

# Llama Learns to Direct: DirectorLLM for Human-Centric Video Generation

Kunpeng Song<sup>1,2,\*</sup>, Tingbo Hou<sup>1</sup>, Zecheng He<sup>1</sup>, Haoyu Ma<sup>1</sup>, Jialiang Wang<sup>1</sup>, Animesh Sinha<sup>1</sup>, Sam Tsai<sup>1</sup>, Yaqiao Luo<sup>1</sup>, Xiaoliang Dai<sup>1</sup>, Li Chen<sup>1</sup>, Xide Xia<sup>1</sup>, Peizhao Zhang<sup>1</sup>, Peter Vajda<sup>1</sup>, Ahmed Elgammal<sup>2</sup>, Felix Juefei-Xu<sup>1</sup>

<sup>1</sup>GenAI at Meta, <sup>2</sup>Rutgers University

\*Work done at Meta


In this paper, we introduce **DirectorLLM**, a novel video generation model that employs a large language model (LLM) as the “director” to simulate human poses within videos. To enhance the authenticity of human motions in text-to-video models, we extend the LLM from a text generator to a video director and human motion simulator. We train the DirectorLLM to generate detailed human poses, to guide video generation, offloading the simulation of human motion from the video generator to an LLM, effectively creating informative outlines for human-centric scenes. These signals are used as conditions by the video renderer, facilitating more realistic and prompt-following video generation. Experiments on automatic evaluation benchmarks and human evaluations show that our model outperforms existing ones in generating videos with higher human motion fidelity, improved prompt faithfulness, and enhanced rendered subject naturalness.

**Date:** September 9, 2025



## 1 Introduction

The fields of image and video generation have witnessed remarkable advancements, largely due to the emergence of diffusion-based models like DALL-E 2 [Ramesh et al. \(2022\)](#), Imagen [Saharia et al. \(2022\)](#), Stable Diffusion [Rombach et al. \(2022\)](#), Emu [Dai et al. \(2023\)](#), eDiff-I [Balaji et al. \(2022\)](#), Kandinsky [Razzhigaev et al. \(2023\)](#), Stable Unclip [AI \(2024\)](#). These models empower users to create diverse and realistic visuals from text prompts. Building on these advances, video generation models such as VideoCrafter [Chen et al. \(2023\)](#) and LaVie [Wang et al. \(2023b\)](#), Sora [OpenAI \(2024\)](#), CogVideoX [Hong et al. \(2022\)](#), and Movie Gen [Polyak et al. \(2024\)](#) have demonstrated unprecedented capacities, achieving longer and high-quality videos across various resolutions and aspect ratios. Methods like VideoComposer [Wang et al. \(2024b\)](#) and DreamVideo [Wei et al. \(2024\)](#) integrate video generation with various control signals. Although these techniques significantly improve content generation and customization, they often face limitations in understanding human poses and motion



**Figure 1** In the real-world movie production, the director works closely with actors to shape their portrayals, providing direction on how to deliver lines, express emotions, and move within a scene. Our vision is to **train** a large language model to assume the role of a **director** for human-centric video generation. The pipeline consists of two main parts: the DirectorLLM and the video renderer (cinematographer). The DirectorLLM, based on Llama 3, handles high-level scene understanding, predicting subject’s location, motion, pose, and interactions within the scene based on the user input text prompt. The video renderer then generates realistic and temporally consistent video frames conditioned on the guidance provided by the DirectorLLM. In this paper, we apply this concept to human-centric text-to-video generation.

dynamics across extended video durations.

Recent advancements in text-to-image and text-to-video generation leverage large language models (LLMs) to improve the spatial and temporal coherence of generated content. Generating complex visual dynamics directly from text prompts remains a significant challenge for diffusion models, especially when maintaining coherent spatial and temporal relationships. To address this, recent research LVD Lian et al. (2023), LayoutGPT Feng et al. (2024), attention-refocusing Phung et al. (2024), and VideoDirectorGPT Lin et al. (2023) have explored using LLMs like GPT-4 Achiam et al. (2023) and PaLM2 Anil et al. (2023) to produce explicit spatiotemporal layouts, such as bounding boxes or frame-by-frame descriptions, derived from a single user prompt. These layouts then serve as a foundation for conditioning diffusion models, enhancing their ability to produce images and videos with coherent object relationships across frames. By transforming the narrative input into precise bounding boxes and layout structures, these methods rely on off-the-shelf LLM to simulate subject dynamics, improving motion consistency and dynamic realism.

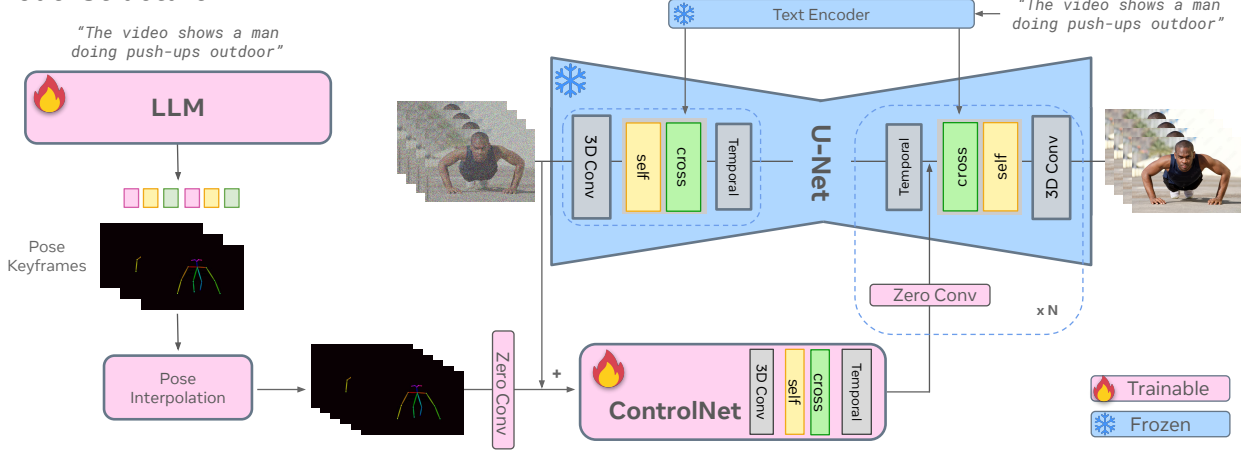
In this work, we present **DirectorLLM**, a novel T2V model designed to tackle the challenge of creating realistic and dynamic human-centric videos with improved human motions. Existing methods Chen et al. (2023); Wang et al. (2023b) generally rely on diffusion models to synthesize videos directly from text prompts, requiring the model to simultaneously manage scene understanding, object interactions, and video rendering, often resulting in anatomically inaccurate human body outputs, such as broken or extra limbs, unnatural joint movements, or abrupt shifts in human motion. To address these limitations, we decouple scene understanding and pose simulation from the video generation process. Our approach consists of three main components: a specialized large language model (LLM) that interprets the text prompt to generate human poses at 1 frame per second (FPS), a compact linear diffusion model that interpolates these sparse poses to a smooth 30 FPS, enhancing temporal consistency, and a video generator (VideoCrafter Chen et al. (2023)) enhanced with pose ControlNet Zhang et al. (2023) to render realistic videos conditioned on both the predicted human poses and text prompt. This structured framework allows the LLM to handle scene understanding and human motion simulation, offloading these tasks from the video generator, which can then focus on creating visually accurate frames. By integrating these specialized components, our model achieves improved realism and coherence in human-centric video generation compared to prior methods.

Our contributions are as follows: (i) To our knowledge, this is the first work to **train** a large language model (e.g., Llama 3 Dubey et al. (2024)), residual VQ-VAE Van Den Oord et al. (2017), and linear diffusion for simulating and generating human poses, specifically tailored for text-to-video generation task. (ii) We enhance diffusion-based text-to-video (T2V) synthesis by integrating predicted human motion, leading to improved realism and coherence in human-centric videos. (iii) Our system advances the current state of T2V synthesis, particularly in terms of human motion realism and temporal consistency.

