arXiv:2603.16871v1 [cs.CV] 17 Mar 2026

WorldCam: Interactive Autoregressive 3D Gaming Worlds with Camera Pose as a Unifying Geometric Representation

Jisu Nam[1*], **Yicong Hong**[2], **Chun-Hao Paul Huang**[2], **Feng Liu**[2], **JoungBin Lee**[1], **Jiyoung Kim**[1], **Siyoon Jin**[1], **Yunsung Lee**[3], **Jaeyoon Jung**[3], **Suhwan Choi**[3], **Seungryong Kim**[1†], **Yang Zhou**[2†]

[1]KAIST, [2]Adobe Research, [3]MAUM AI

[†]Co-corresponding Authors, [*]Work done during Adobe Research internship.

Recent advances in video diffusion transformers have enabled interactive gaming world models that allow users to explore generated environments over extended horizons. However, existing approaches struggle with precise action control and long-horizon 3D consistency. Most prior works treat user actions as abstract conditioning signals, overlooking the fundamental geometric coupling between actions and the 3D world, whereby actions induce relative camera motions that accumulate into a global camera pose within a 3D world. In this paper, we establish camera pose as a unifying geometric representation to jointly ground immediate action control and long-term 3D consistency. First, we define a physics-based continuous action space and represent user inputs in the Lie algebra to derive precise 6-DoF camera poses, which are injected into the generative model via a camera embedder to ensure accurate action alignment. Second, we use global camera poses as spatial indices to retrieve relevant past observations, enabling geometrically consistent revisiting of locations during long-horizon navigation. To support this research, we introduce a large-scale dataset comprising 3,000 minutes of authentic human gameplay annotated with camera trajectories and textual descriptions. Extensive experiments show that our approach substantially outperforms state-of-the-art interactive gaming world models in action controllability, long-horizon visual quality, and 3D spatial consistency.

Date: 2026.03.17

Project Page: https://cvlab-kaist.github.io/WorldCam/

1 Introduction

Recent advances in Video Diffusion Transformers (DiTs) (Yang et al., 2024; Wan et al., 2025; HaCohen et al., 2024; Agarwal et al., 2025) have significantly improved the realism and scalability of video generation. Building on this progress, recent studies (Valevski et al., 2024; Che et al., 2024; Bruce et al., 2024; Zhang et al., 2025; He et al., 2025; Gao et al., 2025; Chen et al., 2025c) take important steps toward *interactive gaming world models*, demonstrating that generative models can simulate playable environments. However, despite plausible visual outputs, they still struggle with precise action control and 3D world consistency, which are prerequisites for a functional gaming engine.

This limitation arises because prior works have overlooked the **fundamental geometric coupling between user actions and the 3D world**. In gaming environments, user actions (e.g., keyboard presses and mouse movements) are not abstract control signals, but instead induce *relative camera motions within a 3D scene*. These relative motions accumulate over time to form the camera's global trajectory, which dictates how the underlying 3D world is projected into 2D observations. As a result, accurate action control and 3D consistency are not independent objectives, but are inherently coupled through the camera pose.

Despite this geometric coupling, existing interactive gaming world models (Valevski et al., 2024; Che et al., 2024; Bruce et al., 2024; Zhang et al., 2025; He et al., 2025; Gao et al., 2025; Chen et al., 2025c) treat actions as abstract conditioning signals by directly injecting raw action inputs into video generative models. This design leads to misaligned camera motion and inconsistent 3D geometry due to the lack of explicit geometric constraints. Several camera-controlled video generation methods (Wang et al., 2024b; He et al., 2024; Wan et al., 2025) address 3D consistency by conditioning generation on camera poses, however, they focus on short videos (e.g., 16 frames) and fail to model action-driven control and long-horizon inference.

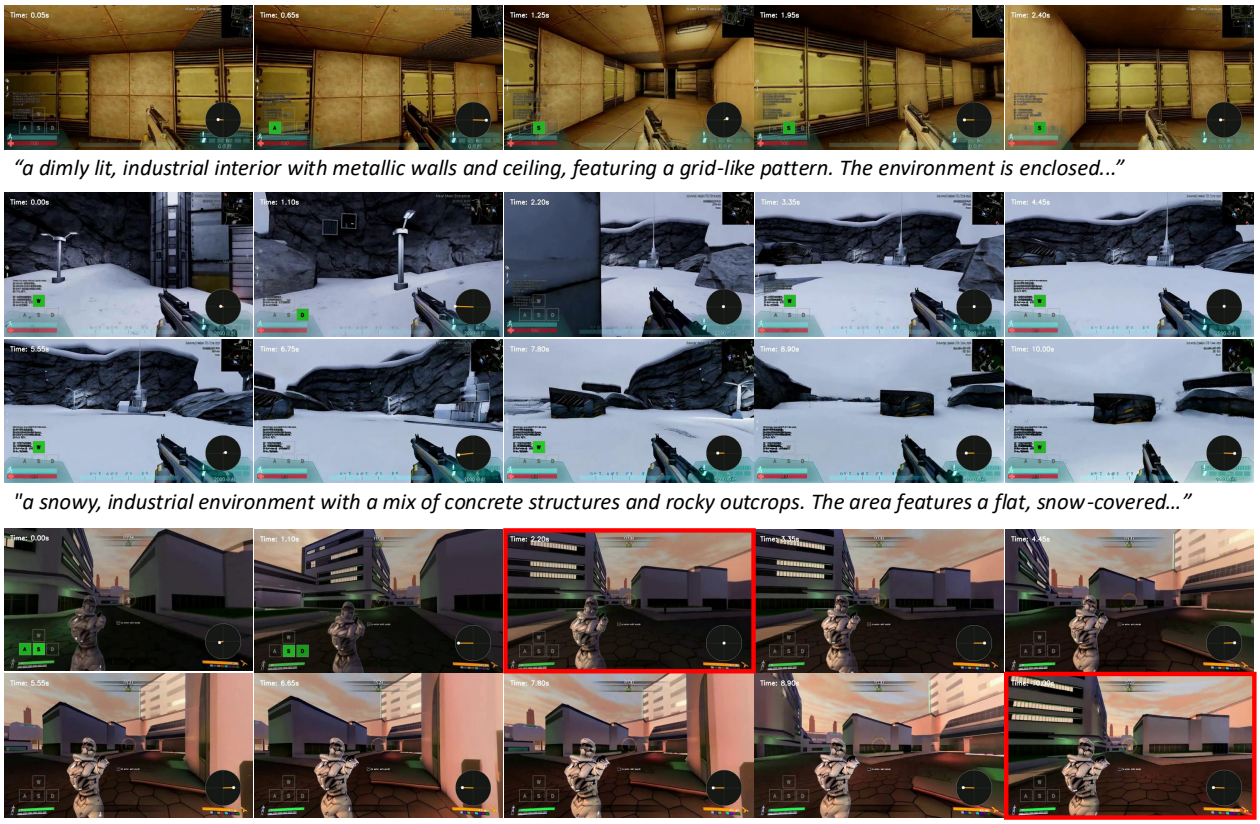

Figure 1 Teaser (Best viewed in color and zoomed in): WorldCam is an interactive 3D gaming model that enables **precise action control under challenging keyboard and mouse inputs (top), supports long-horizon interactions (middle), and preserves consistent 3D geometry across viewpoints (bottom).** Time (seconds at 20 FPS) is visualized in the top-left of each frame, while keyboard and mouse inputs are shown in the bottom-left and bottom-right, respectively. The red box highlights consistent 3D geometry in revisited views.

In this paper, we introduce **WorldCam**, a foundation model for interactive gaming worlds built on a video DiT backbone. WorldCam enables precise action control, long-horizon inference, and consistent 3D world modeling. Our core contribution lies in establishing **camera pose as a unifying geometric representation** that jointly grounds both immediate action control and long-horizon 3D consistency.

Compared to prior works that directly inject raw action signals into video generative models, we define a physics-based continuous action space that translates complex user action inputs (e.g., coupled keyboard and mouse actions) into geometrically accurate camera poses. Unlike prior approaches (Li et al., 2025a) relying on naive, decoupled linear approximations, which inherently fail to capture coupled dynamics such as screw motion, we model user actions as spatial velocities in the Lie algebra (Hall, 2013) and strictly derive precise 6-DoF relative poses. These poses are then encoded as Plücker embeddings (Sitzmann et al., 2021) via a camera embedder and injected into intermediate features of the video DiT, ensuring that the generated outputs adhere to the intended physical motion.

Crucially, the camera pose serves a dual purpose, acting not only as a control signal for actions but also as an explicit geometric cue for long-horizon 3D consistency. Specifically, we retrieve relevant previously generated latents based on camera pose similarity between the current and past camera poses. These latents are concatenated with the current latent sequence, and their associated camera pose embeddings establish geometric correspondences between the current latents and retrieved past latents. By grounding both action control and 3D geometry in a shared camera space, we achieve precise action control and long-horizon 3D world consistency.

