

PoseCrafter: One-Shot Personalized Video Synthesis Following Flexible Pose Control

Yong Zhong^{*1}, Min Zhao^{*1}, Zebin You¹, Xiaofeng Yu², Changwang Zhang²,
and Chongxuan Li^{**1}

¹ Gaoling School of AI, Renmin University of China, Beijing, China

² Huawei Technologies Co., Ltd

yongzhong@ruc.edu.cn, gracezhao1997@gmail.com, zebin@ruc.edu.cn,
yuxiaofeng16@huawei.com, changwangzhang@foxmail.com,
chongxuanli@ruc.edu.cn

Abstract. In this paper, we introduce PoseCrafter, a one-shot method for personalized video generation following the control of flexible poses. Built upon Stable Diffusion and ControlNet, we carefully design an inference process to produce high-quality videos without the corresponding ground-truth frames. First, we select an appropriate reference frame from the training video and invert it to initialize all latent variables for generation. Then, we insert the corresponding training pose into the target pose sequences to enhance faithfulness through a trained temporal attention module. Furthermore, to alleviate the face and hand degradation resulting from discrepancies between poses of training videos and inference poses, we implement simple latent editing through an affine transformation matrix involving facial and hand landmarks. Extensive experiments on several datasets demonstrate that PoseCrafter achieves superior results to baselines pre-trained on a vast collection of videos under 8 commonly used metrics. Besides, PoseCrafter can follow poses from different individuals or artificial edits and simultaneously retain the human identity in an open-domain training video. Our project page is available at <https://ml-gsai.github.io/PoseCrafter-demo/>.

Keywords: Personalized Video Generation · Pose Guidance · One-Shot · Diffusion Models

1 Introduction

The generation of digital humans, particularly through pose-guided techniques that allow for the creation of personalized human videos following specific pose instructions, holds significant promise in both practical and entertainment arenas. At the front of this area, pose-guided image-to-video approaches [6, 20, 41, 45] involve the conversion of a static human photograph into a dynamic video sequence adhering to given pose sequence. Despite its promising prospects, this

^{*} Equal contribution.

^{**} Correspondence to Chongxuan Li.

technique confronts challenges due to the high-dimensional nature of video data, necessitating extensive datasets of high-quality human videos, often beyond public accessibility [20]. Moreover, the computational costs involved in processing such data can be exceptionally daunting.

To circumvent the challenges associated with large-scale data requisites and substantial training expenditures, this paper focuses on a one-shot setting where we aim to generate videos, retaining the identity in a single training video and following the control of flexible poses. Notably, while existing video editing methods [4, 26, 36, 43, 51] have provided valuable insights, their applicability is limited because they significantly rely on real video frames corresponding to the target poses. This prerequisite poses a significant challenge within the realm of video generation, where such frames are unavailable, compounding the difficulty of generating high-quality content. Particularly, our empirical findings, as detailed in Tab. 1 and Tab. 2 vividly demonstrate that a representative approach [51] produces videos with notably inferior quality.

To relieve this dependency, in this paper, we present PoseCrafter, a one-shot method for personalized video generation following the control of flexible poses. Built upon Stable Diffusion [33] and ControlNet [47], PoseCrafter requires merely fine-tuning the pre-trained model on a single open-domain video. Technically, we carefully design an inference process to produce high-quality videos following flexible pose control without the corresponding frames. First, we select an appropriate reference frame from the training video and invert it to initialize all latent variables for generation. Then, we insert the corresponding training pose into the target pose sequences to enhance faithfulness through a trained temporal attention module. Furthermore, to alleviate the face and hand degradation resulting from discrepancies between poses of training videos and inference poses, we implement simple latent editing through an affine transformation matrix involving facial and hand landmarks.

Compared with the aforementioned image-to-video methods [6, 20, 41, 45], PoseCrafter is both computational and data efficient. Besides, leveraging a video as a source offers a richer tapestry of information than a single image, e.g., uncovering occluded details. Moreover, fine-tuning with a video source facilitates a deeper comprehension of individual identities and environmental nuances, significantly enhancing the authenticity and fidelity of the generated video content.

Experimentally, we conduct extensive evaluations on two publicly accessible datasets, TikTok [23] and TED [37], covering real-life dancing and speaking scenarios. Quantitative (under 8 common metrics) and qualitative results (see Sec. 5.2) demonstrate that PoseCrafter produces videos of higher quality than available image-to-video baselines [41, 45], including commercial-grade applications such as GEN-2 [11], and fine-tuning based approaches for specific subjects [41, 51]. Moreover, we illustrate that PoseCrafter can follow flexible pose control including different poses from the same individual, artificially designed poses, and poses of different individuals (see Sec. 5.3). These findings strongly support the potential of one-shot methods for pose-guided video generation.

The key contributions of this paper are summarized as follows:

- We present PoseCrafter, a one-shot pose-guided personalized video generation framework, designed to circumvent the need for large-scale high-quality human video datasets and extensive training resources.
- PoseCrafter introduces reference-frame selection and insertion to improve sample quality, and further alleviates face and hand degradation through simple latent editing.
- PoseCrafter fine-tuned on a single video achieves superior results to baselines pre-trained on a vast collection of videos under 8 commonly used metrics.
- PoseCrafter can follow poses from different individuals or artificial edits and simultaneously retain the human identity in an open-domain training video.