## 1.1 Related Work

**Text-to-video (T2V) generation.** Text-to-video (T2V) generation leverages deep learning models to interpret text input and generate video content, and it’s been evolving rapidly. Notable models, such as ModelScopeT2V Wang et al. (2023a) and LaVie Wang et al. (2023b), enhance video synthesis by integrating temporal layers within spatial frameworks to achieve more cohesive video output. VideoCrafter2 Chen et al. (2023), for instance, tackles the scarcity of labeled video data by making use of high-quality image datasets to improve training. Transformer-based models like Sora OpenAI (2024) have achieved impressive results in generating videos with high visual quality, stable temporal consistency, and varied motion dynamics. Moreover, Movie Gen Polyak et al. (2024) pushes the boundaries further, achieving exceptional quality in both video and audio generation. Recent work has explored the incorporation of explicit motion controls. Notable works include AnimateDiff Guo et al. (2023), VideoComposer Wang et al. (2024b), CameraCtrl He et al. (2024a), Direct-A-Video Yang et al. (2024), and MotionCtrl Wang et al. (2024c), which design specific modules to encode camera trajectories. Trajectory-based motion guidance is used in works such as DragNUWA Yin et al. (2023), MotionCtrl Wang et al. (2024c), Motion-I2V Shi et al. (2024), Tora Zhang et al. (2024), and DragAnything Wu et al. (2025), while box-based motion guidance is used by TrailBlazer Ma et al. (2023) and Boximator Wang et al. (2024a). Additionally, methods like MotionDirector Zhao et al. (2023) use reference videos to extract motion patterns for generation. In contrast to these models, which require user-provided motion signals, our model takes care of generating human poses, enabling fully text-driven human-centric

## Model Structure



**Figure 2** Overall design of our model architecture. The DirectorLLM processes the text prompt to generate instance-level bounding boxes and human skeletons, effectively acting as a layout and motion planner. The linear diffusion module interpolates human skeletons, enhancing smoothness and consistency before final rendering. These outputs guide the video generation process in a U-Net-based video generator which employs ControlNet for fine-grained pose control.

video generation.

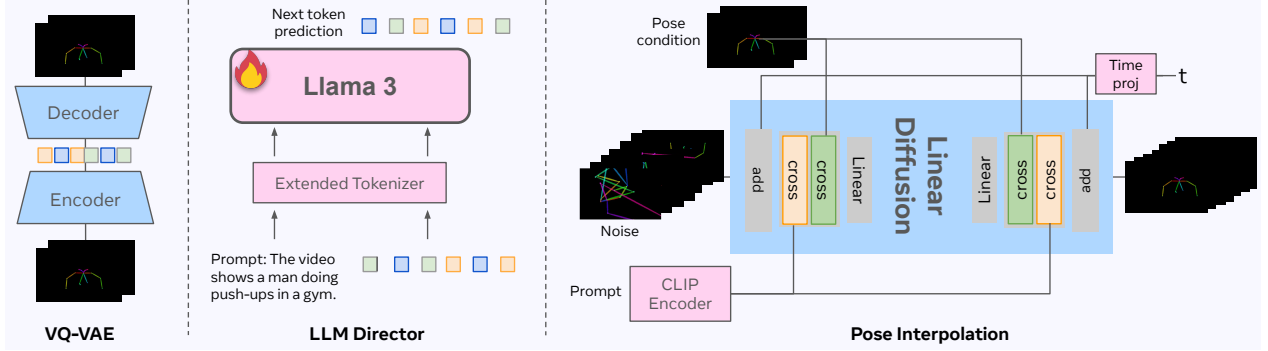
**LLM in image and video generation.** Several recent text-to-image models have incorporated large language models (LLMs) to assist image and video generation. Methods like LayoutGPT [Feng et al. \(2024\)](#) employ a program-guided method for layout-oriented visual planning across various domains. Additionally, Attention Refocusing [Phung et al. \(2024\)](#) introduces innovative loss functions to realign attention maps according to specified layouts. Building upon these developments, recent efforts have extended the use of LLMs to assist text-to-video models. Free-bloom [Huang et al. \(2024\)](#) leverages LLMs to generate detailed frame-by-frame descriptions from a single prompt, enriching the video’s narrative. LLM-grounded video generation [Lian et al. \(2023\)](#) and VideoDirectorGPT [Lin et al. \(2023\)](#) propose a training-free approach to guide diffusion models using LLM-generated frame-by-frame layouts. GPT-4-Motion [Lv et al. \(2024\)](#) uses GPT-4 [Achiam et al. \(2023\)](#) to generate scripts for Blender to produce scene components, which serve as conditions for Stable Diffusion to synthesize videos. Unlike previous work that uses pre-trained and frozen language models to generate layout and bounding boxes from text prompts, our model fine-tunes a large language model, Llama 3 [Dubey et al. \(2024\)](#), to predict vector-quantized tokens to simulate human poses, enabling more accurate motion understanding tailored for video generation.

## 2 Method

### 2.1 Preliminaries

**Text-to-video.** T2V diffusion models are designed to generate video content by adding noise in a latent space and then learning to denoise it. The process starts by encoding an input video into a latent representation,  $\mathbf{z}_0$ . Noise, represented by  $\epsilon$ , is incrementally added to this latent representation over time steps  $t$ , creating a noisy latent  $\mathbf{z}_t$  that corresponds to different noise levels. This progression simulates a reverse Markov chain [Rombach et al. \(2022\)](#) process, where the diffusion model,  $\epsilon_\theta$ , is trained to predict and remove the noise added to  $\mathbf{z}_t$  as it transitions back toward the clean latent representation. The model is trained using a reconstruction loss that minimizes the difference between the actual noise  $\epsilon$  and the predicted noise  $\epsilon_\theta(\mathbf{z}_t, \mathbf{c}, t)$ , where  $\mathbf{c}$  represents a conditional input, such as text or an image used for guidance. This is expressed mathematically as:

$$\mathcal{L} = \mathbb{E}_{\mathbf{z}, \epsilon \sim \mathcal{N}(0, \mathbf{I}), t, \mathbf{c}} \left[ \|\epsilon - \epsilon_\theta(\mathbf{z}_t, \mathbf{c}, t)\|_2^2 \right] \quad (1)$$



**Figure 3** Overview of DirectorLLM. **Left:** The VQ-VAE module encodes human poses into discrete tokens. **Middle:** The Llama 3-based DirectorLLM interprets text prompts to generate bounding boxes and pose tokens for each instance in the scene. **Right:** The linear diffusion model interpolates sparse frames to dense sequences, integrating pose conditions and CLIP-encoded text prompts, producing smooth poses.

**Motion generation.** Prior works on diffusion have explored audio to dance Tseng et al. (2023), text to motion Zhang et al. (2022); Jiang et al. (2023), audio to gestures Alexanderson et al. (2023); Ao et al. (2023); Yu et al. (2023) and speech to human pose Ng et al. (2024). Audio2Photoreal Ng et al. (2024) presents a avatar pose generation systems conditioned on audio inputs. Given speech audio, it outputs multiple possibilities of gestural motion for an individual, including face, body, and hands. Unlike these pose models, our method focuses on generating dynamic human poses directly from text prompts at the context of text to video generation, leveraging a fine-tuned large language model for detailed and contextually accurate motion planning.

## 2.2 Methodology

Our model consists of three main stages: a large language model (*e.g.*, tuned from Llama 3 Dubey et al. (2024)) generates discretized human pose tokens from text prompts at low FPS for human motion planning, followed by a linear diffusion model for pose interpolation. The final stage uses a video generation model (VideoCrafter2 Chen et al. (2023)) enhanced by ControlNet Zhang et al. (2023) to produce realistic human-centric videos.