A major obstacle to building interactive gaming world models is the lack of large-scale, high-fidelity video datasets that capture real human gameplay dynamics. Existing works often rely on Minecraft datasets (Chen et al., 2025c; Guo et al., 2025; Po et al., 2025; Zhang et al., 2025), which are limited by simplified geometry and discrete motion patterns, or on closed-licensed game video datasets (Li et al., 2025a; Tang et al., 2025; Bruce et al., 2024) that are inaccessible for reproducible research.

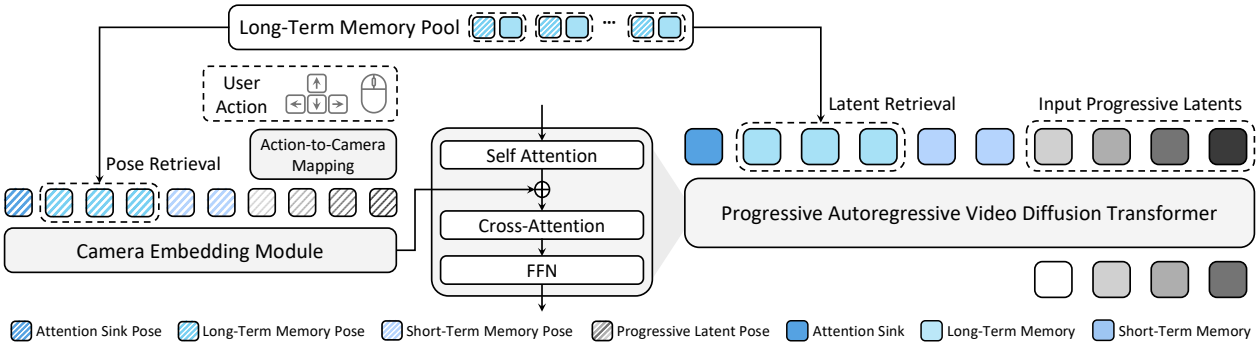

Figure 2 Overall architecture. WorldCam converts user actions into camera poses in Lie algebra and conditions a progressive autoregressive video transformer on these camera poses for precise action control. Retrieved long-term memory latents and camera poses from the memory pool enforce 3D consistency of the generated world, while short-term memory with an attention sink stabilizes long-horizon generation.

Table 1 Comparison with existing interactive gaming world models and camera-controlled models.

	The Matrix	Matrix-game	Matrix-game 2.0	Genie	Yan	GameCraft	CameraCtrl	MotionCtrl	**WorldCam**
Action control	✓	✓	✓	✓	✓	✓	×	×	✓
3D consistency	×	×	×	×	×	×	✓	✓	✓
Long-horizon inference	×	×	✓	✓	✓	✓	×	×	✓

To address this, we introduce **WorldCam-50h**, a large-scale dataset comprising 3,000 minutes of authentic human gameplay collected from one closed-licensed commercial game, *Counter-Strike*, and two open-licensed games, *Xonotic* and *Unvanquished*. To capture diverse human player behaviors, the dataset covers complex scenarios such as general navigation, rapid 360° camera rotations, and reverse traversal across varied geometries. Furthermore, each video is annotated with rich textual descriptions (Yang et al., 2025) and pseudo ground-truth camera poses (Huang et al., 2025).

Through extensive experiments, we demonstrate that our approach simultaneously achieves precise action alignment, sustained long-horizon visual quality, and robust 3D spatial consistency. WorldCam consistently outperforms prior interactive gaming world models (He et al., 2025; Mao et al., 2025; Li et al., 2025a) as well as camera-controllable video generation methods (Wang et al., 2024b; Wan et al., 2025). We validate the effectiveness of our design choices through comprehensive ablation studies. We will publicly release our open-licensed datasets, code, and pretrained models.

2 Related Work

Interactive gaming world models. Prior works typically inject raw action signals into video generative models via cross-attention (Feng et al., 2024; He et al., 2025; Valevski et al., 2024), AdaLN (Xiao et al., 2025), or text (Mao et al., 2025; Chen et al., 2025b). However, by treating actions as abstract conditioning signals, they often fail to model accurate camera motion and achieve 3D consistency, since raw actions lack an understanding of the underlying 3D scene. GameCraft (Li et al., 2025a) improves action control by linearly approximating user inputs into camera poses, but ignores the underlying $SE(3)$ geometry by decoupling translation and rotation. This makes it struggle to capture coupled dynamics such as screw motion arising from entangled keyboard and mouse inputs. Moreover, camera motion in GameCraft is used only for immediate control, lacking a persistent geometric anchor for 3D consistency. In contrast to prior works that fail to simultaneously satisfy precise action control, long-horizon inference, and 3D consistency, we propose a camera-grounded framework that uses camera pose as a unified geometric representation to achieve all three (Table 1).

Interactive gaming datasets. Large-scale gaming video datasets that capture authentic human gameplay motion are crucial for training interactive gaming world models. Prior works (Chen et al., 2025c; Guo et al., 2025; Po et al., 2025; Zhang et al., 2025) commonly rely on Minecraft dataset (Guss et al., 2019), which provides paired action labels but suffers from limited visual diversity and simplified geometry. Recent efforts (Tang et al., 2025; Li et al., 2025a) utilize internal, closed-licensed gameplay datasets, hindering research reproducibility. In contrast, we introduce a large-scale, open-licensed dataset capturing diverse and dynamic human gameplay, fully annotated with pseudo ground-truth camera poses and textual descriptions.

3 WorldCam

Given an initial RGB observation I_0 (a single image or a short video clip), a text prompt c_{text}, and a sequence of user actions $\{A_1, A_2, \ldots\}$, our goal is to autoregressively generate a sequence of video frames $\{I_1, I_2, \ldots\}$ that (i) accurately follow the user actions (Section 3.2 & 3.3), (ii) remain consistent with a single 3D world (Section 3.4), and (iii) maintain high visual quality over long horizons (Section 3.5).

Table 1 compares our method with recent interactive world models (Feng et al., 2024; Zhang et al., 2025; He et al., 2025; Bruce et al., 2024; Ye et al., 2025; Li et al., 2025a) and camera-controlled approaches (Wan et al., 2025; Wang et al., 2024b). Figure 2 provides an overview of the overall architecture.

3.1 Baseline: Video Diffusion Transformer

We build WorldCam on a pretrained video Diffusion Transformer (DiT), Wan-2.1-T2V (Wan et al., 2025), which consists of spatio-temporal self-attention and text cross-attention layers. Given an input video $V \in \mathbb{R}^{F \times H \times W \times 3}$, the VAE encoder maps it to a latent sequence $\mathbf{z}_0 \in \mathbb{R}^{f \times h \times w \times c}$. Here, $(F, H, W, 3)$ denote the number of frames, height, width, and RGB channels of the video, while (f, h, w, c) denote the corresponding dimensions in the latent space.

Gaussian noise $\epsilon \sim \mathcal{N}(0, I)$ is added to z_0 to obtain noisy latents $\mathbf{z}_t \in \mathbb{R}^{f \times h \times w \times c}$ at timestep $t \in [0, T]$. Given a text prompt c_{text}, the DiT v_θ learns to predict the velocity field that transports $\mathbf{z}_t$ toward $\mathbf{z}_0$ (Lipman et al., 2022). The training objective is defined as

$$\boldsymbol{L}_{\text{FM}} = \mathbb{E}_{\mathbf{z}_0, c_{\text{text}}, t}\left[\left\| v_\theta(\mathbf{z}_t, c_{\text{text}}, t) - \frac{\mathbf{z}_0 - \mathbf{z}_t}{1 - t} \right\|_2^2\right]. \tag{1}$$

3.2 Action-to-Camera Mapping

A central challenge in interactive 3D world modeling is to translate user actions into physically consistent camera motion. Prior works often directly inject raw action signals into the generative model (He et al., 2025; Feng et al., 2024; Che et al., 2024), rely on text prompts to describe actions (Mao et al., 2025), or adopt linear action-to-camera approximations (Li et al., 2025a). Such formulations frequently lead to misaligned or geometrically inconsistent camera trajectories, especially under complex motions involving coupled translation and rotation.

To ensure physically accurate and fine-grained action control, we define the action space in the Lie algebra $\mathfrak{se}(3)$. At each transition from frame I_{i-1} to I_i, the user action A_i is represented as a twist vector $A_i = [\mathbf{v}_i; \boldsymbol{\omega}_i] \in \mathbb{R}^6$, where $\mathbf{v}_i = [v_x, v_y, v_z]^\top \in \mathbb{R}^3$ and $\boldsymbol{\omega}_i = [\omega_x, \omega_y, \omega_z]^\top \in \mathbb{R}^3$ denote the linear and angular velocities, respectively. We then derive the corresponding relative camera pose $\Delta P_i \in SE(3)$ via the matrix exponential map:

$$\Delta P_i = \exp(\hat{A}_i) = \begin{bmatrix} \Delta R_i & \Delta t_i \\ \mathbf{0}^\top & 1 \end{bmatrix}, \tag{2}$$

where $\hat{A}_i \in \mathfrak{se}(3)$ denotes the 4×4 matrix of the twist A_i.

Unlike decoupled linear approximations (Li et al., 2025a; Chen et al., 2025b; Li et al., 2025b) that update translation and rotation independently, our formulation jointly integrates linear and angular velocities directly on the $SE(3)$ manifold. This design yields geometrically precise camera trajectories even under complex user actions involving tightly coupled translation and rotation.

