

# MultiAnimate: Pose-Guided Image Animation Made Extensible

Yingcheng Hu<sup>1,2,3\*</sup> Haowen Gong<sup>3</sup> Chuanguang Yang<sup>1</sup> Zhulin An<sup>1,†</sup>  
Yongjun Xu<sup>1</sup> Songhua Liu<sup>3,†</sup>

<sup>1</sup>State Key Laboratory of AI Safety, Institute of Computing Technology,  
Chinese Academy of Sciences <sup>2</sup>ShanghaiTech University <sup>3</sup>Shanghai Jiao Tong University

<sup>†</sup>Corresponding Authors



Figure 1. **Multi-character pose-guided image animation generated by our framework.** Our method performs multi-character image animation with consistent identity and appearance for each character. Notably, our framework, trained only on two-character data, is capable of producing identity-consistent three-person videos and can, in principle, be extended to scenarios with even more participants (e.g., seven characters).

## Abstract

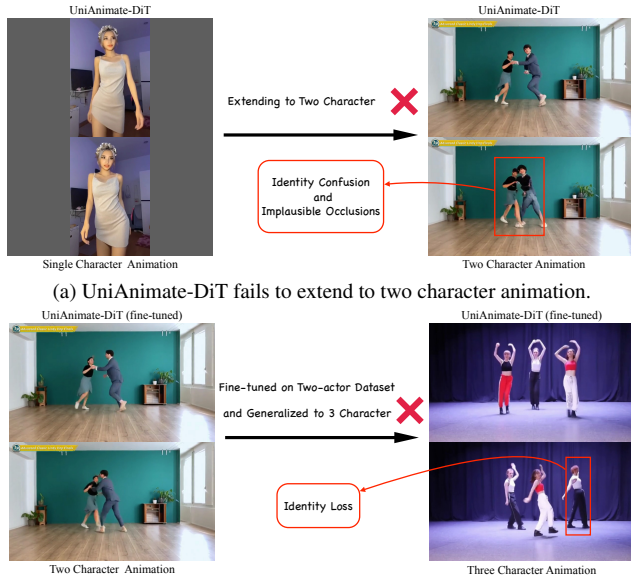
Pose-guided human image animation aims to synthesize realistic videos of a reference character driven by a sequence of poses. While diffusion-based methods have achieved remarkable success, most existing approaches are limited to single-character animation. We observe that naively extending these methods to multi-character scenarios often leads to identity confusion and implausible occlusions between characters. To address these challenges, in this paper, we propose an extensible multi-character image animation framework built upon modern Diffusion Transformers (DiTs) for video generation. At its core, our framework introduces two novel components—Identifier Assigner and Identifier Adapter—which collaboratively capture per-

son positional cues and inter-person spatial relationships. This mask-driven scheme, along with a scalable training strategy, not only enhances flexibility but also enables generalization to scenarios with more characters than those seen during training. Remarkably, trained on only a two-character dataset, our model generalizes to multi-character animation while maintaining compatibility with single-character cases. Extensive experiments demonstrate that our approach achieves state-of-the-art performance in multi-character image animation, surpassing existing diffusion-based baselines. Video demos are available at <https://hyc001.github.io/MultiAnimate/>.

## 1. Introduction

Diffusion-based video generation [2, 10, 14, 18, 34] has witnessed remarkable advances in recent years, providing

\*Work done during a research internship at the School of Artificial Intelligence, Shanghai Jiao Tong University.



(a) UniAnimate-DiT fails to extend to two character animation.  
 (b) Fine-tuned UniAnimate-DiT fails to generalize to cases with unseen numbers of characters.  
 Figure 2. Dilemmas of current methods in multi-character image animation.

powerful tools for creating dynamic visual content. Within this trend, controllable video generation [19, 36, 42, 46, 48] has emerged as a major research focus, as users increasingly seek to create visual content with customized layouts, motions, or identities. A representative task is human image animation, which aims to synthesize a video of a reference character driven by a sequence of poses. This task has attracted significant interest due to its widespread applications, such as film production [39] and digital avatars [40].

Human image animation generally imposes two core requirements: (i) maintaining the consistent identity of the generated person with the given reference character and (ii) aligning the generated motion with the input pose sequence. Various works [17, 19, 36, 48] have been proposed to address these issues and have made impressive progress. However, most existing methods primarily focus on single-character animation, while its natural extension—multi-character animation—remains underexplored.

In fact, although this extension seems conceptually straightforward, it poses significantly greater challenges than it appears. On the one hand, it requires preserving per-character consistency while faithfully modeling the spatial interactions between multiple characters. However, as shown in Fig. 2a, we observe that naively extending single-character animation methods to multi-character scenarios often fails to meet these requirements, leading to identity confusion and implausible occlusions. On the other hand, scalability remains a major challenge. Although some prior works consider multi-character scenarios, existing human image animation models are typically trained with a fixed

number of participants, *e.g.*, two-person animation trained on two-person datasets. As illustrated in Fig. 2b, such models struggle to generalize to cases with unseen numbers of characters. Extending them to more participants requires collecting corresponding multi-person data and retraining or adapting the models accordingly—a process that is highly cumbersome if not infeasible at all.

