ANY TIME, ANY SPACE: A WORLD-STATE-CENTRIC FRAMEWORK FOR AGENTIC VIDEO GENERATION

Anonymous authors
Paper under double-blind review



Fig. 1: **World-state-centric agentic video generation.** ATAS treats a multi-shot film as a sequence whose spatial and temporal information can be reused. Shared scene geometry anchors consistent return views, while time-indexed event evidence supports replay from a new camera. The generated shots form the visible film, and the resulting scene and event information guides subsequent requests.

ABSTRACT

Long-form agentic video generation must turn a screenplay into shots that remain spatially and temporally consistent across nonconsecutive requests. Passing prompts or recent frames between shots does not provide a shared spatial frame for revisiting a scene or a time-indexed record for replaying an earlier event. We present **Any Time, Any Space (ATAS)**, which compiles a screenplay into shot requests with explicit location, camera, and event-time references. Before generation, it retrieves relevant evidence from a scene memory of registered static geometry and an event memory of dynamic observations indexed by scene, event, and time. After generation, ATAS reconstructs and registers the clip. It updates memory only when review and state checks accept the observation, while film delivery remains independent of memory admission. Across four video generators, ATAS improves all reported cross-shot consistency measures and achieves the highest aggregate benchmark score for each backbone. It also outperforms the compared agents and 3D-guided methods on target-image spatial revisit, leads three event-recapture criteria, and receives the highest human ratings on all six criteria. Video results are available at `https://anonymous-submission-1186.github.io/`.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

1 INTRODUCTION

Agentic video generation translates a screenplay into a sequence of video shots by coordinating shot planning, reference selection, and video synthesis. As a story unfolds across locations, times, and viewpoints, its coherence depends on preserving the spatial and temporal relationships established in earlier shots. For example, a story may depict an event in a room, shift to another location, and later return to the room or replay the earlier event from a different viewpoint. The return shot should preserve the established scene layout, while the replay should reproduce the same event, including its actions and temporal progression, under the new camera view. These requirements extend beyond the visual quality of individual clips: they require scenes and events to remain consistently accessible across intervening shots, so that later generations remain grounded in a persistent underlying world, as illustrated in Fig. 1.

Existing agentic video systems typically preserve continuity through recent frames, text summaries, or retrieved reference images. While effective for locally consecutive shots, these cues neither organize past observations in a shared spatial and temporal frame nor support non-local requests to revisit an earlier scene or event. This motivates a persistent world state, but maintaining such a state introduces two challenges. Past observations must remain addressable in both space and time, and the state must evolve without blindly absorbing generated content, whose visual, geometric, or temporal errors could otherwise propagate to later shots. A practical system must therefore resolve each request into a concrete shot specification, retrieve the relevant spatial and temporal evidence, and incorporate the resulting observation into the persistent world state only when it is sufficiently reliable.

To address these challenges, we introduce **Any Time, Any Space (ATAS)**, a world-state-centric framework for agentic video generation. ATAS maintains two complementary forms of persistent evidence: a scene state that preserves spatial structure in a shared coordinate frame, and an event state that retains dynamic observations with their temporal context. For each screenplay request, a world-state compiler grounds the requested shot in this world state, retrieves the relevant spatial and temporal evidence, and conditions a short-video generator on the resulting context.

Each generated clip is then treated as a candidate observation rather than trusted state. ATAS validates the reconstructed observation before incorporating it into the persistent world state, preventing unreliable generated evidence from propagating to subsequent shots. Importantly, whether a clip is suitable for the final video is separated from whether its observation is reliable enough to update the world state. This closed loop allows the world state to evolve across generation while remaining addressable for later spatial revisits and historical event recapture.

Our contributions are threefold:

- We introduce ATAS, a world-state-centric framework that organizes screenplay-driven video generation around persistent scene and event memories, enabling later shots to query evidence from earlier observations in any space and any time.
- We develop complementary spatially registered scene and time-indexed event memories, together with a validation-controlled update mechanism that separates film acceptance from the admission of new evidence into persistent state.
- Experiments against five agent baselines across four generators show that ATAS improves every reported cross-shot consistency measure and attains the highest aggregate score in each backbone block. It also outperforms the evaluated agent and 3D-guided methods on all three spatial-revisit metrics, leads three of five event-recapture criteria, and receives the highest mean human rating on all six criteria.