3.3 Camera-Controlled Video Generation

We design a video generative model conditioned on the camera motion derived from the actions. Given the sequence of relative camera poses $\{\Delta P_i\}_{i=1}^{F}$ which are obtained via action-to-camera mapping (Section 3.2), we accumulate them into global camera poses aligned with the first frame of the window. To provide explicit view-dependent geometric conditioning, the camera poses are converted into Plücker embeddings $\hat{P} \in \mathbb{R}^{F \times 6}$ (Sitzmann et al., 2021).

To inject camera control into the DiT, we introduce a lightweight camera embedding module c_ϕ consisting of two MLP layers. Since the VAE temporally compresses the input by a factor of r, we align camera conditioning with the latent sequence by concatenating r consecutive Plücker embeddings for each latent frame, resulting in $\hat{\mathbf{p}} \in \mathbb{R}^{f \times (6r)}$, where $f = F/r$. The camera embeddings are added to the DiT features $\mathbf{d}$ after each self-attention layer:

$$\mathbf{d} \leftarrow \mathbf{d} + c_\phi(\hat{\mathbf{p}}). \tag{3}$$

3.4 Pose-Anchored Long-Term Memory

Beyond precise action control, we leverage the camera pose derived in Section 3.2 as a geometric anchor to maintain 3D consistency, ensuring spatial coherence when revisiting locations or viewpoints.

Global pose accumulation. Since our action-to-camera mapping strictly adheres to $SE(3)$ geometry, relative motions can be reliably accumulated into global camera poses. Specifically, the global pose of the j-th frame I_j is computed via pose composition:

$$P_j^{\text{global}} = P_{j-1}^{\text{global}} \circ \Delta P_j, \quad P_0^{\text{global}} = \mathbf{I}, \tag{4}$$

where $\mathbf{I}$ is the identity and $\circ$ denotes pose composition.

Pose-indexed memory retrieval. To exploit long-horizon spatial context, we maintain a long-term memory pool $\mathcal{M}$ that stores previously generated latents with their global camera poses. Due to the temporal compression of the video VAE, each latent z_j corresponds to a short clip of r consecutive frames and is therefore associated with a set of r global camera poses. For simplicity, we denote each memory entry as $(P_j^{\text{global}}, z_j)$, where P_j^{global} represents the set of global camera poses associated with the r frames:

$$P_j^{\text{global}} = \begin{bmatrix} R_j & t_j \\ \mathbf{0}^\top & 1 \end{bmatrix}. \tag{5}$$

Regarding the complex layouts and frequent occlusions in our gaming environments, we adopt a hierarchical memory retrieval strategy that uses the global camera pose as a spatial index over the memory bank. We first select the top-K candidates $\mathcal{M}_{\text{trans}}$ whose camera positions are closest to the current position t_i:

$$\mathcal{M}_{\text{trans}} = \text{TopK}_K \left(-\|t_j - t_i\|_2 \; ; \; (P_j^{\text{global}}, z_j) \in \mathcal{M} \right). \tag{6}$$

From $\mathcal{M}_{\text{trans}}$, we further select L entries ($L \leq K$) whose viewing directions are most aligned with the current orientation R_i, measured using the trace of the relative rotation matrix:

$$\mathcal{M}_{\text{rot}} = \text{TopK}_L \left(\text{tr}(R_j^\top R_i) \; ; \; (P_j^{\text{global}}, z_j) \in \mathcal{M}_{\text{trans}} \right). \tag{7}$$

Long-term memory conditioning. The retrieved memory entries $(P_j^{\text{global}}, z_j) \in \mathcal{M}_{\text{rot}}$ provide the geometric context required to enforce spatial consistency during generation. We concatenate the retrieved latents with the current input latent sequence. Their associated camera poses are realigned to the first frame of the current denoising window, embedded via the camera embedding module c_ϕ, and injected into the intermediate features of the DiT (Equation 3). This allows the model to establish explicit geometric correspondences between the current latents and long-term memory latents.

3.5 Progressive Autoregressive Inference

Progressive noise scheduling. To support long-horizon autoregressive video diffusion, we adopt a progressive per-frame noise schedule that assigns monotonically increasing noise levels to latent frames within each denoising window. This design provides a reliable low-noise anchor in early frames, while keeping future frames at higher noise levels and thus correctable. This enables efficient cross-frame conditioning with large window overlap and stable rollout (Xie et al., 2025; Chen et al., 2024, 2025a). During training, conditioning on partially noisy context improves robustness to corrupted frames and reduces the train–test mismatch that would otherwise amplify error accumulation at inference time.

Specifically, we discretize the diffusion process into N inference steps with monotonically increasing noise levels and partition them into S stages, where S equals the number of noisy latent frames maintained in the sequence. The noisy latent sequence $\mathbf{z}_t$ (Section 3.1) is reformulated as a stage-wise $\mathbf{z}_s$, where $s \in \{0, \ldots, S-1\}$ denotes the noise stage, and each latent frame must complete all denoising steps across all stages to be fully denoised. After completing all S stages, the latent sequence is shifted forward: the earliest latent frame is evicted and decoded by the VAE, and a newly initialized pure-noise latent frame is appended to the end of the sequence.

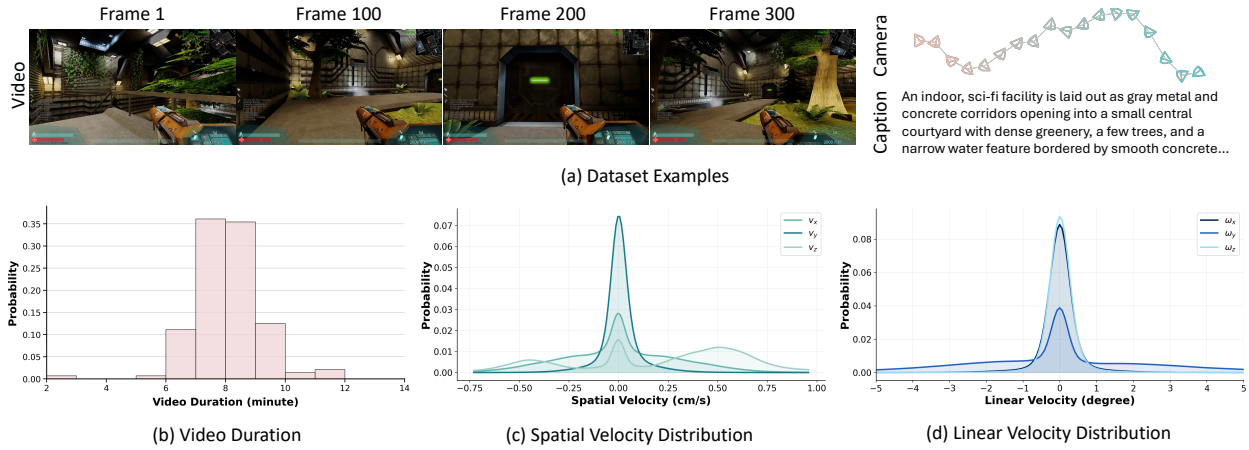

Figure 3 Dataset samples and statistics: (a) Example gameplay frames annotated with camera trajectories, and text captions. (b) Distribution of video durations. (c) Distribution of linear velocities (v_x, v_y, v_z). (d) Distribution of angular velocities $(\omega_x, \omega_y, \omega_z)$. The dataset captures diverse and authentic human gameplay behaviors for training interactive gaming world models.

Attention sink. While progressive noise improves temporal stability, it can still accumulate errors under large and complex gaming motions, leading to visual saturation and distorted UI elements. To mitigate this, we incorporate an attention sink mechanism inspired by StreamingLLM (Xiao et al., 2023), which stabilizes attention by anchoring a small set of initial tokens. During inference, we retain global initial frames as attention anchors, helping preserve frame fidelity, scene style, and UI consistency.

Short-term memory. In practice, we find that providing recently generated latents, referred to as short-term memory, is crucial for reducing error drift during autoregressive generation. We empirically set the number of short-term memory latents to match the number of generated latents, striking a balance between stability and computational efficiency.

4 WorldCam-50h

Large-scale interactive gaming datasets that capture authentic human action dynamics are essential for training foundational gaming world models, yet remain challenging to collect. Prior works (Xiao et al., 2025; He et al., 2025; Guo et al., 2025) often rely on the Minecraft dataset (Guss et al., 2019), which offers limited visual diversity, or on closed-licensed gameplay videos (Li et al., 2025a), which hinder reproducible research.

To address these limitations, we introduce **WorldCam-50h**, a large-scale dataset of human gameplay videos. The dataset is annotated with detailed textual descriptions and camera pose information. Dataset samples and statistics are shown in Figure 3. Data preprocessing pipeline is described in Appendix A.

