

# SlotMemory: Object-Centric KV Memory for Streaming Long-Video Generation

WEIJIA DOU, Fudan University, China

HUI LI, Fudan University, China

JIAHAO CUI, Fudan University, China

LEI ZHOU, Meta super intelligence, China

JINGDONG WANG, Baidu, China

SIYU ZHU\*, Fudan University, China

Streaming video generation models typically rely on temporal-centric memory, which organizes historical context as raw frames, chunk segments, or unclustered tokens. This organization frequently leads to identity drift and semantic inconsistency when entities exit the frame or during interactive prompt transitions. To address these limitations, we propose SlotMemory, an object-centric Key-Value (KV) memory mechanism for streaming video diffusion. Our approach shifts the memory abstraction from “when” an event occurred to “what” is being represented by decomposing the transformer’s KV manifold into discrete, reusable semantic slots. By utilizing these slots as routing addresses to index and store high-fidelity KV tokens, we enable entity-level persistence and prompt-aware retrieval across long horizons. Evaluated on 60-second interactive narratives using the Wan2.1-T2V-1.3B backbone, SlotMemory achieves a state-of-the-art quality score of 81.61 and a 22.8% relative improvement in dynamic consistency over the strongest existing streaming baseline. Our results demonstrate that structured semantic representation, rather than raw temporal capacity, is the essential primitive for persistent long-form video synthesis. Our codes and checkpoints are available in <https://tj12323.github.io/SlotMemory/>.

CCS Concepts: • **Computing methodologies** → **Computer vision; Artificial intelligence**.

Additional Key Words and Phrases: Image & Video Generative AI

## 1 Introduction

The frontier of generative video modeling has rapidly evolved from synthesizing isolated, short-duration clips [Chen et al. 2025; Ha-Cohen et al. 2024; Wan et al. 2025] toward the more formidable objective of producing long-form, persistent narratives. Current streaming frameworks [Henschel et al. 2025; Huang et al. 2026; Yang et al. 2026a] address this challenge via the autoregressive processing of video chunks, utilizing a sliding window of historical context to facilitate extended rollouts. However, these architectures typically rely on a temporal-centric memory organization, where history is represented as unclustered tokens, raw frames, or discrete chunk segments [Gu et al. 2025; Kim et al. 2024; Yin et al. 2025]. This approach introduces a fundamental bottleneck: by treating memory as a linear sequence of temporal slices rather than a structured representation of the scene, the models neglect the semantic hierarchy of the visual field. Consequently, they frequently fail to maintain identity coherence when entities exit and re-enter the frame or when the synthesis requires long-range interaction.

\*Corresponding author.

Authors’ Contact Information: Weijia Dou, Fudan University, Shanghai, China; Hui Li, Fudan University, Shanghai, China; Jiahao Cui, Fudan University, Shanghai, China; Lei Zhou, Meta super intelligence, China; Jingdong Wang, Baidu, China; Siyu Zhu, Fudan University, Shanghai, China. [siyuzhu@fudan.edu.cn](mailto:siyuzhu@fudan.edu.cn).

This temporal-centric memory organization precipitates two fundamental failure modes in long-horizon video synthesis. First, it undermines entity persistence; when a character or object exits the local temporal window and subsequently re-enters, the model often fails to recover its unique attributes, resulting in visual drift or identity switching. Second, existing architectures are ill-suited for interactive, multi-prompt generation. When a prompt transitions—for example, from a brown dog playing in a backyard to returning to a porch—the model must distinguish between ephemeral motion context that should be discarded and persistent semantic entities that must be retained. Existing approaches, which predominantly rely on uniform temporal decay or prompt-aware retrieval [Ji et al. 2025; Yang et al. 2026b; Yu et al. 2025], lack the structural awareness to bridge these semantic gaps. Consequently, they suffer from a cumulative loss of coherence, leading to catastrophic narrative degradation over extended 60-second rollouts.

To address these limitations, we propose **SlotMemory**, an object-centric KV memory mechanism designed to enforce long-term semantic consistency in streaming video diffusion. Our core insight is a shift in memory abstraction from temporal occurrence (“when”) to semantic identity (“what”). We achieve this by decomposing the transformer’s KV manifold into discrete, reusable semantic slots via slot-based representation learning [Locatello et al. 2020; Manasyan et al. 2025]. Figure 1 presents long-form examples where subject identity, key props, and action flow remain coherent across scene changes and temporary occlusions. Unlike traditional object-centric models optimized for unsupervised discovery in simplified environments [Burgess et al. 2019; Greff et al. 2019], SlotMemory bridges structured scene decomposition with high-resolution diffusion. Specifically, it utilizes these slots as routing addresses to index and store original, transformer-compatible KV tokens. This hybrid architecture preserves the high-fidelity detail essential for diffusion while providing a structured, entity-level index for efficient retrieval across extended horizons.

Architecturally, the technical framework of SlotMemory centers on a recursive read-write-update cycle that maintains a persistent semantic memory bank within fixed computational constraints. During the writing phase, we leverage a temporally-initialized Slot Attention mechanism [Manasyan et al. 2025] to partition high-dimensional transformer manifolds into discrete, reusable semantic regions. To facilitate selective retrieval, we introduce a prompt-aware scoring function that calibrates textual alignment against visual relevance; this mechanism ensures that core entities are preserved across narrative transitions while obsolete temporal context



Fig. 1. **Teaser examples of long-form subject persistence generated by SlotMemory.** Each row shows six keyframes from one sequence, covering scene transitions, action changes, and temporary occlusion. Across three different narratives (cat, rider, dog), the main subject identity and key visual anchors remain consistent over time.

is evicted. By embedding this module directly into the attention stack of the diffusion transformer, our architecture enables the model to attend to structured entities rather than unstructured frames, fundamentally stabilizing the synthesis of complex, multi-stage narratives.

We evaluate SlotMemory using the Wan2.1-T2V-1.3B backbone across a comprehensive suite of benchmarks, including a challenging 60-second interactive multi-prompt scenario. Experimental results demonstrate that shifting toward an object-centric memory abstraction yields significant gains in long-term consistency and prompt adherence over state-of-the-art streaming baselines [Ji et al. 2025; Yang et al. 2026a; Yesiltepe et al. 2025]. Specifically, SlotMemory attains a state-of-the-art quality score of 81.61 in interactive narratives and a 74.29 dynamic score at the 30-second horizon—a 22.8% relative improvement over the existing approach. By successfully mitigating cumulative feature drift across 60-second rollouts, this work establishes that structured semantic representation, rather than raw temporal capacity, is the fundamental primitive for persistent and high-fidelity long-form video synthesis.

