

Consistent Video Editing as Flow-Driven Image-to-Video Generation

Ge Wang^{1,*} Songlin Fan^{1,*} Hangxu Liu¹ Qianjian Song² Hewei Wang³ Jinfeng Xu^{4,†}
¹Fudan University ²Xiamen University ³Carnegie Mellon University ⁴The University of Hong Kong

{wangg25, slfan, 23210720103}@fudan.edu.cn

songqj@stu.xmu.edu.cn, heweiw@alumni.cmu.edu, jinfeng@connect.hku.hk

*Equal contribution †Corresponding author

Abstract

With the prosper of video diffusion models, video editing have been significantly promoted without consuming much computational cost. One particular challenge in this task lies at the motion transfer process from the source video to the edited one, where it requires the consideration of the shape deformation in between, meanwhile maintaining the temporal consistency in the generated video sequence. However, existing methods fail to model complicated motion patterns for video editing, and are fundamentally limited to object replacement, where tasks with non-rigid object motions like multi-object and portrait editing are largely neglected. In this paper, we observe that optical flows offer a promising alternative in complex motion modeling, and present FlowV2V to re-investigate video editing as a task of flow-driven Image-to-Video (I2V) generation. Specifically, FlowV2V decomposes the entire pipeline into first-frame editing and conditional I2V generation, and simulates pseudo flow sequence that aligns with the deformed shape, thus ensuring the consistency during editing. Experimental results on DAVIS-EDIT with improvements of 13.67% and 50.66% on DOVER and warping error illustrate the superior temporal consistency and sample quality of FlowV2V compared to existing state-of-the-art ones. Furthermore, we conduct comprehensive ablation studies to analyze the internal functionalities of the first-frame paradigm and flow alignment in the proposed method.

1. Introduction

With the exponential growth of digital content, video has become a central medium for daily recording, where this data modality has prompted multiple applications in the field of computer vision, typically including video editing. To achieve video editing, traditional methods are normally labor-intensive due to the manual operations, and inherently limited in scalability, where automatic solutions are

thus expected. With the development of diffusion models [6, 8–10, 16, 19, 23, 25, 33, 37], video editing has been promoted to a new era, enabled with their strong generative capabilities. To perform the task of video editing, early-proposed methods [14, 15, 18, 26, 35, 39, 43] mainly rely on the motion trajectories of the original video, where they perform the editing process under the assumption that edited objects do not undergo significant shape changes and use techniques like DDIM inversion [13, 15, 18] and one-shot tuning [15, 18]. However, these methods fall short in the scenarios involving significant shape deformation, as is noted by StableV2V [14], where in real-world editing cases, users might expect to modify the geometry of objects—particularly in tasks such as multi-object replacement or facial editing—where preserving rigid structure is thus largely neglected in these studies. Considering the alignment of shape deformation, StableV2V constructs its pipeline based on the first-frame-based methods—which sequentially edits the first video frame and then propagates it to the others through a conditional Image-to-Video (I2V) process—and utilizes depth maps as the intermediate media for motion control. Yet, depth map inherently struggles to handle complex and non-rigid shape transformations, as it fails to represent motion trajectories when complicated patterns like occlusion, rotation, facial expression changes occur. More effective condition to represent complex motion patterns is thus expected for the field of video editing.

In this paper, we propose FlowV2V to re-organize video editing as the flow-guided I2V generation process, where we follow the first-frame-based paradigm, and adopt shape-aligned flow for motion control throughout the editing process. Our method comprises four main components, including First Frame Editing (FFE), Iterative Motion Propagation (IMP), Shape-Consistent Flow Calibration (SCFC), and Flow-Driven Image-to-Video Generation (FD-I2V). Sequentially, FFE processes prompts of various modalities (e.g., text, instruction, reference image, and so on) and edits the first video frame and latter plays as the initial image condition for the I2V model in FD-I2V. IMP performs an it-



Figure 1. Qualitative comparisons of non-rigid motion editing (top), including face rotation, and additional applications (bottom), such as face editing and multi-motion, using our proposed FlowV2V.

erative alignment process, where it obtains a pseudo optical flow sequence with the consideration of shape alignment to the edited objects, and afterward, SCFC further refines the flow sequence to facilitate its preciseness. Eventually, FD-I2V utilizes a flow-driven I2V model to propagate the contents in first edited frame to all others, and outputs the entire edited video sequence in the manner of I2V generation. To prove the validity of FlowV2V, we evaluate our method on DAVIS-EDIT [14], where results show our superior sample quality, controllability, and flexibility, along with comprehensive ablation studies analyzing the impacts of the modality of control signal and various edited frames.

2. Related Works

Video Diffusion Model. Diffusion models have achieved remarkable success in a wide series of visual generation tasks [4, 5, 7, 36]. Two particular video-based topics are primarily studied—Text-to-Video (T2V) and Image-to-Video (I2V) generation. For T2V generation, methods such as WAN [32] and Hunyuan-DiT [12] jointly model spatial and temporal dimensions with the network architecture of Diffusion Transformer (DiT) [33]. For I2V generation, I2VGen-XL [27] and Step-Video¹ are two most relevant studies in recent years. Besides, SVD [2] adopts a spatial-temporal decoupling strategy to balance spatial quality and temporal modeling efficiency, so as to conduct better I2V generation. Generally speaking, the aforementioned models enables a wide series of applications, e.g., style transfer, video editing, and etc., especially without tuning task-specific models, where they serve as the fundamental sup-

port for future development of the visual generation community.

Video Editing. In recent years, video editing has emerged as an important topic. Existing diffusion-based video editing methods can be broadly categorized into four groups—one-shot tuning methods, DDIM inversion-based methods, learning-based methods, and first-frame-based methods [14, 15, 18, 26, 28, 35, 39, 43]. Specifically, one-shot tuning methods [15, 35] adapt model weights for a given video through video-specific tuning, where it enables diverse editing results conditioned on different user-provided text prompts. DDIM inversion-based methods [13, 15, 18] reconstruct latent representations by inverting diffusion trajectories, then capture motion patterns and enforcing temporal consistency across generated frames. Learning-based methods [41, 43] integrate motion modules into pre-trained image diffusion models and tune dedicated video-based inpainting models with a large amount of video-text pairs. Compared to all the aforementioned methods, first-frame-based methods [13, 14, 21] offer a flexible paradigm to process various prompts from different modalities, starting by modifying the first frame and propagating the results to subsequent frames via motion trajectory control, with more details provided in the next paragraph.