2 RELATED WORK

Video Generation. Text-to-video and image-to-video models provide the renderer for long-video systems, improving appearance, motion, identity, and temporal coverage (Ho et al., 2022b; Singer et al., 2022; Ho et al., 2022a; Hong et al., 2022; Bar-Tal et al., 2024; Kondratyuk et al., 2023; Yang et al., 2024; Polyak et al., 2024; Chen et al., 2024; Blattmann et al., 2023). Long-video methods extend temporal context through co-denoising, distributed generation, or frame-level direction (Wang et al., 2023; Tan et al., 2024; Hong et al., 2023), while camera-, sparse-control-, and geometry-guided

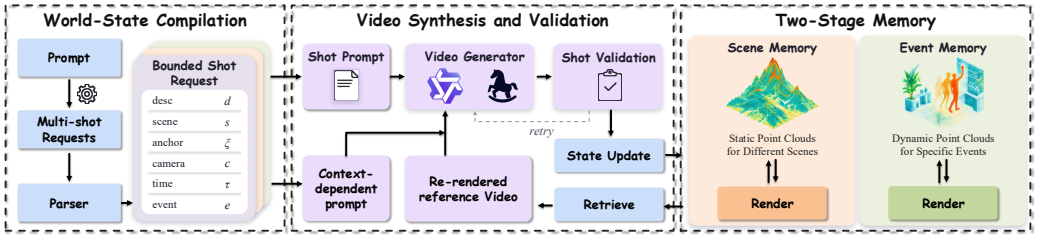

Fig. 2: **Workflow of ATAS.** A high-level prompt is parsed into multi-shot requests and bounded shot requests. Each request contains a description, scene, spatial anchor, camera, time interval, and optional event. The shot prompt and the context-dependent prompt condition the video generator. Scene memory supplies static point-cloud references for the requested scene, while event memory supplies dynamic references for the requested event. Shot validation either requests another candidate or sends the accepted observation to the state update. The generated shot remains available as film independently of whether the observation is added to persistent state.

methods constrain view or motion (Zhou et al., 2024; Zhao et al., 2025a; Khachatryan et al., 2023; Xing et al., 2023; Guo et al., 2023; Wu et al., 2025a). These models improve local rendering and short-term control, but do not specify how observations from multiple clips are registered, retained, or queried in a later shot.

Agentic Video Generation. Agentic video systems add script parsing, shot planning, reference selection, and review to the generation loop (Long et al., 2024; Li et al., 2024; Zheng et al., 2025; Wu et al., 2025c; Lin et al., 2023; Xie et al., 2024; Gong et al., 2023; An et al., 2025; Zhou et al., 2026; Han et al., 2025). Their inter-shot state is usually an editorial context of prompts, reference images, or generated histories. Such context can preserve local appearance without defining a shared coordinate frame, an event-time index, or a rule for admitting a generated observation into future state. ATAS separates request compilation, memory retrieval, shot generation, and state update within the plan-act-observe loop of language agents (Yao et al., 2022).

Explicit Memory for Video Generation. Point-cloud and explicit-memory methods condition future generation on spatial evidence. Neural rendering establishes reconstruction from sparse or multi-view observations (Yu et al., 2021; Chen et al., 2021), and 3D Gaussian Splatting provides an efficient explicit scene representation (Kerbl et al., 2023). Spatia, SPMem, LSM-World, and EvoWorld accumulate or query scene geometry (Zhao et al., 2025b; Wu et al., 2025b; Wang et al., 2026; 2025a). ViewCrafter, GEN3C, and Vista4D study dynamic camera-controlled reshooting (Yu et al., 2024; Ren et al., 2025; Lin et al., 2026). FantasyWorld, WorldSplat, WorldForge, WonderWorld, and HY-World 2.0 explore geometry-consistent or navigable world construction (Dai et al., 2025; Zhu et al., 2025; Song et al., 2025; Yu et al., 2025; Team HY-World et al., 2026). These methods provide geometric memory, dynamic view synthesis, or world assets, but do not expose the same shot-local observation layer, event-time addressing, and state-update decision interface. ATAS separates static scene memory from event-local dynamic evidence and evaluates an observation before incorporating it into persistent state.

3 METHOD

ATAS turns a high-level screenplay into short video shots while keeping the spatial and temporal evidence needed by later shots. The central object is a persistent world state, queried before each shot and updated by observations that pass validation.

3.1 TERMS AND WORKFLOW

The terminology is as follows. A *source observation* o^{src} is the input frame or clip that starts a scene. Registering this observation establishes the scene coordinate frame. The observation and the coordinate frame are different objects. The *scene memory* $\mathcal{S}_t$ stores the accepted static geometry in that frame, while the *event memory* $\mathcal{E}_t$ stores accepted dynamic observations together with their scene,

event, and time information. The *world state* is simply the pair of memories at a particular step,

$$\mathcal{W}_t = (\mathcal{S}_t, \mathcal{E}_t). \tag{1}$$

Equation 1 defines the persistent state used in the method. Scene and event refer to the contents of the two memories, while world state denotes their current pair after the first t requests.

