

# ConditionVideo: Training-Free Condition-Guided Video Generation

Bo Peng<sup>1,2\*</sup>, Xinyuan Chen<sup>2†</sup>, Yaohui Wang<sup>2</sup>, Chaochao Lu<sup>2</sup>, Yu Qiao<sup>2</sup>

<sup>1</sup>Shanghai Jiao Tong University

<sup>2</sup>Shanghai Artificial Intelligence Laboratory

{pengbo, chenxinyuan, wangyaohui, luchaochao, qiaoyu}@pjlab.org.cn

## Abstract

Recent works have successfully extended large-scale text-to-image models to the video domain, producing promising results but at a high computational cost and requiring a large amount of video data. In this work, we introduce ConditionVideo, a training-free approach to text-to-video generation based on the provided condition, video, and input text, by leveraging the power of off-the-shelf text-to-image generation methods (e.g., Stable Diffusion). ConditionVideo generates realistic dynamic videos from random noise or given scene videos. Our method explicitly disentangles the motion representation into condition-guided and scenery motion components. To this end, the ConditionVideo model is designed with a UNet branch and a control branch. To improve temporal coherence, we introduce sparse bi-directional spatial-temporal attention (sBiST-Attn). The 3D control network extends the conventional 2D controlnet model, aiming to strengthen conditional generation accuracy by additionally leveraging the bi-directional frames in the temporal domain. Our method exhibits superior performance in terms of frame consistency, clip score, and conditional accuracy, outperforming compared methods. For the project website, see <https://pengbo807.github.io/conditionvideo-website/>

## 1 Introduction

Diffusion-based models (Song, Meng, and Ermon 2021; Song et al. 2021; Ho, Jain, and Abbeel 2020; Sohl-Dickstein et al. 2015) demonstrates impressive results in large-scale text-to-image (T2I) generation (Ramesh et al. 2022; Saharia et al. 2022; Gafni et al. 2022; Rombach et al. 2022). Much of the existing research proposes to utilize image generation models for video generation. Recent works (Singer et al. 2023; Blattmann et al. 2023; Hong et al. 2023) attempt to inflate the success of the image generation model to video generation by introducing temporal modules. While these methods reuse image generation models, they still require a massive amount of video data and training with significant amounts of computing power. Tune-A-Video (Wu et al. 2022b) extends Stable Diffusion (Rombach et al. 2022) with additional attention and a temporal module for video editing by tuning one given video. It significantly decreases the

training workload, although an optimization process is still necessary. Text2Video (Khachatryan et al. 2023) proposes training-free generation, however, the generated video fails to simulate natural background dynamics. Consequently, the question arises: *How can we effectively utilize image generation models without any optimization process and embed controlling information as well as modeling dynamic backgrounds for video synthesis?*

We propose ConditionVideo, a training-free conditional-guided video generation method that utilizes off-the-shelf text-to-image generation models to generate realistic videos without any fine-tuning. Specifically, aiming at generating dynamic videos, our model disentangles the representation of motion in videos into two distinct components: conditional-guided motion and scenery motion, enabling the generation of realistic and temporally consistent frames. By leveraging this disentanglement, we propose a pipeline that consists of a UNet branch and a control branch, with two separate noise vectors utilized in the sampling process. Each noise vector represents conditional-guided motion and scenery motion, respectively. To further enforce temporal consistency, we introduce sparse bi-directional spatial-temporal attention (sBiST-Attn) and a 3D control branch that leverages bi-directional adjacent frames in the temporal dimension to enhance conditional accuracy. Our ConditionVideo method outperforms the baseline methods in terms of frame consistency, conditional accuracy, and clip score.

Our key contributions are as follows. (1) We propose ConditionVideo, a training-free video generation method that leverages off-the-shelf text-to-image generation models to generate conditional-guided videos with realistic dynamic backgrounds. (2) Our method disentangles motion representation into conditional-guided and scenery motion components via a pipeline that includes a U-Net branch and a conditional-control branch. (3) We introduce sparse bi-directional spatial-temporal attention (sBiST-Attn) and a 3D conditional-control branch to improve conditional accuracy and temporal consistency.

## 2 Related Work

### 2.1 Diffusion Models

Image diffusion models have achieved significant success in the field of generation (Ho, Jain, and Abbeel 2020; Song,

\*Work done as an intern at Shanghai AI Lab.