First-Frame-based Editing Methods. This type of methods leverage the modified first frame to drive the entire editing process. Specifically, one prominent approach is AnyV2V [13], which utilizes image editing models to modify the first frame, and then uses an I2V model to propagate the edited contents via temporal feature injection. Based on AnyV2V, I2VEdit adaptively preserves both the visual

¹<https://github.com/stepfun-ai/Step-Video-Ti2V>

and motion integrity of the source video and adjusts the attention maps considering both global and local modifications. However, due to a lack of effective alignment between the transferred motion and the edited content, video editing cases with significant shape changes normally fail. To address this issue, StableV2V [14] aligns the depth map sequence with the edited objects, and establishes alignment between the transferred motion and the user prompts, thus enabling a shape-consistent video editing process. Even so, depth-map contains its intrinsic problem, where it normally struggle to represent complicated motion patterns like occlusions and rotation, where more effective ways to handle such scenarios are expected, thus promoting our work in this paper.

3. Methods

To implement FlowV2V, we follow previous first frame-based methods [13, 14, 21] to decompose the entire video editing pipeline into two parts, i.e., first frame editing and conditional image-to-video generation, where the overview of this work is shown in Figure 2. As demonstrated, FlowV2V contains four procedures, including First Frame Editing (FFE), Iterative Motion Propagation (IMP), Shape-Consistent Flow Calibration (SCFC), and Flow-Driven Image-to-Video Generation (FD-I2V). Specifically, FFE accepts external text or image prompts to add edited contents into the first frame of the input video. Afterward, IMP extracts the optical flow from the input video, and propagates the averaged motions from the regions of the original object to the ones of edited object, and iteratively perform such operation until all edited frames are processed, resulting in the pseudo optical flow to control I2V generation. To further ensure the precision of obtained pseudo flow, we draw inspiration from existing video inpainting studies [42] to calibrate redundant motion information in SCFC. Eventually, we integrate the aforementioned control and transfer the edited contents in the first video frame to all others in a manner of conditional image-to-video generation. For detailed illustration of the above procedures, we construct our paper in the structures below, with Section 3.1, 3.2, 3.3, and 3.4 illustrating FFE, IMP, SCFC, and FD-I2V, respectively.

3.1. First Frame Editing

Similar to previous first-frame-based studies [13, 14, 21], the first step is to process the external prompt into edited contents in the first video frame. To achieve so, we consider off-the-shelf image editors to process prompts from various modalities accordingly. Let P denotes the external prompt offered by users, we send the first video frame $I_1 \in V = \{I_i\}_{i=1}^N$ into the image editor $\mathcal{G}_{img}(\cdot)$, with V as the input video sequence and N referring to its number of frames, respectively, where we simply write the mathematical formulation of this process:

$$\hat{I}_1 = \mathcal{G}_{img}(I_1, P), \quad (1)$$

where \hat{I}_1 represents the edited first frame. For selection of $\mathcal{G}_{img}(\cdot)$, particularly, we use SD Inpaint [25], Instruct-Pix2Pix [30], and Paint-by-Example [3] to handle external prompts in modalities of text, instruction, and reference image, respectively. For the editing scenarios in the human portrait domain, we adopt a portrait editing toolkit in Jimeng AI².

3.2. Iterative Motion Propagation

Once we obtained the edited first frame \hat{I}_1 , the next step is to perform the conditional image-to-video generation process and generate the entire edited video. Previous studies like StableV2V [14] specifically address the shape misalignment problem in the task of video editing, where it utilizes an iterative alignment algorithm to obtain shape-aligned depth maps as the control signal. Although such problem is addressed to some extents, we observe that depth maps are often suboptimal to represent motions of moving objects, where depth-based methods are fundamentally limited to the editing cases of rigid body motion. When handling non-rigid editing scenarios, especially the ones with object occlusions and rotation (as is shown in Figure 1), depth-based methods struggle to propagate the edited contents, since depth map normally fails to precisely depict object motions in such cases. Optical flow seems to be a promising alternative for such problem, yet obtaining feasible flows with shape alignment to the edited object poses a significant challenge, which motivates our proposal of Iterative Motion Propagation (IMP) in this work.

Optical Flow Extraction. The first step in IMP is to extract the optical flow from the input video sequence V . Given the i -th video frame $I_i \in V = \{I_i\}_{i=1}^N$, we denote the optical flow considering the direction of I_i to I_{i+1} as \vec{F}_i . Then, we use a pre-trained optical flow extractor, i.e., RAFT [29], which predicts \vec{F}_i with the inputs of I_i and I_{i+1} , with the overall process formulated by:

$$\vec{F} = \{\vec{F}_i\}_{i=1}^{N-1} = \{\mathcal{E}_{flow}(I_i, I_{i+1})\}_{i=1}^{N-1} \quad (2)$$

where $\mathcal{E}_{flow}(\cdot)$ denotes the optical flow extractor and \vec{F} represents the entire optical flow sequence.

Shape-Aware Flow Deformation. With the optical flow sequence extracted, we acquire a control signal for the Image-to-Video (I2V) generation model, yet it is not feasible to directly apply \vec{F} on the I2V model for the reason that there exists shape misalignment between the edited object and the original one. To alleviate such misalignment, it is vital to deform the original optical flow sequence with consideration of the shape changes, where we propose Shape-Aware Flow Deformation to conduct this process. Since the