The input is a high-level screenplay u, rather than a list of fully specified video prompts. At step t, the compiler produces a bounded shot request q_t, the memory query produces a conditioning context, and the generator produces a clip V_t. The clip is then converted into a candidate observation and either admitted to the next state or retained as film without a state update:

$$(u, \mathcal{W}_t) \xrightarrow{\text{Compile}} q_t \xrightarrow{\text{Query}} C_t \xrightarrow{G} V_t \xrightarrow{\text{Reconstruct/Register/Validate}} (V_t, \mathcal{W}_{t+1}). \tag{2}$$

The next request uses the resulting state. Fig. 2 summarizes the workflow.

3.2 World-state Compilation

Before generation, the compiler recursively splits the screenplay by scene, spatial and temporal relations, event references, and camera requirements until each request fits the short-video budget. A compiler trace is given in Appendix B.

We write a renderable request as

$$q_{t,k} = (d_{t,k}, s_{t,k}, \xi_{t,k}, c_{t,k}, \tau_{t,k}, e_{t,k}), \tag{3}$$

where $d_{t,k}$ is the local shot description, $s_{t,k}$ identifies the scene, $\xi_{t,k}$ is the spatial anchor of the request in the scene coordinate frame, $c_{t,k}$ is the target camera, $\tau_{t,k}$ is the requested time interval, and $e_{t,k}$ is an optional event identifier.

For a screenplay segment, the compiler produces a request schedule

$$\mathcal{Q}_t = \text{Compile}(u, \mathcal{W}_t) = \{q_{t,1}, \ldots, q_{t,K_t}\}, \tag{4}$$

where K_t is the number of requests currently emitted. Requests are processed in order, and a validated $\mathcal{W}_{t+1}$ conditions the next request. A rejected observation remains in the film but is not state evidence.

3.3 Two-Stage Memory: Scene and Event Memory

Scene memory. The registered source observation o^{src} establishes the scene coordinate frame and initializes $\mathcal{S}_0$. Later accepted observations are registered to that frame and merged as static point-cloud evidence with dynamic regions excluded. A request recovers a previously observed space by rendering these coordinates at its spatial anchor ξ_t and requested camera:

$$C_t^{\text{scene}} = \text{Render}(\mathcal{S}_t, c_t), \tag{5}$$

where C_t^{scene} is the re-rendered scene reference supplied to the generator. Reconstruction happens only after generation, when the clip is considered for a future state update (Zhao et al., 2025b; Wu et al., 2025b; Wang et al., 2025a; Wu et al., 2025a).

Event memory. The event memory keeps dynamic evidence separate from static geometry. It is a set of records,

$$\mathcal{E}_t = \{\epsilon_i\}_{i=1}^{N_t}, \qquad \epsilon_i = (s_i, e_i, b_i, \tau_i, \mathcal{D}_i), \tag{6}$$

where s_i is the scene identifier, e_i is the event identifier, b_i identifies the source observation, τ_i is the source time interval, and $\mathcal{D}_i$ contains the registered frame-level RGB-D, mask, and dynamic point-cloud evidence. If a request refers to event e_t during interval τ_t, the event context is retrieved directly from $\mathcal{E}_t$:

$$C_t^{\text{event}} = \begin{cases} \text{Retrieve}(\mathcal{E}_t, s_t, e_t, \tau_t), & e_t \neq \varnothing, \\ \varnothing, & \text{otherwise.} \end{cases} \tag{7}$$

The result is empty when no compatible record exists. Equation 7 is the event branch of the query that produces C_t^{event}. The memory retains frame-local dynamic evidence rather than an identity-resolved motion model (Ren et al., 2025; Lin et al., 2026). Scene and event references are combined into the context-dependent prompt, while reconstruction remains a post-generation operation.