†Corresponding author

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

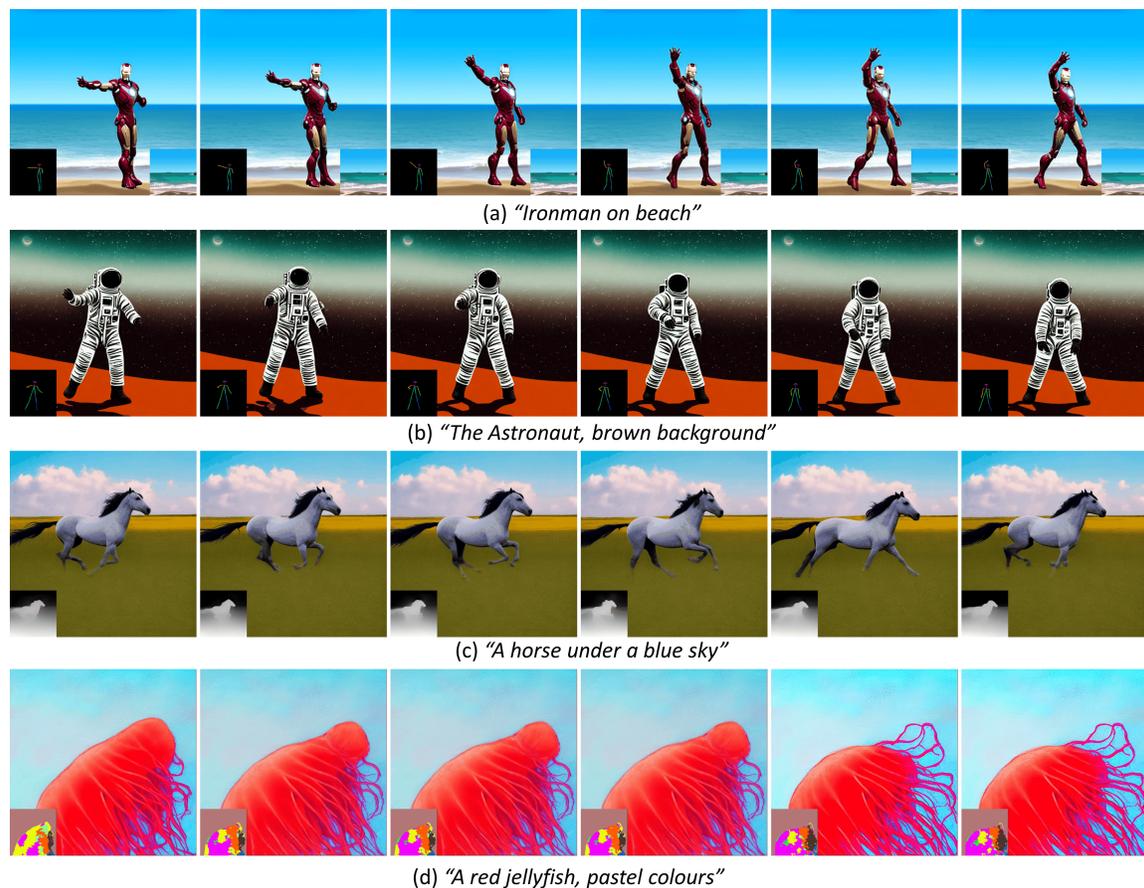


Figure 1: Our training-free method generates videos conditioned on different inputs. In (a), the illustration showcases the process of generation using provided scene videos and pose information, with the background wave exhibiting a convincingly lifelike motion. (b), (c), and (d) are generated based on condition only, which are pose, depth, and segmentation, respectively.

Meng, and Ermon 2021; Song et al. 2021), surpassing numerous generative models that were once considered state-of-the-art (Dhariwal and Nichol 2021; Kingma et al. 2021). With the assistance of large language models (Radford et al. 2021; Raffel et al. 2020), current research can generate videos from text, contributing to the prosperous of image generation (Ramesh et al. 2022; Rombach et al. 2022).

Recent works in video generation (Esser et al. 2023; Ho et al. 2022b; Wu et al. 2022b, 2021, 2022a; Hong et al. 2023; Wang et al. 2023b,c) aim to emulate the success of image diffusion models. Video Diffusion Models (Ho et al. 2022b) extends the UNet (Ronneberger, Fischer, and Brox 2015) to 3D and incorporates factorized spacetime attention (Bertasius, Wang, and Torresani 2021). Imagen Video (Saharia et al. 2022) scales this process up and achieves superior resolution. However, both approaches involve training from scratch, which is both costly and time-consuming. Alternative methods explore leveraging pre-trained text-to-image models. Make-A-Video (Singer et al. 2023) facilitates text-to-video generation through an expanded unCLIP framework. Tune-A-Video (Wu et al. 2022b) employs a one-shot tuning pipeline to generate edited videos from input guided

by text. However, these techniques still necessitate an optimization process. Compared to these video generation methods, our training-free method can yield high-quality results more efficiently and effectively.

## 2.2 Conditioning Generation

Recently, diffusion-based conditional video generation research has begun to emerge, gradually surpassing GAN-based methods (Mirza and Osindero 2014; Wang et al. 2018; Chan et al. 2019; Wang et al. 2019; Liu et al. 2019; Siarohin et al. 2019; Zhou et al. 2022; WANG et al. 2020; Wang et al. 2020, 2022). For the diffusion model-based image generation methods, a lot of works (Mou et al. 2023; Zhang and Agrawala 2023) aim to enhance controllability through the integration of additional annotations. ControlNet (Zhang and Agrawala 2023) duplicates and fixes the original weight of the large pre-trained T2I model. Utilizing the cloned weight, ControlNet trains a conditional branch for task-specific image control.

Recent developments in the field of diffusion-based conditional video generation have been remarkable, branching into two main streams: text-driven video editing, as demon-

strated by (Molad et al. 2023; Esser et al. 2023; Ceylan, Huang, and Mitra 2023; Liu et al. 2023; Wang et al. 2023a; Qi et al. 2023; Hu and Xu 2023), and innovative video creation, featured in works like (Ma et al. 2023; Khachatryan et al. 2023; Hu and Xu 2023; Chen et al. 2023; Zhang et al. 2023). Our work is part of this exciting second stream.

In the realm of video generation, while systems like Follow-Your-Pose (Ma et al. 2023) and Control-A-Video (Chen et al. 2023) are built upon an extensive training process, methods such as Text2Video-Zero (Khachatryan et al. 2023) and ControlVideo (Zhang et al. 2023) align more closely with our approach. A common challenge among these methods, however, is their limited capability in generating dynamic and vibrant backgrounds, a hurdle our methodology overcomes with our unique application of dynamic scene referencing.

### 3 Preliminaries

**Stable Diffusion.** Stable Diffusion employs an autoencoder (Van Den Oord, Vinyals et al. 2017) to preprocess images. An image  $x$  in RGB space is encoded into a latent form by encoder  $\mathcal{E}$  and then decoded back to RGB space by decoder  $\mathcal{D}$ . The diffusion process operates with the encoded latent  $z = \mathcal{E}(x)$ .

For the diffusion forward process, Gaussian noise is iteratively added to latent  $z_0$  over  $T$  iterations (Ho, Jain, and Abbeel 2020):

$$q(z_t | z_{t-1}) = \mathcal{N}\left(z_t; \sqrt{1 - \beta_t}z_{t-1}, \beta_t I\right), \quad (1)$$

$$t = 1, 2, \dots, T,$$

where  $q(z_t | z_{t-1})$  denotes the conditional density function and  $\beta$  is given.