²<https://jimeng.jianying.com/ai-tool>

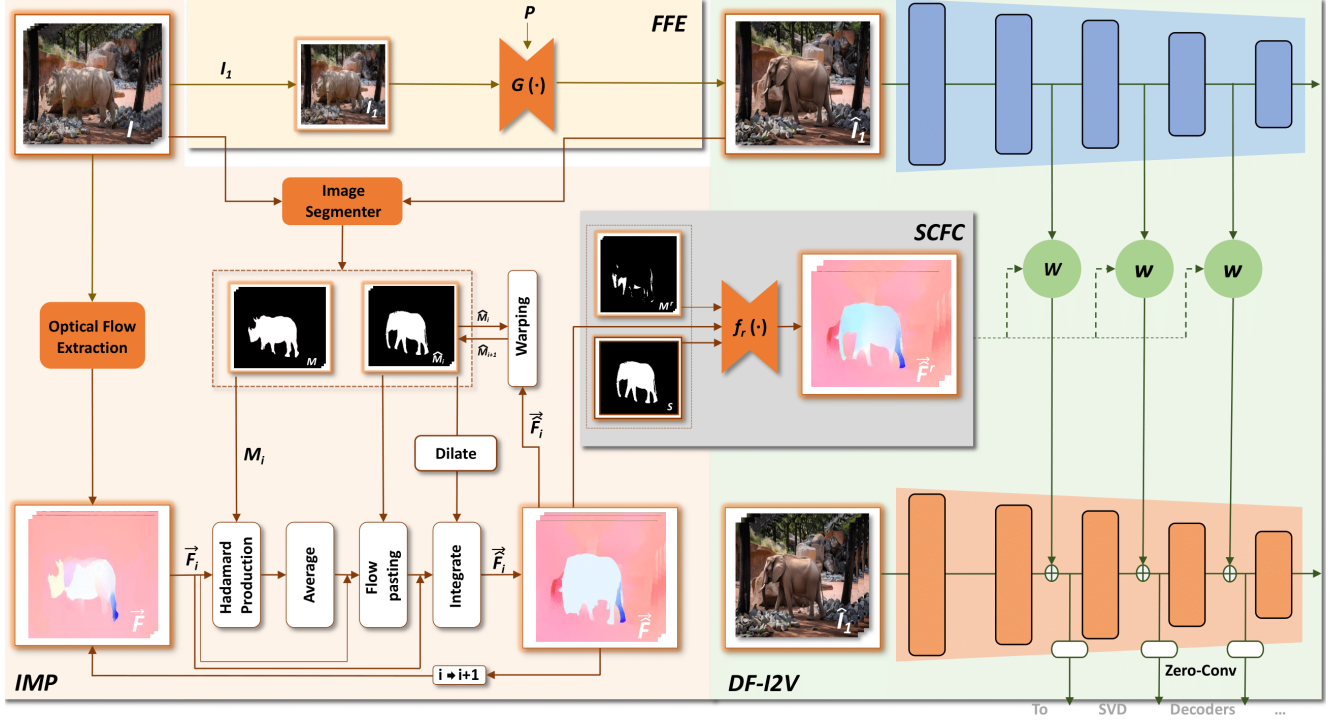


Figure 2. **The overall architecture of our proposed FlowV2V.** It consists of a first-frame editor, an iterative motion propagator, a shape-consistent flow calibration module, a reference encoder capturing multi-scale features from the source image, and a trainable SVD encoder initialized from SVD, serving as the final spatial-temporal feature merger in generation guidance.

edited object and the original one are required to contain consistent motion information under the task of video editing, it is feasible to assume that the average flow values in both regions are theoretically the same. Therefore in IMP, we utilize the binary mask sequence as an intermediate condition in order to transfer the motion information from the edited regions to the original ones. In details, we first adopt a pre-trained image segmentation model (i.e., SAM [11]) to obtain a binary mask sequence $M = \{M_i\}_{i=1}^N$ and the first frame mask \widehat{M}_1 from the input video V and the edited first frame \widehat{I}_1 . Let $\vec{F}(m, n)$ denote the flow value at the m -th column and n -th row, we compute the average flow value termed \bar{F} within M , and assign \bar{F} to each flow point in \widehat{M}_1 , where the aforementioned processes are written as:

$$\begin{aligned}
 M &= \{M_i\}_{i=1}^N = \mathcal{E}_{seg}(I_1, \dots, I_N), \\
 \widehat{M}_1 &= \mathcal{E}_{seg}(\widehat{I}_1), \\
 \bar{F}_1 &= \frac{1}{|M_1|} \sum_{(m,n) \in M_1} \vec{F}(m, n), \\
 \vec{F}_1 &= \bar{F}_1 \odot \widehat{M}_1 + \vec{F}_1(1 - \widehat{M}_1),
 \end{aligned} \tag{3}$$

where \vec{F}_1 represents the pseudo optical flow of the first frame with alignment to the shape of \widehat{I}_1 ; \odot indicates the Hadamard product between matrices; $|M_1|$ denotes the number of pixels in M_1 ; $\mathcal{E}_{seg}(\cdot)$ is the image segmentation model; 1 refers to an all-one matrix. For further il-

lustration of this process, $\vec{F}_1(m, n)$ is the same as the average flow value \bar{F}_1 when $(m, n) \in \widehat{M}_1$, otherwise equals to $\vec{F}_1(m, n)$. One challenge herein is that segmentation masks of other frames, i.e., $\{\widehat{M}_i\}_{i=2}^N$, are not available, but we require them to acquire the entire pseudo optical flow sequence $\vec{F} = \{\vec{F}_i\}_{i=1}^{N-1}$. Based on the binary-valued nature of \widehat{M}_1 , we utilize \vec{F}_1 to obtain \widehat{M}_{i+1} through warping \widehat{M}_1 . Let $\vec{F}_i(m, n) = (\mathbf{u}_i(m, n), \mathbf{v}_i(m, n))$ represents the displacement of each pixel from its position (m, n) in \widehat{M}_i to its next-frame position $(m + \mathbf{u}_i(m, n), n + \mathbf{v}_i(m, n))$ in \widehat{M}_{i+1} , where \mathbf{u}_i and \mathbf{v}_i represents two vectors of i -th frame representing the motion directions along the height and width dimensions, respectively. Then, we use \vec{F}_i to perform the warping operation on \widehat{M}_i as follows:

$$\widehat{M}_{i+1}(m, n) = \widehat{M}_i(m + \mathbf{u}_i(m, n), n + \mathbf{v}_i(m, n)). \tag{4}$$

By implement Equation 4 on 3, we can obtain the segmentation mask \widehat{M}_2 of the second edited frame, which allows us to conduct Equation 3 based on the 2-nd frame elements to obtain \vec{F}_2 . Eventually, we repeat the aforementioned iteration with the elements of the i -th frame starting from $i = 1$ to $i = N - 1$, and eventually get the pseudo optical flow sequence $\vec{F} = \{\vec{F}_i\}_{i=1}^{N-1}$ of the edited video.

Table 1. Quantitative comparison on DAVIS-EDIT [14] with respect to DOVER [34], FVD [31], WE, CLIP-Temporal, and CLIP score [1]. We highlight the scores tested under the original height-width ratio with “*”. \bar{T} represents the average inference time per example in minutes across both DAVIS-EDIT-S and DAVIS-EDIT-C. The optimal and suboptimal results are in **red** and **blue**, respectively.

