MotionEditor: Editing Video Motion via Content-Aware Diffusion

Shuyuan Tu[1,2]   Qi Dai[3] *   Zhi-Qi Cheng[4]   Han Hu[3]   Xintong Han[5]   Zuxuan Wu[1,2] *   Yu-Gang Jiang[1,2]

[1]Shanghai Key Lab of Intell. Info. Processing, School of CS, Fudan University
[2]Shanghai Collaborative Innovation Center of Intelligent Visual Computing
[3]Microsoft Research Asia   [4]Carnegie Mellon University   [5]Huya Inc.
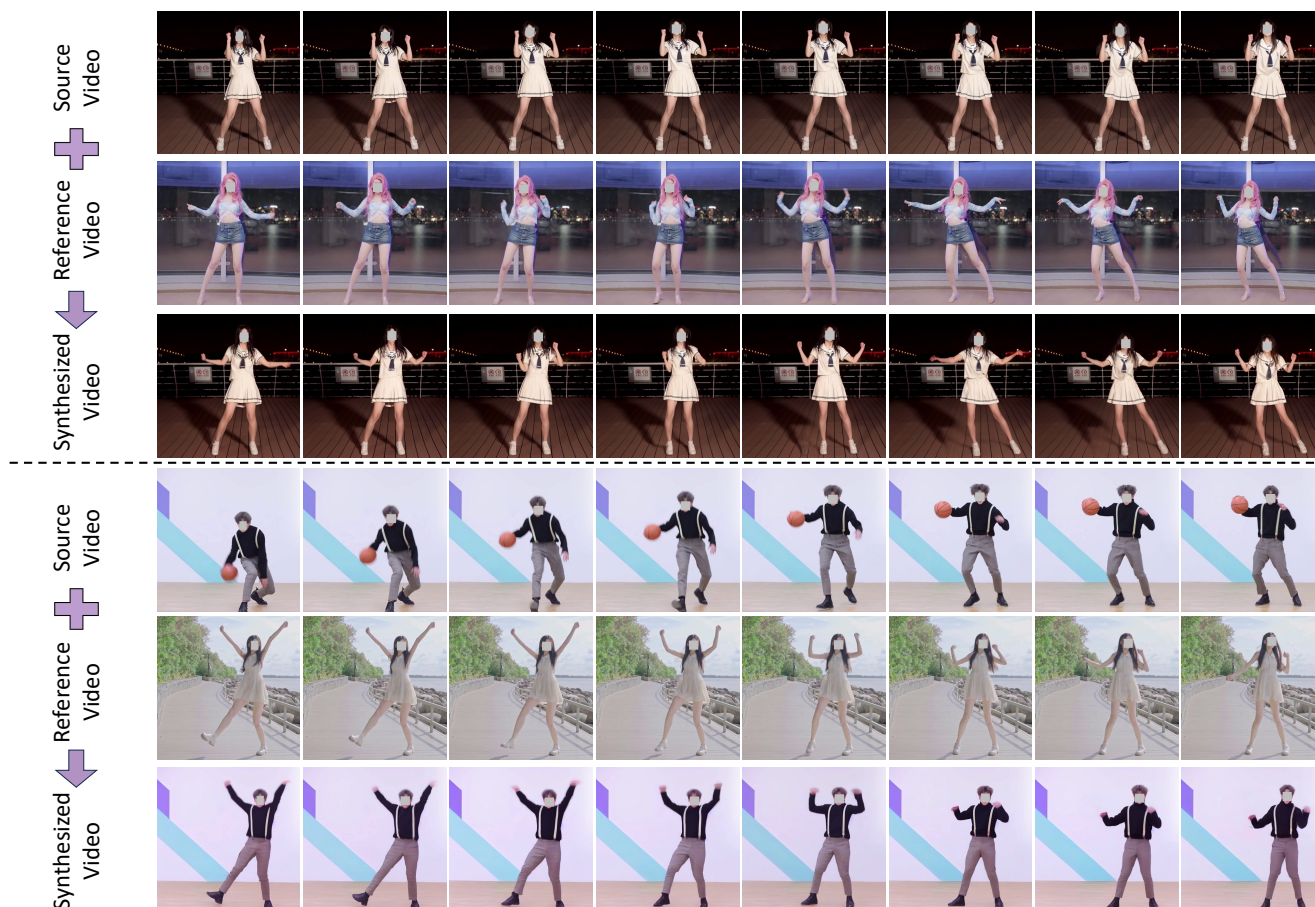https://francis-rings.github.io/MotionEditor

Figure 1. MotionEditor: A diffusion-based video editing method aimed at transferring motion from a reference to a source.

Abstract

Existing diffusion-based video editing models have made gorgeous advances for editing attributes of a source video over time but struggle to manipulate the motion information while preserving the original protagonist's appearance and background. To address this, we propose MotionEditor, the first diffusion model for video motion editing. MotionEditor incorporates a novel content-aware motion adapter into ControlNet to capture temporal motion correspondence. While ControlNet enables direct generation based on skeleton poses, it encounters challenges when modifying the source motion in the inverted noise due to contradictory signals between the noise (source) and the

---

*Corresponding authors.

condition (reference). Our adapter complements Control-Net by involving source content to transfer adapted control signals seamlessly. Further, we build up a two-branch architecture (a reconstruction branch and an editing branch) with a high-fidelity attention injection mechanism facilitating branch interaction. This mechanism enables the editing branch to query the key and value from the reconstruction branch in a decoupled manner, making the editing branch retain the original background and protagonist appearance. We also propose a skeleton alignment algorithm to address the discrepancies in pose size and position. Experiments demonstrate the promising motion editing ability of MotionEditor, both qualitatively and quantitatively. To the best of our knowledge, MotionEditor is the first to use diffusion models specifically for video motion editing, considering the origin dynamic background and camera movement.

