

Matrix-Game 3.5: Enhancing Real-Time Streaming Interactive World Models with Patch Memory

Riemann Dynamics
research@riemandedynamics.ai
Project page: [Matrix-Game-3.5-Homepage](#)




Figure 1: Overview of **Matrix-Game 3.5**. A real-time interactive world model for long-horizon generation with pose-aware representation, unified memory, and progressive distillation, achieving single-GPU 720p real-time generation at 20 FPS.

Abstract

Interactive world models extend video generation from offline clip synthesis toward persistent simulation of interactive virtual worlds, enabling applications in games, robotics, embodied agents, and XR. Achieving stable long-horizon interactive generation, however, remains challenging, as the model must simultaneously preserve scene geometry, dynamic consistency, and camera control while supporting real-time autoregressive generation. Building upon Matrix-Game 3.0, we present **Matrix-Game 3.5**, as shown in Figure 1, which advances real-time interactive world generation toward geometry-aware and long-horizon consistent simulation through three key improvements. First, we propose a unified parameter-free geometry-aware memory framework that integrates Patch Memory and Warped PRoPE. Patch Memory lifts historical observations into 3D patch representations for geometry-aware retrieval, while Warped PRoPE injects camera projection geometry into spatiotemporal rotary position embeddings, unifying long-term memory

and camera control without modifying the backbone architecture. Second, we introduce a disentangled world representation that separately models static scene geometry and dynamic subjects through static Patch Memory and lightweight multi-view Reference Tokens, jointly preserving geometric consistency and subject identity throughout long-horizon generation. Third, we develop a long-horizon real-time distillation framework that progressively transfers a bidirectional diffusion model into a few-step causal generator through Flow Matching and curriculum-based Self-Rollout DMD, enabling minute-long real-time interactive generation. Extensive experiments demonstrate that, with a unified training corpus spanning Unreal simulation environments, open-world games, and Internet videos, Matrix-Game 3.5 achieves strong performance in long-horizon scene recall, precise camera control, subject consistency, prompt-driven world generation, and stable real-time open-world interaction.

1 Introduction

Interactive World Models aim to build persistent virtual environments that continuously evolve in response to user interactions. Unlike conventional video generation models that synthesize short offline clips, interactive world models must causally generate future observations while responding to camera motions, user actions, and high-level prompts. More importantly, the generated world should remain persistent throughout long-horizon interaction: previously visited scenes should be faithfully reconstructed, objects and characters should preserve their identities under viewpoint changes, and user-triggered events should produce coherent and lasting changes to the environment. These capabilities are fundamental for open-world gaming, embodied-agent simulation, robotics, extended reality (XR), and interactive content creation.

Building such persistent worlds remains fundamentally challenging. Interactive world models operate in a closed-loop autoregressive manner, where future predictions continuously become the context for subsequent generation. Consequently, visual and geometric errors inevitably accumulate over time, gradually causing scene drift, identity inconsistency, and degraded controllability. At the same time, the limited context window prevents the model from directly accessing distant observations, making long-term scene revisitation increasingly difficult. Large viewpoint changes further require the model not only to retrieve relevant historical information, but also to determine its correct geometric correspondence under the current camera. Moreover, while static environments should remain persistent over long horizons, dynamic entities are expected to evolve according to user interactions and prompts. Therefore, achieving long-horizon consistency is fundamentally a joint problem of memory, geometry, dynamics, and efficient autoregressive generation.

Existing approaches mainly address this challenge from two complementary perspectives. One line of work compresses historical observations into latent representations or transformer memories (e.g., Matrix-Game 3.0 [1], RELIC [2]), enabling efficient long-context modeling while representing historical information only implicitly. Another line explicitly retrieves historical observations through camera poses or reconstructed geometry (e.g., WorldMem [3], VMem [4]), improving scene revisitation by leveraging explicit spatial cues. Although both paradigms significantly improve long-horizon consistency, neither fully resolves a more fundamental question: *an interactive world model must jointly reason about what historical information should be remembered, where it should reappear under a novel viewpoint, and whether the remembered content should remain persistent or evolve with the world.* As a result, long-term memory, camera control, dynamic-content consistency, and prompt-driven interaction are still largely addressed as separate problems rather than a unified world representation.

In this work, we present **Matrix-Game 3.5**, the latest generation of the Matrix-Game series for real-time long-horizon interactive world generation. Our central idea is to organize historical observations as an explicit geometry-aware persistent memory instead of treating them as a sequence of latent tokens. Based on this perspective, we introduce a unified parameter-free geometry-aware memory framework that integrates Patch Memory [5] and Warped PRoPE, our video adaptation of PRoPE [6]. Patch Memory lifts historical image patches into a persistent 3D memory and retrieves only the regions visible from the target viewpoint, while Warped PRoPE establishes explicit geometric correspondence between retrieved memories and target-view attention. To model dynamic worlds, we further disentangle persistent static environments from evolving dynamic entities, storing stable scene

structures in Patch Memory while representing movable objects and characters using lightweight multi-view Reference Tokens. Finally, we develop a progressive long-horizon distillation framework, building on recent causal video-distillation pipelines [7, 8], that transforms a high-quality bidirectional world model into an efficient few-step fully causal generator, enabling stable real-time interaction without sacrificing long-range consistency.

Extensive experiments demonstrate that, when trained on diverse data collected from Unreal simulation environments, open-world games, and Internet videos, Matrix-Game 3.5 achieves strong performance in long-horizon scene recall, geometry-consistent camera control, subject-consistent generation, prompt-driven world evolution, and stable real-time open-world interaction. These results demonstrate that explicitly unifying memory, geometry, and dynamic world representation provides an effective foundation for scalable, controllable, and persistent interactive world simulation.

Our main contributions are summarized as follows:

- We propose a geometry-aware persistent world representation that unifies long-horizon memory and geometry-consistent camera control without additional backbone parameters. Built on Patch Memory and Warped P_{RoPE}, it enables explicit 3D memory retrieval and geometry-aware attention across viewpoint changes.
- We propose a dynamic-static disentangled world modeling strategy with lightweight multi-view reference tokens for dynamic subject representation. Combined with motion-aware filtering and leakage-resistant subject training, it effectively suppresses ghosting and identity drift during long-horizon interactive generation.
- We develop a progressive long-horizon real-time distillation framework that combines perceptual feature-space Flow Matching with curriculum self-rollout DMD. It transfers a bidirectional world model into a few-step fully causal interactive generator, enabling minute-long real-time generation with long-horizon consistency and accurate camera control.

2 Method

Matrix-Game 3.5 consists of three tightly coupled components. First, we introduce a pose-aware sequence representation that unifies heterogeneous inputs into a single self-attention sequence with shared spatiotemporal and camera-aware coordinates. Second, we develop a unified memory system that combines geometry-aware patch memory with reference-token memory to preserve both scene geometry and dynamic-subject identity over long-horizon interactions. Finally, we present a progressive distillation framework that converts a bidirectional video diffusion model into a few-step fully causal autoregressive generator for real-time interactive generation.

2.1 Pose-aware Sequence Representation

Matrix-Game 3.5 arranges all token types — the anchor frame, the clean context frames, the retrieved memory patches, and the noisy target frames — into a single pose-aware token sequence. Because our way of conditioning on the camera tiles the camera projection over the very channels that already carry the spatiotemporal RoPE, every token receives both a camera pose and a RoPE spatiotemporal coordinate within this one sequence. We first describe how the camera pose is injected into attention, and then how these coordinates are assigned across the sequence.

2.1.1 Camera-aware Attention

Matrix-Game 3.5 first requires explicit camera control over the generated video. Moreover, since the tokens in this sequence come from different cameras and different times, we want the interaction between frames to be pose-aware rather than to rely on the token grid of the native 3D RoPE [9] alone. We build on P_{RoPE} [6], an existing scheme that turns camera geometry into a relative position code: every latent frame is associated with a full world-to-image projection matrix $P = \text{lift}(K) \cdot W$ (with K the intrinsics and W the world-to-camera extrinsics), so the attention score between two frames carries their relative projection $M = P_i P_j^{-1}$, jointly encoding relative rotation, translation, and intrinsics. Native P_{RoPE}, designed for multi-view images, splits the attention head dimension into disjoint blocks — half for the camera projection and the other two for the 2D row and column position — so camera and position never multiply and there is no time axis (Figure 2a). This native

design suits novel-view synthesis, where the inputs are unordered views, but it cannot be transplanted directly to video generation: our backbone relies on the frame axis of its 3D spatiotemporal RoPE to encode temporal ordering, and this temporal channel must be preserved. Our layout therefore departs from the native design and tiles the camera projection across all head channels, on top of the full spatiotemporal RoPE including its frame axis, without carving out a separate subspace (Figure 2b). Concretely, after the standard spatiotemporal RoPE rotation the queries and keys are further multiplied by the camera matrices — queries by $P^\top$, keys and values by P^{-1} — so a single softmax carries both the relative time $\Delta t = i - j$ and the relative pose M , adding no learnable parameters (Figure 2c).




