

Plenoptic Video Generation

Xiao Fu^{1,2}, Shitao Tang¹, Min Shi^{1,3}, Xian Liu¹, Jinwei Gu¹, Ming-Yu Liu¹, Dahua Lin², Chen-Hsuan Lin¹
¹NVIDIA, ²The Chinese University of Hong Kong, ³Georgia Institute of Technology




Figure 1. We present *PlenopticDreamer*, a generative framework that re-renders input video under novel camera trajectories while preserving long-term spatio-temporal memory in hallucinated regions across overlapping views, thereby producing coherent plenoptic functions (see robot’s right side, highlighted in red dashed boxes across three trajectories). Please refer to our [website](#) for more results.

Abstract

Camera-controlled generative video re-rendering methods, such as *ReCamMaster*, have achieved remarkable progress. However, despite their success in single-view setting, these works often struggle to maintain consistency across multi-view scenarios. Ensuring spatio-temporal coherence in hallucinated regions remains challenging due to the inherent stochasticity of generative models. To address it, we introduce *PlenopticDreamer*, a framework that synchronizes generative hallucinations to maintain spatio-temporal memory. The core idea is to train a multi-in-single-out video-conditioned model in an autoregressive manner, aided by a camera-guided video retrieval strategy that adaptively selects salient videos from previous generations as conditional inputs. In addition, our training incorporates progressive context-scaling to improve convergence, self-conditioning to enhance robustness against long-range visual degradation caused by error accumulation, and a long-video conditioning mechanism to support extended video generation. Extensive experiments on the

Basic and Agibot benchmarks demonstrate that PlenopticDreamer achieves state-of-the-art video re-rendering, delivering superior view synchronization, high-fidelity visuals, accurate camera control, and diverse view transformations (e.g., third-person \rightarrow third-person, and head-view \rightarrow gripper-view in robotic manipulation). Project page: <https://research.nvidia.com/labs/dir/plenopticdreamer/>.

1. Introduction

Video generation [3, 9, 20, 21, 23, 26, 34, 45, 50, 51, 53] has become increasingly prevalent in content creation and social media. As video frames result from camera projections of scene radiance, they can be interpreted as discrete samples of the underlying plenoptic function [8, 39]. Consequently, effective control of camera motion [4, 24, 46, 57] is essential for shaping the captured light field, emphasizing visual focus, and guiding the viewer’s attention.

Recently, camera-controlled generative video re-

rendering, which aims to synthesize novel videos along arbitrary camera trajectories while preserving the original content, has attracted significant attention, supporting applications such as immersive content creation and embodied AI. Representative methods, including ReCamMaster [6] and TrajectoryCrafter [65], achieve promising results on their curated real-world or synthetic datasets. However, these methods primarily succeed in the *single-view* setting and struggle in *multi-view* scenarios, which are essential for reconstructing a holistic representation of the scene. Specifically, they fail to maintain consistent spatio-temporal hallucinations in regions unseen from the source view. The inherent stochasticity of diffusion models, combined with their limited long-range spatial memory, leads to geometric misalignment and view desynchronization across different camera-conditioned generations.

To address it, we present *PlenopticDreamer*, a camera-controlled generative video re-rendering framework that explicitly enforces spatio-temporal memory for consistent scene generation. Unlike prior single-shot methods that generate each view independently, *PlenopticDreamer* adopts an *autoregressive*, multi-in–single-out formulation. At each step, it retrieves a set of previously generated video–camera pairs from a memory bank and conditions the next generation on these retrieved contexts. This design enables synchronized hallucinations across time and view-points while preserving scene geometry and motion dynamics. Video context retrieval is guided by a 3D field-of-view (FOV) mechanism that evaluates spatial co-visibility to select the most relevant past video segments.

In addition, *PlenopticDreamer* introduces two training strategies that substantially improve robustness and convergence. First, progressive context-scaling stabilizes optimization by gradually increasing the number of conditioning videos during training, enabling the model to learn context-aware reasoning across short to long temporal horizons. Second, self-conditioned training mitigates error accumulation in autoregressive generation by fine-tuning the model on its own synthesized outputs. Together, these strategies facilitate stable video synthesis while preserving spatial alignment and temporal consistency. We further propose a long-video conditioning mechanism to extend the model’s capability to render longer video sequences.

We evaluate our method on two benchmarks: (1) a *Basic benchmark* covering diverse in-the-wild scenes, and (2) an *Agibot benchmark* [10] focusing on robotic manipulation. Experimental results show that *PlenopticDreamer* achieves state-of-the-art performance in view synchronization while maintaining accurate camera control and high-fidelity visual quality. In summary, our contributions are:

1. We present *PlenopticDreamer*, the first camera-controlled generative video re-rendering framework with long-term spatio-temporal memory.

2. We propose an autoregressive architecture with a 3D FOV–based video retrieval mechanism for scalable, coherent multi-camera generation. We incorporate progressive context-scaling and self-conditioned training strategies to enhance stability and long-term consistency.
3. Extensive experiments show that our method achieves state-of-the-art performance in video re-rendering, including view synchronization, camera accuracy, and visual fidelity. It supports diverse camera transformations (*e.g.*, third-person \rightarrow third-person, head-view \rightarrow gripper-view) and enables long video generation.

2. Related Work

Camera-Controlled Video Generation. Effective control of camera motion has received significant attention in video generation. Existing approaches can be broadly categorized into three directions: (1) *single-view generation*: works rely on explicit 6DoF camera poses [46, 57] or pixel-wise Plücker raymaps [4, 5, 24, 62, 67] to guide text- or image-to-video synthesis, achieving view control. Methods [27, 28, 40, 63] explore training-free strategies for camera manipulation, and a few [19, 56] extend these mechanisms to control object motion via camera trajectories. (2) *multi-view video generation*: methods aim to maintain cross-view consistency, ranging from object-level synthesis [36, 61] to scene-level reconstruction [7, 35]. (3) *video-to-video re-rendering*: some approaches [6, 52] perform implicit re-rendering with minimal 3D supervision, whereas others [22, 46, 60, 61, 65] project video context into 3D representations to synthesize novel views. Despite these advances, none integrate memory mechanisms to maintain long-term spatio-temporal coherence across multiple views. In contrast, our *PlenopticDreamer* introduces the first memory-based framework for generative video-to-video re-rendering, achieving coherent multiview synthesis.