2 Related Work

2.1 Controllable Image Generation

Recently, because of the capacity to produce high-fidelity samples, diffusion models are inherently utilized in controllable image generation [1, 10, 16, 27, 50, 52]. Text-prompt-based methods [32, 33, 35] often struggle with precise control, limited by text-image pair quality and complexity of lengthy inputs. Attention control [10, 16] and latent editing [7, 30] can somewhat enhance alignment of text-to-image generation. Personalized methods [19, 34] capture the common concept across a few images and reproduce that concept in generated images. Furthermore, some research [28, 47] introduces adapters that input various conditions like skeleton images to steer semantics and spatial structures. Our work centers on controllable video generation with pre-trained adapters and latent editing.

2.2 Video Editing and Predicting

Video editing aims to modify given videos to meet specific requirements. Training-based methods [8, 11, 25] involve collecting and training on large-scale videos, incurring high costs. To circumvent this, train-free methods [31, 42, 48] effectively extend existing text-to-image models into the video domain but may compromise on capabilities. One-shot methods [4, 26, 36, 43, 51] offer a balance between computational cost and performance, requiring tuning on just a specific video instance. Our work also focuses on the one-shot setting but surpasses video editing, as it eliminates the need for corresponding original videos of input pose conditions. Video prediction [13, 40, 44] is intended to forecast a segment of a video based on other segments. In our context, when training and inference poses originate from the same individual, our task can be viewed as video prediction.

2.3 Personalization Video Generation

Text-to-video models [17, 24, 38] create a specific subject video through text description. Image-to-video approaches [2, 12, 29] animate a still image to produce a dynamic variation about it. To harness the vast array of personalized image

models in the community, certain techniques [15] transfer these into video generation models by incorporating temporal modeling modules pre-trained on a large-scale video dataset. Concentrating on controllable human-centric generation, some human pose transfer works [37, 49] migrate the motions of a source video to a target subject. However, these methods are inflexible due to using the original video as guidance. More relevant to our work, pose-guided image-to-video generation [6, 20, 41, 45] adapts a given image to match specific pose sequences. Due to the complex distribution of human video data, they often need to train on extensive high-quality human videos that may be unavailable. To circumvent the need for large-scale human video datasets and extensive training resources, we adopt a one-shot setting. Disco [41], which underwent training on the TikTok dataset [23] offers a strategy for fine-tuning on a single video. In contrast, our method eliminates the need for pre-training on video data, requiring tuning only on a single video of the target individual. Similar to our work, Everybody-Dance-Now [5] excludes additional training videos but requires long videos of a target subject (8-17 minutes), while our method merely needs a few video frames (e.g., 8 frames).

3 Preliminary

ControlVideo [51] is a representative one-shot video editing method following the additional control based on ControlNet [47]. For video editing, formally, given a source prompt p_s , a reference video with N frames $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^N$, a target prompt p_t , and extracted condition \mathbf{Q} from \mathbf{X} , it aims to generate a high-quality video $\mathbf{V} = \{\mathbf{v}_i\}_{i=1}^N$. In the following, we present the technical details of ControlVideo.

Model Architecture. ControlVideo consists of two main components: a large-scale latent diffusion model (i.e. Stable Diffusion) for image generation [33], and a condition control model, ControlNet [47], which ensures the generated images adhere to input conditions \mathbf{Q} . To adapt this image generation framework to video applications, ControlVideo integrates two key temporal modules: key-frame attention and temporal attention.

Key-Frame Attention. To achieve an intersection in the temporal axis across frames, ControlVideo advances self-attention as key-frame attention, in which all frames in a video share the key and value of a certain frame. Specifically, for a hidden feature map $x \in \mathbb{R}^{(f,d,c)}$ where f indicates video length and d is the 1D spatial dimension, given a query matrix Q_i , key matrix K_i , and value matrix V_i in a self-attention for the i -th frame, key-frame attention is formulated as:

$$\text{Softmax}\left(\frac{Q_i K_k^T}{\sqrt{c}}\right) V_k, \quad (1)$$

where c is the hidden feature dimension and k is the index of the key frame providing the key and value for all frames.

Temporal Attention. To further improve temporal consistency, ControlVideo introduces temporal attention [18], which computes feature similarity at the same spatial location across all frames. Specifically, it firstly exchanges the frame length and the spatial dimension so that $x \in \mathbb{R}^{(d,f,c)}$. Then, given the Q_j , K_j , and V_j for the j -th spatial location, temporal attention is computed as:

$$\text{Softmax}\left(\frac{Q_j K_j^T}{\sqrt{c}}\right) V_j. \quad (2)$$

Training and Inference. To enable parameter-efficient fine-tuning based on pre-trained weights from large-scale image generation models [33], partial parameters in key-frame attention and temporal attention are trainable. In particular, ControlVideo model ϵ_θ can be trained on a single training video with extracted conditions \mathbf{Q} by the mean-squared error:

$$\mathbb{E}_{t, \mathbf{X}, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} \|\epsilon - \epsilon_\theta(\mathcal{E}(\mathbf{X}_t), t, \mathbf{Q}, p_s)\|_2^2, \quad (3)$$

where \mathbf{X}_t indicates a noise video at the timestep t , and $\mathcal{E}(\cdot)$ is an encoder function compressing \mathbf{X}_t into a lower-dimensional latent variables $\mathbf{Z}_t = \mathcal{E}(\mathbf{X}_t)$ which is reversible through a decoder function $\mathcal{D}(\cdot)$ [33]. During inference, it employs DDIM inversion [39] of \mathbf{X} to create an initial latent \mathbf{Z}_T as guidance, which is a crucial step for detail preservation of \mathbf{X} and is represented as:

$$\mathbf{Z}_T = \text{DDIM-inv}(\mathbf{X}, \mathbf{Q}, p_s). \quad (4)$$

4 PoseCrafter

For crafting pose-guided personalized videos in a one-shot setting, we introduce a holistic framework PoseCrafter in Sec. 4.1 and propose two key techniques, i.e. reference-frame selection and insertion in Sec. 4.2 and latent editing in Sec. 4.3.

