

## X-Dyna: Expressive Dynamic Human Image Animation

Di Chang<sup>1,2</sup> Hongyi Xu<sup>2\*</sup> You Xie<sup>2\*</sup> Yipeng Gao<sup>1\*</sup> Zhengfei Kuang<sup>3\*</sup> Shengqu Cai<sup>3\*</sup>  
 Chenxu Zhang<sup>2\*</sup> Guoxian Song<sup>2</sup> Chao Wang<sup>2</sup> Yichun Shi<sup>2</sup> Zeyuan Chen<sup>2,5</sup>  
 Shijie Zhou<sup>4</sup> Linjie Luo<sup>2</sup> Gordon Wetzstein<sup>3</sup> Mohammad Soleymani<sup>1</sup>

<sup>1</sup> University of Southern California <sup>2</sup> ByteDance <sup>3</sup> Stanford University

<sup>4</sup> University of California Los Angeles <sup>5</sup> University of California San Diego

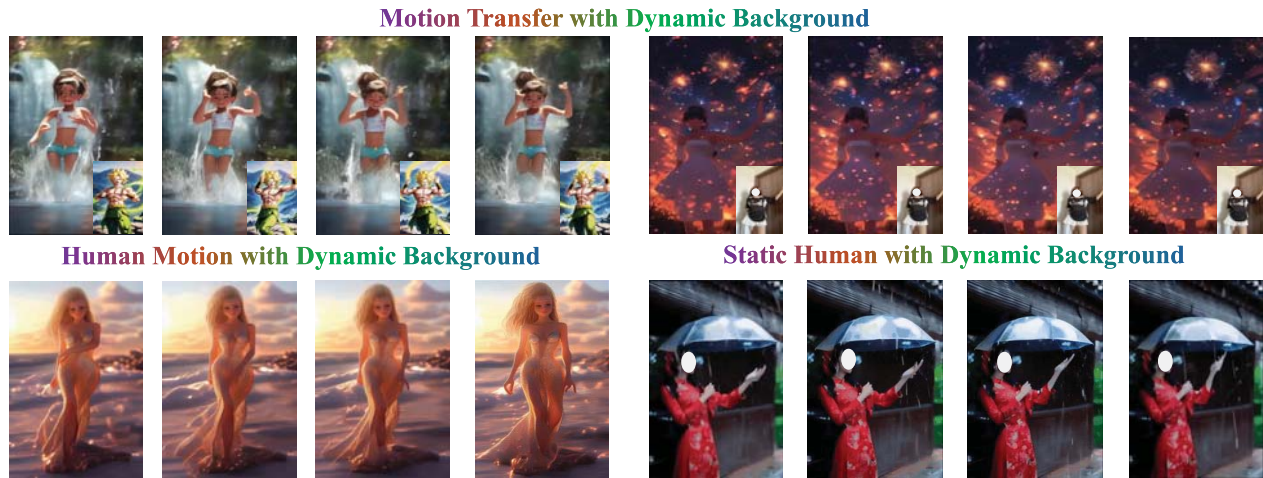


Figure 1. Sampled animations generated by our X-Dyna, including zero-shot motion transfer, and dynamic human image animation with moving or static humans.

### Abstract

We introduce X-Dyna, a novel zero-shot, diffusion-based pipeline for animating a single human image using facial expressions and body movements derived from a driving video, that generates realistic, context-aware dynamics for both the subject and the surrounding environment. Building on prior approaches centered on human pose control, X-Dyna addresses key shortcomings causing the loss of dynamic details, enhancing the lifelike qualities of human video animations. At the core of our approach is the **Dynamics-Adapter**, a lightweight module that effectively integrates reference appearance context into the spatial attentions of the diffusion backbone while preserving the capacity of motion modules in synthesizing fluid and intricate dynamic details. Beyond body pose control, we connect a local control module with our model to capture identity-disentangled facial expressions, facilitating accurate expression transfer for enhanced realism in animated scenes. Together, these components form a unified framework capable of learning physical human motion and nat-

ural scene dynamics from a diverse blend of human and scene videos. Comprehensive qualitative and quantitative evaluations demonstrate that X-Dyna outperforms state-of-the-art methods, creating highly lifelike and expressive animations. The code is available at <https://github.com/bytedance/X-Dyna>

### 1. Introduction

We investigate the task of human video generation, focusing on animating a single human image using body movements and facial expressions derived from a driving video of a different person. This area has garnered growing interest owing to its numerous applications in digital arts, social media and virtual humans. Building on prior research [4, 16, 37, 40, 48, 52, 53], our goal is to advance the field of zero-shot human image animation by not only enhancing the accuracy of pose and expression transfer but also by incorporating vivid human dynamics, e.g., blowing hair and flowing garments, and natural environmental effects, e.g., waterfalls, rain, and fireworks.

Recent approaches have tackled human image animation as a controlled image-to-video diffusion task. These meth-