Memory Mechanism for Video Generation. Building long-term memory is essential for coherent video generation [33, 49]. Existing approaches can be broadly categorized into four types: (1) *frame-level memory*: methods such as [13, 14, 47, 59, 64] store key historical frames and retrieve the top- k relevant ones via camera-pose similarity for conditioning; (2) *latent-level memory*: approaches including [11, 37, 43, 44, 66] maintain hierarchical memory capturing long-term coarse tokens and short-term fine-grained tokens, adaptively retrieving salient features during inference; (3) *3D-level memory*: works like VMem [38] and SPMem [58] reconstruct 3D structures (*e.g.*, surfels or point clouds) to store video context and render geometry-aware representations for novel-view synthesis; (4) *network-level memory*: TTT-Video [16] leverages Test-Time Training (TTT) layers to record input tokens and update model weights. In contrast, our *PlenopticDreamer* introduces an explicit *video-based retrieval* mechanism conditioning gen-

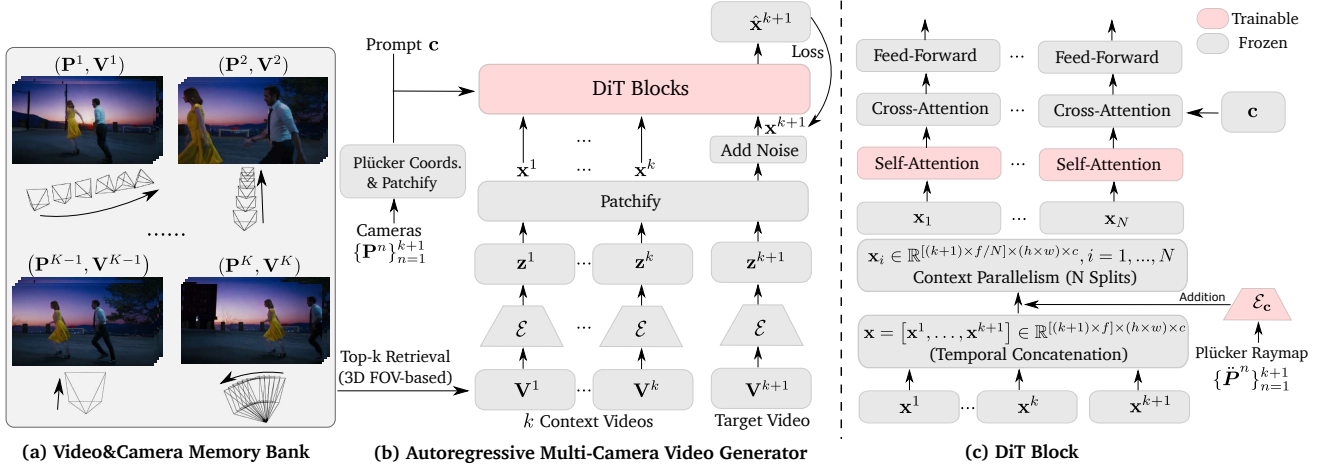


Figure 2. **PlenopticDreamer Framework.** Its core is an autoregressive multi-camera video generator that retrieves k video-camera pairs $\{(\mathbf{P}^n, \mathbf{V}^n)\}_{n=1}^k$ from the memory bank using a 3D FOV-based retrieval strategy. Conditioned on these retrieved pairs and the target camera \mathbf{P}^{k+1} , the model performs noisy scheduling and learnable reconstruction to generate the target video \mathbf{V}^{k+1} . To enable long video generation, a portion of the preceding frames in \mathbf{V}^{k+1} is preserved as clean inputs at a certain ratio during training. Within each DiT block, temporal concatenation is applied to form video tokens \mathbf{x} as in-context condition.

eration on camera-guided selection of past video segments.

3. Method

Our goal is to enable generative video-to-video re-rendering with spatio-temporal memory, interpretable as generating the space-time dependent plenoptic function of a scene. We first introduce the preliminaries and task formulation in Sec. 3.1, followed by our autoregressive modeling paradigm and video retrieval mechanism in Sec. 3.2. Finally, we describe the enhanced training strategies in Sec. 3.3. The overall framework is illustrated in Fig. 2.

3.1. Preliminary and Problem Definition

Flow-based Video Diffusion Transformer (DiT). We conduct experiments using a video diffusion transformer model under the flow-matching paradigm [18, 41]. Given a data sample $\mathbf{x}_0 \sim p(\mathbf{x})$, a noise sampler $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, and a continuous time variable $t \in [0, 1]$, the forward process linearly interpolates between data and noise distribution, *i.e.*,

$$\mathbf{x}_t = (1-t)\mathbf{x}_0 + t\epsilon, \mathbf{v}_t = \epsilon - \mathbf{x}_0 \quad (1)$$

where \mathbf{v}_t is the GT velocity field. The denoising process is solved by an ordinary differential equation (ODE):

$$d\mathbf{x}_t = \mathbf{v}_\Theta(\mathbf{x}_t, t, \mathbf{c}) dt \quad (2)$$

where $\mathbf{v}_\Theta(\cdot)$ represents the predicted velocity function parameterized by a transformer-based network [45], and \mathbf{c} is the conditional signal (*e.g.*, video context). The model is optimized using the following flow-matching objective:

$$\mathcal{L}(\Theta) = \mathbb{E}_{\mathbf{x}, \epsilon, \mathbf{c}, t} \|\mathbf{v}_\Theta(\mathbf{x}_t, t, \mathbf{c}) - \mathbf{v}_t\|^2 \quad (3)$$

During training, the timestep t can be biased toward higher noise levels to encourage robust reconstruction under degraded spatio-temporal correlations.

Task Formulation and Notations. Given a source video $\mathbf{V}_s \in \mathbb{R}^{F \times C \times H \times W}$ and a set of N target camera trajectories $\{\mathbf{P}_t^n\}_{n=1}^N$, each \mathbf{P}_t specified by extrinsic parameters $\mathbf{C}_t = [\mathbf{R}_t; \mathbf{T}_t] \in \mathbb{R}^{F \times 3 \times 4}$ and intrinsics $\mathbf{K}_t \in \mathbb{R}^{3 \times 3}$, our objective is to synthesize N target videos $\{\mathbf{V}_t^n\}_{n=1}^N$, with each $\mathbf{V}_t^n \in \mathbb{R}^{F \times C \times H \times W}$ sharing the same context as the input video while corresponding to a distinct virtual camera trajectory. The generated videos are required to maintain the source content fidelity and exhibit synchronized spatio-temporal consistency across viewpoints, particularly in hallucinated regions. We employ a standard pinhole camera model (zero horizontal and vertical skew) for generation. The overall generative process $f(\cdot)$ is formulated as

$$f(\cdot) : \mathbf{c}, \mathbf{V}_s, \mathbf{P}_s, \{\mathbf{P}_t^n\}_{n=1}^N \rightarrow \{\mathbf{V}_t^n\}_{n=1}^N \quad (4)$$

where \mathbf{c} denotes the video caption and \mathbf{P}_s is the camera trajectory of the source video. Additionally, a variational autoencoder with encoder $\mathcal{E}(\cdot)$ and decoder $\mathcal{D}(\cdot)$ is employed to map the videos between pixel-space and latent-space ($\mathbf{V}_s \leftrightarrow \mathbf{z}_s, \{\mathbf{V}_t^n\}_{n=1}^N \leftrightarrow \{\mathbf{z}_t^n\}_{n=1}^N$, where $\mathbf{z} \in \mathbb{R}^{f \times h \times w \times c}$).