4.1 Settings and Framework

Our goal with a human-centric reference video is to create a high-quality video that follows a specific target pose sequence while preserving the identity from the reference. This task is notably more complex than video editing due to the absence of corresponding videos for the target poses. Applying ControlVideo in this scenario is not straightforward, as it depends on the DDIM inversion of non-existent videos to generate the initial latent \mathbf{Z}_T for guidance (see Eq. (4)).

Formally, given a source prompt p_s , a reference video with N frames $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^N$ and corresponding extracted poses $\mathbf{Q} = \{\mathbf{q}_i\}_{i=1}^N$, along with a target prompt p_t and a pose sequence consisting of M frames $\mathbf{P} = \{\mathbf{p}_i\}_{i=1}^M$, we aim to generate a high-quality video $\mathbf{V} = \{\mathbf{v}_i\}_{i=1}^M$. This video should be of superior fidelity, exhibit exceptional faithfulness to \mathbf{X} , and accurately align with \mathbf{P} .

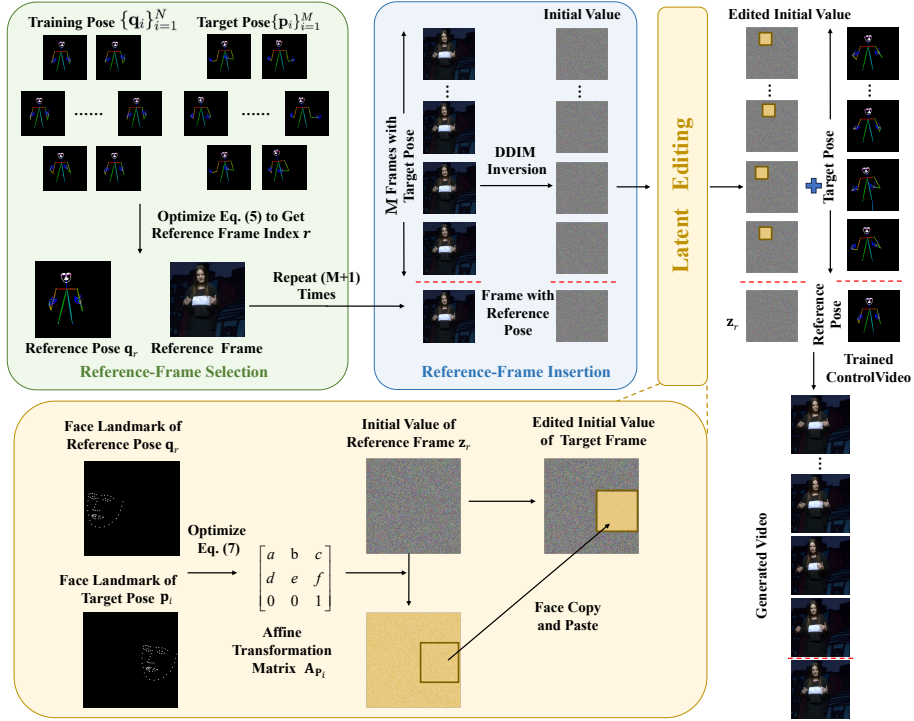


Fig. 1: Inference framework of PoseCrafter. First, we select a frame from the training video to form a pseudo reference video, followed by DDIM inversion. Next, the pose from the reference frame is inserted into the inference poses. Finally, we edit latent to refine the generation of faces and hands by affine transformation.

As illustrated in Fig. 1, we introduce a novel framework PoseCrafter, which shares a similar model and training protocol as ControlVideo yet a distinct inference process to address the aforementioned challenge. In particular, we utilize Stable Diffusion [33] for video generation, coupled with the condition control model ControlNet [47] to accurately dictate specific postures in digital humans (see Sec. 3). The training employs the loss function (see Eq. (3)) on an individual video. Moreover, to mitigate GPU memory constraints on long video training, we randomly sample n frames from a training video along the time dimension as a minimal batch at each training iteration. For the inference phase, to bypass the reliance on videos corresponding to target poses \mathbf{P} , we introduce a reference-frame selection and insertion technique to enhance faithfulness (see Sec. 4.2), and latent editing to refine face and hand details (see Sec. 4.3).

4.2 Reference-Frame Selection and Insertion

Our preliminary experiment suggests that using random noise as the initial latent \mathbf{Z}_T struggles to preserve human identity and background. To tackle this

problem, we select the frame \mathbf{x}_r ($r \in [1, N]$) whose pose coordinates are closest to the inference poses from the training video as the reference frame. Given training poses $\{\mathbf{q}_i\}_{i=1}^N$ and inference poses $\{\mathbf{p}_j\}_{j=1}^M$, we determine the index of the reference frame \mathbf{x}_r as follows:

$$r = \arg \min_{1 \leq i \leq N} \sum_{1 \leq j \leq M} \|\mathbf{q}_i - \mathbf{p}_j\|_2^2, \quad (5)$$

where $\|\cdot\|_2$ denotes the ℓ_2 norm. Subsequently, we replicate \mathbf{x}_r M times to create a pseudo reference video $\{\mathbf{x}_r\}^M$, matching the length of the inference poses. This allows us to use the DDIM inversion [39] of this pseudo video as the initial latent.

Notably, by inverting the reference video $\{\mathbf{x}_r\}^M$ to initial latent $\mathbf{Z}_T = \{\mathbf{z}_r\}^M$ and then generating the video \mathbf{V} guided by inference poses \mathbf{P} , we can endeavor to preserve as much detail as possible in the training video. Previous work [41] uses the same fixed noise latent for all generated frames rather than different random noise latent to maintain better temporal consistency, so we prefer merely one reference frame instead of multiple frames to form the pseudo reference video.