Method	DAVIS-Edit-S					DAVIS-Edit-C					$\bar{T}\downarrow$
	DOVER \uparrow	FVD \downarrow	WE \downarrow	CLIP-T \uparrow	CLIP-S \uparrow	DOVER \uparrow	FVD \downarrow	WE \downarrow	CLIP-T \uparrow	CLIP-S \uparrow	
<i>Text-based Evaluation Settings</i>											
TokenFlow	66.36	17.33	18.58	95.84	24.89	67.47	17.45	18.60	95.61	24.12	5.81
DMT	59.27	19.53	16.65	94.11	24.91	57.45	21.64	19.89	93.58	24.51	8.88
I2VEdit	69.63	15.63	16.94	96.18	24.99	69.57	17.62	16.16	96.50	24.92	10.25
AnyV2V	66.82	14.87	15.35	95.66	25.09	65.01	17.83	18.26	94.36	24.32	8.28
StableV2V	67.78	13.77	15.95	96.34	25.46	70.80	17.18	15.27	96.83	25.68	3.14
FlowV2V	77.20	15.56	5.61	97.05	25.17	76.96	17.28	11.43	96.94	24.98	2.01
FlowV2V*	71.77	17.10	9.31	96.99	24.20	72.49	18.37	9.54	96.93	25.43	4.26
<i>Image-based Evaluation Settings</i>											
AnyV2V	65.83	12.97	24.47	95.89	25.41	64.56	15.25	25.61	96.13	24.79	8.43
StableV2V	67.58	12.36	22.17	96.51	26.24	68.42	14.87	21.23	96.71	26.55	3.23
FlowV2V	80.62	13.98	14.88	97.18	25.51	77.18	14.63	6.02	96.74	24.83	2.23
FlowV2V*	76.36	17.06	11.41	97.19	24.97	70.32	16.63	14.74	96.84	24.98	4.57

3.3. Shape-Consistent Flow Calibration

With both FFE and IMP finished, we are to offer a control signal for the I2V model to propagate the first-frame edited contents to all others, however in real practices, we observe that such pseudo optical flow may contain redundant regions that tend to misguide the I2V generation process. This observation particularly occurs when the edited object has non-overlapped regions with the original one, making it an urgent need to improve the preciseness of the control signal by removing these redundant parts in the pseudo optical flow sequence. On one hand, regarding the requirement of content removal, the task of inpainting fundamentally offers a reference paradigm to solve this requirement, by training neural networks to learn the inpainting process in a self-supervised manner. From another aspect, we expect to preserve the established shape alignment through the flow calibration process, which thus needs specific condition to guide the refinement networks. To this end, we draw inspiration from previous flow-based inpainting and editing studies [22, 42] with a tailored network to perform Shape-Consistent Flow Calibration (SCFC) for FlowV2V. Note that in this paper, we only use the SCFC on the editing cases with obvious shape deformation, otherwise we use the default optical flow from the original video sequence without specified.

Flow Calibration Network. We design our flow calibration network for purpose of shape-consistent flow refinement based on the one from ProPainter [42], where it accepts three inputs, i.e., the corrupted flow sequence \vec{F}^m , the binary mask sequence M^r , and the shape guidance S . To obtain the above elements, we first corrupt the ground-truth optical flow \vec{F} with M^r , and then send the concatenation of the resulting flow sequence along with M^r and S into the flow calibration network $\mathcal{G}_{flow}(\cdot)$, where the afore-

mentioned processes are formulated as:

$$\begin{aligned}\vec{F}^m &= \vec{F} \odot M^r, \\ \vec{F}^r &= \mathcal{G}_{flow}(\vec{F}^m, M^r, S),\end{aligned}\tag{5}$$

where \vec{F}^r denotes the refined optical flow sequence that is used to control the I2V model afterwards. Here, we compute two types of \vec{F} to enhance the temporal modeling ability of the flow calibration network, i.e., forwarding and backward flows, which are termed as \vec{F}_f and \vec{F}_b , respectively. In training, we follow the random masking algorithm in FGT [40] to obtain M^r randomly corrupt \vec{F} . In inference, we use the pseudo flow sequence \vec{F} from IMP instead of \hat{F} in Equation 5 in order to achieve flow calibration. As for S , we use the segmentation mask M_1 of the first video frame I_1 in training, and adopt the one (denoted as \hat{M}_1) of edited first frame \hat{I}_1 in inference.

Network Architecture. We implement a recurrent flow completion network that performs both forward and backward flow estimation to achieve temporally consistent and high-fidelity results. The network follows an encoder-decoder architecture, where the encoder consists of multiple P3D blocks that progressively downsample the input flow sequences (by a factor of 8) and extract hierarchical spatio-temporal representations. In the bottleneck, a dilated convolution module is employed to enlarge the receptive field and capture long-range motion dependencies. To ensure temporal coherence, a Bidirectional Propagation Module with Second-Order Deformable Alignment propagates flow features between adjacent frames in both temporal directions, where the downsampled features are warped according to their estimated optical flows and refined through stacked deformable convolution layers that predict adaptive

offsets. The decoder then reconstructs dense flow fields via multi-stage upsampling and convolutional refinement, incorporating skip connections from the encoder to recover fine-grained motion details. Finally, an auxiliary edge detection branch is adopted to enhance motion boundaries and maintain structural consistency across frames.

Optimization Objective. To optimize the flow calibration network, we compute the reconstruction loss \mathcal{L}_{rec} through the L1 norm-based Euclidean distance between estimated flow sequence and the ground-truth one. Besides, we adopt a second-order smoothness loss \mathcal{L}_{smooth} following UnFlow [17] to promote the smoothness of neighbor regions in the estimated flow³.

3.4. Flow-Driven Image-to-Video Generation

Once the calibrated flow sequence is obtained, we are finally able to produce the final edited video. In details, we use Stable Video Diffusion (SVD) [2] as our foundation generative model, and insert flow-based adapter layers into SVD, similar to the ones in MOFA-Video [20]. To produce the final edited video $\widehat{V} = \{\widehat{I}\}_{i=1}^N$, we send the edited first frame \widehat{I}_1 and the refined optical flow sequence \widehat{F}^r into SVD to propagate the edited contents in \widehat{I}_1 to all others in a conditional I2V manner.