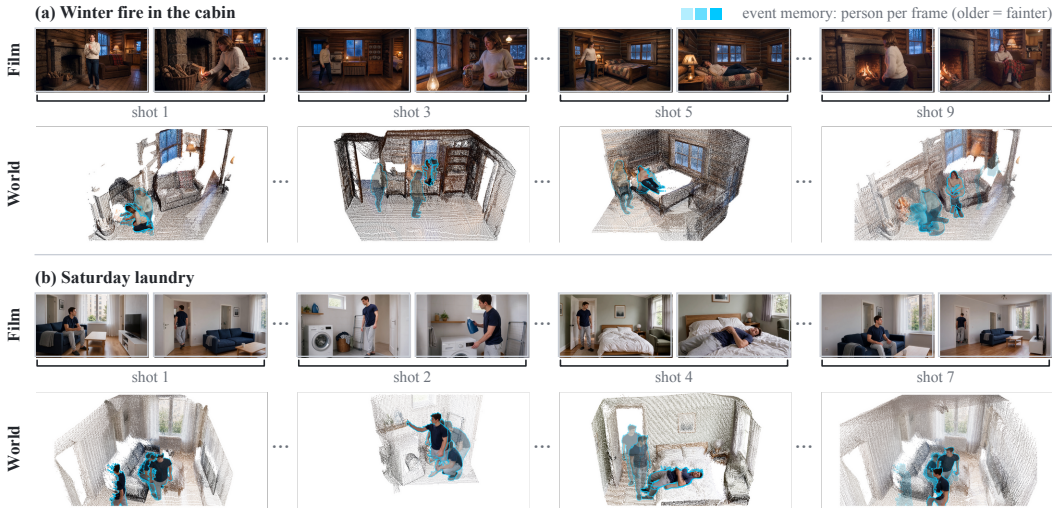

Fig. 3: **Qualitative world-state rollout.** Film frames and registered world views are shown for two story sequences. The examples illustrate how accepted observations extend the scene and event state across intervening shots.

3.4 Shot Generation from Queried State

The local description in q_t forms the shot prompt, while the scene and event references form the context-dependent prompt shown in Fig. 2:

$$C_t = (C_t^{\text{scene}}, C_t^{\text{event}}), \qquad V_t = G(q_t, C_t), \tag{8}$$

where G is the short-video generator. The shot prompt supplies local intent and its place and time, while C_t supplies scene and, when needed, event references. This interface is kept fixed across comparisons (Zhou et al., 2024; Zhao et al., 2025a; Henschel et al., 2024).

3.5 Observation Registration and State Update

The generated clip enters shot validation, which produces two separate decisions. Film review returns $f_t = 1$ when the clip is accepted for delivery. If it is accepted, reconstruction and registration produce a candidate observation:

$$\widetilde{\mathcal{O}}_t = \text{Reconstruct}(V_t), \qquad \mathcal{O}_t = \text{Register}(\widetilde{\mathcal{O}}_t, \mathcal{S}_t). \tag{9}$$

The registration stage returns $r_t = 1$ when the reconstruction aligns to the scene coordinate frame. For the source observation, a valid registration establishes that frame. If film review fails, the system retries or rejects the candidate. If registration fails, the clip remains usable as film but cannot update state. The decisions come from the respective film-review and registration stages (Umeyama, 1991; Wang et al., 2025b;c; Lin et al., 2026).

The state update check returns $a_t = 1$ after the film and registration checks pass and the candidate has valid provenance, scene scope, and parent-state version. For a candidate that passes this check, the proposed next state is

$$\widehat{\mathcal{W}}_{t+1} = \text{Update}(\mathcal{W}_t, \mathcal{O}_t). \tag{10}$$

The actual transition is

$$\mathcal{W}_{t+1} = \begin{cases} \widehat{\mathcal{W}}_{t+1}, & f_t = r_t = a_t = 1, \\ \mathcal{W}_t, & \text{otherwise.} \end{cases} \tag{11}$$

Equation 11 makes the decision rule explicit: f_t controls film delivery, r_t registration, and a_t admission to persistent state. Film and state update are independent outputs, so V_t remains in the film when $f_t = 1$ even if $r_t = 0$ or $a_t = 0$. In that case, the clip is unavailable as evidence for later requests. Its sensitivity is analyzed in Appendix D.

Tab. 1: **Agent-level quantitative evaluation on standard benchmarks.** Models are grouped by source type and rendering backbone. Methods within a block share the generator, prompt package, duration budget, and evaluation protocol. Bold and underlined entries mark the best and second-best values within each backbone block.

