
Addressable Memory for Video World Models

Xindi Wu^{1,2} Sven Elflein^{1,3,4} James Lucas¹ Olga Russakovsky² Laura Leal-Taixé¹ Despoina Paschalidou¹
Jonathan Lorraine¹ Aljosa Osep¹

Project page: research.nvidia.com/labs/sil/projects/WorldTrace/

Abstract

Long-horizon video world models forget: revisiting a previously seen scene yields a plausible-looking but different world. The bottleneck is positional addressability, not content: once cached-key positions exceed the training horizon, temporal RoPE goes out of distribution and the model cannot read stored content regardless of how it was compressed; on the evaluated architecture, any compressed summary at out-of-distribution positions is byte-identical to a sliding window. We propose WORLDTRACE, a training-free framework that anchors every summary slot to a fixed in-distribution position by slot rank, paired with two writers: WORLDTRACE-FIELD (canonical key averaging) for coherence and WORLDTRACE-LANDMARK (frozen canonical landmarks) for recall. On Matrix-Game-2, WORLDTRACE lifts TempSSIM by +15.5% at 24× the training horizon, raises position-aligned CLIP recall by up to +31.6%, and sustains near-perfect scene reconstruction at 256×. We also introduce **LoopMem**, a scripted-navigation benchmark scoring returned-scene fidelity across waypoints, edge lengths, orientation closures, and revisit depths.

1. Introduction

Autoregressive video world models (Ha & Schmidhuber, 2018; Hafner et al., 2025; Brooks et al., 2024; Bruce et al., 2024; He et al., 2025b; Decart et al., 2024) generate scenes block-by-block, attending over a Key-Value (KV) cache of prior context. Such models promise interactive applications such as next-generation game engines and closed-loop robot simulators, where users move freely through a visual world and revisit prior locations. Revisits demand *visual persistence*: when an agent returns, the model must reproduce the original scene, not a plausible-looking alternative.

¹NVIDIA ²Princeton University ³University of Toronto ⁴Vector Institute.

Proceedings of the 43rd International Conference on Machine Learning, From Frames to Stories (F2S) Workshop, Seoul, South Korea. PMLR 306, 2026. Copyright 2026 by the author(s).

In practice, visual persistence degrades sharply once generation exceeds the training context. The standard fix compresses the linearly growing KV cache into a fixed-size memory of recent context plus older summary slots (Zhang et al., 2023; Li et al., 2024; Cai et al., 2025; Yesiltepe et al., 2026; Kim et al., 2026). We find that compression alone cannot restore recall. As the rollout horizon extends past the training window, temporal positional encodings (Su et al., 2024) on cached keys go out of distribution (OOD), and attention silently fails to retrieve what is stored: even a perfect summary becomes effectively inaccessible. *Long-horizon failure is therefore an addressability problem first, and a content problem second.* We make this concrete on the evaluated architecture: under MG2’s local-attention mask (He et al., 2025b), any compressed summary at OOD positions yields outputs byte-identical to a sliding window over recent KV entries, regardless of what was compressed.

We propose WORLDTRACE, a training-free framework that keeps compressed memory *addressable* by assigning each summary slot a fixed in-distribution position by slot rank, stable as generation time grows. With addressability fixed, retention is a clean choice. We organize the cache as a verbatim recent window plus summary slots, with two complementary writers: WORLDTRACE-FIELD aggregates older context in canonical (rotation-invariant) key space for temporal coherence, and WORLDTRACE-LANDMARK freezes canonical keys at detected scene-entry events for episodic recall. Both are drop-in to a frozen generator.

In controlled return-to-origin tasks, a sliding-window cache rapidly loses structural fidelity on the return frame, while WORLDTRACE recovers the original scene appearance (Fig. 1). This extends the effective memory horizon from seconds to minutes-long rollouts: WORLDTRACE-FIELD improves TempSSIM by +15.5% at 24× the training horizon and WORLDTRACE-LANDMARK sustains near-perfect scene reconstruction at 256×, where the strongest compression baseline degrades by over a third. To exercise long-horizon recall under a fixed compute budget, we introduce **LoopMem**, a scripted-navigation benchmark scoring return-leg frames against initial-scene appearance across four difficulty axes.

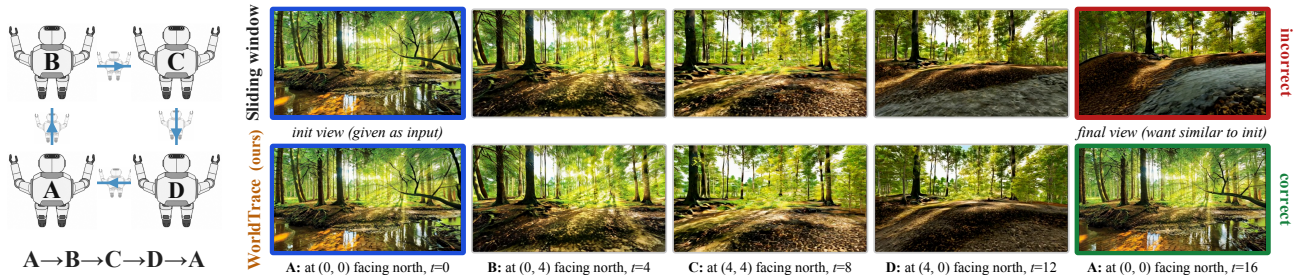


Figure 1. Addressable memory restores long-horizon visual persistence. *Left*: Scripted topology $A \rightarrow B \rightarrow C \rightarrow D \rightarrow A$. *Right*: Representative frames on Matrix-Game-2 along the path. The last column is the return to the initial scene A. The sliding window (top) drifts and mismatches the reference appearance on return (red border). WORLDTRACE with frozen landmark keys (*ours*, green border) matches scene A, illustrating addressable long-horizon recall.

Contributions. (i) We identify *positional addressability*, not summary content, as the binding constraint on long-horizon recall in autoregressive video world models, and show that on the evaluated architecture any compressed summary at OOD positions reduces byte-identically to a sliding window over recent context. (ii) We propose WORLDTRACE, a training-free KV-cache framework that fixes addressability by slot-rank positions, then exposes a two-knob design space for retention: WORLDTRACE-FIELD for coherence and WORLDTRACE-LANDMARK for episodic recall. (iii) We propose **LoopMem**, a scripted return-to-origin benchmark for long-horizon visual persistence, and show WORLDTRACE turns minutes-long generation from forgetting into recall: +15.5% TempSSIM at $24\times$ training horizon, near-perfect PAC at $256\times$, with $< 10\%$ wall-clock overhead and no fine-tuning.

2. WORLDTRACE

WORLDTRACE pairs (i) a slot-rank position scheme that keeps summary entries addressable at any horizon (Sec. 2.2) with (ii) a canonical-key writer: WORLDTRACE-FIELD (Sec. 2.3) for coherence, WORLDTRACE-LANDMARK (Sec. 2.4) for recall (Fig. 2). We assume temporal RoPE on keys and a fixed AR KV cache; single-shot diffusion and content-agnostic encodings are out of scope (App. F).

2.1. Cache Structure

We split the cache into a *recent window* \mathcal{R} of W_r verbatim AR blocks and a *summary cache* \mathcal{S} of N_s compressed slots, with $N_s + W_r = L_{\text{train}}$ (F latent frames per block). Total storage matches a sliding window. The N_s/W_r split trades coherence for recall (App. F).

2.2. Slot-Rank Virtual Positions

During training, attention is restricted to a local window $|q - t| \leq \Delta t_{\text{train}}$, so cached tokens at offsets exceeding Δt_{train} enter RoPE phase regimes the model never inverted at training time. Block-relativistic RoPE (Yesiltepe et al., 2026) caps offsets but collapses every summary slot be-

yond the cap to the same virtual position (Rem. 1); MemRoPE (Kim et al., 2026) sidesteps this only for two slots. To keep N_s slots distinguishable and in-distribution at any horizon, we assign positions by slot rank.

Definition 1 (WORLDTRACE Slot Indexing). The virtual position of slot s (oldest first) is

$$t_v^{(s)} = q - (L_{\text{train}} - 1 - s)F, \quad s = 0, \dots, N_s - 1. \quad (1)$$

Slot offsets depend only on slot rank, not absolute horizon N , so positions stay in-distribution as N grows.

Remark 1 (Block-Rel saturation). *Block-Rel caps look-back at $(L_{\text{train}} - 1)F$, so any slot whose centroid exceeds the cap is clamped to the same position; under uniform averaging all centroids drift past the cap as N grows, becoming indistinguishable. Eq. (1) avoids this: per-slot offsets stay distinct and N -independent.*

With slot positions fixed in-distribution, the remaining choice is the content writer.

2.3. WORLDTRACE-FIELD: Canonical Key Averaging

The most literal compressor for keys is averaging (Bolya et al., 2023; Rae et al., 2020). Naively averaging RoPE-rotated keys $\bar{K}_{\text{naive}}^{(k)} = \frac{1}{M} \sum_m R(\theta_k t_m) K_{\text{cx},m}^{(k)}$ sums vectors at different RoPE phases; once source frames span a RoPE period, components cancel and the summary erases its own signal regardless of content. WORLDTRACE-FIELD avoids this by averaging in canonical (un-rotated) key space and re-rotating to the slot’s virtual position.

Definition 2 (WORLDTRACE-FIELD operator). The compressed key at t_v for RoPE pair k is

$$K_{\text{field}}^{(k)}(t_v) := R(\theta_k t_v) \frac{1}{M} \sum_{m=1}^M R(-\theta_k t_m) K_{t_m}^{(k)}. \quad (2)$$

Each source key is unrotated to canonical content $K_{\text{cx},m}^{(k)}$, averaged, and re-encoded at t_v . Values carry no RoPE; slot values are the per-frame mean of source values. Sources are split into N_s contiguous temporal groups, so M has no separate hyperparameter and grows linearly with N .

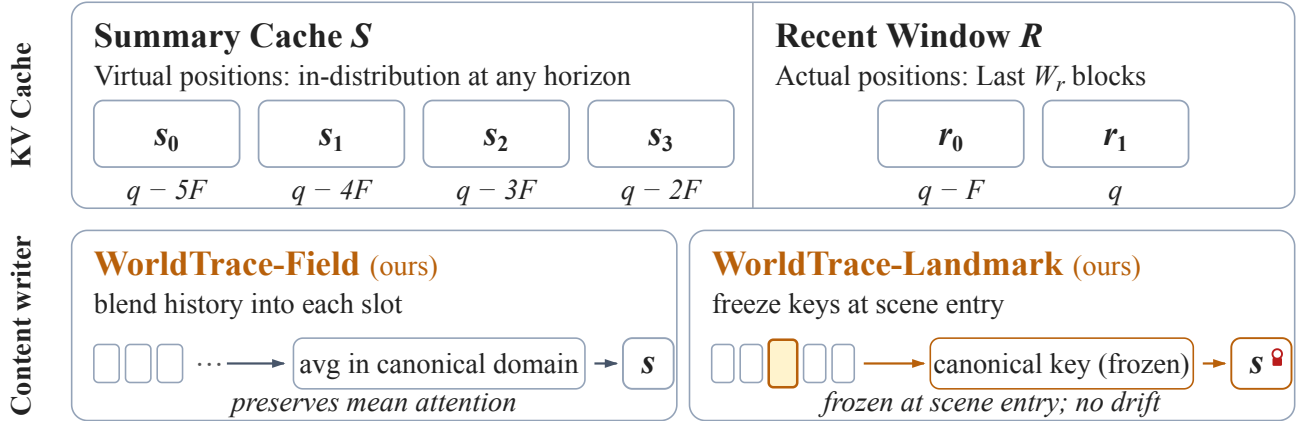


Figure 2. **WORLDTRACE cache structure.** Each AR chunk appends fresh KVs to the verbatim recent window \mathcal{R} (W_r slots). When that window overflows, the evicted block is compressed into the summary cache \mathcal{S} (N_s slots) by one of two writers: WORLDTRACE-FIELD aggregates older context in canonical (un-rotated) key space for trajectory coherence; WORLDTRACE-LANDMARK stores verbatim canonical keys at detected scene-entry events for episodic recall. Every summary slot receives a fixed in-distribution position by slot rank (Eq. (1)), so cached content stays addressable at any horizon, $O(1)$ in N , with $< 10\%$ wall-clock overhead.

Remark 2 (Mean-attention preservation, informal). *The compressed key in Def. 2 preserves the mean attention logit that the source keys would receive after being reassigned to the shared virtual position t_v (logits only; softmax non-linearity precludes the same statement for weights). Formal statement and proof in App. C (Prop. 1).*

2.4. WORLDTRACE-LANDMARK: Frozen Canonical Landmarks

WORLDTRACE-FIELD cannot recover specific scenes. WORLDTRACE-LANDMARK instead stores verbatim frames at detected scene-entry events as frozen canonical keys at slot-rank positions (in the spirit of Landmark Attention (Mohtashami & Jaggi, 2023)). We mark a scene-entry when the cosine distance between consecutive canonical-K signatures (Eq. (2)) spikes above τ ; the most recent fill the N_s slots (FIFO).

Frozen canonical keys. Slots shift toward the older end as new chunks arrive; under standard summary updates each shift unrotates and re-rotates a stored key, accumulating bfloat16 phase error (Wang et al., 2025). WORLDTRACE-LANDMARK stores each landmark’s canonical key once and applies a single fresh rotation to the current virtual position at every shift:

$$K_{\text{land}}^{(k)}(t_v^{(s)}) := R(\theta_k t_v^{(s)}) R(-\theta_k t_{\ell^*}) K_{t_{\ell^*}}^{(k)}, \quad (3)$$

where t_{ℓ^*} is the landmark timestamp and $K_{t_{\ell^*}}^{(k)}$ its source key; the form matches Eq. (2) with $M=1$, with canonical key reused across shifts.

2.5. Position-Content Coupling

Position and content are not independent: writers accumulate statistics at write-time offsets but are queried at read-time offsets, and when these disagree the stored statistic

no longer matches the query. WORLDTRACE handles both: Def. 1 fixes positions and Def. 2 writes in canonical space, which absorbs the source-to-virtual shift; slot content is then orthogonal. Storage stays at L_{train} blocks ($O(1)$ in N , identical to a sliding window) with $< 10\%$ per-chunk overhead up to $N=102$ (App. F, Tab. 15; full procedure in Alg. 1).

3. Experiments

We organize our experiments around three claims. **(Q1, Sec. 3.1)** Is position, not content, the binding constraint? **(Q2, Sec. 3.2)** Does WORLDTRACE-FIELD improve *coherence* over compression baselines? **(Q3, Sec. 3.3)** Does WORLDTRACE-LANDMARK sustain *episodic recall* at extended horizons on **LoopMem**?

Setup. We evaluate on Matrix-Game-2 (MG2-1.3B) (He et al., 2025b), a distilled 1.3B-parameter autoregressive video world model based on Wan 1.3B T2V (Wan Team, 2025) with 3D-RoPE, training-time KV cache extent $L_{\text{train}}=6$ AR blocks ($F=3$ latent frames each), and local attention window of two blocks ($\Delta t_{\text{train}}+1=6$ latent frames). We use $N_s=2$, $W_r=4$ for coherence and $N_s=4$, $W_r=2$ for episodic recall. Baselines: the MG2 default sliding window; Naive+Block-Rel (averaging in rotated key space without RoPE disentanglement); canonical averaging + Block-Rel (canonical-domain compression with Block-Rel positions, isolating position contribution); Landmark+Block-Rel (verbatim-recall ablation); concurrent training-free MemRoPE (Kim et al., 2026) and YaRN (Peng et al., 2024). Metrics: TempSSIM (Wang et al., 2004) and Local Scene Drift (CLIP feature distance to preceding chunk) for coherence; Position-Aligned CLIP (PAC, ViT-H/14 cosine between geometrically paired return- and forward-leg frames) for recall. Full definitions in App. F.5; LingBot-Fast (Robbyant Team et al., 2026) replication in App. D.

LoopMem benchmark. We introduce **LoopMem**, a scripted-navigation benchmark for episodic recall in autoregressive world models. The model executes a scripted path that returns to a previously visited location; the regenerated return frame is scored against the original scene appearance at geometrically matched positions, requiring no external reference. LoopMem organizes difficulty along four axes (App. E): waypoint count K , edge length L , camera-orientation closure, and multi-revisit depth R .