Hierarchical Feature Warping and Aggregation. The flow-based adapter composes a reference encoder and a feature fusion encoder, where the former adopts multi-scale convolutional blocks to extract hierarchical features from the input first frame; the latter warps the extracted features and fuses them to the corresponding layers in a duplicated SVD encoder network, following a similar strategy of ControlNet [16]. Then, the outputted features of the duplicated SVD encoder are aggregated to the SVD decoder, and control the I2V process in producing a motion-aligned edited video sequence.

4. Implementation Details

Evaluation Setting. We follow the evaluation setting of in StableV2V [14], where we utilize its proposed DAVIS-EDIT dataset. Specifically, the dataset consists of two subsets to address editing cases with Similar (S) and Changing (C) shape deformation, i.e., DAVIS-EDIT-S and DAVIS-EDIT-C, respectively. In total, it contains 50 editing instances, with 28 and 22 videos for text-based and image-based evaluations, where we employ its official annotations.

Experiment Setting. We employ RAFT [29] for optical flow estimation, using 20 inference iterations to balance accuracy and computation. Following the estimated flows, we synthesize video frames using pre-trained text-to-image diffusion models. However, due to the resolution constraints

³We illustrate more implementation details of the calibration network in Section A of our Appendix.

of these models (mostly based on Stable Diffusion [25]), we primarily generate video frames under the resolution of 512×512 , where notably, StableV2V also follows this setting. Yet, we discover that videos from DAVIS [24] (the data source of DAVIS-Edit) contain a certain height-width ratio, where the setting of 512×512 resolution fundamentally compress the video to an unnatural size. Therefore, we adopt both the default setting (with height-width ratio of 1 : 1) as well as the original ratio when conducting video editing, where quantitative results under the latter setting are highlighted with *.

Evaluation Metrics. We benchmark methods along three dimensions, i.e., visual fidelity, temporal coherence, text-image alignment, where we adopt DOVER, FVD, inter-frame Warping Error (WE), CLIP-Temporal (CLIP-T), and image-text CLIP score. Besides, we conduct a user study to measure the human preferences of different methods.

5. Experiment

5.1. Comparison with State-of-the-Art Methods

In this section, we demonstrate and compare FlowV2V against several state-of-the-art video editing approaches, including StableV2V, AnyV2V, I2VEdit, TokenFlow and DMT. Methods based on the first frame use the same initial frame to ensure fairness.

Performance Comparison. Table 1 presents the quantitative evaluation of FlowV2V against existing methods on the DAVIS-EDIT [14]. Here, TokenFlow [18] and DMT [38] exhibit deficiencies in maintaining fidelity of the original video and obtain inferior temporal consistency on WE and CLIP-Temporal. Relatively, I2VEdit, AnyV2V and StableV2V achieve improved performance, yet the visual quality of their results are still limited. In contrast, FlowV2V demonstrates superior performance upon all evaluated metrics, particularly achieving outstanding sample quality and temporal consistency on DOVER, WE, and CLIP-Temporal, meanwhile preserving comparable performance on the other metrics, where this comparison confirms the effectiveness of our flow-driven editing paradigm. Notably, StableV2V points out that most studies present poorer on DAVIS-Edit-C rather than DAVIS-Edit-S, since the editing cases with significant shape changes are intrinsically more challenging, where FlowV2V performs consistently owing to the stable control from the aligned flow.

Despite of the classical usage shown in Figure 3, FlowV2V demonstrates superior on other challenging applications with complicated motion patterns, where the results are illustrated in Figure 4. Here, we particularly compare FlowV2V with depth-guided (StableV2V) and latent features-guided methods (AnyV2V), to show the advantages of optical flow under challenging scenarios. In Figure 4, StableV2V and AnyV2V produce still contents when

