

MOFA-Video: Controllable Image Animation via Generative Motion Field Adaptions in Frozen Image-to-Video Diffusion Model

Muyao Niu¹, Xiaodong Cun^{2,*}, Xintao Wang², Yong Zhang²,
Ying Shan², and Yinqiang Zheng^{1,*}

¹ The University of Tokyo

muyao.niu@gmail.com, yqzheng@ai.u-tokyo.ac.jp

² Tencent AI Lab

vinthony@gmail.com, xintaowang@tencent.com,
zhangyong201303@gmail.com, yingsshan@tencent.com

https://myniuuu.github.io/MOFA_Video/




Fig. 1: We present MOFA-Video for controllable image animation. We train MOFA-Adapters for (a) manual trajectories animation, (b) facial landmarks sequences animation (SadTalker [50] is used for audio to landmark generation). These two adaptors can be combined in a zero-shot manner for (c) the animation from both trajectories and human landmarks without retraining.

Abstract. We present MOFA-Video, an advanced controllable image animation method that generates video from the given image using various additional controllable signals (such as human landmarks reference, manual trajectories, and another even provided video) or their combinations. This is different from previous methods which only can work on a specific motion domain or show weak control abilities with diffusion prior. To achieve our goal, we design several domain-aware motion field adapters (*i.e.*, MOFA-Adapters) to control the generated motions

* Corresponding authors

in the video generation pipeline. For MOFA-Adapters, we consider the temporal motion consistency of the video and generate the dense motion flow from the given sparse control conditions first, and then, the multi-scale features of the given image are wrapped as a guided feature for stable video diffusion generation. We naively train two motion adapters for the manual trajectories and the human landmarks individually since they both contain sparse information about the control. After training, the MOFA-Adapters in different domains can also work together for more controllable video generation. Codes available: <https://github.com/MyNiuuu/MOFA-Video>

1 Introduction

Bringing images to life is considered magic in the old days. In the traditional Chinese story “The Magic Brush Ma Liang”, the author imagines a magic pen that can directly draw the living photos. Coincidentally, the story of Harry Potter creates a world where the dead ancestors are living in the wall paintings. Besides fiction, the exploration to make it come true has never stopped. In 1878, Muybridge presented a well-known experiment called “The Horse in Motion”, which shows a series of successive pictures of the running horse consecutively that can be considered a video.

With the development of digital devices, current methods attempt to animate the image using computer vision algorithms [8, 13, 16, 17, 21, 24, 27, 30, 32, 34, 36, 40, 50]. However, it faces several limitations. On the one hand, these methods often focus on limited categories of animated objects such as fluids [16, 24, 25], human hairs [37], and human body/face [6, 8, 9, 13, 17, 27, 32, 34, 36, 50]. Thanks to the domain knowledge of each specific type, these methods often have fully controllable abilities of the scenes. *e.g.*, SadTalker [50] can produce the accurate human face animation by audio and the given face. Text2Cinemagraph [25] produces the natural animation of water using text description. For control abilities, these methods usually follow the rule of learning the video via self-supervised decomposition and then animating via new driving signals. However, due to the limitation of natural animation prior, these methods fail in the general image space because of the diversity of the general domain knowledge.

Unlike previous in-domain image animation, current diffusion-based Image-to-Video (I2V) methods learn to generate the video from the image in an end-to-end fashion. Thanks to the large-scale generative prior of the text-to-image model, *i.e.*, Stable Diffusion [29], these methods [1, 2, 7, 11, 39] have proved the possibility of open-domain image animation. However, their generated contents might be different from the given image [1, 2, 11, 39], and often generate simple motions by text descriptions [1, 2, 39] or just simple idle animation [7]. These drawbacks limit their applications for real-world image animation tasks, where users often need to create more controllable videos as in previous in-domain image animation algorithms.

Taking advantage of both in-domain image animations and image-to-video generations, we are curious: *is there a general image animation framework that*

supports meticulous control abilities over in-the-wild images? We then find that all the animations can be formulated by the motion propagation of the sparse key-points (or key-frame) as the control handle.

To this end, we present MOFA-Video to add the additional different motion control abilities to the generative video diffusion model (Stable Video Diffusion [7] in our case), inspired by ControlNet [48]. In detail, to animate an input image into a video according to the sparse control signals from multiple domains, we design a novel MOFA-Adapter that serves as an additional adapter on the pre-trained video diffusion model so that the motions of the video can be controlled. Different from the previous ControlNet-like Adapter [38, 51] for video generation, the proposed MOFA-Adapter models the frame-wise motion explicitly. In detail, we first utilize the given sparse motion hint to generate the dense motion fields using a sparse-to-dense motion generation network, then, we warp the multi-scale features of the first frame as the conditional features for the diffusion generation process. This sparse-to-dense motion generation provides a good balance between the provided motion guidance and the generation process, providing high-quality animation results with good temporal consistency. We also think about the problems of there containing multiple motion domains. Thus, we train multiple MOFA-Adapters by considering these tasks as sparse control point generation problems, including open-world manual trajectories, human facial animations, *etc.* In addition, since the parameters of the video diffusion model are fixed, we can jointly perform motion control abilities across multiple domains, *e.g.*, the human face and background objects and the camera movement. We give more detailed applications and examples in the experiments.

The contribution of this paper can be summarized as:

- We propose a novel unified framework for controllable image animation in Stable Video Diffusion (SVD).
- We design a novel network structure, *i.e.*, MOFA-Adapter, which utilizes the explicit sparse motion hint for warping and generation.
- Detailed experiments and ablation show the advantage of the proposed method over current ones.