The backward process is accomplished by a well-trained Stable Diffusion model that incrementally denoises the latent variable  $\hat{z}_0$  from the noise  $z_T$ . Typically, the T2I diffusion model leverages a UNet architecture, with text conditions being integrated as supplementary information. The trained diffusion model can also conduct a deterministic forward process, which can be restored back to the original  $z_0$ . This deterministic forward process is referred to as DDIM inversion (Song, Meng, and Ermon 2021; Dhariwal and Nichol 2021). We will refer to  $z_T$  as the noisy latent code and  $z_0$  as the original latent in the subsequent section. Unless otherwise specified, the frames and videos discussed henceforth refer to those in latent space.

**ControlNet.** ControlNet (Zhang and Agrawala 2023) enhances pre-trained large-scale diffusion models by introducing extra input conditions. These inputs are processed by a specially designed conditioning control branch, which originates from a clone of the encoding and middle blocks of the T2I diffusion model and is subsequently trained on task-specific datasets. The output from this control branch is added to the skip connections and the middle block of the T2I model’s UNet architecture.

## 4 Methods

ConditionVideo leverages guided annotation, denoted as *Condition*, and optional reference scenery, denoted as *Video*, to generate realistic videos. We start by introducing our training-free pipeline in Sec. 4, followed by our method for modeling motion in Sec. 4.2. In Sec. 4.3, we present our sparse bi-directional spatial-temporal attention (sBiST-Attn) mechanism. Finally, a detailed explanation of our proposed 3D control branch is provided in Sec. 4.4.

### 4.1 Training-Free Sampling Pipeline

Fig. 2 depicts our proposed training-free sampling pipeline. Inheriting the autoencoder  $\mathcal{D}(\mathcal{E}(\cdot))$  from the pre-trained image diffusion model (Sec. 3), we conduct video transformation between RGB space and latent space frame by frame. Our ConditionVideo model contains two branches: a UNet branch and a 3D control branch. A text description is fed into both branches. Depending on the user’s preference for customized or random background, the UNet branch accepts either the inverted code  $z_T^{INV}$  of the reference background video or the random noise  $\epsilon_b$ . The condition is fed into the 3D control branch after being added with random noise  $\epsilon_c$ . We will further describe this disentanglement input mechanism and random noise  $\epsilon_b, \epsilon_c$  in Sec. 4.2.

Our branch uses the original weight of ControlNet (Zhang and Agrawala 2023). As illustrated on the right side of Fig. 2, we modify the basic spatial-temporal blocks of these two branches from the conditional T2I model by transforming 2D convolution into 3D with  $1 \times 3 \times 3$  kernel and replacing the self-attention module with our proposed sBiST-Attn module (Sec. 4.3). We keep other input-output mechanisms the same as before.

### 4.2 Strategy for Motion Representation

**Disentanglement for Latent Motion Representation** In conventional diffusion models for generation (e.g., ControlNet), the noise vector  $\epsilon$  is sampled from an i.i.d. Gaussian distribution  $\epsilon \sim \mathcal{N}(0, I)$  and then shared by both the control branch and UNet branch. However, if we follow the original mechanism and let the inverse background video’s latent code to shared by two branches, we observe that the background generation results will be blurred (Experiments are shown in Appx. B.). This is because using the same latent to generate both the foreground and the background presumes that the foreground character has a strong relationship with the background. Motivated by this observation, we explicitly disentangle the video motion presentation into two components: the motion of the background and the motion of the foreground. The background motion is generated by the UNet branch whose latent code is presented as background noise  $\epsilon_b \sim \mathcal{N}(0, I)$ . The foreground motion is represented by the given conditional annotations while the appearance representation of the foreground is generated from the noise  $\epsilon_c \sim \mathcal{N}(0, I)$ .

**Strategy for Temporal Consistency Motion Representation** To attain temporal consistency across consecutively generated frames, We investigated selected noise patterns that facilitate the creation of cohesive videos. Consistency

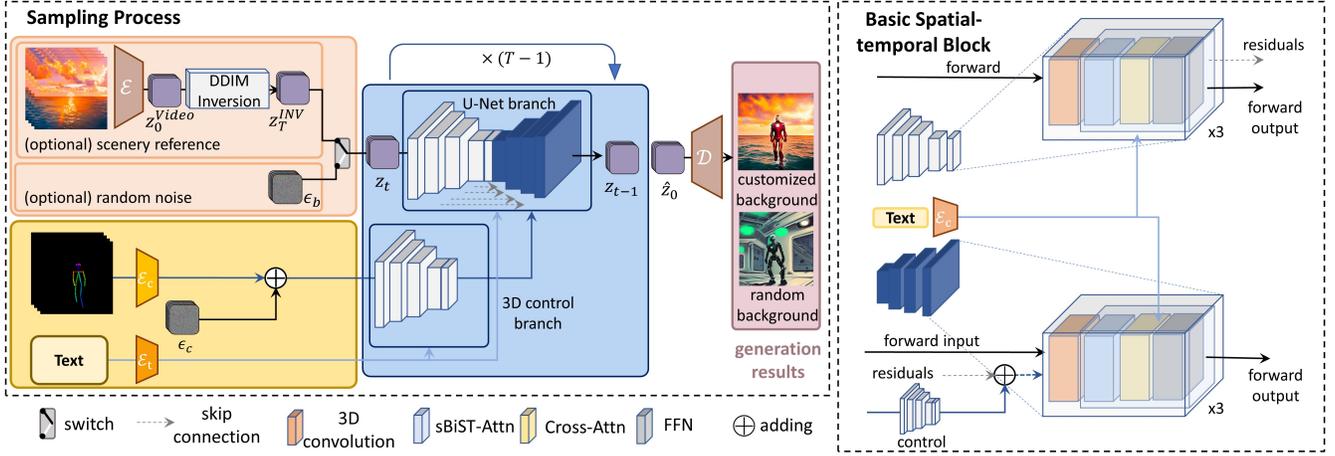


Figure 2: Illustration of our proposed training-free pipeline. (Left) Our framework consists of a UNet branch and a 3D control branch. The UNet branch receives either the inverted reference video  $z_T^{INV}$  or image-level noise  $\epsilon_b$  for background generation. The 3D control branch receives an encoded condition for foreground generation. Text description is fed into both branches. (Right) Illustration of our basic spatial-temporal block. We employ our proposed sBiST-Attn module into the basic block between the 3D convolution block and the cross-attention block. The detail of sBiST-Attn module is shown in Fig. 3