1. Introduction

Diffusion models [5, 9, 11, 12, 22, 25, 31, 39, 40, 42, 46, 55, 56] have achieved remarkable success in image and video generation, which inspired plenty studies in video editing [3, 17, 19, 29, 47, 49, 50]. While significant progress has been made, existing diffusion models for video editing primarily focus on texture editing, such as attribute manipulation for protagonists, background editing, and style editing. The motion information, which is widely studied in video understanding [34, 38, 41, 48] and stands out as the most unique and distinct feature when compared to images, is mostly ignored. This raises the question: Can we manipulate the motion of a video in alignment with a reference video? In this paper, we attempt to explore a novel, higher-level, and more challenging video editing scenario—motion editing. Given a reference video and prompt, we aim to change the protagonist's motion in the source to match the reference video while preserving the original appearance.

To date in literature, researchers have explored human motion transfer [20] and pose-guided video generation [21]. The former focuses on animating still images based on reference skeletons, while the latter tries to generate pose-aligned videos without preserving a desired appearance. Motion editing differs in that it directly modifies the motion in video while preserving other information, such as dynamic per-frame background and camera movement.

Recent advancements in visual editing have mainly emerged through the use of diffusion models [2, 21, 24, 35, 52, 54]. For example, ControlNet [52] enables direct controllable generation conditioned on poses. However, they suffer from severe artifacts when trying to edit the motion according to another pose. We hypothesize the reason is that the control signals are only derived from the reference poses, which cannot properly accommodate the source content, thus resulting in a contradiction. Some methods also

lack the ability to preserve the appearance of the protagonist and background, as well as the temporal consistency. In this paper, we rethink the fundamental efficacy of ControlNet. We argue that it is essential to involve the source latent when generating the controlling signal for the editing task. With the help of source contents, the condition guidance can precisely perceive the entire context and structures, and adjust its distribution to prevent undesired distortion.

To this end, we propose MotionEditor, depicted in Fig. 2, to take one step forward in exploring video motion editing with diffusion models. MotionEditor requires one-shot learning on a source video to preserve the original texture feature. We then introduce a content-aware motion adapter appended to ControlNet for enhancing control capability and temporal modeling. The adapter consists of content-aware blocks and temporal blocks. In particular, content-aware blocks perform cross-attention to incorporate the source frame feature, which significantly boosts the quality of motion control.

At inference time, a skeleton alignment algorithm is devised to counter the size and position disparities between source and reference skeletons. We further propose an attention injection mechanism based on a two-branch architecture (reconstruction and editing branches) to preserve the source appearance of the protagonist and the background. Previous attention fusion strategies directly inject the attention map or key into the editing branch. The direct injection may result in confusion between the edited foreground and background. In some cases, it would bring noise to the editing branch, thereby resulting in the phenomenon of overlapping and shadow flickering. To avoid this, we propose to decouple the keys/values in the foreground and background using a segmentation mask. The keys/values in the editing branch are thus supplemented by those from the foreground and background in the reconstruction branch. As such, the editing branch is able to capture the background details and the geometric structure of the protagonist from the source.

In conclusion, our contributions are as follows: (1) We are the first to explore video diffusion models for motion editing, which is usually ignored by previous video editing works. (2) A novel content-aware motion adapter is proposed to enable the ControlNet to perform consistent and precise motion editing. (3) We propose a high-fidelity attention injection mechanism that preserves the background information and the geometric structure of the protagonist from the source. The mechanism is only active during inference, making it training-free. (4) We conduct experiments on in-the-wild videos, where the results show the superiority of our method compared with the state-of-the-art.

2. Related Work

Diffusion for Image Editing Image generation has been remarkably improved with diffusion models [31], surpass-

ing previous GAN models [7, 8, 27, 28, 44] in quality and editing capabilities. Building on the pioneering work by [31], researchers have made progress in this direction. Meng et al. propose SDEdit, the first approach that enables image editing via inversion and reversion. This allows more precise control over the edits compared with previous GAN-based approaches. Prompt-to-Prompt [9] and Plug-and-Play [42] introduce techniques to generate more coherent hybrid images, tackling the inconsistent noise. Meanwhile, UniTune [43] and Imagic [16] focus on personalized fine-tuning of diffusion models. In short, these approaches enable edits that remain truer to the original image.

To generate diverse contents, researchers have proposed several controllable diffusion models [6, 13, 26, 32, 33, 52]. Zhang et al. introduce ControlNet, which enables Stable Diffusion [31] to embrace multiple controllable conditions for text-to-image synthesis. Cao et al. propose MasaCtrl, which converts existing self-attention into mutual self-attention to achieve non-rigid image synthesis. Overall, these approaches have advanced diffusion-based image editing by handling inconsistent noise, fine-tuning strategies, and inversion-based editing. However, it remains an open challenge to enable precise semantic edits while maintaining perfect fidelity and coherence.

Diffusion Models for Video Editing and Generation Video editing is complex, which demands temporal consistency in editing. Most methods [47, 49] use existing T2I diffusion models with an additional temporal modeling module. Tune-a-Video [47], SimDA [51], and Text2Video-Zero [17] inflate 2D diffusion models to 3D models. FateZero [29] and Vid2Vid-zero [45] use mutual attention to ensure consistency of geometry and color. On the other hand, Text2LIVE [1] and StableVideo [4] decompose video semantics by leveraging layered neural atlas [15]. However, these models are limited to low-level attribute editing and cannot edit complex information such as motion.

Recently, pose-driven video generation has become popular. Follow-Your-Pose [21] extracts coarse temporal motion information by implementing an adapter on the text-to-image diffusion model. ControlVideo [54] introduces full cross-frame attention by extending ControlNet to videos. However, these methods focus on video generation rather than motion editing, resulting in distortion of the protagonist and background. Unlike these models, our MotionEditor aims to perform motion editing while preserving the appearance of the original video.

Human Motion Transfer This task aims to transfer the motion from a video to a target image, enabling animation of the target. Previous GAN-based methods [14, 18, 20, 36, 37] have tackled this task but struggle with complex motions and backgrounds. LWG [20] uses 3D pose projection for motion transfer but cannot model internal motions well. FOMM [36] approximates motion transfer via

boundary key points. MRAA [37] exploits regional features to capture part motion, yet the performance is limited on complex scenes. To address these limitations, we propose MotionEditor for high-quality motion editing on complex motions and backgrounds. In contrast to prior work, it leverages diffusion models capable of *generating consistent details* even for intricate motions and backgrounds.