**DirectorLLM.** Our DirectorLLM, built on Llama 3 Dubey et al. (2024), plays a crucial role in scene understanding and human pose generation planning, serving as a high-level director that interprets the text prompt to produce structured, instance-level layouts for human motion. Given a prompt, the DirectorLLM identifies content of the requested video in the scene and predicts corresponding human pose tokens at 1 fps. To make these outputs compatible with our generative pipeline, we convert continuous human pose vectors into discrete tokens via a residual vector quantized autoencoder (VQ-VAE Van Den Oord et al. (2017)), enabling the LLM to process human pose information as token sequences. Input sequences are  $z = (t, p)$ , in which  $t = (t_1, t_2, \dots, t_n)$  are text prompt tokens, and  $p = (p_1, p_2, \dots, p_n)$  are pose tokens. The model is trained on cross-entropy loss using the standard next-token prediction framework,

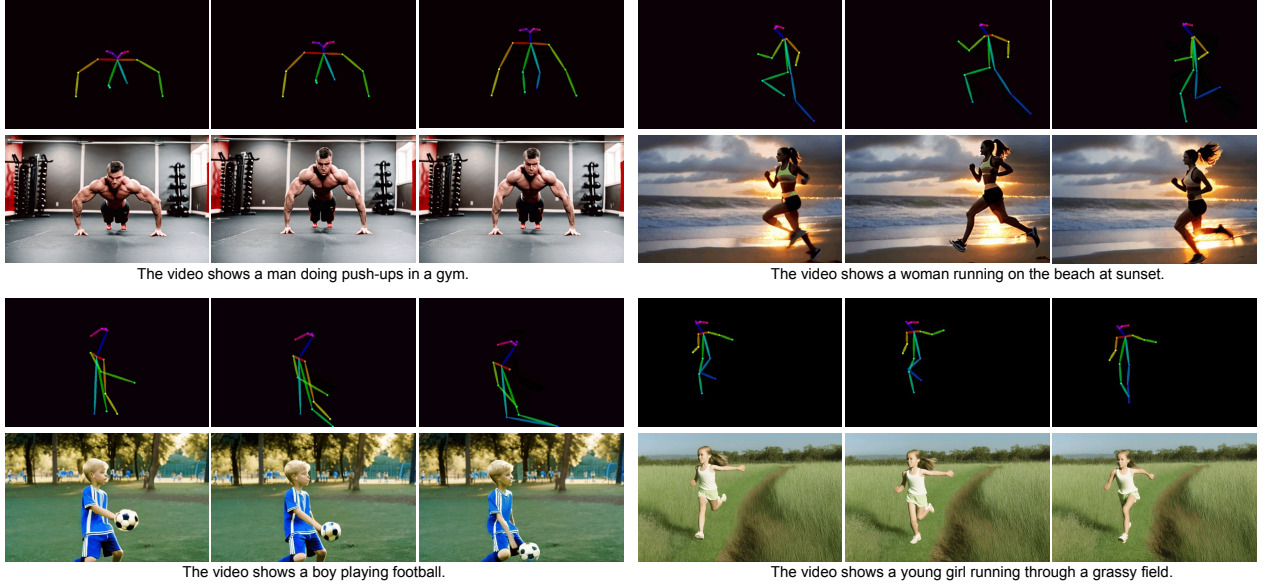
$$L = - \sum_{t=1}^n \log [p(z_{0:t-1}) = z_t] \quad (2)$$

where  $p(z_{0:t-1})$  represents the predicted probability given the preceding tokens.

Our LLM learns human pose dynamics through extensive text and video data. We show its structure in Fig. 8 middle part. By embedding human poses as tokenized predictions, our DirectorLLM effectively captures the relationship between complex human dynamics and text prompts, enabling control over movement and interactions in the generated video. This component predicts 1 FPS human motions that align closely with the intended actions in the text prompt, which will be further enhance by diffusion pose interpolator.

**Diffusion interpolator.** We train a compact linear diffusion model to densify sparse poses generated by LLM into smooth and temporally consistent outputs at 30 FPS. As shown in Fig. 8 right part, this model enhances the initial sparse human poses by generating intermediate frames that maintain natural motion and dynamics.





**Figure 4** Examples of generated poses and videos frames from DirectorLLM. Poses are generated from the LLM and interpolated by pose diffusion model. We show more pose-video sequences in the supplementary material.

We incorporate conditional pose information through cross-attention layers, allowing the model to refine pose sequences frame-by-frame based on both structural and temporal cues. To improve fidelity further, we condition the model on CLIP [Radford et al. \(2021\)](#) embeddings of the text prompt, enhancing alignment between the generated dynamics and the scene description. Additionally, we employ classifier-free guidance [Ho and Salimans \(2022\)](#) by intermittently replacing conditions with zero tensors during training, promoting the model’s ability to synthesize realistic sequences even under partial guidance. This linear diffusion process enables seamless transition from sparse layouts to dense, high-frequency poses, preparing the input for the final video generation phase.

**ControlNet video diffusion.** In our video generator, we build our renderer on VideoCrafter2 [Chen et al. \(2023\)](#), a text-to-video generator with a UNet architecture. We enhance following the design of ControlNet [Zhang et al. \(2023\)](#) on additional human pose signals, ensuring that it adheres closely to the predicted human poses and dynamics throughout the video generation process. We incorporate zero convolution layers between ControlNet and original UNet, which are initialized with zero weights, minimizing its interference with the original U-Net’s operations at the beginning of training. We choose a UNet structured model as our base video generator instead of a diffusion transformer (DiT) due to the substantial computational demands of DiT. Diffusion transformer requires self-attention across all spatial and temporal locations, which makes it very costly. In our design, we offload the high-level motion understanding and planning a dedicated large language model (LLM) and a pose interpolator, reducing the workload for the video generator. This division of tasks enables the UNet model to focus exclusively on rendering details based on the human poses generated by the prior components, appropriate for generating realistic, temporally consistent human-centric video frames, making it a efficient and suitable choice.

## 3 Experiments

### 3.1 Dataset

For data preparation, we start data preparation with the Shutterstock (SSTK) [Shutterstock \(2024\)](#) video dataset, generating detailed text prompts with a video captioning model to provide scene descriptions. As our task focus on human-centric video generation, we filter the dataset according to text prompt and only keep human-centric video samples. We collected videos where the same person appears across all frames following [Polyak et al. \(2024\)](#). Then we apply OpenPose [Cao et al. \(2017\)](#) on each frame to extract precise

pose information. The extracted poses have 18 key points, disregarding all hand and feet points. We limit a maximum sequence length of 240 frames to capture extended scenes. During training, a subsequence of 200 frames is randomly sampled, with shorter sequences zero-padded to reach this length. From the 200 complete frames, we evenly extract 20 key frames. This thorough filtering and annotation process produces a high-quality dataset, well-suited for our model’s structured scene planning and pose conditioning tasks.

### 3.2 Training and Implementation Detail

Our training pipeline consists of four steps: (1) training a 6 depth residual vector-quantized VQ-VAE [Van Den Oord et al. \(2017\)](#) human poses, effectively discretizing the human poses, (2) training the large language model Director on these pose tokens, (3) training a linear diffusion model for smoothing and interpolation, and (4) training a ControlNet [Zhang et al. \(2023\)](#) for the video generator.

As shown in Fig. 8 left part, to convert continuous human poses detected by OpenPose into discrete tokens compatible with our large language model, we train a temporally-aware vector-quantized variational autoencoder (VQ-VAE [Van Den Oord et al. \(2017\)](#)). Our VQ-VAE is structured as a residual linear model with a depth of six, meaning it employs six codebooks, each with a size of 512. It takes in 20 pose key frames and convert them into 120 token ids. This configuration allows for high-fidelity discretization, capturing complex pose dynamics while efficiently encoding temporal information across frames. By transforming continuous pose vectors into discretized tokens, this VQVAE enables our large language model to effectively handle and predict pose sequences as structured data within a standard text generation pipeline.

Our large language model, fine-tuned from Llama 3 [Dubey et al. \(2024\)](#), is adapted to generate human poses tokens based on text prompts, with an extended vocabulary to accommodate newly added pose tokens. We show its structure in Fig. 8 middle part. The model is initialized from Llama 3 8B model checkpoint and all model parameters within its transformer stack are trained at this stage. The model is trained as a standard language generation model with PyTorch Lightning for 8 hours on 8 Nvidia A100 GPUs.

As shown in Fig. 8 right part, our linear interpolator integrates key frame poses and CLIP-encoded text prompts, producing smooth poses. The video ControlNet [Zhang et al. \(2023\)](#) is trained on dense pose frames where we randomly subsample 16 frames and its poses for training. We initialize it by duplicating the encoder structure and parameters of the UNet in VideoCrafter, as shown in Fig. 2. We keep the base UNet frozen during training and fine-tune only the pose ControlNet and zero conv layers. This approach allows our model to learn pose-specific control while maintaining the generative capabilities of the base video model.