2 Related Work

2.1 Controllable Image Animation

Image animation has a long history in computer vision and graphics. Since it is a very ill-posed problem, previous work only focuses on the specific domain. *e.g.*, previous methods to generate the video from the image in an unsupervised and stochastic manner [16, 20, 22, 41] using generative model, *e.g.*, generative adversarial networks (GAN). These methods do not provide the control handle of the generation, limiting its applications on real-world applications. On the other hand, controlling the motion of the generation is also difficult for in-the-wild images. To this end, recent works aim to control the fluids [24, 25], human pose/face [8, 9, 13, 17, 27, 32, 34, 36] only, other than the general scenes. Different

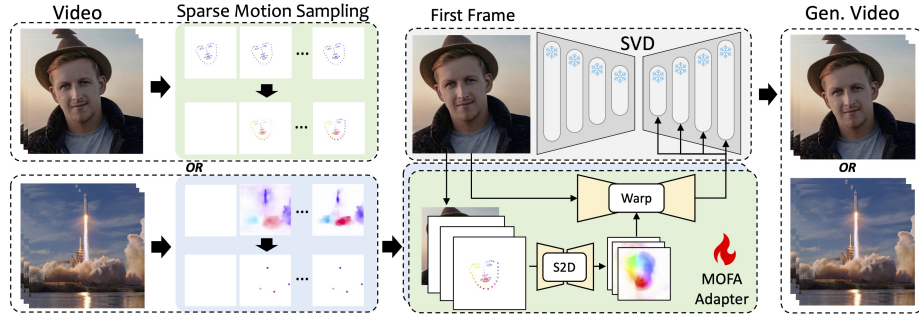


Fig. 2: Overview of MOFA-Video. We design MOFA-Adapters for adapting the motions from different domains with a unified structure on the frozen Video Diffusion Model. It generates the video from a single image and the corresponding sparse motion hints. For training, we generate the sparse motion hints through sparse motion sampling and then train different MOFA-Adapters to generate video via pre-trained SVD [7].

from these methods, we propose a novel and unified framework to model the motions from different domains and make it work singly and together in a pre-trained Video Diffusion Model.

2.2 Image-to-Video Diffusion Models.

In the realm of downstream tasks utilizing VDMs for video-related applications, there exists a category of work known as Image-to-Video Diffusion Models (I2Vs) [1–4, 7, 39, 52]. However, current I2V models only generate the video from the given image and aim to control the motions from the text description, we argue it is somewhat hard for text-based motion generation since they contain limited motion priors. Some related methods also try to control the generation of the motion. *e.g.*, DragNUWA [45] model the trajectory of the generation via adaptive normalization, which shows weak spatial correction. MotionCtrl [35] tries to control the object and camera motions of the T2V model, however, it is hard for the generation process since there is no world coordinate system for text-to-video generation. The most related controllable image animation is the concurrent work MotionI2V [31]. However, it focuses only on the natural motion of the objects, which is hard to fully control the motions from other domains, *e.g.*, the human face or body. Differently, we aim to create a method which is built on a pre-trained base model, *i.e.*, SVD, and also enjoys the motion hints from different domains using different adapters.

3 Method

We aim to generate videos from the given reference image and additional motion control signals in multiple motion domains (*e.g.*, hand-crafted trajectories, human landmark sequences, dense motion flows, *etc.*) in a unified framework, so they can share a unified network structure and jointly work together like Multi-ControlNet [48]. To achieve this goal, as shown in Fig. 2, we design a generative

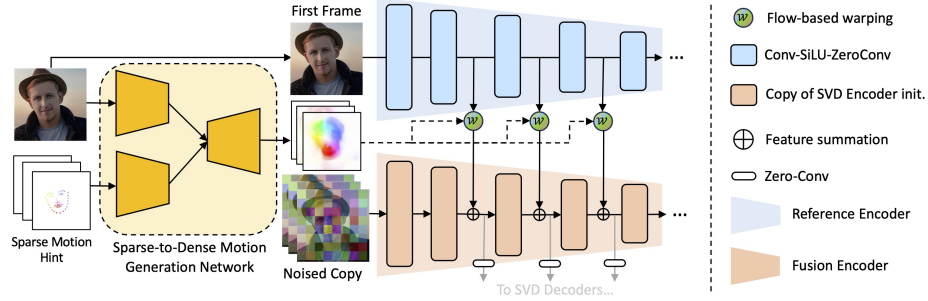


Fig. 3: Detailed Structure of MOFA-Adapter. It contains an S2D Network that accepts the motion hints and produces a dense motion field of the video. A reference encoder that extracts multi-scale features from the source image. A training-able copy of the SVD encoder, which initializes the weights from SVD and serves as the final spatial-temporal feature merging for generation guidance.

motion field adapter (MOFA-Adapter) that can accept sparse motion control signals as the condition and produce detailed control abilities for a frozen Stable Video Diffusion [7] Model. We train the proposed MOFA-Adapter in two different motion domains individually and provide various applications based on each model and their combinations.

In the following, we first introduce the structure of the proposed MOFA-Adapter in Sec. 3.1. Then, we give details of how we train the domain-aware MOFA-Adapter for the video diffusion model in Sec. 3.2. Finally, in Sec. 3.3, we give the inference details of the proposed method and various additional applications.

3.1 Generative Motion Field Adapter (MOFA-Adapter)

Given the sparse motion hints (*e.g.*, hand-crafted trajectories, human pose sequential trajectories), we design a new adapter structure, *i.e.*, MOFA-Adapter, for the pre-trained frozen Video Diffusion Model, so that the generated video of different subjects can be controlled individually or jointly in the generation process. Inspired by ControlNet [48], we consider this network as an additional motion control signal to the frozen denoising UNet of the video diffusion model. Below, we give the detailed network structure.