To further enhance identity faithfulness, we insert the pose \mathbf{q}_r of \mathbf{x}_r into inference poses $\{\mathbf{P}_i\}_{i=1}^M$, creating an extended sequence of $M + 1$ elements, i.e. $\{\mathbf{q}_r, \mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_M\}$ ³. Following this augmented sequence, we generate $M + 1$ frames and subsequently remove the redundant frame after inference. This method effectively deepens the connection between the generated video and the reference frame, i.e. faithfulness, as DDIM inversion and sampling [39] can accurately reconstruct the original image (i.e. the reference frame) while temporal modules (see Sec. 3) encourage generated frames to be consistent.

4.3 Improve Face and Hand Generation by Editing \mathbf{Z}_T

Our initial experiments reveal that a higher number of landmarks for faces and hands, coupled with the demand for finer details in the generated results, necessitates a tailored approach to editing latent variables for facial and hand features to enhance generation quality. Specifically, we draw inspiration from classical methods in face alignment [22], calculating the corresponding affine matrix based on facial landmarks from the reference frame pose and the target frame pose. More precisely, given a face landmark (x, y) and its corresponding target (x', y') , we can achieve that mapping by an affine transformation:

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \mathbf{A} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}, \quad (6)$$

where \mathbf{A} is an affine transformation matrix with six degrees of freedom. For scenarios involving multiple landmarks, we solve the following least square opti-

³ To be compatible with key-frame attention, we place \mathbf{q}_r as the key frame with position k , and in our experiments, we set $k = 1$ by default.

mization problem to get the affine transformation matrix $\mathbf{A}_{\mathbf{p}_i}$:

$$\mathbf{A}_{\mathbf{p}_i} = \arg \min_{\mathbf{A}} \sum_{\mathbf{v} \in F_{\mathbf{q}_r}} \|\mathbf{A}\mathbf{v} - \mathbf{v}_{\mathbf{p}_i}\|_2^2, \quad (7)$$

where $F_{\mathbf{q}_r}$ consists of all face landmarks for the reference pose \mathbf{p}_r and $\mathbf{v}_{\mathbf{p}_i}$ is the landmark in the target pose \mathbf{p}_i corresponding to \mathbf{v} .

As illustrated in Fig. 1, for each inference pose \mathbf{p}_i , we edit its latent variable \mathbf{z}_r , derived from the DDIM inversion process, using the specific affine matrix $\mathbf{A}_{\mathbf{p}_i}$. Specifically, we apply the affine transformation to the entire latent variable \mathbf{z}_r , creating a transformed version $\mathbf{z}^{\mathbf{p}_i}$. We then selectively transfer the facial area of $\mathbf{z}^{\mathbf{p}_i}$, defined as the minimal bounding rectangle encompassing the facial landmarks of \mathbf{p}_i , directly onto the corresponding areas of \mathbf{z}_r for generating target frame \mathbf{v}_i . Moreover, similar operations based on hand landmarks are also performed, resulting in improved hand features in our experiments (see Fig. 6). We believe that this latent editing can be extended to more body parts like legs.

It is noted that the latent variable for each frame is a copy of the latent variable from a single reference frame, hence we perform the DDIM inversion operation just once. Subsequently, different affine transformation matrices are applied to obtain corresponding final latent variables. This approach is more efficient than performing affine transformations on the reference frame followed by separate DDIM inversions for each. Moreover, editing directly in pixel space reduces image quality obtained after DDIM inversion and subsequent sampling.

Notably, PoseCrafter can follow poses from different individuals or artificial edits and simultaneously retain the identity in an open-domain training video.

5 Experiments

5.1 Experimental Settings

Datasets. We train and evaluate our method on test splits of two public datasets TED [37] and TikTok [23]. Due to changes in user privacy and download permissions, we can only download 128 TED test videos from YouTube using download links provided in [37]. Following prior work [6, 41, 45], we use 10 TikTok test videos presented in [41]. We also gather 10 high-quality videos from YouTube, encompassing a variety of scenes such as interviews, movie clips, and talk shows. For each video, we designate 100 frames as test frames and uniformly select N frames from their preceding frames as training frames. We crop and resize all videos to a uniform 512×512 resolution. See Appendix A.1 for more details.

Implementation Details. We initialize the models with the publicly available realistic model, leosamsMoonfilm [9], based on Stable Diffusion 1.5. For accurate pose extraction, we employ DWpose [46], an excellent 2D pose extractor. During training, we set the learning rate as 0.003 and the minimal batch size $n = 8$. Note that our method requires only 2 minutes to complete training on an individual video with 8 frames using an RTX3090 GPU, and all experiments are implemented on RTX3090 GPUs. Please refer to Appendix A.2 for more details.

Evaluation. We evaluate video quality encompassing temporal consistency, faithfulness, fidelity, and pose alignment by 8 common quantitative metrics. Specifically, we measure temporal consistency by calculating the average CLIP similarity (CLIP-T) among all frames of the generated video. To assess faithfulness, we use the average similarity (CLIP-I) between reference images and generated frames. Regarding video fidelity, we adopt FID and FVD. For pose alignment, we compute the mean squared error (MSE-P) between 2D keypoint coordinates of all pose pairs extracted from generated images and real images. Additionally, we consider common reconstitution measures such as SSIM, PSNR, and LPIPS, which calculate the average distance between all pairs of generated images and corresponding ground truth, involving temporal consistency, faithfulness, fidelity, and pose alignment.

