

CustomX: Unified Character, Action, and Scene Customization in Video World Models

Yitong Wang^{1,2*}, Fangyun Wei^{2*}, Hongyang Zhang³,
Bo Dai^{4†}, and Yan Lu²

¹ Fudan University, Shanghai, China

² Microsoft Research, Beijing, China

³ University of Waterloo, Waterloo, ON, Canada

⁴ The University of Hong Kong, Hong Kong SAR, China

wangyitong23@m.fudan.edu.cn {fawe, yanlu}@microsoft.com

hongyang.zhang@uwaterloo.ca bdai@hku.hk

https://snowflakewang.github.io/CustomX_Page/



Fig. 1: CustomX enables users to provide a 3D or multi-view character together with a 3DGS scene, supporting iterative, multi-condition control of character behaviors and active environment exploration through natural language-guided video generation. The system features: (1) *Consistent Environment and Character Fidelity*, ensuring visual and spatial coherence with the user-provided scene and character; (2) a *Rich Action Repertoire* covering a wide range of behaviors, including locomotion, gestures, and object-centric interactions; (3) *Long-Horizon, Temporally Coherent Interaction*, enabling iterative user interaction while maintaining continuity across generated clips; and (4) *Controllable Camera Behavior*, which explicitly incorporates camera control, analogous to navigating 3DGS views, to produce accurate, user-specified viewpoints.

Abstract. Recent advances in world models have greatly enhanced interactive environment simulation. Existing methods mainly fall into two

★ Equal contribution. † Corresponding author.

categories: (1) static world generation models, which construct 3D environments without active agents, and (2) controllable-entity models, which allow a single entity to perform limited actions in an otherwise uncontrollable environment. In this work, we introduce CustomX, leveraging the realism and structural grounding of static world generation while extending controllable-entity models to support user-specified characters capable of performing open-ended actions. Users can provide a 3DGS scene and a character, then use natural language to direct the character to perform diverse behaviors, ranging from basic locomotion to object-centric interactions, while freely exploring the environment. CustomX synthesizes temporally coherent video clips that preserve visual fidelity with the provided scene and character, formulated as a conditional autoregressive video generation problem. Built upon a pre-trained video generator, our training strategy significantly enhances motion dynamics while maintaining generalization across actions and characters. Our evaluation covers a broad range of aspects, including visual quality, character consistency, action controllability, and long-horizon coherence.

Keywords: World Models · Controllable Video Generation · Video Generation Post-Training

1 Introduction

Recent advances in world models have led to substantial progress in simulating dynamic and interactive environments. Existing methods generally fall into two categories: (1) static world generation approaches [10, 56, 85], which construct explorable 3D environments but lack active agents; and (2) controllable-entity approaches [46, 87], which allow a single agent to execute only a limited set of actions, such as steering a vehicle along a predefined path [23, 33, 69], while leaving the environment itself uncontrollable. In this work, we propose an alternative framework that combines the strengths of both paradigms: leveraging the realism and structural grounding of static world generation while extending controllable-entity models to support user-specified characters capable of performing open-ended actions.

Specifically, users can provide a 3D Gaussian Splatting (3DGS) scene [43], which represents either a synthetic environment or a real-world reconstruction, along with a 3D or multi-view character. Through natural language instructions, users can control the character’s behavior and enable active exploration within the environment. At each iteration, the model generates a video clip that captures the evolving states of both the character and the environment, resulting in coherent and temporally consistent generation. We name our method *CustomX*. As illustrated in Fig. 1, CustomX exhibits several key capabilities:

1. **Consistent Environment and Character Fidelity.** The visual contents appearing in the generated video clips exhibit strong consistency in visual identity and spatial layout with the user-provided scene and character.

2. **Rich Action Repertoire.** Unlike previous works [16, 33, 46] that limit the controllable entity to basic locomotion, our model enables the character to perform up to hundreds of distinct actions, encompassing not only navigation behaviors (e.g., “moving forward”, “turning left”) but also body-language gestures (e.g., “waving hands”, “saluting”) and object-centric interactions (e.g., “making a phone call”, “playing a guitar”).
3. **Text Instruction as the Interface.** Users can directly guide the character through natural language commands.
4. **Long-Horizon, Temporally Coherent Interaction.** Users can interact with the model iteratively, generating new video clips that remain temporally consistent with previously produced sequences.
5. **Controllable Camera Behavior.** Our model supports flexible and intuitive camera control, allowing behaviors such as following a character’s trajectory or orbiting around it to achieve user-specified viewpoints. Unlike previous methods [31, 32] that encode camera trajectories into Plücker embeddings [67] and inject them as conditioning signals into the generation network, our approach achieves camera control in a more geometrically grounded manner. Specifically, given a user-provided 3DGS scene and a defined camera path, we directly render a projection scene video along the specified trajectory. This rendered video serves as an explicit conditioning input, enabling the generation model to produce videos that accurately follow the desired camera motion.

We formulate the entire process as a conditional autoregressive video generation problem. Concretely, the objective is to synthesize a video clip at each iteration, conditioned on a set of multi-modal inputs including: (a) the user-provided scene and character, which establish the spatial and visual grounding; (b) a text instruction, which specifies the intended behavior of the character; and (c) the previously generated video clips, which serve as temporal references to ensure consistency across iterations.

Additionally, we adopt a pre-trained video generator [35] as the foundation of our framework. We find that fine-tuning it on a small dataset containing basic locomotion actions across diverse characters not only preserves the generalization ability of the pre-trained model but also enhances overall motion quality compared to the original generator. This phenomenon is analyzed in Sec. 4.2.

In our experiments, we comprehensively evaluate the proposed model from multiple perspectives, including: (a) visual quality, assessed using the WorldScore benchmark [20]; (b) character consistency, which measures the alignment between the character appearing in the generated video and the user-provided reference; (c) action control success rate, which quantifies how accurately the character’s behavior follows the input text instruction across a diverse set of up to around 150 actions; and (d) long-horizon generation quality, which evaluates the model’s ability to maintain temporal coherence and visual fidelity over extended interaction sequences. We compare CustomX with both video generation foundation models [44, 73, 86] and dedicated world models [16, 33, 46]. Experi-

mental results show that our method consistently outperforms both categories across nearly all evaluation metrics.

2 Related Works

Controllable Video Generation. Recent foundation models for video generation [1, 5, 6, 13, 29, 44, 73, 86] have greatly improved modality alignment across text, image, and video, enabling large-scale pre-training for both text-to-video and image-to-video synthesis. Building on these advances, subsequent research has pursued finer-grained controllability by introducing mechanisms such as explicit subject control [15, 18, 22, 39–41, 52, 55] and camera control [12, 24, 26, 28, 57, 65, 76, 82, 83, 92, 93, 98]. For subject control, a typical approach [35, 50] is to extract visual embeddings from reference images and use them within a Multimodal Diffusion Transformer [21, 61] to guide the generated subjects to remain consistent with the reference appearance. For camera control, one common practice [31, 32] is to convert the camera path into Plücker embeddings [67] and inject them into the main network, guiding the synthesized video to follow the specified trajectory. In contrast, our model controls the camera by navigating through 3DGS views: given a 3DGS scene and a camera path, we render a projection video along the path, which conditions the generator to follow the desired motion faithfully.