Figure 2: **Injection of camera projection geometry into spatiotemporal RoPE.** (a) Native PProPE, designed for unordered multi-view images in novel-view synthesis where no time axis is needed, splits the head dimension into disjoint blocks: camera projection takes half, and the other two encode the 2D patch row and column; camera and position occupy disjoint subspaces and never multiply. (b) The proposed overlay, redesigned for video generation, tiles the camera projection over *all* channels on top of the full spatiotemporal RoPE (frame axis included), without allocating a dedicated subspace, preserving the temporal channel that video generation depends on. (c) One attention link between frame i (query) and frame j (key/value): q and k are first RoPE-rotated, then multiplied by $\text{tile}(P_i^\top)$ and $\text{tile}(P_j^{-1})$; a single softmax thus carries both the relative time $\Delta t = i - j$ and the relative pose $M_{ij} = P_i P_j^{-1}$. Values are transported through the same relative pose ($v \leftarrow P^{-1}$, $\text{out} \leftarrow P$), with no learnable parameters.

Since this overlay departs from native PProPE and reuses the RoPE channels—including the temporal axis that video generation depends on—rather than adding a separate branch, we explicitly test the adaptation: Appendix B verifies that introducing our PProPE variant leaves the temporal information encoded by RoPE intact. To address numerical instability caused by large camera transformations, Matrix-Game 3.5 recenters all camera poses with respect to the first target frame of the current generation window and applies a direction-preserving logarithmic compression only to the metric-scale translation component: the translation magnitude is mapped to log scale and then divided by 4, while the translation direction is kept unchanged. This preserves small camera motions almost losslessly while preventing distant historical viewpoints from introducing numerical instability.

2.1.2 Pose-aware Sequence Layout

At each step of long-horizon autoregressive generation, the model must reconcile several heterogeneous sources at once: faithful unwarped observations of the past, the most recent clean frame for local continuity, a persistent representation of the dynamic subject, geometry-aligned memory of previously seen scene content, and the noisy frames currently being denoised. Rather than routing these through separate cross-attention branches, Matrix-Game 3.5 places all of them into the single pose-aware sequence introduced above and lets one self-attention stack relate them directly through their shared RoPE coordinates and camera poses. Because a token’s coordinates already encode where

and when its content belongs, camera control and every form of memory reduce to placing tokens in this one sequence; the memory mechanisms differ only in how they assign those coordinates.

These rules give every token a place in one pose-aware sequence (Figure 3): each block — the object token for the dynamic subject, the context frames, the anchor, the patch-memory tokens, and the noisy target frames — sits at its own RoPE time and, where it corresponds to a view, carries its own camera, and the memory mechanisms in the next section simply populate these coordinates.




Figure 3: **The pose-aware token sequence.** All token types share one sequence, each carrying its own RoPE time and, for the camera-bearing blocks, its own camera. *Top:* the cameras that produce the sequence move along one trajectory around the scene, and the context frames (blue), the anchor (purple), and the noisy target frames (gray) each observe the scene from their own pose; the object token (teal) is a prefix token for the dynamic subject that carries no camera of its own, shown here linked to the character in the scene, with a front view of its appearance drawn below the token. *Bottom:* every camera maps straight down to its position on the shared RoPE timeline. The context frames and the anchor—the most recent clean frame—keep their true past RoPE times, separated from the generation window by a time jump (//). The noisy target frames and the patch-memory tokens reuse the *same* target cameras at the *same* RoPE time: each target camera contributes both a noisy token (on the axis) and a patch-memory token, the latter lying on a parallel memory track that runs alongside the main timeline, with the memory patches placed at projected floating-point spatial coordinates and uncovered patches (holes) dropped.

2.2 Unified Memory System

Built on the pose-aware sequence above, the unified memory system comprises three complementary mechanisms: a geometry-aware patch memory that recalls previously observed scene content, a context frame memory that supplies faithful unwrapped observations, and a static-dynamic decoupled design that separates persistent scene structure from dynamic subjects.

2.2.1 Patch Memory

In autoregressive long-video generation, the current segment must not only remain visually consistent with previously generated content, but also recover scene structures that were observed earlier when the camera revisits the same location. Existing approaches can largely be divided into two categories. Explicit memory methods represent scenes using points, splats, or other geometric primitives, providing strong geometric constraints but primarily targeting static scene reconstruction. Implicit memory methods instead preserve entire video frames or latent features as memory, offering greater generative flexibility, yet relying on the network itself to learn cross-view correspondences. We observe that an intermediate memory granularity—patch-level memory—has received relatively

little attention. Rather than treating an entire frame as an indivisible memory unit or maintaining a complete global 3D representation, we follow MosaicMem [5] and store localized image patches as the fundamental unit of memory, allowing them to be selectively retrieved, spatially aligned, and reused when needed.




Figure 4: **The patch-memory mechanism.** (a) *Lift to 3D*: historical latent patches are back-projected into a common 3D space using metric depth together with camera intrinsics and poses. (b) *Frustum query*: the target camera frustum queries this patch cloud, and a z-buffer under the target view keeps, for every target latent location, only the surface closest to the camera while discarding occluded candidates. (c) *Patch memory frame*: the selected patches are scattered into an aligned memory canvas under the target viewpoint; regions that are occluded or have never been observed are left empty (black) for the diffusion model to synthesize from the current latent, text prompt, and surrounding context.

The key novelty of this patch-level memory lies in its organic combination of the complementary strengths of explicit and implicit memory. Its first stage resembles an explicit-memory pipeline: historical patches are lifted into a common 3D space using metric depth together with camera intrinsics and poses. Its second stage resembles an implicit-memory pipeline: the retrieved and geometrically aligned patches are injected into the diffusion transformer as reference conditions rather than hard reconstruction targets. This design enables the model to exploit accurate geometric correspondence for reliable scene recall while retaining the flexibility to synthesize novel content, dynamic objects, and unseen regions according to the text prompt. Moreover, because memory is represented at the patch level rather than as complete frames or globally reconstructed geometry, the memory naturally supports sparse retrieval, localized editing and manipulation, and robust long-horizon memory updates without accumulating global reconstruction errors over repeated revisits.

Concretely, the patch memory addresses a simple but fundamental question: for every latent patch in the target frame, where should its memory be retrieved from within the previously observed video? Instead of learning this correspondence implicitly, we compute it directly using camera geometry. Every historical frame is associated with its camera-to-world pose, camera intrinsics, and a metric depth map. Given a latent patch from a historical frame, we first recover its center in image space according to the VAE downsampling ratio, then back-project it into the historical camera coordinate system using the corresponding depth and camera intrinsics. The resulting 3D point is transformed into world coordinates through the historical camera pose, and finally projected into the target camera using the target pose and intrinsics. This geometric lifting-and-reprojection procedure establishes an explicit patch-wise correspondence between historical observations and the target viewpoint (Figure 4).

During memory fusion, multiple historical patches may project onto the same target latent location. To resolve occlusions and conflicting observations, we perform a z-buffer selection under the target viewpoint, retaining only the surface closest to the target camera while discarding farther candidates. The selected latent feature is written into an aligned memory canvas under the target view, while a reverse mapping records both the source frame and the original patch location from which the memory originates. Regions that are occluded, geometrically unreliable, or have never been observed leave holes on the memory canvas; these holes are dropped from the token sequence rather than kept as empty placeholder tokens, so the memory track avoids doubling the sequence length, and the diffusion model synthesizes the corresponding regions conditioned on the current noisy latent, text

prompt, and surrounding context. Consequently, memory is reused only where reliable geometric evidence exists, while newly emerging content remains fully generative.

The complete pipeline further begins by selecting a small set of complementary historical frames for each target frame. Simply choosing the nearest camera poses often results in highly redundant viewpoints with poor overall coverage. Instead, we adopt a coverage-aware selection strategy that progressively chooses the historical frames contributing the largest additional coverage over the target latent grid. The selected frames are then processed through geometric reprojection, z-buffer fusion, and latent scattering to construct two outputs: an aligned memory canvas under the target viewpoint and a patch-wise correspondence map recording the origin of every retrieved memory token. The aligned memory canvas is provided to the diffusion transformer as spatial conditioning, while the correspondence map is used to modify the positional encoding of memory tokens, enabling the model to reason about their geometric relationship with the current latent tokens. Combined with dynamic-object masking, metric-scale depth alignment, and visibility consistency checks, this mechanism provides a geometrically reliable, locally controllable, and robust patch-level spatial memory for long-horizon autoregressive video generation. Within the pose-aware sequence, each retrieved patch declares its position by geometry rather than by storage order: it takes the RoPE timestamp of the target frame it supports rather than its original historical time, and, where a reliable projection exists, the floating-point spatial coordinate at which its source content lands in the target view; these fractional coordinates index a higher-resolution RoPE frequency table for sub-grid precision, and fall back to the native integer grid where no reliable projection exists. A patch is thus addressed by where it should appear in the frame being generated rather than by where it was stored.

2.2.2 Context Frame Memory

The patch memory supplies target-aligned local content, but it is warped and excluded from cross-attention; to also give the model faithful, unwarped observations of the scene, the clean prefix carries a small set of context frames placed before the anchor. The anchor is the clean frame immediately preceding the generating window, whereas context frames are drawn from the earlier history of the same clip. Rather than the nearest frames, which tend to be redundant, context frames are selected so that their camera poses cover the viewpoints of the window being generated, providing complementary past observations of the scene. Within the pose-aware sequence, each context frame keeps its own coordinates: it retains its true historical RoPE timestamp — so the relative temporal gaps between the older context, the anchor, and the generating window are preserved rather than collapsed to consecutive indices — together with its native spatial grid and its real camera pose. Context memory and patch memory are thus complementary: context memory contributes whole, unwarped observations at their original time and viewpoint, while patch memory contributes local content warped into the current view.