3.2. Injecting Conditions into Video DiT

To effectively guide the video model with target conditions, including the source video and target camera trajectories, we adopt an in-context conditioning strategy. [6, 25, 32, 64].

Native Solution. A straightforward approach is to enlarge the context window by extending the number of input

videos from 1 to N , following ReCamMaster [6]:

$$\begin{cases} \mathbf{x}_s = \text{patchify}(\mathbf{z}_s), \mathbf{x}_t^n = \text{patchify}(\mathbf{z}_t^n), n = 1, \dots, N \\ \mathbf{x} = [\mathbf{x}_s, \mathbf{x}_1, \dots, \mathbf{x}_N]_{\text{frame-dim}} \in \mathbb{R}^{[(N+1) \times f] \times (h \times w) \times c}, \end{cases} \quad (5)$$

where \mathbf{x} denotes the input video tokens fed into the DiT block. While this strategy can be effective for small N (e.g., 2 or 3) under low-resolution settings ($\leq 480\text{p}$), it rapidly becomes computationally prohibitive and prone to out-of-memory (OOM) failures as N or video resolution increases.

Autoregressive Generation Paradigm. Inspired by the effectiveness of the autoregressive paradigm for long-context modeling [1, 12, 30, 55, 64], we reformulate video generation as a sequential process instead of performing single-shot inference. Specifically, we generate one video at a time and produce all videos in a sequential manner. Accordingly, we rewrite Eq. 4 as:

$$f(\cdot) : \mathbf{c}, \{(\mathbf{P}^n, \mathbf{V}^n)\}_{n=1}^k, \mathbf{P}^{k+1} \rightarrow \mathbf{V}^{k+1}, k = 1, \dots, N-1 \quad (6)$$

where $\{(\mathbf{P}^n, \mathbf{V}^n)\}_{n=1}^k$ denotes previously generated videos and their cameras and $(\mathbf{P}_s, \mathbf{V}_s)$ is regarded as $(\mathbf{P}^1, \mathbf{V}^1)$, and k is the model context size. We adopt temporal concatenation strategy to form conditional video tokens:

$$\mathbf{x} = [\mathbf{x}_1, \dots, \mathbf{x}_{k+1}]_{\text{frame-dim}} \in \mathbb{R}^{[(k+1) \times f] \times (h \times w) \times c} \quad (7)$$

Camera Conditioning. To encode camera information, we represent it using Plücker raymaps [48], mapping pixels to 6D ray representations: $\mathbf{P}^n = (\mathbf{C}^n \in \mathbb{R}^{f \times 3 \times 4}, \mathbf{K} \in \mathbb{R}^{3 \times 3}) \rightarrow \tilde{\mathbf{P}}^n \in \mathbb{R}^{f \times H \times W \times 6}, n = 1, \dots, k+1$. These raymaps are then temporally concatenated and patchified. A camera projection layer $\mathcal{E}_{\text{cam}}(\cdot)$ is introduced to align the raymap dimensionality with that of the video latents. The resulting raymap tokens are channel-wise added to the video tokens before self-attention layer, enabling DiT to integrate camera pose information.

3D FOV-based Video Retrieval. A key challenge in this autoregressive framework lies in selecting the most salient k videos from the previous video pool as conditioning inputs for the next generation step. Given that each video corresponds to a distinct camera trajectory, we adopt a 3D Field-of-View (FOV) retrieval mechanism to identify the most relevant candidates. Specifically, we compute video-level similarity via spatial co-visibility across all frames, and select the top- k context videos. When the number of context videos is less than k , we replicate input video-camera pair $(\mathbf{P}^1, \mathbf{V}^1)$ to match the required context length. Furthermore, when the number of retrieved videos l exceeds the model’s context capacity k , a divide-and-conquer inference strategy is employed to cover as diverse viewpoints as possible and minimize viewpoint overlap, as described in Algorithm 2. Here the trajectory fusion process results in a merged trajectory roughly spanning the FOV of all inputs.

Algorithm 1 Video Retrieval Algorithm

Input:

- Memory bank of K videos $\{(\mathbf{V}^n, \mathbf{P}^n)\}_{n=1}^K$
- Target camera trajectory \mathbf{P}^{K+1}
- Maximum retrieved video number k
- Near/Far plane distances D_n, D_f
- Monte Carlo sampling points P

Output: Top- k retrieved videos

- 1: Initialize similarity set $S \leftarrow \emptyset$
 - 2: **for** $n = 1$ to K **do**
 - 3: Initialize similarity $S_n \leftarrow 0$
 - 4: **for** $f = 1$ to F **do**
 - 5: Construct f -th camera frustum of \mathbf{P}^n and \mathbf{P}^{K+1}
 - 6: Perform Monte Carlo sampling within near/far planes of each frustum
 - 7: Count visible points P_n, P_{K+1} in other’s frustum
 - 8: Update similarity: $S_n \leftarrow S_n + \frac{P_n + P_{K+1}}{2P \times F}$
 - 9: **end for**
 - 10: Append S_n to S
 - 11: **end for**
 - 12: Select indices of the top- k values in S
 - 13: **return** retrieved k videos
-

Algorithm 2 Divide-and-Conquer Inference Algorithm

Input:

- Top- l retrieved videos $\{(\mathbf{V}^n, \mathbf{P}^n)\}_{n=1}^l$ (sorted by ascending camera similarity)
- Model context video size k
- Target camera trajectory \mathbf{P}^{L+1}

Output: Target video \mathbf{V}^{L+1} .