## 2 Related Work

**Long and Streaming Video Generation.** The frontier of video synthesis has rapidly evolved from the generation of isolated, short-duration clips [Bar-Tal et al. 2024; Blattmann et al. 2023; Chen et al. 2025; Guo et al. 2023; HaCohen et al. 2024; Ho et al. 2022a,b; Jin et al. 2025; Singer et al. 2022; Wan et al. 2025; Wang et al. 2023; Wu et al. 2023] to the production of extended, persistent narratives via streaming frameworks. These architectures typically employ autoregressive rollouts to synthesize video chunks sequentially while managing state carry-over under fixed computational budgets [Dalal et al. 2025; Fang et al. 2025; Gu et al. 2025; Henschel et al. 2025; Huang et al. 2026; Kim et al. 2024; Teng et al. 2025; Yesiltepe et al. 2025; Yi et al. 2025; Yin et al. 2025; Zhang et al. 2025a]. To enforce long-range consistency, contemporary streaming methods utilize explicit historical memory banks coupled with prompt-aware retrieval [Cai et al. 2025; Ji et al. 2025; Yang et al. 2026a,b; Yu et al. 2025]. However, a fundamental limitation remains: these approaches uniformly organize memory as a sequence of discrete temporal slices—such

as raw frames, chunk segments, or unclustered tokens—prioritizing local temporal smoothness over global semantic structure. This temporal-centric organization often precludes robust entity-level reusability, leading to identity drift when objects exit the frame or during complex prompt transitions. Our method addresses this bottleneck by shifting the memory abstraction from frame-centric storage to slot-conditioned semantic storage, enabling visual entities to persist independently of the local temporal window.

**Object-Centric Representation Learning.** Object-centric representation learning aims to decompose complex visual scenes into discrete, reusable latent components rather than monolithic global features. Building upon foundational slot-attention mechanisms [Burgess et al. 2019; Engelcke et al. 2020; Greff et al. 2019; Locatello et al. 2020], recent advancements have extended these frameworks to handle intricate temporal dynamics, 3D consistency, and naturalistic video sequences [Chen et al. 2024; Didolkar et al. 2025; Elsayed et al. 2022; Jabri et al. 2024; Kipf et al. 2021; Kori et al. 2024a,b; Majellaro et al. 2025; Manasyan et al. 2025; Rubinstein et al. 2025; Seitzer et al. 2022; Wu et al. 2022; Yang et al. 2025; Zhang et al. 2025b; Zhao et al. 2024]. These developments establish that semantic components serve as more robust primitives for persistent visual states than raw full-frame snapshots. However, a significant gap remains: while prior research focuses on representation quality or unsupervised discovery, these models are not designed to function as transformer-compatible, long-term KV memory for high-resolution diffusion. Conversely, state-of-the-art streaming video generators lack explicit object-centric primitives, relying instead on unclustered temporal histories [Cai et al. 2025; Ji et al. 2025; Yang et al. 2026a,b; Yu et al. 2025]. Our method bridges this divide by utilizing semantic slots as explicit routing addresses to index and store original KV tokens, thereby integrating structured scene decomposition directly into the attention stack of modern diffusion transformers.

## 3 Method

We frame long-form, multi-prompt video synthesis as a bounded-memory streaming generation problem and propose SlotMemory,

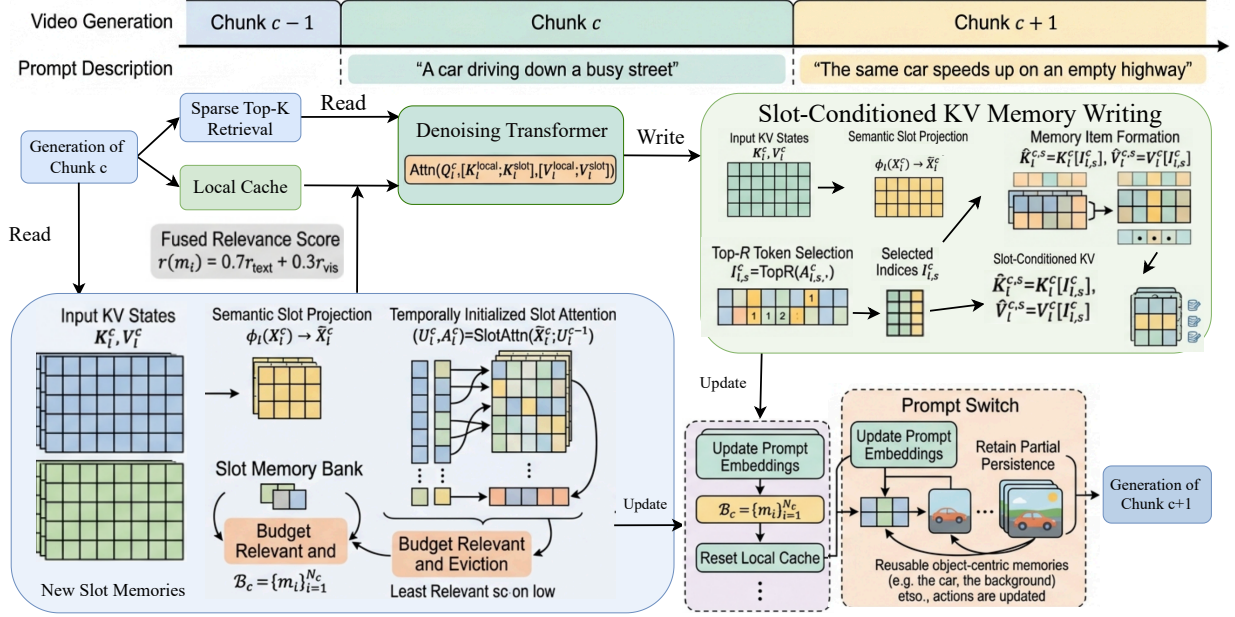


Fig. 2. **Pipeline of SlotMemory in streaming long-video generation.** For each chunk, the model retrieves slot-indexed long-term KV memory, denoises the current chunk using both local and retrieved context, writes new slot-conditioned memory items from current transformer states, and updates the memory bank under a fixed budget. At prompt switches, prompt-aware retrieval retains reusable entities while allowing the generation to adapt to new instructions.

an object-centric KV memory mechanism designed to enforce long-range semantic consistency. A fundamental challenge in streaming architectures is determining which historical information should persist once it exits the local temporal window. Unlike conventional methods that organize memory via raw frames, chunk segments, or unclustered tokens, SlotMemory indexes transformer-compatible KV tokens using discrete semantic slots. This abstraction decouples entity persistence from temporal occurrence, allowing the model to retrieve scene regions based on “what” they represent rather than “when” they appeared. The remainder of this section is organized as follows: Section 3.1 defines the streaming problem setup; Section 3.2 provides an architectural overview of the read-write-update pipeline; Section 3.3 details the slot-conditioned memory writing process; Section 3.4 describes memory retrieval, eviction, and prompt-aware adaptation; and Section 3.5 outlines the training objective and slot regularization.

### 3.1 Problem Setup and Streaming Framework

We formulate long-form, multi-prompt video synthesis as a bounded-memory streaming generation problem. Given a sequence of  $M$  text prompts  $\mathcal{P} = \{p^{(1)}, p^{(2)}, \dots, p^{(M)}\}$  and their associated switch points  $\mathcal{S} = \{s_1, \dots, s_{M-1}\}$ , the objective is to synthesize a video  $V_{1:T}$  in an autoregressive manner. The generation proceeds across  $C$  discrete chunks of length  $T_c$ , where  $c \in \{1, \dots, C\}$  denotes the current chunk index. To ensure temporal and semantic continuity, the generator maintains three distinct states: a short-term local KV cache  $C_c$  for local context, a text-conditioning cache  $C_c^{\text{text}}$  for prompt-dependent features, and a long-term memory bank  $\mathcal{B}_c$  for persistent history.