2.2.3 Static-Dynamic Decoupled Memory




Figure 5: Overview of the static-dynamic decoupled memory system. A motion-aware object filter separates static scene observations from dynamic subjects. Static regions are fused into the patch memory, while dynamic subjects are represented by reference-token memory. Subject-region masks are also used for context masking and an auxiliary training loss.

Long-horizon video generation requires a memory system that can preserve both scene-level stability and object-level temporal consistency. These two requirements correspond to different temporal structures. Static scene content, including background geometry, global layout, and stationary objects, should remain spatially consistent across time. Dynamic objects, in contrast, preserve their identity and appearance while undergoing time-dependent changes in position, pose, and visibility.

The memory system separates these two types of temporal structure. The static branch is built on the patch memory, a geometry-aware mechanism for retrieving and reusing previous observations of the stable scene. The dynamic branch is represented by sequence-level reference tokens, which provide persistent identity and appearance cues for dynamic subjects.

Without such a separation, dynamic objects can be accumulated into a static spatial memory. The same object observed at different time steps may then appear at different locations or with different poses. Memory fusion can conflate these temporal states with independent static observations, leading to duplicated objects, ghosting artifacts, incorrect occlusions, and ambiguous supervision.

Motion-aware Object Filter The motion-aware dynamic object filter serves as the interface between the static and dynamic memory branches. It applies YOLO segmentation and tracking [10] to common movable categories, including humans, vehicles, and animals, and estimates whether each tracked object is static or dynamic within a local video block. Regions classified as static are retained in the patch-memory branch, whereas regions classified as dynamic are masked before memory fusion.

The filter is based on motion rather than semantic category. A stationary foreground object remains a useful spatial reference, whereas an object from the same category should be excluded from static memory when it exhibits motion. Motion is therefore used as the criterion for static-memory validity.

For each track, the filter uses the predicted depth map together with camera intrinsics and extrinsics to test geometric consistency across nearby frames. Let M_i denote the segmentation mask in frame i , D_i the corresponding depth map, and (K_i, E_i) the camera intrinsics and extrinsics. Mask samples are lifted from frame i to 3D and reprojected to frame j , yielding a directional overlap score

$$s_{i \rightarrow j} = \frac{|\Pi_j(\mathcal{U}_i(M_i, D_i, K_i, E_i)) \cap M_j|}{|\Pi_j(\mathcal{U}_i(M_i, D_i, K_i, E_i))|}, \quad (1)$$

where $\mathcal{U}_i(\cdot)$ denotes depth-based unprojection and $\Pi_j(\cdot)$ denotes projection into frame j . The motion score is computed symmetrically as

$$s_{ij} = \min(s_{i \rightarrow j}, s_{j \rightarrow i}). \quad (2)$$

Tracks with sufficient temporal support and consistently high s_{ij} are retained as static observations, while tracks with low s_{ij} are excluded from the patch memory. Uncertain tracks are kept outside the static branch. The same filtering criterion is used in training and inference.

Reference Token Memory The dynamic branch of the memory system is implemented with sequence-level reference tokens. In the current third-person game setting, the dynamic subject refers to the controllable character whose identity and appearance should remain consistent throughout the rollout. Reference tokens serve as persistent representations of this subject, providing long-term identity and appearance memory even when it changes position, pose, or visibility.

Four reference images are encoded through the VAE and mapped to a compact latent token budget. Each reference image uses roughly one quarter of the anchor-context token budget, such that the four-reference prefix has a token count comparable to one anchor-context frame. Each reference patch token receives additive reference-index, type, and local spatial embeddings. The reference views are further assigned fixed-gap temporal RoPE coordinates: the r -th reference is placed at $t_r = -r\Delta_{\text{ref}}$, while the original video timeline remains unchanged. This gap-timing RoPE distinguishes the reference views and keeps their temporal placement independent of the absolute rollout step. For PRoPE, all reference tokens inherit the camera transform of the current clean anchor rather than a fixed identity camera. Since the clean anchor advances with each autoregressive window, the relative camera transform between the reference prefix and the target frames remains local to the current generation window, instead of growing with the accumulated camera motion over the full rollout.

During training, the subject regions in the anchor frame and memory-context frames are masked aggressively before being passed to the model. The masked regions are filled with background patches sampled from the same scene rather than with zeros or blurred noise, which reduces the train-inference gap and prevents direct appearance leakage from the visible conditioning frames. A subject-region auxiliary diffusion loss is computed on the subject mask to increase the gradient weight on the dynamic subject and encourage consistent subject-level modeling:

$$\mathcal{L} = \mathcal{L}_{\text{full}} + \lambda_{\text{sub}} \frac{\sum_p M_{\text{sub}}(p) \|\epsilon_{\theta}(p) - \epsilon(p)\|_2^2}{\sum_p M_{\text{sub}}(p) + \varepsilon}. \quad (3)$$

Here M_{sub} denotes the subject-region mask, and λ_{sub} controls the additional subject-region weight. The prefix-token design avoids an additional cross-attention module. A cross-attention formulation would introduce a separate alignment problem between reference space and video space, while also adding a parameterized conditioning path that can be bypassed through shortcut solutions. The prefix formulation keeps the reference branch inside the original transformer token space and allows it to interact with the generated video through the same self-attention stack.

2.3 Progressive Distillation




Figure 6: **Overview of our two-stage progressive distillation pipeline.** First, causal adaptation uses teacher-forced perceptual flow matching to obtain a high-quality few-step causal initializer. Second, self-rollout DMD directly optimizes inference-time autoregression through a curriculum that progressively distills classifier-free guidance, camera control, and memory-conditioned generation.

We use the two-stage pipeline in Figure 6. A teacher-forced causal-adaptation stage first provides a high-quality few-step initializer, after which self-rollout distribution matching directly optimizes the inference-time autoregressive process.

Causal Adaptation We reformulate generation as chunk-wise causal denoising. For target chunk i , the teacher-forced condition is

$$\mathcal{H}_i^{\text{gt}} = (\mathbf{x}_{<i}^{\text{gt}}, \mathbf{p}_{\leq i}, \mathbf{m}_i^{\text{gt}}, \mathbf{r}_i^{\text{gt}}, \mathbf{c}), \quad (4)$$

where $\mathbf{p}_{\leq i}$ denotes causal camera control, $\mathbf{m}_i^{\text{gt}}$ is the query-aligned patch memory, $\mathbf{r}_i^{\text{gt}}$ is the chunk-local context frame, and $\mathbf{c}$ is the text condition. Neither attention nor memory retrieval can access future chunks. Given noise ϵ , we construct

$$\mathbf{x}_t^i = (1 - t)\mathbf{x}_i + t\epsilon. \quad (5)$$

Following perceptual flow matching (PFM) [11], we recover the clean prediction

$$\hat{\mathbf{x}}_{0,\theta}^i = \mathbf{x}_t^i - t \mathbf{v}_\theta(\mathbf{x}_t^i, t \mid \mathcal{H}_i^{\text{gt}}), \quad (6)$$

and constrain flow matching in a pretrained perceptual feature space rather than regressing velocity only in the VAE latent space. Let $\mathcal{D}$ be the frozen VAE decoder, Φ_ℓ the ℓ -th feature block of a frozen perceptual model, and $\hat{\mathbf{x}}_{0,\theta}^{\leq i} = [\mathbf{x}_{<i}^{\text{gt}}, \hat{\mathbf{x}}_{0,\theta}^i]$ the complete teacher-forced causal chain. We optimize

$$\mathcal{L}_{\text{PFM}} = \mathbb{E}_{\mathbf{x}, i, t, \epsilon} \left[\frac{1}{|\mathcal{S}|} \sum_{\ell \in \mathcal{S}} d(\Phi_\ell(\mathcal{D}(\hat{\mathbf{x}}_{0,\theta}^{\leq i})), \Phi_\ell(\mathcal{D}(\mathbf{x}_{\leq i}^{\text{gt}}))) \right], \quad (7)$$

where d measures feature distance and $\mathcal{S}$ contains the selected perceptual blocks. This single objective simultaneously learns causal denoising and few-step generation, yielding an efficient, high-quality causal initializer.

Self-Rollout Distribution Matching The second stage removes teacher forcing and performs DMD [12, 13] on the student’s own autoregressive trajectory. For chunk i , the online condition becomes

$$\hat{\mathcal{H}}_i^\theta = (\hat{\mathbf{x}}_{<i}^\theta, \mathbf{p}_{\leq i}, \hat{\mathbf{m}}_i^\theta, \hat{\mathbf{r}}_i^\theta, \mathbf{c}), \quad (8)$$

where both patch memory and context frames are retrieved online from the causally visible generated history. The objective matches the student distribution $q_{\theta,t}$ to the bidirectional teacher distribution $p_{\phi,t}$ under this shared condition:

$$\mathcal{L}_{\text{DMD}} = \mathbb{E}_{i,t} \left[w(t) \text{KL} \left(q_{\theta,t}(\cdot \mid \hat{\mathcal{H}}_i^\theta) \parallel p_{\phi,t}(\cdot \mid \hat{\mathcal{H}}_i^\theta) \right) \right]. \quad (9)$$

Applying this objective to autoregressive video requires care because the student and scorers maintain different memory states. The student updates online memory from generated chunks, while feeding the same potentially drifted history to the bidirectional scorers would compromise the supervision. We therefore share only stable external conditions—the initial memory, anchor frame, text prompt, and camera trajectory—while allowing the student to update its memory and keeping scorer memory fixed. This provides reliable scores without forcing incompatible internal trajectories to match.