3. Preliminaries

Diffusion models [11, 25, 39] have recently shown promising results for synthesizing high-quality and diverse contents using iterative denoising operations. They consist of a forward diffusion process and a reverse denoising process. During the forward process, models equipped with a predefined noise schedule α_t add random noise to the source sample $\boldsymbol{x}_0$ at time step t for obtaining a noised sample $\boldsymbol{x}_t$:

$$q(\boldsymbol{x}_{1:T}) = q(\boldsymbol{x}_0) \prod_{t=1}^{T} q(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}),$$

$$q(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}) = \mathcal{N}(\boldsymbol{x}_t; \sqrt{\alpha_t}\boldsymbol{x}_{t-1}, (1-\alpha_t)\mathbf{I}). \quad (1)$$

The original input $\boldsymbol{x}_0$ is inverted into Gaussian noise $\boldsymbol{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$ after T forward steps. The reverse process attempts to predict a cleaner sample $\boldsymbol{x}_{t-1}$ based on $\boldsymbol{x}_t$ by removing noise. The process is depicted as:

$$p_\theta(\boldsymbol{x}_{t-1}|\boldsymbol{x}_t) = \mathcal{N}(\boldsymbol{x}_{t-1}; \boldsymbol{\mu}_\theta(\boldsymbol{x}_t, t), \boldsymbol{\sigma}_t^2 \mathbf{I}), \quad (2)$$

where $\boldsymbol{\mu}_\theta(\boldsymbol{x}_t, t)$ and $\boldsymbol{\sigma}_t^2$ indicate the mean and variance of the sample at the current time step. Only the mean is correlated with the time step and noise, while the variance is constant. The denoising network $\boldsymbol{\varepsilon}_\theta(\boldsymbol{x}_t, t)$ aims to predict the noise $\boldsymbol{\varepsilon}$ by training with a simplified mean squared error:

$$\mathcal{L}_{simple} = \mathbb{E}_{\boldsymbol{x}_0, \boldsymbol{\varepsilon}, t}(\|\boldsymbol{\varepsilon} - \boldsymbol{\varepsilon}_\theta(\boldsymbol{x}_t, t)\|^2). \quad (3)$$

Once the model is trained, we feed $\boldsymbol{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$ to the diffusion model and iteratively perform DDIM sampling for predicting cleaner $\boldsymbol{x}_{t-1}$ from the noise sample $\boldsymbol{x}_t$ of a previous time step. The process is shown as follows:

$$\boldsymbol{x}_{t-1} = \sqrt{\alpha_{t-1}} \frac{\boldsymbol{x}_t - \sqrt{1-\alpha_t}\boldsymbol{\varepsilon}_\theta(\boldsymbol{x}_t, t)}{\sqrt{\alpha_t}} + \sqrt{1-\alpha_{t-1}}\boldsymbol{\varepsilon}_\theta(\boldsymbol{x}_t, t). \quad (4)$$

One can also inject a text prompt $\boldsymbol{p}$ into the prediction model $\boldsymbol{\varepsilon}_\theta(\boldsymbol{x}_t, t, \boldsymbol{p})$ as a condition, where the diffusion model can perform T2I synthesis. Recent work [31] introduces an encoder $\mathcal{E}$ to compress images $\boldsymbol{x}$ into a latent space $\boldsymbol{z} = \mathcal{E}(\boldsymbol{x})$, and a decoder $\mathcal{D}$ to transfer the latent embedding back to the pixel space. In this way, the diffusion process is performed in the latent space.

4. Method

4.1. Architecture Overview

As illustrated in Fig. 2, MotionEditor is based on the commonly used T2I diffusion model (*i.e.* LDM [31]) and Con-
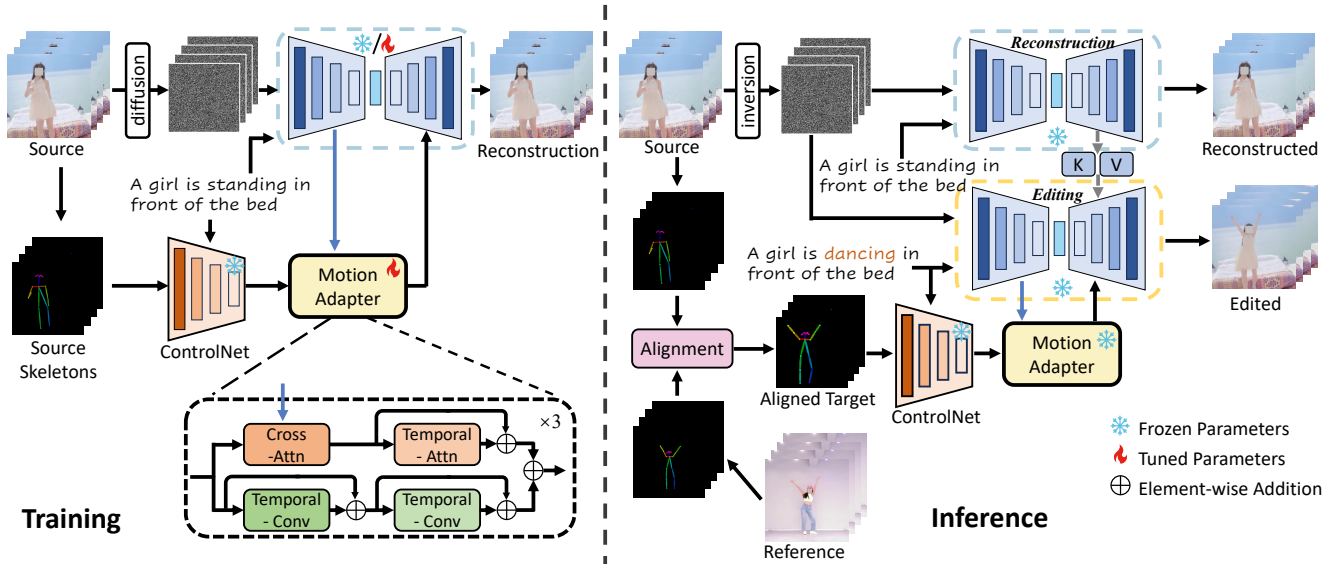
Figure 2. Architecture overview of MotionEditor. In training, only the motion adapter and temporal attention in U-Net are trainable. In inference, we first align the source and reference skeletons through resizing and translation. We then build a two-branch framework: one for reconstruction and the other for editing. Motion adapter enhances the motion guidance of ControlNet by leveraging the information from the source latent. We also inject the key/value in the reconstruction branch into the editing branch to preserve the source appearance.