- 1: **while** $l > k$ **do**
 - 2: Select the first $m = \min(l - k, k)$ videos to form context set V
 - 3: **while** $m < k$ **do**
 - 4: Append $(\mathbf{P}_s, \mathbf{V}_s)$ to V
 - 5: $m \leftarrow m + 1$
 - 6: **end while**
 - 7: Merge trajectories in V to form $\mathbf{P}_{\text{merge}}$
 - 8: Infer merged video $\mathbf{V}_{\text{merge}}$ using $(V, \mathbf{P}_{\text{merge}})$
 - 9: Replace the first m elements with $(\mathbf{V}_{\text{merge}}, \mathbf{P}_{\text{merge}})$
 - 10: $l \leftarrow l - m + 1$
 - 11: **end while**
 - 12: Perform final inference to obtain \mathbf{V}^{L+1}
 - 13: **return** target video \mathbf{V}^{L+1}
-

Autoregressive Long Video Generation. For input video exceeding the model’s temporal window, we partition them into overlapping sub-chunks $\{\mathbf{V}_s^m\}_{m=1}^M$, where consecutive chunks share a set of frames from the latter portion of the previous chunk to preserve temporal continuity. Unlike the formulation in Eq. 6, where \mathbf{V}^{k+1} is generated from pure

noise, we incorporate the overlapping frames as additional conditioning:

$$f(\cdot) : \mathbf{c}, \{(\mathbf{P}^{m,n}, \mathbf{V}^{m,n})\}_{n=1}^k, \mathbf{P}^{m,k+1}, \tilde{\mathbf{V}}^{m,k+1} \rightarrow \mathbf{V}^{m,k+1} \quad (8)$$

where $m = 1, \dots, M, k = 1, \dots, N - 1$. Here $\tilde{\mathbf{V}}^{m,k+1} \in \mathbb{R}^{\tilde{F} \times C \times H \times W}$ contains the \tilde{F} overlapping frames, and m indexes the m -th sub-chunk \mathbf{V}_s^m from the source video. During inference, we sequentially generate chunked video segments, as illustrated in Fig. 3, and concatenate them to form the final output.

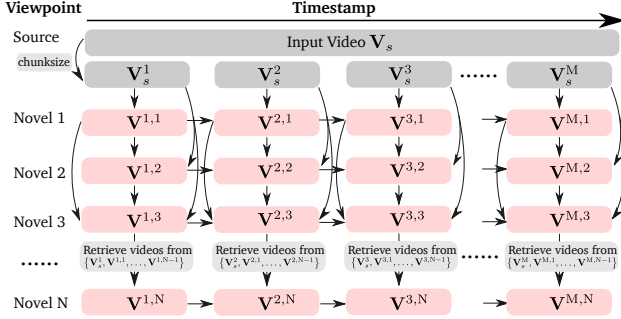


Figure 3. Flowchart of Long Video Re-rendering.

3.3. Training Strategy

Progressive Training. The training objective corresponding to Eq. 6 is defined as:

$$\mathcal{L}(\Theta) = \mathbb{E}_{\epsilon, \mathbf{c}, \mathbf{P}, \mathbf{V}, t} \|\mathbf{v}_{\Theta}(\{(\mathbf{P}^n, \mathbf{V}^n)\}_{n=1}^{k+1}, t, \mathbf{c}) - \mathbf{v}_t\|^2 \quad (9)$$

We also incorporate the extended prediction $\tilde{\mathbf{V}}^{k+1}$ from Eq. 8 into the loss at a certain ratio. In our empirical experiments, we observe that directly training the model with a large context size often leads to unstable convergence. To address this, we adopt a progressive training strategy: the model is first trained with a small context size (e.g., 1) and gradually scaled up as training stabilizes, until reaching the target context size k . This progressive scheme significantly improves convergence stability and accelerates training in later stages with larger contexts.

Self-conditioned Training. When the total generation length N becomes large, multiple inference steps are required, where previously generated videos are repeatedly used as conditioning inputs. This recursive dependency can lead to error accumulation due to propagation of imperfect generations. To alleviate this issue, we adopt a self-conditioned training strategy. Specifically, in the first training stage of Eq. 9, all conditioning videos are ground-truth samples. After convergence, the model is used to generate synthetic outputs from training-set input, which then replace the ground-truth conditions in the second training round. This iterative refinement improves model robustness to imperfect inputs during long-range inference.

4. Experiment

4.1. Experiment Setting

Implementation Details. We adopt Cosmos-Predict2.5-2B [3] as the backbone. The generated videos have a resolution of 432×768 with 93 frames. We employ context parallelism [2, 3] to alleviate memory overhead and set the parallelism size to 8. Finetuning is conducted on 32 NVIDIA H100 GPUs with batch size 1 and a learning rate $2e-5$. During finetuning, only the self-attention layers and camera encoder are updated, while all other parameters remain frozen. We post-train two model variants in Sec. 4.2 and Sec. 4.3.

Evaluation Metrics. We evaluate models from three aspects: 1) *Visual Quality*: PSNR and FVD measure pixel- and frame-level fidelity, respectively. 2) *Camera Accuracy*: TransErr and RotErr [24] quantify translation and rotation errors. Dynamic poses are evaluated with ViPE [29], while static novel views (e.g., azimuth/elevation shifts) assessed with VGGT [54] for relative pose estimation. 3) *Video Synchronization*: RoMa [17] computes the number of matched pixels above a confidence threshold, denoted as Mat. Pix.

Baselines. We compare the proposed PlenopticDreamer with state-of-the-art camera-controlled generative video re-rendering methods: ReCamMaster [6], TrajectoryCrafter [65], and Trajectory-Attention [60]. All baselines are used with their best-performing settings from the official open-sourced models. For a fairer comparison, we also retrain ReCamMaster on Cosmos-Predict2.5 with Plücker raymaps on the same datasets, denoted as ReCamMaster*.

4.2. Experiment on Basic Benchmark

Experiment Details.

- **Functionality:** The model performs third-view to third-view transformations, such as left/right rotations, azimuth and elevation shifts, distance variations, and dynamic focal length changes.
- **Training Dataset:** We use MultiCamVideo [6] and SynCamVideo [7], large-scale synthetic datasets comprising approximately 136K and 34K episodes, respectively, depicting human motion captured under dynamic and static camera trajectories across 40 synthetic 3D environments.
- **Training Details:** The model context size k is set to 4. In the first stage, it is progressively trained for 10K, 4K, 1K, and 1K steps with context sizes 1–4, respectively. In the second stage, the model generates synthetic data from 1,000 scenes and is further trained for 2K steps.
- **Benchmark:** We construct a Basic benchmark of 100 in-the-wild videos and 12 sequential camera trajectories.

Qualitative and Quantitative Comparison. As shown in Fig. 4 and Tab. 1, PlenopticDreamer achieves superior view synchronization with high-fidelity visuals compared to all baselines (see the painting and electrical outlet on the wall in the first example, the traffic light in the sec-

Table 1. **Quantitative Comparison on the Basic Benchmark.** Ours consistently outperforms all baselines in view synchronization across all shots while maintaining high-fidelity visual quality and accurate camera accuracy. ReCamMaster* denotes a retrained version on Cosmos-Predict2.5 with Plücker raymaps, using the same combined datasets (MultiCamVideo + SyncCamVideo) for a fair comparison.