To stabilize training under this interface, we adopt a condition curriculum: first distilling classifier-free guidance and camera control without online memory, and then gradually introducing patch memory and context frames. Following HiAR [14], we keep the autoregressive prefix and chunk-local context at the next noise level during each denoising substep while retaining a clean anchor. This represents imperfect generated history with appropriate uncertainty, reducing long-horizon drift and encouraging the use of geometrically aligned patch memory.

Together, perceptual causal adaptation and curriculum-based self-rollout DMD yield a three-step causal generator that preserves quality and control over minute-long, real-time interactive streams.

3 Data Infrastructure

Memory-aware interactive world models require supervision beyond raw videos. In particular, our model relies on three complementary forms of supervision that are unavailable in large-scale video corpora: geometry supervision for long-horizon memory retrieval and geometry-consistent camera control, semantic supervision for temporally coherent world generation, and identity supervision for persistent dynamic-subject modeling. To this end, we build an offline data infrastructure that augments raw videos with geometry, language, and object-level annotations, followed by quality-aware data curation. The resulting training corpus provides unified supervision for Patch Memory, Warped PRoPE, and lightweight multi-view Reference Tokens.

3.1 Geometry Annotation for Geometry-Aware Generation

Patch Memory and Warped PRoPE require accurate geometric supervision to establish consistent 3D correspondences across long-horizon generation. Since raw videos do not provide camera poses, metric depth, or camera intrinsics, as shown in Figure 7, we build an offline geometry annotation pipeline based on VGGT-Omega [15]. Long videos are processed in overlapping temporal chunks to avoid memory limitations and accumulated reconstruction errors, and neighboring chunks are stitched through Sim(3) alignment to produce continuous camera trajectories.

Because VGGT-Omega predicts geometry only up to an unknown scale, we further introduce the metric branch of Depth Anything 3 [16] as a metric-scale anchor. A global optimization jointly calibrates all chunks while preserving their local geometric consistency, producing globally consistent metric camera trajectories and depth maps. The final annotation consists of camera-to-world poses, camera intrinsics, and metric depth maps, providing complete geometric supervision for geometry-aware memory retrieval and camera control.

3.2 Semantic Annotation for World Generation

Interactive world generation requires temporally coherent semantic supervision describing both persistent scene layouts and dynamic object evolution. Rather than generating independent frame-level captions, we annotate each temporal window using Gemma-4-26B-A4B-it [17], matching the temporal granularity of the downstream generator.

Each prompt jointly describes the static environment together with the identities, motions, trajectories, pose changes, and interactions of all dynamic objects. Degenerate descriptions such as “no motion” are explicitly suppressed, while subjective descriptions unrelated to scene dynamics are removed. The generated window-level prompts are expanded to all frames within the corresponding window and aligned with the geometric annotations, providing temporally stable semantic supervision for world generation.




Figure 7: **Metric-scale geometric annotation.** Left: the annotation pipeline — chunkwise VGGT-Omega inference over overlapping segments, per-segment metric anchoring with the Depth Anything 3 metric branch, and Sim(3) cross-segment alignment with a single global metric scale. Right: a reconstruction example — three views of the same scene, each with its input frame and per-frame estimated metric depth, and the three views back-projected into a shared metric frame (axes in real meters), yielding a colored point cloud, a smoothed camera trajectory, and the camera frusta of the three views. The pipeline outputs metric depth, camera-to-world poses, camera intrinsics, and scale and quality diagnostics.

3.3 Identity Annotation for Dynamic Subject Modeling

To support persistent dynamic-subject modeling, we construct an offline identity annotation pipeline that extracts object detections, segmentation masks, track identities, and multi-view reference images from each video. An offline YOLO-based pipeline [10] first produces temporally consistent object tracks for common dynamic categories. For third-person game videos, the controllable character is automatically identified according to tracking quality, temporal duration, spatial prominence, image-center proximity, and mask quality, eliminating the need for manual identity annotation.

Multi-view reference images are then extracted from diverse timestamps to maximize viewpoint and appearance diversity. During training, DINO features [18] are used to preferentially sample mutually dissimilar references, providing complementary identity and appearance cues for training lightweight multi-view Reference Tokens.

3.4 Scene-Quality Data Curation

Large-scale automatic annotation inevitably introduces inaccurate geometry and low-quality samples. We therefore perform quality-aware data curation before training. Each annotated clip is evaluated from four complementary perspectives: reconstruction quality, camera behavior, visual quality, and scene suitability. The scoring jointly considers reconstruction confidence, trajectory smoothness, image quality, depth validity, camera exploration, and content suitability.

To mitigate the bias introduced by automatic annotation, we additionally incorporate video length and camera motion as regularization terms, preventing short clips or nearly static videos from being over-favored. The final Scene-Quality score jointly evaluates all quality dimensions, and only clips satisfying all criteria are retained in the training corpus.

4 Experiments

4.1 Implementation Details

Camera- and Memory-augmented Base Model. We build Matrix-Game 3.5 upon Wan2.2-TI2V-5B [19] and jointly optimize camera conditioning and the patch memory in a single training stage. Camera intrinsics and extrinsics are injected in every DiT block through our PROPE variant adapted

to the backbone’s 3D spatiotemporal RoPE, while geometrically retrieved historical observations are fused into query-aligned memory latents and processed by the same self-attention stack. We fine-tune the full DiT while keeping the text encoder and VAE frozen. Training is conducted at a resolution of 1280×704 with 21 target latent frames per sample. Each sample contains one clean anchor latent and up to five historical context latents selected by trajectory coverage. With probability 0.8, the anchor is obtained by jointly encoding a 13-frame block; otherwise, a single-frame anchor is used. For each four-frame query group, we retrieve up to five historical candidates and fuse them into target-aligned memory using depth-aware z-buffering. The dense patch memory is enabled for 80% of training iterations and omitted for the remaining 20%, while PROPE camera conditioning is retained in both cases. Dynamic-object regions are filtered from the patch memory, and text conditioning is dropped with probability 0.1. We use AdamW with a constant learning rate of 5×10^{-5} , weight decay of 0.01, and gradient clipping at 0.5. Training uses BF16 precision, DeepSpeed ZeRO-2, gradient checkpointing, and a global batch size of 8.

Distillation Model. For perceptual causal adaptation, training is conducted at 1280×704 . Each causal training window spans 22 latent positions: one clean anchor followed by seven autoregressive chunks of three target latents each. Chunk i can attend to the anchor and preceding ground-truth chunks, whereas future chunks are masked. We query patch memory for every target latent, and one context frame is selected for each three-latent chunk. The memory and context tokens are drawn only from causally visible history and are exposed only to their corresponding target latent or chunk. We convert the predicted flow to a clean latent using the exact scheduler noise level, decode the complete causal chain with the frozen VAE, and average feature distances over all blocks of a frozen InternVideo2-1B encoder [20]. This stage is trained for 10,000 optimization steps using AdamW with a learning rate of 5×10^{-6} and a global batch size of 32 on 32 GPUs.

For self-rollout DMD, the student generates three latent frames per chunk with three denoising steps. At every substep, the rolling autoregressive prefix and selected context frame are noised to the next scheduler level, while the first anchor remains clean. On memory-conditioned iterations, each generated chunk is decoded and published to the online patch memory and context-frame memory before the next chunk is queried. We begin DMD with both memory conditions disabled and progressively increase the probability of enabling them together. The real scorer uses guidance scale 3, while inference uses one conditional student forward per denoising step (CFG= 1). We optimize the student and fake scorer with learning rates of 2×10^{-6} and 4×10^{-7} , respectively, using a global batch size of 64.

4.2 Quantitative Results

Evaluation Protocol and Metrics. Following the one-minute world-model benchmark introduced by SANA-WM [21], we evaluate minute-long generation on the Simple- and Hard-Trajectory splits. For pose accuracy, the estimated and target camera trajectories are aligned by a similarity transform before measuring rotation error (R, in degrees), translation error (T), and camera-motion consistency error (CMC), defined as the discrepancy between the full camera pose matrices; lower values indicate more faithful camera control. VBench Overall [22] aggregates visual quality, temporal consistency, and motion-related dimensions, with higher values being better. Efficiency is reported as peak memory and video throughput on eight H100 GPUs. Revisit consistency compares generated frame pairs captured from nearly identical camera poses: PSNR and SSIM [23] reward pixel and structural agreement, respectively, whereas LPIPS [24] measures perceptual distance. Finally, temporal degradation compares VBench imaging quality in the first and last 10-second windows, with $\Delta IQ = IQ_{0-10} - IQ_{50-60}$; a smaller drop indicates better long-horizon stability.