in foreground generation can be established by ensuring that the control branch produces accurate conditional controls. Consequently, we propose utilizing our control branch input for this purpose:  $C_{cond} = \epsilon_c + \mathcal{E}_c(\text{Condition})$ ,  $\epsilon_{c_i} \in \mathcal{N}(0, I) \subseteq \mathbb{R}^{H \times W \times C}$ ,  $\forall i, j = 1, \dots, F$ ,  $\epsilon_{c_i} = \epsilon_{c_j}$ , where  $H$ ,  $W$ , and  $C$  denote the height, width, and channel of the latent  $z_t$ ,  $F$  represents the total frame number,  $C_{cond}$  denotes the encoded conditional vector which will be fed into the control branch and  $\mathcal{E}_c$  denotes the conditional encoder. Additionally, it's important to observe that  $\epsilon_{c_i}$  corresponds to a single frame of noise derived from the video-level noise denoted as  $\epsilon_c$ . The same relationship applies to  $\epsilon_{b_i}$  and  $\epsilon_b$  as well. When generating backgrounds, there are two approaches we could take. The first is to create the background using background noise  $\epsilon_b$ :  $\epsilon_{b_i} \in \epsilon_b$ ,  $\epsilon_{b_i} \sim \mathcal{N}(0, I) \subseteq \mathbb{R}^{H \times W \times C}$ .  $\epsilon_{b_i} = \epsilon_{b_j}$ ,  $\forall i, j = 1, \dots, F$ . The second approach is to generate the background from an inverted latent code,  $z_T^{INV}$ , of the reference scenery video. Notably, we observed that the dynamic motion correlation present in original video is retained when it undergoes DDIM inversion. So we utilize this latent motion correlation to generate background videos.

During the sampling process, in the first forward step  $t = T$ , we feed the background latent code  $z_T^{INV}$  or  $\epsilon_b$  into the UNet branch and the condition  $C_{cond}$  into our 3D control branch. Then, during the subsequent reverse steps  $t = T - 1, \dots, 0$ , we feed the denoised latent  $z_t$  into the UNet branch while still using  $C_{cond}$  for 3D control branch input. The details of the sampling algorithm are shown in Alg. 1

### 4.3 Sparse Bi-directional Spatial-Temporal Attention (sBiST-Attn)

Taking into account both temporal coherence and computational complexity, we propose a sparse bi-directional spatial-

#### Algorithm 1: Sampling Algorithm

**Input:** *Condition, Text, Video*(Optional)

**Parameter:**  $T$

**Output:**  $\hat{X}_0$ : generated video

```

1: if Video is not None then
2:    $z_0^{Video} \leftarrow \mathcal{E}(\text{Video})$  //encode video
3:    $z_T^{INV} \leftarrow \text{DDIM\_Inversion}(z_0^{Video}, T, \text{UNetBranch})$ 
4:    $z_T \leftarrow z_T^{INV}$  //customize background
5: else
6:    $z_T \leftarrow \epsilon_b$ , //random background
7: end if
8:  $C_{cond} \leftarrow \epsilon_c + \mathcal{E}_c(\text{Condition})$  //encode condition
9:  $C_{text} \leftarrow \mathcal{E}_t(\text{Text})$  //encode input prompt
10: for  $t = T \dots 1$  do
11:    $c_t \leftarrow \text{ControlBranch}(C_{cond}, t, C_{text})$ 
12:    $\hat{z}_{t-1} \leftarrow \text{DDIM\_Backward}(z_t, t, C_{text}, c_t, \text{UNetBranch})$ 
13: end for
14:  $\hat{X}_0 \leftarrow \mathcal{D}(\hat{z}_0)$ 
15: return  $\hat{X}_0$ 
    
```

temporal attention (sBiST-Attn) mechanism, as depicted in Fig. 3. For video latent  $z_t^i$ ,  $i = 1, \dots, F$ , the attention matrix is computed between frame  $z_t^i$  and its bi-directional frames, sampled with a gap of 3. This interval was chosen after weighing frame consistency and computational cost (see Appx. C.1). For each  $z_t^i$  in  $z_t$ , we derive the query feature from its frame  $z_t^i$ . The key and value features are derived from the bi-directional frames  $z_t^{3j+1}$ ,  $j = 0, \dots, \lfloor (F - 1)/3 \rfloor$ . Mathematically, our sBiST-Attn can be expressed as:

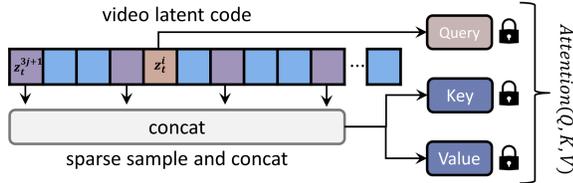


Figure 3: Illustration of ConditionVideo’s sBiST-Attn. The purple blocks signify the frame we’ve selected for concatenation, which can be computed for key and value. The pink block represents the current block from which we’ll calculate the query. The blue blocks correspond to the other frames within the video sequence. Latent features of frame  $z_t^i$ , bi-directional frames  $z_t^{3j+1}$ ,  $j = 0, \dots, \lfloor (F - 1)/3 \rfloor$  are projected to query  $Q$ , key  $K$  and value  $V$ . Then the attention-weighted sum is computed based on key, query, and value. The parameters are the same as the ones in the self-attention module of the pre-trained image model.

$$\begin{cases} \text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d}}\right) \cdot V \\ Q = W^Q z_t^i, K = W^K z_t^{[3j+1]}, V = W^V z_t^{[3j+1]}, \\ j = 0, 1, \dots, \lfloor (F - 1)/3 \rfloor \end{cases} \quad (2)$$

where  $[\cdot]$  denotes the concatenation operation, and  $W^Q, W^K, W^V$  are the weighted matrices that are identical to those used in the self-attention layers of the image generation model.