Model	Visual Quality	Camera Accuracy		View Synchronization (Mat. Pix.(K) \uparrow)			
	FVD \downarrow	TransErr \downarrow	RotErr (rad) \downarrow	3 Shots	6 Shots	9 Shots	12 Shots
Trajectory-Attention [60]	734.1	0.77	0.26	22.7	26.9	28.8	29.1
TrajectoryCrafter [65]	665.9	0.65	0.27	31.2	29.3	35.3	36.2
ReCamMaster [6]	731.6	0.72	0.23	32.1	29.0	30.9	27.6
ReCamMaster* [6]	675.4	0.52	<u>0.22</u>	24.6	20.2	29.7	31.2
PlenopticDreamer (Ours)	425.8	<u>0.54</u>	0.21	41.4	40.8	45.4	41.2

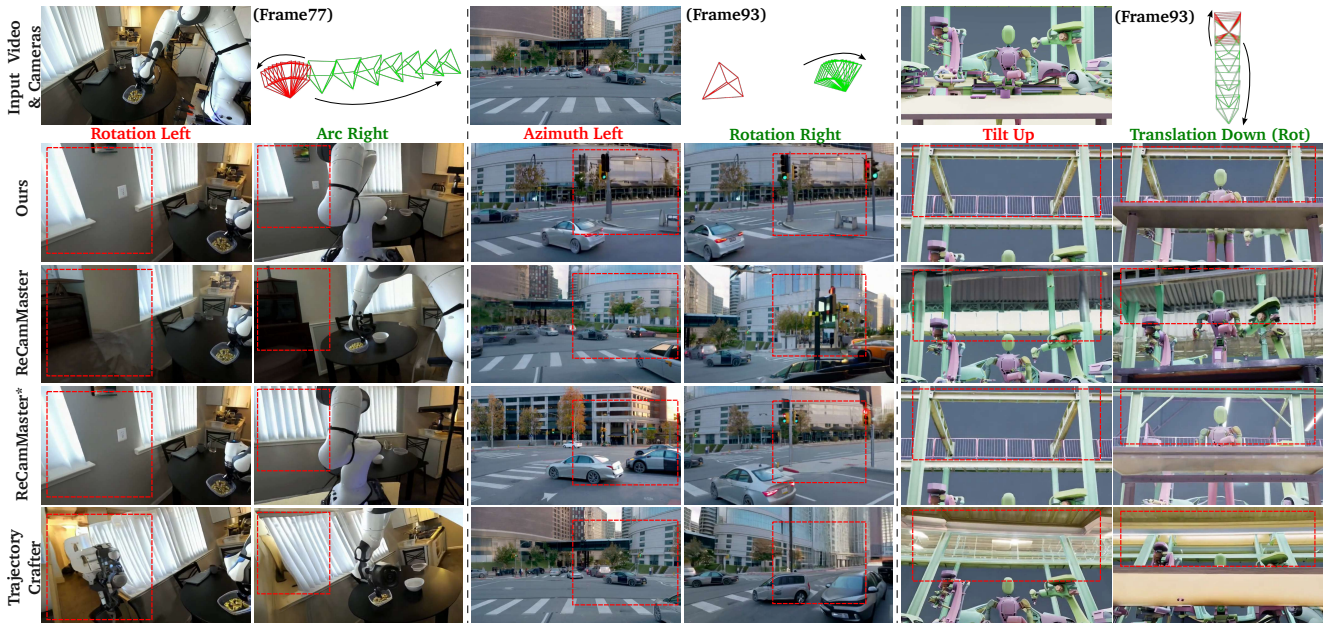


Figure 4. **Qualitative Comparison on the Basic Benchmark.** PlenopticDreamer generates high-fidelity visuals with consistent hallucinations from different camera trajectories. In contrast, ReCamMaster and TrajectoryCrafter fail to preserve spatio-temporal consistency while maintaining visual quality, especially under large-angle viewpoint changes, such as leftward azimuth shifts.

ond, and the eave above the robot in the third). TrajectoryCrafter and Trajectory-Attention leverage 3D point tracking to extract dynamic cues from the source video and feed them as conditional inputs to the generator. However, without updating the 3D memory using newly rendered content, they fail to maintain consistent cross-view synthesis. Moreover, the off-the-shelf checkpoints of these baselines exhibit low camera accuracy, especially in translation, due to poor performance on static novel-view synthesis under large-angle viewpoint changes (*e.g.*, azimuth and elevation shifts). When retrained on the same datasets, ReCamMaster* achieves comparable camera accuracy. Notably, the original ReCamMaster does not employ Plücker raymaps; we integrate them to ensure a fair comparison.

4.3. Experiment on Agibot Benchmark

Experiment Details.

- **Functionality:** The model supports head-view to gripper-view transformations in robotic manipulation.
- **Training Dataset:** We use Agibot [10], a large-scale robotic dataset with about 1M episodes. We sample 145,820 episodes, each containing three synchronized video views (one head-view and two gripper-views) with precise camera pose annotations.
- **Training Details:** The model context size k is set to 2, and it is trained for 15K steps with merely the first stage, requiring ~ 5 days.
- **Benchmark:** We build an Agibot benchmark using 200 test videos, covering head-to-hand and hand-to-hand camera transformations.

Qualitative and Quantitative Comparison. As illustrated in Fig. 5 and Tab. 2, PlenopticDreamer can perform head-view \rightarrow gripper-view and gripper-view \rightarrow gripper-view trans-

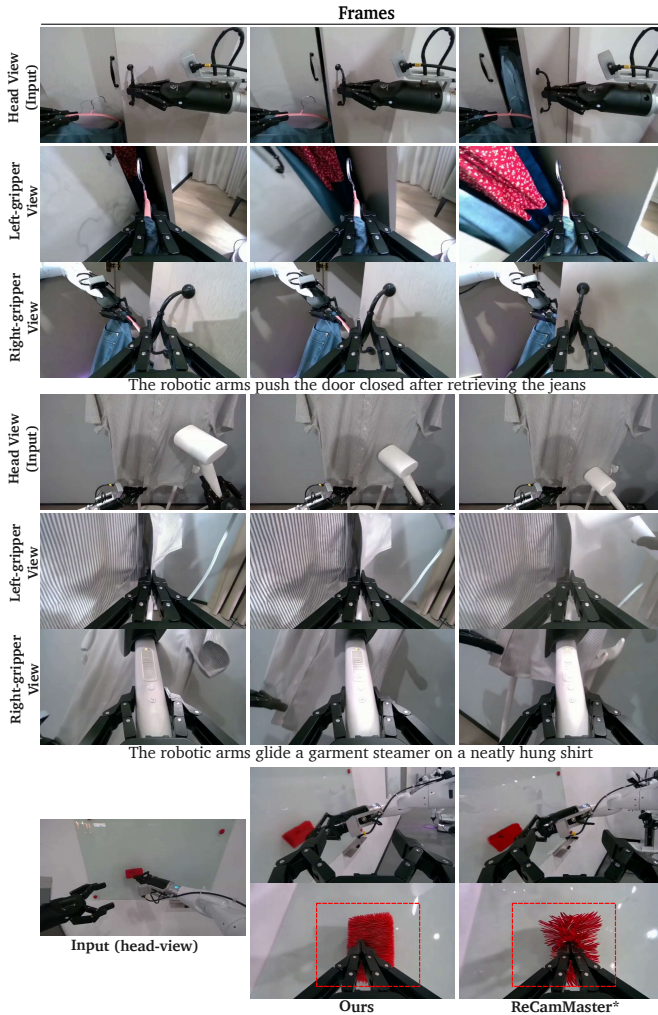


Figure 5. **Qualitative Results on the Agibot Benchmark.** Given a head-view manipulation video in Agibot, PlenopticDreamer-agibot (Ours) can generate temporally consistent videos from the left and right gripper viewpoints.