3.1. Position is the Binding Constraint

Holding the content writer fixed at canonical key averaging (Eq. (2)) and varying only the position scheme, Block-Rel offsets beyond the training window receive zero softmax weight

under MG2’s local-attention mask, collapsing canonical averaging to byte-identical sliding-window eviction; Centroid linear avoids saturation but uses an N -dependent formula that becomes unstable past the training window. Only WORLDTRACE assigns positions by slot rank alone (Eq. (1)), keeping every summary slot in-distribution at every horizon (Tab. 1; +5.9% over Block-Rel at $N=8$, +2.8% at $N=16$).

Table 1. TempSSIM (\uparrow): canonical averaging fixed; only positions vary.

Position	$N=8$	$N=16$
Block-Rel	0.390	0.530
Centroid	0.377	0.479
WORLDTRACE	0.413	0.545

3.2. Coherence: WORLDTRACE-FIELD

WORLDTRACE-FIELD combines WORLDTRACE slot-rank positions with canonical key averaging. Against the sliding-window baseline at $N=48$ ($24\times$ training horizon), WORLDTRACE-FIELD improves TempSSIM from 0.472 to 0.545 (+15.5% relative) and lowers Local Drift from 0.0305 to 0.0295, the only method to lead on both (Tab. 6, App. D.1). Canonical-K averaging variants with N -dependent positions cluster non-monotonically as horizon grows (0.479 to 0.449 at $N=48$), confirming that stacking content-domain heuristics cannot recover what an unstable position scheme loses; qualitative comparison in Fig. 5.

3.3. Episodic Recall on LoopMem

Tab. 2 ($N_s=4$, $W_r=2$) reports the full four-tier sweep; Fig. 3 the long-horizon trend. **Tier 1 (waypoint count).** PAC improves +19.5%/+17.7%/+19.9% on ABA/ABCA/ABCD; verbatim landmarks survive intermediate scenes the agent passes through. **Tier 2 (edge length).** The gap widens with KV-context distance, +7.3% at $N=8$, +19.5% at $N=16$, +31.6% at $N=32$. **Tier 3 (orientation closure).** +3.9%/+16.4% at $90^\circ/180^\circ$ (smaller 360° from approximate yaw scaling). **Tier 4 (multi-revisit depth).** ABABA reaches PAC=0.941 vs. 0.892; repeated visits do not corrupt the stored landmark. At $N=256$ (~ 192 s),




Figure 3. PAC vs. horizon. WORLDTRACE-LANDMARK (top) sustains PAC \approx 0.99 through $N=256$ (~ 192 s of decoded video); the canonical-K anchoring tier (middle) decays past $N=64$; compression-only sliding window (bottom) collapses by $N=32$.

WORLDTRACE-LANDMARK sustains PAC=0.989 while the best baseline drops below 0.65 (App. D).

Concurrent training-free baselines. MemRoPE (Kim et al., 2026) (Block-Rel + dual-rate EMA) survives OOD positions via decay but is beaten by Landmark+Block-Rel in the verbatim tier; transplanting WORLDTRACE positions into MemRoPE degrades it ($p<0.001$, App. D.2), confirming the position-content coupling of Sec. 2.5. YaRN (Peng et al., 2024) (NTK-aware temporal RoPE rescaling) gains +22% over sliding window at $N=32$ ($p<0.001$), but uses $O(N)$ memory and OOMs by $N\approx 100$; WORLDTRACE-LANDMARK exceeds it by 0.964 vs. 0.490 at $N=32$ in $O(1)$ memory and stays near-constant through $N=512$.

Cross-architecture replication. On LingBot-Fast (Robbyant Team et al., 2026), WORLDTRACE-LANDMARK improves PAC from $4\times$ training horizon onward (+8.9%/+14.1%/+7.3% at $4\times/6\times/8\times$), so the addressability bottleneck is not MG2-specific (App. D.3).

Discussion. The addressability framing factorizes long-horizon memory into a horizon-stable position scheme and a content writer that decides what each slot stores. This factorization explains a recurring pattern across prior work: position-only fixes (Block-Rel, YaRN) and content-only fixes (EMA, naive averaging) each plateau because each absorbs only one of the two failure modes. With both axes addressed, the design space along the N_s/W_r split exposes a clean trade-off: more summary slots favour recall-dominant tasks; more recent slots favour coherence-dominant ones. The framework as stated requires (i) temporal RoPE on keys and (ii) a known training context L_{train} ; learned absolute encodings and mixed-window models lie outside its scope. WORLDTRACE-LANDMARK’s scene-entry detector uses a fixed cosine threshold ($\tau=0.15$); learned or policy-conditioned detectors are an open direction (App. G.1).

Ablation. With cache budget fixed at $N_s+W_r=6$, three axes isolate the design choices (Tab. 7, App. D). **Position assignment** is the binding constraint for both writers: canon-

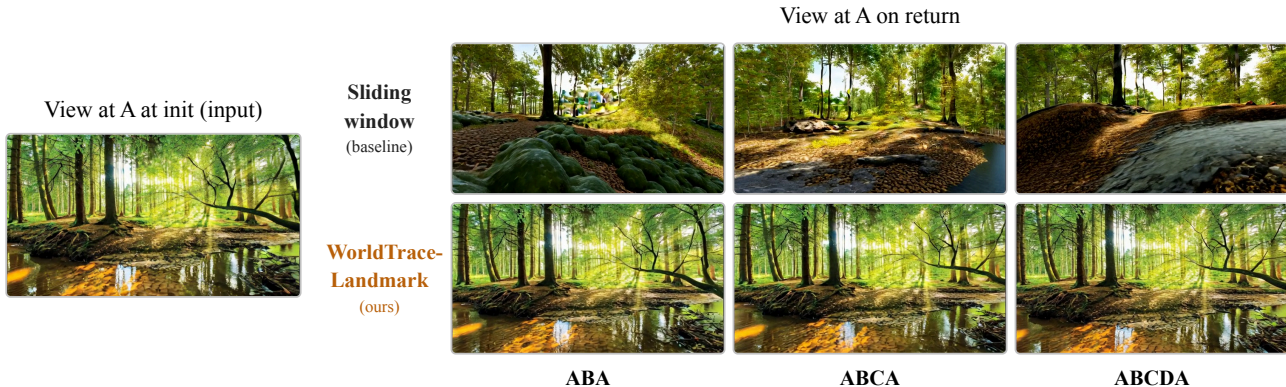


Figure 4. **LoopMem qualitative.** Sliding window mismatches the reference on return for ABA/ABCA/ABCDA; WORLDTRACE-LANDMARK matches scene A in all three.

Table 2. **WORLDTRACE-LANDMARK improves episodic recall across topology, edge length, camera orientation, and multi-revisit depth.** Each tier groups methods (rows) by loop configuration (columns). Primary metric (left of /) is PAC for ABA-topology and multi-revisit. Secondary metric (right of /) is TempSSIM over the return leg.

Tier 1: varying topology (K)			
	ABA ($N=16$)	ABCA ($N=17$)	ABCD ($N=16$)
Sliding window	0.723 \pm 0.013 / 0.761	0.673 \pm 0.014 / 0.740	0.666 \pm 0.013 / 0.748
WORLDTRACE-LANDMARK	0.864 \pm 0.009 / 0.786	0.792 \pm 0.011 / 0.758	0.799 \pm 0.011 / 0.771
Tier 2: varying edge length (L , ABA)			
	ABA short ($L=4$, $N=8$)	ABA ($L=8$, $N=16$)	ABA long ($L=16$, $N=32$)
Sliding window	0.859 \pm 0.007 / 0.780	0.723 \pm 0.013 / 0.761	0.627 \pm 0.013 / 0.771
WORLDTRACE-LANDMARK	0.922 \pm 0.004 / 0.789	0.864 \pm 0.009 / 0.786	0.825 \pm 0.009 / 0.787
Tier 3: camera-orientation (agent fixed at A)			
	Pan 90° ($N=4$)	Pan 180° ($N=8$)	Pan 360° ($N=8$)
Sliding window	0.829 \pm 0.008 / 0.480	0.671 \pm 0.014 / 0.458	0.559 \pm 0.015 / 0.455
WORLDTRACE-LANDMARK	0.861 \pm 0.007 / 0.493	0.781 \pm 0.010 / 0.486	0.577 \pm 0.015 / 0.467
Tier 4: multi-revisit ($R>1$)			
	ABABA ($R=2$, $N=32$)	ABCBA ($R_B=2$, $N=20$)	ABCDAB ($R_B=2$, $N=20$)
Sliding window	0.892 \pm 0.004 / 0.765	0.825 \pm 0.010 / 0.751	0.842 \pm 0.010 / 0.758
WORLDTRACE-LANDMARK	0.941 \pm 0.005 / 0.789	0.863 \pm 0.009 / 0.771	0.876 \pm 0.009 / 0.775

ical averaging with Block-Rel collapses to sliding-window byte-for-byte (OOD offsets suppress summary-slot attention), so the canonical-domain compression is invisible until positions stay in-distribution. **Canonical vs. naive averaging** surfaces on coherence: holding positions fixed at Block-Rel, switching from naive to canonical averaging cuts LatentDiff by -25.3% , and the gap grows as N_s rises ($+0.257$ at $N_s=1$ to $+0.312$ at $N_s=4$ for the naive baseline, while canonical averaging stays flat at 0.224 – 0.257). **Frozen landmark keys** prevent rotation-drift accumulation across summary updates: WORLDTRACE-LANDMARK vs. Landmark+Block-Rel grows from $+0.040$ at $N=64$ to $+0.046$ at $N=128$ (0.986 vs. 0.940), the gap scaling with the number of accumulated unrotate–rerotate shifts a long-running agent would incur.

Runtime overhead. WORLDTRACE adds canonical-domain key transforms and per-slot bookkeeping on top of the sliding-window baseline. At short horizons ($N=8$) all methods are within 6% of baseline because the forward pass dominates and the cache is small. At $N=102$ (~ 77 s of decoded video), WORLDTRACE-FIELD adds $+9.4\%$ per chunk as the canonical unrotate/rerotate of the growing source set takes a larger share of each step, consistent with the $O(T_{\text{old}} \cdot c_t \cdot n_\ell)$ per-token form (App. F). WORLDTRACE-

LANDMARK’s $+27.2\%$ at $N=102$ reflects per-step kernel overhead from scene-entry bookmark accumulation rather than additional FLOP; the per-chunk FLOP cost itself is $O(1)$ in N . Peak cache memory is unchanged from the sliding-window baseline at ~ 2.79 GB per batch element on MG2-1.3B in fp16.

Discussion. The addressability framing factorizes long-horizon memory into a horizon-stable position scheme and a content writer; position-only fixes (Block-Rel, YaRN) and content-only fixes (EMA, naive averaging) each plateau because they absorb only one of the two failure modes. With both axes addressed, the N_s/W_r split exposes a clean trade-off: more summary slots favour revisit-dominant tasks; more recent slots favour smooth local trajectories. The framework requires temporal RoPE on keys and a known training context L_{train} ; learned absolute encodings and mixed-window models lie outside its scope. WORLDTRACE-LANDMARK’s scene-entry detector uses a fixed cosine threshold ($\tau=0.15$); pose-, action-, or task-conditioned detectors that turn the current passive event-triggering into actively-selected retention are a natural next step (App. G.1).

4. Conclusion

Long-horizon visual persistence in autoregressive video world models is bottlenecked by *positional addressability*, not summary content: once RoPE offsets exceed the training window, attention fails to read cached content regardless of how it was compressed. WORLDTRACE fixes addressability by slot-rank positions, paired with two writers: WORLDTRACE-FIELD (canonical averaging) for coherence, WORLDTRACE-LANDMARK (frozen canonical landmarks) for recall. Both are training-free, drop-in, and turn minutes-long generation from forgetting into recall. Limitations and broader impact: App. G.1, App. G.3.

References

- Akulov, D., Sana, M., De Domenico, A., Salem, T. S., Pivovarov, N., and Ayed, F. Kvcompose: Efficient structured kv cache compression with composite tokens. *arXiv preprint arXiv:2509.05165*, 2025.
- Alonso, E., Jelley, A., Micheli, V., Kanervisto, A., Storkey, A., Pearce, T., and Fleuret, F. Diffusion for world modeling: Visual details matter in atari. *Advances in Neural Information Processing Systems*, 37:58757–58791, 2024.
- Barbero, F., Vitvitskiy, A., Perivolaropoulos, C., Pascanu, R., and Veličković, P. Round and round we go! what makes rotary positional encodings useful? In *International Conference on Learning Representations (ICLR)*, 2025. arXiv:2410.06205.
- Beltagy, I., Peters, M. E., and Cohan, A. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- Bettencourt, J., Wu, X., Atzmon, M., Lucas, J., and Lorraine, J. Variance reduction for expectations with diffusion teachers. *arXiv preprint arXiv:2605.21489*, 2026.
- Bolya, D., Fu, C.-Y., Dai, X., Zhang, P., Feichtenhofer, C., and Hoffman, J. Token merging: Your ViT but faster. In *International Conference on Learning Representations (ICLR)*, 2023.
- Brooks, T., Peebles, B., Holmes, C., DePue, W., Guo, Y., Jing, L., Schnurr, D., Taylor, J., Luhman, T., Luhman, E., Ng, C., Wang, R., and Ramesh, A. Video generation models as world simulators. OpenAI Technical Report, <https://openai.com/research/video-generation-models-as-world-simulators>, 2024.
- Bruce, J., Dennis, M. D., Edwards, A., Parker-Holder, J., Shi, Y., Hughes, E., Lai, M., Mavalankar, A., Steigerwald, R., Apps, C., et al. Genie: Generative interactive environments. In *International Conference on Machine Learning (ICML)*, 2024.
- Bulatov, A., Kuratov, Y., and Burtsev, M. S. Recurrent memory transformer. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- Cai, S., Yang, C., Zhang, L., Guo, Y., Xiao, J., Yang, Z., Xu, Y., Yang, Z., Yuille, A., Guibas, L., Agrawala, M., Jiang, L., and Wetzstein, G. Mixture of contexts for long video generation. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2508.21058, OpenReview: <https://openreview.net/forum?id=y6XJZ1EC2x>.
- Cai, Z., Zhang, Y., Gao, B., Liu, Y., Li, Y., Liu, T., Lu, K., Xiong, W., Dong, Y., Hu, J., and Xiao, W. PyramidKV: Dynamic KV cache compression based on pyramidal information funneling. In *Conference on Language Modeling (COLM)*, 2025.
- Chen, B., Martí Monsó, D., Du, Y., Simchowicz, M., Tedrake, R., and Sitzmann, V. Diffusion forcing: Next-token prediction meets full-sequence diffusion. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- Chen, H., Xu, C., Yang, X., Chen, X., and Deng, C. Past and future-informed kv cache policy with salience estimation in autoregressive video diffusion. *arXiv preprint arXiv:2601.21896*, 2026a.
- Chen, J., Bai, C., Hu, J., Xue, X., and Xu, M. Grounded forcing: Bridging time-independent semantics and proximal dynamics in autoregressive video synthesis. *arXiv preprint arXiv:2604.06939*, 2026b.
- Chen, K., Liang, D., Zhou, X., Ding, Y., Liu, X., Wan, P., and Bai, X. Out of sight but not out of mind: Hybrid memory for dynamic video world models. *arXiv preprint arXiv:2603.25716*, 2026c.
- Chen, S., Wong, S., Chen, L., and Tian, Y. Extending context window of large language models via positional interpolation. *arXiv preprint arXiv:2306.15595*, 2023.
- Chen, S., Wei, C., Sun, S., Nie, P., Zhou, K., Zhang, G., Yang, M.-H., and Chen, W. Context forcing: Consistent autoregressive video generation with long context. *arXiv preprint arXiv:2602.06028*, 2026d.
- Chen, T., Hu, X., Ding, Z., and Jin, C. Learning world models for interactive video generation. *arXiv preprint arXiv:2505.21996*, 2025a. Project page <https://sites.google.com/view/vrag>.
- Chen, Y., Lv, A., Luan, J., Wang, B., and Liu, W. HoPE: A novel positional encoding without long-term decay for enhanced context awareness and extrapolation. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2025b. arXiv:2410.21216.
- Corallo, G. and Papotti, P. FINCH: Prompt-guided key-value cache compression for large language models. *Transactions of the Association for Computational Linguistics*, 12:1517–1532, 2024.
- Cui, J., Wu, J., Li, M., Yang, T., Li, X., Wang, R., Bai, A., Ban, Y., and Hsieh, C.-J. Lol: Longer than longer, scaling video generation to hour. *arXiv preprint arXiv:2601.16914*, 2026a.