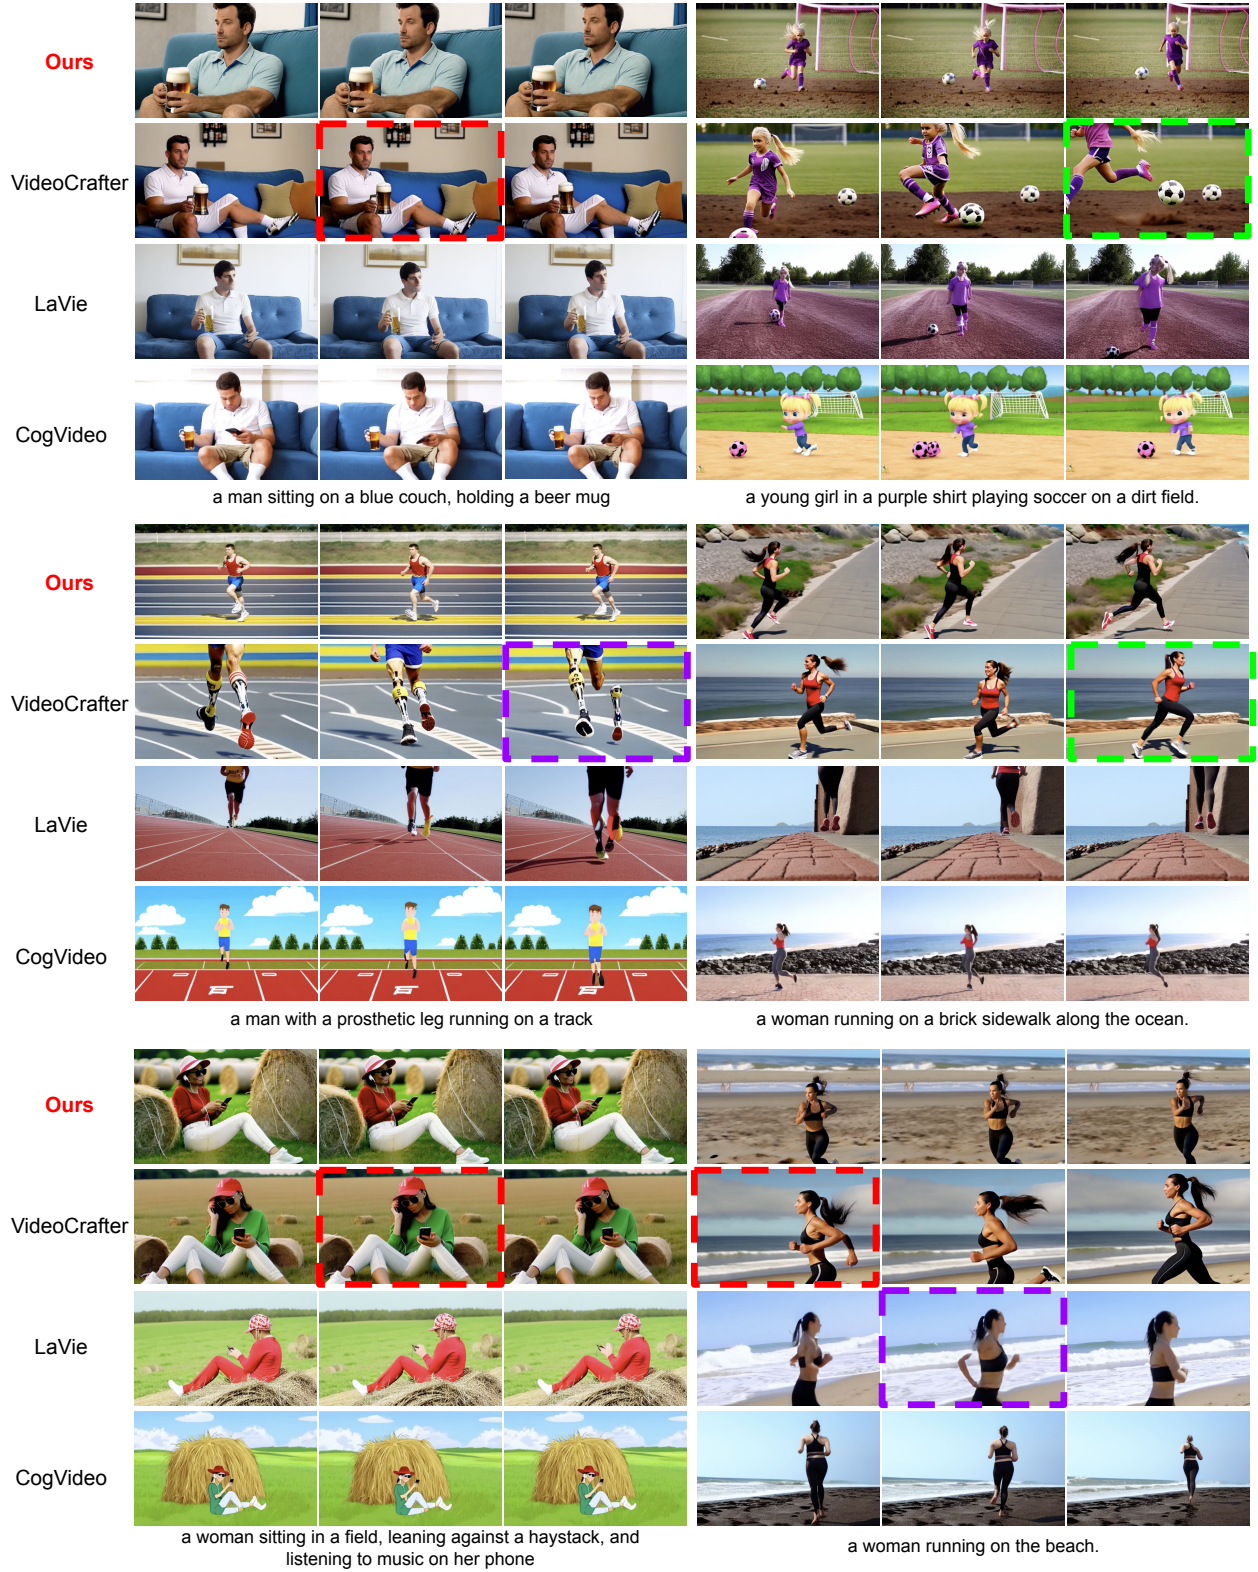
### 3.3 Evaluation

**Baselines and evaluation setup.** Since we are the first work to train a LLM as director for human-centric-video generation, we choose to compare with a few popular text to video generation models, including VideoCrafter [Chen et al. \(2023\)](#), LaVie-2 [Wang et al. \(2023b\)](#), and CogVideo [Hong et al. \(2022\)](#). We craft a test dataset with 150 human-centric text prompts, focusing on cases with dynamic scenes and human poses.


**Qualitative evaluation.** In our qualitative evaluation, we show outputs from each model for six video generation scenarios, as shown in Fig. 5. The baselines sometimes struggle with anatomical realism and temporal consistency, as indicated by several noticeable errors (highlighted in dashed colored boxes), including extra or broken limbs, and abrupt directional changes during movement. These issues reflect the baselines’ limitations in handling complex human dynamics, leading to unrealistic and visually jarring results. In contrast, our model avoids these visual artifacts, showing its improved performance in generating high-quality, human-centric videos. Due to space limits, we show more visual examples of generated pose and video frames from our model in Fig. 4. Further more, our DirectorLLM generates 20 key pose frames which are then interpolated to 200 frames by pose interpolation model, sufficient for a video for as long as 7 seconds. We show examples of generate pose sequences and rendered RGB frames in the supplementary, demonstrating its capability of understand human dynamics and generating poses with high temporal consistency, dynamics and realism over extended durations.

**Quantitative evaluation.**





**Figure 5** Qualitative comparison. Comparison of our model with baseline models (VideoCrafter, LaVie, and CogVideo) across six different video generation scenarios, each described by a text prompt. Our model consistently produces realistic and coherent human poses without visual artifacts, while the baselines frequently exhibit errors highlighted in dashed red boxes. These baseline failures include anatomical inconsistencies like extra limbs, detached body part, and abrupt direction changes.