formation in an autoregressive manner. Specifically, given a head-view manipulation video, it generates temporally consistent videos from both left and right gripper viewpoints across diverse manipulation tasks. In contrast, ReCamMaster* (also retrained on the Agibot dataset) fails to maintain view synchronization and high visual quality (see the blackboard eraser marked in the red dashed box).

Table 2. **Quantitative Comparison on the Agibot Benchmark.** Visual quality and view synchronization are assessed on 2 shots (left and right gripper viewpoints).

	PSNR \uparrow	View Sync. (Mat. Pix.(K) \uparrow)
ReCamMaster*	13.84	13.2
Ours	14.54	15.3

4.4. Ablation Study

Training Strategy. As shown in Fig. 6 and Tab. 3, removing the progressive training strategy (*w/o Progressive*



Figure 6. **Ablation Study.** Qualitative visualization of effects from different training strategies and context retrieval method.

Training) leads to a notable degradation in camera accuracy ($0.54 \rightarrow 0.63$ in TransErr), and an occluded man becomes erroneously visible under the “Rotation Right” case. This indicates that progressively enlarging the context size effectively stabilizes model convergence and enhances camera performance. When the self-conditioned training is removed, the generated videos exhibit pronounced artifacts and over-exposure, particularly in long-shot sequences. Correspondingly, both FVD and IQ (Image Quality in VBench [31]) metrics worsen, verifying that training with imperfect inputs enhances model robustness and mitigates error accumulation over time.

Video Retrieval Strategy. Replacing the proposed retrieval mechanism with random selection leads to a significant decline in view synchronization across all shots, with inconsistent hallucinations (highlighted with red dashed boxes in Fig. 6). We further examine the impact of context size in Tab. 4: increasing the number of retrieved contexts from 4 to 6 enhances multi-view consistency by offering richer spatial cues. However, further enlarging the context brings diminishing gains due to compounded trajectory fusion errors and accumulated generative noise.

4.5. Application

Long Video Generation. With the proposed long-video conditioning strategy, PlenopticDreamer supports coherent long-context video re-rendering, as shown in Fig. 7. Given a leftward rotation trajectory, our method produces temporally consistent long video segments while preserving spatial coherence across adjacent chunks. In contrast, removing this conditioning results in visible inconsistencies, as

Table 3. **Ablation Study on the Basic Benchmark.** Evaluation is conducted on the full set, comprising a total of 1,200 generated videos.

Model	Visual Quality		Camera Accuracy		View Synchronization (Mat. Pix.(K) \uparrow)			
	FVD \downarrow	IQ \uparrow	TransErr \downarrow	RotErr (rad) \downarrow	3 Shots	6 Shots	9 Shots	12 Shots
w/o Self-Cond. Training	464.3	56.7	0.54	0.23	40.9	40.2	<u>45.1</u>	40.7
w/ Random Context Retrieval	520.5	<u>58.3</u>	<u>0.56</u>	0.20	33.6	33.4	36.5	32.4
w/o Progressive Training	<u>453.8</u>	57.2	0.63	0.23	39.6	<u>40.6</u>	43.6	39.4
Full Model	425.8	58.5	0.54	<u>0.21</u>	41.4	40.8	45.4	41.2



Figure 7. **Long Video Generation.** Given a leftward rotation camera trajectory, ours (w/ LVG Cond.) preserves spatial consistency across adjacent video chunks, yielding seamless transitions at their boundaries (highlighted by red dotted lines).

Table 4. **Ablation on Retrieved Context Video Number.**

Video Num.	View Synchronization (Mat. Pix.(K) \uparrow)			
	3 Shots	6 Shots	9 Shots	12 Shots
4	58.1	51.2	52.1	42.7
6	\uparrow	52.1	53.1	43.6
8	\uparrow	\uparrow	50.2	41.0
10	\uparrow	\uparrow	\uparrow	40.8

illustrated in the second row.

Focal Length Effect. Varying the input focal length leads to corresponding depth-of-field changes, as shown in Fig. 8 under a “zoom-in” camera trajectory. This enables finer control over camera behavior and visual emphasis, offering users greater flexibility in camera-aware video generation.

5. Conclusion

We introduce PlenopticDreamer, a camera-controlled video re-rendering framework enforcing spatio-temporal consistency. It employs a multi-in-single-out, *autoregressive* diffusion model conditioned on spatio-temporal memory via a 3D FOV strategy for coherent hallucinations along trajectories. By incorporating context-scaling and self-conditioning, the method improves stability and reduces error accumulation in long-range generation. Evaluation on Basic and Agibot benchmarks shows state-of-the-art view synchronization, high fidelity, and precise camera control.



Figure 8. **Focal Length Effect.** Our method simulates varying depth-of-field effects corresponding to different focal lengths (18mm \rightarrow 100mm) under a “zoom-in” camera trajectory.

Limitations. Despite self-conditioned training, ours exhibits occasional failures, including over-exposure and distortion in long-shot videos. Future work could explore a Self-Forcing-style [15, 30, 42] paradigm. Besides, computational cost scales with memory views. We also observed degraded performance in complex human motions, such as dancing, likely from pretraining data biases in Cosmos.