These observations raise a key question: *Can we design a human image animation framework that is both robust and generalizable?* To this end, in this paper, we propose **MultiAnimate**, a novel multi-actor image animation framework that (1) preserves multi-character consistency and visual fidelity under complex interactions, and (2) exhibits strong scalability to scenarios with more characters than those seen during training.

At the core of our approach lies an effective mechanism that resolves the inherently underdetermined association between poses and identities during spatial interactions, as illustrated in Fig. 3. Specifically, we introduce two key modules: Identifier Assigner and Identifier Adapter—which assign a unique identifier to each character across all frames and embed this identity cue into the feature space of the Diffusion Transformer (DiT) backbone. This explicit identity encoding enables consistent character tracking throughout the video, effectively mitigating confusion in complex multi-character interactions.

While the identifier-based strategy effectively resolves character confusion, a remaining challenge lies in enabling these identifiers to generalize to unseen numbers of characters. To address this problem, we devise a scalable training strategy that randomly samples identifiers from a learnable embedding space at each iteration. By the time training converges, all identifiers have been explicitly trained to be mutually distinguishable. As a result, during inference, even when more characters, along with more identifiers, are introduced than were present during training, the model can naturally differentiate them and maintain consistent character identities. In our experiments, we demonstrate that training solely on two-character data is sufficient to perform multi-character image animation while maintaining compatibility with single-character cases.

Moreover, to further enhance the visual quality of generated videos, we leverage modern video generation models, such as Wan 2.2 [34], to synthesize high-quality multi-character training data as an additional supplement. Extensive experiments demonstrate that our approach surpasses existing diffusion-based methods in multi-character image animation while preserving competitive performance in single-character scenarios. In summary, our contributions are threefold:

- Overall, we propose **MultiAnimate** for multi-character image animation, which is the first extensible framework built upon modern DiT-based video generators, to the

best of our knowledge.

- Technically, we (i) introduce two novel modules—Identifier Assigner and Identifier Adapter—to jointly and explicitly model per-person spatial features and inter-person interactions, and (ii) propose a scalable training strategy that allows our model to handle multi-character image animation with more characters than those present in the training data.
- Empirically, we (i) curate a new high-quality dataset for our experiments to overcome the limitations of prior training data in frame quality, and (ii) perform experiments demonstrating that our framework achieves superior performance in multi-character video generation while remaining competitive in single-character scenarios.

## 2. Related Works

**Diffusion models for video generation.** Diffusion models [6–8, 12, 13, 23, 27, 30, 31, 38] have rapidly advanced video generation [14, 25] in recent years. Many approaches build upon UNet [28], adding temporal modules to image diffusion models [11, 26, 29] for spatio-temporal modeling [2, 10]. More recently, inspired by the success of DiT [24] in image generation, some works [5, 9, 20, 21] replace the UNet with DiT blocks to further improve generation quality. 3D VAE architectures are introduced in Hunyuan-Video [18] and Wan-Video [34] to enhance video compression and reconstruction. These DiT-based frameworks flexibly inject conditions via cross-attention or feature concatenation, and have become increasingly popular foundations for controllable video generation upon which our work is built.

**Pose-guided human image animation.** Pose-guided human image animation typically uses skeleton poses [3, 45] to drive the motion of a target character. Current works introduce additional conditioning signals to improve controllability and visual quality. Champ [49] and Vidvid-Pose [35] incorporate depth and 3D body priors such as SMPL to enhance controllable motion synthesis. StableAnimator [32] introduces a dedicated face encoder combined with face masks to better preserve facial identity. However, these approaches mainly focus on single-character animation. Recently, Follow-Your-Pose v2 [41] leverages optical flow and depth cues to better handle occlusions between interacting people. DanceTogether [4] jointly exploits mask and pose features to model spatial relationships between various characters. Unlike them, we target a scalable multi-character image animation framework that can generalize to scenarios with more characters than those seen in training.

## 3. Methods

For pose-guided multi-character image animation, given a reference image  $I_{\text{ref}} \in \mathbb{R}^{3 \times H \times W}$ , a driving pose sequence




Figure 3. In multi-character image animation, identical pose sequences can lead to multiple plausible motion trajectories.