The primary challenge in this streaming framework is the management of the long-term bank  $\mathcal{B}_c$  under a fixed capacity constraint  $N_{\text{budget}}$ . The system must selectively retain, retrieve, and evict historical states to maximize global consistency and prompt adherence while operating within bounded computational limits. This optimization problem is formulated as follows:

$$\max_{\mathcal{B}_c: |\mathcal{B}_c| \leq N_{\text{budget}}} \mathbb{E}[\text{Consistency}(V_{1:T}) + \text{PromptAdherence}(V_{1:T}, \mathcal{P}) - \lambda \text{Cost}(\mathcal{B}_c)]. \quad (1)$$

We build upon a causal streaming diffusion backbone where each chunk is denoised conditioned on the active prompt and the retrieved historical context. Our approach is modular and agnostic to the specific backbone architecture; we replace the conventional unclustered temporal memory with a structured, object-centric memory bank. By representing history as reusable semantic slots rather than raw temporal slices, the framework decouples entity persistence from the local temporal window, facilitating stable synthesis across extended horizons and complex prompt transitions.

### 3.2 SlotMemory Overview

Conventional streaming frameworks represent historical context through unclustered temporal slices, such as raw frames or chunk-level KV tokens. This organization entangles foreground entities, background geometry, and transient motion within a single temporal unit, which precludes the stable preservation of identities across temporal gaps or prompt transitions. To resolve this bottleneck, SlotMemory redefines the fundamental memory primitive as a

persistent semantic region rather than a temporal frame. By decomposing the transformer’s KV manifold into discrete, reusable slots, the architecture facilitates entity-level addressing while maintaining the high-fidelity token representations necessary for diffusion synthesis.

The system operates via a recursive read-write-update cycle synchronized with the autoregressive generation of video chunks. At each step, the model retrieves semantically relevant items from the memory bank to augment the local KV cache. The diffusion transformer then denoises the current chunk by attending to both local and retrieved contexts. Following generation, new slot-conditioned items are written to the bank, and a relevance-driven eviction policy manages capacity under a fixed budget. This object-centric abstraction is particularly robust during interactive prompt switches; while local motion context is refreshed, the underlying semantic slots allow core visual entities to persist independently of the temporal window. The end-to-end integration of these modules within the streaming pipeline is illustrated in Figure 2.

### 3.3 Slot-Conditioned KV Memory Writing

To transform unclustered transformer manifolds into structured semantic units, we introduce a slot-conditioned writing mechanism that partitions visual tokens into discrete entity regions. For a specific layer  $l$  and chunk  $c$ , let  $K_l^c, V_l^c \in \mathbb{R}^{B \times L \times H \times D}$  represent the key and value tensors, respectively. We derive semantic assignments by first projecting the corresponding token features  $X_l^c$  into a compact latent space via  $\tilde{X}_l^c = \phi_l(X_l^c)$ , followed by the application of temporally-initialized Slot Attention:

$$(U_l^c, A_l^c) = \text{SlotAttn}(\tilde{X}_l^c; U_l^{c-1}), \quad (2)$$

Here,  $U_l^c \in \mathbb{R}^{B \times S \times d_s}$  denotes the set of  $S$  semantic slots, while  $A_l^c \in \mathbb{R}^{B \times S \times L}$  provides the soft assignment weights between slots and tokens. The temporal initialization using the previous slot state  $U_l^{c-1}$  enforces identity persistence across sequential chunks, ensuring that the same slot represents the same semantic entity over time.

A critical distinction of our approach is that the low-dimensional slot embeddings are utilized solely as semantic routing addresses rather than the primary storage medium. To preserve the high-frequency details necessary for high-fidelity diffusion, we utilize these slots to index and store original transformer tokens. Specifically, for each slot  $s$ , we identify the  $R$  tokens with the highest assignment weights,  $I_{l,s}^c = \text{TopR}(A_{l,s,\cdot}^c)$ , and extract their corresponding KV states defined as:

$$\hat{K}_l^{c,s} = K_l^c[I_{l,s}^c], \quad \hat{V}_l^{c,s} = V_l^c[I_{l,s}^c]. \quad (3)$$

We define the resulting memory item as  $m_{l,s}^c = (\hat{K}_l^{c,s}, \hat{V}_l^{c,s}, u_{l,s}^c, \eta_c)$ , where  $u_{l,s}^c$  serves as the retrieval key and  $\eta_c$  encapsulates temporal metadata. This hybrid architecture enables the diffusion transformer to attend to structured, object-centric history without sacrificing the numerical fidelity of the original feature space.

### 3.4 Memory Retrieval, Eviction, and Prompt Switching

Let  $\mathcal{B}_c = \{m_i\}_{i=1}^{N_c}$  denote the long-term memory bank. Before denoising chunk  $c$ , the model retrieves a subset of memory items whose semantic content aligns with the current prompt and visual state.

For each item  $m_i$ , we compute a fused relevance score:

$$r(m_i) = \gamma \cdot r_{\text{text}}(m_i, p_c) + (1 - \gamma) \cdot r_{\text{vis}}(m_i, q_c), \quad (4)$$

where  $p_c$  is the active prompt,  $q_c$  represents the current visual query or denoising state,  $r_{\text{text}}$  measures text-memory alignment via cosine similarity in the joint embedding space,  $r_{\text{vis}}$  measures visual compatibility, and  $\gamma$  regulates their relative contribution. The top- $k$  ranked items are then integrated into the attention mechanism:

$$\text{Attn}(Q_l^c, [K_l^{\text{local}}, K_l^{\text{slot}}], [V_l^{\text{local}}, V_l^{\text{slot}}]). \quad (5)$$

Following the generation of the current chunk, new slot-conditioned items are written to the bank,  $\mathcal{B}'_c = \mathcal{B}_{c-1} \cup \{m_{l,s}^c\}_{l,s}$ . To maintain the fixed capacity  $N_{\text{budget}}$ , the relevance function  $r(m_i)$  is utilized to rank all existing items, and the lowest-scoring entries are evicted. This dual-relevance policy is critical during prompt transitions. When the sequence reaches a switch point  $s_m$ , the text-conditioning cache  $C_c^{\text{text}}$  is refreshed to reflect the new instruction. While ephemeral motion context in the local cache may be updated, the persistent semantic bank uses the updated prompt  $p_{c+1}$  to recalibrate retrieval scores. This mechanism ensures that core entities—such as a specific character or environment—are selectively retained across narrative shifts while obsolete semantic context is suppressed, thereby facilitating long-term narrative coherence.