- Cui, J., Wu, J., Li, M., Yang, T., Li, X., Wang, R., Bai, A., Ban, Y., and Hsieh, C.-J. Self-Forcing++: Towards minute-scale high-quality video generation. In *International Conference on Learning Representations (ICLR)*, 2026b. arXiv:2510.02283.
- Dai, Z., Yang, Z., Yang, Y., Carbonell, J., Le, Q. V., and Salakhutdinov, R. Transformer-XL: Attentive language models beyond a fixed-length context. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 2978–2988, 2019. arXiv:1901.02860.
- Decart, Quevedo, J., McIntyre, Q., Campbell, S., Chen, X., and Wachen, R. Oasis: A universe in a transformer. <https://oasis-model.github.io>, 2024.
- Ding, Y., Zhang, L. L., Zhang, C., Xu, Y., Shang, N., Xu, J., Yang, F., and Yang, M. LongRoPE: Extending LLM context window beyond 2 million tokens. In *International Conference on Machine Learning (ICML)*, 2024.
- Duan, H., Yu, H.-X., Chen, S., Fei-Fei, L., and Wu, J. WorldScore: A unified evaluation benchmark for world generation. In *International Conference on Computer Vision (ICCV)*, 2025.
- Gao, X., Guan, J., Luo, S., Li, W., Tan, G., and Wang, J. Memcam: Memory-augmented camera control for consistent video generation. *arXiv preprint arXiv:2603.26193*, 2026.
- Golovneva, O., Wang, T., Weston, J., and Sukhbaatar, S. Contextual position encoding: Learning to count what’s important. *arXiv preprint arXiv:2405.18719*, 2024.
- Google DeepMind. Genie 2: A large-scale foundation world model. DeepMind Blog, <https://deepmind.google/discover/blog/genie-2-a-large-scale-foundation-world-model>, 2024.
- Google DeepMind. Genie 3: A new frontier for world models. DeepMind Blog, <https://deepmind.google/blog/genie-3-a-new-frontier-for-world-models>, 2025.
- Gu, A. and Dao, T. Mamba: Linear-time sequence modeling with selective state spaces. In *Conference on Language Modeling (COLM)*, 2024.
- Gu, Y., Mao, W., and Shou, M. Z. Long-context autoregressive video modeling with next-frame prediction. *arXiv preprint arXiv:2503.19325*, 2025.
- Guo, J., Ye, Y., He, T., Wu, H., Jiang, Y., Pearce, T., and Bian, J. MineWorld: A real-time and open-source interactive world model on Minecraft. *arXiv preprint arXiv:2504.08388*, 2025.
- Guo, Y., Zhang, Z., Wang, P., Liang, X., Ma, Z., and Zhang, L. Memorize when needed: Decoupled memory control for spatially consistent long-horizon video generation. *arXiv preprint arXiv:2604.18215*, 2026.
- Ha, D. and Schmidhuber, J. Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- Hafner, D., Pasukonis, J., Ba, J., and Lillicrap, T. Mastering diverse control tasks through world models. *Nature*, 640: 647–653, 2025.
- He, J., Xing, J., Wang, N., Xu, R., Wu, S., Zhou, P., Liu, Q., Xue, C. J., and Li, Q. A²ATS: Retrieval-based KV cache reduction via windowed rotary position embedding and query-aware vector quantization. In *Findings of the Association for Computational Linguistics (ACL Findings)*, 2025a. arXiv:2502.12665.
- He, X., Peng, C., Liu, Z., Wang, B., Zhang, Y., Cui, Q., Kang, F., Jiang, B., An, M., Ren, Y., Xu, B., Guo, H.-X., Gong, K., Wu, S., Li, W., Song, X., Liu, Y., Li, Y., and Zhou, Y. Matrix-game 2.0: An open-source, real-time, and streaming interactive world model. *arXiv preprint arXiv:2508.13009*, 2025b.
- Henschel, R., Khachatryan, L., Poghosyan, H., Hayrapetyan, D., Tadevosyan, V., Wang, Z., Navasardyan, S., and Shi, H. StreamingT2V: Consistent, dynamic, and extendable long video generation from text. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. arXiv:2403.14773.
- Heo, B., Park, S., Han, D., and Yun, S. Rotary position embedding for vision transformer. In *European Conference on Computer Vision (ECCV)*, 2024. arXiv:2403.13298.
- Hong, Y., Mei, Y., Ge, C., Xu, Y., Zhou, Y., Bi, S., Hold-Geoffroy, Y., Roberts, M., Fisher, M., Shechtman, E., Sunkavalli, K., Liu, F., Li, Z., and Tan, H. RELIC: Interactive video world model with long-horizon memory. *arXiv preprint arXiv:2512.04040*, 2025.
- Hsieh, C.-P., Sun, S., Kriman, S., Acharya, S., Rekish, D., Jia, F., Zhang, Y., and Ginsburg, B. RULER: What’s the real context size of your long-context language models? In *Conference on Language Modeling (COLM)*, 2024. arXiv:2404.06654.
- Hua, E., Jiang, C., Lv, X., Zhang, K., Sun, Y., Fan, Y., Zhu, X., Qi, B., Ding, N., and Zhou, B. Fourier position embedding: Enhancing attention’s periodic extension for length generalization. In *International Conference on Machine Learning (ICML)*, 2025. arXiv:2412.17739.

- Huang, J., Hu, X., Han, B., Shi, S., Tian, Z., He, T., and Jiang, L. Memory forcing: Spatio-temporal memory for consistent scene generation on Minecraft. *arXiv preprint arXiv:2510.03198*, 2025a.
- Huang, J., Ye, Z., Hu, X., He, T., Zhang, G., Shi, S., Bian, J., and Jiang, L. LIVE: Long-horizon interactive video world modeling. *arXiv preprint arXiv:2602.03747*, 2026a.
- Huang, S., Wu, J., Zhou, Q., Miao, S., and Long, M. Vid2World: Crafting video diffusion models to interactive world models. In *International Conference on Learning Representations (ICLR)*, 2026b. arXiv:2505.14357.
- Huang, X., Li, Z., He, G., Zhou, M., and Shechtman, E. Self forcing: Bridging the train-test gap in autoregressive video diffusion. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025b. Spotlight; arXiv:2506.08009.
- Huang, Z., He, Y., Yu, J., Zhang, F., Si, C., Jiang, Y., Zhang, Y., Wu, T., Jin, Q., Chanpaisit, N., Wang, Y., Chen, X., Wang, L., Lin, D., Qiao, Y., and Liu, Z. VBench: Comprehensive benchmark suite for video generative models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- Hutchins, D., Schlag, I., Wu, Y., Dyer, E., and Neyshabur, B. Block-recurrent transformers. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. arXiv:2203.07852.
- Ji, S., Chen, X., Yang, S., Tao, X., Wan, P., and Zhao, H. MemFlow: Flowing adaptive memory for consistent and efficient long video narratives. *arXiv preprint arXiv:2512.14699*, 2025.
- Katharopoulos, A., Vyas, A., Pappas, N., and Fleuret, F. Transformers are RNNs: Fast autoregressive transformers with linear attention. In *International Conference on Machine Learning (ICML)*, 2020.
- Kazemnejad, A., Padhi, I., Natesan Ramamurthy, K., Das, P., and Reddy, S. The impact of positional encoding on length generalization in transformers. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. arXiv:2305.19466.
- Kim, Y., Hu, Q., Kuo, C.-C. J., and Beerel, P. A. MemRoPE: Training-free infinite video generation via evolving memory tokens. *arXiv preprint arXiv:2603.12513*, 2026.
- Li, D., Fang, Y., Chen, Y., Yang, S., Cao, S., Wong, J., Luo, M., Wang, X., Yin, H., Gonzalez, J. E., Stoica, I., Han, S., and Lu, Y. WorldModelBench: Judging video generation models as world models. In *Advances in Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks Track*, 2025a. arXiv:2502.20694.
- Li, H., Liu, S., Lin, Z., and Chandraker, M. Rolling sink: Bridging limited-horizon training and open-ended testing in autoregressive video diffusion. *arXiv preprint arXiv:2602.07775*, 2026a.
- Li, J., Tang, J., Xu, Z., Wu, L., Zhou, Y., Shao, S., Yu, T., Cao, Z., and Lu, Q. Hunyuan-GameCraft: High-dynamic interactive game video generation with hybrid history condition. *arXiv preprint arXiv:2506.17201*, 2025b.
- Li, J., Fu, X., Peng, X., Chen, W., Zheng, Y., Zhao, T., Wang, J., Chen, F., Wang, X., and So, H. K.-H. Train short, inference long: Training-free horizon extension for autoregressive video generation. *arXiv preprint arXiv:2602.14027*, 2026b.
- Li, K., Shah, M., and Shang, Y. PackCache: A training-free acceleration method for unified autoregressive video generation via compact KV-cache. *arXiv preprint arXiv:2601.04359*, 2026c.
- Li, R., Torr, P., Vedaldi, A., and Jakab, T. VMem: Consistent interactive video scene generation with surfel-indexed view memory. In *International Conference on Computer Vision (ICCV)*, 2025c. arXiv:2506.18903.
- Li, R., Yang, T., Ai, F., Wu, T., Wen, S., Peng, B., and Zhang, L. Long-horizon streaming video generation via hybrid attention with decoupled distillation. *arXiv preprint arXiv:2604.10103*, 2026d.
- Li, W., Pan, W., Luan, P.-C., Gao, Y., and Alahi, A. Stable video infinity: Infinite-length video generation with error recycling. *arXiv preprint arXiv:2510.09212*, 2025d.
- Li, Y., Huang, Y., Yang, B., Venkitesh, B., Locatelli, A., Ye, H., Cai, T., Lewis, P., and Chen, D. SnapKV: LLM knows what you are looking for before generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. arXiv:2404.14469.
- Lian, K., Cai, S., Du, Y., and Liang, Y. Toward memory-aided world models: Benchmarking via spatial consistency. *arXiv preprint arXiv:2505.22976*, 2025. Loop-based Minecraft navigation benchmark for spatial consistency in world models.
- Liu, K., Hu, W., Xu, J., Shan, Y., and Lu, S. Rolling forcing: Autoregressive long video diffusion in real time. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2509.25161.
- Liu, Z., Yuan, J., Jin, H., Zhong, S., Xu, Z., Braverman, V., Chen, B., and Hu, X. KIVI: A tuning-free asymmetric 2-bit quantization for KV cache. In *International Conference on Machine Learning (ICML)*, 2024. arXiv:2402.02750.

- Lorraine, J. *Scalable nested optimization for deep learning*. PhD thesis, University of Toronto, 2024. Ph.D. thesis; arXiv:2407.01526.
- Lorraine, J. and Hossain, S. JacNet: Learning functions with structured Jacobians. In *ICML Workshop on Invertible Neural Networks and Normalizing Flows (INNF)*, 2019.
- Lorraine, J., Anderson, N., Lee, C., De Laroussilhe, Q., and Hassen, M. Task selection for AutoML system evaluation. *arXiv preprint arXiv:2208.12754*, 2022a.
- Lorraine, J., Vicol, P., Parker-Holder, J., Kachman, T., Metz, L., and Foerster, J. Lyapunov exponents for diversity in differentiable games. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2022b.
- Lorraine, J., Xie, K., Zeng, X., Lin, C.-H., Takikawa, T., Sharp, N., Lin, T.-Y., Liu, M.-Y., Fidler, S., and Lucas, J. ATT3D: Amortized text-to-3D object synthesis. In *International Conference on Computer Vision (ICCV)*, 2023.
- Ma, Y., Zheng, X., Xu, J., Xu, X., Ling, F., Zheng, X., Kuang, H., Li, H., Wang, X., Xiao, X., Chao, F., and Ji, R. Flow caching for autoregressive video generation. *arXiv preprint arXiv:2602.10825*, 2026.
- Mao, W., Lin, X., Huang, W., Xie, Y., Fu, T., Zhuang, B., Han, S., and Chen, Y. TriAttention: Efficient long reasoning with trigonometric KV compression. *arXiv preprint arXiv:2604.04921*, 2026a.
- Mao, X., Rui, S., Ying, K., Zheng, B., Li, C., Chi, M., and Zhang, K. PackForcing: Short video training suffices for long video sampling and long context inference. *arXiv preprint arXiv:2603.25730*, 2026b.
- Mehta, N., Lorraine, J., Masson, S., Arunachalam, R., Perwaiz Bhat, Z., Lucas, J., and George Zachariah, A. Improving hyperparameter optimization with checkpointed model weights. In *ECCV Workshop on Efficient Deep Learning Foundation Models (EFM)*, 2024.
- Micheli, V., Alonso, E., and Fleuret, F. Transformers are sample-efficient world models. In *International Conference on Learning Representations (ICLR)*, 2023. Notable Top 5%; arXiv:2209.00588.
- Mohtashami, A. and Jaggi, M. Random-access infinite context length for transformers. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. arXiv:2305.16300.
- Munkhdalai, T., Faruqui, M., and Gopal, S. Leave no context behind: Efficient infinite context transformers with Infini-attention. *arXiv preprint arXiv:2404.07143*, 2024.
- NVIDIA. KVPress: A compression library for transformer KV caches. <https://github.com/NVIDIA/kvpress>, 2024.
- NVIDIA. Cosmos world foundation model platform for physical AI. *arXiv preprint arXiv:2501.03575*, 2025.
- Oshima, Y., Iwasawa, Y., Suzuki, M., Matsuo, Y., and Furuta, H. WorldPack: Compressed memory improves spatial consistency in video world modeling. *arXiv preprint arXiv:2512.02473*, 2025.
- Peng, B., Quesnelle, J., Fan, H., and Shippole, E. YaRN: Efficient context window extension of large language models. In *International Conference on Learning Representations (ICLR)*, 2024. arXiv:2309.00071.
- Press, O., Smith, N. A., and Lewis, M. Train short, test long: Attention with linear biases enables input length extrapolation. In *International Conference on Learning Representations (ICLR)*, 2022.
- Rae, J. W., Potapenko, A., Jayakumar, S. M., Hillier, C., and Lillicrap, T. P. Compressive transformers for long-range sequence modelling. In *International Conference on Learning Representations (ICLR)*, 2020.
- Ranganath, S., Menon, V., and Patnaik, A. KV cache quantization for self-forcing video generation: A 33-method empirical study. *arXiv preprint arXiv:2603.27469*, 2026.
- Richter-Powell, J., Torralba, A., and Lorraine, J. Score distillation sampling for audio: Source separation, synthesis, and beyond. In *ICML Workshop on AI Heard That!*, 2025. arXiv:2505.04621.
- Robbyant Team, Gao, Z., Wang, Q., Zeng, Y., Zhu, J., Cheng, K. L., Li, Y., Wang, H., Xu, Y., Ma, S., et al. Advancing open-source world models. *arXiv preprint arXiv:2601.20540*, 2026.
- Samuel, D., Tzachor, I., Levy, M., Green, M., Chechik, G., and Ben-Ari, R. Fast autoregressive video diffusion and world models with temporal cache compression and sparse attention. *arXiv preprint arXiv:2602.01801*, 2026.
- Sand AI. MAGI-1: Autoregressive video generation at scale. *arXiv preprint arXiv:2505.13211*, 2025.
- Shang, N., Zhang, L. L., Wang, S., Zhang, G., Lopez, G., Yang, F., Chen, W., and Yang, M. LongRoPE2: Near-lossless LLM context window scaling. In *International Conference on Machine Learning (ICML)*, 2025. arXiv:2502.20082.
- Song, K., Chen, B., Simchowitz, M., Du, Y., Tedrake, R., and Sitzmann, V. History-guided video diffusion. In *International Conference on Machine Learning (ICML)*, 2025a. arXiv:2502.06764.