Data collection. We record human gameplay from one closed-licensed game, *Counter-Strike*[1], and two open-licensed titles, *Xonotic*[2] and *Unvanquished*[3], chosen for their complex 3D environments and high interactivity. Note that all game-derived figures and video clips in the paper and supplementary material are from *Xonotic* and *Unvanquished*, licensed under CC BY-SA 2.5 and GPL v3. We focus on single-player exploration of static environments without dynamic objects or other players. To capture realistic human action dynamics, participants are instructed to perform diverse behaviors, including navigation, combined keyboard–mouse inputs, rapid camera movements, and revisiting locations (Figure 3(c) & (d)). We collect over 100 videos per game, each averaging 8 minutes (Figure 3(b)), totaling about 1,000 minutes (≈17 hours) of gameplay per game.

Captioning. While prior works (Bruce et al., 2024; Po et al., 2025; Gao et al., 2025) often discard textual captions during training, we find textual guidance essential for maintaining frame quality and scene style. We therefore generate detailed captions for each training video chunk using Qwen2.5-VL-7B (Yang et al., 2025). Specifically, we prompt the

[1] https://www.counter-strike.net.
[2] https://xonotic.org.
[3] https://unvanquished.net.

model with: *"Describe the static world of this video in one concise paragraph. Focus on the global layout (overall topology, primary regions, spatial arrangement, and key objects), the visual theme (colors, materials, and architectural style), and the ambient environmental conditions (overall lighting and weather)."*

Camera annotation. We further extract global camera pose annotations using ViPE (Huang et al., 2025) to estimate both camera intrinsics and extrinsics for each one-minute segment of gameplay video. Since ViPE can produce erroneous estimates (e.g., unrealistically large translations), we apply additional filtering based on a predefined threshold on translation magnitudes.

5 Experiments

5.1 Implementation Details

We use Wan2.1-1.3B-T2V (Wan et al., 2025) as our video DiT backbone. The spatial resolution is 480×832. During training, we use 8 progressive latents, 8 short-term memory latents, and 4 long-term memory latents. All experiments are conducted on 8 NVIDIA H100 GPUs.

Training is performed in three stages: (1) camera-controlled video generation with short-term memory for 10k iterations with a batch size of 64; (2) progressive autoregressive training with short-term memory for 10k iterations with a batch size of 48; and (3) progressive autoregressive training with both short- and long-term memory for 10k iterations with a batch size of 16 to enforce 3D consistency.

For progressive autoregressive training, we use $N = 64$ and $S = 8$, indicating that each latent is denoised for 8 sampling steps per stage, resulting in a total of 64 sampling timesteps across 8 stages. We use the AdamW optimizer with a learning rate of 1×10^{-5} for all stages.

5.2 Evaluation Settings

We compare our method with state-of-the-art interactive gaming world models, including Yume (Mao et al., 2025), Matrix-Game 2.0 (He et al., 2025), and GameCraft (Li et al., 2025a), as well as camera-controlled video generation models, including CameraCtrl (He et al., 2024) and MotionCtrl (Wang et al., 2024b). Additional evaluation details for the compared baselines are provided in B.

We evaluate three key aspects: action controllability, visual quality, and 3D consistency. To evaluate complex action controllability, we randomly sample 70 action trajectories across 70 starting images from the test split of our dataset. We generate 200-frame sequences for interactive gaming world models (Table 2) and 16-frame sequences for camera-controlled methods, since they do not support long-horizon inference and suffer from severe visual drift under window-based autoregressive rollout (Table 4). For interactive gaming world models, action inputs are carefully mapped to each model's corresponding trained action space. For camera-controlled models, which require explicit camera trajectories, we provide ground-truth camera poses.

For visual quality and 3D consistency, we randomly select 50 starting images from our test dataset and predefine four simple closed-loop trajectories per image, resulting in 200 evaluation cases. We generate 200-frame sequences for interactive gaming world models (Table 3). Here, camera-controlled methods are not included, as they do not support long-horizon video generation and therefore cannot be evaluated for long-horizon 3D consistency.

5.3 Evaluation Metrics

Action controllability. We evaluate how accurately generated videos follow input actions using average Relative Pose Errors in translation ($\text{RPE}_{\text{trans}}$), rotation ($\text{RPE}_{\text{rot}}$), and camera extrinsics ($\text{RPE}_{\text{camera}}$) between camera trajectories estimated from generated videos by ViPE (Huang et al., 2025) and ground-truth camera motion. We apply Sim(3) Umeyama alignment (Umeyama, 2002) to compensate for differences in scale and coordinate frames.

Visual quality. We evaluate long-horizon visual quality using VBench++ (Huang et al., 2024) metrics. We report Aesthetic Quality (Aesth.), Subject Consistency (Subj. Cons.), Background Consistency (Bg. Cons.), Imaging Quality (Img.), Temporal Flickering (Temp.), and Motion Smoothness (Motion.), along with the average score (Avg.) across all categories.

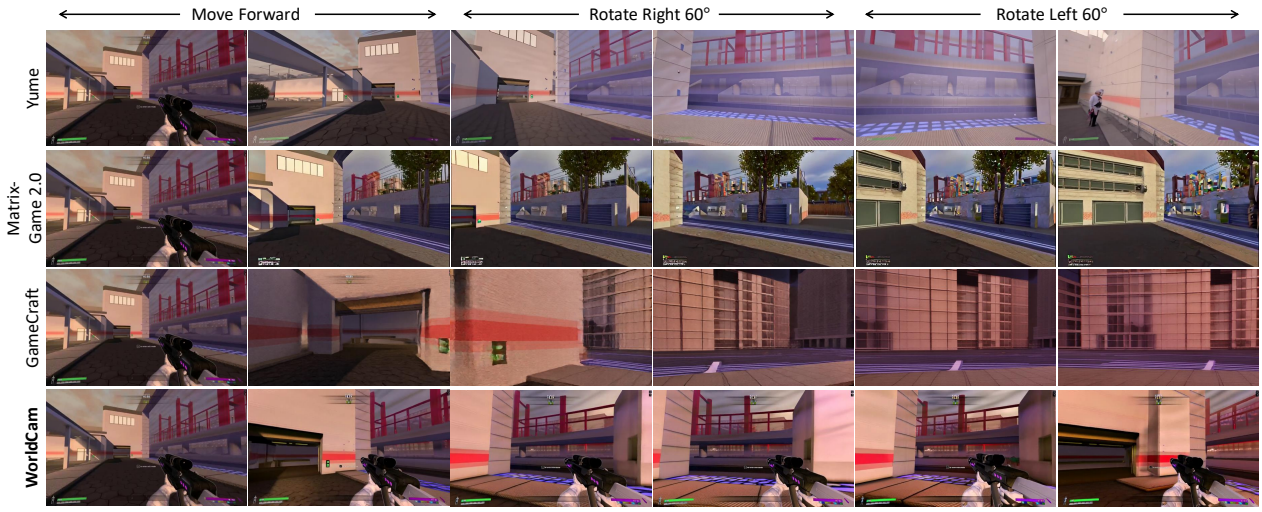

Figure 4 Qualitative comparison with recent interactive gaming world models: Compared to prior works, WorldCam faithfully follows user actions and maintains coherent 3D scene structure with high visual fidelity over long horizons.

Table 2 Quantitative comparison with recent interactive gaming world models in terms of action controllability and visual quality. All methods are evaluated under the same action trajectories with 200-frame video generation.

Method	Action Controllability			Visual Quality						
	RPE_{trans} ↓	RPE_{rot} (°) ↓	RPE_{camera} ↓	Avg.↑	Aesth.↑	Subj. Cons.↑	Bg. Cons.↑	Img.↑	Temp.↑	Motion.↑
Yume (Mao et al., 2025)	0.111	2.222	0.137	0.774	0.476	0.741	0.892	0.600	0.955	**0.986**
Matrix-Game 2.0 (He et al., 2025)	0.098	1.656	0.119	0.766	0.457	0.741	0.843	0.633	0.937	0.981
GameCraft (Li et al., 2025a)	0.086	1.146	0.100	0.781	0.464	0.804	0.850	0.626	0.958	**0.986**
WorldCam	**0.080**	**0.696**	**0.086**	**0.844**	**0.508**	**0.896**	**0.959**	**0.752**	**0.964**	0.984

3D consistency. To evaluate whether the generated environment maintains geometric consistency over long horizons, we measure view consistency using PSNR and LPIPS between all intermediate palindromic frame pairs in 200-frame closed-loop camera trajectories.

To account for the potential effect of blur on pixel-level metrics, we additionally report Sharpness, measured as the variance of the Laplacian (Pertuz et al., 2013), averaged over all frames in each video. Blurry outputs (low sharpness) tend to yield artificially high PSNR and low LPIPS due to reduced pixel variance.

Table 3 Quantitative results for long-horizon 3D consistency.

Method	PSNR↑	LPIPS↓	MEt3R↓	DINO Sim.↑	Sharpness↑
Real Videos	-	-	-	-	577
Yume (Mao et al., 2025)	16.03	0.5629	0.0905	0.4545	95
Matrix-Game 2.0 (He et al., 2025)	13.66	0.4997	0.0662	0.6153	179
GameCraft (Li et al., 2025a)	14.27	0.5749	0.0489	0.5960	201
WorldCam	**16.69**	**0.3277**	**0.0342**	**0.8884**	**656**

We further report MEt3R (Asim et al., 2025), which measures geometric multi-view consistency by reconstructing dense 3D point maps with DUSt3R (Wang et al., 2024a) between consecutive frames, warping DINO features (Caron et al., 2021) across views via the estimated geometry, and computing their similarity.