Comparison with State of the Art. As shown in Table 1, Matrix-Game 3.5 achieves the best camera accuracy on every pose metric and on both trajectory splits. On the Simple split, it obtains R/T/CMC errors of 1.63/1.10/1.11, reducing the rotation error of the next-best method from 4.50 to 1.63. The advantage becomes even more pronounced on the Hard split, where the rotation error is 2.70 compared with 8.34 for the runner-up, while T and CMC remain the best at 1.25 and 1.33. These results show that the generated camera follows the requested trajectory stably even under challenging long-horizon motion. This gain does not come at the expense of appearance quality: our VBench Overall scores are 80.14 and 80.85 on the two splits, competitive with the strongest baselines while generating at 720p.

Table 1: Quantitative comparison on our 1-min benchmark. Bold Res marks 720p. Pose Acc. reports R in degrees, plus T/CMC; VBench reports Overall score only. Revisit consistency is measured by PSNR/SSIM/LPIPS on same-pose pairs, and temporal degradation compares imaging quality between the first and last 10-second windows. Mem/Tput are peak GB and videos/hour on 8 H100s. Blue highlights mark top-three entries, with darker blue indicating a better rank.

Method	Param	Res	#G	Pose Acc. (↓)			VBench (↑)	Efficiency			Revisit consistency			Temporal degradation		
				R	T	CMC	Overall	Mem↓	Tput↑	PSNR↑	SSIM↑	LPIPS↓	IQ ₀₋₁₀ ↑	IQ ₅₀₋₆₀ ↑	ΔIQ↓	
Simple-Trajectory Split																
Infinite-World [25]	1.3B	480p	1	16.55	1.98	2.08	79.18	53.5	5.9	12.60	0.284	0.595	73.93	67.22	6.72	
LingBot-World [26]	14B+14B	480p	8	10.47	2.01	2.05	81.82	454.1	0.6	14.59	0.366	0.394	73.46	73.42	0.04	
HY-WorldPlay [27]	8B	480p	8	17.89	2.36	2.45	68.82	215.5	1.1	12.83	0.321	0.616	70.08	46.50	23.59	
Matrix-Game 3.0 [1]	5B	720p	8	12.96	1.83	1.92	78.53	106.2	3.1	12.29	0.326	0.553	69.07	66.66	2.41	
SANA-WM [21]	2.6B	720p	1	7.59	1.59	1.63	79.29	51.1	24.1	14.16	0.333	0.458	72.63	68.84	3.79	
SANA-WM + refiner [21]	2.6B+17B	720p	1	4.50	1.39	1.41	80.62	74.7	22.0	14.46	0.292	0.479	73.37	72.21	1.17	
Matrix-Game 3.5	5B	720p	1	1.63	1.10	1.11	80.14	77.0	10.9	14.56	0.439	0.439	73.80	70.56	3.24	
Hard-Trajectory Split																
Infinite-World [25]	1.3B	480p	1	41.31	2.49	2.84	79.51	53.5	5.9	12.04	0.248	0.617	73.79	69.63	4.16	
LingBot-World [26]	14B+14B	480p	8	18.99	1.65	1.81	81.89	454.1	0.6	14.08	0.332	0.436	73.66	73.09	0.58	
HY-WorldPlay [27]	8B	480p	8	35.46	2.34	2.64	70.46	215.5	1.1	13.72	0.328	0.654	70.21	44.33	25.88	
Matrix-Game 3.0 [1]	5B	720p	8	18.79	1.67	1.82	78.79	106.2	3.1	12.17	0.317	0.556	69.24	68.92	0.32	
SANA-WM [21]	2.6B	720p	1	10.02	1.66	1.72	79.60	51.1	24.1	14.10	0.327	0.469	72.58	69.49	3.09	
SANA-WM + refiner [21]	2.6B+17B	720p	1	8.34	1.39	1.44	81.89	74.7	22.0	14.80	0.312	0.458	73.34	73.03	0.31	
Matrix-Game 3.5	5B	720p	1	2.70	1.25	1.33	80.85	77.0	10.9	14.14	0.414	0.474	73.52	71.12	2.40	

Matrix-Game 3.5 also achieves the highest revisit SSIM on both splits (0.439 and 0.414) and near-best PSNR (14.56 and 14.14), indicating strong structural recall by the patch memory. Its LPIPS and PSNR do not dominate every setting. This is expected in our dynamic-world setting: revisiting the same camera pose does not imply an identical image because characters, vegetation, weather, and other background content may continue to evolve. Pixel-aligned metrics therefore penalize valid scene dynamics in addition to memory errors. The strong SSIM and pose results together provide a more faithful picture of geometric scene recall. Over one minute, the imaging-quality scores remain at 70.56 and 71.12 in the last 10-second window, with moderate drops of 3.24 and 2.40, further showing that visual quality does not exhibit severe cumulative degradation.

4.3 Qualitative Results

Figures 8 and 9 present long-horizon interactive rollouts in game-style and real-world environments, respectively. Across large viewpoint changes, the patch memory preserves recognizable scene structure and local appearance, while PRoPE produces smooth and stable camera motion. The late-stage frames remain sharp and coherent, and revisited content agrees closely with the initial observations, as highlighted by the red boxes. These results demonstrate that the same memory-augmented model supports persistent scene recall across both stylized game content and realistic visual domains.

Figure 10 further demonstrates prompt-controllable generation. Before each autoregressive rollout section, Gemma observes only current or past rollout context, using pose-ordered visual frames retrieved by the patch memory when available. It combines this context with the actual camera motion and a user-specified event to rewrite the complete video prompt for the next section, carrying the event forward naturally across sections. The four rows show, from top to bottom, the appearance of a volcano, a cat, a train, and a UFO. The generated videos retain the surrounding scene and controllable camera motion while responding clearly to these events. This shows that long-term memory and camera conditioning preserve rather than suppress the pretrained model’s prompt-following ability.

Figure 11 compares generation with and without the dynamic-subject reference prefix. We keep the initial state, camera trajectory, text prompt, and random seed fixed, and remove only the reference tokens at inference time. Without reference conditioning, the controllable subject gradually drifts in identity and local appearance under viewpoint changes and extended rollout. Enabling reference tokens preserves the subject’s characteristic appearance more consistently while leaving the background evolution and camera motion unchanged. This comparison supports the role of reference tokens as a persistent appearance memory for dynamic subjects.




Figure 8: **Long-horizon generation in game-style environments.** Columns show frames at 0, 15, 30, 45, and 60 seconds. Scene layout and local appearance remain coherent throughout extended camera motion. Red boxes in the 60-second frames mark regions that match content observed at 0 seconds, highlighting long-term recall enabled by memory.



Figure 9: **Long-horizon generation in real-world environments.** Columns show frames at 0, 15, 30, 45, and 60 seconds. The model maintains a realistic appearance and stable scene geometry across sustained camera motion. Red boxes in the 60-second frames mark regions that match content observed at 0 seconds, highlighting long-term recall enabled by memory.

Figure 12 shows long-horizon rollouts from the final three-step distilled model. The four examples span indoor architecture, waterfront settlements, autumn countryside, and snowy night scenes. Across approximately one minute of interactive camera motion, the distilled model retains sharp local detail and coherent scene geometry without abrupt scene resets. This demonstrates that our




Figure 10: **Prompt-controllable interactive generation.** Columns show frames at 0, 4, 8, 12, and 16 seconds. From top to bottom, the user events are the appearance of a volcano, a cat, a train, and a UFO. For each rollout section, Gemma uses current or past rollout context, with pose-ordered frames retrieved from the patch memory when available, and combines it with the camera motion and user event to rewrite a complete prompt for the next section. The event is carried forward across sections while scene coherence and camera control are preserved.




Figure 11: **Effect of dynamic-subject reference tokens.** We compare matched rollouts with and without the reference-token prefix while keeping all other conditions fixed. Reference conditioning reduces identity and appearance drift under camera motion over long-horizon generation (300–400 frames). In the first, second, and third rows, respectively, it better preserves the soldier’s tactical vest, the protagonist’s back-mounted weapon, and the vehicle’s silver side trim.

distillation pipeline preserves the base model’s long-horizon controllability under aggressive sampling acceleration.




Figure 12: **Long-horizon generation with the distilled model.** Columns show frames at 0, 15, 30, 45, and approximately 60 seconds. Rows cover four diverse indoor and outdoor environments. The three-step causal student follows sustained camera motion while preserving scene structure, appearance, and fine detail throughout the rollout.

4.4 Real-Time Inference

Following Matrix-Game 3.0 [1], we incorporate three additional inference optimizations: quantization for DiT, MG-LightVAE, and optimized memory retrieval. Specifically, INT8 quantization and PyTorch compilation accelerate DiT inference with negligible impact on generation quality. For memory retrieval, we reduce the number of depth candidates evaluated for each generated chunk and migrate key retrieval operations to the GPU. With DiT inference significantly accelerated, VAE decoding becomes the primary latency bottleneck in high-resolution streaming generation. To address this, we introduce a 75%-pruned MG-LightVAE variant that further improves end-to-end throughput while maintaining high reconstruction quality. Overall, these optimizations enable Matrix-Game 3.5 to achieve real-time interactive generation at up to 20 FPS on a single H200 GPU.