- Song, Y., Lorraine, J., Nie, W., Kreis, K., and Lucas, J. Multi-student diffusion distillation for better one-step generators. In *ICML Workshop on Efficient Systems for Foundation Models (ES-FoMo)*, 2025b. arXiv:2410.23274.
- Stapf, S., Acuaviva, P., Davtyan, A., and Favaro, P. Composition of memory experts for diffusion world models. In *International Conference on Learning Representations (ICLR)*, 2026. OpenReview: <https://openreview.net/forum?id=sUEdpZCHdp>.
- Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., and Liu, Y. RoFormer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.
- Sun, W., Zhang, H., Wang, H., Wu, J., Wang, Z., Wang, Z., Wang, Y., Zhang, J., Wang, T., and Guo, C. World-Play: Towards long-term geometric consistency for real-time interactive world modeling. *arXiv preprint arXiv:2512.14614*, 2025a.
- Sun, Y., Dong, L., Patra, B., Ma, S., Huang, S., Benhaim, A., Chaudhary, V., Song, X., and Wei, F. A length-extrapolatable transformer. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023. arXiv:2212.10554.
- Sun, Y., Li, X., Dalal, K., Xu, J., Vikram, A., Zhang, G., Dubois, Y., Chen, X., Wang, X., Koyejo, S., Hashimoto, T., and Guestrin, C. Learning to (learn at test time): RNNs with expressive hidden states. In *International Conference on Machine Learning (ICML)*, 2025b.
- Tang, J., Zhao, Y., Zhu, K., Xiao, G., Kasikci, B., and Han, S. Quest: Query-aware sparsity for efficient long-context LLM inference. In *International Conference on Machine Learning (ICML)*, 2024. arXiv:2406.10774.
- Tian, Y., Wang, Z., Peng, Y., Yuan, A., Wang, Z., Yi, B., Liu, X., Cui, Y., and Yang, T. KeepKV: Achieving periodic lossless KV cache compression for efficient LLM inference. In *AAAI Conference on Artificial Intelligence*, 2026. arXiv:2504.09936.
- Valevski, D., Leviathan, Y., Arar, M., and Fruchter, S. Diffusion models are real-time game engines. In *International Conference on Learning Representations (ICLR)*, 2025.
- Wan Team. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025.
- Wang, H., Liu, Q., Du, C., Zhu, T., Du, C., Kawaguchi, K., and Pang, T. When precision meets position: BFloat16 breaks down RoPE in long-context training. *Transactions on Machine Learning Research (TMLR)*, 2025. arXiv:2411.13476.
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- Wang, Z., Lorraine, J., Wang, Y., Su, H., Zhu, J., Fidler, S., and Zeng, X. LLaMA-Mesh: Unifying 3D mesh generation with language models. *arXiv preprint arXiv:2411.09595*, 2024.
- Wang, Z., Lin, H., Yoon, J., Cho, J., Zhang, Y., and Bansal, M. AnchorWeave: World-consistent video generation with retrieved local spatial memories. *arXiv preprint arXiv:2602.14941*, 2026a.
- Wang, Z., Liu, Z., Li, J., Huang, K., Xu, B., Kang, F., An, M., Wang, P., Jiang, B., Wei, Y., Xietian, Y., Pei, J., Hu, L., Jiang, B., Xue, H., Wang, Z., Sun, H., Li, W., Ouyang, W., He, X., Liu, Y., Li, Y., and Zhou, Y. Matrix-game 3.0: Real-time and streaming interactive world model with long-horizon memory. *arXiv preprint arXiv:2604.08995*, 2026b.
- Wei, X., Liu, X., Zang, Y., Dong, X., Zhang, P., Cao, Y., Tong, J., Duan, H., Guo, Q., Wang, J., Qiu, X., and Lin, D. VideoRoPE: What makes for good video rotary position embedding? In *International Conference on Machine Learning (ICML)*, 2025. Oral; arXiv:2502.05173.
- Wu, R., He, X., Cheng, M., Yang, T., Zhang, Y., Kang, Z., Cai, X., Wei, X., Guo, C., Li, C., and Cheng, M.-M. Infinite-world: Scaling interactive world models to 1000-frame horizons via pose-free hierarchical memory. *arXiv preprint arXiv:2602.02393*, 2026a.
- Wu, S., Lin, H., Chen, S., Lu, T., Zeng, Y., Zhan, S., Zheng, H.-T., and Kim, H.-G. Semanticache: Efficient kv cache compression via semantic chunking and clustered merging. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 19562–19566. IEEE, 2026b.
- Wu, T., Yang, S., Po, R., Xu, Y., Liu, Z., Lin, D., and Wetzstein, G. Video world models with long-term spatial memory. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. arXiv:2506.05284, project page <https://spmем.github.io>.
- Wu, X., Paschalidou, D., Gao, J., Torralba, A., Leal-Taixé, L., Russakovsky, O., Fidler, S., and Lorraine, J. Motion attribution for video generation. In *International Conference on Machine Learning (ICML)*, 2026c. arXiv:2601.08828.
- Wu, Y., Rabe, M. N., Hutchins, D., and Szegedy, C. Memorizing transformers. In *International Conference on Learning Representations (ICLR)*, 2022.

- Xi, H., Yang, S., Zhao, Y., Li, M., Cai, H., Li, X., Lin, Y., Zhang, Z., Zhang, J., Li, X., Xu, Z., Wu, J., Xu, C., Stoica, I., Han, S., and Keutzer, K. Quant VideoGen: Auto-regressive long video generation via 2-bit KV-cache quantization. In *International Conference on Machine Learning (ICML)*, 2026. arXiv:2602.02958.
- Xiang, C., Liu, J., Zhang, J., Yang, X., Fang, Z., Wang, S., Wang, Z., Zou, Y., Su, H., and Zhu, J. Geometry-aware rotary position embedding for consistent video world model. *arXiv preprint arXiv:2602.07854*, 2026.
- Xiao, C., Zhang, P., Han, X., Xiao, G., Lin, Y., Zhang, Z., Liu, Z., and Sun, M. InfLLM: Training-free long-context extrapolation for LLMs with an efficient context memory. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024a. arXiv:2402.04617.
- Xiao, G., Tian, Y., Chen, B., Han, S., and Lewis, M. Efficient streaming language models with attention sinks. In *International Conference on Learning Representations (ICLR)*, 2024b.
- Xiao, G., Tang, J., Zuo, J., Guo, J., Yang, S., Tang, H., Fu, Y., and Han, S. DuoAttention: Efficient long-context LLM inference with retrieval and streaming heads. In *International Conference on Learning Representations (ICLR)*, 2025a. arXiv:2410.10819.
- Xiao, Z., Lan, Y., Zhou, Y., Ouyang, W., Yang, S., Zeng, Y., and Pan, X. WorldMem: Long-term consistent world simulation with memory. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025b.
- Xie, K., Lorraine, J., Cao, T., Gao, J., Lucas, J., Torralba, A., Fidler, S., and Zeng, X. LATTE3D: Large-scale amortized text-to-enhanced 3D synthesis. In *European Conference on Computer Vision (ECCV)*, 2024.
- Xu, B., Du, Y., Liu, Z., Yang, S., Jiang, Z., Yan, S., Saha, R., Pumarola, A., Wang, W., and Li, P. Sparse forcing: Native trainable sparse attention for real-time autoregressive diffusion video generation. *arXiv preprint arXiv:2604.21221*, 2026a.
- Xu, T.-X., Wang, Z.-X., Wang, G., Hu, L., Zhang, Z., Zhang, P., Zhang, B., and Zhang, S.-H. UCM: Unifying camera control and memory with time-aware positional encoding warping for world models. *arXiv preprint arXiv:2602.22960*, 2026b.
- Yan Team. Yan: Foundational interactive video generation. *arXiv preprint arXiv:2508.08601*, 2025.
- Yang, S., Huang, W., Chu, R., Xiao, Y., Zhao, Y., Wang, X., Li, M., Xie, E., Chen, Y.-C., Lu, Y., Han, S., and Chen, Y. LongLive: Real-time interactive long video generation. In *International Conference on Learning Representations (ICLR)*, 2026a. arXiv:2509.22622.
- Yang, Y., Zhang, T., Huang, W., Chen, J., Wu, B., He, X., Cai, D., Li, B., and Jiang, P.-T. Anchor forcing: Anchor memory and tri-region RoPE for interactive streaming video diffusion. *arXiv preprint arXiv:2603.13405*, 2026b.
- Ye, Y., Lu, X., Jiang, Y., Gu, Y., Zhao, R., Liang, Q., Pan, J., Zhang, F., Wu, W., and Wang, A. J. MIND: Benchmarking memory consistency and action control in world models. *arXiv preprint arXiv:2602.08025*, 2026.
- Yesiltepe, H., Meral, T. H. S., Akan, A. K., Oktay, K., and Yanardag, P. ∞ -RoPE: Action-controllable infinite video generation emerges from autoregressive self-rollout. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026. arXiv:2511.20649.
- Yi, J., Jang, W., Cho, P. H., Nam, J., Yoon, H., and Kim, S. Deep forcing: Training-free long video generation with deep sink and participative compression. *arXiv preprint arXiv:2512.05081*, 2025.
- Yin, T., Zhang, Q., Zhang, R., Freeman, W. T., Durand, F., Shechtman, E., and Huang, X. From slow bidirectional to fast autoregressive video diffusion models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- Yu, J., Bai, J., Qin, Y., Liu, Q., Wang, X., Wan, P., Zhang, D., and Liu, X. Context as memory: Scene-consistent interactive long video generation with memory retrieval. *arXiv preprint arXiv:2506.03141*, 2025a.
- Yu, W., Qian, R., Li, Y., Wang, L., Yin, S., Chakaravarthy, S. S. P., Anthony, D., Ye, Y., Li, Y., Wan, W., and Garg, A. MosaicMem: Hybrid spatial memory for controllable video world models. *arXiv preprint arXiv:2603.17117*, 2026.
- Yu, Y., Wu, X., Hu, X., Hu, T., Sun, Y.-T., Lyu, X., Wang, B., Ma, L., Ma, Y., Wang, Z., and Qi, X. VideoSSM: Autoregressive long video generation with hybrid state-space memory. *arXiv preprint arXiv:2512.04519*, 2025b.
- Zhang, L., Cai, S., Li, M., Wetzstein, G., and Agrawala, M. Frame context packing and drift prevention in next-frame-prediction video diffusion models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025. Spotlight; arXiv:2504.12626.
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. The unreasonable effectiveness of deep features as a perceptual metric. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- Zhang, X., Du, C., Du, C., Pang, T., Gao, W., and Lin, M. SimLayerKV: A simple framework for layer-level KV cache reduction. *arXiv preprint arXiv:2410.13846v1*, 2024.

- Zhang, Z., Sheng, Y., Zhou, T., Chen, T., Zheng, L., Cai, R., Song, Z., Tian, Y., Ré, C., Barrett, C., Wang, Z., and Chen, B. H₂O: Heavy-hitter oracle for efficient generative inference of large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- Zhao, M., He, G., Chen, Y., Zhu, H., Li, C., and Zhu, J. RIFLEx: A free lunch for length extrapolation in video diffusion transformers. In *International Conference on Machine Learning (ICML)*, 2025. arXiv:2502.15894.
- Zhao, Z., Lu, Y., Liu, Z., Song, J., Deng, J., and Patras, I. Relax forcing: Relaxed KV-memory for consistent long video generation. *arXiv preprint arXiv:2603.21366*, 2026.
- Zheng, D., Huang, Z., Liu, H., Zou, K., He, Y., Zhang, F., Gu, L., Zhang, Y., He, J., Zheng, W.-S., Qiao, Y., and Liu, Z. VBench-2.0: Advancing video generation benchmark suite for intrinsic faithfulness. *arXiv preprint arXiv:2503.21755*, 2025.
- Zhou, F., Huang, J., Li, J., Ramanan, D., and Shi, H. PAI-Bench: A comprehensive benchmark for physical AI, 2025a.
- Zhou, Y., Song, S., Liu, B., Xi, Z., Jin, S., Fan, X., Zhang, Z., Li, W., and Huang, X. EliteKV: Scalable KV cache compression via RoPE frequency selection and joint low-rank projection. *arXiv preprint arXiv:2503.01586*, 2025b.

Table 3. **Core notation.** Methods, cache structure, indices, virtual positions, scene-boundary threshold, conventions.

Symbol	Description
<i>Method Names</i>	
WORLDTRACE	Training-free KV-cache framework (Sec. 2); assigns each summary slot a fixed slot-rank virtual position
WORLDTRACE-FIELD	WORLDTRACE variant that compresses history by canonical key averaging (coherence; Sec. 2.3, Def. 2)
WORLDTRACE-LANDMARK	WORLDTRACE variant keeping verbatim scene-entry frames with frozen canonical keys (recall; Sec. 2.4, Eq. (3))
<i>Cache Structure</i>	
N	Total number of AR chunks in a rollout (generation length)
N_s	Number of summary slots in the KV cache
W_r	Number of recent-window slots (verbatim, newest W_r chunks)
L_{train}	Training context length in AR blocks; $N_s + W_r = L_{\text{train}}$
F	Latent frames per AR block (chunk size)
M	Source frames compressed into one summary slot under WORLDTRACE-FIELD
T_{old}	Frames outside the recent window: $T_{\text{old}} = (N - W_r)F$ for $N > L_{\text{train}}$ (linear in N)
s	Summary slot index, $s = 0, \dots, N_s - 1$ (0 = oldest)
\mathcal{R}	Recent-window cache: ordered set of verbatim KVs for the newest W_r blocks
\mathcal{S}	Summary cache: N_s compressed slots indexed by s
K_*	Block just popped from \mathcal{R} when a new chunk is appended; $K_{\text{cx}, K_*, f}^{(k)}$ denotes its canonical key at intra-block frame f
SE	Scene-entry indicator: SE = 1 if cosine-distance spike on canonical- K exceeds τ
J	Number of detected scene-entry landmarks at the current horizon ($J \leq N_s$)
<i>Indices (dummy / iteration)</i>	
k	RoPE frequency-pair index
m	Source-frame iteration index (inside $\sum_{m=1}^M$)
n	AR-chunk iteration index
f	Intra-block frame index, $f \in \{0, \dots, F-1\}$ (per-frame slot of an AR block; Alg. 1)
μ	Algorithm mode selector ($\mu \in \{\text{WORLDTRACE-FIELD}, \text{WORLDTRACE-LANDMARK}\}$, Alg. 1)
<i>Virtual Position Assignment (Slot Indexing, Def. 1)</i>	
q	Absolute timestamp of the current query frame
t	Absolute timestamp of a cached key frame
t_{min}^v	In-distribution lower bound: $\max(0, q - (L_{\text{train}} - 1)F)$
t_{max}^v	In-distribution upper bound for summary slots: $q - W_r F$
t_v	Generic virtual position (no slot index)
$t_v^{(s)}$	Virtual position of slot s : $q - (L_{\text{train}} - 1 - s)F$
<i>Scene-Boundary Detection</i>	
τ	Scene-boundary detection threshold ($\tau=0.15$; used by WORLDTRACE-LANDMARK)
<i>Experimental Conventions</i>	
n	Number of distinct initial scenes (test instances) per evaluation condition
seed	Global random seed controlling initial scene selection and denoising noise
$N_s + W_r = L_{\text{train}}$	Capacity constraint: total cache always fills the training window

Appendices

A. Notation

Symbols used throughout the paper. Tab. 3 lists model and cache notation, indices, and virtual positions; Tab. 4 lists RoPE and key representations. Evaluation metrics: Tab. 16 (App. F.5); LoopMem parameters: App. E.

Table 4. **RoPE and key representations.** Parameters, rotations, and key/query variants used across the equations.