Memory Mechanism in Video Generation. Recent works incorporate memory mechanisms to improve long-term spatial and temporal consistency in video generation. These approaches retrieve generation history to localize relevant content across modalities such as RGB [90] and depth [16], often using surfel-indexed view selection [47] or camera FOV overlap [79, 90]. Other methods [36, 37, 77, 89] maintain a global point cloud map during generation, enabling the model to identify and reuse the most relevant spatial regions, thereby maintaining coherence across continuously generated scenes. In our work, the video generation with memory mechanism is realized by conditioning on both the character and the 3DGS scene. The 3DGS scene serves as a spatial memory that explicitly encodes the geometric and appearance information of the environment, while the character provides dynamic cues for motion and interaction.

World Models for Static Scene Creation. Existing world models that generate static yet explorable environments can be broadly categorized into two types. The first type [4, 11, 14, 19, 23, 33, 42, 46, 58, 69, 71, 72, 81, 87, 91, 96] stores the world implicitly within neural networks, using video generation models [5, 44, 73, 86] to visualize it. Users provide navigation commands (e.g., “camera forward”) drawn from a predefined set of camera trajectories, and the model synthesizes new frames along this path while maintaining spatial consistency with past generations. The second type [49, 66] explicitly constructs a 3DGS world, where multi-view images are optimized to form a manipulable 3D representation, allowing users to render novel views from arbitrary camera poses. Further developments extend this paradigm by using panoramic inputs [85, 97] or directly generating 3DGS representations from text or a single image through feed-forward networks [25, 51, 56, 70, 84], while others [17, 30, 68, 74, 95] integrate video generation




Fig. 2: (a) Each training sample consists of a 3D character and a video depicting the character performing an action described by a short text. Through segmentation and inpainting, we obtain the corresponding scene video and character mask sequence. The VAE encoder is then applied to encode these inputs into tokens. (b) CustomX predicts target video tokens conditioned on scene, mask, text, and multi-view character tokens within a Multi-Modal Diffusion Transformer, trained using Flow Matching [53]. Refer to Fig. 3 for the training process of the auto-regressive mode, which enables iterative interaction with CustomX, and Fig. 4 for the inference.

to streamline and accelerate 3DGS creation. In this work, users can either create or specify a 3DGS scene before generation. When users do not provide one, we adopt Marble [10] to automatically generate a static 3DGS world.

3 Method

In this section, we illustrate the detailed design of our method. Given a pre-generated or real-world reconstructed 3DGS scene $\mathcal{S}$ and a user-specified 3D character $\mathcal{C}$, **CustomX** enables users to iteratively control the character $\mathcal{C}$ via text instructions $\mathcal{T}$ within the scene $\mathcal{S}$, generating long-horizon, temporally coherent video clips that remain visually consistent with both $\mathcal{S}$ and $\mathcal{C}$.

We formulate this process as a conditional autoregressive video generation problem. At each iteration, CustomX synthesizes the current video clip conditioned on multiple multi-modal inputs: the previous generated clip, character representations, scene representations, and the current text instruction. Overviews of the training and inference pipelines are shown in Fig. 2 and Fig. 4, respectively.

3.1 Training Data Pre-Processing

Training Set Construction. As shown in Fig. 2(a), our training data is GTA-V [8], a game where players can control a character to perform basic actions such as “run forward”. We record gameplay sequences and segment them into short video clips, ensuring that each clip (1) contains only a single action and (2) has a fixed length of 129 frames. For each clip $\mathbf{V}$, we apply the following steps:

1. *Character Segmentation.* We use Grounded-SAM-2 [64] to segment characters and extract their bounding-box mask sequences, denoted as $\mathbf{M}$.
2. *Scene Inpainting.* We remove the segmented characters and apply DiffuEraser [48] to fill the missing regions, yielding the inpainted scene video $\mathbf{S}$.
3. *Action Labeling.* Each clip is then annotated with a concise text label $\mathbf{T}$ describing the action performed by the character, such as “The character is running forward”.

GTA-V also provides access to 3D character models used in the game. To ease the character modeling, we represent each character using four viewpoint renders [94]—front, left, right, and back—denoted as $\mathbf{C} = \{\mathbf{C}_F, \mathbf{C}_L, \mathbf{C}_R, \mathbf{C}_B\}$.

Finally, each processed training sample is represented as a tuple $\{\mathbf{V}, \mathbf{S}, \mathbf{M}, \mathbf{T}, \mathbf{C}\}$, forming a structured dataset.

Token Extraction. As shown in Fig. 2(a), given a training sample $\{\mathbf{V}, \mathbf{S}, \mathbf{M}, \mathbf{T}, \mathbf{C}\}$, we adopt the video VAE encoder from HunyuanCustom [35] to extract tokens for the video $\mathbf{V}$, scene $\mathbf{S}$, and mask $\mathbf{M}$. The resulting token sequences are denoted as $\mathcal{T}_V$, $\mathcal{T}_S$ and $\mathcal{T}_M$, respectively. The video VAE encoder operates with a spatial downsampling rate of 8 and a temporal downsampling rate of 4.

Note that the video VAE encoder can also be applied to single images. Therefore, for the multi-view character $\mathbf{C} = \{\mathbf{C}_F, \mathbf{C}_L, \mathbf{C}_R, \mathbf{C}_B\}$, we use the same encoder to extract tokens from each view, resulting in the multi-view character token set $\mathcal{T}_C = \{\mathcal{T}_{C_F}, \mathcal{T}_{C_L}, \mathcal{T}_{C_R}, \mathcal{T}_{C_B}\}$.

Finally, to extract text tokens, following HunyuanCustom [35], we employ the multi-modal encoder LLaVA [54], which takes both the text instruction $\mathbf{T}$ and character tokens $\mathcal{T}_C$ as input. The resulting encoded text tokens are denoted as $\mathcal{T}_T$. Implementation details are provided in the appendix.

Now, the training sample $\{\mathbf{V}, \mathbf{S}, \mathbf{M}, \mathbf{T}, \mathbf{C}\}$ is fully encoded into the latent space as $\{\mathcal{T}_V, \mathcal{T}_S, \mathcal{T}_M, \mathcal{T}_T, \mathcal{T}_C\}$.

3.2 Architecture

Training Objective. Fig. 2(b) illustrates the architecture of CustomX, whose backbone consists of a stack of full-attention Transformer blocks. We adopt Flow Matching [53] for model training, conditioned on multiple inputs (i.e., $\mathcal{T}_S$, $\mathcal{T}_M$, $\mathcal{T}_T$, and $\mathcal{T}_C$), to guide the generation process from pure noise to $\mathcal{T}_V$.