\*Equally contributed as second authors

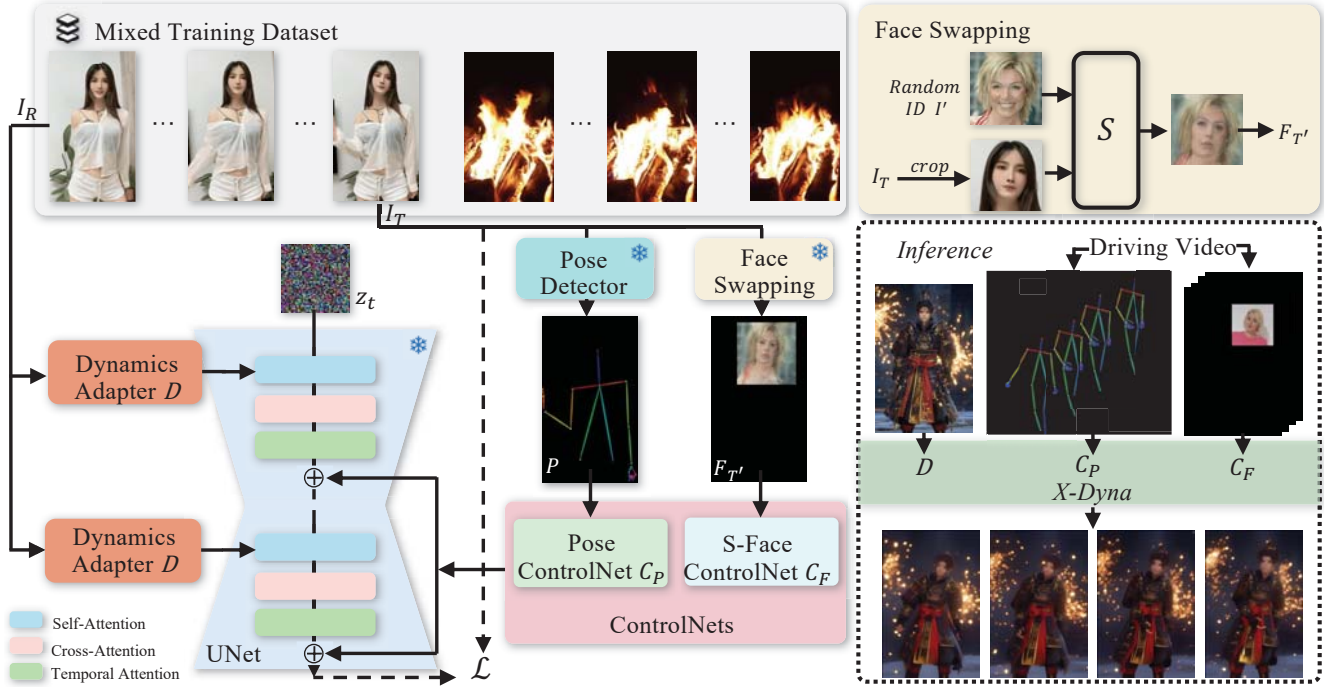


Figure 2. We leverage a pretrained diffusion UNet backbone for controlled human image animation, enabling expressive dynamic details and precise motion control. Specifically, we introduce a dynamics adapter  $D$  that seamlessly integrates the reference image context as a trainable residual to the spatial attention, in parallel with the denoising process, while preserving the original spatial and temporal attention mechanisms within the UNet. In addition to body pose control via a ControlNet  $C_P$ , we introduce a local face control module  $C_F$  that implicitly learns facial expression control from a synthesized cross-identity face patch. We train our model on a diverse dataset of human motion videos and natural scene videos simultaneously. Our model achieves remarkable transfer of body poses and facial expressions, as well as highly vivid and detailed dynamics for both the human and the scene.

ods typically employ a parallel UNet to incorporate reference appearance through mutual self-attentions [3], while body motion cues (e.g., 2D skeletons [4, 37, 52] and DensePose [48]) are integrated as spatial guidance through frameworks like ControlNet [51] or PoseGuider [16]. Temporal modules, such as AnimateDiff [10] and Align-Your-Latents [1, 2], have been introduced to the diffusion backbone, trained from large-scale videos to enhance consistency and dynamics in visual sequence generation. Despite improvements in control precision and generation realism, the combined modules for human image animation often fall short in capturing intricate visual dynamics, leading to static backgrounds and rigid human motions. This shortcoming, rooted in both network design and training data distribution, ultimately compromises the lifelike quality of the generated videos.

To this end, we propose **X-Dyna**, a diffusion-based human image animation pipeline that achieves the accurate transfer of pose and facial expressions along with consistent and vivid human and background dynamics. We observe that the loss of dynamic details primarily arises from the strong appearance constraints on spatial attentions imposed by the appearance reference modules, typically formulated as a trainable copy of a parallel UNet. To ad-

dress this, we introduce a lightweight cross-frame attention module, *Dynamics-Adapter*, which seamlessly propagates the reference appearance context to the denoising process by feeding the denoised reference image in parallel with noised sequences to the model. It integrates with the diffusion backbone via a trainable copy of the query projector and zero-initialized output projector, ensuring the backbone’s spatial and temporal generation capability stays intact. Unlike the standard I2V settings [1] that generate subsequent frames from the reference image, our design maintains appearance consistency from varying poses, in coordination with pose control modules. Notably, beyond body pose control, we employ a local control module to capture identity-disentangled facial expressions, enhancing realism with accurate expression transfer. While prior image animation models are primarily trained on human videos with static backgrounds, our dynamics adapter enables the learning of subtle human dynamics and fluid environmental effects, in addition to body and facial expression controls, from a diverse mixture of human and scene videos.

Trained on a curated dataset of 900h human dancing and natural scene videos, our method excels at accurately transferring the body poses and facial expressions while generating lifelike human and scene dynamics consistent with

the reference image context. We comprehensively evaluate our model on challenging benchmarks [17, 27, 29], and X-Dyna outperforms state-of-the-art human image animation baselines both quantitatively and qualitatively, demonstrating superior dynamics expressiveness, identity preservation and visual quality. Our main contributions are:

- A zero-shot diffusion-based human image animation model for both pose control and dynamics synthesis, trained on a mixture of human and natural scene videos;
- An efficient Dynamics-Adapter module that effectively incorporates reference appearance while maintaining the foundation model’s capability in generating high-quality dynamics;
- A local implicit face control module that enables refined, identity-disentangled facial expression control; and
- Demonstration of captivating zero-shot controllable human image animations and live photos with vivid dynamics.

## 2. Related Works

### 2.1. Diffusion Models for Human Video Animation

Recent advancements [30] in latent diffusion models [33] have greatly advanced human image animation. Previous approaches [4, 16, 41, 48] commonly employed a two-stage training paradigm: in the first stage, a pose-driven image model is trained on individual video frames paired with corresponding pose images; in the second stage, a temporal module is introduced to capture temporal dynamics, while the image generation model remains fixed. Following this framework, these methods integrate ReferenceNet [16] with a UNet architecture to extract appearance features from reference characters. With progress in the video foundation models, recent works [40, 52] have simplified this process by directly fine-tuning Stable Video Diffusion [1], effectively replacing the two-stage training approach. As mentioned in Sec. 1, there are several human video animation methods [4, 16, 18, 25, 26, 37, 40, 41, 45, 48, 52, 53], including CLIP [32] embedding with ControlNet [51], ReferenceNet [4, 16, 48, 53] with ControlNet [51], and SVD [1] with Pose Encoder [43]. However, these methods are not capable of capturing dynamics-related semantic information from the reference image and cannot provide a vivid animation of physical details from a natural background and human foreground.