$P \in \mathbb{R}^{T \times 3 \times H \times W}$ , and a set of per-person tracking masks  $\{M_i\}_{i=1}^n$ ,  $M_i \in \mathbb{R}^{T \times 1 \times H \times W}$ , where  $T$  denotes the number of frames and  $n$  is the number of tracked persons, our goal is to generate a target video  $V_{\text{tar}} \in \mathbb{R}^{T \times 3 \times H \times W}$ , that (1) preserves each identity in  $I_{\text{ref}}$ , (2) aligns with the motion in pose sequence  $P$ , and (3) aligns with the spatial relationships among per-person tracking masks  $\{M_i\}_{i=1}^n$ .

In Sec. 3.1, we present the design of our pipeline, which models per-person positional cues and inter-person spatial relationships through the proposed Identifier Assigner and Identifier Adapter. In Sec. 3.2, we describe our training strategy that enables multi-character image animation with more characters than those seen during training.

### 3.1. MultiAnimate

In this section, we first analyze the problem in pose-guided multi-character image animation and present our pipeline.

#### 3.1.1. A Multi-Character Motion Ambiguity

Consider simply extending single-character to multi-character pose-guided image animation. In single-character setting, the mapping between the reference image and the pose sequence is unique. However, when multiple persons are generated simultaneously, an ambiguity arises. As illustrated in Fig. 3, the two people first rotate  $180^\circ$  clockwise and exchange their positions. From this intermediate state, there exist two visually plausible continuations: (1) both characters continue rotating clockwise by another  $180^\circ$  to return to their initial positions, or (2) both characters rotate counterclockwise by  $180^\circ$  back to the same configuration. These two trajectories can be yielded from nearly identical pose sequences, making the output non-unique unless additional spatial cues are provided. Therefore, in multi-character pose-guided image animation, per-person tracking masks are necessary to fix a unique correspondence between each character and its motion trajectory.

A straightforward way to incorporate these cues is to extract pose and tracking mask features for each person individually, bind them together, and then aggregate all per-person representations by summation, as done in concurrent

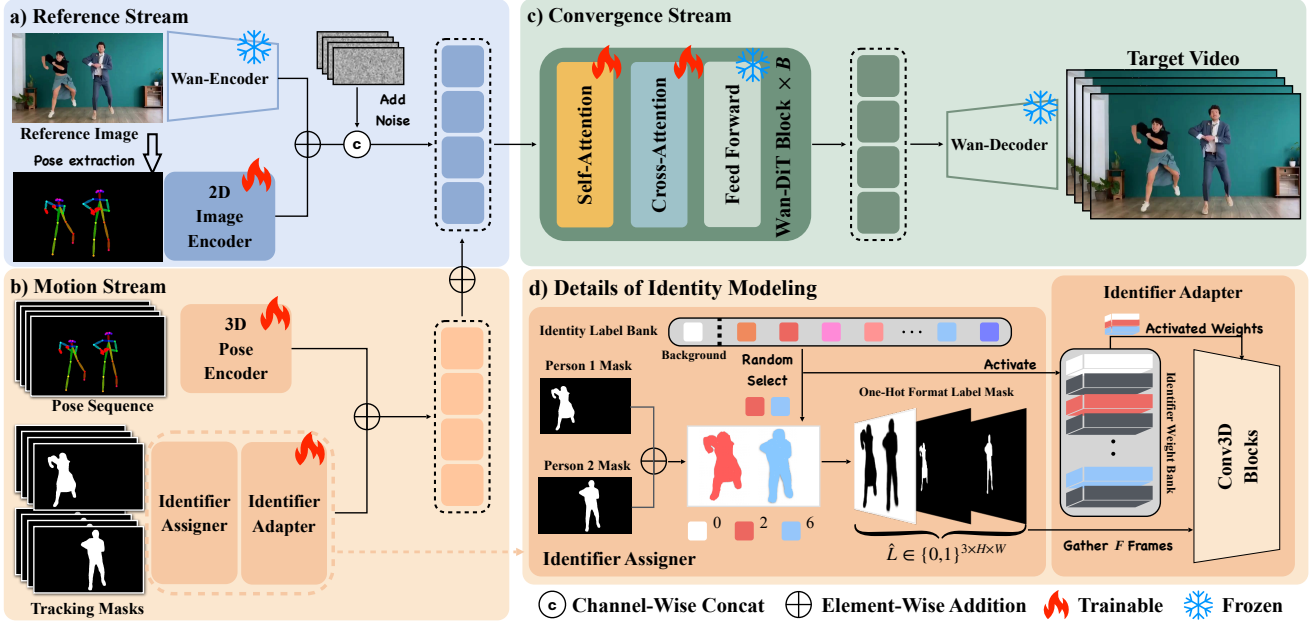


Figure 4. **Overview of our framework.** Our pipeline contains two main streams: the *reference stream*, which encodes the reference image and its pose to capture appearance information, and the *motion stream*, which encodes multi-character pose sequences and tracking masks to model motion and spatial conditions. The two streams are fused through element-wise addition of latent tokens. The Identifier Assigner unifies per-person tracking masks into a structured label representation, preserving spatial relationships and interactions among multiple characters. This representation is converted to the feature space of the DiT backbone by the Identifier Adapter.

work [4]. However, such a design requires additional data processing to extract a pose sequence for each character, and direct summation cannot ensure that spatial relationships between different people are adequately preserved.