Concretely, given target video tokens $\mathcal{T}_V$, we first sample $t \in [0, 1]$ from a logit-normal distribution and initialize the noise $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ following Gaussian distribution. The intermediate sample $\mathbf{x}_t = (1 - t)\mathbf{x}_0 + t\mathcal{T}_V$ is then obtained via linear interpolation. The model is trained to predict the velocity $\mathbf{u}_t = d\mathbf{x}_t/dt$ by




Fig. 3: Illustration of the auto-regressive mode. The only difference from the original architecture in Fig. 2 is the addition of an extra conditioning input, i.e., the preceding video tokens. Note that a misalignment exists between training and inference: during training, the preceding video tokens are derived from ground-truth videos, whereas during inference, they are generated by the model itself. To mitigate this discrepancy, we add a small Gaussian noise to the preceding video tokens during training and refer to the resulting tokens as augmented preceding video tokens.

minimizing the mean squared error between the predicted velocity $\mathbf{v}_t$ and the ground-truth velocity $\mathbf{u}_t$:

$$\mathcal{L} = \mathbb{E}_{t, \mathbf{x}_0, \mathcal{T}_V} \|\mathbf{v}_t - \mathbf{u}_t\|^2. \quad (1)$$

Condition Modeling. We incorporate multiple conditioning signals to guide the learning process, including text tokens $\mathcal{T}_T$, multi-view character tokens $\mathcal{T}_C$, scene tokens $\mathcal{T}_S$, and mask tokens $\mathcal{T}_M$. As illustrated in Fig. 2(b), to inject scene and mask priors, we directly fuse $\mathcal{T}_S$ and $\mathcal{T}_M$ into the noisy target tokens $\mathbf{x}_t$ via: $\mathbf{x}'_t = \mathbf{x}_t + \text{Projector}([\mathcal{T}_S; \mathcal{T}_M])$, where $[\ ; \]$ denotes channel-wise concatenation. The Projector maps the concatenated tokens $[\mathcal{T}_S; \mathcal{T}_M]$ to the same dimensionality as $\mathbf{x}_t$ using a lightweight linear layer. We refer to the resulting $\mathbf{x}'_t$ as the conditional noisy tokens.

At last, to integrate the text tokens $\mathcal{T}_T$ and multi-view character tokens $\mathcal{T}_C$, we concatenate $\mathcal{T}_T$, $\mathcal{T}_C$, and $\mathbf{x}'_t$ along the sequence dimension. The concatenated sequence is then fed into the backbone, which is implemented as a stack of full-attention Transformer blocks, to denoise $\mathbf{x}_t$ under the supervision of Eq. (1).

Positional Embeddings. Following HunyuanCustom [35], no positional embeddings are added to the text tokens $\mathcal{T}_T$. A standard 3D-RoPE (over time, height, and width dimensions) is applied to the conditional noisy tokens $\mathbf{x}'_t$. For the multi-view character tokens $\mathcal{T}_C = \{\mathcal{T}_{C_F}, \mathcal{T}_{C_L}, \mathcal{T}_{C_R}, \mathcal{T}_{C_B}\}$, each $\mathcal{T}_{C_F}$, $\mathcal{T}_{C_L}$, $\mathcal{T}_{C_R}$, and $\mathcal{T}_{C_B}$ represents single-view character tokens. For each view, a shifted-3D-RoPE is applied. Taking $\mathcal{T}_{C_F}$ as an example,

$$\text{PE}_{\mathcal{T}_{C_F}}(i, j) = \text{3D-RoPE}(-1, i + w, j + h), \quad (2)$$

where (w, h) denotes the spatial size of $\mathcal{T}_{C_F}$, and the shifts along the temporal and spatial dimensions are -1 and (w, h) , respectively. For $\mathcal{T}_{C_L}$, $\mathcal{T}_{C_R}$ and $\mathcal{T}_{C_B}$, the spatial shift remains the same, while the temporal shifts are set to -2 , -3 , and -4 , respectively.




Fig. 4: Inference of CustomX. (a) Users first specify the inputs, including the character, 3DGS scene, virtual camera location, and character anchor. (b) The user-provided text instruction is parsed, and a corresponding camera path is generated. Applying this path to the 3DGS scene produces a rendered scene video. (c) CustomX then takes multiple inputs as conditions to generate the final output. Steps (b) and (c) can be performed iteratively, enabling temporally consistent, long-horizon interactions.

3.3 Auto-Regressive Mode

CustomX supports multi-round user interaction while maintaining temporal continuity and semantic coherence between adjacent video clips. To achieve this, we extend CustomX into an auto-regressive mode. Specifically, we divide the target video tokens $\mathcal{T}_V$ along the temporal dimension into two parts: the first quarter, denoted as $\mathcal{T}_{V_1}$, serves as the preceding video tokens, and the remaining three quarters, denoted as $\mathcal{T}_{V_{[2:4]}}$, serve as the *new* target video tokens.

The model is trained to generate $\mathcal{T}_{V_{[2:4]}}$ conditioned on both the preceding video tokens $\mathcal{T}_{V_1}$ and the other conditioning signals introduced in Sec. 3.2. As shown in Fig. 3, to incorporate the newly added condition $\mathcal{T}_{V_1}$, we prepend it to the conditional noisy tokens. The fusion strategy for the other conditioning signals remains unchanged, as described in Sec. 3.2.

3.4 Inference and Acceleration

Inference. Fig. 4 illustrates the inference pipeline of CustomX, which consists of three main stages:

1. *User Configuration.* Users first specify a character and a 3DGS scene. They can place a virtual camera at any desired viewpoint within the scene and define a single-frame bounding-box mask (i.e., an anchor) to indicate where the character should appear in the generated video. This anchor is automatically propagated to all subsequent frames, forming per-frame anchors and keeping the manual effort minimal. Users may also leverage existing tools [2, 7, 9, 38, 75, 78] to create the 3D character (e.g., Hunyuan3D [45]) or the 3DGS scene (e.g., World Labs Marble [10]).
2. *Scene Video Rendering.* Next, users provide a text instruction. The instruction is parsed into four categories: (a) locomotion, (b) gesture, (c) object-centric action, and (d) camera behavior. Each category determines a different

camera path. For (a), CustomX generates a camera trajectory consistent with the motion described in the text, for example, for “The character is running forward”, the camera follows a forward-moving path. For (b) and (c), the camera remains stationary. For (d), CustomX generates a trajectory matching the specified camera motion, for instance, for “The camera orbits around the character”, the camera follows a circular path around the character. CustomX then renders a scene video clip along the corresponding camera trajectory.

3. *CustomX Inference.* Finally, CustomX encodes the text instruction, character, scene, character anchor, and the previously generated clip (optional, for auto-regressive mode) into tokens using the encoders illustrated in Fig. 2. These tokens are then fed into the CustomX model to generate the final output.

Note that Steps 2 and 3 can be performed iteratively, enabling temporally consistent, long-horizon generation.

Acceleration. To accelerate inference, we adopt DMD2 [88] to distill the original 30-step denoising model into a more efficient 4-step version.

4 Experiment

4.1 Experimental Settings

Training Details. Our model is initialized from HunyuanCustom [35], which contains 13B parameters. We freeze the LLaVA encoder and the scene condition projector, and apply LoRA-style [34] fine-tuning to the backbone with a rank of 64. Two separate models are trained for 360P and 720P data using the AdamW optimizer with a learning rate of $1e-4$, each for 5,000 steps under the ZeRO-2 strategy. The 360P model is trained on $8 \times$ NVIDIA H100 GPUs, while the 720P model, owing to its higher resolution, is trained on $8 \times$ NVIDIA B200 GPUs. Further details on training and acceleration are provided in the appendix.

Training Data. Following the data curation procedure described in Sec. 3.1, we construct a training set comprising 2,084 video samples featuring five characters. Each sample is annotated with text labels describing either locomotion actions ({“run forward”, “run to the left”, “run to the right”, “run backward”}) or camera behaviors ({“orbit”, “follow” }). A key observation of CustomX is that post-training on such simple locomotion and camera-behavior data can substantially enhance pre-trained models, improving both action dynamics and camera control capability. This is analogous to large language models [27, 80], where extensive pre-training endows rich world knowledge, while post-training, which is conducted with a much smaller dataset, serves to activate specific capabilities (e.g., in GPT-3 [60], post-training data is orders of magnitude smaller than the pre-training corpus). For each sample, we prepare two resolutions, 360P and 720P, to train models of different quality levels. Unless otherwise specified, the 360P version is used by default. Additional details are provided in the appendix.