5 Related Work

5.1 Interactive Video World Models

Latent diffusion and Diffusion Transformer backbones provide the scalable visual priors behind modern video generators [28, 29]. Systems such as Sora, HunyuanVideo, and Wan demonstrate high-fidelity generation through large-scale data and model scaling [30–32], but are primarily designed for offline clip synthesis rather than closed-loop interaction. Interactive world models instead predict future observations conditioned on actions, camera motion, or high-level instructions. Genie and Genie 2 demonstrated action-controllable environment generation from video data [33, 34]; GameGen-X and GameFactory focus on interactive game-video generation [35, 36]; and OASIS and Matrix-Game 2.0 advanced open-source, real-time autoregressive interaction [37, 38]. More recent systems have pushed toward minute-scale and higher-resolution rollouts. WorldPlay introduces geometric memory for long-term camera-controllable generation [27], LingBot-World targets real-time minute-long simulation [26], SANA-WM improves the efficiency of 720p minute-scale generation with hybrid linear attention [21], and Matrix-Game 3.0 combines camera-aware memory, self-corrective training, and multi-segment DMD for real-time 720p streaming [1]. These systems establish the feasibility of long-horizon interactive generation, but jointly maintaining accurate cross-view recall, dynamic-subject identity, prompt-driven evolution, and real-time causal rollout remains difficult.

5.2 Long-Horizon Memory and Geometric Consistency

Long-horizon generation is limited by finite context windows and autoregressive error accumulation. Context as Memory retrieves historical observations as conditioning [39], while WorldMem and RELIC maintain learned or camera-aware latent memories for distant scene recall [2, 3]. Explicit geometric memories provide stronger spatial grounding: VMem indexes observations with surfels [4], and WorldPlay uses geometry-aware retrieval to improve revisitation consistency [27]. MosaicMem bridges explicit and implicit memory by lifting localized patches into 3D and composing them in the queried view, while leaving unseen regions generative [5]. Camera control has developed in parallel, from learned camera-conditioning modules to Projective Positional Encoding (PRoPE), which represents relative camera intrinsics and extrinsics directly in attention [6, 40]. Matrix-Game 3.5 builds on these directions by coupling patch-level geometric retrieval with a video-adapted PRoPE overlay, and by separating static scene memory from reference-token memory for dynamic subjects.

Dynamic-Subject Control and Static-Dynamic Memory. Object-controllable video generation has primarily followed two complementary directions. Trajectory-conditioned methods represent boxes, masks, or tracklets with instance-aware spatial tokens or control features, enabling local motion control and reducing identity swaps under occlusion [41–43]. Reference-conditioned methods instead extract appearance from one or more subject images and inject it through reference attention, shared token spaces, or reconstruction objectives to preserve identity across motion and viewpoint changes [44–46]. In interactive world models, recent methods additionally explore training-free concept insertion, persistent object-state tokens, explicit orchestration of 3D object trajectories, or even code based method [47–50]. These approaches improve either object motion control or long-range identity, but generally do not decide which observations should become part of a static geometric world and which should remain independently evolving entities. We make this separation explicit: a motion-aware 3D consistency filter admits only stable observations into Mosaic Memory, whereas the controllable subject is represented by compact multi-view reference tokens in the native, pose-aware DiT sequence. Leakage-resistant context masking and a subject-region objective further encourage the model to recover subject appearance from the reference memory rather than copy it from recent frames.

5.3 Real-Time Autoregressive Video Generation and Distillation

Streaming video generation reduces response latency by producing frames or chunks causally instead of denoising an entire clip at once. StreamingT2V uses short- and long-term memory for extendable generation [51], MAGI-1 scales chunk-wise autoregressive video diffusion [52], and LTX-Video explores efficient latent diffusion for real-time synthesis [53]. In the adjacent video restoration setting, FlashVSR demonstrates that one-step diffusion, sparse local attention, and a lightweight decoder can support real-time streaming video super-resolution [54]. These works motivate causal computation and systems-level optimization, but a high-quality bidirectional video model must still be transferred carefully into a few-step autoregressive generator.

Distribution Matching Distillation (DMD) enables one-step or few-step sampling by matching the student’s generated distribution to that of a diffusion teacher [12]. CausVid applies distribution matching to convert a bidirectional video diffusion model into a causal autoregressive generator [55]. Self-Forcing further trains on autoregressive student rollouts to reduce exposure bias, and Self-Forcing++ extends this strategy to minute-scale generation [56, 57]. Causal Forcing identifies the architectural mismatch between bidirectional teachers and causal students and initializes the student from an autoregressive teacher before DMD [7]. Causal Forcing++ replaces costly trajectory-based initialization with causal consistency distillation, enabling scalable frame-wise one- or two-step generation [8]. Our distillation pipeline follows this causal-distillation line while bringing geometry-conditioned memory into the autoregressive rollout: perceptual feature-space flow matching first adapts the causal student, and curriculum self-rollout DMD then trains the few-step model on its own long-horizon memory and context distribution.

6 Conclusion

We presented Matrix-Game 3.5, a real-time streaming interactive world model that advances the Matrix-Game series toward geometry-aware and long-horizon consistent world simulation. At its

core is a unified parameter-free geometry-aware memory framework: Patch Memory lifts historical observations into a persistent 3D patch representation and retrieves only the regions visible from the target viewpoint, while Warped PROPE overlays camera projection geometry onto the backbone’s spatiotemporal RoPE, so that long-term memory and camera control are handled by the same self-attention stack without any architectural modification. On top of this representation, a dynamic–static disentangled world modeling strategy stores stable scene structure in the static Patch Memory and carries movable subjects through lightweight multi-view Reference Tokens, suppressing ghosting and identity drift during extended rollouts. Finally, a progressive long-horizon distillation framework, combining perceptual feature-space Flow Matching with curriculum-based Self-Rollout DMD, converts the bidirectional world model into a few-step fully causal generator that sustains minute-long real-time interaction. Trained on a unified corpus spanning Unreal simulation environments, open-world games, and Internet videos, Matrix-Game 3.5 attains the best camera-pose accuracy on every metric of our benchmark, strong revisit consistency and subject preservation, and stable prompt-driven world evolution at 720p, demonstrating that explicitly unifying memory, geometry, and dynamic world representation is an effective foundation for persistent interactive world generation.

While Matrix-Game 3.5 substantially advances long-horizon interactive generation, several directions remain open. The current design persists scene structure through Patch Memory and dynamic-subject identity through Reference Tokens; a natural next step is to give dynamic entities persistent states of their own, so that subjects leaving the camera’s view keep evolving coherently and re-emerge seamlessly when the viewpoint returns. Likewise, the geometry grounding the patch memory currently comes from an external estimator; making the generator itself geometry-aware — jointly predicting the depth and camera pose of its own generated content — would let geometric grounding and generation reinforce each other. Extending the framework toward richer embodied action spaces and physical interaction, and scaling to larger backbones and longer horizons, are further promising directions. We hope Matrix-Game 3.5 serves as a solid foundation for future research on persistent, fully interactive world simulation.

A Text-Prompt Annotation Examples

Driving a vintage car



Frame 162 Frame 189 Frame 215 Frame 242

A black vintage car, viewed over its glossy reflective hood, drives forward along a wide asphalt road toward an intersection, the hood mirroring the sky and trees as it advances. In the middle ground, another black vintage car travels from left to right across the intersection. The background has a low stone wall with grey pillars, dense green trees, utility poles and power lines against a blue sky, and distant mountains. The road surface shifts in perspective with the car's motion.

Pedestrians on a shopping street



Frame 0 Frame 27 Frame 53 Frame 80

A woman in a bright pink jacket walks toward the camera along the sidewalk at a steady pace. Behind her, two women walk away — one in black, the other in maroon — while a man in dark clothing heads the same way to their left. Further back, several more pedestrians come and go along the pavement. On the street, a black car enters from the bottom right and moves left, partly blocking the view. On the far left, a person in a white shirt passes a storefront, and bicycles line the right sidewalk. Red-brick buildings with ground-floor storefronts line both sides of the street, Union Jack flags hang above the roadway, and a zebra crossing marks the foreground.


Galloping through a meadow



Frame 81 Frame 108 Frame 134 Frame 161

A first-person view from the back of a horse galloping through a lush green meadow toward a dense forest under a bright sky. The white horse, with a dark mane running between its ears, is the primary moving element: its ears twitch and its head bobs rhythmically with the gallop, conveying forward momentum. As the horse advances, the view shifts continuously along a straight path through the tall grass. The meadow is filled with green grass and scattered wildflowers, ringed by dark coniferous and deciduous trees, while fluffy white clouds drift across the blue sky.

Sunlit futuristic city canyon



Frame 0 Frame 27 Frame 53 Frame 80

The camera moves forward through a sunlit futuristic urban canyon of dense, multi-level buildings and neon signage. On the left, a steep grey road with a red border climbs past glowing neon signs, where two pedestrians in dark baggy clothing walk slowly away from the viewer. On the right, a concrete sidewalk runs past large green trash containers and a building with a pink 'Fuyutsuki' neon sign. A tall utility pole stands in the foreground as power lines crisscross the sky.