trolNet [52]. We first inflate the spatial transformer in the U-Net of the LDM to a 3D transformer by appending a temporal self-attention layer. We also propose Consistent-Sparse attention, described in Sec. 4.3, to replace the spatial self-attention. To achieve precise motion manipulation and temporal consistency, we design a content-aware motion adapter, operating on features from the U-Net and conditional pose information from the ControlNet. Following [47], we perform one-shot training to compute the weights for the temporal attention module and the motion adapter to reconstruct the input source video.

During inference, given a reference target video X_{rf}, MotionEditor aims to transfer the motion of X_{rf} to the source while preserving the appearance information of the source. To this end, we first develop a skeleton alignment algorithm to narrow the gap between source skeleton S_{sr} and reference skeleton S_{rf} by considering the position and size, and produce a refined target skeleton $\bar{S}_{tg}$. We then employ the DDIM inversion [39] on pixel values of the video to produce a latent noise that serves as the starting point for sampling. More importantly, we introduce a high-fidelity attention injection module exploring a carefully designed two-branch network. One is dedicated to reconstruction, while the other one focuses on editing. More specifically, the editing branch takes features from the refined target skeleton as inputs, transferring the motion information from the reference to the source. Meanwhile, critical appearance information encoded in the reconstruction branch is further injected into the editing branch so that the appearance and the background are preserved. Below, we introduce the proposed components in detail.

4.2. Content-Aware Motion Adapter

Our goal is to manipulate the body motion in videos under the guidance of pose signals. While ControlNet [52] enables direct controllable generation based on conditions, it has difficulties modifying the source motion from the inverted noise. The motion signals injected by the ControlNet could conflict with the source motion, thereby resulting in pronounced ghosting and blurring effects or even the loss of controlling ability. Furthermore, the model is derived from an image, lacking the ability to generate temporal consistent contents. Therefore, we propose a temporal content-aware motion adapter that enhances motion guidance as well as facilitates temporal consistency, as shown in Fig. 2.

Our motion adapter takes as input the feature output by ControlNet, which has been observed to achieve promising spatial control. Instead of inserting temporal layers into ControlNet, we leave it as is to prevent prejudicing its inherent modeling capability. The adapter consists of two parallel paths corresponding to different perception granularity. One is the global modeling path, including a content-aware cross-attention block and a temporal attention block. The other is the local modeling path, which uses two temporal convolution blocks to capture local motion features. Specifically, our cross-attention involves the latent feature from the U-Net to model the pose feature, where the query comes from the pose feature m_i, and the key/value is from the corresponding frame latent z_i produced by the U-Net:

$$Q = W_c^Q m_i, K = W_c^K z_i, V = W_c^V z_i, \qquad (5)$$

where W_c^Q, W_c^K and W_c^V are projection matrices. Our cross-attention enables the motion adapter to concentrate on
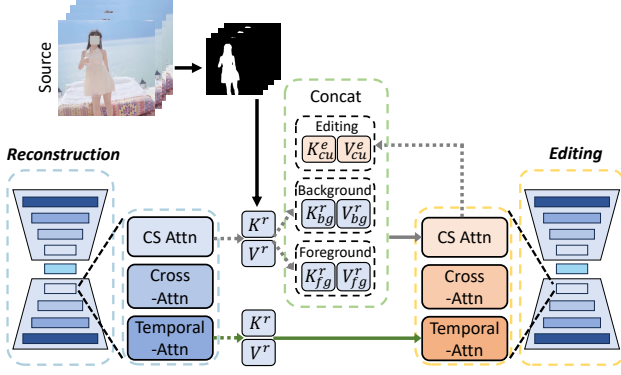
Figure 3. Illustration of high-fidelity attention injection during inference. We leverage the source foreground masks to guide the decoupling of key/value in the Consistent-Sparse Attention.

correlated motion clues in the video latent space, which significantly enhances the control capability. By building up a bridge between them, the model can thereby manipulate the source motion seamlessly, without contradiction.

4.3. High-Fidelity Attention Injection

Although our motion adapter can accurately capture body poses, it may undesirably alter the appearance of the protagonist and the background. Consequently, we propose a high-fidelity attention injection from the reconstruction branch to the editing branch, which preserves the details of the subject and background in the synthesized video. While previous attention fusion paradigms [2, 29, 45] have used attention maps or keys/values for editing, they suffer from severe quality degradation in ambiguous regions, *i.e.* motion areas, due to context confusion. To solve the issue, we decouple the keys and values into foreground and background through semantic masks. By injecting separated keys and values into the editing branch from the reconstruction branch, we can reduce the confusion hindering the editing. The pipeline is shown in Fig. 3.

Before introducing the injection, we first present the details of attention blocks in the model. Each attention block in U-Net consists of our designed Consistent-Sparse Attention (CS Attention), Cross Attention, and Temporal Attention. The Consistent-Sparse Attention, as a sparse causal attention, replaces the spatial attention in the original U-Net. It aims to perform spatiotemporal modeling with little additional computational overhead. Specifically, taking the reconstruction branch as an example, the query in CS Attention is derived from the current frame z_i^r, while the key/value is obtained from the preceding and current frames z_{i-1}^r, z_i^r. This design can improve the frame consistency:

$$\boldsymbol{Q}^r = \boldsymbol{W}^Q \boldsymbol{z}_i^r, \boldsymbol{K}^r = \boldsymbol{W}^K[\boldsymbol{z}_{i-1}^r, \boldsymbol{z}_i^r], \boldsymbol{V}^r = \boldsymbol{W}^V[\boldsymbol{z}_{i-1}^r, \boldsymbol{z}_i^r],$$
(6)

where $[\cdot]$ refers to concatenation. $\boldsymbol{W}^Q$, $\boldsymbol{W}^K$ and $\boldsymbol{W}^V$ are projection matrices. It is worth noting that we do not employ the sparse attention in [47], in which the key/value is

from the first and preceding frames z_0, z_{i-1}. We find that it may force the synthesized motion to favor the first frame excessively, resulting in flickering.