Table 1: WorldScore metrics for evaluating generation quality, categorized into three groups: (1) controllability, (2) quality, and (3) dynamics. The static and dynamic scores are computed by aggregating metrics from these three groups. † denotes dedicated world models. *Ctrl*: Controllability; *Align*: Alignment; *Const*: Consistency; *Photo*: Photometric; *Acc*: Accuracy; *Mag*: Magnitude; *Smooth*: Smoothness.

Method	WorldScore		Controllability			Quality				Dynamics		
	Static	Dyna- mic	Camera Ctrl	Object Ctrl	Content Align	3D Const	Photo Const	Style Const	Subjective Quality	Motion Acc	Motion Mag	Motion Smooth
CogVideoX1.5-I2V (5B) [86]	60.08	56.77	42.13	100.00	31.12	68.65	81.43	77.86	19.39	54.92	24.58	67.62
HunyuanVideo-I2V (13B) [44]	56.43	55.14	27.30	100.00	13.70	58.24	93.22	91.97	10.58	59.36	24.46	72.58
Wan2.1-I2V (14B) [73]	57.91	55.87	37.32	98.33	26.34	81.88	83.77	65.03	12.73	59.59	28.60	65.07
Wan2.2-I2V (14B) [73]	54.52	51.74	24.79	100.00	24.03	57.95	59.44	<u>98.91</u>	16.51	<u>59.60</u>	<u>38.93</u>	37.26
Wan-VACE (14B) [40]	51.54	52.03	21.29	100.00	30.39	27.53	54.18	99.02	17.39	56.73	34.78	67.98
Wan-VACE (1.3B) [40]	50.75	49.38	31.12	100.00	11.54	29.66	61.00	97.67	24.26	55.87	29.48	53.17
HunyuanCustom (13B) [35]	62.64	<u>61.11</u>	47.19	100.00	<u>72.07</u>	48.06	31.40	97.47	<u>25.84</u>	59.04	24.05	<u>89.56</u>
DeepVerse† [16]	52.63	47.63	52.48	75.00	18.80	35.47	83.16	92.39	11.09	33.30	33.61	40.97
Hunyuan-GameCraft† [46]	<u>69.92</u>	57.77	<u>77.45</u>	83.33	51.16	85.91	82.12	93.39	16.11	16.65	31.83	39.77
Matrix-Game-2.0† [33]	52.26	43.98	15.10	<u>99.17</u>	12.38	60.29	68.35	97.60	12.92	3.07	47.36	23.59
CustomX (Ours)	84.64	77.22	88.91	100.00	75.73	<u>83.57</u>	<u>87.68</u>	<u>98.91</u>	57.72	61.08	24.12	94.47

Evaluation. We evaluate our model across four aspects: (1) visual quality, (2) camera controllability, (3) action control capability and generalization to novel actions, and (4) character consistency on novel characters. For (1) and (2), we adopt the WorldScore [20] metrics to assess the generated samples. For (3), we measure the control success rate via human evaluation and report the CLIP [62] text-to-image similarity score, covering both the four seen locomotion actions and up to 142 novel actions. For (4), we assess character similarity between the ground-truth character and the generated one using 30 novel characters, evaluated by DINOv2 [59] and CLIP [62] scores. By default, we use the 360P version of our model. At each iteration, our model generates 96 frames, using the previously generated 33 frames as conditions when available. Unless otherwise noted, evaluations are conducted on the first generated clip.

4.2 Main Results

Visual Quality Evaluation. We adopt the metrics proposed by WorldScore [20] to evaluate the visual quality of generated videos. WorldScore defines a suite of metrics categorized into three groups: *Controllability*, *Quality*, and *Dynamics*. The first two primarily assess visual fidelity in static regions, while the last evaluates motion quality in dynamic regions (i.e., the character in our case).

We compare our model with dedicated world models, including DeepVerse [16], Hunyuan-GameCraft [46], and Matrix-Game-2.0 [33], that focus on controlling the main entity, as well as with video generation base models listed in Tab. 1.

For each model, we use the first two metric groups to evaluate 60 generated videos covering 30 different characters and two camera behaviors, {“orbit”, “follow”}. The third metric group is evaluated on 146 videos, where each video features the same character performing a distinct action, either one of four locomotion actions {“run forward”, “run to the right”, “run to the left”, “run backward”} or one of 142 novel actions unseen during training. Note that (1) to

Table 2: Action control and character consistency evaluation on general foundation models and dedicated world models[†]. Locomotion actions include “run forward”, “run to the left”, “run to the right”, and “run backward”, while richer actions encompass 142 gesture and object-centric actions. Note that the three world models restrict action control to locomotion only.

Method	Action Control				Character Consistency	
	Success Rate (%)		CLIP Text-Image Score		DINOv2 Score	CLIP Score
	Locomotion Actions	Richer Actions	Locomotion Actions	Richer Actions		
CogVideoX1.5-I2V (5B) [86]	23.3	21.1	0.261	0.273	0.594	0.611
HunyuanVideo-I2V (13B) [44]	26.7	35.2	0.272	0.293	0.645	<u>0.709</u>
Wan2.1-I2V (14B) [73]	26.7	64.8	0.267	0.302	0.627	0.678
Wan2.2-I2V (14B) [73]	53.3	<u>74.6</u>	0.272	<u>0.303</u>	<u>0.650</u>	0.704
Wan-VACE (14B) [40]	43.3	73.2	0.270	<u>0.303</u>	0.398	0.541
Wan-VACE (1.3B) [40]	26.7	13.4	0.261	0.269	0.504	0.548
HunyuanCustom (13B) [35]	<u>56.7</u>	51.4	<u>0.273</u>	0.297	0.558	0.665
DeepVerse [†] [16]	6.7	-	0.259	-	0.291	0.523
Hunyuan-GameCraft [†] [46]	16.7	-	0.255	-	0.329	0.529
Matrix-Game-2.0 [†] [33]	3.3	-	0.255	-	0.339	0.524
CustomX (Ours)	100.0	80.7	0.279	0.305	0.698	0.721

the best of our knowledge, our model is the first framework to jointly integrate 3D or multi-view characters, action text prompts, scenes, and long-horizon generation, which necessitates an input configuration that naturally diverges from established baselines; (2) both foundation and dedicated models are evaluated using their original control mechanisms for character and camera, without any modification; (3) the three dedicated world models only support locomotion actions, thus they are evaluated solely on those; and (4) for models requiring an initial image input, we use Google Gemini [3] to render the corresponding character within the scene as the initialization. The results are shown in Tab. 1.

Action Control and Generalization. In Tab. 2, we evaluate the model’s action control capability on both seen actions (four locomotion actions) and 142 novel actions (referred to as “richer actions”). The novel actions cover both gesture-based behaviors (e.g., “*wave hands*”) and object-centric interactions (e.g., “*play a guitar*”). For each model, we conduct 30 evaluations on locomotion actions and 142 evaluations on distinct novel actions using the same set of characters. We report the action control success rate via human evaluation and the CLIP text-image similarity score, computed as the average frame-wise text-image score.