#### 4.4 3D Control Branch

Frame-wise conditional guidance is generally effective, but there may be instances when the network doesn’t correctly interpret the guide, resulting in an inconsistent conditional output. Given the continuous nature of condition movements, ConditionVideo propose enhancing conditional alignment by referencing neighboring frames. If a frame isn’t properly aligned due to weak control, other correctly aligned frames can provide more substantial conditional alignment information. In light of this, we design our control branch to operate temporally, where we choose to replace the self-attention module with the sBiST-Attn module and inflate 2D convolution to 3D. The replacing attention module can consider both previous and subsequent frames, thereby bolstering our control effectiveness.

### 5 Experiments

#### 5.1 Implementation Details

We implement our model based on the pre-trained weights of ControlNet (Zhang and Agrawala 2023) and Stable Diffusion (Rombach et al. 2022) 1.5. We generate 24 frames with a resolution of  $512 \times 512$  pixels for each video. During inference, we use the same sampling setting as Tune-A-Video (Wu et al. 2022b). More details can be found in Appx. D at <https://arxiv.org/abs/2310.07697>.

#### 5.2 Main results

In Fig. 1, we display the success of our training-free video generation technique. The generated results from ConditionVideo, depicted in Fig. 1 (a), imitate moving scenery videos and show realistic waves as well as generate the correct character movement based on posture. Notably, the style of the backgrounds is distinct from the original guiding videos, while the motion of the backgrounds is the same. Furthermore, our model can generate consistent backgrounds when sampling  $\epsilon_b$  from Gaussian noise based on conditional information, as shown in Fig.1 (b),(c),(d). These videos showcase high temporal consistency and rich graphical content.

#### 5.3 Comparison

**Compared Methods** We compare our method with Tune-A-Video (Wu et al. 2022b), ControlNet (Zhang and Agrawala 2023), and Text2Video-Zero (Khachatryan et al. 2023). For Tune-A-Video, we first fine-tune the model on the video from which the condition was extracted, and then sample from the corresponding noise latent code of the condition video.

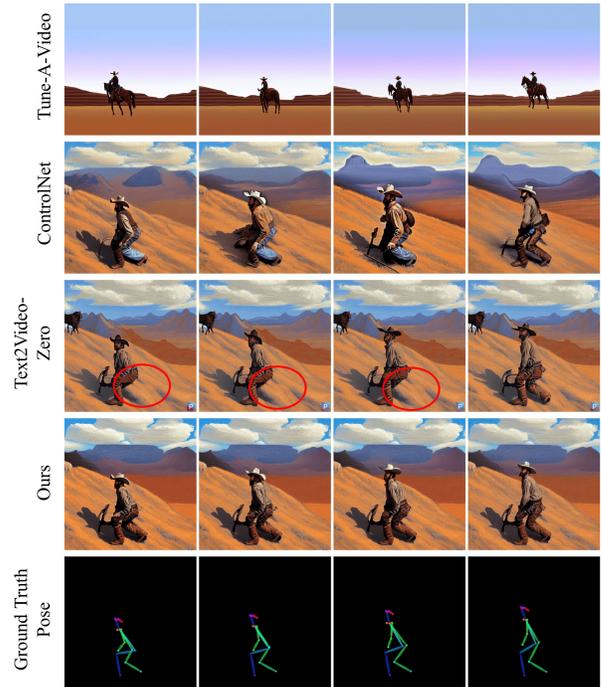


Figure 4: Qualitative comparison condition on the pose. “The Cowboy, on a rugged mountain range, Western painting style”. Our result outperforms in both temporal consistency and pose accuracy, while others have difficulty in maintaining either one or both of the qualities.

**Qualitative Comparison** Our visual comparison conditioning on pose, canny, and depth information is presented in Fig. 4, 5, and 6. Tune-A-Video struggles to align well with our given condition and text description. ControlNet demonstrates improvement in condition-alignment accuracy

but suffers from a lack of temporal consistency. Despite the capability of Text2Video to produce videos of exceptional quality, there are still some minor imperfections that we have identified and indicated using a red circle in the figure. Our model surpasses all others, showcasing outstanding condition-alignment quality and frame consistency.

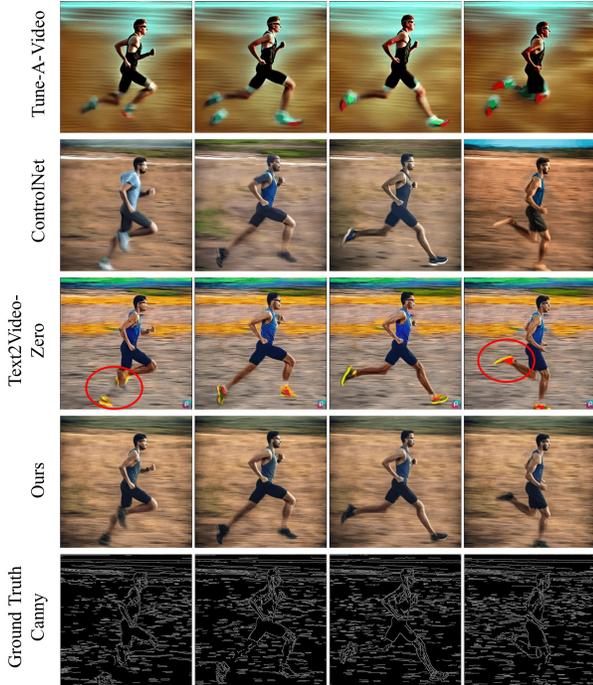


Figure 5: Qualitative comparison condition on canny. “A man is runnin”. Tune-A-Video experiences difficulties with canny-alignment, while ControlNet struggles to maintain temporal consistency. Though Text2Video surpasses these first two approaches, it inaccurately produces parts of the legs that don’t align with the actual human body structure, and the colors of the shoes it generates are inconsistent.