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 4
- [2] 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. 5
- [3] Arslan Ali, Junjie Bai, Maciej Bala, Yogesh Balaji, Aaron Blakeman, Tiffany Cai, Jiaxin Cao, Tianshi Cao, Elizabeth Cha, Yu-Wei Chao, et al. World simulation with video foundation models for physical ai. *arXiv preprint arXiv:2511.00062*, 2025. 1, 5
- [4] Sherwin Bahmani, Ivan Skorokhodov, Guocheng Qian, Aliaksandr Siarohin, Willi Menapace, Andrea Tagliasacchi, David B Lindell, and Sergey Tulyakov. Ac3d: Analyzing and improving 3d camera control in video diffusion transformers. In *CVPR*, 2025. 1, 2
- [5] Sherwin Bahmani, Ivan Skorokhodov, Aliaksandr Siarohin, Willi Menapace, Guocheng Qian, Michael Vasilkovsky, Hsin-Ying Lee, Chaoyang Wang, Jiaxu Zou, Andrea Tagliasacchi, et al. Vd3d: Taming large video diffusion transformers for 3d camera control. In *ICLR*, 2025. 2
- [6] Jianhong Bai, Menghan Xia, Xiao Fu, Xintao Wang, Lianrui Mu, Jinwen Cao, Zuozhu Liu, Haoji Hu, Xiang Bai, Pengfei Wan, et al. Recammaster: Camera-controlled generative rendering from a single video. In *ICCV*, 2025. 2, 3, 4, 5, 6
- [7] Jianhong Bai, Menghan Xia, Xintao Wang, Ziyang Yuan, Xiao Fu, Zuozhu Liu, Haoji Hu, Pengfei Wan, and Di Zhang. Syncammaster: Synchronizing multi-camera video generation from diverse viewpoints. In *ICLR*, 2025. 2, 5
- [8] James R Bergen and Edward H Adelson. The plenoptic function and the elements of early vision. *Computational models of visual processing*, 1991. 1
- [9] Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, et al. Video generation models as world simulators. *OpenAI Blog*, 1(8):1, 2024. 1
- [10] Qingwen Bu, Jisong Cai, Li Chen, Xiuqi Cui, Yan Ding, Siyuan Feng, Shen Yuan Gao, Xindong He, Xuan Hu, Xu Huang, et al. Agibot world colosseum: A large-scale manipulation platform for scalable and intelligent embodied systems. In *IROS*, 2025. 2, 6
- [11] Shengqu Cai, Ceyuan Yang, Lvmin Zhang, Yuwei Guo, Junfei Xiao, Ziyang Yang, Yinghao Xu, Zhenheng Yang, Alan Yuille, Leonidas Guibas, et al. Mixture of contexts for long video generation. *arXiv preprint arXiv:2508.21058*, 2025. 2
- [12] Boyuan Chen, Diego Martí Monsó, Yilun Du, Max Simchowitz, Russ Tedrake, and Vincent Sitzmann. Diffusion forcing: Next-token prediction meets full-sequence diffusion. In *NeurIPS*, 2024. 4
- [13] 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*, 2025. 2
- [14] Taiye Chen, Xun Hu, Zihan Ding, and Chi Jin. Learning world models for interactive video generation. *arXiv preprint arXiv:2505.21996*, 2025. 2
- [15] Justin Cui, Jie Wu, Ming Li, Tao Yang, Xiaojie Li, Rui Wang, Andrew Bai, Yuanhao Ban, and Cho-Jui Hsieh. Self-forcing++: Towards minute-scale high-quality video generation. *arXiv preprint arXiv:2510.02283*, 2025. 8
- [16] Karan Dalal, Daniel Kocejka, Jiarui Xu, Yue Zhao, Shihao Han, Ka Chun Cheung, Jan Kautz, Yejin Choi, Yu Sun, and Xiaolong Wang. One-minute video generation with test-time training. In *CVPR*, 2025. 2
- [17] Johan Edstedt, Qiyu Sun, Georg Bökman, Márten Wadenbäck, and Michael Felsberg. Roma: Robust dense feature matching. In *CVPR*, 2024. 5
- [18] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *ICML*, 2024. 3
- [19] Xiao Fu, Xian Liu, Xintao Wang, Sida Peng, Menghan Xia, Xiaoyu Shi, Ziyang Yuan, Pengfei Wan, Di Zhang, and Dahua Lin. 3dtrajmaster: Mastering 3d trajectory for multi-entity motion in video generation. In *ICLR*, 2025. 2
- [20] Xiao Fu, Xintao Wang, Xian Liu, Jianhong Bai, Runsen Xu, Pengfei Wan, Di Zhang, and Dahua Lin. Learning video generation for robotic manipulation with collaborative trajectory control. *arXiv preprint arXiv:2506.01943*, 2025. 1
- [21] Yu Gao, Haoyuan Guo, Tuyen Hoang, Weilin Huang, Lu Jiang, Fangyuan Kong, Huixia Li, Jiashi Li, Liang Li, Xiaojie Li, et al. Seedance 1.0: Exploring the boundaries of video generation models. *arXiv preprint arXiv:2506.09113*, 2025. 1
- [22] Zekai Gu, Rui Yan, Jiahao Lu, Peng Li, Zhiyang Dou, Chenyang Si, Zhen Dong, Qifeng Liu, Cheng Lin, Ziwei Liu, et al. Diffusion as shader: 3d-aware video diffusion for versatile video generation control. In *SIGGRAPH*, 2025. 2
- [23] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. In *ICLR*, 2024. 1
- [24] Hao He, Yinghao Xu, Yuwei Guo, Gordon Wetzstein, Bo Dai, Hongsheng Li, and Ceyuan Yang. Cameractrl: Enabling camera control for video diffusion models. In *ICLR*, 2025. 1, 2, 5
- [25] Xuanhua He, Quande Liu, Zixuan Ye, Weicai Ye, Qulin Wang, Xintao Wang, Qifeng Chen, Pengfei Wan, Di Zhang, and Kun Gai. Fulldit2: Efficient in-context conditioning for video diffusion transformers. *arXiv preprint arXiv:2506.04213*, 2025. 3
- [26] Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. In *ICLR*, 2023. 1
- [27] Chen Hou and Zhibo Chen. Training-free camera control for video generation. *arXiv preprint arXiv:2406.10126*, 2024. 2