As shown in Fig. 3, the proposed adapter is based on a reference encoder, a sparse-to-dense (S2D) motion generator for sampled motion hints, and the feature fusion encoder to add the warped feature back to the pre-trained video diffusion model. In detail, the reference image encoder is a multi-scale convolutional feature encoder, which extracts the multi-scale features of the first frame for warping, with each stage built by Conv-SiLU-ZeroConv [48] as the basic block. For the sparse-to-dense motion generator, we use the same network structure as CMP [47] for adaption. This network is also a convolutional neural network which accepts the first frame image and the sparse hint of the motion, and produces the dense motion fields. More details of the network structure can be founded in the original paper [47]. When the dense motion field is generated,

we warp the referenced features and then add them to the feature map of the corresponding levels of a copied SVD encoder, which is then added to the feature space of the decoder of the pre-trained SVD similar to ControlNet [48].

3.2 Training MOFA-Adapters on Stable Video Diffusion

We use Stable Video Diffusion [7] as our base image-to-video diffusion model which accepts the image as input and generates video with idle animations. It is a latent diffusion model [29] which firstly compresses the reference image into the latent space using a pre-trained auto-encoder, and then, the video is generated via the sampled Gaussian noise, the condition image, and the diffusion process [15].

Fig. 2 illustrates the training pipeline of our whole framework. Given a video clip $V \in \mathbb{R}^{L \times 3 \times H \times W}$ in L frames, we first extract the sparse motion vectors which serve as input to the S2D network. *e.g.*, for open-domain, we handle the motion hint as the sparse motion vectors sampled from the extracted dense optical flow. As for human motion, we generate the motion hints from structural key-points, *e.g.*, facial landmarks. In the following, we give the details of each specific type:

Sparse Motion Vectors from Dense Optical Flow. By considering the dense optical flow as a general motion representation between video frames, we first utilize Unimatch [42] to extract the forward flow as $F \in \mathbb{R}^{(L-1) \times 2 \times H \times W}$, where $F_i \in \mathbb{R}^{2 \times H \times W}$ represents the optical flow from the 0-th to the $(i+1)$ -th frames. Based on the flow sequence F , we sample n spatial points for each frame F_i using the watershed sampling strategy [47]. Specifically, we first obtain a sparse mask $M^s \in \mathbb{R}^{H \times W}$, where the value for the sampled spatial points is set to 1, and other points are set to 0. We then calculate the sparse motion vectors $F^s \in \mathbb{R}^{(L-1) \times 2 \times H \times W}$ as:

$$F_{::,i,j}^s = \begin{cases} F_{::,i,j} & \text{if } M_{i,j}^s = 1, \\ 0 & \text{if } M_{i,j}^s = 0. \end{cases} \quad (1)$$

Sparse Motion Vectors from Structural Human Key-Points. Different from natural motion fields, human key-points provide concise while semantically meaningful representations. In our approach, we consider the movement of a group of key-points as a special case of the sparse motion vectors mentioned earlier. This unified representation simplifies our framework and allows us to share the mutual prior information of the S2D model. Specifically, given a series of 2D facial landmarks $P \in \mathbb{R}^{L \times K \times 2}$ extracted from an L -frames portrait video, we consider the motion difference between the landmarks of reference (first) frame \hat{P} and P , calculating point-wise sparse flow F^s via:

$$F^s[l-1, :, \hat{P}[k, 0], \hat{P}[k, 1]] = P[l, k, :] - \hat{P}[k, :], \quad (2)$$

where $l \in \{1, 2, \dots, L-1\}$. $k \in \{0, 1, \dots, K-1\}$. This motion mapping enables us to incorporate key-point information into our framework effectively as Fig. 2.

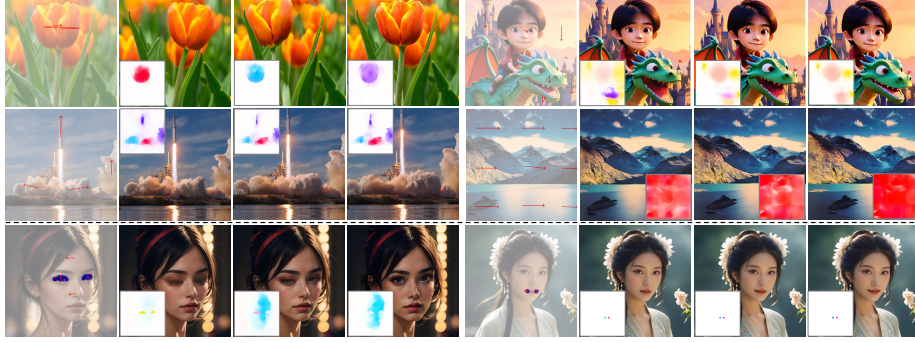


Fig. 4: Trajectory-based Animation. Image animation results from different trajectories. Results below the dashes are the fine-grained results using motion brushes. Intermediate optical flow results are also visualized.

After unifying the motion guidance from different domains, following the training process of stable video diffusion [7], our model is expected to reconstruct the video clips V using the first frame I and the sampled sparse motion vector F^s . Formally, the diffusion-based method first encodes video V into the latent space through $\mathcal{V} = \mathcal{E}(V)$, then progressively adds noise $\{\epsilon_t\}_{t=0}^{T-1}$ to \mathcal{V} to produce $\{\mathcal{V}_t\}_{t=0}^{T-1}$, where t represents the time step of the added noise. With the pre-trained Stable Video Diffusion \mathcal{S} and the MOFA-Adapter \mathcal{M} , we optimize the weights of the MOFA-Adapter $\theta_{\mathcal{M}}$ via :

$$\mathcal{L} = \|\mathcal{S}(\mathcal{V}_t, t, \mathcal{M}(\mathcal{V}_t, t, I, F^s; \theta_{\mathcal{M}})) - \mathcal{V}\|^2, \quad (3)$$

where \mathcal{L} is the overall learning objective. After that, the video diffusion model is run T time steps to recover the clear latent from pure sampled Gaussian noise. Finally, the recovered latent is decoded via the pre-trained auto-encoder to the sequence of the image as the produced video.