### 2.2. Dynamics Generation

Dynamics generation has become a critical area in video generation, focusing on creating realistic motion and temporal consistency. GAN-based methods such as TGAN [34] and MoCoGAN [38] pioneered the decomposition of motion and content, allowing for better temporal coherence. However, GANs often struggle with complex motion

scenes, and artifacts may appear due to difficulties in modeling long-term dependencies. Later models, such as Progressive Growing of GANs [19], introduced gradual increases in resolution, achieving more stable results in video synthesis. Diffusion models have emerged as powerful alternatives for video generation, with methods like Video Diffusion Models [14] and AnimateDiff [10] incorporating temporal conditioning to ensure consistency across frames. AnimateDiff, for instance, applies temporal attention to produce smoother, continuous animations in human-centered videos. Similarly, Stable Video Diffusion [1] employs temporal modeling strategies to enhance dynamic texture quality, often surpassing GAN-based approaches in long-term coherence and photorealism. These works inspired most recent diffusion-based methods [7, 24] to further improve the ability for dynamics generation.

## 3. Method

Given a single reference image  $I_R$ , the objective of X-Dyna is to reanimate the human subject with a pose and expression sequence  $P_i$  derived from a driving video, where  $i = 1, \dots, T$  denotes the frame index. Most prior approaches decompose this task into two main sub-tasks: (1) transferring the appearance of the individual and background from the reference image and (2) controlling the video frames based on the pose and expression sequence  $P_i$ . X-Dyna not only focuses on generating temporally smooth image sequences but also aims to enhance lifelike dynamics realism. We achieve this by creating vivid and expressive dynamics for both the foreground human and background scenes in an end-to-end fashion, eliminating the need for any foreground and background disentanglement pre- or post-processing steps.

In Section. 3.1, we first examine existing network designs for transferring reference appearance and background, and identify their underlying causes for the loss of dynamic details. We then introduce our dynamics-adapter, which achieves accurate transfer of reference appearance with minimal impact on the diffusion backbone’s dynamics synthesis capability. To further enhance expression transfer and identity preservation, we integrate an additional local control module using synthetic cross-driven face images, as elaborated in Section 3.3. Our model design enables us to effectively learn human dynamics and environmental effects simultaneously from a diverse fusion of human motion and natural scenes videos (Section 3.4). Our pipeline is illustrated in Figure. 2.

### 3.1. Preliminary

**Latent Diffusion Model.** Facilitated by a pretrained auto-encoder, latent diffusion models [33] are a class of diffusion models [13, 35, 36] that synthesize desired samples in the image latent space, starting from Gaussian noise  $z_T \sim$



Figure 3. a) IP-Adapter [49] can generate vivid texture from the reference image but fails to preserve the appearance. b) Though ReferenceNet [16] can preserve the identity from the human reference, it generates a static background without any dynamics. c) Dynamics-Adapter provides both expressive details and consistent identities.

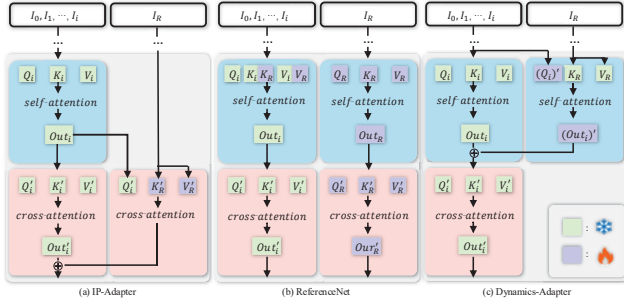


Figure 4. a) IP-Adapter [49] encodes the reference image as an image CLIP embedding and injects the information into the cross-attention layers in SD as the residual. b) ReferenceNet [16] is a trainable parallel UNet and feeds the semantic information into SD via concatenation of self-attention features. c) Dynamics-Adapter encodes the reference image with a partially shared-weight UNet. The appearance control is realized by learning a residual in the self-attention with trainable query and output linear layers. All other components share the same frozen weight with SD.

$\mathcal{N}(0, 1)$  and refining through  $T$  denoising steps. During training, latent representations of images are progressively corrupted by Gaussian noise  $\epsilon$ , following the Denoising Diffusion Probabilistic Model (DDPM) framework [13]. A UNet-based denoising backbone network containing intervened layers of convolutions and attentions, is trained to learn the reverse denoising process.

For our task, we employ a pretrained text-to-image (T2I) diffusion model Stable Diffusion (SD) as the generative backbone, with the addition of a ControlNet [51] module to incorporate 2D skeletal pose control as in recent human image animation work [4, 16, 37], as well as temporal modules [10] for enhanced consistency across generated video frames.

**Appearance Reference.** Previous research has introduced various strategies to maintain appearance consistency with a given reference image. Early approaches such as [41] represent reference appearance features using CLIP image embeddings, which are injected into the text-conditioned cross-attention layers of the diffusion backbone. More recently, IP-Adapter [49] introduced a novel approach where image CLIP embeddings are incorporated into the diffusion model via new cross-attention layers, which learn to predict a residual over the original cross-attention latents, as illustrated in (Figure 4 (a)). However, due to limitations in the CLIP image embeddings’ ability to capture detailed appearance information, this approach often results in noticeable identity loss and inconsistencies. The latest human image animation models [4, 16, 48, 53] have addressed these shortcomings by employing a ReferenceNet module for appearance control. ReferenceNet, a parallel and trainable duplication of the entire diffusion UNet, captures rich, detailed appearance features from a single reference image and interconnects with the diffusion UNet’s self-attention layers through feature concatenation (Fig. 4 (b)). Although this method effectively transfers appearance features to the denoising process, the full set of trainable parameters in ReferenceNet often imposes a strong and strict influence over all spatial pixels, resulting in static backgrounds and rigid dynamics, as visualized in Fig. 3.