### 3.1.2. Proposed Pipeline

**Data flow.** Following previous work [36], we build our framework upon the Wan2.1 [34] I2V architecture. As illustrated in Fig. 4, our pipeline consists of two main processing streams: one for encoding the reference appearance and the other for encoding the motion and spatial conditions.

- **Reference stream.** The reference image  $I_{\text{ref}} \in \mathbb{R}^{3 \times H \times W}$  is first passed through a VAE Encoder to obtain its latent representation. The reference pose extracted from  $I_{\text{ref}}$  is processed by an Image Encoder composed of stacked 2D convolutional layers. Their outputs are added, concatenated with random noise, and further patchified to serve as the input tokens.
- **Motion stream.** The pose sequence  $P \in \mathbb{R}^{T \times 3 \times H \times W}$  is encoded by a 3D convolutional Pose Encoder to capture temporal motion dynamics. Meanwhile, the per-person tracking masks  $\{M_i\}_{i=1}^n$  are first processed by the Identifier Assigner to unify and encode all masks into a structured label representation. The resulting feature maps are then fed into the Identifier Adapter, which models per-person positional features and inter-person spatial relationships. The encoded pose and mask features are first added and then fused with patchified latent tokens from

the reference stream through element-wise addition.

**Identifier Assigner.** Given the set of per-person tracking masks  $\{M_i\}_{i=1}^n$ , we first perform a preprocessing step, referred to as the Identifier Assigner in Fig. 4. This module unifies all per-person masks into a single label map  $\mathcal{L} \in \mathbb{R}^{H \times W}$ , where background pixels are labeled as 0. Consider the case where only a two-character dataset is available. For the two characters A and B, all pixels in their tracking masks with values greater than 0 are assigned labels  $a$  and  $b$  in  $\mathcal{L}$ , respectively, where  $a$  and  $b$  are two distinct non-zero identifiers from Identity Label Bank. As a result, we obtain a label map with values in  $\{0, a, b\}$ , denoted as  $\mathcal{L} \in \{0, a, b\}^{H \times W}$ . We then apply one-hot encoding to  $\mathcal{L}$ , producing a binary tensor  $\hat{\mathcal{L}} \in \{0, 1\}^{3 \times H \times W}$ , where the three channels correspond to the background and the two characters, each encoding its spatial occupancy.

The Identifier Adapter, consisting of stacked 3D convolutional layers, then takes  $\hat{\mathcal{L}}$  as input. The design of the Identifier Adapter, together with the associated training strategy, will be detailed in the following subsection. Our design preserves spatial relationships among multiple characters and captures interpersonal interactions, such as proximity and occlusion. This provides an explicit spatial prior for generating videos of multiple people.

### 3.2. Training Strategy

In this section, we first discuss the challenges encountered during training and then present our strategy.




Figure 5. Our framework performs well at early training stages, but inconsistencies emerge when the person-assigned labels at inference differ from those seen during training.

### 3.2.1. Symmetry Issue and Generalization Ability

Our pipeline can achieve promising results at early stage of training, as shown in Fig. 5. However, a *symmetry problem* arises during training. Specifically, if the person-assigned labels during inference differ from those used in training, the generated results become inconsistent. This suggests that during training, each person’s position is associated with a particular channel in the Identifier Adapter, as opposed to being linked to its corresponding tracking mask. Such channel-identity binding contradicts our design goal, which is to let the model associate each person with its tracking mask, not with a fixed label or channel.

Additionally, current pose-guided human image animation frameworks are typically trained with a fixed number of participants, *e.g.*, two-person animation trained on two-person datasets. As collecting multi-person data is costly, such a setup severely limits generalization of existing frameworks to more complex multi-character scenarios. We target a flexible framework that generalizes to a varying number of participants, even with more characters than those seen during training.

### 3.2.2. Proposed Strategy

Due to our mask-driven design, the above two goals can be effectively achieved. Consider that the model is expected to support up to  $n$ -character image animation during inference. In the first Conv3D layer of the Identifier Adapter, we introduce an Identifier Weight Bank storing  $n$  identifier channels. Under a two-character training setting, only three channels, including identifiers for the two characters and the background, are activated in each training iteration.

Specifically, in each training iteration, we randomly assign identity labels to the two characters from the Identity Label Bank of size  $n$ . These assigned labels are also used to activate the corresponding weights in the Identifier Weight Bank. By the time training converges, although the training dataset contains only two-character videos, all the  $n$  channels in the Identity Weight Bank are seen and trained to be mutually distinct. In this way, our framework can robustly generate videos with up to  $n$  characters during inference. Moreover, since labels are randomly assigned, the

network learns to associate each character with its spatial mask rather than a fixed channel, which helps alleviate the symmetry problem.