3.3 Inferences

After training, our method can generate video from a single image and the given control signals, *e.g.*, handcrafted trajectories, facial key-points, *etc.* In the following parts, we introduce several inference pipelines and applications of our method, including hand-crafted trajectory-based animation, motion brush masking, and portrait animation with key-points, *etc.*

Hand-Crafted Trajectories for Objects/Camera Movement. Given an input image $I \in \mathbb{R}^{H \times W \times 3}$, the users can draw multiple motion trajectories on the image, each of which can be represented as P joint points $\mathcal{T} \in \mathbb{R}^{P \times 2} = [(x_0, y_0), (x_1, y_1), \dots, (x_{P-1}, y_{P-1})]$. Take one trajectory for example, we bicubically interpolate this trajectory to L points, *i.e.*, $\hat{\mathcal{T}} \in \mathbb{R}^{L \times 2}$ and calculate the sparse motion hints $F^s \in \mathbb{R}^{(L-1) \times 2 \times H \times W}$ between the frames as:

$$F^s[l-1, :, \hat{\mathcal{T}}[0, 0], \hat{\mathcal{T}}[0, 1]] = \hat{\mathcal{T}}[l, :] - \hat{\mathcal{T}}[0, :], \quad (4)$$

where $l \in \{1, 2, \dots, L-1\}$.




Fig. 5: Facial Landmarks based Animation. We produce facial landmarks from the driven audio using SadTalker [50]. Then, the portrait animation can be produced by the facial MOFA-Adapter. Intermediate optical flow results are also visualized.

As a result, we can create several trajectories to indicate the motions of the specific object individually. As shown in Fig. 4, we can handle trajectories in both linear and non-linear fashion since our model is able to handle both these two types of motion via optical flows. In addition, by adding relatively dense guidance on the whole images, our method also has the capability to control the motions of the entire scene to perform as “camera motion”.

Regional Motion Brushes. Besides simple hand-craft trajectories, we can generate the regional motion animation based on the trajectory model. This feature is different from current implicitly drag-based methods (DragNUWA [45]), where the animated region is determined by the model itself instead of the user. Specifically, the user could additionally provide a binary motion mask $M \in \mathbb{R}^{H \times W}$ to control the animation region. Given this mask, we divide the trajectory set \mathcal{T} into two groups: \mathcal{T}_{in} and \mathcal{T}_{out} . For each trajectory, we distribute it to \mathcal{T}_{in} if its start points lie in the mask, otherwise to \mathcal{T}_{out} . We then respectively obtain the dense optical flow F_{in} and F_{out} via the S2D network, and mask out all the values that are located outside the mask. We then obtain the final dense optical flow by combining the masked F_{in} and F_{out} . We demonstrate the corresponding results in Fig. 4.

Portrait Image Animation. Besides natural object movement, we also consider the motion domain of the human, specifically the simple face movement, which is more detailed and structural. For inference animation, we first use an audio-driven talking face generation method, *i.e.*, SadTalker [50], to produce facial landmarks from a single image, then, as introduced in Sec. 3.1, we generate the sparse motion hints and use the trained MOFA-Adapter for motion generation. Fig. 5 demonstrates image animation results via audio-driven facial key-points generation [50] and a single image. Notice that, since we use facial landmarks as the intermediate representation, we can also drive the animation with another facial video or manual facial blendshapes animation.

Longer Animation. Despite the significant generation ability provided by the SVD, the frame length of its generated video is extremely limited to the specific frames (*e.g.*, 14 in our case). The computational complexity of temporal attention scales quadratically with the number of frames, making it difficult to directly generate ultra-long videos. On the other hand, ensuring consistency becomes more challenging with the increase in the number of frames. To solve this

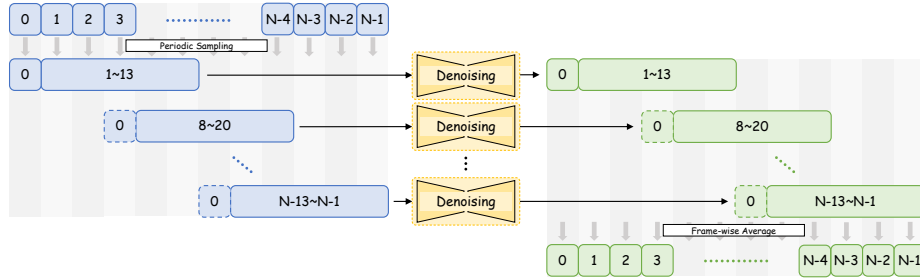


Fig. 6: Periodic Sampling for Longer Animation. The first frame (0) is used as the condition and we produce latent by frame-wise average over sequential 7 frames.




Fig. 7: Animation via Multi MOFA-Adapters. We show the combination of the hand-craft trajectories and audio-driven facial landmarks-based MOFA-adapters for more complex animations.

limitation, we propose the periodic sampling strategy to resolve the frame number issue and generate longer videos, inspired by Gen-L-Video [33]. As illustrated in Fig. 6, in each diffusion step, we periodically sample 14 frames as one group and denoise them using our SVD model. Each latent group has an overlapping of 7 frames to provide temporal smoothness. After all the groups are sampled, we average the predicted noise for each frame and obtain the final latent for this diffusion step.

Plugin Combinations. Our careful design enables us to animate the given image by simultaneously using different control signals from multiple modalities without any retraining. This is achieved by combining multiple MOFA-Adapters for the animation similar to Multi-ControlNet [48]. In Fig. 7, we show an example of controlling the facial expression of one portrait while animating other objects such as backgrounds via hand-craft trajectories. Please refer to the supplementary materials for more experiment results and technical details.