### 3.2. Dynamics Adapter

To address the aforementioned limitations in existing reference appearance control designs, we introduce a *Dynamics-Adapter* module which effectively transfers human appearance and background context from the reference image to the diffusion backbone, without compromising its genera-

tive capability for dynamic motion synthesis. Inspired by the attention mechanism in the I2V-Adapter [9] which generates subsequent video frames from the given reference image guided by text prompts, our dynamics adapter  $\mathcal{D}$  is tailored for explicit cross-driven pose and expression control. Unlike I2V, our task accommodates motions that may differ significantly from the pose and expression of the reference image, often originating from subjects with distinct identities and body characteristics. To achieve this,  $\mathcal{D}$  is designed as a shared-weight, parallel UNet branch that injects layer-by-layer self-attention guidance of reference appearance features.

The self-attention calculation in the transformer blocks of the diffusion UNet can be represented as:

$$\mathbf{A}_i = \text{softmax}\left(\frac{\mathbf{Q}_i \mathbf{K}_i^\top}{\sqrt{d}}\right) \mathbf{V}_i, \quad (1)$$

where  $\mathbf{Q}_i, \mathbf{K}_i, \mathbf{V}_i$  are query, key, and value of the  $i$ th latent noise frame, respectively, and  $d$  is the dimension of the key and query. To introduce reference appearance guidance through our dynamics adapter, we capitalize on the prior capabilities of the original UNet to generate the key  $\mathbf{K}_R$  and value  $\mathbf{V}_R$  from the denoised latent map of the reference image  $\mathbf{I}_R$ . Additionally a trainable copy of query projector forms new query matrices  $\mathbf{Q}'_i$  from the latent noise of generation frame  $\mathbf{I}_i$ . This enables a cross-frame attention, computed as:

$$\mathbf{A}'_i = \text{softmax}\left(\frac{\mathbf{Q}'_i \mathbf{K}_R^\top}{\sqrt{d}}\right) \mathbf{V}_R. \quad (2)$$

We combine these two attention outputs with separate output projection matrices,  $\mathbf{W}_O$  and  $\mathbf{W}'_O$ , as follows:

$$\mathbf{Out}_i = (\mathbf{A}_i \mathbf{W}_O) + (\mathbf{A}'_i \mathbf{W}'_O), \quad (3)$$

where  $\mathbf{W}'_O$  is a trainable output projector.

This residual term enriches the original spatial attentions with correlated and detailed appearance information derived from the reference image. To implement this seamlessly, we initialize the query projector weights from the original UNet and zero-initialize the output projection layer  $\mathbf{W}'_O$ , ensuring that the model begins with no effect from these modifications, thus preserving its pre-existing behavior. Our design keeps the generative diffusion backbone untouched, effectively disentangling appearance control from motion generation. This separation allows the diffusion backbone to focus exclusively on pose control and dynamic synthesis, supported by ControlNet [51] and temporal modules [10], while the dynamics adapter manages appearance consistency across frames.

### 3.3. Implicit Local Face Expression Control

In human video synthesis, natural variations in facial expressions significantly enhance realism and expressiveness.

While many human image animation models offer robust control over full body poses, there has been limited efforts in simultaneously controlling facial expressions. Previous approaches to representing head motion often use simplified face landmark maps, capturing only key points such as the neck, nose, eyes, and ears. However, these simplified signals lack the detail needed for expressive facial animation. Moreover, even a basic facial skeleton encodes identity clues such as face shapes, which inadvertently influence face identity during cross-identity motion transfer. To address these limitations, we introduce **S-Face ControlNet**, a control module in addition to body control, designed for identity-disentangled control over facial expressions and head poses, enabling more expressive and adaptable human video synthesis.

Inspired by X-Portrait [46], instead of using an explicit face landmarks map from  $\mathbf{I}_i$ , we crop the face patch and utilize a pre-trained portrait reenactment network  $\mathcal{S}$  like FaceVid2Vid [42] to transfer facial expressions onto a randomly selected subject with different facial attributes. This results in an identity-swapped face patch with close expressions, which is then reinserted at the original position of  $\mathbf{I}_i$ , with other pixels masked as blank, and used as the conditional input to an additional expression ControlNet  $\mathcal{C}_F$  (Figure. 2). Unlike explicit motion control signals, this cross-identity training approach enables  $\mathcal{C}_F$  to learn identity-disentangled facial expressions and head movements implicitly from  $\mathbf{I}_T$ , reducing appearance leakage from the driving signal. Notably, we bypass the need for  $\mathcal{S}$  during inference, allowing expression control directly from the driving video.

### 3.4. Harmonic Data Fusion Training

Prior human image animation models, especially those utilizing ReferenceNet for appearance control, generally mandate static backgrounds in training videos, which limits the capture of dynamic environmental details. On the other hand, collecting video data with both moving human and dynamic backgrounds for training is challenging. We therefore introduce a mixed data training strategy, facilitating the diffusion backbone along with the temporal module to learn both human dynamics and background scene effects. Specifically, we integrate natural scene videos, such as waterfall, fireworks and wind, alongside real human motion videos for training. For videos without human, we leave the conditional inputs to the Pose ControlNet  $\mathcal{C}_P$  and S-Face ControlNet  $\mathcal{C}_F$  blank, enabling the model to generalize background motion independently. By using this mixed data, our model not only achieves more realistic dynamic details than those trained solely on human videos but also reduces unintended effects of ControlNets on background motion from the blank region of pose and expression conditional map.

Table 1. **Quantitative comparisons of X-Dyna with the recent state-of-the-art (SOTA) methods on dynamics texture generation.** A downward-pointing arrow indicates that lower values are better and vise versa. **DTFVD** [5] is calculated by replacing the FVD pre-trained backbone with one trained on DTDB [11]. **FG-DTFVD** denotes the DTFVD is running on the **foreground** parts of the videos after segmentation, and **BG-DTFVD** denotes the DTFVD of the **background** parts.