We now present the injection of keys and values from the reconstruction branch to those of the editing branch, operating on both Consistent-Sparse Attention (CS Attention) and Temporal Attention. The injection is only active in the decoder of U-Net. For CS Attention, we leverage a source foreground mask $\boldsymbol{M}$, obtained from an off-the-shelf segmentation model, to decouple the foreground and background information. Given the key $\boldsymbol{K}^r$ and value $\boldsymbol{V}^r$ in the reconstruction branch, we separate them into the foreground ($\boldsymbol{K}_{fg}^r$ and $\boldsymbol{V}_{fg}^r$) and background ($\boldsymbol{K}_{bg}^r$ and $\boldsymbol{V}_{bg}^r$),

$$\begin{aligned} \boldsymbol{K}_{fg}^r &= \boldsymbol{K}^r \odot \boldsymbol{M}, \ \ \boldsymbol{V}_{fg}^r = \boldsymbol{V}^r \odot \boldsymbol{M}, \\ \boldsymbol{K}_{bg}^r &= \boldsymbol{K}^r \odot (\boldsymbol{1} - \boldsymbol{M}), \ \ \boldsymbol{V}_{bg}^r = \boldsymbol{V}^r \odot (\boldsymbol{1} - \boldsymbol{M}). \end{aligned}$$
(7)

The decoupling operation introduces an explicit distinction between background and foreground. It encourages the model to focus more on the individual appearance instead of mixing both, ensuring the high fidelity of the subject and background. Notably, simply replacing the key/value in the editing branch with the above ones would cause a large number of abrupt motion changes, as the model is significantly influenced by the source motion. Instead, we combine them to maintain the target motion precisely. Therefore, in the editing branch, the key and value of CS Attention are updated by the injected key $\boldsymbol{K}_{inj}$ and value $\boldsymbol{V}_{inj}$:

$$\begin{aligned} \boldsymbol{K}^e &= \boldsymbol{W}^K[\boldsymbol{z}_{i-1}^e, \boldsymbol{z}_i^e] = [\boldsymbol{W}^K \boldsymbol{z}_{i-1}^e, \boldsymbol{W}^K \boldsymbol{z}_i^e] \coloneqq [\boldsymbol{K}_{pr}^e, \boldsymbol{K}_{cu}^e], \\ \boldsymbol{V}^e &= \boldsymbol{W}^V[\boldsymbol{z}_{i-1}^e, \boldsymbol{z}_i^e] = [\boldsymbol{W}^V \boldsymbol{z}_{i-1}^e, \boldsymbol{W}^V \boldsymbol{z}_i^e] \coloneqq [\boldsymbol{V}_{pr}^e, \boldsymbol{V}_{cu}^e], \\ \boldsymbol{K}_{inj} &= [\boldsymbol{K}_{fg}^r, \boldsymbol{K}_{bg}^r, \boldsymbol{K}_{cu}^e], \\ \boldsymbol{V}_{inj} &= [\boldsymbol{V}_{fg}^r, \boldsymbol{V}_{bg}^r, \boldsymbol{V}_{cu}^e], \end{aligned}$$
(8)

where z_i^e, z_{i-1}^e indicate the current and preceding frames in the editing branch. $\boldsymbol{K}^e, \boldsymbol{V}^e$ are the original key and value in CS attention of the editing branch. $\boldsymbol{K}_{cu}^e, \boldsymbol{V}_{cu}^e$ are the key and value from the current frame.

The injection in temporal attention is much simpler than in CS Attention since it already performs the local operation concerning the spatial region. We directly inject the $\boldsymbol{K}^r, \boldsymbol{V}^r$ of reconstruction branch into the editing branch.

4.4. Skeleton Signal Alignment

Given the source and reference videos, there always exists a gap between the source and the target protagonists due to different sizes and coordinated positions. The discrepancy may affect the performance of editing. Therefore, we propose an alignment algorithm for addressing this issue.

Our algorithm contains two steps, namely a resizing operation, and a translation operation. Concretely, given the source video $\boldsymbol{X}_{sr}$ and a reference video $\boldsymbol{X}_{rf}$, we first extract the source skeleton $\boldsymbol{S}_{sr}$ and a foreground mask $\boldsymbol{M}_{sr}$, as well as the reference skeleton $\boldsymbol{S}_{rf}$ and its mask $\boldsymbol{M}_{rf}$,

Figure 4. Motion editing results of our MotionEditor. More examples can be found in the supplementary material and anonymous website.

using off-the-shelf models. We then perform edge detection on the masks to obtain rectangular outlines of the foreground. Based on the area of two rectangular outlines, we scale S_{rf} to the same size as the source. Regarding the foreground position, we calculate the average coordinate of the foreground pixels in each mask, which denotes the center of the protagonist. An offset vector is computed by calculating the difference between the two centers. With this vector, we obtain an affine matrix for the translation operation, which is further applied to the resized reference skeleton. Finally, the target skeleton $\bar{S}_{tg}$ is generated. The details of alignment are depicted in the supplementary material.

5. Experiments

5.1. Implementation Details

Our proposed MotionEditor is based on the Latent Diffusion Model [31] (Stable Diffusion). The person segmentation and skeleton estimation methods refer to SAM and OpenPose. We evaluate our model on YouTube videos and videos from the TaichiHD [36] dataset, in which each video includes a protagonist of at least 70 frames. The resolution of frames is unified to 512×512. We perform one-shot

learning for the motion adapter for 300 steps with a constant learning rate of 3×10^{-5}. We employ DDIM inversion [39] and null-text optimization [23] with classifier-free guidance [10] during inference. Due to the usage of DDIM inversion and null-text optimization, MotionEditor needs 10 minutes to perform motion editing for each video on a single NVIDIA A100 GPU.