Symbol	Description
<i>RoPE Parameters</i>	
θ	RoPE base frequency ($\theta=10000$ for MG2)
θ_k	RoPE angular frequency for temporal head-dimension pair k ; $\theta_k = \theta^{-k/c_t}$
c_t	Number of temporal RoPE complex pairs per head
c_h	Number of height-spatial RoPE complex pairs per head
c_w	Number of width-spatial RoPE complex pairs per head
n_ℓ	Number of transformer layers
$R(\alpha)$	Rotation matrix by angle α (applied per frequency pair)
<i>Key / Query Variants</i>	
$K_t^{(k)}$	RoPE-rotated key at absolute position t_m , frequency pair k
$K_{\text{cx},m}^{(k)}$	Canonical (unrotated) key content: $R(-\theta_k t_m) K_t^{(k)}$
$\bar{K}_{\text{cx}}^{(k)}$	Canonical mean across M source frames (implicit operator output of WORLDTRACE-FIELD)
$\bar{K}_{\text{naive}}^{(k)}$	Naive RoPE-space average of M rotated keys
$K_{\text{field}}^{(k)}(t_v)$	WORLDTRACE-FIELD compressed key at virtual position t_v (Def. 2)
$K_{\text{land}}^{(k)}(t_v^{(s)})$	WORLDTRACE-LANDMARK frozen canonical key at slot s (Eq. (3))
t_{ℓ^*}	Original timestamp of a selected landmark frame
$K_t^{(k)}$	Landmark source key at t_{ℓ^*} , frequency pair k
$Q_q^{(k)}$	RoPE-rotated query at q , frequency pair k
<i>Attention Quantities</i>	
$\delta_{q,t}$	Query-key temporal offset: $q - t$
Δt_{train}	Largest temporal offset within the local attention window during training (latent frames; e.g., $\Delta t_{\text{train}}=5$ for MG2). Distinct from the training cache extent $(L_{\text{train}} - 1)F$
$\ell_k(q, t)$	Attention-logit contribution from frequency pair k at offset $\delta_{q,t}$
$A_{q,t}^{(k)}$	Query-key content term in canonical coordinates for $\ell_k(q, t)$
<i>Math Operators</i>	
$\text{Re}(\cdot)$	Real part of a complex expression (used in $\ell_k(q, t)$)
e, i	Euler’s number and imaginary unit (italic; complex exponentials in Eq. (2))
sp	Spatial-axis label superscript ($e^{i\theta^{\text{sp}} \cdot (h,w)}$, App. C)

A.1. Core Notation

A.2. RoPE and Key Representations

B. Related Work

B.1. Memory-Augmented Transformers

WORLDTRACE inherits the two-tier (recent verbatim window plus compressed older history) pattern from a lineage of memory-augmented transformers. Transformer-XL (Dai et al., 2019) pairs segment-level recurrence with a relative position encoding so recurrence does not corrupt absolute indices; Compressive Transformers (Rae et al., 2020) add a compressed memory of older tokens on top, and the two-tier structure we adopt. Memorizing Transformers (Wu et al., 2022) attach an external k NN-retrieved memory; Recurrent Memory Transformer (Bulatov et al., 2022) and Block-Recurrent Transformers (Hutchins et al., 2022) pass learned memory tokens between segments; Infini-attention (Munkhdalai et al., 2024) fuses a sliding window with a compressive long-term memory in one block. Landmark Attention (Mohtashami & Jaggi, 2023) introduces *landmark* tokens that gate cross-block retrieval; this antecedent shares vocabulary with our WORLDTRACE-LANDMARK variant but differs on three axes: their landmarks are *trained* representatives that gate attention to off-cache blocks at 1D position, while WORLDTRACE-LANDMARK is training-free, stores verbatim canonical-frame keys inside the $O(1)$ cache, and freezes them against unrotate–rerotate drift in 3D RoPE. The broader full-attention-to-compressed-state spectrum spans linear attention (Katharopoulos et al., 2020), Mamba (Gu & Dao, 2024), and Test-Time Training (Sun et al., 2025b); we target the fixed-budget, training-free regime in AR video diffusion.

B.2. KV Cache Compression for Large Language Models

Window and eviction methods. Window methods (Beltagy et al., 2020) retain the most recent w tokens; StreamingLLM (Xiao et al., 2024b) augments windows with initial “sink” tokens. Token eviction methods select high-importance tokens via accumulated attention mass (H₂O (Zhang et al., 2023)), key L2-norm (KnormPress (NVIDIA, 2024)), pooled attention (SnapKV (Li et al., 2024)), per-layer budget (PyramidKV (Cai et al., 2025)), lazy-layer drop (SimLayerKV (Zhang et al., 2024)), or query-aware page criticality (Quest (Tang et al., 2024)).

Token merging methods. Merging methods combine similar tokens rather than evicting them: ToMe (Bolya et al., 2023), KVCompose (Akulov et al., 2025) (composite tokens), SemantiCache (Wu et al., 2026b) (semantic clusters), and KeepKV (Tian et al., 2026) (lossless via attention-score adjustment). Pre-RoPE Q and K vectors concentrate around stable directional centers in LLMs; TriAttention (Mao et al., 2026a) scores each token by angular alignment plus magnitude and retains the top-scoring ones verbatim.

Quantization, architectural, and retrieval alternatives. KIVI (Liu et al., 2024) quantizes the cache (orthogonal to and composable with our position-content axis); on the video side, Quant VideoGen (Xi et al., 2026) achieves 2-bit KV quantization with semantic-aware smoothing for AR video diffusion, and Ranganath et al. (2026) systematically benchmark 33 quantization variants under self-forcing rollouts; quantization is empirically decoupled from the position-OD failure mode WORLDTRACE addresses. DuoAttention (Xiao et al., 2025a) partitions heads into retrieval (full-cache) and streaming (constant-size + sinks) families, parallel to our verbatim-recent vs. compressed-summary split; InFLM (Xiao et al., 2024a) combines sliding-window attention with memory-unit retrieval for training-free long-context extrapolation. RULER (Hsieh et al., 2024) is the standard synthetic retrieval/aggregation benchmark on the LLM side; our LoopMem (App. E) extends this paradigm to closed-loop video generation.

Position handling in KV compression. Most LLM-side methods operate only on content and do not address what virtual position to assign a summary representing multiple merged tokens. EliteKV (Zhou et al., 2025b) identifies per-head frequency preferences for low-rank compression but compresses individual tokens, so position reassignment does not arise; A²ATS (He et al., 2025a) decouples positional dependency for retrieval rather than reassigning positions for summaries. The unrotate–rerotate primitive appears in FINCH (Corallo & Papotti, 2024) and the KeyRotationPress module of KVPress (NVIDIA, 2024), both applied per-retained-token after compression; our contribution is its use as an *averaging* operator over multiple canonical keys (WORLDTRACE-FIELD), coupled with the slot-rank positions of WORLDTRACE.

Why KV compression differs in our setting. The KV context spans separately denoised AR blocks rather than a single token stream, and the RoPE-OD failure mode is benign for LLM deployments operating within the trained window (the orthogonal long-context-extension problem is in the next subsection). The angular-concentration prerequisite of TriAttention does not hold under denoising-timestep-conditioned Q : on MG2, TriAttention collapses to a norm-based eviction that matches canonical key averaging or underperforms both it and sliding-window eviction. Layer-allocation methods (PyramidKV, SimLayerKV) reallocate budget across layers but do not reassign positions.

B.3. Position Extrapolation for RoPE

LLM-side analysis. A long line of work addresses the same mechanism we exploit: RoPE (Su et al., 2024) is not robust to relative offsets beyond the training distribution. Position Interpolation (Chen et al., 2023) down-scales position indices; YaRN (Peng et al., 2024) refines this with NTK-by-parts scaling; LongRoPE (Ding et al., 2024) reaches 2M-token context via non-uniform per-dimension interpolation; LongRoPE2 (Shang et al., 2025) attributes residual OOD to undertrained higher RoPE dimensions, aligned with the per-frequency severity we measure in App. F.2. ALiBi (Press et al., 2022) replaces RoPE with a fixed linear distance bias; xPos (Sun et al., 2023) adds an exponential-decay term; Kazemnejad et al. (2023) (NoPE) show position-free transformers can outperform RoPE/ALiBi/APE on length generalization (clarifying that our RoPE-OD argument concerns models *already* trained with RoPE). CoPE (Golovneva et al., 2024) makes positions content-conditional; Barbero et al. (2025) dissect which RoPE frequencies carry position vs. semantics. HoPE (Chen et al., 2025b) and FoPE (Hua et al., 2025) perform per-frequency RoPE-OD analysis at the LLM scale via cascade-failure and Non-Uniform Discrete Fourier Transform theory, respectively, anchoring our per-frequency view (App. F.2).

Vision and video. RoPE-ViT (Heo et al., 2024) adapts RoPE to vision transformers via a 2D RoPE-Mixed split; VideoRoPE (Wei et al., 2025) introduces a 3D RoPE structure with low-frequency temporal allocation and the V-NIAH-D long-context retrieval test for video LLMs; ViewRoPE (Xiang et al., 2026) replaces screen-space coordinates with camera-ray

geometry to maintain loop-closure consistency in pose-conditioned video world models; RIFLEx (Zhao et al., 2025) achieves training-free length extrapolation in video diffusion via per-frequency intrinsic-frequency reduction.

Position-content coupling. Wang et al. (2025) show position couples to numerical precision (bf16 deviates RoPE from its intended relative encoding over long contexts), consistent with our position-content coupling thesis (Sec. 2.5). On the video side, UCM (Xu et al., 2026b) reassigns 3D positional encodings via time-aware warping for camera-controlled world models, and MosaicMem (Yu et al., 2026) introduces “Warped RoPE” that reprojects positional encodings of memory patches into the target view. Both are virtual-position-assignment operators in our sense, but their warping is camera/geometry-driven and paired with trained content writers; ours is slot-rank-driven and training-free.

Why RoPE extrapolation differs in our setting. These works adapt position embeddings for LLMs that read a single growing sequence (or for video LLMs operating on a fixed-length input); our setting differs in two structural ways: (i) inference proceeds across many independently denoised AR video blocks rather than a single autoregressive token stream, so the cache must be *compressed* rather than just *re-mapped*, and (ii) compressing M source frames into one slot raises the question of *which* virtual position to assign that summary, a question that does not arise when each cached token retains its own identity. WORLDTRACE addresses (ii) by assigning each slot a fixed slot-rank offset.

B.4. Autoregressive Video World Models

Lineage and training paradigms. Recent autoregressive video diffusion models (Brooks et al., 2024; Bruce et al., 2024; Google DeepMind, 2024; 2025; He et al., 2025b; Decart et al., 2024; Valevski et al., 2025; Guo et al., 2025; Li et al., 2025b; Yan Team, 2025) generate interactively via AR chunk prediction. The dominant training paradigm Self-Forcing (Huang et al., 2025b) (built on Diffusion Forcing (Chen et al., 2024), DFoT (Song et al., 2025a), and CausVid (Yin et al., 2025)) rolls out on the student’s own KV cache, inheriting a local attention window not designed for arbitrarily long inference. Self-Forcing++ (Cui et al., 2026b) extends this to multi-minute video via long-rollout self-distillation; FAR (Gu et al., 2025) reformulates AR video as next-frame prediction over an unbounded cache; Vid2World (Huang et al., 2026b) causalizes bidirectional video diffusion into interactive world models; LIVE (Huang et al., 2026a) stabilizes long horizons via a cycle-consistency objective. Our work is complementary, operating purely at inference time on top of an already-trained model. Open foundation backbones (Wan Team, 2025; NVIDIA, 2025; Sand AI, 2025) provide the capacity that makes long-horizon AR rollout viable.

World-model and RL lineage. The world-model framing traces to Ha & Schmidhuber (2018) through the Dreamer line (Hafner et al., 2025); transformer-based world models like IRIS (Micheli et al., 2023) and DIAMOND (Alonso et al., 2024) establish the autoregressive transformer as a sample-efficient world-model backbone, and open foundation platforms (NVIDIA, 2025; Wan Team, 2025) provide the capacity our inference-time intervention runs on.

B.5. Memory and KV Cache for Video World Models

Existing memory approaches. Existing systems use one of three cache approaches. *Sliding-window KV caches* (MG2 (He et al., 2025b), LongLive (Yang et al., 2026a)) bound memory at $O(1)$ but discard history and leave cached positions OOD. *Memory re-encoding* (MG3 (Wang et al., 2026b)) selects past frames via field-of-view similarity and re-encodes at each AR step, sidestepping both problems at the cost of roughly doubled per-step latency; FramePack (Zhang et al., 2025) compresses input frame contexts by frame-wise importance, and WorldMem (Xiao et al., 2025b) attaches a state-aware memory of pose-tagged frames re-injected via auxiliary attention. *Architectural memory modules* include StreamingT2V (Henschel et al., 2025), which pairs a short conditional-attention window with a long-term appearance-preservation module anchored on the first chunk.

Concurrent training-free cache work. Closest to ours, four concurrent training-free works share our diagnosis (temporal RoPE OOD as the long-horizon bottleneck) but pair it with a different content writer: Infinity-RoPE (Yesiltepe et al., 2026) uses Block-Rel offsets with KV Flush (sink + most recent block only); FAR (Gu et al., 2025) introduces FlexRoPE for temporal-decay extrapolation on an unbounded $O(T)$ cache; Deep Forcing (Yi et al., 2025) dedicates $\sim 50\%$ of the window to “Deep Sink” tokens with re-aligned RoPE phases; MemRoPE (Kim et al., 2026) pairs Block-Rel positions with a dual-rate EMA summary cache. The first three either operate on existing tokens (verbatim or sink-aligned) or scale to unbounded caches; WORLDTRACE, instead, assigns positions to *compressed summaries* that represent multiple merged frames within a fixed budget. Deep Forcing’s importance-pruning is content-side and composes with slot-rank. Tab. 8 shows that transplanting WORLDTRACE positions into MemRoPE without retuning its EMA content regime regresses

performance, confirming position and content must be jointly designed. PackCache (Li et al., 2026c) is also training-free and uses a “spatially preserving position embedding” for cache removal; it applies virtual position assignment per-token rather than per-summary-slot, paired with cross-frame decay rather than canonical-key averaging or verbatim landmarks. FlowCache (Ma et al., 2026) introduces chunk-wise denoising-step caching and an importance–redundancy KV compression scheme for autoregressive video; its goal is per-step compute reduction rather than long-horizon recall, but the chunk-aware cache structure is similar to the recent–summary split we adopt. TempCache (Samuel et al., 2026) merges near-duplicate cached keys across AR frames via approximate-nearest-neighbor temporal correspondence, paired with sparse self- and cross-attention; merging operates at per-token similarity level rather than reassigning summary-slot virtual positions in canonical space.

Closest concurrent training-time analogs. Two contemporaneous works share WORLDTRACE’s diagnosis and propose related fixes through training-time mechanisms. Grounded Forcing (Chen et al., 2026b) introduces a Dual Memory KV Cache (Local Temporal Memory plus Global Consistency Memory) coupled with Dual-Reference RoPE Injection that stores raw pre-RoPE keys and injects fixed $t=0$ for global anchors and relative offsets for local frames; their dual-reference indexing is the training-time analog of our slot-rank virtual positions. Anchor Forcing (Yang et al., 2026b) partitions the cache into sink, junction, and local regions, each with its own RoPE reference origin capped at the pretrained limit, and learns the assignment via RoPE re-alignment distillation. WORLDTRACE differs in being inference-time-only with no distillation or retraining, in decomposing retention into orthogonal coherence (WORLDTRACE-FIELD) and verbatim recall (WORLDTRACE-LANDMARK) operators, selectable per task, and in pinning summary-slot position rather than content as the binding constraint. Composition of Memory Experts (Stapf et al., 2026), the long-term spatial-memory framework of Wu et al. (2025), and Mixture of Contexts (Cai et al., 2026) provide training-time decompositions: a contrastive product-of-experts over short-term, long-term episodic, and spatial memory; point-cloud spatial memory plus sparse episodic keyframes; and sparse top- k routing with mandatory text/intra-shot anchors. Context Forcing (Chen et al., 2026d) addresses the same student-teacher mismatch by training a long-context student under a long-context teacher with a Slow-Fast Memory architecture that bounds attention cost while preserving 20 s+ effective context. Stable Video Infinity (Li et al., 2025d) bridges the train-test hypothesis gap via Error-Recycling Fine-Tuning that injects historical errors as supervisory prompts so the DiT learns to correct its own drift. Hybrid Forcing (Li et al., 2026d) combines a linear-attention summary state for evicted tokens with block-sparse local attention through a decoupled-distillation pipeline. All three are training-time analogs that complement WORLDTRACE’s inference-time-only operation; WORLDTRACE can, in principle, compose with any of them at the cache level.