Method	FG-DTFVD ↓	BG-DTFVD ↓	DTFVD ↓
MagicAnimate [48]	<u>1.753</u>	2.142	2.601
Animate-Anyone [16]	1.789	2.034	<u>2.310</u>
MagicPose [4]	1.846	<u>1.901</u>	2.412
MimicMotion [52]	2.639	3.274	3.590
X-Dyna	<b>0.900</b>	<b>1.101</b>	<b>1.518</b>

## 4. Experiments

### 4.1. Implementation Details

**Dataset** For animation of human videos, we train our model using a custom dataset including monocular camera recordings of 30-second human motions from 107,546 videos (900 hours in total) with both indoor and outdoor scenes. All the data were processed with a cropped resolution of  $896 \times 512$ . Sequences of low quality were filtered out with [15]. All videos feature real subjects showcasing a diverse range of motions and expressions in various scenes. For data processing, we follow the approach outlined in DisCo [4, 41] but enlarge the cropping region to include the full body. For Harmonic Data Fusion Training, we use Skyscape [47] dataset, which contains 3000 time-lapse videos of dynamic sky scenes, e.g., cloudy skies and night scenes with moving stars.

**Model Training and Inference** We utilize SD 1.5 as our generative backbone, and freeze its weights during the entire training phase. Prior to training,  $\mathcal{C}_P$ ,  $\mathcal{C}_F$ , and trainable parameters in  $\mathcal{D}$  are initialized using SD 1.5, whereas the motion module is initialized with the weight of Animate-Diff v1 [10]. Our training is conducted in stages, where we first train  $\mathcal{D}$ ,  $\mathcal{C}_P$ , and motion module with Harmonic Data Fusion Training for five epochs. Then, we freeze these modules and train  $\mathcal{C}_F$  for two epochs using human video data only.

An AdamW optimizer is utilized with a learning rate of  $10^{-5}$  to train all modules. Each module undergoes training with 16 video frames in each step. During inference, we do not rely on the face-swapping network and directly feed the cropped local face patches from the driving video into S-Face ControlNet.

### 4.2. Evaluations and Comparisons

**Metrics.** We use three different groups of data for evaluation. 1) To evaluate the overall human video generation quality, we use the test set split in TikTok [17] proposed by DisCo [41], and report quantitative metrics **PSNR**, **SSIM**,

**L1**, **LPIPS**, **FID** and **cd-FVD** for human foreground generation, and **FID** and **cd-FVD** for background generation. **cd-FVD** denotes **content-debiased-FVD**, a better Frechet Video Distance (FVD) [39] metric to reflect the overall generation quality, proposed by [8]. We also report Face Cosine Similarity (**Face-Cos**) to reflect the face identity preserving ability, following MagicPose [4]. This metric is designed to gauge the model’s capability to preserve the identity information of the reference image input. To compute this metric, we first align and crop the facial region in both the generated image and the ground truth. Subsequently, we calculate the cosine similarity between the extracted feature by AdaFace [20], frame by frame of the same subject in the test set, and report the averaged value. These metrics has been widely used in previous work [4, 41, 48, 52]. In addition, we report the rate of detected faces among all frames in percentage, denoted as **Face-Det**. 2) To evaluate the dynamics detail generation quality, we use a self-collected test dataset from Pexels [29] with around 100 videos of 2 seconds each and report Dynamic Texture Frechet Video Distance (**DTFVD**) proposed by [5]. DTFVD is calculated by replacing the pre-trained backbone network in FVD with one trained on Dynamics Texture Database (DTDB) [11] for classification. This metric has also been widely used in other work [5, 24] to evaluate dynamics texture quality. We report DTFVD for both the whole videos and the background part of the videos after running human segmentation. 3) To further confirm the effectiveness of X-Dyna in dynamics detail generation, we conduct a comprehensive user study for comparison to other previous work [4, 52]. We collect 50 static real and synthetic reference images from Pexels [29] and generated by MidJourney [27] to feed the model. We then ask users to judge the (1) dynamics quality of background nature, e.g., waterfall, fireworks, cloud, ocean, raining, snowing, grass, etc. (2) dynamics quality of human foreground, e.g., hair, clothes, etc. (3) appearance and identity preservation ability.

**Quantitative Comparison** We compare our method to the state-of-the-art diffusion model-based human video animation methods, including 1) ReferenceNet-based methods Animate-Anyone [16], MagicPose [4], and MagicAnimate [48]; and 2) SVD-based method MimicMotion [52]. The main focus of this work is to improve the dynamics details generation quality. Tab. 1 presents a quantitative analysis of such quality. X-Dyna achieves significant improvements across different baseline models, indicating that the proposed method generates vivid expressiveness of dynamics. Following previous work, we used sequences 335 to 340 from the TikTok [17] dataset and additional self-collected videos by DisCo [41] to test the animation ability of human subjects. Note that the TikTok test set only contains indoor scenes, whose background is all static without any motions. Tab. 2 presents a quantitative analysis

Table 2. **Quantitative comparisons of X-Dyna with the recent SOTA methods on human video animation.** A downward-pointing arrow indicates that lower values are better and vice versa. **Face-Cos** represents the cosine similarity of the extracted feature by AdaFace [20] of face area between generation and ground truth image. **Face-Det** denotes the percentage rate of detected valid faces among all frames. \* denotes the method is not open-sourced; hence, we used the unofficial implementation from [28] to run their method for inference.

Method	Foreground								Background	
	L1 ↓	PSNR ↑	LPIPS ↓	SSIM ↑	Face-Cos ↑	Face-Det ↑	FID ↓	cd-FVD ↓	FID ↓	cd-FVD ↓
MagicAnimate [4]	<u>7.42e-05</u>	<u>17.143</u>	<b>0.228</b>	<b>0.739</b>	0.297	<u>92.1%</u>	31.97	237.59	38.86	<b>176.17</b>
Animate-Anyone* [16]	11.8e-05	13.411	0.338	0.605	<u>0.402</u>	89.0%	33.75	<u>233.39</u>	34.27	203.59
MagicPose [4]	13.7e-05	12.639	0.345	0.618	<u>0.396</u>	85.5%	<b>18.52</b>	<u>537.96</u>	<b>24.43</b>	480.14
MimicMotion [52]	9.78e-05	14.903	0.278	0.647	0.193	92.0%	45.67	<b>150.01</b>	60.32	<u>194.17</u>
X-Dyna	<b>7.15e-05</b>	<b>17.201</b>	<u>0.249</u>	<u>0.724</u>	<b>0.497</b>	<b>94.8%</b>	<u>22.56</u>	325.35	<u>25.59</u>	281.78

of human foreground subjects and background scenes from various methods, with segmentation using Segment Anything [21]. The proposed methods achieve competitive performance across previous state-of-the-art methods, which indicates that the proposed method generates high-quality videos that align with human reference.