4 Experiment

Implementation Details. For training, we train the proposed trajectories-based MOFA-Adapter on WebVid-10M [5] using 8 NVIDIA A100 GPU. The base model is Stable Video Diffusion [7]. We use AdamW [23] as an optimizer. During training, we randomly sample 14 video frames with a stride of 4. The learning rate is set to 2×10^{-5} with a resolution of 256×256 . We first train the model as a flow-based reconstruction model by removing the S2D motion generator and directly taking the first frame together with the estimated optical flow from Unimatch [42]. After that, we add the S2D network and finetune the whole adapter. For portrait-based MOFA-Adapter, we train our model on a self-

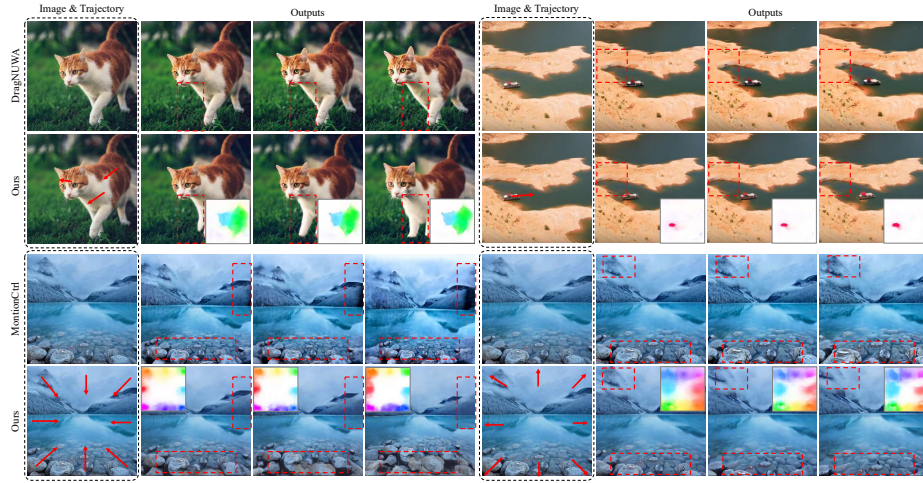


Fig. 8: Comparing trajectory-based control with DragNUWA [45] and Camera control with MotionCtrl [35] using proposed trajectory-based MOFA-Adapter.

Methods	Metrics				User Preference	
	Fra.Con.↑	LPIPS↓	FID↓	FVD↓	Ctrl.Pre.↑	Vis.Qua.↑
DragNUWA	0.9302	0.2705	19.66	91.38	2.76	3.18
Ours	0.9390	0.2274	16.82	86.76	3.58	3.42

Table 1: Quantitative comparison and user study results for trajectory-based image animation.

Methods	Metrics		User Preference		
	CPBD↑	ID↑	Fide.↑	Natur.↑	Vis.Qua.↑
SadTalker	0.3218	0.9188	4.15	3.12	3.97
StyleHEAT	0.2577	0.7993	3.26	3.65	3.7
Ours	0.4075	0.9293	4.8	3.97	4.52

Table 2: Quantitative comparison and user study results for portrait image animation.

collected human video dataset. The pre-trained DWPose [43] is used to extract the facial landmark and other training settings are the same as the trajectories-based model. Please refer to the supplementary for more details.

Evaluation Metrics. For the trajectory-based MOFA-Adapter, we use 1000 samples from the test set of WebVid and respectively calculate LPIPS [49], FID [14], and FVD [46]. We also calculate the cosine similarity of the CLIP [28] embedding of the consecutive generated frames to evaluate the frame consistency of the results. To evaluate the keypoint-based MOFA-Adapter, we choose 40 generated video examples, each of which contains 196 frames, and calculate the following two metrics inspired by previous work [50]: 1) Cumulative Probability Blur Detection (CPBD) [26] to evaluate the sharpness of the generated frames, and 2) identity embedding from ArcFace [12] between the source images and the generated frames to evaluate the fidelity of the generated results. Apart from the quantitative metrics, we also conduct user studies to perceptually evaluate the result of different methods.

4.1 Comparison with Other State-of-the-Art Methods

In this paper, we propose a method that can perform various applications from different modalities by using a unified architecture. In this section, we demonstrate and compare our results with the state-of-the-art methods that specialize in each application field.

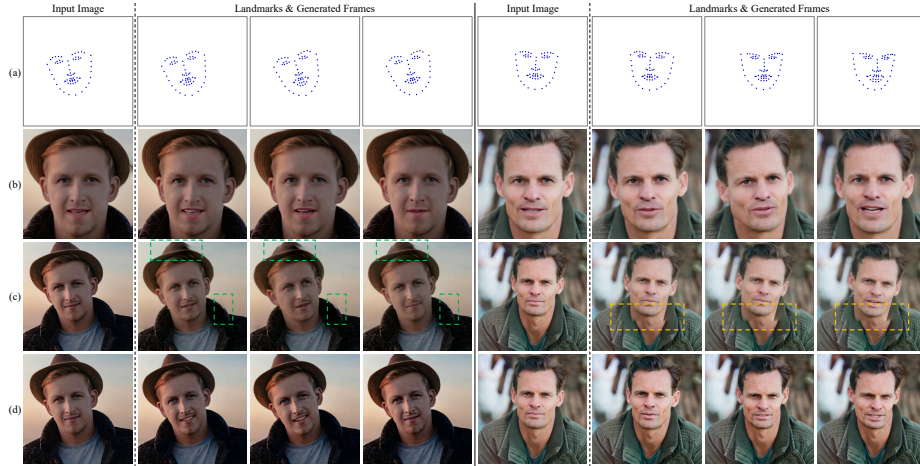


Fig. 9: Qualitative Comparisons for Portrait Animation Using Audio. (a) Generated 3D Landmarks from SadTalker, (b) StyleHEAT [44], (c) SadTalker [50], (d) Ours.

Trajectory-based Image Animation. For pure trajectory-based image animation, we compare our approach with DragNUWA [45]. Their method integrates the sparse motion vector into the diffusion model without employing explicit flow and warping. Consequently, this makes it entirely uninterpretable and potentially causing various limitations. As depicted in the right part of Fig. 8, DragNUWA incorrectly extends the influence of the user annotation to the background, causing the beach to move to the right even though the user’s intention was to control the ship’s cruise direction. This issue is particularly prone to occur in DragNUWA when the motion prior to specific parts is overly strong. In comparison, our method avoids this problem by explicitly determining the moving region through dense optical flow, thus yielding superior results. The quantitative results are shown on Tab. 1, where the proposed method shows much better performance in terms of all metrics.