Method	FC(%)	CS	PA (%)
Tune-A-Video	95.84	30.74	26.13
ControlNet	94.22	32.97	79.51
Text2Video-Zero	98.82	32.84	78.50
Ours	<b>99.02</b>	<b>33.03</b>	<b>83.12</b>

Table 1: Quantitative comparisons condition on pose. FC, CS, PA represent *frame consistency*, *clip score* and *pose-accuracy*, respectively

**Quantitative Comparison** We evaluate all the methods using three metrics: *frame consistency* (Esser et al. 2023; Wang et al. 2023a; Radford et al. 2021), *clip score* (Ho et al. 2022a; Hessel et al. 2021; Park et al. 2021), and *pose accuracy* (Ma et al. 2023). As other conditions are hard to evaluate, we use pose accuracy for conditional consistency only. The results on different conditions are shown in Tab. 1 and

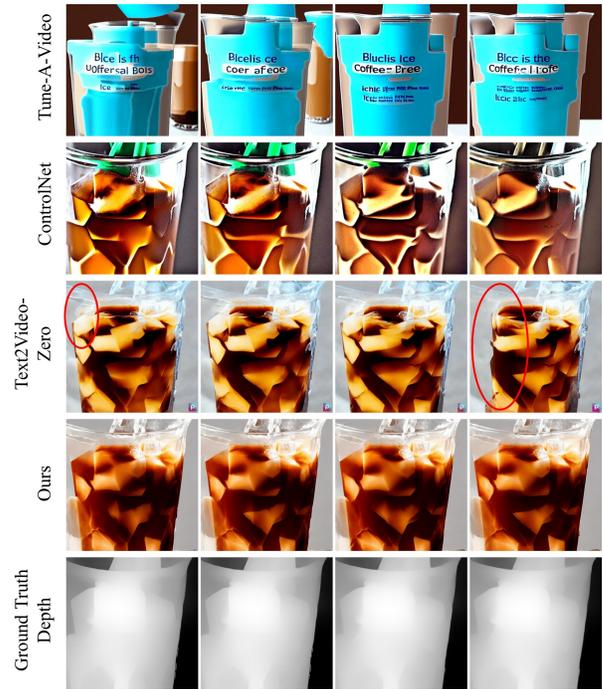


Figure 6: Qualitative comparison condition on depth. “ice coffee”. All three methods used for comparison have the problem of changing the appearance of the object when the viewpoint is switched, and only our method ensures the consistency of the appearance before and after.

2. We achieve the highest frame consistency, and clip score in all conditions, indicating that our method exhibits the best time consistency and text alignment. We also have the best pose-video alignment among the other three techniques. The conditions are randomly generated from a group of 120 different videos. For more information please see Appx. D.2.

### 5.4 Ablation Study

We conduct an ablation study on the pose condition, temporal module, and 3D control branch. The qualitative result is visualized in Fig. 7. In our research, we modify each element individually for comparative analysis, ensuring that all other settings remain constant.

**Ablation on Pose Condition** We evaluate performance with and without using pose, as shown in Fig. 7. Without pose conditioning, the video is fixed as an image, while the use of pose control allows for the generation of videos with certain temporal semantic information.

**Ablation on Temporal Module** Training-free video generation heavily relies on effective spatial-temporal modeling. To evaluate the efficacy of our temporal attention module, We remove our sBiST-attention mechanism and replace it with a non-temporal self-attention mechanism, a Sparse-Causal attention mechanism (Wu et al. 2022b) and a dense attention mechanism (Wang et al. 2023a) which attends to all frames for key and value.

Method	Condition	FC(%)	CS
Tune-A-Video	-	95.84	30.74
ControlNet	Canny	90.53	29.65
Text2Video-Zero	Canny	97.44	28.76
Ours	Canny	<b>97.64</b>	<b>29.76</b>
ControlNet	Depth	90.63	30.16
Text2Video-Zero	Depth	97.46	29.38
Ours	Depth	<b>97.65</b>	<b>30.54</b>
ControlNet	Segment	91.87	31.85
Ours	Segment	<b>98.13</b>	<b>32.09</b>

Table 2: Quantitative comparisons condition on canny, depth and segment.

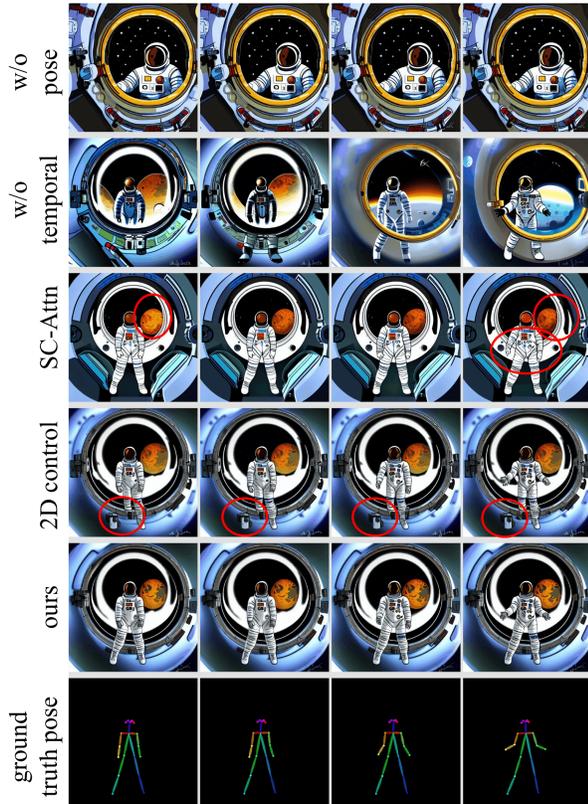


Figure 7: Ablations of each component, generated from image-level noise. “The astronaut, in a spacewalk, sci-fi digital art style”. 1st row displays the generation result without pose conditioning. 2nd and 3rd rows show the results after replacing our sBiST-Attn with self-Attn and SC-Attn (Wu et al. 2022b). 4th row presents the result with the 2D condition-control branch.