**Qualitative Comparison** We qualitatively compare the dynamics texture generation of X-Dyna with previous methods [4, 52] in Figure 5 and pose & facial expressions control in Figure 6. Note that MagicPose, MagicAnimate, and Animate-Anyone are recent representative works using **ReferenceNet**, and MimicMotion uses **SVD** for human video animation. Both MagicPose [4] and MimicMotion [52] exhibit limited expressiveness, with most of their generated dynamics appearing almost static. Please refer to additional video examples provided in the supplementary materials for clearer observations and further comparison.

**User Study** We provide a user study to compare X-Dyna with previous works [4, 52]. We collect reference images, pose conditions, and animation results from previous works and X-Dyna of 50 subjects as mentioned in Sec. 4.1. For each subject, we visualize different human poses and facial expressions and ask 100 users to rate the methods (from 0-5) according to the following three criteria: (1) dynamics quality of background nature (BG-Dyn) (2) dynamics quality of human foreground (FG-Dyn) (3) appearance and identity preservation ability (ID). We present the result of the average vote in Table 3. We observe that the user prefers X-Dyna more than ReferenceNet and SVD based works [4, 52], especially in terms of dynamic texture generation. More details can be found in the supplementary material.

### 4.3. Ablation Analysis

In this section, a comprehensive ablation analysis of X-Dyna is presented. To confirm the effectiveness of our Dynamics-Adapter and Harmonic Data Fusion Training, we quantitatively evaluate the DTFVD on our self-collected data and present the result in Tab. 4. Note that each row indicates individual ablations of each module. More visual-

Table 3. **User study of X-Dyna.** We collect the ratings (0-5) from 100 participants for 50 test cases in the test set. We ask them to rate the generation in terms of Foreground Dynamics (**FG-Dyn**), Background Dynamics (**BG-Dyn**) and Identity Preserving (**ID**).

Method	FG-Dyn	BG-Dyn	ID	Overall
MagicAnimate [48]	2.34	2.78	3.45	2.86
Animate-Anyone [16]	2.21	2.57	<u>3.89</u>	<u>2.89</u>
MagicPose [4]	2.23	2.18	3.85	2.75
MimicMotion [52]	2.02	2.63	2.79	2.48
X-Dyna	<b>3.87</b>	<b>4.26</b>	<b>4.14</b>	<b>4.09</b>

Table 4. **Ablation Analysis of X-Dyna on dynamics texture generation and local facial expressions generation.** w/RefNet denotes Dynamics-Adapter is replaced by a ReferenceNet. w/IP-A denotes Dynamics-Adapter is replaced by an IP-Adapter. w/lmk denotes S-Face ControlNet is not used for fine-tuning and face landmarks are used together with the pose skeleton. w/face denotes S-Face ControlNet is not used for fine-tuning. w/fusion denotes Harmonic Data Fusion Training is not used for disentangled animation.

Method	FG-DTFVD ↓	BG-DTFVD ↓	DTFVD ↓	Face-Cos ↑
w/RefNet	2.137	2.694	2.823	0.466
w/IP-A	3.738	4.702	4.851	0.292
w/lmk	0.914	1.125	1.589	0.406
wo/face	<u>0.912</u>	<b>1.098</b>	<u>1.550</u>	0.442
wo/fusion	1.301	1.467	1.652	0.495
X-Dyna	<b>0.900</b>	<u>1.101</u>	<b>1.518</b>	<b>0.497</b>

izations can be found on our supplementary website.

### 4.4. Limitations and Future Works

Despite its effectiveness in the dynamic expressiveness generation of human video animation, our X-Dyna has certain limitations, particularly in scenarios where the target pose significantly deviates from the reference human. For instance, during extreme zooming in or out, the appearance and identity may not be perfectly preserved. Additionally, our method struggles to generate perfect hand poses. We believe that these challenges can be addressed by collecting more high-quality data and employing advanced hand pose representations as input.

