S. Huang. 2019. Free-Form Image Inpainting With Gated Convolution. In *Int. Conf. Comput. Vis.* 4470–4479. <https://doi.org/10.1109/ICCV.2019.00457>

Wangbo Yu, Jinbo Xing, Li Yuan, Wenbo Hu, Xiaoyu Li, Zhipeng Huang, Xiangjun Gao, Tien-Tsin Wong, Ying Shan, and Yonghong Tian. 2024. ViewCrafter: Taming Video Diffusion Models for High-fidelity Novel View Synthesis. *CoRR* abs/2409.02048 (2024).

Yu Zeng, Zhe Lin, and Vishal M. Patel. 2022. SketchEdit: Mask-Free Local Image Manipulation with Partial Sketches. In *IEEE Conf. Comput. Vis. Pattern Recog.* 5941–5951.

Junyi Zhang, Charles Herrmann, Junhwa Hur, Varun Jampani, Trevor Darrell, Forrester Cole, Deqing Sun, and Ming-Hsuan Yang. 2024b. MonST3R: A Simple Approach for Estimating Geometry in the Presence of Motion. *CoRR* abs/2410.03825 (2024).

Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2023. Adding Conditional Control to Text-to-Image Diffusion Models. In *Int. Conf. Comput. Vis.* 3813–3824.

Tingyang Zhang, Qingzhe Gao, Weiyu Li, Libin Liu, and Baoquan Chen. 2024a. BAGS: Building Animatable Gaussian Splatting from a Monocular Video with Diffusion Priors. *CoRR* abs/2403.11427 (2024).

Yabo Zhang, Xinpeng Zhou, Yihan Zeng, Hang Xu, Hui Li, and Wangmeng Zuo. 2025. FramePainter: Endowing Interactive Image Editing with Video Diffusion Priors. *CoRR* abs/2501.08225 (2025).

Zhixing Zhang, Bichen Wu, Xiaoyan Wang, Yaqiao Luo, Luxin Zhang, Yinan Zhao, Peter Vajda, Dimitris N. Metaxas, and Licheng Yu. 2024c. AVID: Any-Length Video Inpainting with Diffusion Model. In *IEEE Conf. Comput. Vis. Pattern Recog.* 7162–7172.

Guangcong Zheng, Teng Li, Rui Jiang, Yehao Lu, Tao Wu, and Xi Li. 2024a. CamI2V: Camera-Controlled Image-to-Video Diffusion Model. *CoRR* abs/2410.15957 (2024).

Wangguandong Zheng, Haifeng Xia, Rui Chen, Libo Sun, Ming Shao, Siyu Xia, and Zhengming Ding. 2024c. Sketch3D: Style-Consistent Guidance for Sketch-to-3D Generation. In *ACM Int. Conf. Multimedia*. 3617–3626.

Xin-Yang Zheng, Hao Pan, Peng-Shuai Wang, Xin Tong, Yang Liu, and Heung-Yeung Shum. 2023. Locally Attentional SDF Diffusion for Controllable 3D Shape Generation. *ACM Trans. Graph.* 42, 4 (2023), 91:1–91:13.

Zangwei Zheng, Xiangyu Peng, Tianji Yang, Chenhui Shen, Shenggui Li, Hongxin Liu, Yukun Zhou, Tianyi Li, and Yang You. 2024b. *Open-Sora: Democratizing Efficient Video Production for All.* <https://github.com/hpcatech/Open-Sora>

Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. 2018. Stereo magnification: learning view synthesis using multiplane images. *ACM Trans. Graph.* 37, 4 (2018), 65.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In *Int. Conf. Comput. Vis.* 2242–2251.



Fig. 5. The ablation study of point cloud alignment. (a) An input video. (b, d) Edited point cloud rendering results. (c, e) Edited videos. Our method generates a more reasonable edited point cloud than the baseline without depth-guided alignment, resulting in better video editing results with reasonable geometry and component interaction. Original Video ©DL3DV-10K.