Furthermore, our model can effectively control camera motion by employing various trajectory patterns by the suitable trajectories, including zooming in, zooming out, panning left, panning right, rotating clockwise, rotating counter-clockwise, *etc.* We compare our approach with MotionCtrl [45]. As shown in Fig. 8, our model achieves comparable results with MotionCtrl, with fewer visual artifacts and larger motion magnitudes. Notably, since our method immediately uses the optical flow to control the video diffusion model, we can also control the camera motion directly using fixed optical flow patterns. Please refer to the supplementary for more details.

Portrait Image Animation. We evaluate our keypoint-based MOFA-Adapter in comparison to prior state-of-the-art techniques such as StyleHEAT [44], SadTalker [50]. The qualitative results are shown in Fig. 9. As StyleHEAT [44] requires aligned images as input, we first crop and align the input image, then display the resulting output in the aligned format. We can see that Style-

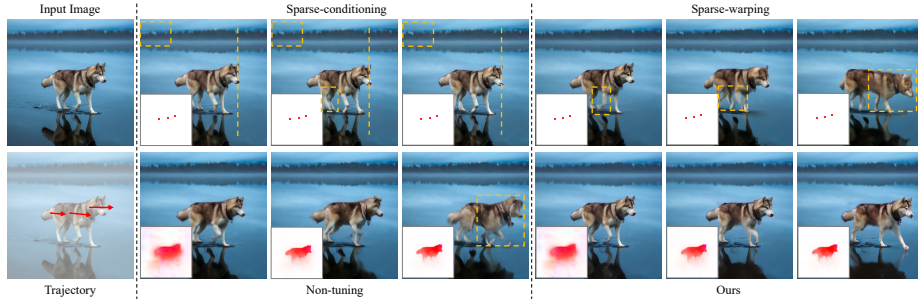


Fig. 10: Ablation on Network Design. Our full model achieves both the best controllability and synthesis quality compared to three baseline methods: 1) Sparse-conditioning model (w/o warping), 2) Sparse-warping model (w/o S2D network), and 3) Non-tuning model (w/o tuning).

HEAT [44] tends to generate overly smooth outcomes and struggles with identity preservation issues due to the inherent limitations of latent inversion via StyleGAN [18, 19]. SadTalker [50] manages to process unaligned images using the paste-back technique, but it creates noticeable unnatural boundary artifacts. Additionally, it only controls head movement, leaving the target’s body static during head motion, and may produce low-resolution outputs (256×256), leading to a resolution discrepancy between facial and other parts. In comparison, our method effectively adapts to unaligned portrait images, removing noticeable paste-back artifacts, synchronizing body movements with head motion, and enabling high-resolution synthesis (up to 1024×1024) results. The corresponding quantitative results are reported in Tab. 2. Our method achieves the best results on all metrics, demonstrating its superiority over existing methods.

User Studies. We also conduct user studies to evaluate the effectiveness of the proposed method compared with baseline methods. For trajectory-based animation, we compare with DragNUWA using the same trajectories. For facial animation, we produce the same landmarks as SadTalker [50] using the same audio. For each task, we choose 20 samples and invite 10 volunteers to evaluate the results. For trajectory-based animation, we consider the visual qualities and the control precision in the score range of 1 to 5 (higher the better). For facial animation, we consider the fidelity to the input image, the naturalness of the motion, and the visual qualities. The results are reported in Tab. 1 and Tab. 2. We can see that the volunteers prefer the proposed method in terms of all the aspects.

4.2 Ablation Studies

Network Structure of MOFA-Adapter. To evaluate the effectiveness of our network design, we compare the results of our method with three different variants:

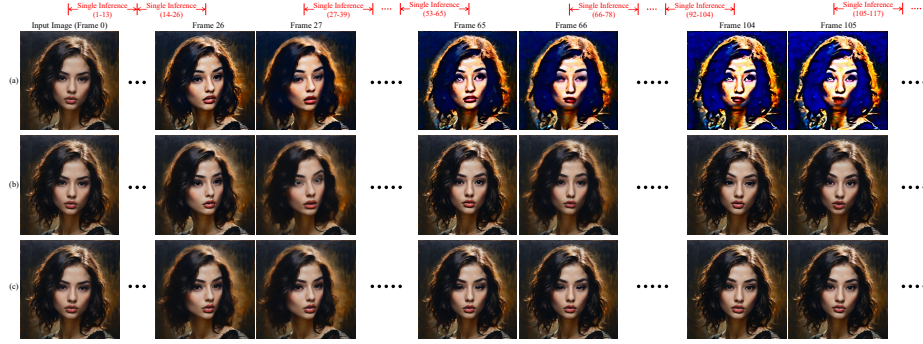


Fig. 11: Ablation on Longer Video Inference. (a) Condition on the *latest frame of previous inference* will show over-exposure results (Dynamic-conditional Naive Separation). (b) Condition on the *input frame and no frame overlap* will cause “sudden changes” between inferences clip (Zero-conditional Naive Separation). (c) Our proposed periodic sampling strategy can successfully generate the whole long sequences.

- **Sparse-conditioning model (w/o warping).** We directly take the sparse trajectories with masks as conditions for Motion-aware ControlNet and do not use spatial warping in the feature space.
- **Sparse-warping model (w/o S2D network).** We directly uses the sparse optical flow to spatially warp the deep features without using the proposed S2D model to turn it into a dense flow representation.
- **Non-tuning model (w/o tuning).** Directly use the flow-based reconstruction model, without further finetuning.