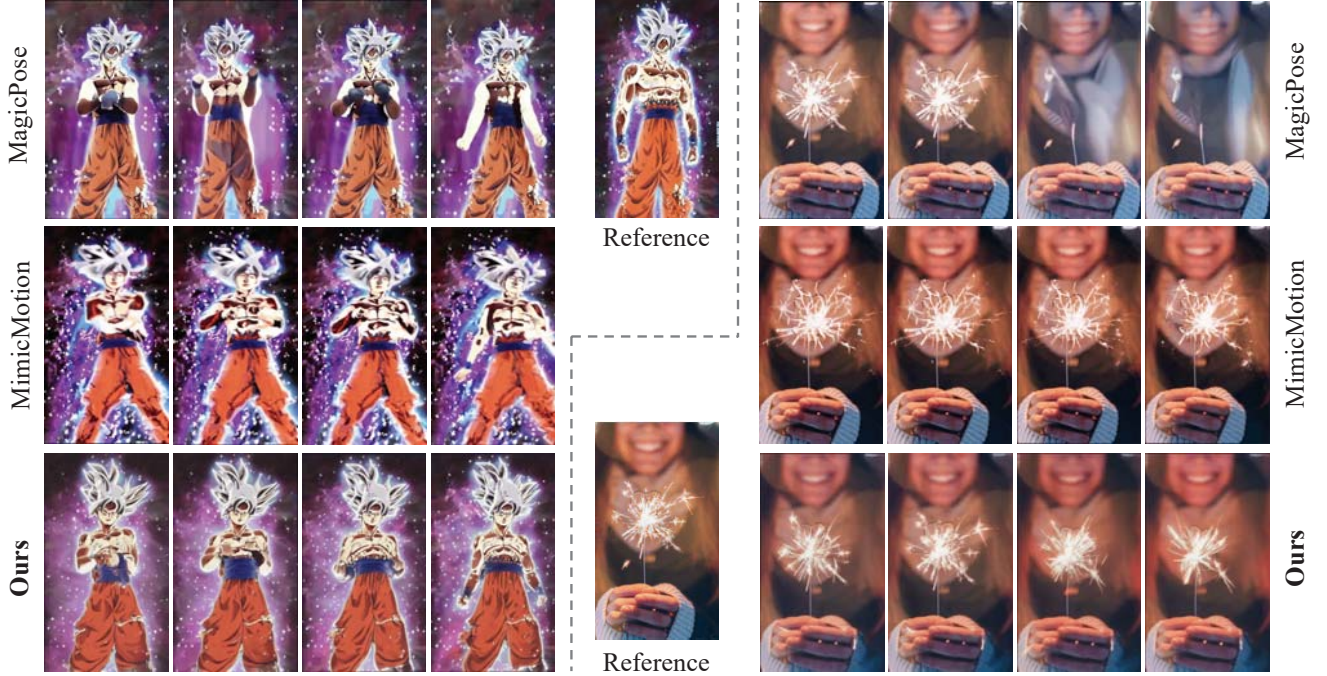


Figure 5. **Qualitative Comparison on Human in Dynamic Scene.** While existing SOTA methods struggle to generate consistent and realistic scene dynamics involving humans, our method successfully produces dynamic human-scene interactions while preserving the structure of the reference image.



Figure 6. **Qualitative Comparison on Poses and Face Expressions Control.** We show each method on test cases using the same reference image and pose skeleton. For improved visualization, a zoomed-in view of the face area is also provided. Our method produces results that most closely match the ground truth and best preserve face identity.

In the future, we will explore applying Dynamics-Adapter to more powerful base image and video diffusion models, such as SVD [1], SDXL [31], Stable Diffusion 3 [6], and Flux [23] to achieve better performance. Moreover, we will investigate adding the camera trajectory or drag control proposed in [12, 22, 44, 50] to our model so that we have a more user-friendly condition.

## 5. Conclusion

In this work, we propose *X-Dyna*, a photorealistic human video animation pipeline with the ability of consistent motion control and vivid dynamics details generation. We propose an efficient *Dynamics-Adapter* module to preserve the human appearance reference while maintaining the foundation model’s ability to generate high-quality dynamics.

To boost the dynamics modeling capability further, we propose a *Harmonic Data Fusion Training* strategy, mixing the training data from real-human and natural scene videos. Moreover, we incorporate two plug-in modules, an *S-Face ControlNet* for facial expressions editing and a *Pose ControlNet* for human motion control. Finally, all proposed modules can be treated as extensions to SD and used for customized pre-trained weights of SD-UNet. Extensive evaluation of various models also validates the effectiveness and generalizability of our model.

**Ethics Statement.** Our work aims to improve human image animation from a technical perspective and is not intended for malicious use like fake videos. Therefore, synthesized videos should clearly indicate their artificial nature.

## Acknowledgement

Soleymani's work was sponsored by the Army Research Office and was accomplished under Cooperative Agreement Number W911NF-20-2-0053. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

## References

- [1] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023. 2, 3, 8
- [2] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22563–22575, 2023. 2
- [3] 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. *arXiv preprint arXiv:2304.08465*, 2023. 2
- [4] Di Chang, Yichun Shi, Quankai Gao, Hongyi Xu, Jessica Fu, Guoxian Song, Qing Yan, Yizhe Zhu, Xiao Yang, and Mohammad Soleymani. Magicpose: Realistic human poses and facial expressions retargeting with identity-aware diffusion. In *Forty-first International Conference on Machine Learning*, 2023. 1, 2, 3, 4, 6, 7
- [5] Michael Dorkenwald, Timo Milbich, Andreas Blattmann, Robin Rombach, Konstantinos G Derpanis, and Bjorn Ommer. Stochastic image-to-video synthesis using cinns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3742–3753, 2021. 6
- [6] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024. 8
- [7] Haiwen Feng, Zheng Ding, Zhihao Xia, Simon Niklaus, Victoria Abrevaya, Michael J Black, and Xuaner Zhang. Explorative inbetweening of time and space. *arXiv preprint arXiv:2403.14611*, 2024. 3
- [8] Songwei Ge, Aniruddha Mahapatra, Gaurav Parmar, Jun-Yan Zhu, and Jia-Bin Huang. On the content bias in fr chet video distance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 6
- [9] Xun Guo, Mingwu Zheng, Liang Hou, Yuan Gao, Yufan Deng, Chongyang Ma, Weiming Hu, Zhengjun Zha, Haibin Huang, Pengfei Wan, et al. I2v-adapter: A general image-to-video adapter for video diffusion models. *arXiv preprint arXiv:2312.16693*, 2023. 5
- [10] Yuwei Guo, Ceyuan Yang, Anyi Rao, Yaohui Wang, Yu Qiao, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. *arXiv preprint arXiv:2307.04725*, 2023. 2, 3, 4, 5, 6
- [11] Isma Hadji and Richard P Wildes. A new large scale dynamic texture dataset with application to convnet understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 320–335, 2018. 6
- [12] Hao He, Yinghao Xu, Yuwei Guo, Gordon Wetzstein, Bo Dai, Hongsheng Li, and Ceyuan Yang. Cameractrl: Enabling camera control for text-to-video generation. *arXiv preprint arXiv:2404.02101*, 2024. 8
- [13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *arXiv preprint arXiv:2006.11239*, 2020. 3, 4
- [14] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–8646, 2022. 3
- [15] Vlad Hosu, Hanhe Lin, Tamas Sziranyi, and Dietmar Saupe. Koniq-10k: An ecologically valid database for deep learning of blind image quality assessment. *IEEE Transactions on Image Processing*, 29:4041–4056, 2020. 6
- [16] Li Hu. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8153–8163, 2024. 1, 2, 3, 4, 6, 7
- [17] Yasamin Jafarian and Hyun Soo Park. Learning high fidelity depths of dressed humans by watching social media dance videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12753–12762, 2021. 3, 6
- [18] Johanna Karras, Aleksander Holynski, Ting-Chun Wang, and Ira Kemelmacher-Shlizerman. Dreampose: Fashion image-to-video synthesis via stable diffusion. *arXiv preprint arXiv:2304.06025*, 2023. 3
- [19] Tero Karras. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 3
- [20] Minchul Kim, Anil K Jain, and Xiaoming Liu. Adaface: Quality adaptive margin for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 6, 7
- [21] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Doll r, and Ross Girshick. Segment anything. *arXiv:2304.02643*, 2023. 7
- [22] Zhengfei Kuang, Shengqu Cai, Hao He, Yinghao Xu, Hongsheng Li, Leonidas Guibas, and Gordon Wetzstein. Collaborative video diffusion: Consistent multi-video generation with camera control. *arXiv preprint arXiv:2405.17414*, 2024. 8
- [23] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024. 8