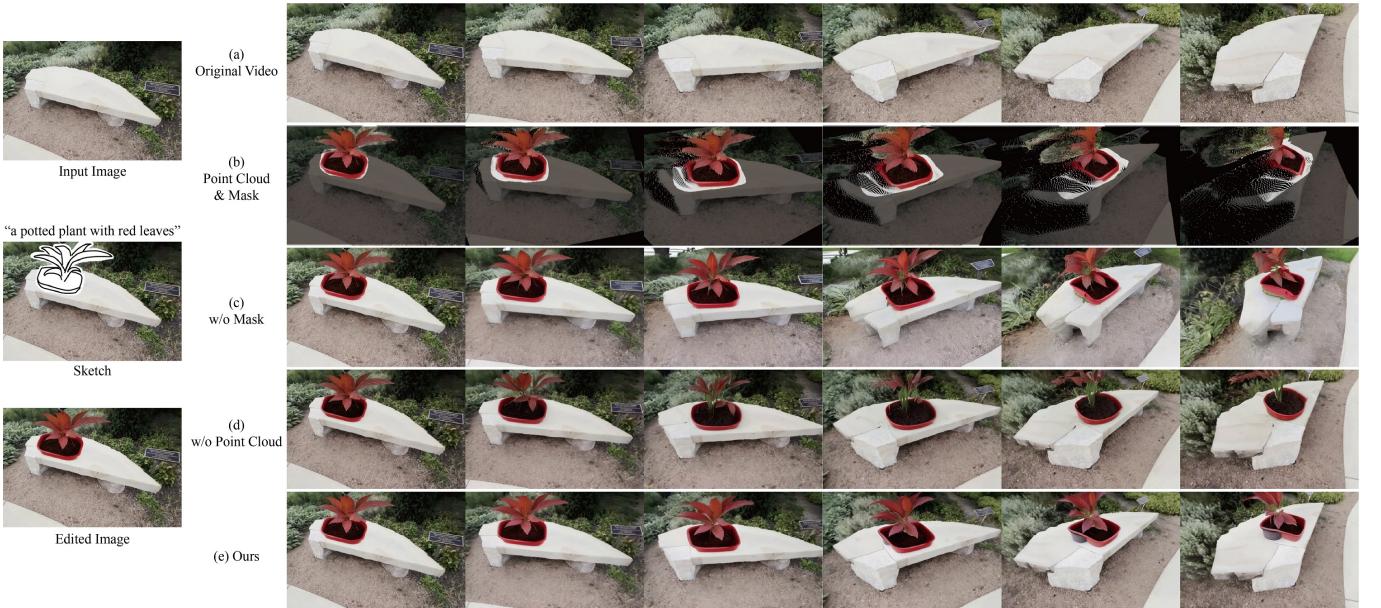


Fig. 6. The ablation study of 3D-aware mask and point cloud rendering condition. Given the original video (a) and editing operations (Left), our method produces the edited point cloud rendering and 3D-aware masks (b). Without the masks (c), the unedited regions (e.g., the shape of stone) are not preserved. Without the point cloud rendering (d), the shape of red leaves is changed in different frames. Our method (e) generates the best results. Original Video ©DL3DV-10K.



Fig. 7. The application of camera-controllable video generation. Given a single input image, users can control the camera viewpoints and generate a video with novel view results. Local region editing can be further added to achieve detailed, controllable video generation. Input Image ©Cats Coming.