In addition, we report DINO Similarity, defined as the cosine similarity between DINOv2 (Oquab et al., 2023) features of palindrome-corresponding frame pairs, directly measuring appearance consistency at the same camera pose.

Human evaluation. Due to the absence of comprehensive benchmarks for interactive gaming world models, we supplemented our quantitative results with a user study focusing on action controllability, visual quality, and 3D consistency. We recruited 30 participants, each assessing 20 test cases drawn randomly from a large evaluation pool. The evaluation was conducted in a blind and randomized manner. Participants rated each model on a scale of 1 (worst) to 5 (best) based on detailed evaluation criteria.

5.4 Comparison

Quantitative results. Table 2 and 3 summarize the quantitative comparison between our method and state-of-the-art interactive gaming world models (Mao et al., 2025; He et al., 2025; Li et al., 2025a). Our approach consistently outperforms existing baselines in terms of action controllability, visual quality, and 3D consistency on 200-frame long-horizon sequences by a large margin.

Table 4 Comparison with camera-controlled models under 16-frame generation.

Method	RPE_{trans} ↓	RPE_{rot} (°) ↓	RPE_{camera} ↓
CameraCtrl (He et al., 2024)	0.071	0.943	0.083
MotionCtrl (Wang et al., 2024b)	0.080	1.271	0.102
WorldCam	**0.026**	**0.386**	**0.030**

In Table 2, while prior methods relying on text prompts (Mao et al., 2025), raw action signals (He et al., 2025), or linearly approximated cameras (Li et al., 2025a) suffer from suboptimal action accuracy, our model achieves the lowest errors across all pose metrics, with RPE_{camera} improved by 16.3% (from 0.086 for the second-best GameCraft to 0.100). In Table 4, our method achieves the best camera accuracy in short-horizon comparisons with camera-controlled methods, improving RPE_{camera} by 36.1% (0.030 vs. 0.083 for the second-best CameraCtrl).

In terms of long-horizon visual quality, our method achieves the highest VBench average score (0.844), outperforming the second-best GameCraft (0.781) by a relative margin of 8.1%. This improvement is attributed to progressive noise scheduling, attention sinks, and longer short-term latents, further analyzed in Section 5.5.

Regarding 3D consistency in Table 3, our method achieves the best performance on all evaluation metrics. While Yume (Mao et al., 2025) achieves comparable PSNR, this is primarily due to blurrier outputs, as indicated by its substantially lower Sharpness score (95 vs. 577 for real videos). Overall, our method demonstrates the strongest 3D consistency across all metrics.

Qualitative results. Figures 4, 5, and 6 present qualitative results. As shown in Figure 4, our method accurately follows user actions while maintaining high visual quality and consistent 3D geometry over long horizons. In contrast, prior methods often struggle with challenging action inputs, exhibit quality degradation over time, and fail to preserve 3D consistency beyond the denoising window.

Figure 5(a) demonstrates that our method effectively handles challenging coupled keyboard and mouse inputs, while Figure 5(b) shows that it can generate plausible long-horizon outputs with newly emerging scenes while remaining faithful to user actions. Figure 6 further highlight that our method preserves consistent 3D geometry when users revisit previously seen locations and viewpoints beyond the denoising window.

Human evaluation. Table 5 further validates the effectiveness of our method across all three evaluation aspects, where our approach consistently outperforms prior methods by a clear margin, achieving relative improvements of 14.0%, 29.8%, and 29.8% over the second-best baselines in action controllability, visual quality, and 3D consistency, respectively.

Table 5 Human evaluation results. We report average user ratings on a scale of 1 (worst) to 5 (best).

Method	Action Controllability↑	Visual Quality↑	3D Consistency↑
Yume (Mao et al., 2025)	2.47	2.83	1.44
Matrix-2.0 (He et al., 2025)	3.78	3.42	2.75
GameCraft (Li et al., 2025a)	2.55	3.34	3.36
WorldCam	**4.31**	**4.44**	**4.36**

5.5 Ablation Study

Table 6 Ablation on the number of long-term memory latents for visual quality and 3D consistency.

#	Avg.↑	Aesth.↑	Subj. Cons.↑	Bg. Cons.↑	Img.↑	Temp.↑	Motion.↑	PSNR↑	LPIPS↓
0	0.840	**0.511**	0.876	0.955	0.751	**0.961**	**0.984**	12.163	0.591
1	0.840	**0.511**	0.879	0.955	**0.752**	**0.961**	**0.984**	12.624	0.573
4	**0.841**	0.510	**0.881**	**0.956**	0.750	**0.961**	**0.984**	**12.950**	**0.554**

Table 7 Ablation on the number of short-term memory latents for visual quality.

#	Avg.↑	Aesth.↑	Subj. Cons.↑	Bg. Cons.↑	Img.↑	Temp.↑	Motion.↑
1	0.749	0.317	0.873	0.947	0.414	0.959	0.983
4	0.836	**0.514**	0.873	0.947	0.737	0.959	0.983
8	**0.840**	0.511	**0.876**	**0.955**	**0.751**	**0.961**	**0.984**

Component analysis. Table 6 demonstrates that increasing the number of long-term memory latents improves PSNR and SSIM, thereby enhancing 3D consistency. Under our progressive noise scheduling, different noise levels require distinct historical contexts, consequently, a larger long-term memory provides richer cues to preserve spatial structures. Notably, visual quality remains stable even in the absence of long-term memory (0.841 Avg. with 4 frames vs. 0.840 without), indicating the model's robustness in generating new scenes without severe error drift, even when spatial constraints are relaxed.

Tables 7 and 8 ablate the number of short-term latents and the inclusion of an attention sink regarding long-horizon visual quality. As discussed in Section 3.5, increasing short-term latents and introducing an attention sink reduce long-horizon error drift, ultimately improving the Avg. score.

Table 8 Ablation on the effect of attention sinks for visual quality.

Method	Avg.↑	Aesth.↑	Subj. Cons.↑	Bg. Cons.↑	Img.↑	Temp.↑	Motion.↑
w/o attention sink	0.840	**0.511**	0.876	0.955	0.751	0.961	**0.984**
with attention sink	**0.841**	0.502	**0.883**	**0.956**	**0.753**	**0.964**	**0.984**

Action-to-camera mapping. Table 9 compares linear and Lie algebra-based approximations for action-to-camera mapping in our model. WorldCam (Linear) achieves performance comparable to GameCraft (Li et al., 2025a), while improving the rotational error. WorldCam (Lie) further improves all metrics and achieves the best overall performance. These results demonstrate that our method provides strong action controllability, and that the Lie algebra-based approximation more accurately models camera motion than the linear approximation.

Table 9 Ablation study on action-to-camera mapping.

Method	$\text{RPE}_{\text{trans}}$ ↓	RPE_{rot} (°) ↓	$\text{RPE}_{\text{camera}}$ ↓
Yume (Mao et al., 2025)	0.111	2.222	0.137
Matrix-Game 2.0 (He et al., 2025)	0.098	1.656	0.119
GameCraft (Li et al., 2025a)	<u>0.086</u>	1.146	<u>0.100</u>
WorldCam (Linear)	0.093	<u>0.962</u>	0.102
WorldCam (Lie)	**0.080**	**0.696**	**0.086**

Memory retrieval. Table 10 compares our memory retrieval strategy with random and temporal retrieval. The former randomly selects long-term latents from the memory pool, while the latter retrieves the most recent latents. Our camera pose-based retrieval significantly outperforms both baselines. For this ablation study, we evaluate on 100 generated videos, each consisting of 200 palindrome frames.

Table 10 Ablation study on long-term memory retrieval.

Method	PSNR↑	LPIPS↓	MEt3R↓
Random	<u>15.76</u>	<u>0.3645</u>	0.041
Temporal	15.18	0.3867	<u>0.040</u>
Ours	**16.42**	**0.3496**	**0.038**

Note that our framework is agnostic to the retrieval strategy. In this work, we adopt a lightweight method based on camera position and orientation. Field-of-View (FOV) overlap is a natural alternative used in prior works (Yu et al., 2025), which typically assumes constant depth to unproject a generated 2D image into 3D and reproject it to a target camera pose to estimate viewpoint overlap. However, this approximation becomes unreliable in complex game environments with frequent occlusions and narrow corridors. Accurate FOV computation would require precise depth maps for generated frames, introducing additional overhead. Therefore, to keep retrieval efficient and compatible with our latent space memory design, we adopt a camera pose-based retrieval strategy.

6 Conclusion

We introduce WorldCam, a foundational interactive gaming world model that enables precise action control, long-horizon generation, and consistent 3D world modeling by using camera pose as a unifying geometric anchor for both immediate action control and long-horizon 3D world consistency. We also present WorldCam-50h, a large-scale dataset of human gameplay that includes open-licensed games, supporting research reproducibility. Extensive experiments demonstrate clear improvements over prior interactive world models and camera-controlled baselines in terms of action accuracy, long-horizon visual quality, and 3D consistency.