The results in Tab. 2 reveal that our model outperforms the base model, HunyuanCustom [35]. This phenomenon can be interpreted through the lens of post-training in large language models [27, 60, 63, 80], where fine-tuning typically does not disrupt the pre-trained representation space; rather, it adjusts the *response style*, for example, to make the outputs more helpful or harmless, while preserving the extensive knowledge acquired during pre-training. In our case, the structurally simple fine-tuning data, composed primarily of fundamental



Fig. 5: Screenshot visualizations of videos generated by CustomX, showcasing different characters performing various novel actions across two scenes. Additional examples are provided in the appendix and supplementary videos.

locomotion behaviors, serve to refine *motion dynamics* and align *human embodiment representations*, rather than to redefine the model’s generative space. Consequently, our fine-tuning strategy enhances motion fidelity and behavioral coherence while maintaining the broad semantic and generative generalization inherited from large-scale pre-training.

Character Consistency Evaluation. A key feature of our model is its ability to maintain consistent visual identity between the provided character and the one appearing in the generated videos. In Tab. 2, we evaluate character consistency using DINOv2 and CLIP scores. Both metrics measure the similarity between the generated and ground-truth characters. During evaluation, we crop the character region from each generated frame and compute the similarity score; the final result is averaged across all frames. Since our method takes multi-view inputs, each frame is scored by its maximum similarity to the ground-truth views. Each model is evaluated 30 times with different characters performing locomotion.

Visualization. Fig. 5 shows screenshots from the generated videos, featuring different characters performing various actions across diverse scenes. Fig. 6 illustrates a long-horizon generation example in which users interact with the model four times. The resulting video follows the specified character, scene, and text instructions, while maintaining temporal coherence across generated clips. Additional examples are provided in the appendix and supplementary videos.

4.3 Ablation Study

Multi-View Character Condition. We compare our four-view character-condition model with two baselines: a single-view model and a front-back-view model. To evaluate character consistency, we generate videos by instructing characters to run toward the front, left, right, and back, which naturally reveals their appearance from multiple viewpoints. We then compute DINOv2 and CLIP scores following the protocol in Sec. 4.2. We use the same dataset as that used in Tab. 2. As shown in Tab. 3, character consistency improves as more view inputs are used.

Character Anchor Condition. As illustrated in Fig. 2, we introduce an additional condition, namely a character mask, to enable the model to distinguish




Fig. 6: Screenshot visualizations of a long video generated by CustomX, showcasing a character iteratively exploring a scene with diverse actions. Additional long-form results are available in the appendix and supplementary videos.

Table 3: Using multi-view character inputs improves character consistency. By default, we employ all four views.

Character View	DINOv2 Score	CLIP Score
Front	0.556	0.628
+Back	0.613	0.678
+Right and Left	0.698	0.721

Table 4: Using per-frame character anchors helps the model distinguish dynamic entities from the static scene, leading to higher DINOv2 and CLIP character consistency scores.

Character Anchor Type	DINOv2 Score	CLIP Score
w/o Anchor	0.477	0.529
w/ First-Frame Anchor	0.597	0.645
w/ Per-Frame Anchor	0.698	0.721

dynamic entities from static scenes. In the default training setting, a bounding-box mask, referred to as an anchor, is extracted for each frame. During inference, as shown in Fig. 4, users only need to specify a single-frame anchor, which is then propagated to all subsequent frames to form per-frame anchors. Tab. 4 compares our default model, which uses per-frame anchors, with two variants: (1) a model without any anchor during both training and inference, and (2) a model that uses only the first-frame anchor, without propagating it to subsequent frames, throughout training and inference.

Visual Conditions Enhance Long-Horizon Generation. We introduce two types of visual conditions, multi-view character and 3DGS scene, to ensure both character and scene consistency during generation. Beyond improving spatial coherence, these visual conditions also alleviate the issue of visual quality degra-




Fig. 7: Using both visual conditions, namely the 3DGS scene and the multi-view character, significantly improves long-horizon visual quality across diverse video clips.




Fig. 8: Evaluation of the hybrid data training strategy. (a) Training solely on game data causes the model to inherit a game-engine rendering style in synthesized characters. (b) Incorporating real-world data improves photorealism and enables the model to capture physical effects, such as dynamic clothing wrinkles, which are absent in GTA-V.

dation over long-horizon generation. To validate this, we compare CustomX operating in the auto-regressive mode (see Sec. 3.3) against two variants: (1) using only the multi-view character condition, where the 3DGS scene condition is replaced with textual scene descriptions; and (2) using only the 3DGS scene condition, where the multi-view character condition is replaced with textual character descriptions. Fig. 7 presents a comparison with the two variants, using the CLIP-Aesthetic and DINOv2 scores as metrics to evaluate visual quality and character consistency, respectively. The evaluation protocol consists of 10 trials per model. In each trial, a different character is used to autoregressively generate 10 video clips, forming an iterative long-horizon generation process.

Game-Real Hybrid Data Enhances Real-World Character Fidelity and Physics. CustomX trained solely on GTA-V videos tends to inherit the game-engine rendering style, leading to stylized outputs even when conditioned on real-world characters. We therefore investigate whether incorporating real-world data into training can improve generation realism. Specifically, we collect additional 400 real-world videos using the same data collection pipeline as for GTA-V, except that the videos are captured from the real world rather than rendered by a game engine. As shown in Fig. 8, the model trained on the hybrid dataset achieves higher realism and captures physical effects present in the real world but absent in games, such as dynamic clothing wrinkles. Additional details of the real-world dataset and further results are provided in the appendix.

Acceleration. Our 13B-parameter base model generates a 93-frame 360P video in 121s on a single NVIDIA H100 using a 30-step denoising schedule. By applying DMD2 [88], we distill it into a 4-step version, reducing latency to 21s with

only slight drops in DINOv2 (0.698→0.669) and CLIP-Aesthetic (5.665→5.583) scores. Latencies for higher resolutions are provided in the appendix.

5 Conclusion

We present *CustomX*, a novel framework that allows users to provide a character and a 3DGS scene, enabling iterative interaction for both character control and world exploration. Unlike prior controllable-entity models that restrict the agent to a small set of predefined actions, CustomX supports open-ended control over a wide range of behaviors through natural language commands. CustomX delivers substantial improvements in motion dynamics and character consistency over its base model, as validated across a broad set of metrics including visual quality, action controllability, character fidelity, and long-horizon generation capability.

Limitations. As a controllable video generation framework, CustomX shares common limitations with existing video generation models. Although inference has been accelerated, it still does not operate in real time. Furthermore, the current model is limited to single-subject scenarios and cannot generate videos involving multiple simultaneously active characters.

Acknowledgements

This work is supported in part by Microsoft Research, the HKU Startup Fund, and HKU Musketeers Foundation Institute of Data Science.