5.2. Motion Editing Results

We validate the motion editing superiority of our proposed MotionEditor extensively. Here, we demonstrate several cases in Figure 4 and the rest of the cases are in the supplementary materials. We can see that our MotionEditor can accomplish a wide range of motion editing while simultaneously preserving the original protagonist's appearance and background information.

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

Competitors. We compare our MotionEditor against recent approaches to validate the superiority of our model. The competitors are depicted as follows: (1) Models based on GANs for human motion transfer, including LWG [20] that disentangles the pose and shape, and MRAA [37]

Figure 5. Qualitative comparison between our MotionEditor and other state-of-the-art video editing models. Source prompt: "a girl in a black dress is dancing." Target prompt: "a girl in a black dress is practicing tai chi." Our method exhibits accurate motion editing and appearance preservation.

Table 1. Quantitative comparisons on 20 in-the-wild cases. L-S, L-N, and L-T indicate LPIPS-S, LPIPS-N, LPIPS-T respectively.

Method	CLIP (↑)	L-S (↓)	L-N (↓)	L-T (↓)
LWG [20]	25.35	0.431	0.194	0.203
MRAA [37]	26.80	0.462	0.269	0.353
Tune-A-Video [47]	27.71	0.345	0.169	0.157
Follow-Your-Pose [21]	26.55	0.337	0.144	0.183
ControlVideo [54]	26.87	0.428	0.228	0.311
MasaCtrl [2]	27.14	0.372	0.236	0.177
FateZero [29]	28.07	0.308	0.176	0.124
MotionEditor	**28.86**	**0.273**	**0.124**	**0.082**

that learns semantic object parts. (2) Tune-A-Video [47], which inflates Stable Diffusion in a one-shot learning manner. (3) Follow-Your-Pose [21], which proposes a spatial pose adapter for pose-guided video generation. (4) ControlVideo [54], which designs fully cross-frame attention, attempting to stitch video frames into one large image. (5) MasaCtrl [2], which conducts mask-guided mutual self-attention fusion, in which the masks are derived from text cross attention. (6) FateZero [29], which designs inversion attention fusion for retaining source structure. It is worth noting that Tune-A-Video and Masactrl are equipped with a conditional T2I model (*i.e.* ControlNet) for controllable video editing. We also combine FateZero with ControlNet to enable pose-guided editing. More details on implementation are depicted in the supplementary material.

Qualitative results. We conduct a qualitative comparison of our MotionEditor against several competitors in Fig. 5. The results of LWG [20] and MRAA [37] are provided in the supplementary material. By analyzing the results, we have the following observations: (1) Tune-A-Video [47], ControlVideo [54], Masactrl [2] and FateZero [29] demonstrate capabilities in editing motion to a certain extent. Nev-

ertheless, their edited videos exhibit a considerable degree of ghosting, manifested as overlapping frames of the protagonist's heads and legs. (2) Follow-Your-Pose [21] fails to perform motion editing and preserve the source appearance. The ghosting effect also exists in the result. The plausible reason is that it has difficulties handling the motion conflict between inverted noise and input reference pose. (3) Tune-A-Video [47] and ControlVideo [54] both have dramatic appearance changes, *e.g.*, hairstyle, texture of clothes and background. Due to the constraint of a single-branch architecture, they are unable to interact with the original video features when additional pose conditions are introduced. This leads to a gradual loss of the source appearance along with the increasing denoising steps. (4) Masactrl [2] generates an excessive amount of blurry noise. The possible reason is that the masks generated by cross-attention maps are unreliable and inconsistent across time, thus introducing blurring noise to the result. (5) The edited result of FateZero [29] shows overlapping frames of the protagonist's legs. It demonstrates that the attention fusion strategy in FateZero may not be suitable for motion editing. (6) Finally, our MotionEditor can effectively perform motion editing while preserving the original background and appearance compared with previous approaches, highlighting its great potential.

Quantitative results. To our best knowledge, there is still no widely recognized metric for evaluating the performance of video editing. In this paper, we conduct quantitative comparisons against previous methods through several perceptual metrics [53] and a user study on edited videos. The detailed metrics are as follows: (1) CLIP score [30]: Target textual faithfulness. (2) LPIPS-S: Learned Perceptual Image Patch Similarity (LPIPS) [53] between edited frames and source frames. (3) LPIPS-N: LPIPS between edited neighboring frames. (4) LPIPS-T: We split a long video into two segments. The first segment is used for the source, while the second segment is for the reference. We compute LPIPS between the edited video and the second segment. The results are shown in Table 1. We observe that MotionEditor surpasses the competitors by a large margin.

We also conduct a user study to evaluate the human preference between our method and the competitors. For each

Table 2. User preference ratio of MotionEditor when comparing with each method. Higher indicates the users prefer more to our MotionEditor. M-A, A-A and T-A indicate motion alignment, appearance alignment, and textual alignment, respectively.

Method	M-A	A-A	T-A
LWG [20]	91.9%	95.7%	90.0%
MRAA [37]	94.8%	98.4%	92.6%
Tune-A-Video [47]	87.6%	90.8%	79.2%
Follow-Your-Pose [21]	96.3%	85.0%	84.6%
ControlVideo [54]	94.1%	98.8%	86.0%
MasaCtrl [2]	89.4%	94.6%	87.8%
FateZero [29]	78.9%	74.5%	74.7%

case, participants are first presented with the source and target videos, as well as the prompts. We then show two motion-edited videos; one is generated by our method and the other is from a competitor, in random order. Participants are asked to answer the following questions: "which one has better motion alignment with reference", "which one has better appearance alignment with source", and "which one has better content alignment with the prompt." The total number of cases is 20, and the participants are mainly university students. Table 2 shows that our method outperforms other methods in terms of subjective evaluation.