"a futuristic, metallic corridor with a high-tech aesthetic. The walls and ceiling feature a grid-like pattern of panels and lights, creating a sterile, industrial atmosphere..."

"an industrial interior with a large, open room featuring metal walls and floors, illuminated by overhead lights. The space is divided into sections by vertical beams, creating a grid-like pattern..."

(a) Action Controllability

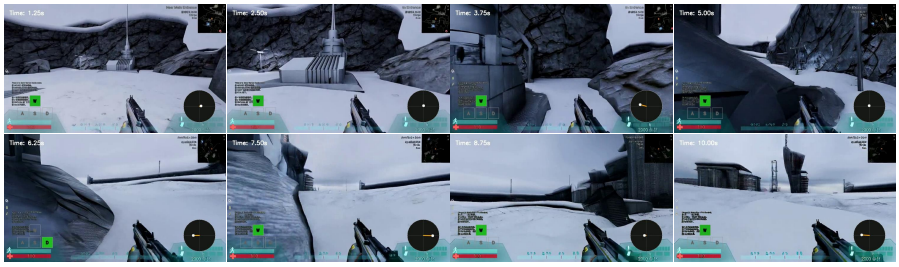

"a snowy, industrial environment with a central building surrounded by rocky terrain. The scene is dominated by shades of gray and white, with a cold, overcast sky suggesting a winter setting.

"The scene is set in a futuristic urban environment with a hexagonal cobblestone street leading to a modern building featuring a green globe logo. The architecture is sleek and contemporary, with clean lines and ..."

"a futuristic, metallic interior with a high-tech aesthetic. The environment features a corridor leading to an observatory, marked by blue and yellow signage. The walls are clad in smooth, reflective panels..."

(b) Long-Horizon Inference

Figure 5 Qualitative Results (Best viewed in color and zoomed in): (a) WorldCam enables fine-grained action control coupled with simultaneous keyboard and mouse inputs. (b) WorldCam generates long-horizon videos exceeding 10 seconds (20 FPS) without error drift. Time (seconds at 20 FPS) is visualized in the top-left of each frame, while keyboard and mouse inputs are shown in the bottom-left and bottom-right, respectively.

"a futuristic, open-air environment with a mix of concrete structures and natural elements like trees and rocks. The architecture features modern, angular designs with large glass panels.."

"The scene is set in a modern, urban environment with a mix of high-rise buildings and low-rise structures. The architecture features clean lines and a predominantly white color scheme with occasional green accents..."

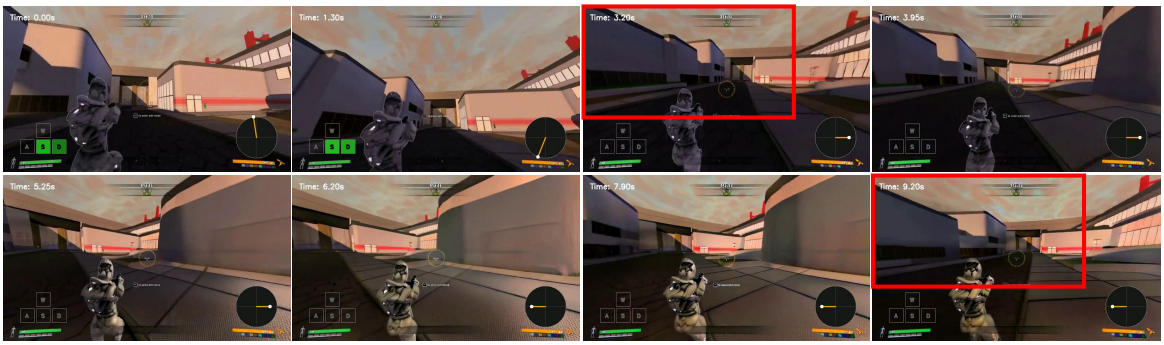

"a futuristic urban environment with modern, angular buildings featuring red accents and large windows. The ground is paved with hexagonal tiles, and the sky is a soft pinkish hue..."

"A modern, industrial outdoor plaza with a mix of dark cobblestone paths and sleek architectural structures. The architecture features a prominent red-framed elevated walkway alongside.."

Figure 6 Qualitative Results (Best viewed in color and zoomed in): Our model preserves the underlying 3D scene consistently over long-horizon beyond the denoising window. Time (seconds at 20 FPS) is visualized in the top-left of each frame, while keyboard and mouse inputs are shown in the bottom-left and bottom-right, respectively. The red box highlights consistent 3D geometry in revisited views.

Appendix Overview

- Section A: Data preprocessing pipeline.
- Section B: Detailed comparison with existing works and evaluation details for each comparison baseline.
- Section C: Detailed explanation of linear and Lie algebra-based action-to-camera approximation.
- Section D: Limitations and runtime comparison.

A Data Preprocessing

We segment the collected videos (with an average duration of 8 minutes) into 1-minute clips (1,800 frames at 30 FPS) and extract global camera extrinsic and intrinsic parameters for each clip using ViPE (Huang et al., 2025). The 1,800-frame clips provide sufficient temporal context for selecting long-term memory frames. Each 1,800-frame clip is further divided into segments of 64 frames, consisting of 32 frames for short-term memory and 32 frames for generation.

To train with long-term memory latents, long-term memory frames are selected from the same 1,800-frame clip based on camera translation and rotation similarity, following Section 3.4. Specifically, using the camera pose of the last frame of the 64-frame segment as a spatial index, we select 4 long-term memory clips (4 frames each) from the remaining 1,736 frames, excluding the current 64-frame segment. To reduce training time, we precompute all VAE latents and store them in a cache.

B Comparison Details

B.1 Comparison with Interactive Gaming World Models

In this section, we highlight the novelty of our approach and provide evaluation details along with comparisons to recent interactive gaming world models, including Matrix-Game 2.0 (He et al., 2025), GameCraft (Li et al., 2025a), and Yume (Mao et al., 2025). As summarized in Table 1 of the main paper, our model satisfies three key requirements for an interactive 3D gaming world model: precise action control, 3D consistency, and long-horizon inference.

Many prior works focus on embedding raw action inputs directly into the model (The Matrix (Feng et al., 2024), Matrix-Game (Zhang et al., 2025), Matrix-Game 2.0 (He et al., 2025), Genie (Bruce et al., 2024), Yan (Ye et al., 2025)). However, these approaches cannot ensure 3D consistency, since raw actions cannot be directly used as spatial indices for memory retrieval. GameCraft (Li et al., 2025a) employs camera poses for action control, but it does not leverage them to maintain long-term 3D consistency. Camera-controlled video generation models (CameraCtrl (He et al., 2024), MotionCtrl (Wang et al., 2024b)) can achieve short-horizon 3D consistency (*e.g.*, within 16 frames), but they are not designed for world modeling, which limits both action control and long-horizon inference.

Consequently, none of the prior works satisfy all three requirements for interactive gaming world models. In contrast, we adopt camera poses as a unified representation for both precise action control and long-horizon 3D consistency.

B.2 Evaluation Details of Comparison Works

Matrix-Game 2.0 (He et al., 2025) is trained on an internal 1,200-hour dataset collected from Unreal Engine and GTA5 environments. Keyboard actions (W, A, S, D) are embedded via cross-attention mechanisms, while mouse movements (dx, dy) are injected into intermediate feature representations. The model adopts a few-step distillation strategy based on a causal architecture to enable low-latency streaming video generation.

For evaluation, we convert our action space $(v_x, v_z, \omega_x, \omega_y)$ into the Matrix-Game 2.0 pre-trained action space (W, A, S, D, dx, dy). Since Matrix-Game does not control the speed of each action, we carefully select thresholds to activate each action. Specifically, linear velocities v_x and v_z are thresholded to remove noise and mapped to binary key inputs W, S (forward/backward) and A, D (left/right), while angular velocities ω_x and ω_y are thresholded and converted into mouse movements dy and dx, respectively.

While Matrix-Game 2.0 directly injects raw action inputs into the generative model, our method maps input actions to camera poses and leverages these poses for precise action control and long-horizon 3D consistency.

GameCraft (Li et al., 2025a) employs a chunk-wise autoregressive video diffusion model with a convolutional camera embedder that encodes Plücker embeddings and injects them into video latents. The model is trained on more than 100 AAA games and synthetic datasets, with camera poses extracted using Monst3r (Zhang et al., 2024). During inference, discrete actions (*e.g.*, forward, backward, left, right, turn left/right, tilt up/down) and their associated speeds are converted into camera poses via linear approximations.

For evaluation, we map our action space $(v_x, v_z, \omega_x, \omega_y)$ into their predefined action space (*e.g.*, mapping ω_y to turn left/right) and convert the corresponding velocities into speeds via scaling. We carefully select a speed scaling factor to match our velocity magnitudes with the speed settings used in the pre-trained GameCraft model.

While GameCraft performs chunk-wise generation, our approach adopts autoregressive progressive noise scheduling for latent-by-latent generation, enabling interactive world modeling. Moreover, GameCraft relies on linear approximations for action-to-camera mappings, whereas our method introduces a Lie algebra-based action-to-camera mapping for fine-grained action control and further leverages predicted camera poses to support long-horizon world consistency.