References

1. ByteDance Seed: Seedance 2.0. https://seed.bytedance.com/en/seedance2_0 (2026), accessed: 2026-06-25
2. Deemos Technologies: Rodin. <https://hyper3d.ai/> (2026), accessed: 2026-06-25
3. Google: Gemini. <https://gemini.google.com/app> (2026), accessed: 2026-06-25
4. Google: Genie 3: A new frontier for world models. <https://deepmind.google/blog/genie-3-a-new-frontier-for-world-models/> (2026), accessed: 2026-06-25
5. Google: Veo. <https://deepmind.google/models/veo/> (2026), accessed: 2026-06-25
6. Kuaishou: Kling ai. <https://klingai.com/global/> (2026), accessed: 2026-06-25
7. Meshy: Meshy ai. <https://www.meshy.ai/> (2026), accessed: 2026-06-25
8. Rockstar Games: Grand theft auto v. <https://www.rockstargames.com/gta-v> (2026), accessed: 2026-06-25
9. VAST AI Research: Tripo. <https://www.tripo3d.ai/> (2026), accessed: 2026-06-25
10. World Labs: Marble. <https://marble.worldlabs.ai/> (2026), accessed: 2026-06-25
11. World Labs: Rtfm: A real-time frame model. <https://rtfm.worldlabs.ai/> (2026), accessed: 2026-06-25
12. Bai, J., Xia, M., Fu, X., Wang, X., Mu, L., Cao, J., Liu, Z., Hu, H., Bai, X., Wan, P., et al.: Recammaster: Camera-controlled generative rendering from a single video. arXiv preprint arXiv:2503.11647 (2025)

13. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023)
14. Bruce, J., Dennis, M.D., Edwards, A., Parker-Holder, J., Shi, Y., Hughes, E., Lai, M., Mavalankar, A., Steigerwald, R., Apps, C., et al.: Genie: Generative interactive environments. In: Forty-first International Conference on Machine Learning (2024)
15. Chen, J., Li, X., Bai, X., Ma, T., Zhang, P., Chen, Z., Li, G., Liu, L., Zhao, S., Li, B., et al.: Omniinsert: Mask-free video insertion of any reference via diffusion transformer models. arXiv preprint arXiv:2509.17627 (2025)
16. Chen, J., Zhu, H., He, X., Wang, Y., Zhou, J., Chang, W., Zhou, Y., Li, Z., Fu, Z., Pang, J., et al.: Deepverse: 4d autoregressive video generation as a world model. arXiv preprint arXiv:2506.01103 (2025)
17. Chen, L., Zhou, Z., Zhao, M., Wang, Y., Zhang, G., Huang, W., Sun, H., Wen, J.R., Li, C.: Flexworld: Progressively expanding 3d scenes for flexible-view synthesis. arXiv preprint arXiv:2503.13265 (2025)
18. Chen, T.S., Siarohin, A., Menapace, W., Fang, Y., Lee, K.S., Skorokhodov, I., Aberman, K., Zhu, J.Y., Yang, M.H., Tulyakov, S.: Multi-subject open-set personalization in video generation. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 6099–6110 (2025)
19. Decart, E., McIntyre, Q., Campbell, S., Chen, X., Wachen, R.: Oasis: A universe in a transformer (2024)
20. Duan, H., Yu, H.X., Chen, S., Fei-Fei, L., Wu, J.: Worldscore: A unified evaluation benchmark for world generation. arXiv preprint arXiv:2504.00983 (2025)
21. Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., et al.: Scaling rectified flow transformers for high-resolution image synthesis. In: Forty-first international conference on machine learning (2024)
22. Fei, Z., Li, D., Qiu, D., Wang, J., Dou, Y., Wang, R., Xu, J., Fan, M., Chen, G., Li, Y., et al.: Skyreels-a2: Compose anything in video diffusion transformers. arXiv preprint arXiv:2504.02436 (2025)
23. Feng, R., Zhang, H., Yang, Z., Xiao, J., Shu, Z., Liu, Z., Zheng, A., Huang, Y., Liu, Y., Zhang, H.: The matrix: Infinite-horizon world generation with real-time moving control. arXiv preprint arXiv:2412.03568 (2024)
24. Feng, W., Liu, J., Tu, P., Qi, T., Sun, M., Ma, T., Zhao, S., Zhou, S., He, Q.: I2vcontrol-camera: Precise video camera control with adjustable motion strength. arXiv preprint arXiv:2411.06525 (2024)
25. Go, H., Park, B., Jang, J., Kim, J.Y., Kwon, S., Kim, C.: Splatflow: Multi-view rectified flow model for 3d gaussian splatting synthesis. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 21524–21536 (2025)
26. Gu, Z., Yan, R., Lu, J., Li, P., Dou, Z., Si, C., Dong, Z., Liu, Q., Lin, C., Liu, Z., et al.: Diffusion as shader: 3d-aware video diffusion for versatile video generation control. In: Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers. pp. 1–12 (2025)
27. Guo, D., Yang, D., Zhang, H., Song, J., Zhang, R., Xu, R., Zhu, Q., Ma, S., Wang, P., Bi, X., et al.: Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. arXiv preprint arXiv:2501.12948 (2025)
28. Guo, Y., Yang, C., Rao, A., Liang, Z., Wang, Y., Qiao, Y., Agrawala, M., Lin, D., Dai, B.: Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. arXiv preprint arXiv:2307.04725 (2023)

29. HaCohen, Y., Chiprut, N., Brazowski, B., Shalem, D., Moshe, D., Richardson, E., Levin, E., Shiran, G., Zabari, N., Gordon, O., Panet, P., Weissbuch, S., Kulikov, V., Bitterman, Y., Melumian, Z., Bibi, O.: Ltx-video: Realtime video latent diffusion. arXiv preprint arXiv:2501.00103 (2024)
30. Hao, J., Wang, P., Wang, H., Zhang, X., Guo, Z.: Gaussvideodreamer: 3d scene generation with video diffusion and inconsistency-aware gaussian splatting. arXiv preprint arXiv:2504.10001 (2025)
31. He, H., Xu, Y., Guo, Y., Wetzstein, G., Dai, B., Li, H., Yang, C.: Cameractrl: Enabling camera control for text-to-video generation. arXiv preprint arXiv:2404.02101 (2024)
32. He, H., Yang, C., Lin, S., Xu, Y., Wei, M., Gui, L., Zhao, Q., Wetzstein, G., Jiang, L., Li, H.: Cameractrl ii: Dynamic scene exploration via camera-controlled video diffusion models. arXiv preprint arXiv:2503.10592 (2025)
33. He, X., Peng, C., Liu, Z., Wang, B., Zhang, Y., Cui, Q., Kang, F., Jiang, B., An, M., Ren, Y., et al.: Matrix-game 2.0: An open-source, real-time, and streaming interactive world model. arXiv preprint arXiv:2508.13009 (2025)
34. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: Lora: Low-rank adaptation of large language models. ICLR **1**(2), 3 (2022)
35. Hu, T., Yu, Z., Zhou, Z., Liang, S., Zhou, Y., Lin, Q., Lu, Q.: Hunyuancustom: A multimodal-driven architecture for customized video generation. arXiv preprint arXiv:2505.04512 (2025)
36. Huang, J., Hu, X., Han, B., Shi, S., Tian, Z., He, T., Jiang, L.: Memory forcing: Spatio-temporal memory for consistent scene generation on minecraft. arXiv preprint arXiv:2510.03198 (2025)
37. Huang, T., Zheng, W., Wang, T., Liu, Y., Wang, Z., Wu, J., Jiang, J., Li, H., Lau, R.W., Zuo, W., Guo, C.: Voyager: Long-range and world-consistent video diffusion for explorable 3d scene generation. arXiv preprint arXiv:2506.04225 (2025)
38. Huang, Z., Guo, Y.C., Wang, H., Yi, R., Ma, L., Cao, Y.P., Sheng, L.: Mv-adapter: Multi-view consistent image generation made easy. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16377–16387 (2025)
39. Jiang, Y., Wu, T., Yang, S., Si, C., Lin, D., Qiao, Y., Loy, C.C., Liu, Z.: Video-booth: Diffusion-based video generation with image prompts. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6689–6700 (2024)
40. Jiang, Z., Han, Z., Mao, C., Zhang, J., Pan, Y., Liu, Y.: Vace: All-in-one video creation and editing. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 17191–17202 (2025)
41. Ju, X., Ye, W., Liu, Q., Wang, Q., Wang, X., Wan, P., Zhang, D., Gai, K., Xu, Q.: Fulldit: Multi-task video generative foundation model with full attention. arXiv preprint arXiv:2503.19907 (2025)
42. Kanervisto, A., Bignell, D., Wen, L.Y., Grayson, M., Georgescu, R., Valcarcel Macua, S., Tan, S.Z., Rashid, T., Pearce, T., Cao, Y., et al.: World and human action models towards gameplay ideation. Nature **638**(8051), 656–663 (2025)
43. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics **42**(4) (July 2023)
44. Kong, W., Tian, Q., Zhang, Z., Min, R., Dai, Z., Zhou, J., Xiong, J., Li, X., Wu, B., Zhang, J., et al.: Hunyuanvideo: A systematic framework for large video generative models. arXiv preprint arXiv:2412.03603 (2024)