Baselines. We select four representative methods as baselines for a comprehensive comparison. (1) Pose-guided image-to-video methods: Disco [41] and MagicAnimate [45], which are pre-trained on a combination of large-scale image data and smaller video data. We utilize their provided public checkpoints trained on TikTok videos (around 350 videos with 10-12 seconds and 30 fps)⁴. (2) Pose-guided one-shot fine-tuning methods: ControlVideo and fine-tuned Disco. For ControlVideo, we inverse the 1-st frame in the training video and edit it guided by inference poses. For Disco, we use their introduced methods to fine-tune it on each video. More details are provided in Appendix A.3.

5.2 Main Results

We demonstrate that PoseCrafter fine-tuned on a single video achieves superior quantitative and qualitative results to baselines pre-trained on a vast collection of videos. Note that, for a fair comparison, we unify image resolution, test dataset, and evaluation codes which may differ across all methods.

Quantitative Comparisons. Tab. 1 and Tab. 2 present the quantitative comparison on TikTok and TED, separately. Although training on 8 frames, by the comprehensive framework and sophisticated inference algorithm, PoseCrafter achieves competitive results on all 8 metrics to the state-of-the-art baseline MagicAnimate on TikTok, but superior results than it on TED which indicates its limited generalization capability. Moreover, our method significantly improves all metrics than ControlVideo and one-shot Disco, demonstrating the superiority of our training and inference framework. Remarkably, PoseCrafter, which requires only about 12 minutes of training on 32 frames with a single RTX3090 GPU, significantly outperforms MagicAnimate and excels in all quantitative metrics than all baselines on TikTok. This underscores the potential of one-shot methods, which, even with limited training, can surpass approaches relying on extensive training for crafting high-quality pose-guided personalized videos.

Additionally, fine-tuned Disco degrades many metrics on TikTok. We hypothesize that Disco fine-tunes its cross-attention layers in UNet and full parameters

⁴ Disco additionally collects 250 internal TikTok-style videos for training.

Table 1: Quantitative results on TikTok test dataset. NFs represent the number of available frames for the corresponding method. Image-GT signifies that these image-level metrics are calculated between each generated image and its corresponding ground truth image. † represents we fine-tune corresponding methods using their tailored strategies for specific subjects.

Method	NFs	Image-GT				Image		Video	
		SSIM ↑	PSNR ↑	LPIPS ↓	MSE-P ↓	FID ↓	CLIP-I ↑	FVD ↓	CLIP-T ↑
DisCo [41]	1	0.704	15.16	0.358	9.11E-3	76.40	0.827	689.23	0.908
MagicAnimate [45]	1	0.756	17.95	0.265	5.48E-3	57.60	0.846	374.40	0.918
Disco† [41]	8	0.683	15.33	0.371	8.86E-3	65.73	0.813	807.35	0.886
ControlVideo [51]	8	0.738	16.93	0.311	6.78E-3	56.50	0.827	489.30	0.912
PoseCrafter (ours)	8	0.765	17.36	0.275	5.76E-3	48.09	0.840	440.49	0.921
PoseCrafter (ours)	16	0.776	17.87	0.252	5.23E-3	42.09	0.864	397.19	0.919
PoseCrafter (ours)	32	0.786	18.56	0.233	4.05E-3	39.65	0.854	362.09	0.922

Table 2: Quantitative results on TED test dataset. NFs represent the number of available frames for the corresponding method. Image-GT signifies that these image-level metrics are calculated between each generated image and its corresponding ground truth image. † represents we fine-tune corresponding methods using their tailored strategies for specific subjects.

Method	NFs	Image-GT				Image		Video	
		SSIM ↑	PSNR ↑	LPIPS ↓	MSE-P ↓	FID ↓	CLIP-I ↑	FVD ↓	CLIP-T ↑
DisCo [41]	1	0.550	17.22	0.327	3.20E-3	51.57	0.794	309.62	0.909
MagicAnimate [45]	1	0.498	13.32	0.338	1.97E-3	46.96	0.801	194.03	0.912
Disco† [41]	8	0.551	17.98	0.373	3.93E-3	60.56	0.827	244.62	0.909
ControlVideo [51]	8	0.771	22.32	0.196	1.62E-3	29.01	0.866	109.36	0.940
PoseCrafter (ours)	8	0.810	23.92	0.142	1.56E-3	20.40	0.906	80.01	0.954

in ControlNet (294.95 M in total), and thus is easy to overfit to a single training video. As a comparison, our method only requires fine-tuning 68.79M trainable parameters (see Sec. 3), which is a more efficient fine-tuning strategy. Moreover, given the variations in evaluation codes, generated image resolutions and test frame counts among methods, we avoid directly citing their reported results. Instead, we standardize these elements to ensure a fair comparison. A comparison with their originally reported results is provided in Appendix B.1.

Qualitative Comparisons. Fig. 2 shows the qualitative comparisons of our method to all baselines on TED and TikTok. Disco and MagicAnimate rely on OpenPose [3] or DensePose [14] without hand and face poses and thus do not precisely control hand and face details such as facial expression. Moreover, due to the absence of similar reference videos in their training corpus, they struggle to preserve high-level human details such as illumination and clothing on TED. Although fine-tuned Disco on each individual improves detail maintenance to training videos and ControlVideo uses the DDIM inversion, they still do not supersede PoseCrafter, emphasizing the importance of our training and inference

framework. As time progresses, the characters generated by a commercial-grade application (i.e. GEN-2) increasingly diverge from the reference image.

Additionally, with more training frames ($N = 32$), PoseCrafter generates the highest quality videos among all methods on TikTok, which is consistent with quantitative results. Overall, PoseCrafter yields videos most closely aligning with the ground truth among all approaches on TED and TikTok. Remarkably, when the extracted poses are precise, PoseCrafter is capable of generating videos of superior quality compared to the original ground truth, exemplified by our hands in the last image of the TED examples in Fig. 2.