Model	Method	VBench	ST-Bench	ViStory-Self	ViStory-Cross	Avg.
Open-Source Models						
MiniMax-H3	I2V Extension	0.828 ± 0.013	0.572 ± 0.020	0.601 ± 0.028	0.418 ± 0.027	0.605
	VGoT	0.834 ± 0.014	0.612 ± 0.017	0.638 ± 0.023	0.468 ± 0.030	0.638
	Mora	0.836 ± 0.012	0.598 ± 0.019	0.628 ± 0.026	0.452 ± 0.024	0.628
	MovieAgent	0.852 ± 0.012	0.608 ± 0.016	0.642 ± 0.025	0.476 ± 0.023	0.644
	ReCA	0.844 ± 0.011	0.628 ± 0.018	0.658 ± 0.021	0.508 ± 0.022	0.659
	ATAS	**0.860** ± 0.010	**0.691** ± 0.014	**0.710** ± 0.020	**0.547** ± 0.021	**0.702**
Proprietary Models						
Wan 2.7	I2V Extension	0.778 ± 0.014	0.486 ± 0.022	0.512 ± 0.029	0.348 ± 0.031	0.531
	VGoT	0.786 ± 0.011	0.536 ± 0.019	0.564 ± 0.026	0.412 ± 0.028	0.575
	Mora	0.791 ± 0.015	0.528 ± 0.017	0.558 ± 0.031	0.388 ± 0.033	0.566
	MovieAgent	0.808 ± 0.013	0.542 ± 0.021	0.572 ± 0.024	0.402 ± 0.026	0.581
	ReCA	**0.854** ± 0.009	0.568 ± 0.018	0.598 ± 0.022	0.428 ± 0.029	0.612
	ATAS	0.846 ± 0.012	**0.625** ± 0.016	**0.666** ± 0.021	**0.486** ± 0.032	**0.656**
HappyHorse 1.0	I2V Extension	0.836 ± 0.012	0.558 ± 0.023	0.592 ± 0.029	0.408 ± 0.031	0.599
	VGoT	0.822 ± 0.014	0.588 ± 0.018	0.628 ± 0.025	0.458 ± 0.027	0.624
	Mora	0.824 ± 0.011	0.574 ± 0.020	0.618 ± 0.027	0.442 ± 0.026	0.614
	MovieAgent	0.856 ± 0.013	0.598 ± 0.016	0.648 ± 0.021	0.472 ± 0.024	0.643
	ReCA	0.841 ± 0.010	0.621 ± 0.019	0.668 ± 0.024	0.496 ± 0.025	0.656
	ATAS	**0.862** ± 0.009	**0.668** ± 0.015	**0.706** ± 0.019	**0.540** ± 0.030	**0.694**
Wan 3.0	I2V Extension	0.836 ± 0.013	0.548 ± 0.021	0.582 ± 0.027	0.401 ± 0.030	0.592
	VGoT	0.834 ± 0.010	0.571 ± 0.019	0.622 ± 0.024	0.461 ± 0.026	0.622
	Mora	0.838 ± 0.012	0.562 ± 0.023	0.614 ± 0.022	0.438 ± 0.031	0.613
	MovieAgent	0.858 ± 0.011	0.578 ± 0.018	0.628 ± 0.025	0.452 ± 0.024	0.629
	ReCA	0.848 ± 0.009	0.602 ± 0.017	0.661 ± 0.020	0.481 ± 0.027	0.648
	ATAS	**0.877** ± 0.011	**0.719** ± 0.015	**0.756** ± 0.019	**0.584** ± 0.028	**0.734**

3.6 Two Persistence Queries

The same workflow supports two queries. For spatial revisit, the compiler emits the earlier scene identifier and the requested return anchor and camera. The system renders $\mathcal{S}_t$ and compares the generated return image with the target image. For historical event recapture, the request includes an earlier event identifier and time interval. The system retrieves the matching records from $\mathcal{E}_t$, combines them with the scene context, and generates the event from the requested camera. The first query tests recovery of static spatial evidence, and the second tests whether a time-local event remains addressable after intervening shots. The event protocol is source-grounded and does not assume a synchronized target video.

4 Experiments

4.1 Experimental Setup

Standard benchmarks assess complementary aspects of long-form video generation. We use the official inputs and evaluators for VBench (Huang et al., 2023), ST-Bench (Zhang et al., 2025), and ViStory-Self and ViStory-Cross (Zhuang et al., 2025). VBench measures aesthetics, subject and background consistency, temporal flicker, motion smoothness, and dynamic degree. ST-Bench evaluates aesthetics, global and single-shot prompt following, and cross-shot consistency with C-All and C-Top10. ViStory-Self measures identity, appearance, and temporal consistency within a story, while ViStory-Cross evaluates identity, appearance, and scene consistency across shots. Methods

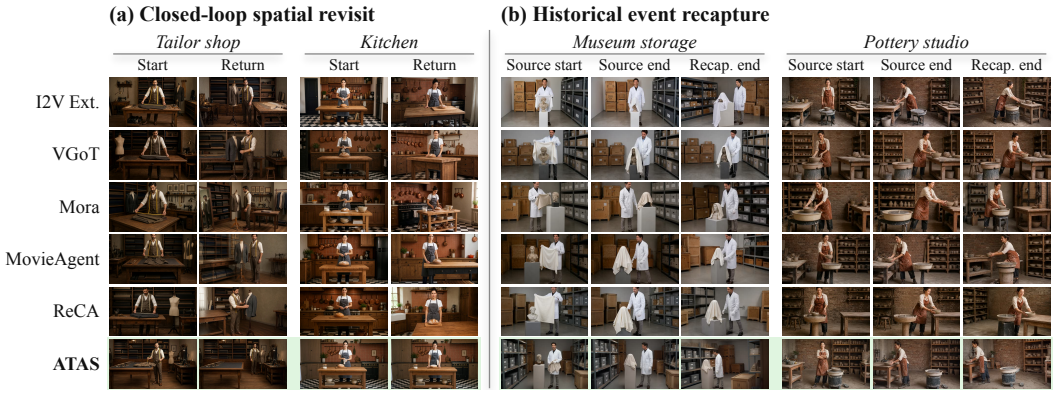

Fig. 4: **Spatial revisit and historical event recapture.** The top row uses a target image for spatial comparison. The bottom row uses source-grounded event evidence under a requested camera.