The results are presented in Tab. 3. A comparison of temporal and non-temporal attention underlines the importance of temporal modeling for generating time-consistent videos. By comparing our method with Sparse Causal attention, we demonstrate the effectiveness of ConditionVideo’s sBiST at-

Method	FC(%)	Time
w/o Temp-Attn	94.22	<b>31s</b>
S-C Attn	98.77	43s
sBiST-Attn	99.02	1m30s
Full-Attn	<b>99.03</b>	3m37s

Table 3: Ablations on temporal module. Time represents the duration required to generate a 24-frame video with a size of 512x512.

Method	FC(%)	CS (%)	PA (%)
2D control	<b>99.03</b>	<b>33.11</b>	81.26
3D control	99.02	33.03	<b>83.12</b>

Table 4: Ablation on 3D control branch. FC, CS, PA represent frame consistency, clip score, and pose-accuracy, respectively.

tention module, proving that incorporating information from bi-directional frames improves performance compared to using only previous frames. Furthermore, we observe almost no difference in frame consistency between our method and dense attention, despite the latter requiring more than double our generation duration.

**Ablations on 3D Control Branch** We compare our 3D control branch with a 2D version that processes conditions frame-by-frame. For the 2D branch, we utilize the original ControlNet conditional branch. Both control branches are evaluated in terms of frame consistency, clip score, and pose accuracy. Results in Tab. 4 show that our 3D control branch outperforms the 2D control branch in pose accuracy while maintaining similar frame consistency and clip scores. This proves that additional consideration of bi-directional frames enhances pose control.

## 6 Discussion and Conclusion

In this paper, we propose ConditionVideo, a training-free method for generating videos with vivid motion. This technique leverages a unique motion representation, informed by background video and conditional data, and utilizes our sBiST-Attn mechanism and 3D control branch to enhance frame consistency and condition alignment. Our experiments show that ConditionVideo can produce high-quality videos, marking a significant step forward in video generation and AI-driven content creation.

During our experiments, we find that our method is capable of generating long videos. Moreover, this approach is compatible with the hierarchical sampler from ControlVideo (Zhang et al. 2023), which is used for generating long videos. Despite the effectiveness of condition-based and temporal attention in maintaining video coherence, challenges such as flickering in videos with sparse conditions like pose data were noted. To address this issue, a potential solution would involve incorporating more densely sampled control inputs and additional temporal-related structures.

## Acknowledgements

This work was jointly supported by the National Key R&D Program of China(NO.2022ZD0160100) and the National Natural Science Foundation of China under Grant No. 62102150.

## References

- Bertasius, G.; Wang, H.; and Torresani, L. 2021. Is space-time attention all you need for video understanding? In *ICML*, volume 2.
- Blattmann, A.; Rombach, R.; Ling, H.; Dockhorn, T.; Kim, S. W.; Fidler, S.; and Kreis, K. 2023. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Ceylan, D.; Huang, C. P.; and Mitra, N. J. 2023. Pix2Video: Video Editing using Image Diffusion. *CoRR*, abs/2303.12688.
- Chan, C.; Ginosar, S.; Zhou, T.; and Efros, A. A. 2019. Everybody dance now. In *Proceedings of the IEEE/CVF international conference on computer vision*.
- Chen, W.; Wu, J.; Xie, P.; Wu, H.; Li, J.; Xia, X.; Xiao, X.; and Lin, L. 2023. Control-A-Video: Controllable Text-to-Video Generation with Diffusion Models. *CoRR*, abs/2305.13840.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34.
- Esser, P.; Chiu, J.; Atighehchian, P.; Granskog, J.; and Germanidis, A. 2023. Structure and content-guided video synthesis with diffusion models. *arXiv preprint arXiv:2302.03011*.
- Gafni, O.; Polyak, A.; Ashual, O.; Sheynin, S.; Parikh, D.; and Taigman, Y. 2022. Make-a-scene: Scene-based text-to-image generation with human priors. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XV*. Springer.
- Hessel, J.; Holtzman, A.; Forbes, M.; Le Bras, R.; and Choi, Y. 2021. CLIPScore: A Reference-free Evaluation Metric for Image Captioning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Ho, J.; Chan, W.; Saharia, C.; Whang, J.; Gao, R.; Gritsenko, A.; Kingma, D. P.; Poole, B.; Norouzi, M.; Fleet, D. J.; et al. 2022a. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33.
- Ho, J.; Salimans, T.; Gritsenko, A. A.; Chan, W.; Norouzi, M.; and Fleet, D. J. 2022b. Video Diffusion Models. In *NeurIPS*.
- Hong, W.; Ding, M.; Zheng, W.; Liu, X.; and Tang, J. 2023. CogVideo: Large-scale Pretraining for Text-to-Video Generation via Transformers. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Hu, Z.; and Xu, D. 2023. VideoControlNet: A Motion-Guided Video-to-Video Translation Framework by Using Diffusion Model with ControlNet. *arXiv:2307.14073*.
- Khachatryan, L.; Movsisyan, A.; Tadevosyan, V.; Henschel, R.; Wang, Z.; Navasardyan, S.; and Shi, H. 2023. Text2Video-Zero: Text-to-Image Diffusion Models are Zero-Shot Video Generators. *arXiv preprint arXiv:2303.13439*.
- Kingma, D.; Salimans, T.; Poole, B.; and Ho, J. 2021. Variational diffusion models. *Advances in neural information processing systems*, 34.
- Liu, S.; Zhang, Y.; Li, W.; Lin, Z.; and Jia, J. 2023. Video-P2P: Video Editing with Cross-attention Control. *arXiv preprint arXiv:2303.04761*.
- Liu, W.; Piao, Z.; Min, J.; Luo, W.; Ma, L.; and Gao, S. 2019. Liquid warping gan: A unified framework for human motion imitation, appearance transfer and novel view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Ma, Y.; He, Y.; Cun, X.; Wang, X.; Shan, Y.; Li, X.; and Chen, Q. 2023. Follow Your Pose: Pose-Guided Text-to-Video Generation using Pose-Free Videos. *arXiv preprint arXiv:2304.01186*.
- Mirza, M.; and Osindero, S. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*.
- Molad, E.; Horwitz, E.; Valevski, D.; Acha, A. R.; Matias, Y.; Pritch, Y.; Leviathan, Y.; and Hoshen, Y. 2023. Dreamix: Video diffusion models are general video editors. *arXiv preprint arXiv:2302.01329*.
- Mou, C.; Wang, X.; Xie, L.; Zhang, J.; Qi, Z.; Shan, Y.; and Qie, X. 2023. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *arXiv preprint arXiv:2302.08453*.
- Park, D. H.; Azadi, S.; Liu, X.; Darrell, T.; and Rohrbach, A. 2021. Benchmark for compositional text-to-image synthesis. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- Qi, C.; Cun, X.; Zhang, Y.; Lei, C.; Wang, X.; Shan, Y.; and Chen, Q. 2023. Fatezero: Fusing attentions for zero-shot text-based video editing.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. PMLR.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1).

- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*. Springer.
- 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. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35.
- Siarohin, A.; Lathuilière, S.; Tulyakov, S.; Ricci, E.; and Sebe, N. 2019. First order motion model for image animation. *Advances in Neural Information Processing Systems*, 32.
- Singer, U.; Polyak, A.; Hayes, T.; Yin, X.; An, J.; Zhang, S.; Hu, Q.; Yang, H.; Ashual, O.; Gafni, O.; Parikh, D.; Gupta, S.; and Taigman, Y. 2023. Make-A-Video: Text-to-Video Generation without Text-Video Data. In *The Eleventh International Conference on Learning Representations*.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*. PMLR.
- Song, J.; Meng, C.; and Ermon, S. 2021. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations*.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.
- Van Den Oord, A.; Vinyals, O.; et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*.
- Wang, T.-C.; Liu, M.-Y.; Tao, A.; Liu, G.; Kautz, J.; and Catanzaro, B. 2019. Few-shot Video-to-Video Synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Wang, T.-C.; Liu, M.-Y.; Zhu, J.-Y.; Liu, G.; Tao, A.; Kautz, J.; and Catanzaro, B. 2018. Video-to-Video Synthesis. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Wang, W.; Xie, K.; Liu, Z.; Chen, H.; Cao, Y.; Wang, X.; and Shen, C. 2023a. Zero-Shot Video Editing Using Off-The-Shelf Image Diffusion Models. *arXiv preprint arXiv:2303.17599*.
- Wang, Y.; Bilinski, P.; Bremond, F.; and Dantcheva, A. 2020. G3AN: Disentangling Appearance and Motion for Video Generation. In *CVPR*.
- WANG, Y.; Bilinski, P.; Bremond, F.; and Dantcheva, A. 2020. ImaGINator: Conditional Spatio-Temporal GAN for Video Generation. In *WACV*.
- Wang, Y.; Chen, X.; Ma, X.; Zhou, S.; Huang, Z.; Wang, Y.; Yang, C.; He, Y.; Yu, J.; Yang, P.; et al. 2023b. LAVIE: High-Quality Video Generation with Cascaded Latent Diffusion Models. *arXiv preprint arXiv:2309.15103*.
- Wang, Y.; Ma, X.; Chen, X.; Dantcheva, A.; Dai, B.; and Qiao, Y. 2023c. LEO: Generative Latent Image Animator for Human Video Synthesis. *arXiv preprint arXiv:2305.03989*.
- Wang, Y.; Yang, D.; Bremond, F.; and Dantcheva, A. 2022. Latent Image Animator: Learning to Animate Images via Latent Space Navigation. In *ICLR*.
- Wu, C.; Huang, L.; Zhang, Q.; Li, B.; Ji, L.; Yang, F.; Sapiro, G.; and Duan, N. 2021. Godiva: Generating open-domain videos from natural descriptions. *arXiv preprint arXiv:2104.14806*.
- Wu, C.; Liang, J.; Ji, L.; Yang, F.; Fang, Y.; Jiang, D.; and Duan, N. 2022a. Nüwa: Visual synthesis pre-training for neural visual world creation. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVI*. Springer.
- Wu, J. Z.; Ge, Y.; Wang, X.; Lei, W.; Gu, Y.; Hsu, W.; Shan, Y.; Qie, X.; and Shou, M. Z. 2022b. Tune-A-Video: One-Shot Tuning of Image Diffusion Models for Text-to-Video Generation. *arXiv preprint arXiv:2212.11565*.
- Zhang, L.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. *arXiv preprint arXiv:2302.05543*.
- Zhang, Y.; Wei, Y.; Jiang, D.; Zhang, X.; Zuo, W.; and Tian, Q. 2023. ControlVideo: Training-free Controllable Text-to-Video Generation. *arXiv preprint arXiv:2305.13077*.
- Zhou, X.; Yin, M.; Chen, X.; Sun, L.; Gao, C.; and Li, Q. 2022. Cross attention based style distribution for controllable person image synthesis. In *European Conference on Computer Vision*, 161–178. Springer.