## 4. Experiments

**Datasets.** We used the Swing Dance dataset [22], which comprises approximately 30 hours of swing dance videos featuring 680 unique couple subjects. This dataset is utilized to train the model to learn complex interactions and spatial relationships between individuals. Also, we generated 2079 five-second videos using the Wan 2.2 Video Generator [34], denoted as the *Gen-dataset*. It contains various scenes with two or three characters and is primarily used to assess the adaptability of the model to varying scenarios. Optionally, it can also be incorporated into training to further enhance robustness of the model. In addition, the TiKToK dataset [16], along with over 50 minutes of unseen dance videos featuring 3-7 characters from social media, serves solely to evaluate the generalization capability of our framework.

**Evaluation Metrics and Baselines.** The quality of generated videos is evaluated using FVD [33] and FID-VID [1]. To further evaluate the frame generation quality, we report frame-wise PSNR [15], SSIM [37], L1 loss and LISPIS [47]. We compare our method with other open-source pose-guided human image animation models, including UniAnimate-DiT [36], VACE [17], MimicMotion [48], and DisPose [19].

**Implementation Details.** We utilize DWPose [43] and Sa2VA [44] to extract skeleton poses and tracking masks. We initialized the Image Encoder, Pose Encoder and Lora with pre-trained weights in UniAnimate-DiT [36]. Our pipeline is first trained on the training split of the Swing Dance dataset for 40 epochs, 7000 training steps, and supports up to 3-character animation, referred to as the *Stage 1 model*. Followed by 3 epochs, totally 2400 steps training on the training split of the Gen-dataset, referred to as the *Stage 2 model*. Additionally, we extended the *Stage 1 model* to support up to 7-character image animation. This version was trained on the training split of the Swing Dance dataset for 24 epochs (4,200 steps) and is referred to as the *Extended model*. The training is performed on 2 NVIDIA A100 80GB GPUs with a batch size of 1 per GPU and the learning rate is set to  $1 \times 10^{-4}$ .

### 4.1. Comparison with Concurrent Work

**Quantitative Results.** We evaluate *Stage 1 model* and *Extended model*, which are trained only on the training split of Swing Dance dataset. As shown in Tab. 1, we compare our method with state-of-the-art approaches on the test splits of the Swing Dance dataset, the Gen-dataset, and on unseen dance videos.

Table 1. Quantitative comparison with other SOTA methods on the test split of Swing Dance, Gen-dataset and Unseen dance videos.

Dataset	Method	PSNR $\uparrow$	SSIM $\uparrow$	L1(E-3) $\downarrow$	LPIPS $\downarrow$	FVD $\downarrow$	FID-VID $\downarrow$
Swing Dance	Mimicmotion [48]	13.47	0.390	145.10	0.501	780.31	49.69
	DisPose [19]	14.08	0.444	137.69	0.466	<u>746.29</u>	46.24
	UniAnimate-DiT [36]	<u>16.15</u>	<u>0.619</u>	<u>98.72</u>	<u>0.427</u>	891.89	<u>27.71</u>
	VACE [17]	11.15	0.311	216.98	0.563	763.75	29.88
	Ours ( <i>Stage 1 model</i> )	<b>19.40</b>	<b>0.687</b>	<b>64.94</b>	<b>0.335</b>	<b>648.84</b>	<b>22.50</b>
Gen-dataset	Mimicmotion [48]	11.96	0.286	173.14	0.444	865.36	86.99
	DisPose [19]	15.33	0.488	105.41	0.266	712.61	53.70
	UniAnimate-DiT [36]	<u>18.10</u>	<u>0.630</u>	<u>80.52</u>	0.262	<u>565.50</u>	<b>32.10</b>
	VACE [17]	16.39	0.559	97.14	<b>0.231</b>	657.23	38.86
	Ours ( <i>Stage 1 model</i> )	<b>19.05</b>	<b>0.659</b>	<b>73.01</b>	<u>0.251</u>	<b>454.92</b>	<u>33.57</u>
Unseen dance videos	Mimicmotion [48]	13.53	0.376	145.73	0.381	1634.12	145.67
	DisPose [19]	15.88	0.694	92.63	0.304	1133.57	97.23
	UniAnimate-DiT [36]	<u>17.94</u>	<u>0.751</u>	<u>70.07</u>	0.286	<u>624.45</u>	<u>71.24</u>
	VACE [17]	17.24	0.714	84.34	<u>0.279</u>	922.66	78.93
	Ours ( <i>Extended model</i> )	<b>23.24</b>	<b>0.857</b>	<b>33.40</b>	<b>0.185</b>	<b>358.74</b>	<b>43.12</b>




Figure 6. **Quality comparison with state-of-the-art methods on the Swing Dance dataset.** More examples are provided in the supplementary material.