Concurrent trained video-memory architectures. A second cluster of concurrent works addresses the same long-horizon memory problem but with trained components. RELIC (Hong et al., 2025) stores compressed historical latents in the KV cache with both relative-action and absolute-camera-pose annotations (camera-aware memory); MosaicMem (Yu et al., 2026) pairs Warped RoPE (geometry-driven virtual position assignment) with Warped Latent injection; UCM (Xu et al., 2026b) reassigns 3D positional encodings via time-aware warping for camera-controlled world models; PackForcing (Mao et al., 2026b) partitions the cache into sink/mid-compressed/recent and applies a Continuous Temporal RoPE Adjustment to re-align position gaps left by dropped tokens (zero-shot or 5s-clip-trained). Compared to these, WORLDTRACE pursues the strict zero-fine-tune regime, depends on slot-rank rather than camera/geometry signals, and decomposes the cache update into orthogonal coherence (WORLDTRACE-FIELD) and recall (WORLDTRACE-LANDMARK) operators that can be selected per task without retraining.

Other concurrent video-cache and memory-bank work. Other concurrent training-free work spans position-side fixes (LoL (Cui et al., 2026a), FLEX (Li et al., 2026b)), content-side selection (PaFu-KV (Chen et al., 2026a), Rolling Sink (Li et al., 2026a), Relax Forcing (Zhao et al., 2026), Rolling Forcing (Liu et al., 2026)), world-model architectures (WorldPack (Oshima et al., 2025), WorldPlay (Sun et al., 2025a), VideoSSM (Yu et al., 2025b)), and memory-bank designs (Memory Forcing (Huang et al., 2025a), VMem (Li et al., 2025c), Context as Memory (Yu et al., 2025a), AnchorWeave (Wang et al., 2026a), Infinite-World (Wu et al., 2026a), HyDRA (Chen et al., 2026c), MemFlow (Ji et al., 2025), Memorize-When-Needed (Guo et al., 2026), VRAG (Chen et al., 2025a)). Sparse Forcing (Xu et al., 2026a) learns native block-sparse attention plus persistent spatiotemporal anchors via a Persistent Block-Sparse Attention kernel, observing the same “implicit spatiotemporal memory” on persistent KV slots that motivates WORLDTRACE-LANDMARK’s scene-entry detection but training the sparsity end-to-end. MemCam (Gao et al., 2026) pairs a context-compression module with co-visibility-based historical-frame selection for camera-controlled long video generation; the selection is geometry-driven rather than canonical-K-spike-driven, and the framework is trained. MemFlow retrieves the most relevant historical frames per chunk and activates only top- k tokens in attention; Memorize-When-Needed adds a decoupled memory branch with camera-

Table 5. **KV cache comparison.** WORLDTRACE makes compressed memory addressable; WORLDTRACE-FIELD uses a field vs. WORLDTRACE-LANDMARK retaining selected landmarks.

Method	Fixes pos. OOD?	Bounded memory?	Preserves history?	Training-free?
Full KV	No	No ($O(N)$)	Yes (exact)	Yes
Sliding window	No	Yes ($O(1)$)	Recent only	Yes
KV Flush + Block-Rel (Yesiltepe et al., 2026) (KV-Flush variant)	Yes	Yes ($O(1)$)	No (amnesic)	Yes
Naive + Block-Rel	Partial	Yes ($O(1)$)	Corrupted (Sec. 1)	Yes
MemRoPE (Kim et al., 2026)	Partial	Yes ($O(1)$)	EMA-diluted	Yes
FAR (FlexRoPE) (Gu et al., 2025)	Yes	No ($O(T)$)	Yes (exact)	Yes
Deep Forcing (Yi et al., 2025)	Partial	Yes ($O(1)$)	Sink + importance-pruned recent	Yes
KnormEvict + WORLDTRACE	Yes	Yes ($O(1)$)	Partial (1 token/slot)	Yes
WORLDTRACE-FIELD (ours)	Yes	Yes ($O(1)$)	Yes (compressed field)	Yes
WORLDTRACE-LANDMARK (ours)	Yes	Yes ($O(1)$)	Yes (verbatim landmarks)	Yes

aware gating that conditions generation on memory only when meaningful historical references exist; VRAG augments interactive video generation with explicit global-state retrieval to reduce compounding errors. The memory-bank designs store geometry-anchored or retrieval-indexed frames *outside* the standard KV cache and re-inject them via auxiliary attention; WORLDTRACE instead compresses history into a constant-size summary *within* the standard KV cache, supporting both coherence (WORLDTRACE-FIELD) and verbatim recall (WORLDTRACE-LANDMARK) under a single $O(1)$ budget without re-encoding or retrieval.

Evaluation benchmarks for world generation. WorldScore (Duan et al., 2025), MIND (Ye et al., 2026), and VBench-2.0 (Zheng et al., 2025) provide unified long-horizon evaluation; WorldModelBench (Li et al., 2025a) explicitly positions video generation as world modeling. Closer to our setting, Lian et al. (2025) constructs a Minecraft loop-navigation benchmark that assesses spatial consistency during revisits to previously seen locations. Our LoopMem benchmark (App. E) targets the orthogonal axis of *closed-loop* episodic recall (return-to-origin trajectories) on top of pretrained autoregressive video world models, complementing rather than replacing these broader quality benchmarks.

Differentiation summary. Tab. 5 summarizes four axes: whether positional OOD is fixed, whether memory is bounded, whether intermediate history is preserved, and whether retraining is required.

C. WORLDTRACE-FIELD: Additional Properties

This section expands Eq. (2), summarised at the end of Sec. 2.3.

Mean attention preservation. The WORLDTRACE-FIELD compression preserves the mean attention logit: a single summary token $K_{\text{field}}^{(k)}(t_v)$ reproduces exactly the average of the logits that the M individual source keys would produce if they were all relocated to the same virtual position t_v . This holds for every query $Q_q^{(k)}$ and every temporal RoPE frequency k , so no query can distinguish a slot that was built from one frame versus many.

Proposition 1 (Mean attention preservation). *Fix a temporal RoPE pair k and a virtual position t_v . For any query $Q_q^{(k)}$,*

$$\langle Q_q^{(k)}, K_{\text{field}}^{(k)}(t_v) \rangle = \frac{1}{M} \sum_{m=1}^M \langle Q_q^{(k)}, R(\theta_k t_v) K_{\text{cx},m}^{(k)} \rangle.$$

By Eq. (2), $K_{\text{field}}^{(k)}(t_v) = R(\theta_k t_v) \bar{K}_{\text{cx}}^{(k)}$ with $\bar{K}_{\text{cx}}^{(k)} = \frac{1}{M} \sum_m K_{\text{cx},m}^{(k)}$; *linearity of the inner product in its second argument then gives $\langle Q_q^{(k)}, K_{\text{field}}^{(k)}(t_v) \rangle = \langle Q_q^{(k)}, R(\theta_k t_v) \bar{K}_{\text{cx}}^{(k)} \rangle = \frac{1}{M} \sum_m \langle Q_q^{(k)}, R(\theta_k t_v) K_{\text{cx},m}^{(k)} \rangle$. The summary token contributes the mean attention logit that the source keys would produce if their content were first moved to the shared virtual position t_v .*

Concretely, $K_{\text{field}}^{(k)}(t_v)$ acts as if all M source frames were simultaneously repositioned to t_v : no individual timestamp t_m appears in the compressed representation, only the canonical content average $\bar{K}_{\text{cx}}^{(k)}$. This fully decouples the summary content from its temporal placement, enabling WORLDTRACE to freely reassign virtual positions (Def. 1) without distorting the stored signal.

Characterization. Eq. (2) is the linear compression operator that (a) maps M RoPE-encoded keys to a single token at virtual position t_v and (b) satisfies the mean attention preservation property of Prop. 1: unrotate each key to canonical space, average, and re-encode at t_v . The content of the compressed token is fully determined by property (b); any linear operator satisfying it must produce the same canonical average $\bar{K}_{\text{cx}}^{(k)}$, since the property pins the output attention logit for all queries. Only t_v remains free, and WORLDTRACE chooses it (Sec. 2.2). Pinning a learned operator by a structural property of its derivative (here, the block-diagonal phase rotation of RoPE that Eq. (2) inverts) has antecedents in structured-Jacobian

Table 6. **WORLDTRACE-FIELD leads on TempSSIM at both horizons and on Local Drift.** At $N=32$, WORLDTRACE-FIELD has the highest TempSSIM (0.613) though centroid variants have lower Drift; at $N=48$, WORLDTRACE-FIELD leads on both metrics. *Centroid* replaces Block-Rel’s saturating offset by linearly mapping each slot’s mean source-frame timestamp into $[t_{\min}^v, t_{\max}^v]$; *norm* rescales the canonical mean to preserve per-frame norm; *geom* uses geometric group sizes; *knorm* weights frames by their canonical-key L2 norm.

Method	$N=32$		$N=48$	
	TempSSIM \uparrow	Local Drift \downarrow	TempSSIM \uparrow	Local Drift \downarrow
Sliding window (baseline)	0.571	0.0229	0.472	0.0305
Canonical averaging + Block-Rel (uniform)	0.585	0.0215	0.530	0.0339
Canonical averaging + Centroid (uniform)	0.573	0.0211	0.479	0.0297
WORLDTRACE-FIELD (ours)	0.613	0.0250	0.545	0.0295




Figure 5. **WORLDTRACE-FIELD qualitative comparison at $N=48$.** Four position-encoding methods generated from the same $t=1$ conditioning frame: sliding window, Block-Rel, Centroid, and WORLDTRACE-FIELD (ours). Once a baseline diverges, it does not recover. Diverged frames are bordered in red. WORLDTRACE-FIELD retains memory of earlier chunks and remains coherent across the horizon. The camera path across the seven columns is plotted in Fig. 9.

parameterizations (Lorraine & Hossain, 2019); WORLDTRACE-FIELD is the closed-form, training-free analog for the temporal RoPE block. Non-linear alternatives (e.g., selecting the highest-norm token verbatim) satisfy neither property and lose phase coherence across compressed frames. We confirm this empirically: KnormEvict+WORLDTRACE (single highest-norm canonical key per slot) achieves PAC 0.553 vs. 0.574 for WORLDTRACE-FIELD (-3.7% , $p=0.074$, $n=10$, paired at $N=16$), with TempSSIM also significantly lower ($p=0.007$), consistent with information loss from discarding all history except the single highest-norm frame.

Spatial RoPE is unaffected. Spatial RoPE (encoding height h and width w) is fixed across all frames at the same spatial position. The spatial phase $e^{i\theta^{sp}h}$ is identical for tokens (f_1, h, w) and (f_2, h, w) ; it factors out of the sum in Eq. (2) and is preserved exactly. The temporal unrotation and re-rotation in Eq. (2) touch only the first $2c_t$ head dimensions, at 2 FLOP per key per channel pair.

Norm collapse. Averaging M uncorrelated canonical keys reduces the norm by $\sim 1/\sqrt{M}$, suppressing the attention weight of summary tokens in the softmax. A norm-preserving rescale can counteract this; we evaluate it as canonical key averaging+Norm in App. D. The rescale does not improve our method at long horizons, so we omit it.

D. Additional Experimental Results

D.1. Coherence and Ablation Tables

Tab. 6 reports TempSSIM and Local Scene Drift for WORLDTRACE-FIELD and three canonical-averaging baselines at $N=32$ and $N=48$. Fig. 5 shows a qualitative four-method comparison along a fixed camera path at $N=48$ on Matrix-Game-2. Tab. 7 isolates the three design choices of WORLDTRACE on the ABA recall protocol.

Table 7. **Ablation.** PAC (\uparrow) on ABA recall with $N_s+W_r=6$ slots. Top section (Axes 1 & 2): isolates position assignment and canonical vs. naive compression. Bottom section (Axes 1 & 3): isolates frozen landmark keys and position assignment in the verbatim tier. Field averaging with Block-Rel matches sliding-window eviction byte-for-byte (OOD offsets suppress summary-slot attention).

Method	PAC ($N=64$) \uparrow	PAC ($N=128$) \uparrow	PAC ($N=256$) \uparrow
<i>Axes 1 & 2: position assignment & canonical compression (compression tier)</i>			
Sliding window (baseline)	0.412	0.504	0.631
+ Naive averaging (Block-Rel)	0.443	0.486	0.598
+ Field averaging (Block-Rel)	0.412	0.504	0.631
+ WORLDTRACE (WORLDTRACE-FIELD, ours)	0.442	0.495	0.602
<i>Axis 3: frozen vs. unfrozen landmark keys (verbatim recall):</i>			
Landmark + Block-Rel (<i>unfrozen</i>)	0.936	0.940	0.965
WORLDTRACE-LANDMARK (ours, frozen)	0.976	0.986	0.989

Table 8. **Concurrent training-free baselines.** PAC (\uparrow) at two horizons (ABA). YaRN (Peng et al., 2024) uses $O(N)$ memory; only at $N=32$. MemRoPE+WORLDTRACE swaps Block-Rel for WORLDTRACE; dual-rate EMA unchanged.

Method	Position	Content	$N=32$ ($16\times$)	$N=48$ ($24\times$)
<i>$O(N)$ cache:</i>				
Sliding window (baseline)	actual	Sliding window	0.401	0.388
YaRN (Peng et al., 2024)	NTK	Sliding window	0.490	0.412
<i>$O(1)$ cache, Block-Rel positions:</i>				
MemRoPE (Kim et al., 2026)	Block-Rel	dual EMA	0.651	0.706
Landmark + Block-Rel	Block-Rel	verbatim	0.929	0.934
<i>$O(1)$ cache, WORLDTRACE positions:</i>				
MemRoPE + WORLDTRACE positions	WORLDTRACE	dual EMA	0.592	0.662
WORLDTRACE-LANDMARK (ours)	WORLDTRACE	verbatim	0.964	0.972

D.2. Concurrent Training-Free Baselines

Tab. 8 compares MemRoPE (Kim et al., 2026) and YaRN (Peng et al., 2024) on the Sec. 3.3 ABA loops, factoring the position scheme (Block-Rel vs. WORLDTRACE) from the content writer at two horizons.

MemRoPE (Kim et al., 2026) (Block-Rel + dual-rate EMA with $\alpha_{\text{long}}=0.01$ and $\alpha_{\text{short}}=0.1$, matching the original release) survives OOD positions via EMA decay but is still beaten by Landmark+Block-Rel: verbatim canonical keys retain stronger query-key similarity than smoothed averages, and WORLDTRACE-LANDMARK adds in-distribution positions on top. MemRoPE+WORLDTRACE (Block-Rel swapped for WORLDTRACE, EMA kept) degrades ($p<0.001$): write/read offsets disagree.

YaRN (Peng et al., 2024) rescales temporal RoPE frequencies to bring offsets back in-distribution, gaining +22% over the sliding-window baseline at $N=32$ ($p<0.001$) and confirming position is the primary bottleneck; but it needs $O(N)$ memory (OOM by $\sim N=100$) and the rescale diverges with horizon. WORLDTRACE-LANDMARK exceeds it by a wide margin at $N=32$ (0.964 vs. 0.490) in $O(1)$ memory and stays near-constant through $N=512$.

D.3. Cross-Architecture Experiments

Tab. 9 reports PAC at $2\times-8\times$ training horizon on LingBot-Fast (Robbyant Team et al., 2026) with WORLDTRACE applied. WORLDTRACE-FIELD matches or exceeds the sliding-window baseline at all horizons: Plücker camera conditioning already supplies the recall signal that KV-cache content provides on MG2, leaving no room for canonical averaging to add. WORLDTRACE-LANDMARK improves over sliding-window retention from $4\times$ onward (+8.9%, +14.1%, +7.3% at $4\times/6\times/8\times$), with no significant gain at $2\times$ (-0.006 absolute, $p>0.1$): verbatim key injection provides a complementary recall signal that strengthens as camera-pose priors alone become insufficient at longer horizons. A camera-ablation corroborates this: zeroing Plücker embeddings raises sliding-window PAC by +7.0% while depressing WORLDTRACE-FIELD by -23.2% , confirming that canonical averaging dilutes scene-origin content when the pose-based path is removed.

Table 9. **LingBot-Fast (Robbyant Team et al., 2026) PAC at extended horizons.** WORLDTRACE-LANDMARK improves over the sliding-window baseline from 4× onward; WORLDTRACE-FIELD matches the sliding window at all horizons, consistent with Plücker conditioning supplying the primary recall signal.