Yume (Mao et al., 2025) employs a chunk-wise autoregressive masked video diffusion transformer trained on the Sekai-Real-HQ dataset (Li et al., 2025c). Camera motion is discretized into predefined actions (*e.g.*, forward, backward, left, right, turn left/right, still) and incorporated into text captions. To maintain spatial consistency, Yume adopts FramePack (Zhang and Agrawala, 2025), which compresses historical latents via downsampling.

For evaluation, since the model generates chunks of eight latents at a time conditioned on text that encodes actions, we average our actions across the eight latents, apply thresholding, and convert them into Yume's predefined textual action space (*e.g.*, mapping v_z to "*Person moves forward*"). Because Yume encodes actions through text prompts, the resulting action control is relatively coarse. We attempt to match the motion speed by specifying speed-related cues in the text prompts.

While Yume encodes actions through textual prompts and compresses historical latents using FramePack, our model is trained directly on camera poses derived from action inputs and retrieves only relevant historical latents from a long-term memory pool.

C Action-to-Camera Approximation

In this section, we provide a detailed comparison between linear and Lie algebra-based action-to-camera approximations. The key difference lies in whether translation and rotation are updated independently or jointly integrated on the $SE(3)$ manifold. This distinction is particularly important for interactive 3D world modeling, where user actions often induce coupled camera motions.

Let the action at time step i be represented as a 6-DoF velocity vector

$$A_i = \begin{bmatrix} \mathbf{v}_i \\ \boldsymbol{\omega}_i \end{bmatrix} = \left[v_x, v_y, v_z, \omega_x, \omega_y, \omega_z\right]^\top, \tag{8}$$

where $\mathbf{v}_i \in \mathbb{R}^3$ and $\boldsymbol{\omega}_i \in \mathbb{R}^3$ denote the linear and angular velocities, respectively.

Linear approximation. A common approximation used in prior works (Li et al., 2025a,a) is to update translation and rotation separately. Given the camera pose $P_{i-1} = \begin{bmatrix} R_{i-1} & t_{i-1} \\ 0^\top & 1 \end{bmatrix} \in SE(3)$, the linear approximation computes

$$t_i = t_{i-1} + R_{i-1}\mathbf{v}_i, \qquad R_i = R_{i-1}\exp(\widehat{\boldsymbol{\omega}}_i), \tag{9}$$

where $\widehat{\boldsymbol{\omega}}_i \in \mathfrak{so}(3)$ is the skew-symmetric matrix of $\boldsymbol{\omega}_i$. The resulting pose is then

$$P_i = \begin{bmatrix} R_i & t_i \\ 0^\top & 1 \end{bmatrix}. \tag{10}$$

Although simple, this formulation updates translation and rotation independently, failing to capture coupled rigid-body motions. Under screw motion (*e.g.* moving forward while turning right), the camera should follow a curved trajectory

induced by simultaneous translation and rotation, whereas the linear approximation applies these updates separately, leading to accumulated trajectory drift.

Lie algebra-based approximation. In contrast, our formulation represents the action directly as a twist in $\mathfrak{se}(3)$. Specifically, we construct

$$\widehat{A}_i = \begin{bmatrix} \widehat{\boldsymbol{\omega}}_i & \mathbf{v}_i \\ 0^\top & 0 \end{bmatrix} \in \mathfrak{se}(3), \tag{11}$$

and compute the relative camera motion via the matrix exponential

$$\Delta P_i = \exp(\widehat{A}_i) = \begin{bmatrix} \Delta R_i & \Delta t_i \\ 0^\top & 1 \end{bmatrix}. \tag{12}$$

The camera pose is then updated as

$$P_i = P_{i-1} \Delta P_i. \tag{13}$$

Unlike Equation 9, Equation 13 jointly integrates linear and angular velocities on the $SE(3)$ manifold. Consequently, the translational component Δt_i is not treated as a simple vector addition, but as part of a geometrically consistent rigid transformation coupled with rotation. This allows the model to accurately represent screw motions and other complex camera trajectories arising from simultaneous translation and rotation.

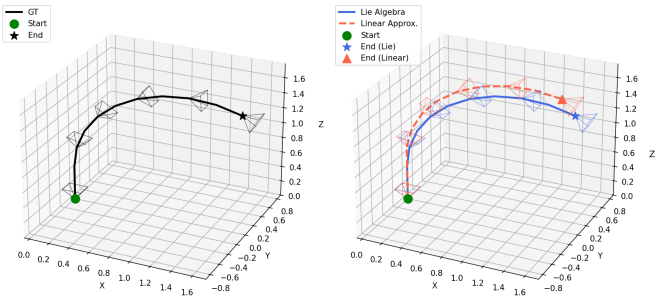

Figure 7 Visualization of the comparison between linear and Lie algebra-based action-to-camera mappings.

Table 11 Comparison between linear and Lie algebra-based action-to-camera approximations. All values are scaled by 10^3 for better visibility.

Method	$\text{RPE}_{\text{trans}} \downarrow$	$\text{ATE}_{\text{avg}} \downarrow$	$\text{ATE}_{\text{final}} \downarrow$
Linear	0.564	29.239	91.149
Lie algebra	**0.001**	**0.005**	**0.015**

Why screw motion matters. This difference is visualized in Figure 7, which considers a simple screw motion example: moving forward while turning right. In this case, the ground-truth camera trajectory follows a smooth curved path. The linear approximation fails to reproduce this motion because it decouples translation from rotation, resulting in inaccurate relative motion and accumulated trajectory drift. In contrast, the Lie algebra-based formulation closely matches the ground-truth trajectory, since it preserves the coupling between translation and rotation at every step.

Table 11 quantitatively confirms this behavior. For evaluation, we generate 50 trajectories of length 200 frames and compute Relative Pose Error ($\text{RPE}_{\text{trans}}$) and Absolute Trajectory Error (ATE). Specifically, ATE_{avg} and $\text{ATE}_{\text{final}}$ measure the drift between the approximated and ground-truth camera trajectories over the full sequence and at the final frame, respectively. Linear approximations exhibit large RPE and accumulated ATE, indicating inaccurate relative camera motion and substantial long-term drift. In contrast, the Lie algebra–based approximation maintains near-zero error across all metrics. These results show that our formulation provides a stable and accurate camera representation, which is crucial for precise action control as well as stable long-term memory retrieval and spatial consistency (Section 3.4).

D Limitation

While our approach improves generation quality and controllability, inference efficiency remains a limitation. This work primarily focuses on camera pose as a unifying geometric representation that jointly enables immediate action control and long-horizon, 3D consistent memory. Improving runtime is therefore orthogonal to our main contribution and could likely be achieved by integrating existing acceleration and distillation techniques. In particular, diffusion distillation methods (Yin et al., 2025; Geng et al., 2025; Song et al., 2023; Yin et al., 2023) could reduce the number of sampling timesteps per latent, potentially collapsing multi-step sampling into few-step generation, consistent with recent progress toward real-time video synthesis.

In Table 12, we compare the inference time complexity of our model with several recent approaches. For a fair comparison, all methods are evaluated at a spatial resolution of 480×832 on a single H100 GPU. Yume (Mao et al., 2025) and GameCraft (Li et al., 2025a) are chunk-wise autoregressive models that generate 32 and 33 frames per chunk, respectively. While this design improves throughput per chunk, it is not well-suited for interactive world models that require low-latency interaction. In particular, chunk-wise generation requires a sequence of actions corresponding to all frames within a chunk before generation, which limits responsiveness in interactive settings.

Table 12 Inference time comparison. All methods are evaluated at a spatial resolution of 480×832 on a single H100 GPU.

Method	Latent-wise AR	Steps	Frames/Chunk	Time/Chunk (s)	Time/Step (s)
Yume (Mao et al., 2025)	×	4	32	**2.91**	<u>0.73</u>
GameCraft (Li et al., 2025a)	×	8	33	109.4	13.68
Matrix-Game 2.0 (He et al., 2025)	✓	3	12	<u>3.09</u>	1.03
WorldCam	✓	8	4	4.17	**0.52**

In contrast, Matrix-Game 2.0 (He et al., 2025) and our method adopt latent-wise autoregressive generation, enabling low-latency interaction. Matrix-Game 2.0 generates three latent tokens every three sampling steps, whereas our method generates a single latent token using eight sampling steps. Notably, Matrix-Game 2.0 employs distillation to reduce the number of sampling steps, whereas our method does not use any distillation technique. Despite this, our method achieves a comparable time per chunk to Matrix-Game 2.0 while achieving the lowest time per step among all compared approaches. The low time per step suggests that further acceleration through distillation could potentially enable real-time interactive generation.

References

Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chattopadhyay, Yongxin Chen, Yin Cui, Yifan Ding, et al. Cosmos world foundation model platform for physical ai. *arXiv preprint arXiv:2501.03575*, 2025.

Mohammad Asim, Christopher Wewer, Thomas Wimmer, Bernt Schiele, and Jan Eric Lenssen. Met3r: Measuring multi-view consistency in generated images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6034–6044, 2025.

Jake Bruce, Michael D Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, et al. Genie: Generative interactive environments. In *Forty-first International Conference on Machine Learning*, 2024.

Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021.

Haoxuan Che, Xuanhua He, Quande Liu, Cheng Jin, and Hao Chen. Gamegen-x: Interactive open-world game video generation. *arXiv preprint arXiv:2411.00769*, 2024.

Boyuan Chen, Diego Martí Monsó, Yilun Du, Max Simchowitz, Russ Tedrake, and Vincent Sitzmann. Diffusion forcing: Next-token prediction meets full-sequence diffusion. *Advances in Neural Information Processing Systems*, 37:24081–24125, 2024.

Guibin Chen, Dixuan Lin, Jiangping Yang, Chunze Lin, Junchen Zhu, Mingyuan Fan, Hao Zhang, Sheng Chen, Zheng Chen, Chengcheng Ma, et al. Skyreels-v2: Infinite-length film generative model. *arXiv preprint arXiv:2504.13074*, 2025a.

Junyi Chen, Haoyi Zhu, Xianglong He, Yifan Wang, Jianjun Zhou, Wenzheng Chang, Yang Zhou, Zizun Li, Zhoujie Fu, Jiangmiao Pang, et al. Deepverse: 4d autoregressive video generation as a world model. *arXiv preprint arXiv:2506.01103*, 2025b.

Taiye Chen, Xun Hu, Zihan Ding, and Chi Jin. Learning world models for interactive video generation. *arXiv preprint arXiv:2505.21996*, 2025c.

Ruili Feng, Han Zhang, Zhantao Yang, Jie Xiao, Zhilei Shu, Zhiheng Liu, Andy Zheng, Yukun Huang, Yu Liu, and Hongyang Zhang. The matrix: Infinite-horizon world generation with real-time moving control. *arXiv preprint arXiv:2412.03568*, 2024.

Shenyuan Gao, Siyuan Zhou, Yilun Du, Jun Zhang, and Chuang Gan. Adaworld: Learning adaptable world models with latent actions. *arXiv preprint arXiv:2503.18938*, 2025.

Zhengyang Geng, Mingyang Deng, Xingjian Bai, J Zico Kolter, and Kaiming He. Mean flows for one-step generative modeling. *arXiv preprint arXiv:2505.13447*, 2025.

Junliang Guo, Yang Ye, Tianyu He, Haoyu Wu, Yushu Jiang, Tim Pearce, and Jiang Bian. Mineworld: a real-time and open-source interactive world model on minecraft. *arXiv preprint arXiv:2504.08388*, 2025.

William H Guss, Brandon Houghton, Nicholay Topin, Phillip Wang, Cayden Codel, Manuela Veloso, and Ruslan Salakhutdinov. Minerl: A large-scale dataset of minecraft demonstrations. *arXiv preprint arXiv:1907.13440*, 2019.

Yoav HaCohen, Nisan Chiprut, Benny Brazowski, Daniel Shalem, Dudu Moshe, Eitan Richardson, Eran Levin, Guy Shiran, Nir Zabari, Ori Gordon, et al. Ltx-video: Realtime video latent diffusion. *arXiv preprint arXiv:2501.00103*, 2024.

Brian C Hall. Lie groups, lie algebras, and representations. In *Quantum Theory for Mathematicians*, pages 333–366. Springer, 2013.

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*, 2024.

Xianglong He, Chunli Peng, Zexiang Liu, Boyang Wang, Yifan Zhang, Qi Cui, Fei Kang, Biao Jiang, Mengyin An, Yangyang Ren, et al. Matrix-game 2.0: An open-source, real-time, and streaming interactive world model. *arXiv preprint arXiv:2508.13009*, 2025.

Jiahui Huang, Qunjie Zhou, Hesam Rabeti, Aleksandr Korovko, Huan Ling, Xuanchi Ren, Tianchang Shen, Jun Gao, Dmitry Slepichev, Chen-Hsuan Lin, et al. Vipe: Video pose engine for 3d geometric perception. *arXiv preprint arXiv:2508.10934*, 2025.

Ziqi Huang, Fan Zhang, Xiaojie Xu, Yinan He, Jiashuo Yu, Ziyue Dong, Qianli Ma, Nattapol Chanpaisit, Chenyang Si, Yuming Jiang, et al. Vbench++: Comprehensive and versatile benchmark suite for video generative models. *arXiv preprint arXiv:2411.13503*, 2024.

Jiaqi Li, Junshu Tang, Zhiyong Xu, Longhuang Wu, Yuan Zhou, Shuai Shao, Tianbao Yu, Zhiguo Cao, and Qinglin Lu. Hunyuan-gamecraft: High-dynamic interactive game video generation with hybrid history condition. *arXiv preprint arXiv:2506.17201*, 2025a.

Runjia Li, Philip Torr, Andrea Vedaldi, and Tomas Jakab. Vmem: Consistent interactive video scene generation with surfel-indexed view memory. *arXiv preprint arXiv:2506.18903*, 2025b.

Zhen Li, Chuanhao Li, Xiaofeng Mao, Shaoheng Lin, Ming Li, Shitian Zhao, Zhaopan Xu, Xinyue Li, Yukang Feng, Jianwen Sun, et al. Sekai: A video dataset towards world exploration. *arXiv preprint arXiv:2506.15675*, 2025c.

Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

Xiaofeng Mao, Shaoheng Lin, Zhen Li, Chuanhao Li, Wenshuo Peng, Tong He, Jiangmiao Pang, Mingmin Chi, Yu Qiao, and Kaipeng Zhang. Yume: An interactive world generation model. *arXiv preprint arXiv:2507.17744*, 2025.

Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.

Said Pertuz, Domenec Puig, and Miguel Angel Garcia. Analysis of focus measure operators for shape-from-focus. *Pattern Recognition*, 46(5):1415–1432, 2013.

Ryan Po, Yotam Nitzan, Richard Zhang, Berlin Chen, Tri Dao, Eli Shechtman, Gordon Wetzstein, and Xun Huang. Long-context state-space video world models. *arXiv preprint arXiv:2505.20171*, 2025.

Vincent Sitzmann, Semon Rezchikov, Bill Freeman, Josh Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. *Advances in Neural Information Processing Systems*, 34:19313–19325, 2021.

Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. 2023.

Junshu Tang, Jiacheng Liu, Jiaqi Li, Longhuang Wu, Haoyu Yang, Penghao Zhao, Siruis Gong, Xiang Yuan, Shuai Shao, and Qinglin Lu. Hunyuan-gamecraft-2: Instruction-following interactive game world model. *arXiv preprint arXiv:2511.23429*, 2025.

Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on pattern analysis and machine intelligence*, 13(4):376–380, 2002.

Dani Valevski, Yaniv Leviathan, Moab Arar, and Shlomi Fruchter. Diffusion models are real-time game engines. *arXiv preprint arXiv:2408.14837*, 2024.

Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025.

Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 20697–20709, 2024a.

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, 2024b.

Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*, 2023.

Zeqi Xiao, Yushi Lan, Yifan Zhou, Wenqi Ouyang, Shuai Yang, Yanhong Zeng, and Xingang Pan. Worldmem: Long-term consistent world simulation with memory. *arXiv preprint arXiv:2504.12369*, 2025.

Desai Xie, Zhan Xu, Yicong Hong, Hao Tan, Difan Liu, Feng Liu, Arie Kaufman, and Yang Zhou. Progressive autoregressive video diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6322–6332, 2025.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.

Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024.

Deheng Ye, Fangyun Zhou, Jiacheng Lv, Jianqi Ma, Jun Zhang, Junyan Lv, Junyou Li, Minwen Deng, Mingyu Yang, Qiang Fu, et al. Yan: Foundational interactive video generation. *arXiv preprint arXiv:2508.08601*, 2025.

Tianwei Yin, Michael Gharbi, Richard Zhang, Eli Shechtman, Frédo Durand, William T Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. 2024 ieee. In *CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6613–6623, 2023.

Tianwei Yin, Qiang Zhang, Richard Zhang, William T Freeman, Fredo Durand, Eli Shechtman, and Xun Huang. From slow bidirectional to fast autoregressive video diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 22963–22974, 2025.

Jiwen Yu, Jianhong Bai, Yiran Qin, Quande Liu, Xintao Wang, Pengfei Wan, Di Zhang, and Xihui Liu. Context as memory: Scene-consistent interactive long video generation with memory retrieval. In *Proceedings of the SIGGRAPH Asia 2025 Conference Papers*, pages 1–11, 2025.

Junyi Zhang, Charles Herrmann, Junhwa Hur, Varun Jampani, Trevor Darrell, Forrester Cole, Deqing Sun, and Ming-Hsuan Yang. Monst3r: A simple approach for estimating geometry in the presence of motion. *arXiv preprint arXiv:2410.03825*, 2024.

Lvmin Zhang and Maneesh Agrawala. Packing input frame context in next-frame prediction models for video generation. *arXiv preprint arXiv:2504.12626*, 2025.

Yifan Zhang, Chunli Peng, Boyang Wang, Puyi Wang, Qingcheng Zhu, Fei Kang, Biao Jiang, Zedong Gao, Eric Li, Yang Liu, et al. Matrix-game: Interactive world foundation model. *arXiv preprint arXiv:2506.18701*, 2025.