Tab. 2: **Spatial revisit.** The upper block lists multi-shot agent-based methods and the lower block lists 3D-guided view-synthesis methods. ATAS is shown separately. All methods are evaluated against target images under a common-visibility mask.

Method	$PSNR_C$	$SSIM_C$	$LPIPS_C$
I2V Extension	13.601	0.293	0.339
VGoT	14.444	0.313	0.300
Mora	13.796	0.318	0.288
MovieAgent	14.543	0.352	0.296
ReCA	13.870	0.268	0.306
FlexWorld	14.740	0.343	0.315
ViewCrafter	15.255	0.375	0.293
ATAS	**15.968**	**0.385**	**0.284**

Tab. 3: **Source-grounded event recapture.** Five criteria compare generated clips with source-event evidence and the requested camera. The rubric does not use synchronized target videos. Its axes assess event identity, action and state, temporal order, character and object consistency, and viewpoint realization. Scores are averaged across axes and then across event requests.

Method	Event id.	Action state	Temporal order	Char. object	Viewpoint realization
I2V Extension	0.450	0.433	0.433	0.817	0.200
VGoT	0.517	0.533	0.400	0.850	0.400
Mora	0.286	0.339	0.357	0.768	0.250
MovieAgent	0.500	0.450	0.500	0.783	0.350
ReCA	0.733	**0.700**	0.617	**0.917**	0.500
ATAS	**0.750**	0.688	**0.688**	**0.917**	**0.650**

within each backbone block share the generator, prompt package, duration budget, and aggregation protocol.

Spatial revisit and event recapture use task-specific evidence. Spatial revisit compares a generated return view with a target image under a common-visibility mask using $PSNR_C$, $SSIM_C$, and $LPIPS_C$. The comparison includes multi-shot agent-based and 3D-guided view-synthesis methods. Event recapture uses a source-grounded VLM judge given sampled source-event frames, an event description and interval, the requested camera, and the generated clip. Its rubric scores event identity, action and state, temporal order, character and object consistency, and viewpoint realization. Binary decisions map to zero or one, ordinal ratings from one to five are normalized to zero through one, and scores are averaged across axes and then across event requests. No synchronized target video is used.

Controlled analyses and human ratings measure system behavior beyond benchmark scores. Memory ablation fixes the ATAS agent and Wan 3.0 while varying stored evidence. Component ablation adds scene and event resolution, recursive compilation, observation validation, and state-update checks under matched requests. Efficiency measurements report per-story runtime, real-time factor, candidate-review outcomes, retries per shot, retry-adjusted delivery time, throughput, and memory overhead across six memory settings. Human raters compare anonymized whole-video outputs from six methods on six criteria using a five-point scale.

4.2 Main Results

Qualitative rollouts illustrate continuity across nonconsecutive shots. Fig. 3 pairs film frames with reconstructed world views for two stories. In both examples, the story moves through other

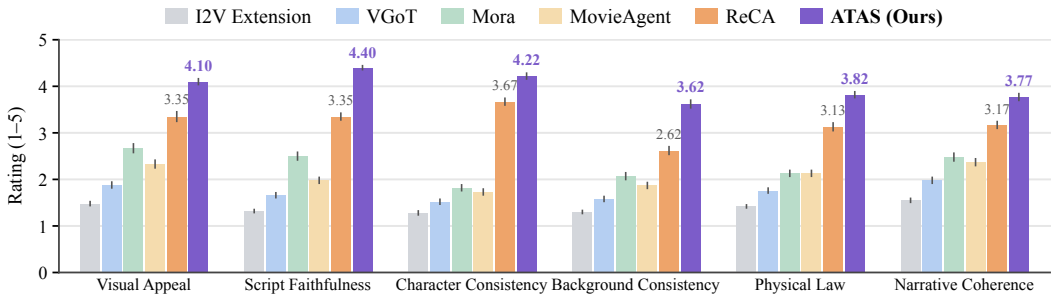

Fig. 5: **Human evaluation of agent methods.** Grouped bars show mean ratings and standard-error whiskers for six methods across six criteria.

Tab. 4: **Agent-component ablation.** All rows use Wan 3.0 and the same requests, cameras, candidate budget, and retry policy. Components are added cumulatively from shot rendering through memory resolution, recursive compilation, observation validation, and state-update checking.