N	Horizon	Sliding window	WORLDTRACE-FIELD	WORLDTRACE-LANDMARK
14	2×	0.657	0.668	0.651
28	4×	0.624	0.619	0.680
42	6×	0.591	0.620	0.674
56	8×	0.632	0.648	0.678

D.4. Slot Allocation

Summary/recent split sensitivity. The main experiments use $N_s=4$ summary slots and $W_r=2$ recent-window slots (total $N_s+W_r=6$, matching the training context size). To test sensitivity to this split, we swept five allocations ($N_s+W_r=6$, ABA loops at $N=32$ and $N=64$); results are in Tab. 10. $N_s \leq 2$ collapses toward the sliding-window baseline: with one slot the B→A detector overwrites the scene-A landmark; with two slots the B-side traversal evicts it before the return. $N_s=3$ is intermediate (0.652/0.693 at $N=32/64$) but still evicts the landmark. $N_s=4$ retains the landmark through the B-side return; $N_s=5$ adds only +0.009/+0.005, confirming a four-slot plateau.

Table 10. **Slot Allocation.** PAC at $N=32$ and 64 as a function of N_s (total budget fixed at $N_s+W_r=6$). $N_s \leq 2$ collapses toward the sliding-window baseline; four slots is the critical boundary for retaining the scene-A landmark through the B-side traversal.

N_s	W_r	PAC $N=32$	PAC $N=64$
Sliding window (reference)		0.401	0.412
1	5	0.413	0.437
2	4	0.419	0.426
3	3	0.652	0.693
4	2	0.964	0.976
5	1	0.973	0.981

Per-slot depth ablation. To isolate which slot carries the recall signal, we run WORLDTRACE-LANDMARK with only one summary slot active at a time (inactive summary slots revert to sliding-window eviction); results are in Tab. 11. Slot 3 alone recovers most of the four-slot PAC at $N=32$ and $N=48$ (Tab. 11: 0.821 vs. 0.964 at $N=32$; the shortfall vs. all four slots reaches 0.184 at $N=48$). Slot 1 only sits on a shallow plateau (0.630, 0.589); Slot 2 only improves moderately (0.628, 0.554) but remains far below Slot 3, with the gap widening over horizon (+0.193 at $N=32$ and +0.234 at $N=48$, Slot 3 only minus Slot 2 only).

Table 11. **Per-slot depth ablation.** PAC when exactly one summary slot runs WORLDTRACE-LANDMARK and the remaining summary slots revert to sliding-window eviction. Slot 1 only still collapses recall; Slot 2 only improves slightly yet remains far below full recall at longer horizons; **Slot 3 only** recovers most of the gain, and the full four-slot WORLDTRACE-LANDMARK achieves the best recall at every horizon.

Method	$N=32$	$N=48$
Slot 1 only	0.630	0.589
Slot 2 only	0.628	0.554
Slot 3 only	0.821	0.788
WORLDTRACE-LANDMARK (all 4 slots)	0.964	0.972

D.5. Phase Cancellation Validation

Tab. 12 verifies that canonical key averaging avoids the phase-cancellation failure of Sec. 1: as N_s grows, naive averaging in RoPE-rotated space loses progressively more low-frequency content (measured by LatentDiff ↓), while canonical averaging in unrotated space is unaffected.

Table 12. **Phase cancellation: Naive vs. canonical key averaging.** LatentDiff (↓) at $N=16$ for varying N_s . Canonical averaging avoids the phase cancellation that corrupts low frequencies under naive RoPE-space averaging.

N_s	Naive LatentDiff ↓	Canonical LatentDiff ↓
1	0.257	0.224
2	0.261	0.257
4	0.312	0.233

E. LoopMem: Scripted Navigation Benchmark

LoopMem evaluates episodic spatial recall in autoregressive world models: a model executes a scripted navigation path that returns to a previously visited location, and the generated return frame is scored against the original scene appearance at geometrically matched positions. Unlike global video-quality metrics, LoopMem requires no external reference; the forward path generates the target, and the return path is scored against it.

The benchmark organizes difficulty along four axes: (1) **waypoint count** (K): number of intermediate locations between departure and return to A (ABA: $K=1$; ABCDA: $K=3$); (2) **edge length** (L): AR chunks per directed edge, controlling KV-context distance; (3) **camera orientation**: heading on arrival at A relative to departure, testing viewpoint-invariant recall; and (4) **multi-revisit depth** (R): how many times a waypoint is re-entered within a single rollout. The four-axis decomposition follows the broader principle that benchmark suites should isolate task properties capable of producing distinguishable rankings between candidate systems (Lorraine et al., 2022a), here projected onto the position–content axes that the cache mechanism is hypothesized to probe. Fig. 6 shows representative configurations; those evaluated here are marked *.

F. Implementation Details


F.1. Model Architecture and Self-Forcing Background

We evaluate on MG2-1.3B, a 1.3B-parameter distilled autoregressive video world model based on Wan 1.3B T2V (Wan Team, 2025), with 30 transformer layers, 12 attention heads, and a head dimension of 128. Video is generated autoregressively in AR blocks of 3 latent frames each at spatial resolution 44×80 (352×640 pixels), yielding 880 tokens per frame \times 3 frames = 2640 tokens per AR block per layer. Each new block is denoised via a multi-step flow-matching process conditioned on the KV cache of all prior blocks. During training with Self-Forcing (Huang et al., 2025b) (rollout on the student’s own KV cache rather than teacher-forced context; cf. Diffusion Forcing (Chen et al., 2024), CausVid (Yin et al., 2025), and one-step diffusion distillation (Song et al., 2025b)), attention is restricted to a local window of `local_attn_size=6` frames (2 AR blocks), so cross-frame temporal offsets stay $\leq \Delta t_{\text{train}}=5$. At inference, the rolling KV cache grows without bound; with a constant $O(1)$ budget, our method caps it at $L_{\text{train}} \times 2640 \times 12 \times 256 \times 30 \times 2 \approx 2.79$ GB per batch element ($L_{\text{train}}=6$ chunks, 2640 tokens/chunk, fp16; cf. Tab. 14), independent of generation length. At long horizons, the OOD mismatch is not the presence of cached keys at large offsets, but the temporal RoPE rotation through which they are read: keys are stored verbatim, but the model was never trained to invert the phase rotations they accumulate beyond Δt_{train} .


F.2. Rotary Position Embeddings in 3D Video Attention

We inherit the 3D RoPE split of the Wan 2.1 backbone (Wan Team, 2025): the 128-dimensional per-head embedding is split across three position axes, $2c_t=44$ dimensions for temporal position, and $2c_h=2c_w=42$ dimensions each for spatial height and width, with shared base frequency $\theta = 10000$. The k -th temporal frequency component follows $\theta_k = \theta^{-k/c_t}$ for $k = 0, \dots, c_t-1$, giving $\theta_0 = 1.0$ (fastest rotating) down to $\theta_{c_t-1} \approx 1.5 \times 10^{-4}$ (slowest). The attention score contribution from frequency k between the query at temporal position q and the key at t is proportional to $\cos(\theta_k(q-t))$. When $\theta_k|\delta_{q,t}| \ll \pi$, the cosine is near 1, and the component is position-invariant, carrying semantic content; when $\theta_k|\delta_{q,t}| \gg \pi$, it oscillates incoherently and constitutes positional noise. This wavelength-vs-context view follows YaRN’s NTK-by-parts framing (Peng et al., 2024) and is consistent with the mechanistic finding that high-frequency RoPE components carry positional structure while low-frequency components carry semantics (Barbero et al., 2025). Viewed as a long-horizon discrete dynamical system, the rapid-oscillation regime is the position-side analog of the local-divergence behavior studied via Lyapunov exponents in iterated maps (Lorraine et al., 2022b): once $\theta_k|\delta_{q,t}| \gg \pi$ the cosine kernel ceases to be Lipschitz in $\delta_{q,t}$ at training-scale resolution, so neighboring offsets receive uncorrelated attention weights and the recall signal decoheres across the rollout. At training max offset $\Delta t_{\text{train}} = 5$, components $k \geq 10$ satisfy $\theta_k \times 5 < 0.08$ rad and act as near-semantic carriers; at inference max offset $|\delta_{q,t}| = 30$, the three fastest components ($k \leq 2$) all exceed 3π rad, with $k=3$ at 8.5 rad and $k=5$ at 3.7 rad.


Row 1: varying topology (K intermediate waypoints)



Row 2: varying edge length (L chunks per leg, ABA topology)



Row 3: camera-orientation pans (blue = pan away, orange dashed = return)



Row 4: multi-revisit patterns ($R > 1$; solid/lighter arcs = 1st/2nd traversal)




Figure 6. **LoopMem benchmark gallery.** Blue solid: outward; orange dashed: return or palindrome; violet dashed: shortcut omitting intermediate waypoints. **Row 1:** topology (ABA/ABCA/ABCDA). **Row 2:** ABA edge length (longer edges push RoPE further OOD). **Row 3:** camera orientation (agent at A; blue = pan away, orange = return). **Row 4:** multi-revisit ($R > 1$): palindromes (ABCBA), shortcuts (ABCDBA).



Figure 7. **ABA qualitative results.** Each sample group of three frames shows trajectory keyframes at chunks 0 (A) / 7 (B) / 15 (A return) for one initial frame; the two rows compare the sliding-window baseline (top) vs. WORLDTRACE-LANDMARK (bottom). The sliding window drifts away from the scene-A appearance on the return leg in all four samples; WORLDTRACE-LANDMARK anchors the return to the original scene.




Figure 8. **Pan 90° qualitative results.** Camera-orientation Tier 3 ($N=4$): the agent is fixed at A while the camera pans right by $\sim 90^\circ$ then returns. Each sample group of four frames shows keyframes at chunks 0 (0°) / 1 ($\sim 45^\circ$) / 2 ($\sim 90^\circ$) / 3 (back $\sim 45^\circ$); the two rows compare the sliding-window baseline (top) vs. WORLDTRACE-LANDMARK (bottom). WORLDTRACE-LANDMARK restores the initial-view appearance on the return half-pan; the sliding window does not.

Table 13. **RoPE OOD severity by frequency.** Per-frequency RoPE phase at inference $|\delta_{q,t}|=30$ vs. training max $|\delta_{q,t}|\leq 5$ (MG2-1.3B). Low frequencies ($k\leq 2$) hit noise-level phases ($> 3\pi$ rad); high frequencies ($k\geq 10$) stay near in-distribution.

Freq k	θ_k	Train max ($ \delta_{q,t} \leq 5$)	Inference ($ \delta_{q,t} =30$)	Status
$k=0$	1.000	5.0 rad	30.0 rad (9.5π)	Random noise
$k=5$	0.123	0.62 rad	3.70 rad (1.2π)	OOD
$k=10$	0.0152	0.076 rad	0.46 rad	Mild OOD
$k=15$	1.9×10^{-3}	0.0094 rad	0.056 rad	Near in-dist
$k=18$	5.4×10^{-4}	0.0027 rad	0.016 rad	Fully in-dist
$k=21$	1.5×10^{-4}	7.6×10^{-4} rad	0.0046 rad	Fully in-dist

Table 14. **MG2-1.3B model and generation statistics.** OOM horizon: shortest length where full-KV inference OOMs on one A100 80 GB; safe horizon: longest degradation-study length without OOM.

Property	Value
<i>Architecture</i>	
Transformer layers	30
Attention heads	12
Head dimension	128
Temporal RoPE complex pairs (c_t)	22
<i>Video & Latent Space</i>	
Decoded output FPS	16
VAE temporal compression	4×
VAE spatial compression	8× (each axis)
Latent resolution	44 × 80
Spatial tokens per latent frame	22 × 40 = 880
Latent frames per chunk (F)	3
Decoded frames per chunk	12
Chunk duration	0.75 s
<i>Training Context</i>	
Training context length (L_{train})	6 chunks
Training context in latent frames	18
Training window duration	≈ 4.5 s
Tokens per chunk	3 × 880 = 2,640
KV memory per chunk (fp16, all layers)	≈ 465 MB
<i>Inference Horizon</i>	
Safe inference horizon (no compression)	≈ 31 chunks (≈ 23 s)
OOM horizon (full KV cache, A100 80 GB)	≈ 150 chunks (≈ 112 s)
$N=256$ horizon	≈ 192 s (≈ 3 min)
$N=512$ horizon	≈ 384 s (≈ 6 min)

E.3. WORLDTRACE Algorithm and Architectural Baselines

WORLDTRACE cache update. Each autoregressive chunk appends fresh keys and values to the recent window; when that window overflows, the evicted block is merged into the summary tier using either uniform-bucket averaging (WORLDTRACE-FIELD) or scene-entry landmark insertion (WORLDTRACE-LANDMARK), as defined in the main text. Alg. 1 presents the full control flow in pseudocode, including the shared attention-time step that applies per-slot rotations to canonical summary keys using the virtual positions from Def. 1.

Default MG2 KV cache behavior. The stock MG2 inference pipeline implements a fixed-size sliding-window KV cache: each transformer layer maintains a pre-allocated buffer of size $L_{\text{train}} \times 880$ tokens, and when new tokens overflow the buffer, the oldest tokens are evicted without position correction or content compression (first-in-first-out over time). This sliding-window baseline is what we compare against throughout the paper.

E.4. Hyperparameters, Compute, and Evaluation Protocol

Inference hyperparameters. MG2-1.3B uses 3 distilled denoising steps, CFG scale 5.0, a single conditioning frame, and $F=3$ latent frames per AR block. The first chunk uses a clean KV context; later chunks receive the accumulated compressed cache.

WORLDTRACE-FIELD hyperparameters. Short-horizon experiments ($N=8$) use $N_s=2$ summary slots and $W_r=4$ recent-window slots ($N_s + W_r = 6 = L_{\text{train}}$). Long-horizon PAC and ablation experiments (Secs. 3.3–3.3) use $N_s=4$, $W_r=2$ (same total capacity) for both WORLDTRACE-FIELD and WORLDTRACE-LANDMARK, as described in the Setup section. The scripted loop evaluation (Sec. 3.3) uses $N_s=5$, $W_r=1$ for WORLDTRACE-LANDMARK (one additional

Algorithm 1 WORLDTRACE cache update (per AR chunk). \mathcal{S} stores canonical (unrotated) keys; $R(\theta_k t_v^{(s)})$ is applied per slot at attention time via Def. 1, so positions recompute automatically as q advances (shared by both variants). Index $f \in \{0, \dots, F-1\}$ ranges over the intra-block frames of an AR chunk; prev denotes the previously-popped block.

Require: recent cache \mathcal{R} , summary cache \mathcal{S} (canonical), new chunk’s KVs, mode $\mu \in \{\text{WORLDTRACE-FIELD}, \text{WORLDTRACE-LANDMARK}\}$

Ensure: updated $(\mathcal{R}, \mathcal{S})$

```

1: Append new KVs to  $\mathcal{R}$  {recent window is verbatim}
2: if  $|\mathcal{R}| > W_r$  then
3:    $K_* \leftarrow \text{POPOLDEST}(\mathcal{R})$ 
4:   if  $\mu = \text{WORLDTRACE-FIELD}$  then
5:      $s^* \leftarrow$  summary slot whose source bucket now contains  $K_*$  {uniform temporal grouping}
6:      $\mathcal{S}[s^*] \leftarrow$  canonical mean of  $s^*$ ’s source frames {Def. 2}
7:   else if  $\mu = \text{WORLDTRACE-LANDMARK}$  then
8:      $\text{SE} \leftarrow \bigvee_{f=0}^{F-1} [\text{cosdist}(K_{\text{cx}, K_*, f}^{(k)}, K_{\text{cx}, \text{prev}, f}^{(k)}) > \tau]$  {per-frame check inside popped block}
9:     if SE then
10:       $\mathcal{S} \leftarrow \text{SHIFTLEFT}(\mathcal{S})$  {drop  $\mathcal{S}[0]$ }
11:       $\mathcal{S}[N_s-1] \leftarrow K_{\text{cx}, K_*}^{(k)}$  {canonical landmark; cf. Eq. (3)}
12:     end if
13:     (Init.) If fewer than  $N_s$  landmarks are stored, fill empty slots with the oldest.
14:   end if
15: end if
16: Attention time: for  $s=0, \dots, N_s-1$ , apply  $R(\theta_k t_v^{(s)})$  to  $\mathcal{S}[s]$  with  $t_v^{(s)}$  from Def. 1.