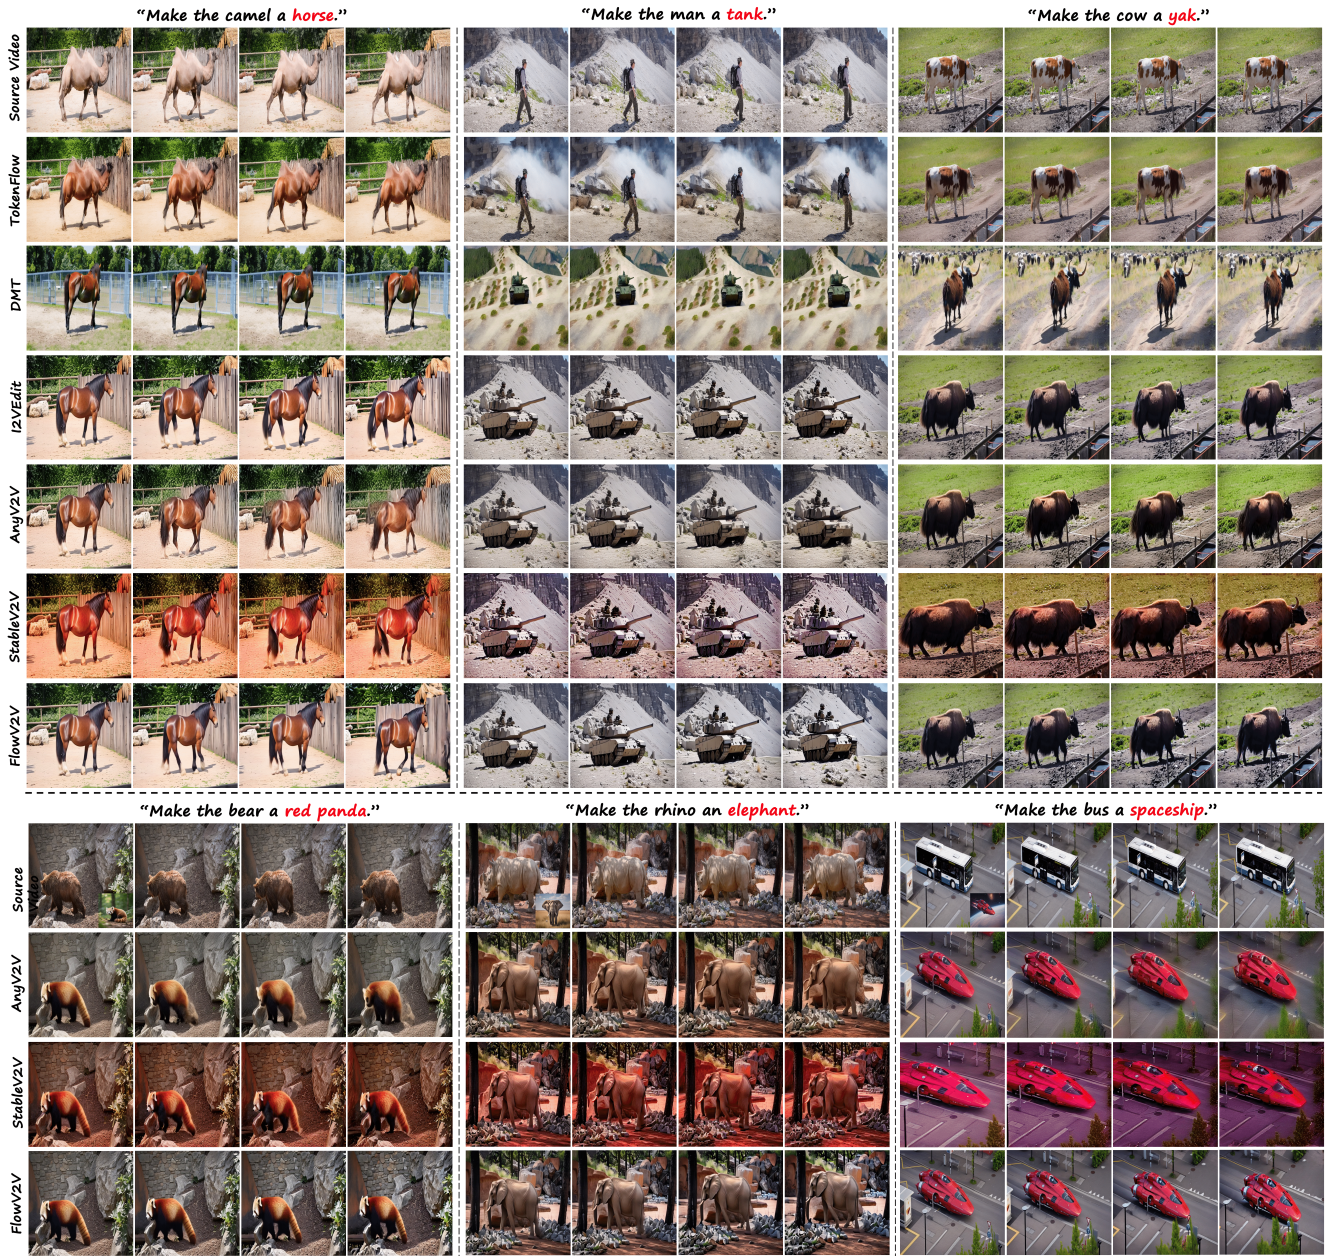


Figure 3. Qualitative comparison of FlowV2V with state-of-the-art methods.

editing human portrait, since both depth map sequence and latent features are fundamentally limited to model such complex motion patterns. Conversely, FlowV2V consistently edits all frames, especially when shape deformation occurs, confirming that the superior ability and robustness of our method. Similar phenomenon is observed under multi-object editing, where the compared methods produce videos with artifacts, and such observation is eliminated in our results⁴.

⁴We present more results in Section B of our Appendix.

5.2. Ablation Studies

Evaluating First Frame Editors for Video Editing. We evaluated the impact of various first-frame editors on FlowV2V, where results are presented in Figure 5. Specifically, we applied SD Inpaint [25], Paint-by-Example [3], InstructPix2Pix [30] to explore the impacts of text, image, and instruction prompts, respectively. Although different editors exhibit some visual differences in the edited frames, the overall quality of the final edited videos remains consistent, with no significant degradation or variation caused by the selection of editor. It is worth mentioning that aligning the motions offers a potential improvement on visual



Figure 4. Human portrait and multi-object editing results.

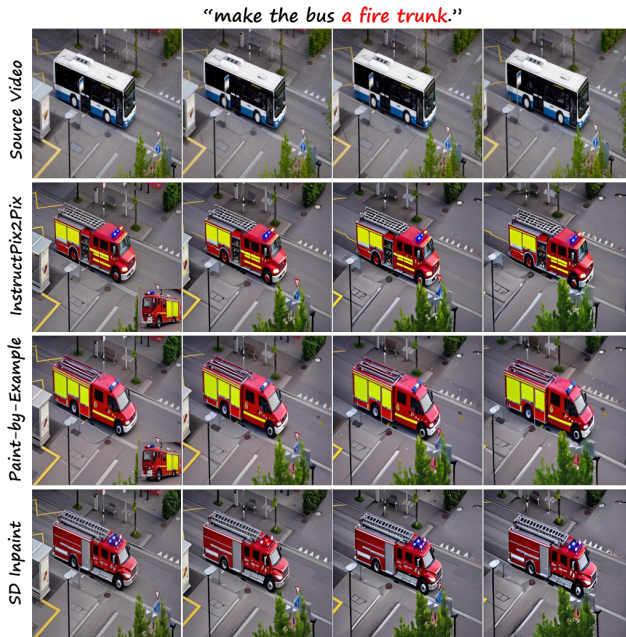


Figure 5. Qualitative results of different first-frame inputs.

quality, when editing a bus into a fire ambulance with SD Inpaint, where FlowV2V produces plausible with photorealistic details like flashing lights.

Evaluating the Impact of Different Control Signals. To verify the effectiveness of our optical flow-based condition-

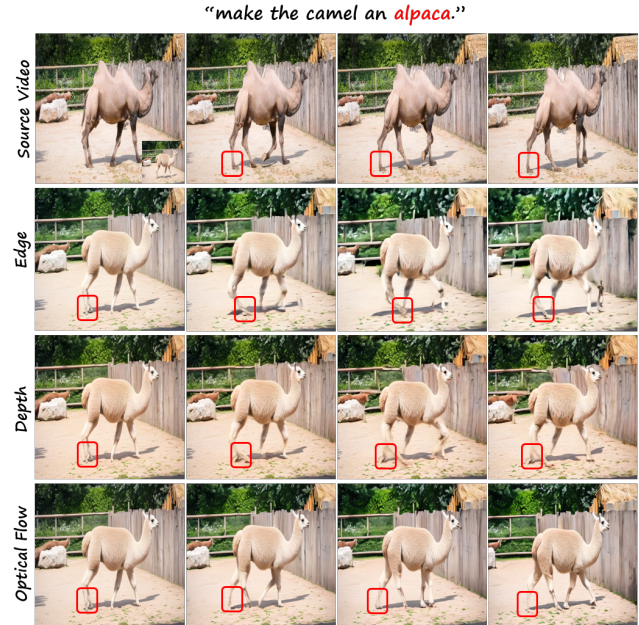


Figure 6. Qualitative results of different condition types.

ing strategy in the I2V task, we compare two different variants for the optical flow: depth and edge, as shown in Figure 6. Here, we highlight the local details by different methods so as to straightforward show their effects. Specifically, in scenarios involving relatively small deformations, using depth or edge as control inputs often leads to object dragging artifacts and motion blur in certain regions of the generated videos. As the deformation becomes more severe, artifacts become increasingly prominent, affecting larger areas and further degrading the video quality. For further illustration, depth and edge serve as two extreme cases in offering control signal for I2V models, where the former condition is natively coarse for fine-grained motion modeling; the latter condition might cause over-control problem⁵.