45. Lai, Z., Zhao, Y., Liu, H., Zhao, Z., Lin, Q., Shi, H., Yang, X., Yang, M., Yang, S., Feng, Y., et al.: Hunyuan3d 2.5: Towards high-fidelity 3d assets generation with ultimate details. arXiv preprint arXiv:2506.16504 (2025)
46. Li, J., Tang, J., Xu, Z., Wu, L., Zhou, Y., Shao, S., Yu, T., Cao, Z., Lu, Q.: Hunyuan-gamecraft: High-dynamic interactive game video generation with hybrid history condition. arXiv preprint arXiv:2506.17201 (2025)
47. Li, R., Torr, P., Vedaldi, A., Jakab, T.: Vmem: Consistent interactive video scene generation with surfel-indexed view memory. arXiv preprint arXiv:2506.18903 (2025)
48. Li, X., Xue, H., Ren, P., Bo, L.: Diffueraser: A diffusion model for video inpainting. arXiv preprint arXiv:2501.10018 (2025)
49. Li, X., Lai, Z., Xu, L., Qu, Y., Cao, L., Zhang, S., Dai, B., Ji, R.: Director3d: Real-world camera trajectory and 3d scene generation from text. *Advances in neural information processing systems* **37**, 75125–75151 (2024)
50. Li, Z., Qian, D., Su, K., Diao, Q., Xia, X., Liu, C., Yang, W., Zhang, T., Yuan, Z.: Bindweave: Subject-consistent video generation via cross-modal integration. arXiv preprint arXiv:2510.00438 (2025)
51. Liang, H., Cao, J., Goel, V., Qian, G., Korolev, S., Terzopoulos, D., Plataniotis, K.N., Tulyakov, S., Ren, J.: Wonderland: Navigating 3d scenes from a single image. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 798–810 (2025)
52. Liang, S., Yu, Z., Zhou, Z., Hu, T., Wang, H., Chen, Y., Lin, Q., Zhou, Y., Li, X., Lu, Q., et al.: Omniv2v: Versatile video generation and editing via dynamic content manipulation. arXiv preprint arXiv:2506.01801 (2025)
53. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. arXiv preprint arXiv:2210.02747 (2022)
54. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. *Advances in neural information processing systems* **36**, 34892–34916 (2023)
55. Liu, L., Ma, T., Li, B., Chen, Z., Liu, J., Li, G., Zhou, S., He, Q., Wu, X.: Phantom: Subject-consistent video generation via cross-modal alignment. arXiv preprint arXiv:2502.11079 (2025)
56. Liu, Y., Min, Z., Wang, Z., Wu, J., Wang, T., Yuan, Y., Luo, Y., Guo, C.: World-mirror: Universal 3d world reconstruction with any-prior prompting. arXiv preprint arXiv:2510.10726 (2025)
57. Luo, Y., Bai, J., Shi, X., Xia, M., Wang, X., Wan, P., Zhang, D., Gai, K., Xue, T.: Camclonemaster: Enabling reference-based camera control for video generation. arXiv preprint arXiv:2506.03140 (2025)
58. Mao, X., Lin, S., Li, Z., Li, C., Peng, W., He, T., Pang, J., Chi, M., Qiao, Y., Zhang, K.: Yume: An interactive world generation model. arXiv preprint arXiv:2507.17744 (2025)
59. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193 (2023)
60. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al.: Training language models to follow instructions with human feedback. *Advances in neural information processing systems* **35**, 27730–27744 (2022)
61. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 4195–4205 (2023)

62. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PmLR (2021)
63. Rafailov, R., Sharma, A., Mitchell, E., Manning, C.D., Ermon, S., Finn, C.: Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems* **36**, 53728–53741 (2023)
64. Ren, T., Liu, S., Zeng, A., Lin, J., Li, K., Cao, H., Chen, J., Huang, X., Chen, Y., Yan, F., et al.: Grounded sam: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159* (2024)
65. Ren, X., Shen, T., Huang, J., Ling, H., Lu, Y., Nimier-David, M., Müller, T., Keller, A., Fidler, S., Gao, J.: Gen3c: 3d-informed world-consistent video generation with precise camera control. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 6121–6132 (2025)
66. Shriram, J., Trevithick, A., Liu, L., Ramamoorthi, R.: Realmdreamer: Text-driven 3d scene generation with inpainting and depth diffusion. *arXiv preprint arXiv:2404.07199* (2024)
67. Sitzmann, V., Rezkikov, S., Freeman, B., Tenenbaum, J., Durand, F.: Light field networks: Neural scene representations with single-evaluation rendering. *Advances in Neural Information Processing Systems* **34**, 19313–19325 (2021)
68. Sun, W., Chen, S., Liu, F., Chen, Z., Duan, Y., Zhang, J., Wang, Y.: Dimensionx: Create any 3d and 4d scenes from a single image with controllable video diffusion. *arXiv preprint arXiv:2411.04928* (2024)
69. Sun, W., Wei, F., Zhao, J., Chen, X., Chen, Z., Zhang, H., Zhang, J., Lu, Y.: From virtual games to real-world play. *arXiv preprint arXiv:2506.18901* (2025)
70. Szymanowicz, S., Zhang, J.Y., Srinivasan, P., Gao, R., Brussee, A., Holynski, A., Martin-Brualla, R., Barron, J.T., Henzler, P.: Bolt3d: Generating 3d scenes in seconds. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 24846–24857 (2025)
71. Team, R., Gao, Z., Wang, Q., Zeng, Y., Zhu, J., Cheng, K.L., Li, Y., Wang, H., Xu, Y., Ma, S., Chen, Y., Liu, J., Cheng, Y., Yao, Y., Zhu, J., Meng, Y., Zheng, K., Bai, Q., Chen, J., Shen, Z., Yu, Y., Zhu, X., Shen, Y., Ouyang, H.: Advancing open-source world models. *arXiv preprint arXiv:2601.20540* (2026)
72. Valevski, D., Leviathan, Y., Arar, M., Fruchter, S.: Diffusion models are real-time game engines. *arXiv preprint arXiv:2408.14837* (2024)
73. Wan, T., Wang, A., Ai, B., Wen, B., Mao, C., Xie, C.W., Chen, D., Yu, F., Zhao, H., Yang, J., Zeng, J., Wang, J., Zhang, J., Zhou, J., Wang, J., Chen, J., Zhu, K., Zhao, K., Yan, K., Huang, L., Feng, M., Zhang, N., Li, P., Wu, P., Chu, R., Feng, R., Zhang, S., Sun, S., Fang, T., Wang, T., Gui, T., Weng, T., Shen, T., Lin, W., Wang, W., Wang, W., Zhou, W., Wang, W., Shen, W., Yu, W., Shi, X., Huang, X., Xu, X., Kou, Y., Lv, Y., Li, Y., Liu, Y., Wang, Y., Zhang, Y., Huang, Y., Li, Y., Wu, Y., Liu, Y., Pan, Y., Zheng, Y., Hong, Y., Shi, Y., Feng, Y., Jiang, Z., Han, Z., Wu, Z.F., Liu, Z.: Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314* (2025)
74. Wang, H., Liu, F., Chi, J., Duan, Y.: Videoscene: Distilling video diffusion model to generate 3d scenes in one step. In: *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 16475–16485. IEEE (2025)
75. Wang, Y., Xu, X., Ma, L., Wang, H., Dai, B.: Boosting 3d object generation through pbr materials. In: *SIGGRAPH Asia 2024 Conference Papers*. pp. 1–11 (2024)