Fig. 2: Qualitative Comparisons of all methods ($M = 100$). Left: TED. Right: TikTok. Time progresses from left to right. PoseCrafter yields the highest quality videos.

5.3 Applications

PoseCrafter, trained in an open-domain video such as interviews, movie clips, and speeches, can follow flexible pose control: (1) poses from the same individual; (2) artificially crafted poses; and (3) poses from different individuals.

Inference with Poses from the Same Individual. Fig. 3 depicts the results of PoseCrafter guided by poses of the same individual as in training frames. PoseCrafter adeptly reconstructs intricate background elements like picture frames and logo texts, and human details including hairstyles and wrinkles. By tweaking the target prompt p_t , it allows for modifications in human attributes, such as changing hair color from black to red, as illustrated in the final row of Fig. 3. In cases where partial segments in a video are of low quality yet yield reliable poses, we can utilize PoseCrafter trained on remaining segments to regenerate the deficient clips, serving as a form of video inpainting.

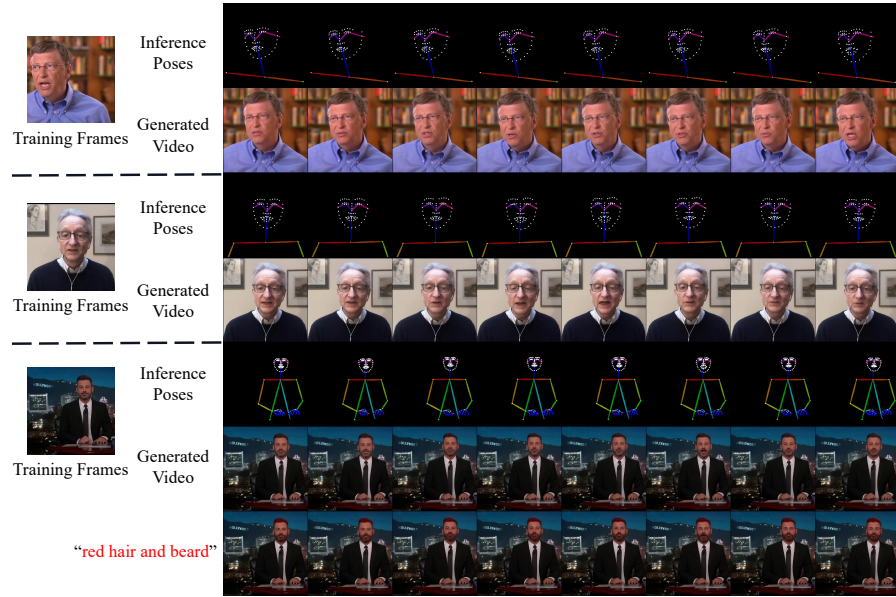


Fig. 3: Inference from poses of the same individual ($N = 100$ and $M = 100$). Time progresses from left to right.

Inference with Artificially Designed Poses. Using pose editors [21], we can alter the pose of a specific training frame to create “dream pose sequences” for guiding PoseCrafter. Illustrated in Fig. 4, we demonstrate this by having a character in entertainment videos progressively blink their right eye, a movie character gradually turn their head to the right, and a speaker lift and wave their left hand and finally sway their head.

Inference with Poses from Different Individuals. As depicted in Fig. 5, PoseCrafter effectively utilizes the poses of an individual different from that in the training frames, successfully creating consistent and faithful videos. However, there may be variations in the facial and body shapes of the created characters, owing to the use of poses from figures with varying facial and body proportions. Moreover, more results with larger and more complex movements are provided in Appendix B.2.

5.4 Ablation

We investigate the significance of reference-frame selection, integration, and latent editing. As demonstrated in the Fig. 6, the automatic selection of a reference frame in original videos leads to a notable improvement in video quality and faithfulness compared to ControlVideo. Its insertion into inference poses further enhances detail preservation in training videos. In addition to refining face and hand generation, latent editing strengthens temporal consistency in generated videos. The quantitative analysis, detailed in Appendix C.1, corroborates the effectiveness of these distinct stages in our methodology. In summary, these critical