6. Conclusion

In this paper, we propose FlowV2V, a novel method to perform video editing through a flow-driven image-to-video generation process, along with a shape-aware flow alignment and a refinement procedures to ensure the preciseness of propagated motions during editing. Experiments on DAVIS-Edit validates the effectiveness and superior sample quality of FlowV2V, with further ablation studies offering more potential insights of it. Yet, we observe that FlowV2V might struggle to handle some editing scenarios, since its performance upper bound is fundamentally restricted by the generative power of the used I2V backbone, where we expect to investigate more possible applications with FlowV2V as stronger foundation models merge.

⁵We analyze the impact of RAFT’s estimation error in Section C, conduct an efficiency study in Section D, and discuss the limitations in Section E of our Appendix.

References

- [1] 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. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*, pages 8748–8763, 2021. 5
- [2] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, Varun Jampani, and Robin Rombach. Stable Video Diffusion: Scaling Latent Video Diffusion Models to Large Datasets. *arXiv*, 2023. 2, 6
- [3] Binxin Yang, Shuyang Gu, Bo Zhang, Ting Zhang, Xuejin Chen, Xiaoyan Sun, Dong Chen, and Fang Wen. Paint by Example: Exemplar-based Image Editing with Diffusion Models. In *CVPR*, pages 18381–18391, 2023. 3, 7
- [4] Yuren Cong, Mengmeng Xu, Christian Simon, Shoufa Chen, Jiawei Ren, Yanping Xie, Juan-Manuel Pérez-Rúa, Bodo Rosenhahn, Tao Xiang, and Sen He. FLATTEN: Optical FLOW-guided ATTENTION for Consistent Text-to-Video Editing. In *ICLR*, 2024. 2
- [5] Duygu Ceylan, Chun-Hao Paul Huang, and Niloy J. Mitra. Pix2Video: Video Editing using Image Diffusion. In *ICCV*, pages 23149–23160, 2023. 2
- [6] Xun Guo, Mingwu Zheng, Liang Hou, Yuan Gao, Yufan Deng, Pengfei Wan, Di Zhang, Yufan Liu, Weiming Hu, Zhengjun Zha, Haibin Huang, and Chongyang Ma. I2V-Adapter: A General Image-to-Video Adapter for Diffusion Models. In *SIGGRAPH*, page 112, 2024. 1
- [7] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. AnimateDiff: Animate Your Personalized Text-to-Image Diffusion Models without Specific Tuning. In *ICLR*, 2024. 2
- [8] Xuanhua He, Quande Liu, Shengju Qian, Xin Wang, Tao Hu, Ke Cao, Keyu Yan, and Jie Zhang. Id-animator: Zero-shot Identity-Preserving Human Video Generation. *arXiv*, 2024. 1
- [9] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising Diffusion Implicit Models. In *ICLR*, 2021.
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. In *NeurIPS*, 2020. 1
- [11] 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. Segment Anything. In *ICCV*, pages 3992–4003, 2023. 4
- [12] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, Kathrina Wu, Qin Lin, Junkun Yuan, Yanxin Long, Aladdin Wang, Andong Wang, Changlin Li, Duojun Huang, Fang Yang, Hao Tan, Hongmei Wang, Jacob Song, Jiawang Bai, Jianbing Wu, Jinbao Xue, Joey Wang, Kai Wang, Mengyang Liu, Pengyu Li, Shuai Li, Weiyan Wang, Wenqing Yu, Xincheng Deng, Yang Li, Yi Chen, Yutao Cui, Yuanbo Peng, Zhentao Yu, Zhiyu He, Zhiyong Xu, Zixiang Zhou, Zunnan Xu, Yangyu Tao, Qinglin Lu, Songtao Liu, Dax Zhou, Hongfa Wang, Yong Yang, Di Wang, Yuhong Liu, Jie Jiang, and Caesar Zhong. Hunyuanvideo: A systematic framework for large video generative models, 2025. 2
- [13] Max Ku, Cong Wei, Weiming Ren, Harry Yang, and Wenhu Chen. AnyV2V: A Tuning-Free Framework For Any Video-to-Video Editing Tasks. *arXiv*, 2024. 1, 2, 3
- [14] Chang Liu, Rui Li, Kaidong Zhang, Yunwei Lan, and Dong Liu. StableV2V: Stabilizing Shape Consistency in Video-to-Video Editing. *arXiv*, 2024. 1, 2, 3, 5, 6
- [15] Shaoteng Liu, Yuechen Zhang, Wenbo Li, Zhe Lin, and Jiaya Jia. Video-P2P: Video Editing with Cross-Attention Control. In *CVPR*, pages 8599–8608, 2024. 1, 2
- [16] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding Conditional Control to Text-to-Image Diffusion Models. In *ICCV*, pages 3813–3824, 2023. 1, 6
- [17] Simon Meister, Junhwa Hur, and Stefan Roth. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In *AAAI*, 2018. 6
- [18] Michal Geyer, Omer Bar-Tal, Shai Bagon, and Tali Dekel. TokenFlow: Consistent Diffusion Features for Consistent Video Editing. In *ICLR*, pages 1–13, 2024. 1, 2, 6
- [19] 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*, pages 22500–22510, 2023. 1
- [20] Muyao Niu, Xiaodong Cun, Xintao Wang, Yong Zhang, Ying Shan, and Yinqiang Zheng. Mofa-video: Controllable image animation via generative motion field adaptations in frozen image-to-video diffusion model. In *European Conference on Computer Vision*, pages 111–128, 2024. 6
- [21] Wenqi Ouyang, Yi Dong, Lei Yang, Jianlou Si, and Xingang Pan. I2VEdit: First-Frame-Guided Video Editing via Image-to-Video Diffusion Models. *arXiv*, 2024. 2, 3
- [22] Elia Peruzzo, Vidit Goel, Dejjia Xu, Xingqian Xu, Yifan Jiang, Zhangyang Wang, Humphrey Shi, and Nicu Sebe. VASE: Object-Centric Appearance and Shape Manipulation of Real Videos. *arXiv*, 2024. 5
- [23] Pierfrancesco Ardino and Marco De Nadai and Bruno Lepri and Elisa Ricci and Stéphane Lathuilière. Click to Move: Controlling Video Generation with Sparse Motion. In *ICCV*, pages 14729–14738, 2021. 1
- [24] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 DAVIS Challenge on Video Object Segmentation. *arXiv*, 2018. 6
- [25] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution Image Synthesis with Latent Diffusion Models. In *CVPR*, pages 10674–10685, 2022. 1, 3, 6, 7
- [26] Ruoyu Feng, Wenming Weng, Yanhui Wang, Yuhui Yuan, Jianmin Bao, Chong Luo, Zhibo Chen, and Baining Guo. CCEdit: Creative and Controllable Video Editing via Diffusion Models. *arXiv*, 2023. 1, 2
- [27] Shiwei Zhang, Jiayu Wang, Yingya Zhang, Kang Zhao, Hangjie Yuan, Zhiwu Qin, Xiang Wang, Deli Zhao, and Jingren Zhou. I2VGen-XL: High-Quality Image-to-Video Synthesis via Cascaded Diffusion Models. *arXiv*, 2023. 2