### 3.5 Training Objective and Slot Regularization

To align the memory module with its inference-time behavior, we train SlotMemory using a streaming rollout trajectory. In each training iteration, the model sequentially synthesizes video chunks, writes slot-conditioned features to the memory bank, and retrieves these states to condition subsequent generation steps. Training sequences incorporate prompt switches to force the module to learn both the preservation of persistent entities and the suppression of obsolete context. The global optimization objective is defined as follows:

$$\mathcal{L}_{\text{gen}} = \mathcal{L}_{\text{DMD}} + \lambda_{\text{slot}} \mathcal{L}_{\text{slot}}, \quad (6)$$

where  $\mathcal{L}_{\text{DMD}}$  represents the distribution matching distillation loss for adapting the diffusion backbone, and  $\lambda_{\text{slot}}$  serves as the balancing coefficient for slot regularization. The auxiliary slot loss  $\mathcal{L}_{\text{slot}}$  is partitioned into two components:  $\mathcal{L}_{\text{con}}$  enforces temporal identity consistency and  $\mathcal{L}_{\text{rec}}$  facilitates token-level grounding during the second training stage, formulated as

$$\mathcal{L}_{\text{slot}} = \mathcal{L}_{\text{con}} + \mathbb{I}_{\text{stage2}} \alpha \mathcal{L}_{\text{rec}}. \quad (7)$$

To maintain the semantic identity of slots across sequential chunks, we employ a contrastive objective. Let  $U^{c-1}$  and  $U^c \in \mathbb{R}^{B \times S \times d_s}$  denote the  $\ell_2$ -normalized slot embeddings from consecutive chunks. With  $N = BS$ , the contrastive loss is defined as:

$$\mathcal{L}_{\text{con}} = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(\text{sim}(u_i^{c-1}, u_i^c)/\tau)}{\sum_{j=1}^N \exp(\text{sim}(u_i^{c-1}, u_j^c)/\tau)}, \quad (8)$$

where  $\text{sim}(\cdot, \cdot)$  denotes cosine similarity and  $\tau$  is the temperature parameter. This objective encourages the model to map identical entities to the same slot index over time, effectively stabilizing the memory routing mechanism across long horizons and preventing identity drift.

Table 1. **Quantitative comparison of interactive long video generation.** Overall Quality scores represent the full 60s sequence, and segmented CLIP scores evaluate prompt compliance at 10s intervals ( $\uparrow$  higher is better). Our method outperforms existing streaming autoregressive models across all temporal boundaries, demonstrating superior long-horizon consistency and reduced semantic degradation. Best results are shown in bold.

Method	Quality Score $\uparrow$	CLIP Score $\uparrow$					
		0–10 s	10–20 s	20–30 s	30–40 s	40–50 s	50–60 s
Infinity-RoPE [Yesiltepe et al. 2025]	79.98	23.87	22.62	22.16	21.82	22.20	22.20
LongLive [Yang et al. 2026a]	79.09	26.49	25.60	24.89	24.48	24.81	24.54
MemFlow [Ji et al. 2025]	78.57	26.29	24.09	23.15	23.93	23.68	23.39
<b>SlotMemory</b>	<b>81.61</b>	<b>26.81</b>	<b>26.11</b>	<b>25.18</b>	<b>25.33</b>	<b>24.91</b>	<b>25.25</b>

To ensure that slots accurately represent the underlying visual manifold, we introduce a reconstruction regularization  $\mathcal{L}_{\text{rec}}$ . We utilize a hard top-1 slot assignment combined with straight-through estimation to reconstruct adapted token features  $\tilde{x}_{b,t}$  from their assigned slot embeddings. Defining a confidence mask

$$v_{b,t} = \mathbb{I}(\max_s A_{b,s,t}^c > 0.5), \quad (9)$$

the reconstruction loss is:

$$\mathcal{L}_{\text{rec}} = \frac{\sum_{b,t} v_{b,t} (1 - \cos(\hat{x}_{b,t}, \text{sg}(\tilde{x}_{b,t})))}{\sum_{b,t} v_{b,t} + \varepsilon}, \quad (10)$$

where  $\hat{x}_{b,t}$  is the reconstructed feature and  $\text{sg}(\cdot)$  denotes the stop-gradient operator. This stage-II loss prevents semantic drift by grounding abstract slot addresses in the high-fidelity transformer feature space, ensuring that the stored KV tokens remain representative of the original scene geometry.

## 4 Experiments

### 4.1 Experiment Setups

**Implementation Details.** We instantiate SlotMemory using the pretrained Wan2.1-T2V-1.3B backbone [Wan et al. 2025], which generates 5-second clips at 16 FPS and a resolution of  $832 \times 480$ . Following the optimization strategy of LongLive [Yang et al. 2026a] and MemFlow [Ji et al. 2025], the training process is divided into two phases. First, the base model is adapted into a few-step causal-attention generator via a Self-Forcing Distribution Matching Distillation pipeline [Huang et al. 2026] on the VidProM dataset [Wang and Yang 2024], utilizing short-window attention and first-frame sink tokens to stabilize autoregressive rollouts. Second, we freeze the backbone weights and initialize the slot adapter and memory modules for streaming long-tuning. This phase utilizes 60-second multi-prompt sequences where each iteration continues the model’s own rollout. Prompt switches are sampled uniformly between 5 and 55 seconds to train the module to perform both entity preservation and context suppression. The memory bank and sparse activation modules are optimized for 3,000 steps using AdamW (learning rates of  $1.0 \times 10^{-5}$  for the actor and  $2.0 \times 10^{-6}$  for the critic). Training is completed in approximately 12 hours on 8 NVIDIA H200 GPUs.

**Benchmarks and Datasets.** To evaluate temporal robustness across varying horizons, we establish benchmarks for three distinct scenarios: short-video generation (5s), long single-prompt generation (30s), and interactive multi-prompt narratives (60s). For the 5-second and 30-second single-prompt evaluations, we adopt the standard protocols defined by Self-Forcing [Huang et al. 2026] and LongLive [Yang

Table 2. **Quantitative comparisons under single-prompt 5-second setting.** This table presents standard VBench metrics to confirm the preservation of base model capabilities ( $\uparrow$  higher is better). Our approach maintains strong visual quality and significantly outperforms recent streaming baselines in both dynamic motion recovery and semantic accuracy, demonstrating that our long-video memory mechanisms do not degrade short-context generation. The best value in each metric column is shown in bold.

Model	#Params	Resolution	Dynamic $\uparrow$	Evaluation scores $\uparrow$		
				Total	Quality	Semantic
<i>Diffusion models</i>						
LTX-Video [HaCohen et al. 2024]	1.9B	768×512	<b>75.00</b>	82.33	85.20	70.83
Wan2.1 [Wan et al. 2025]	1.3B	832×480	73.61	<b>84.27</b>	<b>85.24</b>	80.43
<i>Autoregressive models</i>						
SkyReels-V2 [Chen et al. 2025]	1.3B	960×540	48.61	81.72	82.93	76.90
CausVid [Yin et al. 2025]	1.3B	832×480	63.89	82.96	83.94	79.02
Pyramid Flow [Jin et al. 2025]	2B	640×384	68.06	81.56	83.89	72.24
Self-Forcing [Huang et al. 2026]	1.3B	832×480	63.89	83.83	84.60	80.79
Infinity-RoPE [Yesiltepe et al. 2025]	1.3B	832×480	58.33	83.55	84.34	80.43
DeepForcing [Yi et al. 2025]	1.3B	832×480	63.89	83.85	84.61	80.84
MemFlow [Ji et al. 2025]	1.3B	832×480	50.00	81.76	83.74	73.86
LongLive [Yang et al. 2026a]	1.3B	832×480	37.50	82.81	83.26	81.01
<b>SlotMemory</b>	1.3B	832×480	61.11	83.76	84.40	<b>81.24</b>

et al. 2026a]. The 60-second interactive benchmark follows the methodology introduced in MemFlow [Ji et al. 2025], consisting of 100 narrative scripts. Each script contains six consecutive 10-second prompts, requiring the model to maintain semantic consistency across five discrete transition points.