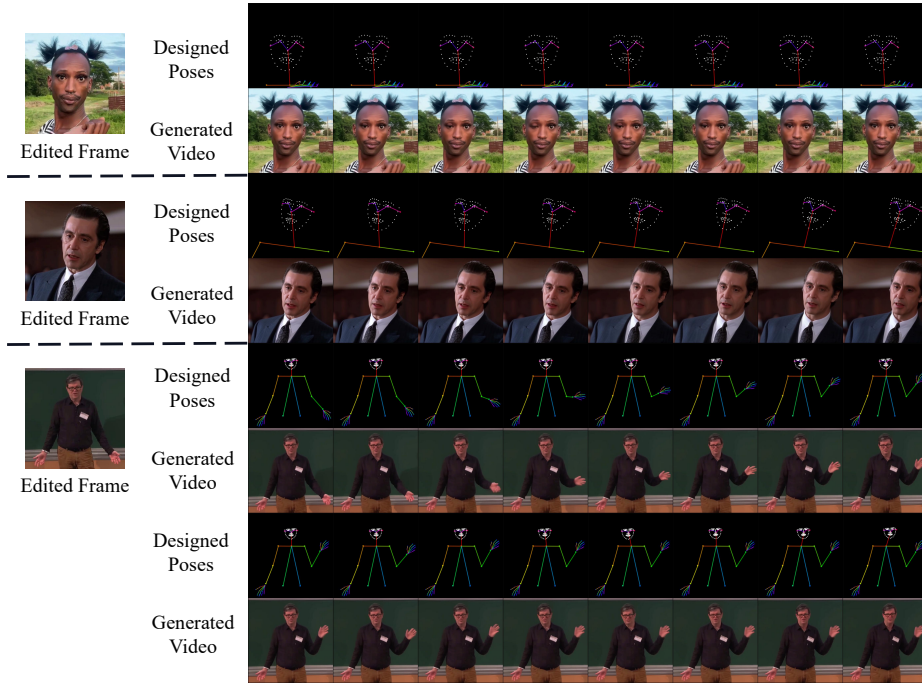


Fig. 4: Inference from artificially designed poses ($N = 100$ and $M = 8, 8, 16$). Time progresses from left to right. From top to bottom: the first pose sequence features a right eye blink; the second pose sequence involves tilting the head rightward; the third pose sequence is waving an arm followed by shaking the head.

design elements contribute to the enhancement of video quality. Furthermore, the roles of training frame quantity and sample timestep choice in latent editing are explored in Appendices C.2 and C.3, respectively.

6 Conclusions

In this paper, we introduce PoseCrafter, a one-shot tuning method for crafting pose-guided personalized videos without the need for extensive data. To address dependence on original videos of given pose sequences in video editing, we introduce two crucial algorithms to enhance video quality: reference frame selection and insertion, and latent editing. Furthermore, PoseCrafter, trained in an open-domain video, demonstrates adaptability to a range of pose sources, including poses from distinct individuals.

Limitations. The video quality of PoseCrafter depends on the ControlNet and latent diffusion models, which struggle with complex poses like crossed fingers, presenting a limitation in our approach. When using poses from others, the body proportions of our generated characters change accordingly. Moreover, large differences between training and inference poses can degrade video quality. Specific failure cases are presented in Appendix B.4.




Fig. 5: Inference poses of other individuals ($N = 50$ and $M = 50$). Time progresses from left to right.




Fig. 6: Qualitative ablation of PoseCrafter.

Acknowledgements

This work was supported by NSF of China (No. 62076145), Major Innovation & Planning Interdisciplinary Platform for the “Double-First Class” Initiative, Renmin University of China, the Fundamental Research Funds for the Central Universities, the Research Funds of Renmin University of China (No. 22XNKJ13), and the Huawei Group Research Fund. C. Li was also sponsored by Beijing Nova Program (No. 20220484044). The work was partially done at the Engineering Research Center of Next-Generation Intelligent Search and Recommendation, Ministry of Education.

References

1. Bhunia, A.K., Khan, S., Cholakkal, H., Anwer, R.M., Laaksonen, J., Shah, M., Khan, F.S.: Person image synthesis via denoising diffusion model. In: CVPR (2023)
2. Blattmann, A., Dockhorn, T., Kulal, S., Mendeleevitch, 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)
3. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7291–7299 (2017)
4. Chai, W., Guo, X., Wang, G., Lu, Y.: Stablevideo: Text-driven consistency-aware diffusion video editing. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 23040–23050 (2023)
5. Chan, C., Ginosar, S., Zhou, T., Efros, A.A.: Everybody dance now. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 5933–5942 (2019)
6. Chang, D., Shi, Y., Gao, Q., Fu, J., Xu, H., Song, G., Yan, Q., Yang, X., Soleymani, M.: Magicedance: Realistic human dance video generation with motions & facial expressions transfer. arXiv preprint arXiv:2311.12052 (2023)
7. Chefer, H., Alaluf, Y., Vinker, Y., Wolf, L., Cohen-Or, D.: Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. ACM Transactions on Graphics (TOG) **42**(4), 1–10 (2023)
8. Chen, W., Wu, J., Xie, P., Wu, H., Li, J., Xia, X., Xiao, X., Lin, L.: Control-a-video: Controllable text-to-video generation with diffusion models. arXiv preprint arXiv:2305.13840 (2023)
9. Civitai: leosamsMoonfilm. <https://civitai.com/models/43977?modelVersionId=113623> (2023)
10. Epstein, D., Jabri, A., Poole, B., Efros, A., Holynski, A.: Diffusion self-guidance for controllable image generation. Advances in Neural Information Processing Systems **36** (2024)
11. Esser, P., Chiu, J., Atighehchian, P., Granskog, J., Germanidis, A.: Structure and content-guided video synthesis with diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7346–7356 (2023)
12. Girdhar, R., Singh, M., Brown, A., Duval, Q., Azadi, S., Rambhatla, S.S., Shah, A., Yin, X., Parikh, D., Misra, I.: Emu video: Factorizing text-to-video generation by explicit image conditioning. arXiv preprint arXiv:2311.10709 (2023)