On the Swing Dance dataset, which emphasizes complex motion interactions between two characters, our method consistently achieves better scores than competing methods for all metrics. The advantage in FVD and FID-VID indicates that our approach produces temporally coherent and perceptually plausible videos. Other evaluation metrics suggest that our model produce high frame-wise image quality.

On the Gen-dataset, where motions are relatively simple but scenes are more diverse, our model-without any additional training on this dataset-still achieves strong perfor-

mance. A slight disadvantage in LPIPS and FID-VID may suggest our method is less adapted to highly diverse environments.

On the unseen dance videos, where character interactions are more frequent and complex and the scenes are more challenging. Nevertheless, our method maintains a clear advantage, demonstrating strong generalization ability from two-actor training data to more complex multi-actor scenarios.

**Qualitative Results.** We present several visual results from the Swing Dance dataset. As shown in Fig. 6, both DisPose




Figure 7. Training on the Gen-dataset enhances the models ability to maintain temporal consistency and adapt to diverse motionscene interaction.

and MimicMotion suffer identity confusion at the beginning of image animation, and the background appear blurry. Although UniAnimate-DiT and VACE maintain better identity consistency in the early stage, they fail to preserve it after complex inter-person interactions, leading to identity swaps and unnatural occlusions. In contrast, our method produces visually coherent results, maintaining consistent identities and clear spatial relationships throughout the video.

## 4.2. Various Animation Scenarios

To further enhance generalization, we additionally trained our model on the Gen-dataset. The examples shown in Fig. 7 are sampled from the Gen-dataset. As illustrated, the first example highlights (in the red box) the weapon held by the character in the generated video. Without training on the Gen-dataset, the *Stage 1 model* fails to maintain subject consistency: the weapon disappears during motion generation. In contrast, the *Stage 2 model*, preserves the weapon throughout the video, demonstrating improved temporal consistency. In the second example, the red box highlights the background region. The *Stage 1 model* produces a nearly static background, showing limited scene dynamics. While the *Stage 2 model* generates a background that changes naturally with the character’s movement, making the camera view transitions more realistic. In general, training Gen-dataset enhances the models ability to capture motion-scene interactions in multi-person video generation.

## 4.3. Ablation Studies

**Mask-driven Design vs. Addition-driven Design.** An alternative to multi-characters image animation is to aggregate per-characters representations through a simple sum-




Figure 8. Addition-driven design works for two characters but fails to generalize to more characters.

mation strategy. Based on UniAnimate-DiT, we first isolate per-person poses using tracking masks, extract character-specific features for each individual, and then sum these features into a single combined representation. This approach largely avoids the *symmetry issue*. As shown in Fig. 8, with sufficient training on a two-characters dataset, the addition-driven design can handle two-character animation; However, it fails to generalize to settings with more characters than those seen during training. In contrast, our mask-driven design preserves clear per-characters spatial organization and achieves much stronger scalability in multi-characters scenarios.

**Identifier Assigner & Identifier Adapter.** We evaluate the effectiveness of the mask-driven design. As shown in Fig. 9, the results are obtained with and without the mask-driven design after 40 epochs of training on the Swing Dance dataset [22]. The first example, sampled from the Swing




Figure 9. **Ablation on Identifier Assigner & Identifier Adapter.** The combination of Identifier Assigner and Identifier Adapter improves identity consistency and enhances the framework’s extensibility, enabling three-characters image animation.




Figure 10. Our framework remains compatible with single-person image animation, even though the introduction increasing training complexity.

Dance dataset, shows that even after sufficient training, the highlighted character in the generated video struggles to maintain temporal consistency. In contrast, our method produces more stable and coherent results. The second example is sampled from our collected unseen videos. It indicates that the mask-driven design facilitates the model’s capacity to discern three individuals in the generated video, despite being trained exclusively on two-person dataset. This indicates that our design improves both identity consistency and generalizability of our model to unseen multi-person scenarios.

In addition, we use several examples from the TiKToK [16] dataset (which is not used in our model training) to evaluate the compatibility of our model with single-person video generation. As shown in Fig. 10, even though the introduction of the mask-driven design increases the training complexity, our method still demonstrates strong performance in motion control and identity consistency for single-

person videos.

Table 2. **Quantitative results on TikTok dataset.** Our approach maintains comparable performance to other models which target at single character animation.

Method	PSNR $\uparrow$	SSIM $\uparrow$	L1(E-3) $\downarrow$	LPIPS $\downarrow$	FVD $\downarrow$	FID-VID $\downarrow$
DisPose [19]	17.17	0.691	74.97	0.261	615.27	64.74
UniAnimate-DiT [36]	17.76	0.781	68.16	0.337	649.30	50.16
Ours	<b>23.68</b>	<b>0.867</b>	<b>33.65</b>	<b>0.250</b>	<b>342.48</b>	<b>41.85</b>

The quantitative results in Tab. 2 further show that our approach maintains comparable performance to other state-of-the-art models on single-person generation tasks.