```

landmark slot) and $N_s=4$, $W_r=2$ for WORLDTRACE-FIELD. Compression uses uniform temporal grouping: the T_{old} oldest frames are split into N_s equal groups; each group’s keys are unrotated to canonical space (fp64 precision), averaged, and re-rotated at the group’s virtual position in the original dtype. **Position assignment:** Tab. 12 reports canonical key averaging with Block-Rel virtual positions (Yesiltepe et al., 2026), where each summary slot is assigned position $\max(0, q - (L_{\text{train}} - 1)F)$, for a fair comparison against the Naive+Block-Rel baseline. Tab. 6 reports the long-horizon conditions (Sliding Window, Block-Rel, centroid linear, fullcombo) at $N=16$. The recent-window slot j keeps its absolute position across schemes.

Evaluation protocol. Each method is evaluated on 100 videos generated from initial frames. TempSSIM is computed on decoded RGB frames using SSIM (Wang et al., 2004) between consecutive frames. LatentDiff is the mean squared difference between consecutive latent frames (pre-VAE-decode), used in long-horizon ablations where VAE decoding is expensive. Multi-seed experiments use seeds $\{0, 42, 123, 456, 789\}$. All experiments run on a single NVIDIA A100 80 GB GPU.

Camera trajectory for the coherence qualitative panel. The qualitative comparison in Fig. 5 (Sec. 3.2) is generated along a single fixed camera path on Matrix-Game 2, played out over $N=48$ AR chunks (~ 36 s of decoded video). Fig. 9 shows that path as a top-down plan: the camera leaves an initial scene (start, $t=1$), explores a roughly counter-clockwise loop through novel territory across $t \in \{8, 18, 24, 32, 40\}$, and arrives at a distinct end pose ($t=48$). Because the world model is run autoregressively on the same control inputs across baselines, every method in Fig. 5 is evaluated at exactly the same intended camera pose at each timestep; differences between columns therefore reflect the cache mechanism, not the trajectory.

Compute. Generating one 8-chunk video (24 latent frames) with WORLDTRACE-FIELD takes approximately 7.3 s at batch size 1; VAE decode adds ~ 2.5 s per chunk. The full ablation set required approximately 100 GPU-hours on single A100 80 GB GPUs. At ~ 250 W typical A100 draw and ~ 0.4 kg $\text{CO}_2\text{-eq/kWh}$, this is approximately 10 kg $\text{CO}_2\text{-eq}$.

Memory and compute. Cache storage is L_{train} blocks of KVs (~ 2.79 GB in fp16 for MG2-1.3B; constant in N , identical to the sliding-window baseline), while full-KV exhausts an 80 GB A100 at $N \approx 150$ chunks (~ 112 s ≈ 1.9 min of decoded video; cf. Tab. 14). The asymptotic per-token update cost is $O(T_{\text{old}} \cdot c_t \cdot n_\ell)$ ($\approx 1.7 \times 10^8$ FLOP at $N=100$); the FLOP cost itself is $O(1)$ in N . Wall-clock overhead per chunk grows modestly with horizon as bookmark accumulation and per-step kernel overhead take a larger share of each step: +3.8% ($N=8$) to +9.4% ($N=102$) for WORLDTRACE-FIELD, and +5.2% to +27.2% for WORLDTRACE-LANDMARK (Tab. 15). WORLDTRACE-LANDMARK replaces the averaging




Figure 9. **Camera trajectory used for Fig. 5.** Top-down plan of the camera (x, z) path executed by the AR rollout for the WORLDTRACE-FIELD qualitative panel at $N=48$ (~ 36 s). Green: start ($t=1$); dark red: end ($t=48$). Orange ticks mark the chunk indices shown as columns in Fig. 5. The path leaves the initial scene, traverses novel territory, and lands at a different pose, so every column $t>1$ is out-of-window for the sliding-window baseline.

loop with a single unrotate-then-store per scene-entry frame; its growing per-chunk wall-clock at long horizons reflects scene-entry-bookmark accumulation rather than added FLOP.

Per-chunk runtime overhead. Tab. 15 reports wall-clock time per AR chunk on one A100 80 GB at batch size 1 (3 distilled denoising steps, 3 latent frames per chunk, 352×640 resolution). At short horizons ($N=8$), all methods are within 6% of the sliding-window baseline because the forward pass dominates and the cache is small. At $N=102$ (~ 77 s of decoded video), WORLDTRACE-FIELD adds +9.4% per chunk as the canonical unrotate/rotate of the growing source set takes a larger share of each step, consistent with the $O(T_{\text{old}} \cdot c_t \cdot n_\ell)$ per-token form above. WORLDTRACE-LANDMARK’s +27.2% at $N=102$ reflects per-step kernel overhead from bookmark accumulation rather than additional FLOP. The long-horizon column uses the `skip_decode` setting (no VAE decode during generation), so absolute times are lower than the $N=8$ column with VAE decode.

Table 15. **Runtime per chunk.** Wall-clock time on one A100 80 GB (batch 1, 3 distilled denoising steps, 352×640). Short: $N=8$ with VAE decode. Long: $N=102$, skip-decode. Δ vs. sliding window. WORLDTRACE-LANDMARK stores per-layer canonical-K bookmarks at detected scene-entry events; the per-chunk overhead grows with horizon as more bookmarks accumulate.

Method	$N=8$ (short)		$N=102$ (long)	
	s/chunk	Δ	s/chunk	Δ
Sliding window	0.95	–	0.57	–
Naive	0.97	+2.3%	0.60	+4.7%
WORLDTRACE-FIELD	0.99	+3.8%	0.63	+9.4%
WORLDTRACE-LANDMARK	1.00	+5.2%	0.73	+27.2%

F.5. Evaluation Metrics

Open-ended video world models have no unique ground truth: a generated rollout legitimately diverges from any reference within a few chunks, so reference-based metrics (PSNR, SSIM vs. GT, LPIPS) and single-window benchmarks (VBench (Huang et al., 2024), PAI-Bench (Zhou et al., 2025a)) penalize valid creative divergence and do not test long-horizon recall. We therefore evaluate with a suite of reference-free metrics organized into two groups, *quality* (frame-to-frame coherence) and *consistency* (long-range episodic recall), with one diagnostic (LatentDiff) as a cross-check (Tab. 16).

Table 16. **Metric definitions.** \uparrow : higher is better; \downarrow : lower is better.

Metric	Role	Dir.	Definition
TempSSIM	Coherence	\uparrow	SSIM (Wang et al., 2004) between consecutive decoded frames, averaged over the rollout.
Local Scene Drift (SceneDrift)	Coherence	\downarrow	Mean per-chunk CLIP feature distance to the preceding chunk.
PAC	Recall	\uparrow	CLIP-ViT-H/14 cosine similarity between geometrically paired return- and forward-leg frames in ABA loops; PAC averages the final $N/8$ return chunks closest to scene A (Sec. 3.3).
Return CLIP	Recall	\uparrow	CLIP cosine similarity between the final return frame and the scene-A reference frame. Used for non-ABA topologies where position alignment is undefined.
LatentDiff	Diagnostic	\downarrow	MSE between consecutive latent frames (pre-decode). Confounded: favors slowly-varying output. Used alongside pixel-domain metrics as a fast sanity check only.

The suite is designed so that no single metric can be trivially gamed: TempSSIM alone rewards frozen output (a model that repeats the same frame scores perfectly), LatentDiff alone rewards degenerate slowly-varying output (sliding-window eviction achieves the lowest LatentDiff despite the worst TempSSIM), and SceneDrift alone could miss longer-range consistency failures. Together, the metrics triangulate actual quality: high TempSSIM (local coherence), low SceneDrift (the model generates dynamic content without scene wandering), and high PAC or Return CLIP (long-range episodic recall). All metrics are defined from first principles.

G. Discussion

G.1. Limitations

Ego-motion and screen-coordinate aliasing. WORLDTRACE-FIELD decouples temporal from spatial RoPE: the unrotate/rotate in Eq. (2) touches only the first $2c_t$ head dimensions and preserves the spatial-RoPE factors of Eq. (2) exactly (App. C). The canonical mean is exact *at fixed screen coordinates*, but under complex ego-motion (panning, strafing, rapid yaw), visually distinct objects can pass through the same (h, w) at different timestamps and be blended, losing object identity along motion paths. WORLDTRACE-LANDMARK sidesteps this by storing verbatim canonical keys, and pose-conditioned generators (e.g., LingBot-Fast’s Plücker, App. F) supply ego-motion as side information.

Architecture and orthogonal cache mechanisms. WORLDTRACE assumes temporal RoPE on keys, and a fixed AR KV cache; single-shot generators and combinations with LLM-side eviction methods such as SnapKV (Li et al., 2024) or H₂O (Zhang et al., 2023) are out of scope here. The null PAC result for WORLDTRACE-FIELD on LingBot-Fast is consistent with our recall/coherence split (Sec. 2.5): Plücker conditioning already supplies the recall signal, but a matched return-SSIM comparison under Plücker would test whether pose conditioning also saturates MG2’s coherence gains.

Metrics and headline numbers. Episodic recall is measured with paired CLIP cosine similarity on scripted return legs (Sec. 3); this rewards semantic alignment rather than pixel-faithful identity, and headline PAC sweeps inherit the metric’s geometric window scaling with horizon. Coherence is measured with frame-level TempSSIM, a local pixel-statistics measure that does not penalize a coherent rollout that has settled into the wrong scene; perceptually-aligned distances grow under such drift (Zhang et al., 2018). Human studies and pixel-level visitation metrics could further provide more insights.

Training paradigm and comparative scope. MG2 inherits a Self-Forcing-style autoregressive training distribution with bounded local attention at write time (App. F.1); inference-time KV stitching does not alter that mismatch if future models train with different cache semantics. “Training-free” here excludes fine-tuning the generator (Sec. 2.1); methods that distill long-rollout consistency, enlarge context by training, or change attention masks lie outside this protocol. Multi-student distillation (Song et al., 2025b) and other one-step generator regimes change the per-block denoising budget and therefore the effective length over which an AR rollout accumulates RoPE OOD; how cache-side interventions like WORLDTRACE compose with such distillation pipelines is an interesting open direction.

Deployment overhead. WORLDTRACE adds canonical-domain key transforms and bookkeeping relative to sliding-window eviction; although peak cache memory scales as $O(1)$ in horizon (Sec. 2), wall-clock latency and bandwidth to move updated keys through fused attention kernels are not modeled here (Tab. 15 reports per-configuration timings under our reference stack).

Domain generalization. Evaluations emphasize game-engine and navigation-conditioned rollouts with strong layout and lighting structure; how WORLDTRACE-FIELD and WORLDTRACE-LANDMARK behave on natural video with thin scene boundaries, film cuts, or rapid appearance changes is not established, and the scene-entry heuristics may need different thresholds.

G.2. Future Directions

The position/content factorization lets several research threads extend WORLDTRACE without retraining the underlying video model.

Geometry-aware canonical keys. The canonical mean of Eq. (2) aggregates by screen coordinate. Coupling the unrotate/rotate primitive with camera-pose warping (e.g., the Plücker-conditioned and Warped-RoPE writers of MosaicMem (Yu et al., 2026) and UCM (Xu et al., 2026b)) would let the same operator average over scene coordinates instead, removing the ego-motion aliasing above and unifying the strict zero-fine-tune regime with the trained camera-aware family.

Learned scene-entry policies. Replacing the canonical-key spike or gradient-onset rule with a small policy trained on action discontinuities, agent-pose deltas, or scene-segmentation logits would enable WORLDTRACE-LANDMARK to commit landmarks during continuous motion. Active recall, in which the policy decides *when* to commit and *which* of multiple stored landmarks to re-rotate at a given query, is a natural extension. The cache layout itself (N_s , W_r , scene-entry threshold) is fixed in our experiments; treating it as an autotunable schedule under a checkpoint-conditioned surrogate (Mehta et al., 2024) could let downstream practitioners adapt it per backbone or per horizon without re-running the full slot-sensitivity sweep of App. D.4.

Composition with content-side eviction. Because slot-rank positions are independent of which canonical content fills the slot, eviction heuristics from the LLM literature (H_2O (Zhang et al., 2023), SnapKV (Li et al., 2024), KnormPress (NVIDIA, 2024)) can be layered on top of WORLDTRACE-FIELD’s canonical averages and WORLDTRACE-LANDMARK’s frozen keys without re-deriving the position scheme. The MemRoPE (Kim et al., 2026) comparison in App. D.2 suggests the two interact, so principled co-design is open. A complementary direction for WORLDTRACE-LANDMARK is to spend an extra summary slot on a residual WORLDTRACE-FIELD-style canonical mean of evicted landmarks, retaining a coarse coherence trace of the discarded scene-entry frames at no additional position-side cost.

Downstream consumers of pretrained generative teachers. WORLDTRACE targets the inference-time cache of an autoregressive video generator, but pretrained diffusion and AR generators feed a broader pipeline ecosystem whose budgets and biases interact with the cache design. Score-distillation pipelines for amortized 3D synthesis (ATT3D (Lorraine et al., 2023), LATTE3D (Xie et al., 2024)) and LLM-conditioned mesh generation (LLaMA-Mesh (Wang et al., 2024)) consume teacher gradients whose Monte Carlo variance, rather than long-horizon recall, dominates compute; compute-aware estimators for those gradients (CARV (Bettencourt et al., 2026)) and the corresponding analyses for non-vision teachers (Richter-Powell et al., 2025) are orthogonal axes to the position–content factorization studied here. On the data side, motion-attribution methods (Wu et al., 2026c) ask which training clips improved the temporal dynamics of a generator, complementing the cache-side question of which past slots a generator can still address at inference. Establishing whether the slot-rank virtual-position primitive transfers to these adjacent regimes is an open direction.

Fine-tuning extensions. WORLDTRACE stays within the training context length L_{train} (Sec. 2.1) precisely because it is training-free: the constraint $N_s + W_r = L_{\text{train}}$ keeps every summary slot at an in-distribution offset the generator was trained on. Two light fine-tuning paths would relax this budget without retraining the generator from scratch. (i) *Context-extension fine-tuning* on synthetic long rollouts would let WORLDTRACE allocate either more recent slots for coherence or more summary slots for longer recall horizons at the same in-distribution attention cost. (ii) *Position-aware fine-tuning* that exposes the model to the slot-rank offsets of Eq. (1) would tighten the canonical-mean approximation underlying WORLDTRACE-FIELD (Rem. 2), narrowing the residual long-horizon PAC gap. Both paths fit the nested-optimization template (Lorraine, 2024) of a frozen large inner model paired with a light outer adapter for the cache schedule, so the position/content factorization is preserved through the outer loop. Both are compatible with the WORLDTRACE cache layout and leave the position/content factorization intact.

Multi-tier and cross-architecture transfer. WORLDTRACE uses a two-tier split (recent verbatim, summary canonical). Adding intermediate tiers, with progressively larger source buckets and progressively deeper slot-rank offsets, could trade sharper recall for longer effective horizons. Porting the same factorization to KV-cache-bearing variants of MG3 (Wang et al., 2026b), Genie3 (Google DeepMind, 2025), or LingBot-Fast (Robbyant Team et al., 2026) would test whether the position/content split is architecture-specific or generic.

Scaling LoopMem. LoopMem (App. E) exercises all four difficulty tiers on MG2-1.3B; camera-orientation Tier 3 uses MG2’s mouse-yaw control, where chunk counts are nominal magnitudes rather than calibrated angles. Pose-conditioned generators with explicit pitch/roll would extend Tier 3, and broader cross-architecture leaderboards on Tiers 1, 2, and 4 would let the community compare position/content trade-offs at controlled compression ratios, complementing the broad video-generation benchmarks of WorldScore (Duan et al., 2025), MIND (Ye et al., 2026), and VBench-2.0 (Zheng et al., 2025). The Pan 360° failure of WORLDTRACE-LANDMARK on visually continuous trajectories (Tab. 2) suggests learned or rule-based scene-entry policies as a concrete next step.

G.3. Broader Impact

This is an inference-time KV-cache modification: it does not alter training data, modify weights, or expand the base model’s capabilities, so it introduces no new dual-use risks beyond those inherent to the underlying video model. The capability shift is minute-scale long-horizon generation at $O(1)$ peak cache memory; longer coherent clips raise the importance of provenance metadata and watermarking, and the lower memory footprint broadens both research access and the surface for misuse, so releases should follow the base model’s content-policy guidance.