The corresponding results are in Fig. 10. The sparse-conditioning model is unable to precisely manipulate the target object’s trajectory, as the dog does not move to the right, and the background slightly shifts to the left. We argue that this issue arises from the spatial misalignment caused by not employing the spatial warping operation. While the sparse-warping model can successfully manage the target’s trajectory, it struggles to effectively transform the deep feature from the initial frame due to the lack of dense feature control provided by dense optical flow, resulting in significant artifacts in the output. The non-tuning model, however, experiences over-control from the dense optical flow and is unable to generate suitable content using the strong prior of the SVD. In comparison, our comprehensive model achieves optimal results.

Longer Video Inference. Additionally, we evaluate our periodic sampling approach against two alternative strategies:

- **Dynamic-conditional Naive Separation** separates all frames into groups of 13, generating each group based on the *generated last frame of the preceding group*, and then combines all the frames.
- **Zero-conditional Naive Separation** separates all frames into groups of 13, generating each group based on the *first frame*, and then combines all the frames.

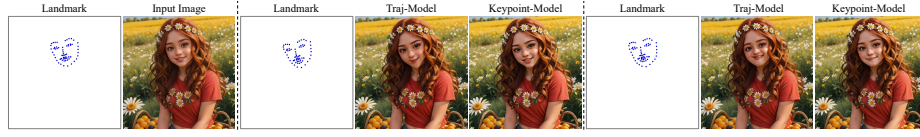


Fig. 12: Ablation on Domain-Aware MOFA-Adapter. Using the model in the trajectory domain without specific tuning for portrait animation will lead to unsatisfactory results due to the domain gap.

As shown in Fig. 11, the Dynamic-conditional Naive Separation strategy (a) leads to significant error accumulation, resulting in unsatisfactory results in the later stages of the long generation process. Although the Zero-conditional Naive Separation strategy (b) mitigates error accumulation by consistently conditioning on the first frame, noticeable flickering effects and temporal inconsistency are apparent between the first frame of the current group and the last frame of the previous group, such as frames 26 and 27, 65 and 66, and 104 and 105 in Fig. 11. In comparison, our periodic sampling algorithm (c) effectively addresses both error accumulation and temporal inconsistency by blending features from neighboring groups in the latent space. Please refer to the supplementary video results, which provide a more intuitive and clear demonstration of the superiority of periodic sampling.

The effectiveness of Domain-Aware MOFA-Adapter. As introduced in Sec. 3, we train different MOFA-Adapters for different motion domains. To evaluate the effectiveness and the necessity of domain-specific tuning, we directly use our trajectory-based model to perform landmark-based portrait image animation. As shown in Fig. 12, directly applying the model trained on open-domain will cause unnatural facial expressions.

4.3 Limitations

Although our method achieves remarkable progress in controllable image animation, it still has some limitations. First, different from SORA [10], our method is hard to control (generate) new content which is far from the given image since the current video diffusion model only trains on limited video clips. Second, our model may suffer from visual artifacts including blurriness or structure loss under large motion guidance. Visual examples are presented in the supplementary.

5 Conclusion

In this paper, we introduce a novel pipeline for generic and controllable image-to-video animation from multiple motion domains (*e.g.*, handcrafted trajectories, dense flows, and human key-points). To achieve this goal, we design the sparse-to-dense MOFA-Adapter to control the generated motions in the video generation pipeline. Powered by the proposed framework, we can achieve controllable video generation via fine-grained control which is unified as flow motion fields. The experiments show the advantage of the proposed framework over current state-of-the-art methods for various applications.

Acknowledgements

This research was supported in part by JSPS KAKENHI Grant Numbers 24K22318, 22H00529, 20H05951, JST-Mirai Program JPMJMI23G1.

References

1. Gen-2. <https://runwayml.com/ai-magic-tools/gen-2/> (2023)
2. Genmo. <https://www.genmo.ai/> (2023)
3. I2vgen-xl. <https://modelscope.cn/models/damo/Image-to-Video/summary> (2023)
4. Pika labs. <https://www.pika.art/> (2023)
5. Bain, M., Nagrani, A., Varol, G., Zisserman, A.: Frozen in time: A joint video and image encoder for end-to-end retrieval. In: IEEE International Conference on Computer Vision (2021)
6. Bertiche, H., Mitra, N.J., Kulkarni, K., Huang, C.H.P., Wang, T.Y., Madadi, M., Escalera, S., Ceylan, D.: Blowing in the wind: Cyclenet for human cinemagraphs from still images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
7. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023)
8. Blattmann, A., Milbich, T., Dorkenwald, M., Ommer, B.: ipoke: Poking a still image for controlled stochastic video synthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021)
9. Blattmann, A., Milbich, T., Dorkenwald, M., Ommer, B.: Understanding object dynamics for interactive image-to-video synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
10. Brooks, T., Peebles, B., Holmes, C., DePue, W., Guo, Y., Jing, L., Schnurr, D., Taylor, J., Luhman, T., Luhman, E., Ng, C., Wang, R., Ramesh, A.: Video generation models as world simulators (2024), <https://openai.com/research/video-generation-models-as-world-simulators>
11. Chen, H., Xia, M., He, Y., Zhang, Y., Cun, X., Yang, S., Xing, J., Liu, Y., Chen, Q., Wang, X., Weng, C., Shan, Y.: Videocrafter1: Open diffusion models for high-quality video generation (2023)
12. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4690–4699 (2019)
13. Dorkenwald, M., Milbich, T., Blattmann, A., Rombach, R., Derpanis, K.G., Ommer, B.: Stochastic image-to-video synthesis using cinns. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
14. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* **30** (2017)
15. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)