- [24] Zhengqi Li, Richard Tucker, Noah Snavely, and Aleksander Holynski. Generative image dynamics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24142–24153, 2024. 3, 6
- [25] Jinlin Liu, Kai Yu, Mengyang Feng, Xiefang Guo, and Miaomiao Cui. Disentangling foreground and background motion for enhanced realism in human video generation. *arXiv preprint arXiv:2405.16393*, 2024. 3
- [26] Yue Ma, Yingqing He, Xiaodong Cun, Xintao Wang, Siran Chen, Xiu Li, and Qifeng Chen. Follow your pose: Pose-guided text-to-video generation using pose-free videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4117–4125, 2024. 3
- [27] Midjourney. midjourney. <https://www.midjourney.com>, 2024. 3, 6
- [28] MooreThreads. Moorethreads/moore-animateanyone. <https://github.com/MooreThreads/Moore-AnimateAnyone>., 2024. 7
- [29] Pexels. pexels. <https://www.pexels.com/>, 2024. 3, 6
- [30] Ryan Po, Wang Yifan, Vladislav Golyanik, Kfir Aberman, Jonathan T Barron, Amit Bermano, Eric Chan, Tali Dekel, Aleksander Holynski, Angjoo Kanazawa, et al. State of the art on diffusion models for visual computing. In *Computer Graphics Forum*, page e15063. Wiley Online Library, 2024. 3
- [31] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 8
- [32] 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. 3
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 3
- [34] Masaki Saito, Eiichi Matsumoto, and Shunta Saito. Temporal generative adversarial nets with singular value clipping. In *Proceedings of the IEEE international conference on computer vision*, pages 2830–2839, 2017. 3
- [35] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 3
- [36] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 3
- [37] Zhengyan Tong, Chao Li, Zhaokang Chen, Bin Wu, and Wenjiang Zhou. Musepose: a pose-driven image-to-video framework for virtual human generation. *arxiv*, 2024. 1, 2, 3, 4
- [38] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1526–1535, 2018. 3
- [39] Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018. 6
- [40] Qilin Wang, Zhengkai Jiang, Chengming Xu, Jiangning Zhang, Yabiao Wang, Xinyi Zhang, Yun Cao, Weijian Cao, Chengjie Wang, and Yanwei Fu. Vividpose: Advancing stable video diffusion for realistic human image animation. *arXiv preprint arXiv:2405.18156*, 2024. 1, 3
- [41] Tan Wang, Linjie Li, Kevin Lin, Chung-Ching Lin, Zhengyuan Yang, Hanwang Zhang, Zicheng Liu, and Lijuan Wang. Disco: Disentangled control for referring human dance generation in real world. *arXiv preprint arXiv:2307.00040*, 2023. 3, 4, 6
- [42] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021. 5
- [43] Huawei Wei, Zejun Yang, and Zhisheng Wang. Aniportrait: Audio-driven synthesis of photorealistic portrait animation. *arXiv preprint arXiv:2403.17694*, 2024. 3
- [44] Weijia Wu, Zhuang Li, Yuchao Gu, Rui Zhao, Yefei He, David Junhao Zhang, Mike Zheng Shou, Yan Li, Tingting Gao, and Di Zhang. Draganything: Motion control for anything using entity representation. In *European Conference on Computer Vision*, pages 331–348. Springer, 2025. 8
- [45] Zhiqiang Xia, Zhaokang Chen, Bin Wu, Chao Li, Kwok-Wai Hung, Chao Zhan, Yingjie He, and Wenjiang Zhou. Musev: Infinite-length and high fidelity virtual human video generation with visual conditioned parallel denoising. *arxiv*, 2024. 3
- [46] You Xie, Hongyi Xu, Guoxian Song, Chao Wang, Yichun Shi, and Linjie Luo. X-portrait: Expressive portrait animation with hierarchical motion attention. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 5
- [47] Wei Xiong, Wenhan Luo, Lin Ma, Wei Liu, and Jiebo Luo. Learning to generate time-lapse videos using multi-stage dynamic generative adversarial networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 6
- [48] Zhongcong Xu, Jianfeng Zhang, Jun Hao Liew, Hanshu Yan, Jia-Wei Liu, Chenxu Zhang, Jiashi Feng, and Mike Zheng Shou. Magicanimate: Temporally consistent human image animation using diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1481–1490, 2024. 1, 2, 3, 4, 6, 7
- [49] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapt: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*, 2023. 4
- [50] Shengming Yin, Chenfei Wu, Jian Liang, Jie Shi, Houqiang Li, Gong Ming, and Nan Duan. Dragnuwa: Fine-grained control in video generation by integrating text, image, and trajectory. *arXiv preprint arXiv:2308.08089*, 2023. 8
- [51] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. 2, 3, 4, 5

- [52] Yuang Zhang, Jiayi Gu, Li-Wen Wang, Han Wang, Junqi Cheng, Yuefeng Zhu, and Fangyuan Zou. Mimicmotion: High-quality human motion video generation with confidence-aware pose guidance. *arXiv preprint arXiv:2406.19680*, 2024. [1](#), [2](#), [3](#), [6](#), [7](#)
- [53] Shenhao Zhu, Junming Leo Chen, Zuozhuo Dai, Yinghui Xu, Xun Cao, Yao Yao, Hao Zhu, and Siyu Zhu. Champ: Controllable and consistent human image animation with 3d parametric guidance. *arXiv preprint arXiv:2403.14781*, 2024. [1](#), [3](#), [4](#)