Figure 13: **Prompt-annotation examples.** Each case shows one annotation window as four evenly spaced frames (left to right) with its window-level prompt below. Across game and real-world sources, the prompt describes the static scene and names each dynamic object together with its motion.

B Preservation of RoPE’s Temporal Encoding under the Camera Overlay

Our injection overlays the camera transform onto the channels that already carry the spatiotemporal RoPE; this appendix verifies that the overlay re-weights, rather than overwrites, the temporal information those channels encode. The concern comes from how the two designs differ. Native P_{RoPE} was built for novel-view synthesis, where the inputs are unordered multi-view images and the positional code carries no time axis; video generation instead depends on the frame ordering encoded along the temporal axis of RoPE, so sharing one set of channels risks the camera term corrupting the temporal one.

We quantify the interaction with a matched-content probe: two frames are assigned identical content, and we read the attention score they exchange, so that any difference in the score is attributable purely to the positional and camera encodings rather than to image content. Figure 14 shows the resulting attention kernels on two controlled trajectories. The overlay leaves the temporal structure of the RoPE-only kernel essentially untouched, and modulates attention only through relative viewpoint — raising it between frames that see the scene from similar directions, lowering it between frames that do not. The temporal code is thus re-weighted by viewpoint, not overwritten.




Figure 14: P_{RoPE} and the native RoPE act together in one attention kernel: temporal structure is preserved, and pose similarity is rewarded. We use two synthetic trajectories here to illustrate the preservation (bottom row), and calculate frame-by-frame matched-content attention logits for heatmap visualization (top row). (a) RoPE only (P_{RoPE} off): the purely temporal kernel, banded in frame offset. (b) Straight walk with P_{RoPE} overlaid: the map is nearly indistinguishable from (a) — the maximum change is empirically 1.3% of (a)’s dynamic range — adding the camera does not disrupt the frame-to-frame attention structure. (c) S-curve with P_{RoPE} overlaid: frames with similar heading light up in blocks far from the diagonal, and opposite-heading frames darken. The camera raises attention exactly where views are geometrically close, while the temporal ordering of (a) survives untouched. This analysis was performed on the model fine-tuned from Wan 2.2 TI2V 5B

C Camera Collision Avoidance via a Progressive 3D Occupancy Map

Matrix-Game 3.5 uses the camera trajectory as its generation control signal: given a sequence of camera poses, the model faithfully advances along that trajectory, while the control signal itself does not reason about collisions between the trajectory and the scene structure. When a planned trajectory passes into an already-generated wall, object, or other scene surface, the model must still keep generating the content behind the obstacle and complete new scene structure for that region. Because this newly generated structure is only weakly constrained by the geometric surfaces observed

so far, it is prone to 3D inconsistency; such inconsistency further accumulates section by section over the subsequent rollout and can eventually cause the generation to collapse. To address this, we introduce a camera collision-avoidance mechanism based on a progressive 3D occupancy map in our patch-memory long-video evaluation, so that camera motion can adapt to the scene structure that has already been generated.

At the beginning of each section, the system estimates metric depth and camera pose for the anchor (or for the generated frames of the previous section), back-projects the depth pixels into a unified world coordinate frame, and aggregates them into 3D voxels according to their spatial positions; once a voxel accumulates enough surface observations, it is marked as occupied. Section 0 initializes the occupancy map from the anchor, and each subsequent section keeps adding geometric observations from previously generated content, so that the map expands progressively as the rollout proceeds and forms a cumulative representation of the observed scene structure.

Building on this cumulative scene representation, we apply a frame-by-frame geometric correction to the originally planned camera trajectory before generating each section. The planned trajectory is first aligned to the registered frame in which the occupancy map lives. Then, starting from the corrected camera center of the previous frame, the current displacement is decomposed into a radial component pointing toward the nearby occupied region and a tangential component moving along its boundary. The correction suppresses the radial component that would drive the camera further toward the obstacle while preserving the tangential component and any component that moves away from the obstacle; when a candidate camera center enters a preset safety margin, it is further pushed back to the safety boundary along the direction away from the occupied region. In this way, camera motion that originally headed straight into a scene surface is turned into a trajectory that advances laterally along the scene structure. The whole procedure adjusts only the camera position, keeps the original rotation and camera intrinsics unchanged, and recomputes the translation term of the world-to-camera extrinsics from the new camera center.

The corrected trajectory is used for both the patch-memory gather and PROPE, keeping the geometric reprojection of the historical frames, the camera-position conditioning, and the current generation trajectory mutually consistent. Once the current section has been generated, its frames are re-registered with depth and pose estimation in the next section and continue to update the 3D occupancy map, forming a closed loop of scene registration, occupancy update, trajectory correction, and video generation (Figure 15). This method lets camera motion adaptively adjust using the scene geometry that keeps accumulating during rollout, while maintaining spatial consistency between the patch memory and the camera conditioning.




Figure 15: **Camera collision avoidance via a progressive 3D occupancy map.** (a) The generated frames of a section, with the planned input trajectory (blue) heading straight toward the statue. (b) Metric depth and camera pose lift the observed surfaces into a unified 3D world frame and aggregate them into voxels. (c) Voxels with enough surface observations are marked as occupied (red); the planned trajectory (blue) that would pass into this occupied region is corrected into an adjusted trajectory (green) that advances laterally along the scene structure while staying outside a preset safety margin.

References

- [1] Zile Wang, Zexiang Liu, Jiaying Li, Kaichen Huang, Baixin Xu, Fei Kang, Mengyin An, Peiyu Wang, Biao Jiang, Yichen Wei, et al. Matrix-game 3.0: Real-time and streaming interactive world model with long-horizon memory. *arXiv preprint arXiv:2604.08995*, 2026.

- [2] Yicong Hong, Yiqun Mei, Chongjian Ge, Yiran Xu, Yang Zhou, Sai Bi, Yannick Hold-Geoffroy, Mike Roberts, Matthew Fisher, Eli Shechtman, et al. Relic: Interactive video world model with long-horizon memory. *arXiv preprint arXiv:2512.04040*, 2025.
- [3] 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.
- [4] Runjia Li, Philip Torr, Andrea Vedaldi, and Tomas Jakab. Vmem: Consistent interactive video scene generation with surfel-indexed view memory. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 25690–25699, 2025.
- [5] Wei Yu, Runjia Qian, Yumeng Li, Liquan Wang, Songheng Yin, Dennis Anthony, Yang Ye, Yidi Li, Weiwei Wan, Animesh Garg, et al. Mosaicmem: Hybrid spatial memory for controllable video world models. *arXiv preprint arXiv:2603.17117*, 2026.
- [6] Ruilong Li, Brent Yi, Junchen Liu, Hang Gao, Yi Ma, and Angjoo Kanazawa. Cameras as relative positional encoding. *Advances in Neural Information Processing Systems*, 38:15984–16009, 2026.
- [7] Hongzhou Zhu, Min Zhao, Guande He, Hang Su, Chongxuan Li, and Jun Zhu. Causal forcing: Autoregressive diffusion distillation done right for high-quality real-time interactive video generation. *arXiv preprint arXiv:2602.02214*, 2026.
- [8] Min Zhao, Hongzhou Zhu, Kaiwen Zheng, Zihan Zhou, Bokai Yan, Xinyuan Li, Xiao Yang, Chongxuan Li, and Jun Zhu. Causal forcing++: Scalable few-step autoregressive diffusion distillation for real-time interactive video generation. *arXiv preprint arXiv:2605.15141*, 2026.
- [9] Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *CoRR*, abs/2104.09864, 2021.
- [10] Ranjan Sapkota and Manoj Karkee. Ultralytics yolo evolution: An overview of yolo26, yolo11, yolov8 and yolov5 object detectors for computer vision and pattern recognition. *arXiv preprint arXiv:2510.09653*, 2025.
- [11] Chuyang Zhao, Yifei Song, Hongfa Wang, Jianlong Yuan, Yuan Zhang, Siming Fu, Zhineng Chen, Huilin Deng, Haoyang Huang, and Nan Duan. Perceptual flow matching for few-step generative modeling. *arXiv preprint arXiv:2607.03524*, 2026.
- [12] Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6613–6623, 2024.
- [13] Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and Bill Freeman. Improved distribution matching distillation for fast image synthesis. *Advances in neural information processing systems*, 37:47455–47487, 2024.
- [14] Kai Zou, Dian Zheng, Hongbo Liu, Tiankai Hang, Bin Liu, and Nenghai Yu. Hiar: Efficient autoregressive long video generation via hierarchical denoising. *arXiv preprint arXiv:2603.08703*, 2026.
- [15] Jianyuan Wang, Minghao Chen, Shangzhan Zhang, Nikita Karaev, Johannes Schönberger, Patrick Labatut, Piotr Bojanowski, David Novotny, Andrea Vedaldi, and Christian Rupprecht. VGGT- Ω . *arXiv preprint arXiv:2605.15195*, 2026.
- [16] Haotong Lin, Sili Chen, Junhao Liew, Donny Y Chen, Zhenyu Li, Guang Shi, Jiashi Feng, and Bingyi Kang. Depth anything 3: Recovering the visual space from any views. *arXiv preprint arXiv:2511.10647*, 2025.
- [17] Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024.
- [18] Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, Francisco Massa, Daniel Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana, Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé Jégou, Patrick Labatut, and Piotr Bojanowski. DINOv3, 2025.
- [19] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, Jianyuan Zeng, Jiayu Wang, Jingfeng Zhang, Jingren Zhou, Jinkai Wang, Jixuan Chen, Kai Zhu, Kang Zhao, Keyu Yan, Lianghua Huang, Mengyang Feng, Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang, Tianxing Wang, Tianyi Gui, Tingyu Weng, Tong Shen, Wei Lin, Wei Wang, Wei Wang, Wenmeng Zhou, Wenten Wang, Wenting Shen, Wenyan Yu, Xianzhong Shi, Xiaoming Huang, Xin Xu, Yan Kou, Yangyu Lv, Yifei Li, Yijing Liu, Yiming Wang, Yingya Zhang, Yitong Huang, Yong Li, You Wu, Yu Liu, Yulin Pan, Yun Zheng, Yuntao Hong, Yupeng Shi, Yutong Feng, Zeyinzi Jiang, Zhen Han, Zhi-Fan Wu, and Ziyu Liu. Wan: Open and advanced large-scale video generative models, 2025.