**Figure 6** Human evaluations show user preference towards our method.

We run a quantitative comparison on the entire test dataset and show the results as follows. In our experiments, we utilize VideoScore [He et al. \(2024b\)](#) to evaluate our text-to-video model. VideoScore is an automatic video quality assessment metric trained on VideoFeedback, a large-scale dataset containing human-provided scores for 37.6K synthesized videos from various generative models. Given its high correlation with human judgments, VideoScore serves as a reliable proxy for human evaluation, allowing us to comprehensively assess our model’s performance across these crucial aspects. We show the results in Fig. 1 for key evaluation dimensions: Visual Quality (**VC**), Temporal Consistency (**TC**), Dynamic Degree (**DD**), Text-to-Video Alignment (**CLIPvid**), and Factual Consistency (**FC**). Through quantitative evaluations, we show that our proposed method outputs realistic motion more diverse than baselines, especially in Text-to-Video Alignment (CLIPvid) score.

Metric	Ours	VideoCrafter2	LaVie	CogVideo
VC ↑	<b>2.5205</b>	2.5171	2.2748	2.2097
TC ↑	<b>2.4578</b>	2.4452	2.1527	2.0874
DD ↑	<b>3.2663</b>	3.2294	3.0343	3.0551
CLIPvid ↑	<b>3.2471</b>	3.1998	2.8947	2.9583
FC ↑	<b>2.3598</b>	2.3238	2.1117	2.1591
average	<b>2.7703</b>	2.7430	2.4936	2.4939

**Table 1** Comparison of different models across various metrics. Evaluated on human-centric dynamic prompts.

**Human evaluation.** Since there is no automatic video metric on the quality of human pose dynamic, we apply human evaluation specifically for this purpose. We compare the best results from each method and make a questionnaire for human evaluators. Specifically, we run blind A/B comparisons between videos generated from each model and ask questions on: (1) overall quality (**Overall**), (2) prompt alignment for the human pose (**Pose Alignment**), (3) prompt alignment for the video (**Text Alignment**), (4) generation quality for the dynamic human motion (**Human Motion**), and (5) generation quality for static human balance (**Human Balance**). For Text Alignment, annotators are instructed to consider all prompt specifications and the better one would be one that achieves all specifications in the prompt. Human Motion evaluates the realism of dynamic human movement within the generated video, focusing on natural and coherent motion patterns. Human Balance assesses the static realism of human figures, ensuring accurate proportions and anatomical balance, such as correct head-to-body ratios and plausible body alignment. For each pair of comparisons, we would randomize A/B and use 3 raters per comparison, and use the majority vote as the final decision. The human evaluation setup was inspired by the video generation model [Polyak et al. \(2024\)](#).

Fig. 6 shows the results of human evaluation comparing our model to baselines (LaVie [Wang et al. \(2023b\)](#), CogVideo [Hong et al. \(2022\)](#), and VideoCrafter [Chen et al. \(2023\)](#)) across five categories. Our model consistently outperformed baselines in key metrics, particularly in Pose Alignment and Human Motion, demonstrating its superior capability in generating human-centric and text-aligned video content. This aligns with our core motivation to improve the realism and accuracy of human representations in generated videos by offloading high-level scene understanding and pose prediction to a dedicated language model. This enhances both the static and dynamic realism of human figure, fulfilling our objective to produce better human-centric videos.

### 3.4 Ablation and Analysis

To better understand the contributions of each component in our model, we conducted ablation studies on three key aspects: vector quantization (VQ-VAE), the use of basic interpolation versus our linear diffusion model, and the impact of text conditioning in the linear diffusion model.

**Residual vector quantization.** In this experiment, we ablate our design choice of residual pose VQ-VAE. The goal is to evaluate whether our 6 depth residual autoencoding benefits LLM’s predictive performance and video quality. As shown in Fig. 2, without residual autoencoding, the VQ-VAE couldn’t accurately reconstruct human pose, significantly degrading pose prediction accuracy and introducing inconsistencies in generated videos. Residual structure in our VQ-VAE, with a depth of 6, plays a crucial role in preserving pose fidelity, directly enhancing both LLM prediction accuracy and overall video quality.

	VC ↑	TC ↑	DD ↑	CLIPvid ↑	FC ↑
w/o Residual VQ	2.380	1.815	3.007	2.773	1.874
w/o Pose Diffusion	2.393	2.197	3.268	2.861	2.025
w/o Text Condition	2.445	2.382	3.255	3.203	2.331
Ours	<b>2.520</b>	<b>2.457</b>	<b>3.266</b>	<b>3.247</b>	<b>2.359</b>

**Table 2** Comparison of different models across various metrics. Our ablation demonstrates the importance of each component. Residual vector quantization improves pose reconstruction and LLM predictions, the linear diffusion model ensures temporal consistency and smooth motion, and text conditioning enhances alignment with prompts.

**Linear diffusion model.** To assess the impact of our diffusion-based pose interpolation model, we replaced it with direct pose interpolation by linearly interpolating the 20 key frames to generate all 200 frames. These frames are then rendered to videos with our video generator. As shown in Table. 2, the direct interpolation lacks authentic human dynamics and affects scores such as dynamic degree. Linear diffusion model outperformed basic interpolation in maintaining temporal consistency and generating smooth motion transitions, particularly in scenes with dynamic human movements.

**Text condition in linear diffusion model.** We also tested the effect of adding text-conditioned CLIP embeddings in the linear diffusion model by training it with and without the text conditions. As shown in Table. 2, without text conditioning, the generated videos exhibited reduced alignment with the original text prompts, leading to inconsistencies in scene layout and subject interaction.

## 4 Discussions

As far as our best knowledge goes, we are the first work that have showcased the effectiveness of offloading human motions simulation to an LLM for human centric T2V generation. Our approach addresses the broader challenge of human video generation with dedicated motion reasoning. We believe this opens a way to motion-simulated video generators, where the motion and interaction of the main characters are driven by dedicated AIs.

## 5 Conclusions

We introduce DirectorLLM, a text-conditioned video generation model that excels in creating temporally consistent, realistic videos, especially in complex scenes involving human interactions and detailed poses. By delegating scene understanding tasks to a specialized large language model that predicts bounding boxes and poses, our approach allows the video generator to focus solely on rendering. Our three-part design—consisting of the LLM, a linear diffusion model for smooth interpolation, and an enhanced VideoCrafter generator—enables precise, structured video synthesis that outperforms existing methods, particularly in scenarios requiring complex subject dynamics. This model highlights the benefits of separating scene comprehension from video generation, advancing the potential for controlled, high-quality human-centric video generation.



## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Stability AI. Unclip model documentation, 2024. <https://github.com/Stability-AI/stablediffusion/blob/main/doc/UNCLIP.MD>. Accessed: 2024-10-31.
- Simon Alexanderson, Rajmund Nagy, Jonas Beskow, and Gustav Eje Henter. Listen, denoise, action! audio-driven motion synthesis with diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4):1–20, 2023.
- Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*, 2023.
- Tenglong Ao, Zeyi Zhang, and Libin Liu. Gesturediffuclip: Gesture diffusion model with clip latents. *ACM Transactions on Graphics (TOG)*, 42(4):1–18, 2023.
- Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Qinsheng Zhang, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, et al. ediff-i: Text-to-image diffusion models with an ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*, 2022.
- 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.
- Haoxin Chen, Menghan Xia, Yingqing He, Yong Zhang, Xiaodong Cun, Shaoshu Yang, Jinbo Xing, Yaofang Liu, Qifeng Chen, Xintao Wang, et al. Videocrafter1: Open diffusion models for high-quality video generation. *arXiv preprint arXiv:2310.19512*, 2023.
- Xiaoliang Dai, Ji Hou, Chih-Yao Ma, Sam Tsai, Jialiang Wang, Rui Wang, Peizhao Zhang, Simon Vandenhende, Xiaofang Wang, Abhimanyu Dubey, et al. Emu: Enhancing image generation models using photogenic needles in a haystack. *arXiv preprint arXiv:2309.15807*, 2023.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Weixi Feng, Wanrong Zhu, Tsu-jui Fu, Varun Jampani, Arjun Akula, Xuehai He, Sugato Basu, Xin Eric Wang, and William Yang Wang. Layoutgpt: Compositional visual planning and generation with large language models. *Advances in Neural Information Processing Systems*, 36, 2024.
- Jia Guo and Jiankang Deng. insightface, 2024. [https://github.com/deepinsight/insightface/tree/master/examples/person\\_detection](https://github.com/deepinsight/insightface/tree/master/examples/person_detection). Accessed: 2024-10-31.
- 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.
- Hao He, Yinghao Xu, Yuwei Guo, Gordon Wetzstein, Bo Dai, Hongsheng Li, and Ceyuan Yang. Cameractrl: Enabling camera control for text-to-video generation. *arXiv preprint arXiv:2404.02101*, 2024a.
- Xuan He, Dongfu Jiang, Ge Zhang, Max Ku, Achint Soni, Sherman Siu, Haonan Chen, Abhranil Chandra, Ziyang Jiang, Aaran Arulraj, et al. Videoscore: Building automatic metrics to simulate fine-grained human feedback for video generation. *arXiv preprint arXiv:2406.15252*, 2024b.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *arXiv preprint arXiv:2205.15868*, 2022.
- Hanzhuo Huang, Yufan Feng, Cheng Shi, Lan Xu, Jingyi Yu, and Sibe Yang. Free-bloom: Zero-shot text-to-video generator with llm director and ldm animator. *Advances in Neural Information Processing Systems*, 36, 2024.
- Biao Jiang, Xin Chen, Wen Liu, Jingyi Yu, Gang Yu, and Tao Chen. Motiongpt: Human motion as a foreign language. *Advances in Neural Information Processing Systems*, 36:20067–20079, 2023.
- Jihwan Kim, Junoh Kang, Jinyoung Choi, and Bohyung Han. Fifo-diffusion: Generating infinite videos from text without training. *arXiv preprint arXiv:2405.11473*, 2024.



- Long Lian, Boyi Li, Adam Yala, and Trevor Darrell. Llm-grounded diffusion: Enhancing prompt understanding of text-to-image diffusion models with large language models. *arXiv preprint arXiv:2305.13655*, 2023.
- Han Lin, Abhay Zala, Jaemin Cho, and Mohit Bansal. Videodirectorgpt: Consistent multi-scene video generation via llm-guided planning. *arXiv preprint arXiv:2309.15091*, 2023.
- Jiaxi Lv, Yi Huang, Mingfu Yan, Jiancheng Huang, Jianzhuang Liu, Yifan Liu, Yafei Wen, Xiaoxin Chen, and Shifeng Chen. Gpt4motion: Scripting physical motions in text-to-video generation via blender-oriented gpt planning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1430–1440, 2024.
- Wan-Duo Kurt Ma, John P Lewis, and W Bastiaan Kleijn. Trailblazer: Trajectory control for diffusion-based video generation. *arXiv preprint arXiv:2401.00896*, 2023.
- Evonne Ng, Javier Romero, Timur Bagautdinov, Shaojie Bai, Trevor Darrell, Angjoo Kanazawa, and Alexander Richard. From audio to photoreal embodiment: Synthesizing humans in conversations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1001–1010, 2024.
- OpenAI. Sora, 2024. <https://openai.com/index/sora/>. Accessed: 2024-10-31.
- Quynh Phung, Songwei Ge, and Jia-Bin Huang. Grounded text-to-image synthesis with attention refocusing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7932–7942, 2024.
- Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, David Yan, Dhruv Choudhary, DingKang Wang, Geet Sethi, Guan Pang, Haoyu Ma, Ishan Misra, Ji Hou, Jialiang Wang, Kiran Jagadeesh, Kunpeng Li, Luxin Zhang, Mannat Singh, Mary Williamson, Matt Le, Matthew Yu, Mitesh Kumar Singh, Peizhao Zhang, Peter Vajda, Quentin Duval, Rohit Girdhar, Roshan Sumbaly, Sai Saketh Rambhatla, Sam Tsai, Samaneh Azadi, Samyak Datta, Sanyuan Chen, Sean Bell, Sharadh Ramaswamy, Shelly Sheynin, Siddharth Bhattacharya, Simran Motwani, Tao Xu, Tianhe Li, Tingbo Hou, Wei-Ning Hsu, Xi Yin, Xiaoliang Dai, Yaniv Taigman, Yaqiao Luo, Yen-Cheng Liu, Yi-Chiao Wu, Yue Zhao, Yuval Kirstain, Zecheng He, Zijian He, Albert Pumarola, Ali Thabet, Artsiom Sanakoyeu, Arun Mallya, Baishan Guo, Boris Araya, Breena Kerr, Carleigh Wood, Ce Liu, Cen Peng, Dmitry Vengertsev, Edgar Schonfeld, Elliot Blanchard, Felix Juefei-Xu, Fraylie Nord, Jeff Liang, John Hoffman, Jonas Kohler, Kaolin Fire, Karthik Sivakumar, Lawrence Chen, Licheng Yu, Luya Gao, Markos Georgopoulos, Rashel Moritz, Sara K. Sampson, Shikai Li, Simone Parmeggiani, Steve Fine, Tara Fowler, Vladan Petrovic, and Yuming Du. Movie gen: A cast of media foundation models. *arXiv preprint arXiv:2410.13720*, 2024.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- 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.
- Anton Razhigayev, Arseniy Shakhmatov, Anastasia Maltseva, Vladimir Arkhipkin, Igor Pavlov, Ilya Ryabov, Angelina Kuts, Alexander Panchenko, Andrey Kuznetsov, and Denis Dimitrov. Kandinsky: an improved text-to-image synthesis with image prior and latent diffusion. *arXiv preprint arXiv:2310.03502*, 2023.
- 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.
- 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.
- Xiaoyu Shi, Zhaoyang Huang, Fu-Yun Wang, Weikang Bian, Dasong Li, Yi Zhang, Manyuan Zhang, Ka Chun Cheung, Simon See, Hongwei Qin, et al. Motion-i2v: Consistent and controllable image-to-video generation with explicit motion modeling. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024.
- Shutterstock. Shutterstock, 2024. <https://www.shutterstock.com/>. Accessed: 2024-10-31.
- Jonathan Tseng, Rodrigo Castellon, and Karen Liu. Edge: Editable dance generation from music. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 448–458, 2023.
- Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.