Configuration	ViStory average ↑	C-Top10 ↑	LPIPS_C ↓	Event Recapture Score ↑
ATAS w/ shot rendering	0.478	0.528	0.341	0.441
ATAS w/ scene and event resolution	0.612	0.568	0.334	0.508
ATAS w/ recursive compilation	0.638	0.684	0.288	0.711
ATAS w/ observation validation	0.651	0.702	0.286	0.726
ATAS w/ state update check	**0.670**	**0.719**	**0.284**	**0.739**

shots before returning to a previously visited place. The paired views make visible whether the scene layout and object state persist across that interval. Appendix G presents complete results.

ATAS improves cross-shot consistency across generators while retaining competitive local quality. In Tab. 1, it leads every reported aggregate on MiniMax-H3, Wan 3.0, and HappyHorse 1.0. With Wan 2.7 it remains ahead on ST-Bench and both ViStory consistency measures despite trailing on VBench. This pattern across open-source and proprietary backbones suggests that the gain comes from persistent multi-shot state rather than a single renderer. We compare with I2V Extension, VGoT (Zheng et al., 2025), Mora (Yuan et al., 2024), MovieAgent (Wu et al., 2025c), and ReCA (Liu et al., 2026) under matched backbone blocks.

ATAS improves both target-view recovery and source-grounded event recapture. The examples in Fig. 4 show return views and event replays from different cameras. ATAS reaches 15.968 PSNR_C, 0.385 SSIM_C, and 0.284 LPIPS_C, outperforming the compared multi-shot and 3D-guided methods (Tab. 2). In Tab. 3, it leads event identity (0.750), temporal order (0.688), and viewpoint realization (0.650), ties ReCA on character and object consistency (0.917), and remains slightly lower on action-state preservation (0.688 versus 0.700). The remaining difficulty is exact action-state reproduction rather than retrieval of the requested occurrence or view.

Human raters give ATAS the highest mean score on all six criteria. Fig. 5 reports ratings for visual appeal, script faithfulness, character and background consistency, physical law, and narrative coherence under the anonymized protocol and criteria of Liu et al. (2026). Ratings are highest for script faithfulness (4.40), character consistency (4.22), and visual appeal (4.10), while background consistency (3.62) and narrative coherence (3.77) remain lower. Participants therefore most clearly perceive faithful shot content and character carryover. Environment-wide consistency and whole-story coherence remain less resolved.

4.3 Ablation Study and Experimental Analysis

Scene and event memory improve different forms of continuity. Frame history improves all three measures over no memory. Adding static scene evidence lowers LPIPS_C by 0.030, and updating that scene raises C-Top10 by 0.041 with little change in event score. Event memory gives the largest recapture gain with a smaller spatial change. Across the full ablation, C-Top10 rises from 0.536 to 0.719, event score from 0.452 to 0.739, and LPIPS_C falls from 0.341 to 0.284.

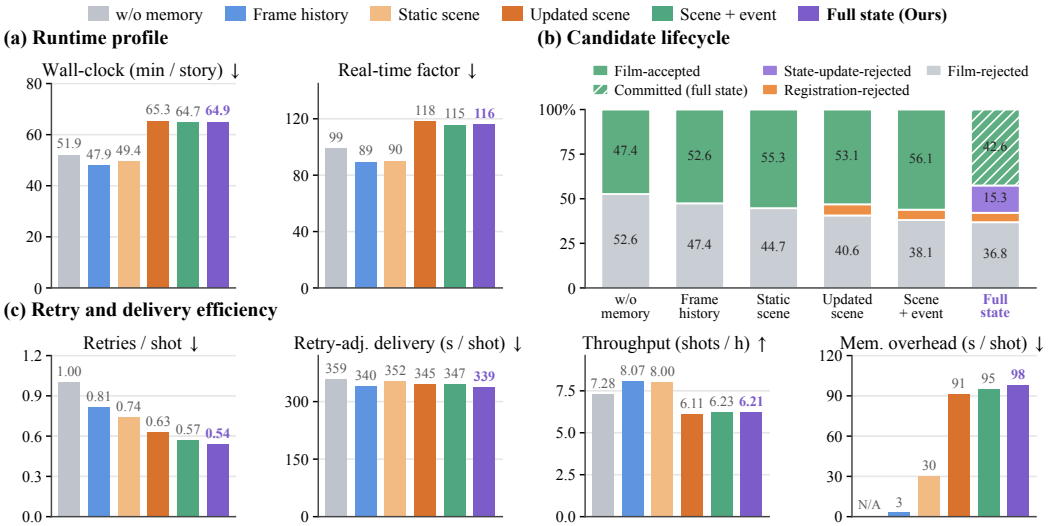

Fig. 6: **Efficiency comparison across memory settings.** Panels (a), (b), and (c) show runtime, candidate lifecycle, and retry and delivery measures across six memory settings.