- [28] Teng Hu, Jiangning Zhang, Ran Yi, Yating Wang, Hongrui Huang, Jieyu Weng, Yabiao Wang, and Lizhuang Ma. Motionmaster: Training-free camera motion transfer for video generation. *arXiv preprint arXiv:2404.15789*, 2024. 2
- [29] 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. 5
- [30] Xun Huang, Zhengqi Li, Guande He, Mingyuan Zhou, and Eli Shechtman. Self forcing: Bridging the train-test gap in autoregressive video diffusion. In *NeurIPS*, 2025. 4, 8
- [31] Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qingyang Jin, Nattapol Chanpaisit, et al. Vbench: Comprehensive benchmark suite for video generative models. In *CVPR*, 2024. 7
- [32] Xuan Ju, Weicai Ye, Quande Liu, Qiulin Wang, Xintao Wang, Pengfei Wan, Di Zhang, Kun Gai, and Qiang Xu. Fulldit: Multi-task video generative foundation model with full attention. In *ICCV*, 2025. 3
- [33] Anssi Kanervisto, Dave Bignell, Linda Yilin Wen, Martin Grayson, Raluca Georgescu, Sergio Valcarcel Macua, Shan Zheng Tan, Tabish Rashid, Tim Pearce, Yuhan Cao, et al. World and human action models towards gameplay ideation. *Nature*, 638(8051):656–663, 2025. 2
- [34] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*, 2024. 1
- [35] Zhengfei Kuang, Shengqu Cai, Hao He, Yinghao Xu, Hongsheng Li, Leonidas J Guibas, and Gordon Wetzstein. Collaborative video diffusion: Consistent multi-video generation with camera control. In *NeurIPS*, 2024. 2
- [36] Bing Li, Cheng Zheng, Wenxuan Zhu, Jinjie Mai, Biao Zhang, Peter Wonka, and Bernard Ghanem. Vivid-zoo: Multi-view video generation with diffusion model. In *NeurIPS*, 2024. 2
- [37] 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*, 2025. 2
- [38] Runjia Li, Philip Torr, Andrea Vedaldi, and Tomas Jakab. Vmem: Consistent interactive video scene generation with surfel-indexed view memory. In *ICCV*, 2025. 2
- [39] Zhengqi Li, Wenqi Xian, Abe Davis, and Noah Snavely. Crowdsampling the plenoptic function. In *ECCV*, 2020. 1
- [40] Pengyang Ling, Jiazi Bu, Pan Zhang, Xiaoyi Dong, Yuhang Zang, Tong Wu, Huaian Chen, Jiaqi Wang, and Yi Jin. Motionclone: Training-free motion cloning for controllable video generation. *arXiv preprint arXiv:2406.05338*, 2024. 2
- [41] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *ICLR*, 2023. 3
- [42] Kunhao Liu, Wenbo Hu, Jiale Xu, Ying Shan, and Shijian Lu. Rolling forcing: Autoregressive long video diffusion in real time. *arXiv preprint arXiv:2509.25161*, 2025. 8
- [43] Zhiheng Liu, Xueqing Deng, Shoufa Chen, Angtian Wang, Qiushan Guo, Mingfei Han, Zeyue Xue, Mengzhao Chen, Ping Luo, and Linjie Yang. Worldweaver: Generating long-horizon video worlds via rich perception. *arXiv preprint arXiv:2508.15720*, 2025. 2
- [44] 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. 2
- [45] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *ICCV*, 2023. 1, 3
- [46] Xuanchi Ren, Tianchang Shen, Jiahui Huang, Huan Ling, Yifan Lu, Merlin Nimier-David, Thomas Müller, Alexander Keller, Sanja Fidler, and Jun Gao. Gen3c: 3d-informed world-consistent video generation with precise camera control. In *CVPR*, 2025. 1, 2
- [47] Manuel-Andreas Schneider, Lukas Höllein, and Matthias Nießner. Worldexplorer: Towards generating fully navigable 3d scenes. *arXiv preprint arXiv:2506.01799*, 2025. 2
- [48] Vincent Sitzmann, Semon Rezhikov, Bill Freeman, Josh Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. 2021. 4
- [49] Kiwhan Song, Boyuan Chen, Max Simchowitz, Yilun Du, Russ Tedrake, and Vincent Sitzmann. History-guided video diffusion. *arXiv preprint arXiv:2502.06764*, 2025. 2
- [50] Google Deepmind Veo Team. Veo 3. 2025. 1
- [51] Kling Team. Kling ai 2.5 turbo. 2025. 1
- [52] Basile Van Hoorick, Rundi Wu, Ege Ozguroglu, Kyle Sargent, Ruoshi Liu, Pavel Tokmakov, Achal Dave, Changxi Zheng, and Carl Vondrick. Generative camera dolly: Extreme monocular dynamic novel view synthesis. In *ECCV*, 2024. 2
- [53] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwei Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025. 1
- [54] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *CVPR*, 2025. 5
- [55] Yuqing Wang, Tianwei Xiong, Daquan Zhou, Zhijie Lin, Yang Zhao, Bingyi Kang, Jiashi Feng, and Xihui Liu. Loong: Generating minute-level long videos with autoregressive language models. *arXiv preprint arXiv:2410.02757*, 2024. 4
- [56] Zhouxia Wang, Yushi Lan, Shangchen Zhou, and Chen Change Loy. Objctrl-2.5 d: Training-free object control with camera poses. *arXiv preprint arXiv:2412.07721*, 2024. 2
- [57] 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 *SIGGRAPH*, 2024. 1, 2
- [58] Tong Wu, Shuai Yang, Ryan Po, Yinghao Xu, Ziwei Liu, Dahua Lin, and Gordon Wetzstein. Video world models with long-term spatial memory. In *NeurIPS*, 2025. 2

- [59] Zeqi Xiao, Yushi Lan, Yifan Zhou, Wenqi Ouyang, Shuai Yang, Yanhong Zeng, and Xingang Pan. Worldmem: Long-term consistent world simulation with memory. In *NeurIPS*, 2025. [2](#)
- [60] Zeqi Xiao, Wenqi Ouyang, Yifan Zhou, Shuai Yang, Lei Yang, Jianlou Si, and Xingang Pan. Trajectory attention for fine-grained video motion control. In *ICLR*, 2025. [2](#), [5](#), [6](#)
- [61] Yiming Xie, Chun-Han Yao, Vikram Voleti, Huaizu Jiang, and Varun Jampani. Sv4d: Dynamic 3d content generation with multi-frame and multi-view consistency. In *ICLR*, 2025. [2](#)
- [62] Dejjia Xu, Weili Nie, Chao Liu, Sifei Liu, Jan Kautz, Zhangyang Wang, and Arash Vahdat. Camco: Camera-controllable 3d-consistent image-to-video generation. *arXiv preprint arXiv:2406.02509*, 2024. [2](#)
- [63] Shiyuan Yang, Liang Hou, Haibin Huang, Chongyang Ma, Pengfei Wan, Di Zhang, Xiaodong Chen, and Jing Liao. Direct-a-video: Customized video generation with user-directed camera movement and object motion. In *SIGGRAPH*, 2024. [2](#)
- [64] 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 *SIGGRAPH Asia*, 2025. [2](#), [3](#), [4](#)
- [65] Mark YU, Wenbo Hu, Jinbo Xing, and Ying Shan. Trajectorycrafter: Redirecting camera trajectory for monocular videos via diffusion models. In *ICCV*, 2025. [2](#), [5](#), [6](#)
- [66] Lvmin Zhang and Maneesh Agrawala. Packing input frame context in next-frame prediction models for video generation. *arXiv preprint arXiv:2504.12626*, 2025. [2](#)
- [67] Guangcong Zheng, Teng Li, Rui Jiang, Yehao Lu, Tao Wu, and Xi Li. Cami2v: Camera-controlled image-to-video diffusion model. *arXiv preprint arXiv:2410.15957*, 2024. [2](#)