- [28] Quanjian Song, Mingbao Lin, Wengyi Zhan, Shuicheng Yan, Liujuan Cao, and Rongrong Ji. UniVST: A Unified Framework for Training-free Localized Video Style Transfer. *arXiv*, 2025. 2
- [29] Zachary Teed and Jia Deng. RAFT: Recurrent All-Pairs Field Transforms for Optical Flow. In *ECCV*, pages 402–419, 2020. 3, 6
- [30] Tim Brooks, Aleksander Holynski, and Alexei A. Efros. InstructPix2Pix: Learning to Follow Image Editing Instructions. In *CVPR*, pages 18392–18402, 2023. 3, 7
- [31] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphaël Marinier, Marcin Michalski, and Sylvain Gelly. FVD: A New Metric for Video Generation. In *ICLR Workshop*, 2019. 5
- [32] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, Jianyuan Zeng, Jiayu Wang, Jingfeng Zhang, Jingren Zhou, Jinkai Wang, Jixuan Chen, Kai Zhu, Kang Zhao, Keyu Yan, Lianghua Huang, Mengyang Feng, Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang, Tianxing Wang, Tianyi Gui, Tingyu Weng, Tong Shen, Wei Lin, Wei Wang, Wei Wang, Wenmeng Zhou, Wenteng Wang, Wenting Shen, Wenyuan Yu, Xianzhong Shi, Xiaoming Huang, Xin Xu, Yan Kou, Yangyu Lv, Yifei Li, Yijing Liu, Yiming Wang, Yingya Zhang, Yitong Huang, Yong Li, You Wu, Yu Liu, Yulin Pan, Yun Zheng, Yuntao Hong, Yupeng Shi, Yutong Feng, Zeyinzi Jiang, Zhen Han, Zhi-Fan Wu, and Ziyu Liu. Wan: Open and advanced large-scale video generative models, 2025. 2
- [33] William Peebles and Saining Xie. Scalable Diffusion Models with Transformers. In *ICCV*, pages 4172–4182, 2023. 1, 2
- [34] Haoning Wu, Erli Zhang, Liang Liao, Chaofeng Chen, Jingwen Hou, Annan Wang, Wenxiu Sun, Qiong Yan, and Weisi Lin. Exploring Video Quality Assessment on User Generated Contents from Aesthetic and Technical Perspectives. In *ICCV*, pages 20087–20097, 2023. 5
- [35] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaoohu Qie, and Mike Zheng Shou. Tune-A-Video: One-Shot Tuning of Image Diffusion Models for Text-to-Video Generation. In *ICCV*, pages 7589–7599, 2023. 1, 2
- [36] Xin Ma, Yaohui Wang, Gengyun Jia, Xinyuan Chen, Ziwei Liu, Yuan-Fang Li, Cunjian Chen, and Yu Qiao. Latte: Latent Diffusion Transformer for Video Generation. *arXiv*, 2024. 2
- [37] Yash Jain, Anshul Nasery, Vibhav Vineet, and Harkirat Behl. PEEKABOO: Interactive Video Generation via Masked-Diffusion. *arXiv*, 2023. 1
- [38] Danah Yatim, Rafail Fridman, Omer Bar-Tal, Yoni Kasten, and Tali Dekel. Space-Time Diffusion Features for Zero-Shot Text-Driven Motion Transfer. In *CVPR*, pages 8466–8476, 2024. 6
- [39] Yutao Cui, Xiaotong Zhao, Guozhen Zhang, Shengming Cao, Kai Ma, and Limin Wang. StableDrag: Stable Dragging for Point-based Image Editing. *arXiv*, 2024. 1, 2
- [40] Kaidong Zhang, Jingjing Fu, and Dong Liu. Flow-Guided Transformer for Video Inpainting. In *ECCV*, pages 74–90, 2022. 5
- [41] Zhixing Zhang, Bichen Wu, Xiaoyan Wang, Yaqiao Luo, Luxin Zhang, Yinan Zhao, Peter Vajda, Dimitris N. Metaxas, and Licheng Yu. AVID: Any-Length Video Inpainting with Diffusion Model. In *CVPR*, pages 7162–7172, 2024. 2
- [42] Shangchen Zhou, Chongyi Li, Kelvin C. K. Chan, and Chen Change Loy. ProPainter: Improving Propagation and Transformer for Video Inpainting. In *ICCV*, pages 10443–10452, 2023. 3, 5
- [43] Bojia Zi, Shihao Zhao, Xianbiao Qi, Jianan Wang, Yukai Shi, Qianyu Chen, Bin Liang, Kam-Fai Wong, and Lei Zhang. CoCoCo: Improving Text-Guided Video Inpainting for Better Consistency, Controllability and Compatibility. *arXiv*, 2024. 1, 2