**Evaluation Metrics.** We perform a multidimensional quantitative assessment using the VBench [Huang et al. 2024] and VBench-Long [Huang et al. 2025] frameworks. Short-video performance is measured via Total, Quality, Semantic, and Dynamic scores. For the 30-second horizon, we extend this to include Imaging Quality and Temporal Style to capture long-range perceptual stability. In the 60-second interactive setting, we report the aggregate Quality score to reflect overall visual fidelity. To specifically analyze the impact of prompt transitions, we compute segment-wise CLIP scores [Wang et al. 2022] for each 10-second interval, providing a granular view of how well semantic identities are preserved as the narrative evolves.

### 4.2 Main Experiments

**Interactive Multi-Prompt Generation.** To test the efficacy of prompt-aware retrieval, we evaluate 60-second interactive narratives comprising six distinct prompts. As shown in Table 1, SlotMemory achieves a state-of-the-art Quality Score of 81.61. Granular analysis via segment-wise CLIP scores reveals that the performance gap between our method and existing baselines widens as the sequence progresses; the relative gain is most substantial in the 30–60s interval. This trend demonstrates that SlotMemory successfully suppresses irrelevant historical context during prompt transitions while selectively retaining core semantic entities. Qualitative results in Figure 8 further confirm these findings, showing that SlotMemory preserves stable astronaut-rover spatial relationships and clear action trajectories—such as returning to base—where baselines exhibit significant narrative degradation.

**Long-Video Generation.** The advantages of structured semantic representation become pronounced at the 30-second horizon. In this

Table 3. **Quantitative comparisons under a 30-second single-prompt setting.** We evaluate extended-horizon generation capabilities against recent autoregressive streaming models ( $\uparrow$  higher is better). While existing baselines experience severe feature drift and memory instability over time—resulting in degraded motion dynamics—our method maintains robust temporal coherence. This yields a state-of-the-art Dynamic Score and superior overall Imaging Quality without sacrificing semantic alignment. The best value in each metric column is shown in bold.

Model	Total Score $\uparrow$	Quality Score $\uparrow$	Semantic Score $\uparrow$	Dynamic Score $\uparrow$	Imaging Quality $\uparrow$	Temporal Style $\uparrow$
Self-Forcing [Huang et al. 2026]	82.21	83.38	77.54	51.85	68.17	23.31
Infinity-RoPE [Yesiltepe et al. 2025]	83.16	84.17	79.13	57.04	67.94	23.79
DeepForcing [Yi et al. 2025]	82.26	83.20	78.50	52.86	67.65	23.51
MemFlow [Ji et al. 2025]	82.95	83.86	79.31	60.46	67.68	23.89
LongLive [Yang et al. 2026a]	82.77	83.31	<b>80.64</b>	40.19	68.96	23.97
<b>SlotMemory</b>	<b>84.28</b>	<b>85.23</b>	80.49	<b>74.29</b>	<b>72.26</b>	<b>24.34</b>

setting, SlotMemory surpasses or matches all autoregressive baselines, establishing state-of-the-art results across nearly all VBench-Long dimensions (Table 3). We achieve a Dynamic Score of 74.29—a significant improvement over the 60.46 achieved by MemFlow and the 40.19 of LongLive. The simultaneous gains in Imaging Quality (72.26) and Temporal Style (24.34) indicate that slot-conditioned storage effectively mitigates the cumulative feature drift typical of unclustered temporal memory. By isolating entities within discrete slots, the model maintains consistent motion dynamics and appearance across extended rollouts, where traditional frame-centric methods often suffer from visual collapse or identity switching.

**Short-Video Generation.** We first evaluate SlotMemory on the 5-second single-prompt benchmark to ensure that the integration of the object-centric memory module does not degrade the generative priors of the Wan2.1-T2V-1.3B backbone. As reported in Table 2, SlotMemory achieves a Total Score of 83.76 and a Quality Score of 84.40, performing competitively with the original Wan2.1 and LTX-Video. Compared to streaming-memory baselines such as LongLive and MemFlow, our method yields superior Semantic Scores (81.24), suggesting that the slot-based abstraction enhances text-entity association even within short temporal windows. These results validate that structured memory provides a robust semantic foundation without compromising short-range visual fidelity.

**Summary of Scaling Performance.** Across the 5s, 30s, and 60s evaluations, a consistent performance trajectory emerges: the marginal utility of slot-conditioned memory scales with the duration and complexity of the synthesis task. While maintaining parity with high-fidelity diffusion models in short clips, SlotMemory provides increasingly dominant advantages in long-range consistency and interactive adherence.

### 4.3 Ablation Studies

We conduct a series of ablation experiments under the 60-second interactive narrative setting to isolate the contribution of each architectural component and assess hyperparameter sensitivity.

**Component Analysis.** The progressive integration of our proposed modules, detailed in Table 4, reveals that structured memory requires explicit regularization to outperform unclustered baselines. The introduction of the Slot Module without temporal constraints leads to a decline in the Dynamic Score, suggesting that naive feature compression disrupts motion continuity. However, the addition

Table 4. **Progressive component integration on interactive long-video generation.** Each row adds one component or objective on top of the previous row. Higher is better for all metrics.

Progressive Setting	CLIP $\uparrow$	Quality $\uparrow$	Dynamic $\uparrow$
Baseline (w/o Slot Compression)	24.09	78.57	36.16
+ Slot Module (w/o Slot Regularization)	24.62	78.61	30.61
+ Contrastive Loss (Stage-I Only)	26.06	81.29	79.68
+ Reconstruction Loss (Stage-I + Stage-II)	25.60	81.61	75.73

Table 5. **Hyperparameter sensitivity.** We evaluate the impact of regularization weight (left) and memory bank capacity (right) on 60s interactive generation. The default configuration ( $\lambda_{\text{slot}} = 0.2$ , size=6) provides the optimal balance between semantic adherence and visual quality.

Slot Regularization weight			Memory Bank Size		
Value	Avg. CLIP $\uparrow$	Quality $\uparrow$	Size	Avg. CLIP $\uparrow$	Quality $\uparrow$
0.0	24.62	78.61	3	25.94	81.42
<b>0.2 (default)</b>	<b>25.60</b>	<b>81.61</b>	<b>6 (default)</b>	25.60	<b>81.61</b>
0.4	25.21	79.68	12	<b>26.04</b>	81.35

of the Stage-I Contrastive Loss provides the most significant performance gain, with the Dynamic Score increasing to 79.68 and CLIP alignment reaching 26.06. This underscores that temporal consistency within the slot manifold is the primary driver of long-term stability. The final inclusion of the Stage-II Reconstruction Loss further refines Quality to 81.61, albeit with a marginal trade-off in motion intensity. This two-stage objective ensures that semantic slots remain grounded in high-fidelity transformer tokens rather than drifting into abstract latent space.