## 5. Conclusion

In conclusion, in this paper, we present an extensible mask-driven framework for multi-character image animation built upon modern Diffusion Transformers. By introducing two novel modules-**Identifier Assigner** and **Identifier Adapter**-and a dedicated training strategy for identity modeling, our method explicitly models spatial relationships among multiple characters, allowing the framework to extend to varying numbers of participants. Furthermore, our training strategy enables the model to generate videos with more characters than those seen during training, significantly improving data efficiency and demonstrating strong practical value. Extensive experiments verify that our model, trained only on two-character videos, successfully performs multi-character image animation and achieves superior performance compared with state-of-the-art baselines, while maintaining compatibility with single-person scenarios.

**Acknowledgment:** This work is partially supported by the National Natural Science Foundation of China under Grant Number 62476264 and 62406312, the Beijing Natural Science Foundation under Grant Number 4244098, and the Science Foundation of the Chinese Academy of Sciences.

## References

- [1] Yogesh Balaji, Martin Renqiang Min, Bing Bai, Rama Chelappa, and Hans Peter Graf. Conditional gan with discriminative filter generation for text-to-video synthesis. In *IJCAI*, page 2, 2019. 5
- [2] 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. 1, 3
- [3] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7291–7299, 2017. 3
- [4] Junhao Chen, Mingjin Chen, Jianjin Xu, Xiang Li, Junting Dong, Mingze Sun, Puhua Jiang, Hongxiang Li, Yuhang Yang, Hao Zhao, Xiaoxiao Long, and Ruqi Huang. Dance-together! identity-preserving multi-person interactive video generation. 2025. 3, 4
- [5] Shoufa Chen, Mengmeng Xu, Jiawei Ren, Yuren Cong, Sen He, Yanping Xie, Animesh Sinha, Ping Luo, Tao Xiang, and Juan-Manuel Perez-Rua. Gentron: Diffusion transformers for image and video generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6441–6451, 2024. 3
- [6] Zigeng Chen, Gongfan Fang, Xinyin Ma, Ruonan Yu, and Xinchao Wang. dparallel: Learnable parallel decoding for dllms. In *International Conference on Learning Representations*, 2026. 3
- [7] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- [8] Weilun Feng, Chuanguang Yang, Zhulin An, Libo Huang, Boyu Diao, Fei Wang, and Yongjun Xu. Relational diffusion distillation for efficient image generation. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 205–213, 2024. 3
- [9] Weilun Feng, Chuanguang Yang, Haotong Qin, Xiangqi Li, Yu Wang, Zhulin An, Libo Huang, Boyu Diao, Zixiang Zhao, Yongjun Xu, et al. Q-vdit: Towards accurate quantization and distillation of video-generation diffusion transformers. In *Proceedings of the 42nd International Conference on Machine Learning*, 2025. 3
- [10] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. *arXiv preprint arXiv:2307.04725*, 2023. 1, 3
- [11] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 3
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 3
- [13] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research*, 23(47):1–33, 2022. 3
- [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. 1, 3
- [15] Alain Hore and Djemel Ziou. Image quality metrics: Psnr vs. ssim. In *2010 20th international conference on pattern recognition*, pages 2366–2369. IEEE, 2010. 5
- [16] 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. 5, 8
- [17] Zeyinzi Jiang, Zhen Han, Chaojie Mao, Jingfeng Zhang, Yulin Pan, and Yu Liu. Vace: All-in-one video creation and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17191–17202, 2025. 2, 5, 6
- [18] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*, 2024. 1, 3
- [19] Hongxiang Li, Yaowei Li, Yuhang Yang, Junjie Cao, Zhihong Zhu, Xuxin Cheng, and Chen Long. Dispose: Disentangling pose guidance for controllable human image animation. *arXiv preprint arXiv:2412.09349*, 2024. 2, 5, 6, 8
- [20] Bin Lin, Yunyang Ge, Xinhua Cheng, Zongjian Li, Bin Zhu, Shaocong Wang, Xianyi He, Yang Ye, Shenghai Yuan, Lihuan Chen, et al. Open-sora plan: Open-source large video generation model. *arXiv preprint arXiv:2412.00131*, 2024. 3
- [21] Xin Ma, Yaohui Wang, Gengyun Jia, Xinyuan Chen, Ziwei Liu, Yuan-Fang Li, Cunjian Chen, and Yu Qiao. Latte: Latent diffusion transformer for video generation. *arXiv preprint arXiv:2401.03048*, 2024. 3
- [22] Vongani H Maluleke, Lea Miller, Jathushan Rajasegaran, Georgios Pavlakos, Shiry Ginosar, Angjoo Kanazawa, and Jitendra Malik. Synergy and synchrony in couple dances. *arXiv preprint arXiv:2409.04440*, 2024. 5, 7
- [23] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pages 8162–8171. PMLR, 2021. 3
- [24] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023. 3
- [25] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 3