5.4. Ablation Study

To validate the importance of the core components in MotionEditor, we conduct an ablation study. The results are illustrated in Fig. 6. It is worth noting that we replace our proposed CS Attention with the previous Sparse Attention [47] in (c). The results in row (c) indicate that Sparse Attention attempts to force the frames to be aligned with the first frame, resulting in unreliable motion. Rows (d) and (e) both fail to accomplish motion editing and background preservation. It demonstrates that the original ControlNet has weak constraints on the motion without additional content-aware modeling. It also forces the model to retain background information. In row (f), model *w/o* high-fidelity attention injection loses the original background details as the road sign behind the girl has disappeared. It validates that our proposed mechanism can promote the model to preserve the source background. Model *w/o* skeleton alignment suffers from appearance change due to the misalignment, as in row (g). The misalignment of the skeletons may potentially introduce unexpected noise to the content latent, thus destroying source data distribution. The ablation results show that our core components certainly contribute to the promising motion editing capability of MotionEditor.

6. Limitations

Temporal inconsistencies sometimes occur in the foreground subjects. Our model only conducts one-shot learning on a single video without any large-scale video pertaining. We will tackle the above issues in the future.
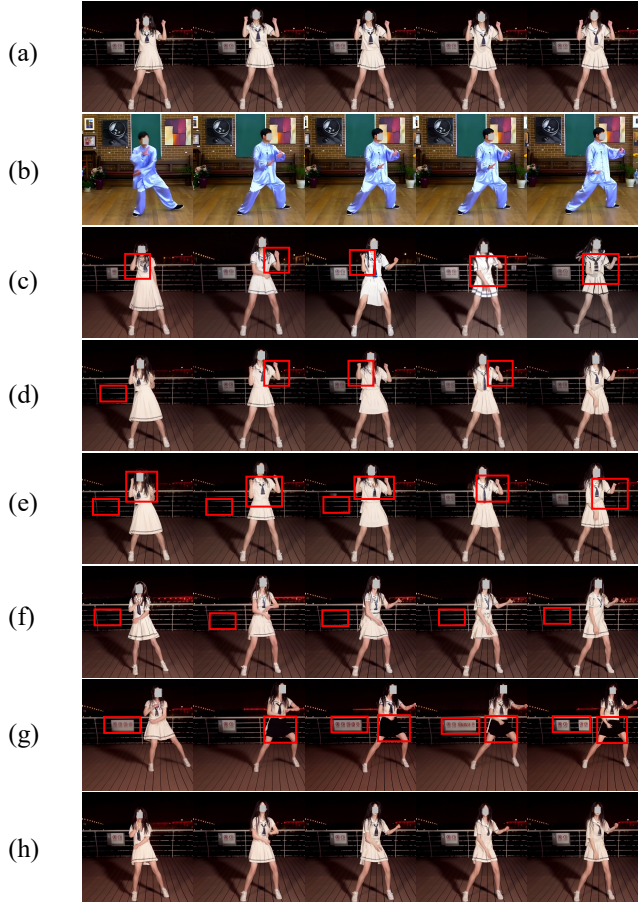

Figure 6. Ablations on core components of MotionEditor. Rows in the figure are: (a) source, (b) reference, (c) *w/o* CS Attention, (d) *w/o* cross attention in motion adapter, (e) *w/o* motion adapter, (f) *w/o* high-fidelity attention injection, (g) *w/o* skeleton alignment, and (h) MotionEditor. Source prompt: "A girl is dancing." Target prompt: "A girl is practicing tai chi."

7. Conclusion

In this paper, we proposed MotionEditor for tackling video motion editing challenges, which is rated as high-level video editing compared with previous video attribute editing. To enhance the motion controllability, a content-aware motion adapter was designed to build up a relationship with the source content, enabling seamless motion editing as well as temporal modeling. We further proposed a high-fidelity attention injection for preserving the source appearance of the background and protagonist. To alleviate the misalignment problem of skeleton signals, we presented a simple yet effective skeleton alignment to normalize the target skeletons. To our knowledge, MotionEditor is the first diffusion model to explore the video motion editing task, encouraging more studies in this challenging scenario.

References

[1] Omer Bar-Tal, Dolev Ofri-Amar, Rafail Fridman, Yoni Kasten, and Tali Dekel. Text2live: Text-driven layered image and video editing. In *ECCV*, 2022. 3

[2] Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiaohu Qie, and Yinqiang Zheng. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *ICCV*, 2023. 2, 3, 5, 7, 8

[3] Duygu Ceylan, Chun-Hao P Huang, and Niloy J Mitra. Pix2video: Video editing using image diffusion. In *ICCV*, 2023. 2

[4] Wenhao Chai, Xun Guo, Gaoang Wang, and Yan Lu. Stablevideo: Text-driven consistency-aware diffusion video editing. In *ICCV*, 2023. 3

[5] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, 2021. 2

[6] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022. 3

[7] Rinon Gal, Or Patashnik, Haggai Maron, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. Stylegan-nada: Clip-guided domain adaptation of image generators. *ACM Transactions on Graphics*, 41(4):1–13, 2022. 3

[8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 3

[9] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 2, 3

[10] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 6

[11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 2, 3

[12] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *JMLR*, 2022. 2

[13] Lianghua Huang, Di Chen, Yu Liu, Yujun Shen, Deli Zhao, and Jingren Zhou. Composer: Creative and controllable image synthesis with composable conditions. *arXiv preprint arXiv:2302.09778*, 2023. 3

[14] Zhichao Huang, Xintong Han, Jia Xu, and Tong Zhang. Few-shot human motion transfer by personalized geometry and texture modeling. In *CVPR*, 2021. 3

[15] Yoni Kasten, Dolev Ofri, Oliver Wang, and Tali Dekel. Layered neural atlases for consistent video editing. *ACM Transactions on Graphics*, 40(6):1–12, 2021. 3

[16] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. In *CVPR*, 2023. 3

[17] Levon Khachatryan, Andranik Movsisyan, Vahram Tadevosyan, Roberto Henschel, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Text2video-zero: Text-to-image diffusion models are zero-shot video generators. In *ICCV*, 2023. 2, 3

[18] Hongyu Liu, Xintong Han, Chengbin Jin, Lihui Qian, Huawei Wei, Zhe Lin, Faqiang Wang, Haoye Dong, Yibing Song, Jia Xu, et al. Human motionformer: Transferring human motions with vision transformers. *ICLR*, 2023. 3