This trend is also visible qualitatively in Figure 5, which follows the same progressive settings as Table 4. Without slot compression, the sequence exhibits the largest instability, including duplicated cake/candle-like objects and weaker role consistency when the mother enters. Adding the slot module improves object-count and layout stability of the boy-cake-candle-mother structure, but action completion is still weak. Adding contrastive loss improves interaction semantics (e.g., the mother joining and accompanying the boy) and narrative readability across the mid/late timestamps. Adding reconstruction loss gives the most coherent local details (facial rendering, candlelight appearance, and affective continuity), consistent with its best Quality score in Table 4.

**Hyperparameter Sensitivity and Capacity.** The efficacy of SlotMemory is sensitive to the regularization weight  $\lambda_{\text{slot}}$  and the memory bank capacity. As shown in Table 5, a weight of 0.2 provides the optimal balance between enforcing object-centric structure and preserving the generative priors of the diffusion backbone; exceeding this threshold (0.4) degrades visual quality. Regarding memory capacity, Figure 4 demonstrates that while increasing the bank size from 3 to 12 slots improves late-segment prompt adherence, the gains in CLIP Score plateau after 6 slots. Furthermore, overall Quality peaks at a capacity of 6, suggesting that excessively large memory banks may introduce historical noise that interferes with the current denoising state. Consequently, we identify a bank size of 6 as the optimal bottleneck for persistent streaming synthesis.




Fig. 3. **Qualitative comparison on an interactive Mars mission script.** Each row is one method and each column is a key moment in the six-stage narrative (surface exploration, rover approach, collaborative scan/drill, sample handling, and return toward base). Our method (bottom row) maintains more stable subject identity, astronaut-rover spatial relations, and action readability across the full sequence.




Fig. 4. **Segment-wise CLIP comparison across memory designs and bank sizes.** We compare Frame Sink, NAM, and SlotMemory with different bank sizes ( $b \in \{3, 6, 9, 12\}$ ). SlotMemory consistently outperforms Frame Sink and NAM across all 10-second segments, and larger banks improve mid/late-segment CLIP with diminishing returns beyond moderate capacity.

**Comparative Memory Architectures.** We evaluate SlotMemory against alternative historical representations, including Frame Sink and Normalized Attention Memory (NAM). As illustrated in the segment-wise analysis in Figure 4, our object-centric approach consistently maintains higher CLIP Scores across all 10-second intervals. Unlike frame-centric methods that suffer from rapid semantic decay as the narrative progresses, SlotMemory effectively preserves entity-level information through multiple prompt transitions, confirming that decomposing the KV manifold into discrete semantic addresses is superior to raw temporal caching.

#### 4.4 Limitations and Future Work

**Failure Cases and Limitations.** Despite the advancements in long-term narrative coherence, certain architectural bottlenecks

persist within the SlotMemory framework. As shown in Figure 6, in scenarios characterized by high entity density or visual similarity, the model occasionally experiences attribute leakage, where fine-grained properties—such as texture or color—migrate between adjacent semantic slots. Furthermore, the transition logic inherited from the baseline KV-recache mechanism can induce transient scene regression artifacts; in these instances, the synthesis may momentarily revert to a prior environmental state before successfully adapting to the new prompt instructions. These failure modes suggest that while slot-based abstraction provides a robust foundation for entity persistence, maintaining strict semantic boundaries under complex spatial conditions remains a non-trivial challenge.

**Future Work.** Future research will focus on three primary avenues to enhance the robustness and scalability of object-centric streaming synthesis. First, to mitigate attribute leakage in dense or cluttered environments, we intend to investigate more rigorous disentanglement objectives and hierarchical slot architectures that enforce stricter semantic boundaries between proximate visual entities. Second, to address scene regression observed during prompt transitions, we aim to replace heuristic KV-recaching with learned, manifold-aware transition operators capable of fluidly bridging disparate semantic states while maintaining global narrative continuity. Finally, we envision extending the framework toward dynamic slot allocation mechanisms, enabling the model to adaptively scale its representational capacity in response to the varying semantic complexity of the synthesized scene.

## 5 Conclusion

We present SlotMemory, a novel object-centric KV memory mechanism that addresses the fundamental limitations of temporal-centric organization in long-form video synthesis. By shifting the memory abstraction from temporal occurrence to semantic identity, our framework enables the persistent representation of visual entities across extended 60-second horizons and complex, multi-prompt

transitions. Both quantitative and qualitative results demonstrate that SlotMemory achieves state-of-the-art performance in both long-range consistency and prompt adherence, establishing structured semantic decomposition—rather than raw temporal capacity—as the critical primitive for high-fidelity streaming diffusion.