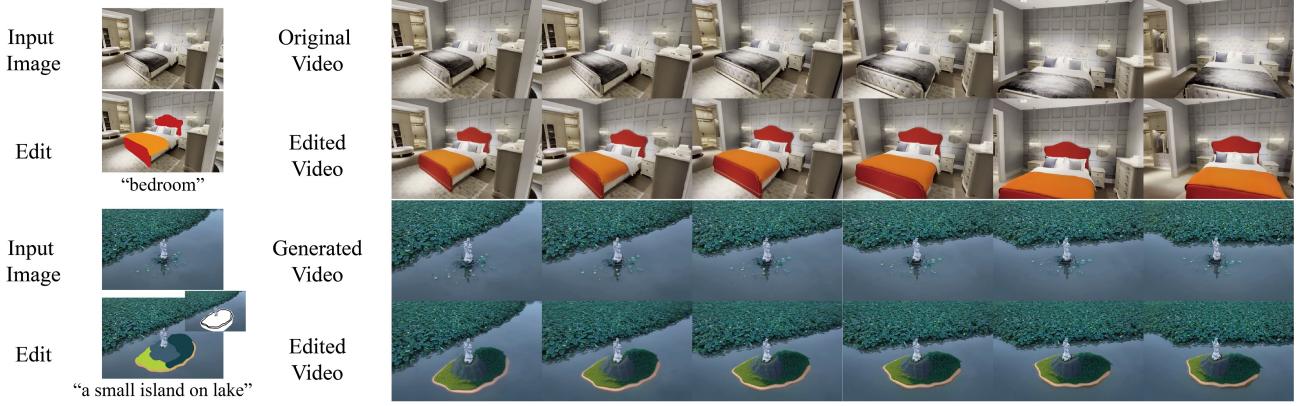


Fig. 8. The application of color stroke-based video editing. The appearance of existing content (bed in the 1st example) can be edited by using color strokes. Users can also control the appearance of newly generated objects, with the editing sketch shown in the top right corner in the 2nd example. Original Video ©DL3DV-10K.

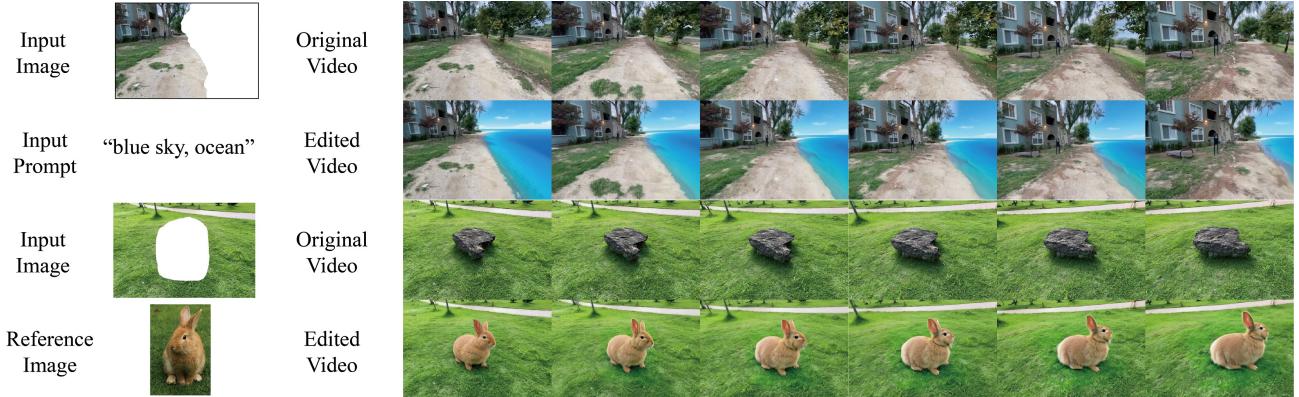


Fig. 9. The application of image inpainting-based video editing. These examples utilize the text-based [Alibaba 2024] and reference-based [Yang et al. 2023a] image inpainting to editing the first frame. The editing operations are well propagated into videos. Original Video ©DL3DV-10K, Reference image ©Pixabay.



Fig. 10. Failure cases. For the video that DUSt3R [Wang et al. 2024a] estimates wrong point clouds and/or camera parameters, the edited point cloud and mask are also mistakenly rendered. The quality of the generated video is severely affected. Original Image ©stein egil liland.