76. Wang, Z., Yuan, Z., Wang, X., Li, Y., Chen, T., Xia, M., Luo, P., Shan, Y.: Motionctrl: A unified and flexible motion controller for video generation. In: ACM SIGGRAPH 2024 Conference Papers. pp. 1–11 (2024)
77. Wu, T., Yang, S., Po, R., Xu, Y., Liu, Z., Lin, D., Wetzstein, G.: Video world models with long-term spatial memory. arXiv preprint arXiv:2506.05284 (2025)
78. Xiang, J., Chen, X., Xu, S., Wang, R., Lv, Z., Deng, Y., Zhu, H., Dong, Y., Zhao, H., Yuan, N.J., et al.: Native and compact structured latents for 3d generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14419–14429 (2026)
79. Xiao, Z., Lan, Y., Zhou, Y., Ouyang, W., Yang, S., Zeng, Y., Pan, X.: Worldmem: Long-term consistent world simulation with memory. arXiv preprint arXiv:2504.12369 (2025)
80. Yang, A., Li, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Gao, C., Huang, C., Lv, C., et al.: Qwen3 technical report. arXiv preprint arXiv:2505.09388 (2025)
81. Yang, M., Li, J., Fang, Z., Chen, S., Yu, Y., Fu, Q., Yang, W., Ye, D.: Playable game generation. arXiv preprint arXiv:2412.00887 (2024)
82. Yang, S., Hou, L., Huang, H., Ma, C., Wan, P., Zhang, D., Chen, X., Liao, J.: Direct-a-video: Customized video generation with user-directed camera movement and object motion. In: ACM SIGGRAPH 2024 Conference Papers. pp. 1–12 (2024)
83. Yang, X., Xu, J., Luan, K., Zhan, X., Qiu, H., Shi, S., Li, H., Yang, S., Zhang, L., Yu, C., et al.: Omnicam: Unified multimodal video generation via camera control. arXiv preprint arXiv:2504.02312 (2025)
84. Yang, Y., Shao, J., Li, X., Shen, Y., Geiger, A., Liao, Y.: Prometheus: 3d-aware latent diffusion models for feed-forward text-to-3d scene generation. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 2857–2869 (2025)
85. Yang, Z., Ge, W., Li, Y., Chen, J., Li, H., An, M., Kang, F., Xue, H., Xu, B., Yin, Y., et al.: Matrix-3d: Omnidirectional explorable 3d world generation. arXiv preprint arXiv:2508.08086 (2025)
86. Yang, Z., Teng, J., Zheng, W., Ding, M., Huang, S., Xu, J., Yang, Y., Hong, W., Zhang, X., Feng, G., et al.: Cogvideox: Text-to-video diffusion models with an expert transformer. arXiv preprint arXiv:2408.06072 (2024)
87. Ye, D., Zhou, F., Lv, J., Ma, J., Zhang, J., Lv, J., Li, J., Deng, M., Yang, M., Fu, Q., et al.: Yan: Foundational interactive video generation. arXiv preprint arXiv:2508.08601 (2025)
88. Yin, T., Gharbi, M., Park, T., Zhang, R., Shechtman, E., Durand, F., Freeman, W.T.: Improved distribution matching distillation for fast image synthesis. In: NeurIPS (2024)
89. Yu, H.X., Duan, H., Hur, J., Sargent, K., Rubinstein, M., Freeman, W.T., Cole, F., Sun, D., Snavely, N., Wu, J., et al.: Wonderjourney: Going from anywhere to everywhere. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6658–6667 (2024)
90. Yu, J., Bai, J., Qin, Y., Liu, Q., Wang, X., Wan, P., Zhang, D., Liu, X.: Context as memory: Scene-consistent interactive long video generation with memory retrieval. arXiv preprint arXiv:2506.03141 (2025)
91. Yu, J., Qin, Y., Wang, X., Wan, P., Zhang, D., Liu, X.: Gamefactory: Creating new games with generative interactive videos. arXiv preprint arXiv:2501.08325 (2025)
92. YU, M., Hu, W., Xing, J., Shan, Y.: Trajectorycrafter: Redirecting camera trajectory for monocular videos via diffusion models. arXiv preprint arXiv:2503.05638 (2025)

93. Yu, W., Xing, J., Yuan, L., Hu, W., Li, X., Huang, Z., Gao, X., Wong, T.T., Shan, Y., Tian, Y.: Viewcrafter: Taming video diffusion models for high-fidelity novel view synthesis. arXiv preprint arXiv:2409.02048 (2024)
94. Zhang, L., Wang, Z., Zhang, Q., Qiu, Q., Pang, A., Jiang, H., Yang, W., Xu, L., Yu, J.: Clay: A controllable large-scale generative model for creating high-quality 3d assets. *ACM Transactions on Graphics (TOG)* **43**(4), 1–20 (2024)
95. Zhang, S., Li, J., Fei, X., Liu, H., Duan, Y.: Scene splatter: Momentum 3d scene generation from single image with video diffusion model. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 6089–6098 (2025)
96. Zhang, Y., Peng, C., Wang, B., Wang, P., Zhu, Q., Kang, F., Jiang, B., Gao, Z., Li, E., Liu, Y., et al.: Matrix-game: Interactive world foundation model. arXiv preprint arXiv:2506.18701 (2025)
97. Zhang, Z., Hold-Geoffroy, Y., Hašan, M., Chen, Z., Luan, F., Dorsey, J., Hu, Y.: Generating 360° video is what you need for a 3d scene. arXiv preprint arXiv:2504.02045 (2025)
98. Zhou, J., Gao, H., Voleti, V., Vasishta, A., Yao, C.H., Boss, M., Torr, P., Rupprecht, C., Jampani, V.: Stable virtual camera: Generative view synthesis with diffusion models. arXiv preprint arXiv:2503.14489 (2025)