- Jiawei Wang, Yuchen Zhang, Jiaxin Zou, Yan Zeng, Guoqiang Wei, Liping Yuan, and Hang Li. Boximator: Generating rich and controllable motions for video synthesis. *arXiv preprint arXiv:2402.01566*, 2024a.
- Jiuniu Wang, Hangjie Yuan, Dayou Chen, Yingya Zhang, Xiang Wang, and Shiwei Zhang. Modelscope text-to-video technical report. *arXiv preprint arXiv:2308.06571*, 2023a.
- Xiang Wang, Hangjie Yuan, Shiwei Zhang, Dayou Chen, Jiuniu Wang, Yingya Zhang, Yujun Shen, Deli Zhao, and Jingren Zhou. Videocomposer: Compositional video synthesis with motion controllability. *Advances in Neural Information Processing Systems*, 36, 2024b.
- Yaohui Wang, Xinyuan Chen, Xin Ma, Shangchen Zhou, Ziqi Huang, Yi Wang, Ceyuan Yang, Yinan He, Jiashuo Yu, Peiqing Yang, et al. Lavie: High-quality video generation with cascaded latent diffusion models. *arXiv preprint arXiv:2309.15103*, 2023b.
- Zhouxia Wang, Ziyang Yuan, Xintao Wang, Yaowei Li, Tianshui Chen, Menghan Xia, Ping Luo, and Ying Shan. Motionctrl: A unified and flexible motion controller for video generation. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024c.
- Yujie Wei, Shiwei Zhang, Zhiwu Qing, Hangjie Yuan, Zhiheng Liu, Yu Liu, Yingya Zhang, Jingren Zhou, and Hongming Shan. Dreamvideo: Composing your dream videos with customized subject and motion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6537–6549, 2024.
- Weijia Wu, Zhuang Li, Yuchao Gu, Rui Zhao, Yefei He, David Junhao Zhang, Mike Zheng Shou, Yan Li, Tingting Gao, and Di Zhang. Draganything: Motion control for anything using entity representation. In *European Conference on Computer Vision*, pages 331–348. Springer, 2025.
- Shiyuan Yang, Liang Hou, Haibin Huang, Chongyang Ma, Pengfei Wan, Di Zhang, Xiaodong Chen, and Jing Liao. Direct-a-video: Customized video generation with user-directed camera movement and object motion. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–12, 2024.
- Shengming Yin, Chenfei Wu, Jian Liang, Jie Shi, Houqiang Li, Gong Ming, and Nan Duan. Dragnuwa: Fine-grained control in video generation by integrating text, image, and trajectory. *arXiv preprint arXiv:2308.08089*, 2023.
- Zhentao Yu, Zixin Yin, Deyu Zhou, Duomin Wang, Finn Wong, and Baoyuan Wang. Talking head generation with probabilistic audio-to-visual diffusion priors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7645–7655, 2023.
- 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.
- Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*, 2022.
- Zhenghao Zhang, Junchao Liao, Menghao Li, Long Qin, and Weizhi Wang. Tora: Trajectory-oriented diffusion transformer for video generation. *arXiv preprint arXiv:2407.21705*, 2024.
- Rui Zhao, Yuchao Gu, Jay Zhangjie Wu, David Junhao Zhang, Jiawei Liu, Weijia Wu, Jussi Keppo, and Mike Zheng Shou. Motiondirector: Motion customization of text-to-video diffusion models. *arXiv preprint arXiv:2310.08465*, 2023.

# Appendix

## A. Implementation Details

### A1. Pre-Processing Details

Our video caption model generates highly structured prompts. Below is an example of full text prompt generated by video captioning model:

*The video shows a woman running through a grassy field. The woman has fair skin and long brown hair pulled back into a ponytail. She is wearing a yellow tank top and black shorts. She is running towards the right side of the frame. The background is a bright sky and a grassy hill. The woman is running through a grassy field. She is wearing a watch on her left wrist. The camera is static.*

The **first sentence**, known as the **"Subject-Level Prompt"**, provides an overall video description with key information on human motions and poses. We use it as input to DirectorLLM transformer and Pose Linear Interpolator. We use full prompt as input to ControlNet Diffusion generator as it offers detailed descriptions of appearance and scene elements. We use OpenPose to extract per-frame human-poses for 25K videos. For videos with more than one human pose detected, we run a light weight human detection model InsightFace [Guo and Deng \(2024\)](#) to keep the one with the largest area and ignore all others.

### A2. Model and Training Details

**Pose representation.** Keypoints for face, hand and feet are ignored, resulting in a 2D representation with 18 keypoints per human pose. The coordinates of keypoints are flattened to 36 dimensional vectors, on which our Pose VQ-VAE and interpolation model work. Missing or un-detected keypoints are replaced with 0 values and ignored in loss function. We only keep 200 frames for each video, with shorter ones padded by zero.

**Residual VQ-VAE.** Our residual VQ-VAE is designed to capture fine-grained details by employing a cascade of 6 codebooks of 512 tokens, progressively refining approximations for each pose. Both the encoder and decoder consist of a series of 1D convolutions with a kernel size of 2, providing a total receptive field of 8. Each pose is represented by a sequence of 6 VQ tokens, enabling detailed and structured encoding. For each video, we evenly sample 20 key frames from a total of 200 frames of human poses. These key frames can be converted to  $6 \times 20 = 120$  tokens by VQ-VAE encoder. The model is compact, requiring about 2 hours of training on a single A100 GPU.

**DirectorLLM.** We leverage the advanced pre-existing knowledge from pretrained large language model to aid in model convergence. Specifically, we initialize our DirectorLLM with Llama 3 -8B checkpoint [Dubey et al. \(2024\)](#) from HuggingFace. We extend its tokenizer and embeddings to an additional  $6 \times 512$  tokens, each of which representing a pose token from VQ-VAE codebooks. We attach  $6 \times 20 = 120$  pose tokens from VQ-VAE to the "subject-level prompt", forming a sequence as long as 300 tokens. Training code is built with Pytorch-Lightning and the model is trained as a standard language generation model for 8 hours on 8 Nvidia A100 GPUs.

**Pose Interpolation.** To densify human motion from 1 fps to 30 fps, we construct an conditional diffusion model. It is built on top of the Audio2PhotoReal [Ng et al. \(2024\)](#). We follow the standard DDPM definition of diffusion, where we add noise to pose vectors during forward pass and train the network to remove it during as the denoising pass. It takes in noise and predict 200 pose vectors, each for one frame, conditioned on 20 key frame poses. To incorporate the text prompt information, we use CLIP text encoder and apply a cross attention layer. Timestep information is incorporated with a feature-wise linear modulation (FiLM) layer, similar to Audio2PhotoReal [Ng et al. \(2024\)](#). All parameters are optimized. Training is finished in 8 hours on one single A100 GPU.

**Text-to-video ControlNet.** We build our video renderer model on top of VideoCrafter2 [Chen et al. \(2023\)](#). A pose ControlNet [Zhang et al. \(2023\)](#) module is initialized by copying its UNet encoder structure and parameters. We convert poses for 16 frames from 32 dimensional vector to image, and feed it to ControlNet. Connection layers are initialized with zero convolution for better converging. We lock the parameters of base

UNet and optimize only our pose ControlNet. We apply large gradient accumulation steps (200) and clip grad norm for smoother training. Training is finished with AdamW optimizer and a learning rate of  $1^{-5}$  on 32 A100 GPUs in 48 hours.

### A3. Evaluation

**Evaluation prompts.** To ensure a meaningful comparison with baseline models, we evaluate our approach using 150 human-centric text prompts randomly sampled from the test dataset. These prompts are carefully filtered to focus on scenarios involving complex scenes and dynamic human poses, selecting only those containing **keywords** indicative of large motions, including "walking", "hiking", "running", "yoga", "jogging", "workout", "dance", "jump", and "cycling". This filtering process is essential as it emphasizes the model's ability to handle challenging cases where realistic motion dynamics and pose accuracy are critical. We notice minimal performance gain with stationary prompts. By targeting prompts with significant motion, we highlight the strengths of our model in generating human-centric videos that require high fidelity in pose transitions and temporal consistency, areas where many baseline methods often struggle. This approach ensures that the evaluation reflects the practical utility and robustness of our model in human-centric, motion-intensive scenarios.

**Generation pipeline.** Generation process begins with a text prompt, which is passed to DirectorLLM, to generate  $6 \times 20 = 120$  human pose tokens, representing key frame poses at 1 frame per second (fps). Pose tokens are then decoded to pose vectors with the decoder of VQ-VAE. The sparse pose vectors are then processed by a linear diffusion model. This model interpolates the 20 key frame sparse poses into 200 frames dense sequences, ensuring temporal consistency and smooth motion transitions. The linear diffusion model is also conditioned on text embeddings from CLIP to align motion dynamics with the prompt. The video generator, based on VideoCrafter and ControlNet, takes dense pose along with the full text prompt as input, producing high-quality and anatomically accurate video frames. However, due to limitations, only 16 frames are generated per inference time.

We use FIFO [Kim et al. \(2024\)](#) to bridge the gap between short-frame limitations and long-sequence requirements, which will be covered in section C.