## References

- Omer Bar-Tal, Hila Chefer, Omer Tov, Charles Herrmann, Roni Paiss, Shiran Zada, Ariel Ephrat, Junhwa Hur, Guanghui Liu, Amit Raj, et al. 2024. Lumiere: A space-time diffusion model for video generation. In *SIGGRAPH Asia 2024 Conference Papers*. 1–11.
- Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. 2023. Stable video diffusion: Scaling latent video diffusion models to large datasets.
- Christopher P Burgess, Loic Matthey, Nicholas Watters, Rishabh Kabra, Irina Higgins, Matt Botvinick, and Alexander Lerchner. 2019. Monet: Unsupervised scene decomposition and representation.
- Shengqu Cai, Ceyuan Yang, Lvmin Zhang, Yuwei Guo, Junfei Xiao, Ziyang Yang, Yinghao Xu, Zhenheng Yang, Alan Yuille, Leonidas Guibas, et al. 2025. Mixture of contexts for long video generation.
- Guibin Chen, Dixuan Lin, Jiangping Yang, Chunze Lin, Junchen Zhu, Mingyuan Fan, Hao Zhang, Sheng Chen, Zheng Chen, Chengcheng Ma, et al. 2025. Skyreels-v2: Infinite-length film generative model.
- Honglin Chen, Wanhee Lee, Hong-Xing Yu, Rahul Mysore Venkatesh, Joshua B Tenenbaum, Daniel Bear, Jiajun Wu, and Daniel LK Yamins. 2024. Unsupervised 3D scene representation learning via movable object inference. *Transactions on Machine Learning Research* (2024).
- Karan Dalal, Daniel Kocaja, Jiarui Xu, Yue Zhao, Shihao Han, Ka Chun Cheung, Jan Kautz, Yejin Choi, Yu Sun, and Xiaolong Wang. 2025. One-minute video generation with test-time training. In *Proceedings of the Computer Vision and Pattern Recognition Conference*. 17702–17711.
- Aniket Rajiv Didolkar, Andrii Zadaianchuk, Anirudh Goyal, Michael Curtis Mozer, Yoshua Bengio, Georg Martius, and Maximilian Seitzer. 2025. On the Transfer of Object-Centric Representation Learning. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=bSq0XGS3kW>
- Gamaleldin Elsayed, Aravindh Mahendran, Sjoerd Van Steenkiste, Klaus Greff, Michael C Mozer, and Thomas Kipf. 2022. Savi++: Towards end-to-end object-centric learning from real-world videos. *Advances in Neural Information Processing Systems* 35 (2022), 28940–28954.
- Martin Engelcke, Adam R Kosiorek, Oiwi Parker Jones, and Ingmar Posner. 2020. GENESIS: Generative Scene Inference and Sampling with Object-Centric Latent Representations. In *International Conference on Learning Representations (ICLR)*.
- Xueji Fang, Liyuan Ma, Zhiyang Chen, Mingyuan Zhou, and Guo-jun Qi. 2025. Inlvg: Reinforce inference-time consistent long video generation with grpo.
- Klaus Greff, Raphaël Lopez Kaufman, Rishabh Kabra, Nick Watters, Christopher Burgess, Daniel Zoran, Loic Matthey, Matthew Botvinick, and Alexander Lerchner. 2019. Multi-object representation learning with iterative variational inference. In *International conference on machine learning*. PMLR, 2424–2433.
- Yuchao Gu, Weijia Mao, and Mike Zheng Shou. 2025. Long-Context Autoregressive Video Modeling with Next-Frame Prediction. arXiv:2503.19325 [cs.CV] <https://arxiv.org/abs/2503.19325>
- Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. 2023. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning.
- Yoav HaCohen, Nisan Chiprut, Benny Brazowski, Daniel Shalem, Dudu Moshe, Eitan Richardson, Eran Levin, Guy Shiran, Nir Zabari, Ori Gordon, et al. 2024. Ltx-video: Realtime video latent diffusion.
- Roberto Henschel, Levon Khachatryan, Hayk Poghosyan, Daniil Hayrapetyan, Vahram Tadevosyan, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. 2025. Streamingt2v: Consistent, dynamic, and extendable long video generation from text. In *Proceedings of the Computer Vision and Pattern Recognition Conference*. 2568–2577.
- Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. 2022a. Imagen video: High definition video generation with diffusion models.
- Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. 2022b. Video diffusion models. *Advances in neural information processing systems* 35 (2022), 8633–8646.
- Xun Huang, Zhengqi Li, Guande He, Mingyuan Zhou, and Eli Shechtman. 2026. Self forcing: Bridging the train-test gap in autoregressive video diffusion. *Advances in Neural Information Processing Systems* 38 (2026), 167283–167308.
- Ziqi Huang, Yanan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qingyang Jin, Nattapol Chanpaisit, Yaohui Wang, Xinyuan Chen, Limin Wang, Dahua Lin, Yu Qiao, and Ziweli Liu. 2024. VBench: Comprehensive Benchmark Suite for Video Generative Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Ziqi Huang, Fan Zhang, Xiaojie Xu, Yanan He, Jiashuo Yu, Ziyue Dong, Qianli Ma, Nattapol Chanpaisit, Chenyang Si, Yuming Jiang, Yaohui Wang, Xinyuan Chen, Ying-Cong Chen, Limin Wang, Dahua Lin, Yu Qiao, and Ziweli Liu. 2025. VBench++: Comprehensive and Versatile Benchmark Suite for Video Generative Models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2025). doi:10.1109/TPAMI.2025.3633890
- Allan Jabri, Sjoerd van Steenkiste, Emiel Hoogeboom, Mehdi SM Sajjadi, and Thomas Kipf. 2024. DORSal: Diffusion for Object-centric Representations of Scenes. In *International Conference on Learning Representations*, Vol. 2024. 43030–43052.
- Sihui Ji, Xi Chen, Shuai Yang, Xin Tao, Pengfei Wan, and Hengshuang Zhao. 2025. Mem-flow: Flowing adaptive memory for consistent and efficient long video narratives.
- Yang Jin, Zhicheng Sun, Ningyuan Li, Kun Xu, Hao Jiang, Nan Zhuang, Quzhe Huang, Yang Song, Yadong Mu, and Zhouchen Lin. 2025. Pyramidal flow matching for efficient video generative modeling. In *International Conference on Learning Representations*, Vol. 2025. 23378–23402.
- Jihwan Kim, Junoh Kang, Jinyoung Choi, and Bohyung Han. 2024. Fifo-diffusion: Generating infinite videos from text without training. *Advances in Neural Information Processing Systems* 37 (2024), 89834–89868.
- Thomas Kipf, Gamaleldin F Elsayed, Aravindh Mahendran, Austin Stone, Sara Sabour, Georg Heigold, Rico Jonschkowski, Alexey Dosovitskiy, and Klaus Greff. 2021. Conditional object-centric learning from video.
- Avinash Kori, Francesco Locatello, Fabio De Sousa Ribeiro, Francesca Toni, and Ben Glocker. 2024a. Grounded object-centric learning. In *International Conference on Learning Representations*, Vol. 2024. 42559–42603.
- Avinash Kori, Francesco Locatello, Aankaran Santhirasekaram, Francesca Toni, Ben Glocker, and Fabio De Sousa Ribeiro. 2024b. Identifiable object-centric representation learning via probabilistic slot attention. *Advances in Neural Information Processing Systems* 37 (2024), 93300–93335.
- Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. 2020. Object-centric learning with slot attention. *Advances in neural information processing systems* 33 (2020), 11525–11538.
- Riccardo Majellaro, Jonathan Collu, Aske Plaatt, and Thomas M. Moerland. 2025. Explicitly Disentangled Representations in Object-Centric Learning. *Transactions on Machine Learning Research* (2025). <https://openreview.net/forum?id=r8UFp9oLQ0>
- Anna Manasyan, Maximilian Seitzer, Filip Radovic, Georg Martius, and Andrii Zadaianchuk. 2025. Temporally consistent object-centric learning by contrasting slots. In *Proceedings of the Computer Vision and Pattern Recognition Conference*. 5401–5411.
- Alexander Rubinstein, Ameya Prabhu, Matthias Bethge, and Seong Joon Oh. 2025. Rethinking Object-Centric Representations in the Era of Foundational Segmentation Models. In *Workshop on Spurious Correlation and Shortcut Learning: Foundations and Solutions*. <https://openreview.net/forum?id=gGhRHtVZDo>
- Maximilian Seitzer, Max Horn, Andrii Zadaianchuk, Dominik Zietlow, Tianjun Xiao, Carl-Johann Simon-Gabriel, Tong He, Zheng Zhang, Bernhard Schölkopf, Thomas Brox, et al. 2022. Bridging the gap to real-world object-centric learning.
- Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gefni, et al. 2022. Make-a-video: Text-to-video generation without text-video data.
- Hansi Teng, Hongyu Jia, Lei Sun, Lingzhi Li, Maolin Li, Mingqiu Tang, Shuai Han, Tianning Zhang, WQ Zhang, Weifeng Luo, et al. 2025. Magi-1: Autoregressive video generation at scale.
- Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Fei Wu Yu, Haiming Zhao, Jianxiao Yang, et al. 2025. Wan: Open and advanced large-scale video generative models.
- Jiunui Wang, Hangjie Yuan, Dayou Chen, Yingya Zhang, Xiang Wang, and Shiweli Zhang. 2023. Modelscope text-to-video technical report.
- Wenhao Wang and Yi Yang. 2024. VidProM: A Million-scale Real Prompt-Gallery Dataset for Text-to-Video Diffusion Models. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*. <https://openreview.net/forum?id=pYNI76onJL>
- Yi Wang, Kunchang Li, Yizhuo Li, Yanan He, Bingkun Huang, Zhiyu Zhao, Hongjie Zhang, Jilan Xu, Yi Liu, Zun Wang, Sen Xing, Guo Chen, Junting Pan, Jiashuo Yu, Yali Wang, Limin Wang, and Yu Qiao. 2022. InternVideo: General Video Foundation Models via Generative and Discriminative Learning. arXiv:2212.03191 [cs.CV] <https://arxiv.org/abs/2212.03191>
- Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. 2023. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *Proceedings of the IEEE/CVF international conference on computer vision*. 7623–7633.
- Ziyi Wu, Nikita Dvornik, Klaus Greff, Thomas Kipf, and Animesh Garg. 2022. Slotformer: Unsupervised visual dynamics simulation with object-centric models.
- Shuai Yang, Wei Huang, Ruihang Chu, Yicheng Xiao, Yuyang Zhao, Xianbang Wang, Muiyang Li, Enze Xie, Ying-Cong Chen, Yao Lu, Song Han, and Yukang Chen. 2026a. LongLive: Real-time Interactive Long Video Generation. In *The Fourteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=>