[19] Shaoteng Liu, Yuechen Zhang, Wenbo Li, Zhe Lin, and Jiaya Jia. Video-p2p: Video editing with cross-attention control. *arXiv preprint arXiv:2303.04761*, 2023. 2

[20] Wen Liu, Zhixin Piao, Jie Min, Wenhan Luo, Lin Ma, and Shenghua Gao. Liquid warping gan: A unified framework for human motion imitation, appearance transfer and novel view synthesis. In *ICCV*, 2019. 2, 3, 6, 7, 8

[21] Yue Ma, Yingqing He, Xiaodong Cun, Xintao Wang, Ying Shan, Xiu Li, and Qifeng Chen. Follow your pose: Pose-guided text-to-video generation using pose-free videos. *arXiv preprint arXiv:2304.01186*, 2023. 2, 3, 7, 8

[22] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *ICLR*, 2021. 2, 3

[23] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. In *CVPR*, 2023. 6

[24] Chong Mou, Xintao Wang, Jiechong Song, Ying Shan, and Jian Zhang. Dragondiffusion: Enabling drag-style manipulation on diffusion models. *arXiv preprint arXiv:2307.02421*, 2023. 2

[25] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *ICML*, 2021. 2, 3

[26] Yaniv Nikankin, Niv Haim, and Michal Irani. Sinfusion: Training diffusion models on a single image or video. *arXiv preprint arXiv:2211.11743*, 2022. 3

[27] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, pages 2337–2346, 2019. 3

[28] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *ICCV*, 2021. 3

[29] Chenyang Qi, Xiaodong Cun, Yong Zhang, Chenyang Lei, Xintao Wang, Ying Shan, and Qifeng Chen. Fatezero: Fusing attentions for zero-shot text-based video editing. In *ICCV*, 2023. 2, 3, 5, 7, 8

[30] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 7

[31] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2, 3, 6

[32] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *CVPR*, 2023. 3

[33] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *SIGGRAPH*, 2022. 3

[34] Baifeng Shi, Qi Dai, Yadong Mu, and Jingdong Wang. Weakly-supervised action localization by generative attention modeling. In *CVPR*, 2020. 2

[35] Yujun Shi, Chuhui Xue, Jiachun Pan, Wenqing Zhang, Vincent YF Tan, and Song Bai. Dragdiffusion: Harnessing diffusion models for interactive point-based image editing. *arXiv preprint arXiv:2306.14435*, 2023. 2

[36] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. *NeurIPS*, 2019. 3, 6

[37] Aliaksandr Siarohin, Oliver J Woodford, Jian Ren, Menglei Chai, and Sergey Tulyakov. Motion representations for articulated animation. In *CVPR*, 2021. 3, 6, 7, 8

[38] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. *NeurIPS*, 2014. 2

[39] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021. 2, 3, 4, 6

[40] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2021. 2

[41] Shuyuan Tu, Qi Dai, Zuxuan Wu, Zhi-Qi Cheng, Han Hu, and Yu-Gang Jiang. Implicit temporal modeling with learnable alignment for video recognition. In *ICCV*, 2023. 2

[42] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. In *CVPR*, 2023. 2, 3

[43] Dani Valevski, Matan Kalman, Yossi Matias, and Yaniv Leviathan. Unitune: Text-driven image editing by fine tuning an image generation model on a single image. *arXiv preprint arXiv:2210.09477*, 2022. 3

[44] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *CVPR*, 2018. 3

[45] Wen Wang, Kangyang Xie, Zide Liu, Hao Chen, Yue Cao, Xinlong Wang, and Chunhua Shen. Zero-shot video editing using off-the-shelf image diffusion models. *arXiv preprint arXiv:2303.17599*, 2023. 3, 5

[46] Yanhui Wang, Jianmin Bao, Wenming Weng, Ruoyu Feng, Dacheng Yin, Tao Yang, Jingxu Zhang, Qi Dai, Zhiyuan Zhao, Chunyu Wang, Kai Qiu, Yuhui Yuan, Chuanxin Tang, Xiaoyan Sun, Chong Luo, and Baining Guo. Microcinema: A divide-and-conquer approach for text-to-video generation. In *CVPR*, 2024. 2

[47] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *ICCV*, 2023. 2, 3, 4, 5, 7, 8

[48] Zhen Xing, Qi Dai, Han Hu, Jingjing Chen, Zuxuan Wu, and Yu-Gang Jiang. Svformer: Semi-supervised video transformer for action recognition. In *CVPR*, 2023. 2

[49] Zhen Xing, Qi Dai, Zihao Zhang, Hui Zhang, Han Hu, Zuxuan Wu, and Yu-Gang Jiang. Vidiff: Translating videos via multi-modal instructions with diffusion models. *arXiv preprint arXiv:2311.18837*, 2023. 2, 3

[50] Zhen Xing, Qijun Feng, Haoran Chen, Qi Dai, Han Hu, Hang Xu, Zuxuan Wu, and Yu-Gang Jiang. A survey on video diffusion models. *arXiv preprint arXiv:2310.10647*, 2023. 2

[51] Zhen Xing, Qi Dai, Han Hu, Zuxuan Wu, and Yu-Gang Jiang. Simda: Simple diffusion adapter for efficient video generation. In *CVPR*, 2024. 3

[52] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, 2023. 2, 3, 4

[53] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 7

[54] Yabo Zhang, Yuxiang Wei, Dongsheng Jiang, Xiaopeng Zhang, Wangmeng Zuo, and Qi Tian. Controlvideo: Training-free controllable text-to-video generation. *arXiv preprint arXiv:2305.13077*, 2023. 2, 3, 7, 8

[55] Zhongwei Zhang, Fuchen Long, Yingwei Pan, Zhaofan Qiu, Ting Yao, Yang Cao, and Tao Mei. Trip: Temporal residual learning with image noise prior for image-to-video diffusion models. In *CVPR*, 2024. 2

[56] Rui Zhu, Yingwei Pan, Yehao Li, Ting Yao, Zhenglong Sun, Tao Mei, and Chang Wen Chen. Sd-dit: Unleashing the power of self-supervised discrimination in diffusion transformer. In *CVPR*, 2024. 2