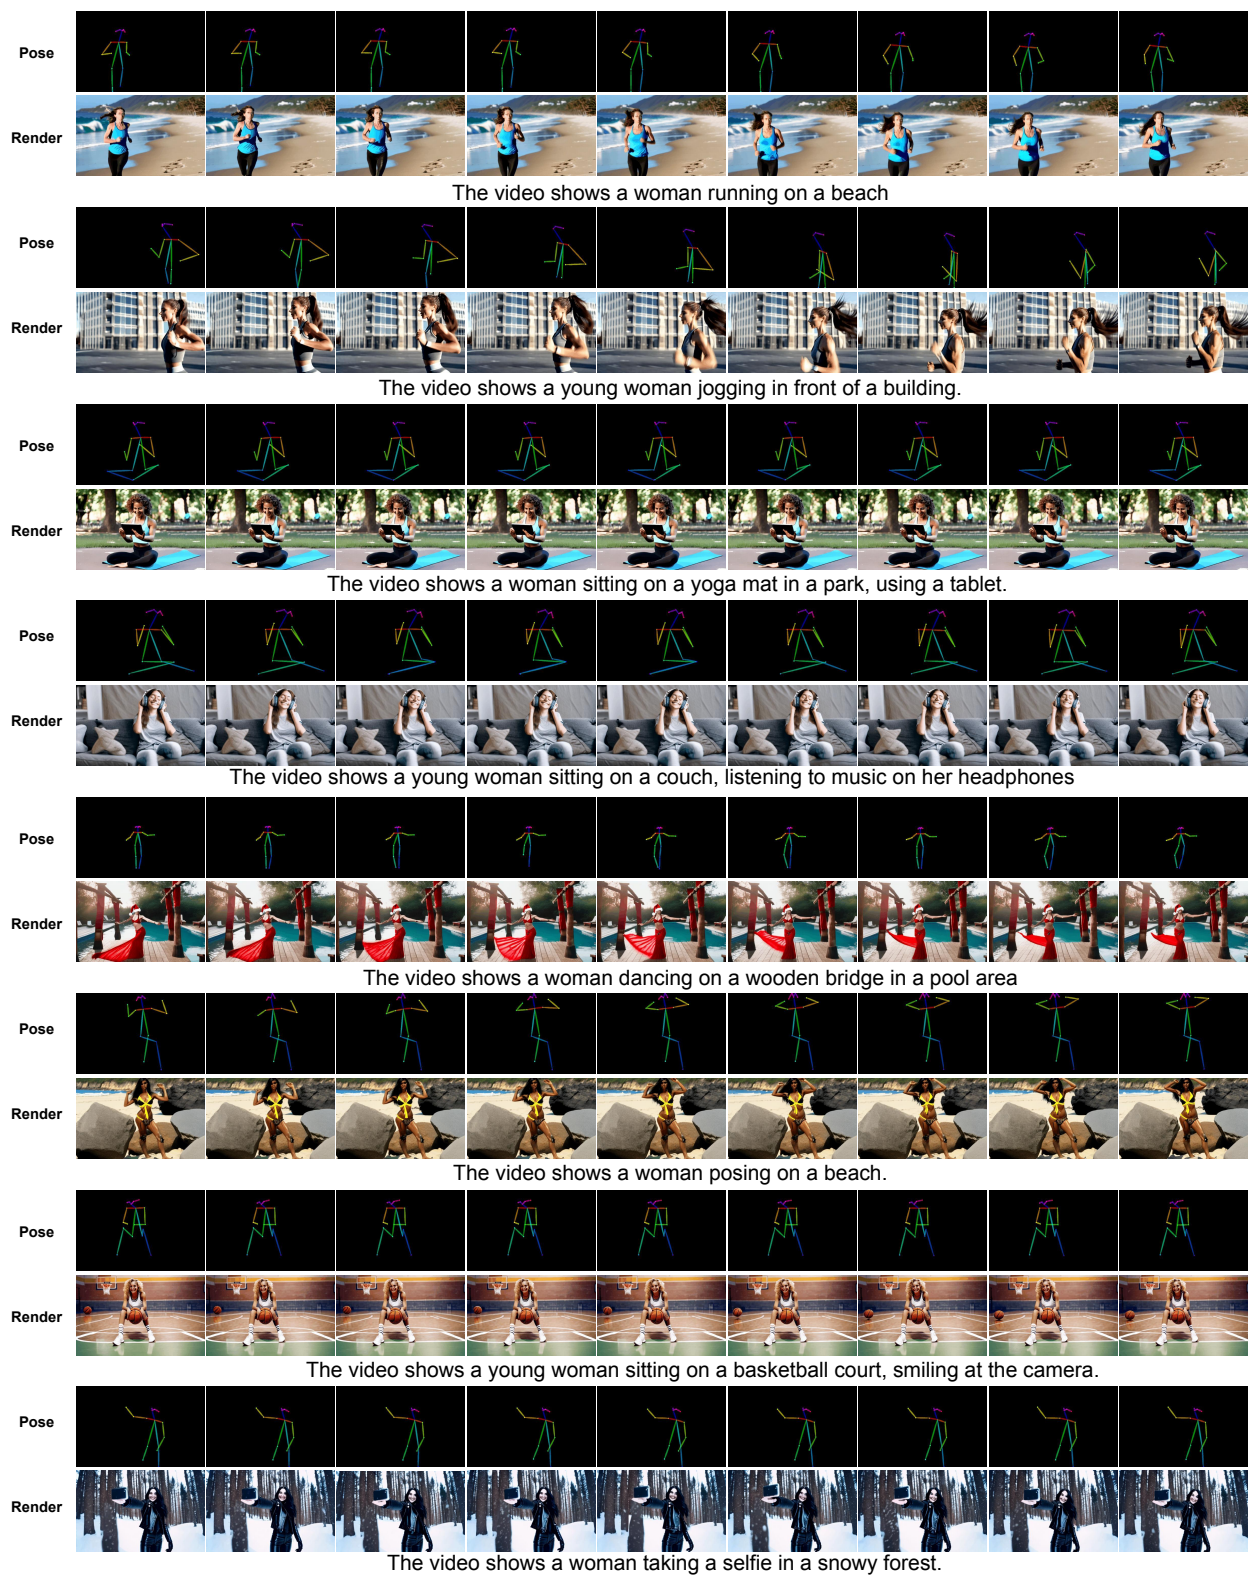
## B. More Results from Our Model

We present additional visual results to demonstrate the effectiveness of our model in this section. [figure 8](#) upper panel showcases the sequence of human poses predicted by the LLM Director, highlighting its ability to generate accurate and temporally coherent pose layouts based on text prompts. [figure 7](#) pairs these pose layouts with their corresponding video renderings, illustrating how the poses guide the video generator to produce realistic, high-quality human motions and interactions. The alignment between the predicted poses and the rendered visuals demonstrates the strength of our multi-stage pipeline


## C. Extending Video Length

Our pose interpolation model generates 200 pose frames, while VideoCrafter with ControlNet is limited to generating 16 frames at a time. To extend the length of the generated video and render all 200 frames, we apply the FIFO-Diffusion [Kim et al. \(2024\)](#) approach. This method enables infinite-length video generation by employing diagonal denoising, which processes consecutive frames in a queue. During this process, fully denoised frames at the head of the queue are dequeued, while new noise frames are enqueued at the tail for processing. To mitigate the training-inference gap caused by this iterative strategy, FIFO incorporates latent partitioning and lookahead denoising techniques [Kim et al. \(2024\)](#), ensuring consistency across extended sequences. FIFO uses the exact same model checkpoint as our base model. Simply adding ControlNet modules to the FIFO code allows us to generate long videos with pose condition awareness, leveraging VideoCrafter's capabilities to produce human-centric videos for as long as 200 frames. A sample is attached in [figure 8](#) lower panel.







**Figure 7** Generated poses and video frames.




The video shows a young woman sitting on a couch



The video shows a young boy running through a sunflower field



The video shows a young woman enacting in a voga routine



The video shows a woman posing on a beach.

**Figure 8** Generated samples from DirectorLLM after pose interpolation. Each prompt showcases four representative samples to highlight the model can generate various poses given the same prompt.



## D. About User Study

We design a complex and complete evaluation pipeline with 5 questions. Specifically, we run blind A/B comparisons between videos generated from each model and ask questions on: (1) overall quality (**Overall**), (2) prompt alignment for the human pose (**Pose Alignment**), (3) prompt alignment for the video (**Text Alignment**), (4) generation quality for the dynamic human motion (**Human Motion**), and (5) generation quality for static human balance (**Human Balance**). To reduce potential biases, we recruit annotators from various countries with diverse backgrounds. Each annotator is presented with randomly paired video samples without any knowledge of the generating model to maintain impartiality. In total, we collect over 5K ratings, providing a robust and comprehensive assessment of the comparative performance of the models.