- [26] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022. 3
- [27] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3
- [28] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 3
- [29] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494, 2022. 3
- [30] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 3
- [31] 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
- [32] Shuyuan Tu, Zhen Xing, Xintong Han, Zhi-Qi Cheng, Qi Dai, Chong Luo, and Zuxuan Wu. Stableanimator: High-quality identity-preserving human image animation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 21096–21106, 2025. 3
- [33] 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. 5
- [34] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, Jianyuan Zeng, Jiayu Wang, Jingfeng Zhang, Jingtong Zhou, Jinkai Wang, Jixuan Chen, Kai Zhu, Kang Zhao, Keyu Yan, Lianghua Huang, Mengyang Feng, Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang, Tianxing Wang, Tianyi Gui, Tingyu Weng, Tong Shen, Wei Lin, Wei Wang, Wei Wang, Wenmeng Zhou, Wente Wang, Wenting Shen, Wenyuan Yu, Xianzhong Shi, Xiaoming Huang, Xin Xu, Yan Kou, Yangyu Lv, Yifei Li, Yijing Liu, Yiming Wang, Yingya Zhang, Yitong Huang, Yong Li, You Wu, Yu Liu, Yulin Pan, Yun Zheng, Yuntao Hong, Yupeng Shi, Yutong Feng, Zeyinzi Jiang, Zhen Han, Zhi-Fan Wu, and Ziyu Liu. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025. 1, 2, 3, 4, 5
- [35] 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. 3
- [36] Xiang Wang, Shiwei Zhang, Changxin Gao, Jiayu Wang, Xiaoqiang Zhou, Yingya Zhang, Luxin Yan, and Nong Sang. Unianimate: Taming unified video diffusion models for consistent human image animation. *Science China Information Sciences*, 2025. 2, 4, 5, 6, 8
- [37] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 5
- [38] Zeqing Wang, Gongfan Fang, Xinyin Ma, Xingyi Yang, and Xinchao Wang. Sparsed: Sparse attention for diffusion language models. In *International Conference on Learning Representations*, 2026. 3
- [39] Weijia Wu, Mingyu Liu, Zeyu Zhu, Xi Xia, Haoen Feng, Wen Wang, Kevin Qinghong Lin, Chunhua Shen, and Mike Zheng Shou. Moviebench: A hierarchical movie level dataset for long video generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 28984–28994, 2025. 2
- [40] Zunnan Xu, Yukang Lin, Haonan Han, Sicheng Yang, Ronghui Li, Yachao Zhang, and Xiu Li. Mambatak: Efficient holistic gesture synthesis with selective state space models. *Advances in Neural Information Processing Systems*, 37:20055–20080, 2024. 2
- [41] Jingyun Xue, Hongfa Wang, Qi Tian, Yue Ma, Andong Wang, Zhiyuan Zhao, Shaobo Min, Wenzhe Zhao, Kaihao Zhang, Heung-Yeung Shum, et al. Follow-your-pose v2: Multiple-condition guided character image animation for stable pose control. *arXiv e-prints*, pages arXiv–2406, 2024. 3
- [42] Han Yang, Chuanguang Yang, Qiuli Wang, Zhulin An, Weilun Feng, Libo Huang, and Yongjun Xu. Multi-party collaborative attention control for image customization, 2025. 2
- [43] Zhendong Yang, Ailing Zeng, Chun Yuan, and Yu Li. Effective whole-body pose estimation with two-stages distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4210–4220, 2023. 5
- [44] Haobo Yuan, Xiangtai Li, Tao Zhang, Zilong Huang, Shilin Xu, Shunping Ji, Yunhai Tong, Lu Qi, Jiashi Feng, and Ming-Hsuan Yang. Sa2va: Marrying sam2 with llava for dense grounded understanding of images and videos. *arXiv*, 2025. 5
- [45] Fan Zhang, Shulin Tian, Ziqi Huang, Yu Qiao, and Ziwei Liu. Evaluation agent: Efficient and promptable evaluation framework for visual generative models. *arXiv preprint arXiv:2412.09645*, 2024. 3
- [46] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3836–3847, 2023. 2
- [47] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 5
- [48] 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. In *International Conference on Machine Learning*, 2025. [2](#), [5](#), [6](#)

- [49] Shenhao Zhu, Junming Leo Chen, Zuozhuo Dai, Zilong Dong, Yinghui Xu, Xun Cao, Yao Yao, Hao Zhu, and Siyu Zhu. Champ: Controllable and consistent human image animation with 3d parametric guidance. In *European Conference on Computer Vision*, pages 145–162. Springer, 2024.

[3](#)