- [20] Yi Wang, Kunchang Li, Xinhao Li, Jiashuo Yu, Yinan He, Chenting Wang, Guo Chen, Baoqi Pei, Ziang Yan, Rongkun Zheng, Jilan Xu, Zun Wang, Yansong Shi, Tianxiang Jiang, Songze Li, Hongjie Zhang, Yifei Huang, Yu Qiao, Yali Wang, and Limin Wang. Internvideo2: Scaling foundation models for multimodal video understanding. *arXiv preprint arXiv:2403.15377*, 2024.
- [21] Haoyi Zhu, Haozhe Liu, Yuyang Zhao, Tian Ye, Junsong Chen, Jincheng Yu, Tong He, Song Han, and Enze Xie. Sana-wm: Efficient minute-scale world modeling with hybrid linear diffusion transformer. *arXiv preprint arXiv:2605.15178*, 2026.
- [22] 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 *Computer Vision and Pattern Recognition*, pages 21807–21818, 2024.
- [23] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [24] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.
- [25] Zelin Gao, Qiuyu Wang, Jiapeng Zhu, Jingye Chen, Zichen Liu, Qingyan Bai, Jiahao Wang, Yufeng Yuan, Hanlin Wang, Yichong Lu, et al. Infinite worlds with versatile interactions. *arXiv preprint arXiv:2607.07534*, 2026.
- [26] Robbyant Team, Zelin Gao, Qiuyu Wang, Yanhong Zeng, Jiapeng Zhu, Ka Leong Cheng, Yixuan Li, Hanlin Wang, Yinghao Xu, Shuailei Ma, et al. Advancing open-source world models. *arXiv preprint arXiv:2601.20540*, 2026.
- [27] Wenqiang Sun, Haiyu Zhang, Haoyuan Wang, Junta Wu, Zehan Wang, Zhenwei Wang, Yunhong Wang, Jun Zhang, Tengfei Wang, and Chunchao Guo. Worldplay: Towards long-term geometric consistency for real-time interactive world modeling. *arXiv preprint arXiv:2512.14614*, 2025.
- [28] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [29] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *International Conference on Computer Vision*, pages 4195–4205, 2023.
- [30] OpenAI. Sora: Video generation models as world simulators. <https://openai.com/index/video-generation-models-as-world-simulators/>, 2024.
- [31] 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.
- [32] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Fei Wu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025.
- [33] 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 *International Conference on Machine Learning*, 2024.
- [34] Jack Parker-Holder, Philip Ball, Jake Bruce, Vibhavari Dasagi, Kristian Holsheimer, Christos Kaplanis, Alexandre Moufaret, Guy Scully, Jeremy Shar, Jimmy Shi, Stephen Spencer, Jessica Yung, Michael Dennis, Sultan Kenjeyev, Shangbang Long, Vlad Mnih, Harris Chan, Maxime Gazeau, Bonnie Li, Fabio Pardo, Luyu Wang, Lei Zhang, Frederic Besse, Tim Harley, Anna Mitenkova, Jane Wang, Jeff Clune, Demis Hassabis, Raia Hadsell, Adrian Bolton, Satinder Singh, and Tim Rocktäschel. Genie 2: A large-scale foundation world model. 2024.
- [35] 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.
- [36] Jiwen Yu, Yiran Qin, Xintao Wang, Pengfei Wan, Di Zhang, and Xihui Liu. Gamefactory: Creating new games with generative interactive videos. In *International Conference on Computer Vision*, 2025.
- [37] Decart. Oasis: A universe in a transformer. 2024.
- [38] 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.
- [39] 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.

- [40] 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.
- [41] Jiawei Wang, Yuchen Zhang, Jiaxin Zou, Yan Zeng, Guoqiang Wei, Liping Yuan, and Hang Li. Boximator: Generating rich and controllable motions for video synthesis. *arXiv preprint arXiv:2402.01566*, 2024.
- [42] Pengxiang Li, Kai Chen, Zhili Liu, Ruiyuan Gao, Lanqing Hong, Dit-Yan Yeung, Huchuan Lu, and Xu Jia. Trackdiffusion: Tracklet-conditioned video generation via diffusion models. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 3539–3548. IEEE, 2025.
- [43] Guofeng Zhang, Angtian Wang, Jacob Zhiyuan Fang, Liming Jiang, Haotian Yang, Bo Liu, Yiding Yang, Guang Chen, Longyin Wen, Alan Yuille, et al. Tgt: Text-grounded trajectories for locally controlled video generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22028–22037, 2026.
- [44] Yujie Wei, Shiwei Zhang, Hangjie Yuan, Xiang Wang, Haonan Qiu, Rui Zhao, Yutong Feng, Feng Liu, Zhizhong Huang, Jiaxin Ye, et al. Dreamvideo-2: Zero-shot subject-driven video customization with precise motion control. *arXiv preprint arXiv:2410.13830*, 2024.
- [45] Ziyang Mai and Yu-Wing Tai. Contextanyone: Context-aware diffusion for character-consistent text-to-video generation. *arXiv preprint arXiv:2512.07328*, 2025.
- [46] Xuan Wei, Longbin Ji, Guan Wang, Xiangrui Liu, Zhenyu Zhang, Shuohuan Wang, Yu Sun, and Qingqi Hong. Memento: Reconstruct to remember for consistent long video generation. *arXiv preprint arXiv:2606.14667*, 2026.
- [47] Kiyomet Akdemir and Pinar Yanardag. From zero to hero: Training-free custom concept spawning in world models. *arXiv preprint arXiv:2606.02575*, 2026.
- [48] Hanlin Wang, Hao Ouyang, Qiuyu Wang, Wen Wang, Qingyan Bai, Ka Leong Cheng, Yue Yu, Yixuan Li, Yihao Meng, Zichen Liu, et al. Worlddirector: Building controllable world simulators with persistent dynamic memory. *arXiv preprint arXiv:2607.02517*, 2026.
- [49] Zhexiao Xiong, Yizhi Song, Hao Kang, Qing Yan, Liming Jiang, Jenson Yang, Zhoujie Fu, Stathi Fotiadis, Angtian Wang, Zichuan Liu, et al. Actworld: From explorable to interactive world model via action-aware memory. *arXiv preprint arXiv:2606.17730*, 2026.
- [50] Puyi Wang, Yuhao Wang, Linjie Li, Zhengyuan Yang, Kevin Qinghong Lin, Yangguang Li, and Yu Cheng. Scenecode: Executable world programs for editable indoor scenes with articulated objects. *arXiv preprint arXiv:2605.19587*, 2026.
- [51] Roberto Henschel, Levon Khachatryan, Hayk Poghosyan, Daniil Hayrapetyan, Vahram Tadevosyan, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Streamingt2v: Consistent, dynamic, and extendable long video generation from text. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 2568–2577, 2025.
- [52] Hansi Teng, Hongyu Jia, Lei Sun, Lingzhi Li, Maolin Li, Mingqiu Tang, Shuai Han, Tianning Zhang, WQ Zhang, Weifeng Luo, et al. Magi-1: Autoregressive video generation at scale. *arXiv preprint arXiv:2505.13211*, 2025.
- [53] Yoav HaCohen, Nisan Chiprut, Benny Brazowski, Daniel Shalem, Dudu Moshe, Eitan Richardson, Eran Levin, Guy Shiran, Nir Zabari, Ori Gordon, Poriya Panet, Sapir Weissbuch, Victor Kulikov, Yaki Bitterman, Zeev Melumian, and Ofir Bibi. Ltx-video: Realtime video latent diffusion. *arXiv preprint arXiv:2501.00103*, 2024.
- [54] Junhao Zhuang, Shi Guo, Xin Cai, Xiaohui Li, Yihao Liu, Chun Yuan, and Tianfan Xue. Flashvsr: Towards real-time diffusion-based streaming video super resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 43482–43493, 2026.
- [55] 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.
- [56] Xun Huang, Zhengqi Li, Guande He, Mingyuan Zhou, and Eli Shechtman. Self forcing: Bridging the train-test gap in autoregressive video diffusion. *arXiv preprint arXiv:2506.08009*, 2025.
- [57] 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.