nCAODkpsPJ

- Yuwei Yang et al. 2025. Object-Centric Noise for Slot-Based Learning.
- Yang Yang, Tianyi Zhang, Wei Huang, Jinwei Chen, Boxi Wu, Xiaofei He, Deng Cai, Bo Li, and Peng-Tao Jiang. 2026b. Anchor Forcing: Anchor Memory and Tri-Region RoPE for Interactive Streaming Video Diffusion.
- Hidir Yesiltepe, Tuna Han Salih Meral, Adil Kaan Akan, Kaan Oktay, and Pinar Yanardag. 2025. Infinity-rope: Action-controllable infinite video generation emerges from autoregressive self-rollout.
- Jung Yi, Wooseok Jang, Paul Hyunbin Cho, Jisu Nam, Heeji Yoon, and Seungryong Kim. 2025. Deep forcing: Training-free long video generation with deep sink and participative compression.
- Tianwei Yin, Qiang Zhang, Richard Zhang, William T Freeman, Fredo Durand, Eli Shechtman, and Xun Huang. 2025. From slow bidirectional to fast autoregressive video diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 22963–22974.
- Jiwen Yu, Jianhong Bai, Yiran Qin, Quande Liu, Xintao Wang, Pengfei Wan, Di Zhang, and Xihui Liu. 2025. Context as memory: Scene-consistent interactive long video generation with memory retrieval. In *Proceedings of the SIGGRAPH Asia 2025 Conference Papers*. 1–11.
- Lvmin Zhang, Shengqu Cai, Muiyang Li, Gordon Wetzstein, and Maneesh Agrawala. 2025a. Frame Context Packing and Drift Prevention in Next-Frame-Prediction Video Diffusion Models. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Weipu Zhang, Adam Jelley, Trevor McInroe, and Amos Storkey. 2025b. Objects Matter: Object-Centric World Models Improve Reinforcement Learning in Visually Complex Environments. In *Reinforcement Learning and Video Games Workshop @ RLC 2025*. <https://openreview.net/forum?id=C6aSY4DKDa>
- Yanpeng Zhao, Siyu Gao, Yunbo Wang, and Xiaokang Yang. 2024. DynaVol: Unsupervised learning for dynamic scenes through object-centric voxelization. In *International Conference on Learning Representations*, Vol. 2024. 4009–4028.




Fig. 5. **Qualitative ablation comparison on a multi-stage “birthday candle” narrative.** Columns show keyframes at 12s, 24s, 36s, 48s, and 53s from the same prompt (dim kitchen, boy, cake/candle, and mother joining later). Rows correspond to progressive settings: w/o Slot Compression, + Slot Module, + Contrastive Loss, and + Reconstruction Loss. As components are added, object-count stability, character interaction clarity, and emotional continuity improve.



Fig. 6. **Failure Cases of SlotMemory: Attribute Leakage and Scene Regression Artifacts.**




Fig. 7. **Qualitative comparison on an interactive script.**



Fig. 8. More qualitative results on an interactive script.

## A Additional Implementation Details

### A.1 Training Recipe

Our implementation follows a streaming DMD setup with chunk-wise autoregressive rollout. We use `distribution_loss=dmd_switch`, `num_frame_per_block=3`, `streaming_chunk_size=21`, `streaming_max_length=240`, `context_noise=0`, and denoising steps `{1000, 750, 500, 250}` (with warped scheduler indexing enabled).

For the final SlotMemory training schedule, we use two consecutive stages with fixed total update budget:

- **Stage-I (2000 steps)**. Set `freeze_slot_grouper=True` to freeze the slot grouper/initializer/dynamics predictor and optimize adapters plus diffusion-side parameters under the DMD objective with slot regularization.
- **Stage-II (1000 steps)**. Resume immediately from the Stage-I checkpoint, set `freeze_slot_grouper=False`, and continue end-to-end slot optimization with reconstruction-enabled slot regularization.

Thus, the full training uses **3000 steps** in total (Stage-I: 2000, Stage-II: 1000), matching the protocol used by our reported main results and ablations. Optimizer settings follow the repository defaults for long training: AdamW, generator learning rate  $1 \times 10^{-5}$ , critic learning rate  $2 \times 10^{-6}$ ,  $\beta_1 = 0.0$ ,  $\beta_2 = 0.999$ , batch size = 1, gradient accumulation = 1, EMA weight = 0.99, and EMA warmup starting at step 200. For switch-aware streaming supervision, the prompt switch position is sampled from a fixed candidate set (`switch_mode=random_choice`) with frame indices `{21, 39, 57, ..., 201}`, ensuring coverage of early/mid/late switching regimes during training.

### A.2 Memory-Module Hyperparameters

Table 6 summarizes the default long-video memory settings used in our experiments.

Table 6. **Default memory and slot hyperparameters in SlotMemory.**

Item	Value
Local attention window ( <code>local_attn_size</code> )	12 frames
Sink size ( <code>sink_size</code> )	3
Memory bank size ( <code>bank_size</code> )	6 blocks
Bank update interval ( <code>record_interval</code> )	3 blocks
SMA flag ( <code>SMA</code> )	False (default)
Slot dimension ( <code>slot_dim</code> )	64
Slots per chunk ( <code>num_slots_per_chunk</code> )	7
Tokens per slot ( <code>tokens_per_slot</code> )	64
Slot regularization weight ( $\lambda_{\text{slot}}$ )	0.2
Text/visual retrieval mixing (historical slot scoring)	0.7/0.3

In code, slot losses are accumulated in `MemoryCompressor` and injected into the generator objective as

$$\mathcal{L} = \mathcal{L}_{\text{DMD}} + \lambda_{\text{slot}} \mathcal{L}_{\text{slot}},$$

where  $\mathcal{L}_{\text{slot}}$  includes contrastive terms in both stages and reconstruction terms in Stage-II.

### A.3 Inference Protocol Details

We evaluate both single-prompt and interactive multi-prompt generation with the same memory backbone:

- **Single-prompt inference**: `num_output_frames=120`, chunked autoregressive generation with block size 6, and persistent KV cache plus bank.
- **Interactive inference**: `num_output_frames=240` with switch points at frames `{40, 80, 120, 160, 200}`, corresponding to six equal temporal segments.

At each prompt switch, we reset cross-attention cache and recache recent context frames under the new prompt; this preserves short-term temporal continuity while aligning conditional context to the new instruction. Memory-bank reads/writes remain bounded by the same fixed budget used in training.