16. Holynski, A., Curless, B.L., Seitz, S.M., Szeliski, R.: Animating pictures with eulerian motion fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
17. Karras, J., Holynski, A., Wang, T.C., Kemelmacher-Shlizerman, I.: Dreampose: Fashion video synthesis with stable diffusion. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
18. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4401–4410 (2019)
19. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8110–8119 (2020)
20. Li, Y., Fang, C., Yang, J., Wang, Z., Lu, X., Yang, M.H.: Flow-grounded spatial-temporal video prediction from still images. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 600–615 (2018)
21. Li, Z., Tucker, R., Snavely, N., Holynski, A.: Generative image dynamics. arXiv preprint arXiv:2309.07906 (2023)
22. Logacheva, E., Suvorov, R., Khomenko, O., Mashikhin, A., Lempitsky, V.: Deep-landscape: Adversarial modeling of landscape videos. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII 16. pp. 256–272. Springer (2020)
23. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
24. Mahapatra, A., Kulkarni, K.: Controllable animation of fluid elements in still images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
25. Mahapatra, A., Siarohin, A., Lee, H.Y., Tulyakov, S., Zhu, J.Y.: Text-guided synthesis of eulerian cinemagraphs (2023)
26. Narvekar, N.D., Karam, L.J.: A no-reference image blur metric based on the cumulative probability of blur detection (cpbd). *IEEE Transactions on Image Processing* **20**(9), 2678–2683 (2011)
27. Ni, H., Shi, C., Li, K., Huang, S.X., Min, M.R.: Conditional image-to-video generation with latent flow diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
28. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning (ICML). PMLR (2021)
29. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
30. Rott Shaham, T., Dekel, T., Michaeli, T.: Singan: Learning a generative model from a single natural image. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2019)
31. Shi, X., Huang, Z., Wang, F.Y., Bian, W., Li, D., Zhang, Y., Zhang, M., Cheung, K.C., See, S., Qin, H., et al.: Motion-i2v: Consistent and controllable image-to-video generation with explicit motion modeling. arXiv preprint arXiv:2401.15977 (2024)
32. Siarohin, A., Woodford, O.J., Ren, J., Chai, M., Tulyakov, S.: Motion representations for articulated animation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)

33. Wang, F.Y., Chen, W., Song, G., Ye, H.J., Liu, Y., Li, H.: Gen-l-video: Multi-text to long video generation via temporal co-denoising. arXiv preprint arXiv:2305.18264 (2023)
34. Wang, Y., Yang, D., Bremond, F., Dantcheva, A.: Latent image animator: Learning to animate images via latent space navigation. In: International Conference on Learning Representations (2022)
35. Wang, Z., Yuan, Z., Wang, X., Chen, T., Xia, M., Luo, P., Shan, Y.: Motionctrl: A unified and flexible motion controller for video generation. In: arXiv preprint arXiv:2312.03641 (2023)
36. Weng, C.Y., Curless, B., Kemelmacher-Shlizerman, I.: Photo wake-up: 3d character animation from a single photo. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
37. Xiao, W., Liu, W., Wang, Y., Ghanem, B., Li, B.: Automatic animation of hair blowing in still portrait photos. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
38. Xing, J., Xia, M., Liu, Y., Zhang, Y., Zhang, Y., He, Y., Liu, H., Chen, H., Cun, X., Wang, X., et al.: Make-your-video: Customized video generation using textual and structural guidance. arXiv preprint arXiv:2306.00943 (2023)
39. Xing, J., Xia, M., Zhang, Y., Chen, H., Wang, X., Wong, T.T., Shan, Y.: Dynamicrafter: Animating open-domain images with video diffusion priors (2023)
40. Xiong, W., Luo, W., Ma, L., Liu, W., Luo, J.: Learning to generate time-lapse videos using multi-stage dynamic generative adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
41. Xiong, W., Luo, W., Ma, L., Liu, W., Luo, J.: Learning to generate time-lapse videos using multi-stage dynamic generative adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2364–2373 (2018)
42. Xu, H., Zhang, J., Cai, J., Rezatofighi, H., Yu, F., Tao, D., Geiger, A.: Unifying flow, stereo and depth estimation. IEEE Transactions on Pattern Analysis and Machine Intelligence (2023)
43. Yang, Z., Zeng, A., Yuan, C., Li, Y.: Effective whole-body pose estimation with two-stages distillation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4210–4220 (2023)
44. Yin, F., Zhang, Y., Cun, X., Cao, M., Fan, Y., Wang, X., Bai, Q., Wu, B., Wang, J., Yang, Y.: Styleheat: One-shot high-resolution editable talking face generation via pre-trained stylegan. In: European conference on computer vision. pp. 85–101. Springer (2022)
45. Yin, S., Wu, C., Liang, J., Shi, J., Li, H., Ming, G., Duan, N.: Dragnuwa: Fine-grained control in video generation by integrating text, image, and trajectory. arXiv preprint arXiv:2308.08089 (2023)
46. Yu, S., Tack, J., Mo, S., Kim, H., Kim, J., Ha, J.W., Shin, J.: Generating videos with dynamics-aware implicit generative adversarial networks. In: International Conference on Learning Representations (2022), <https://openreview.net/forum?id=Czsdv-S4-w9>
47. Zhan, X., Pan, X., Liu, Z., Lin, D., Loy, C.C.: Self-supervised learning via conditional motion propagation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1881–1889 (2019)
48. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models (2023)

- 49. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018)
- 50. Zhang, W., Cun, X., Wang, X., Zhang, Y., Shen, X., Guo, Y., Shan, Y., Wang, F.: Sadtalker: Learning realistic 3d motion coefficients for stylized audio-driven single image talking face animation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8652–8661 (2023)
- 51. Zhang, Y., Wei, Y., Jiang, D., Zhang, X., Zuo, W., Tian, Q.: Controlvideo: Training-free controllable text-to-video generation. arXiv preprint arXiv:2305.13077 (2023)
- 52. Zhang, Z., Long, F., Pan, Y., Qiu, Z., Yao, T., Cao, Y., Mei, T.: Trip: Temporal residual learning with image noise prior for image-to-video diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8671–8681 (June 2024)