13. Gu, X., Wen, C., Song, J., Gao, Y.: Seer: Language instructed video prediction with latent diffusion models. arXiv preprint arXiv:2303.14897 (2023)
14. Güler, R.A., Neverova, N., Kokkinos, I.: Densepose: Dense human pose estimation in the wild. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7297–7306 (2018)
15. Guo, Y., Yang, C., Rao, A., Wang, Y., Qiao, Y., Lin, D., Dai, B.: Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. arXiv preprint arXiv:2307.04725 (2023)
16. Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-Or, D.: Prompt-to-prompt image editing with cross attention control. arXiv preprint arXiv:2208.01626 (2022)
17. Ho, J., Chan, W., Saharia, C., Whang, J., Gao, R., Gritsenko, A., Kingma, D.P., Poole, B., Norouzi, M., Fleet, D.J., et al.: Imagen video: High definition video generation with diffusion models. arXiv preprint arXiv:2210.02303 (2022)
18. Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., Fleet, D.J.: Video diffusion models. arXiv:2204.03458 (2022)
19. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. arXiv preprint arXiv:2106.09685 (2021)
20. Hu, L., Gao, X., Zhang, P., Sun, K., Zhang, B., Bo, L.: Animate anyone: Consistent and controllable image-to-video synthesis for character animation. arXiv preprint arXiv:2311.17117 (2023)
21. Huchenlei: sd-webui-openpose-editor. <https://github.com/huchenlei/sd-webui-openpose-editor> (2023)
22. Jabberi, M., Wali, A., Chaudhuri, B.B., Alimi, A.M.: 68 landmarks are efficient for 3d face alignment: what about more? 3d face alignment method applied to face recognition. Multimedia Tools and Applications pp. 1–35 (2023)
23. Jafarian, Y., Park, H.S.: 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. pp. 12753–12762 (2021)
24. Khachatryan, L., Movsisyan, A., Tadevosyan, V., Henschel, R., Wang, Z., Navasardyan, S., Shi, H.: Text2video-zero: Text-to-image diffusion models are zero-shot video generators. arXiv preprint arXiv:2303.13439 (2023)
25. Liew, J.H., Yan, H., Zhang, J., Xu, Z., Feng, J.: Magicedit: High-fidelity and temporally coherent video editing. arXiv preprint arXiv:2308.14749 (2023)
26. Liu, S., Zhang, Y., Li, W., Lin, Z., Jia, J.: Video-p2p: Video editing with cross-attention control. arXiv preprint arXiv:2303.04761 (2023)
27. Ma, L., Gao, T., Jiang, H., Shen, H., Huang, K.: Waveipt: Joint attention and flow alignment in the wavelet domain for pose transfer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7215–7225 (2023)
28. Mou, C., Wang, X., Xie, L., Zhang, J., Qi, Z., Shan, Y., Qie, X.: T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. arXiv preprint arXiv:2302.08453 (2023)
29. 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. pp. 18444–18455 (2023)
30. Nie, S., Guo, H.A., Lu, C., Zhou, Y., Zheng, C., Li, C.: The blessing of randomness: Sde beats ode in general diffusion-based image editing. arXiv preprint arXiv:2311.01410 (2023)

31. Qi, C., Cun, X., Zhang, Y., Lei, C., Wang, X., Shan, Y., Chen, Q.: Fatezero: Fusing attentions for zero-shot text-based video editing. arXiv preprint arXiv:2303.09535 (2023)
32. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125 1(2), 3 (2022)
33. 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. pp. 10684–10695 (2022)
34. Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dream-booth: Fine tuning text-to-image diffusion models for subject-driven generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22500–22510 (2023)
35. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. Advances in Neural Information Processing Systems **35**, 36479–36494 (2022)
36. Shin, C., Kim, H., Lee, C.H., Lee, S.g., Yoon, S.: Edit-a-video: Single video editing with object-aware consistency. arXiv preprint arXiv:2303.07945 (2023)
37. 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. pp. 13653–13662 (2021)
38. Singer, U., Polyak, A., Hayes, T., Yin, X., An, J., Zhang, S., Hu, Q., Yang, H., Ashual, O., Gafni, O., et al.: Make-a-video: Text-to-video generation without text-video data. arXiv preprint arXiv:2209.14792 (2022)
39. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
40. Voleti, V., Jolicoeur-Martineau, A., Pal, C.: Mcvd-masked conditional video diffusion for prediction, generation, and interpolation. Advances in Neural Information Processing Systems **35**, 23371–23385 (2022)
41. Wang, T., Li, L., Lin, K., Lin, C.C., Yang, Z., Zhang, H., Liu, Z., Wang, L.: Disco: Disentangled control for referring human dance generation in real world. arXiv preprint arXiv:2307.00040 (2023)
42. Wang, W., Jiang, Y., Xie, K., Liu, Z., Chen, H., Cao, Y., Wang, X., Shen, C.: Zero-shot video editing using off-the-shelf image diffusion models. arXiv preprint arXiv:2303.17599 (2023)
43. Wu, J.Z., Ge, Y., Wang, X., Lei, S.W., Gu, Y., Shi, Y., Hsu, W., Shan, Y., Qie, X., Shou, M.Z.: Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7623–7633 (2023)
44. Xing, Z., Feng, Q., Chen, H., Dai, Q., Hu, H., Xu, H., Wu, Z., Jiang, Y.G.: A survey on video diffusion models. arXiv preprint arXiv:2310.10647 (2023)
45. Xu, Z., Zhang, J., Liew, J.H., Yan, H., Liu, J.W., Zhang, C., Feng, J., Shou, M.Z.: Magicanimate: Temporally consistent human image animation using diffusion model. arXiv preprint arXiv:2311.16498 (2023)
46. 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)
47. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3836–3847 (2023)

- 48. 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)
- 49. Zhao, J., Zhang, H.: Thin-plate spline motion model for image animation. In: CVPR (2022)
- 50. Zhao, M., Bao, F., Li, C., Zhu, J.: Egsde: Unpaired image-to-image translation via energy-guided stochastic differential equations. *Advances in Neural Information Processing Systems* **35**, 3609–3623 (2022)
- 51. Zhao, M., Wang, R., Bao, F., Li, C., Zhu, J.: Controlvideo: Adding conditional control for one shot text-to-video editing. arXiv preprint arXiv:2305.17098 (2023)
- 52. Zhou, X., Yin, M., Chen, X., Sun, L., Gao, C., Li, Q.: Cross attention based style distribution for controllable person image synthesis. In: *European Conference on Computer Vision*. pp. 161–178. Springer (2022)