The modest final increment from scene-plus-event to full state suggests that validation preserves the main capabilities rather than adding another retrieval effect (Tab. 5).

Tab. 5: **Memory analysis.** ATAS with Wan 3.0 is fixed. Rows vary stored evidence and state checks.

Config.	C-Top10↑	LPIPS$_C$ ↓	Recap.↑
ATAS w/o memory	0.536	0.341	0.452
ATAS w/ frame history	0.598	0.319	0.548
ATAS w/ static scene	0.617	0.289	0.583
ATAS w/ updated scene	0.658	0.286	0.585
ATAS w/ scene + event	0.703	0.285	0.728
ATAS w/ full state	**0.719**	**0.284**	**0.739**

Scene and event memory establish gains that recursive compilation amplifies. Adding memory raises ViStory average from 0.478 to 0.612, C-Top10 from 0.528 to 0.568, and event score from 0.441 to 0.508. Recursive compilation then raises the event score to 0.711 and C-Top10 to 0.684 (Tab. 4).

Persistent memory reduces retry demand at a measurable compute cost. In Fig. 6, retries fall from 1.00 to 0.54 per shot and retry-adjusted delivery time declines from 359 to 339 seconds per shot. Per-story runtime rises from 51.9 to 64.9 minutes, memory overhead reaches 98 seconds per shot, and throughput falls from 7.28 to 6.21 shots per hour. The benefit is fewer retries and faster retry-adjusted delivery, rather than lower total compute.

Film-review rejection decreases as memory evidence accumulates. The candidate-lifecycle panel shows a monotonic decline from 52.6% with no memory to 36.8% with full state, with each intermediate setting following the same direction. This trend indicates that the improvement is not confined to the final configuration. It also measures whether candidates can be delivered, complementing scores that evaluate the selected outputs.

Film delivery and state admission. Some clips are retained in the film even when registration or the state update is rejected. This lets the sequence continue without using an unaligned observation to condition later requests. The status breakdown separates these delivered clips from observations admitted to memory. Read with the runtime profile, it shows that persistent evidence reduces repeated review cycles while adding processing overhead, rather than making total generation faster.

5 Conclusion

We presented ATAS, which compiles screenplays into shot requests grounded in explicit scene, camera, and event-time references, then updates persistent state only from validated observations. Across four video generators, ATAS improves cross-shot consistency, spatial revisit, event recapture, and human ratings. The ablations show complementary contributions from scene memory, event memory, and recursive compilation, while the efficiency analysis shows fewer retries at added compute cost. By keeping locations and prior events addressable throughout a story, ATAS provides a foundation for agentic video tools in interactive filmmaking, episodic content creation, and previsualization.

AI USE STATEMENT

Generative AI tools were used to assist with manuscript preparation, including format migration, organization of draft materials and notes, and language editing. The authors reviewed all AI-assisted work and remain responsible for the final text, claims, citations, experimental results, and artifacts in this submission.

REPRODUCIBILITY STATEMENT

Appendix B details the request schema and recursive compilation procedure, including a compiler trace. Appendix C specifies scene and event record construction and retrieval. Appendix A documents benchmark inputs, metric aggregation, and source-grounded event scoring. Appendix D reports memory and agent-component analyses. Appendix E gives the anonymized rating protocol, score distributions, and agreement results. Additional qualitative examples and complete long-horizon story sequences appear in Appendices F and G. No component is tuned or adapted for the reported evaluation.

REFERENCES

Zhaochong An, Menglin Jia, Haonan Qiu, Zijian Zhou, Xiaoke Huang, Zhiheng Liu, Weiming Ren, Kumara Kahatapitiya, Ding Liu, Sen He, et al. OneStory: Coherent multi-shot video generation with adaptive memory. *arXiv.org*, 2025. doi: 10.48550/arXiv.2512.07802.

Omer Bar-Tal, Hila Chefer, Omer Tov, Charles Herrmann, Roni Paiss, Shiran Zada, Ariel Ephrat, Junhwa Hur, Yuanzhen Li, T. Michaeli, et al. Lumiere: A space-time diffusion model for video generation, 2024. URL https://arxiv.org/abs/2401.12945.

Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *Computer Vision and Pattern Recognition*, pp. 22563–22575. IEEE, 2023. doi: 10.1109/CVPR52729.2023.02161.

Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *IEEE/CVF International Conference on Computer Vision*, pp. 14124–14133, 2021.

Haoxin Chen, Yong Zhang, Xiaodong Cun, Menghan Xia, Xintao Wang, Chao Weng, and Ying Shan. Videocrafter2: Overcoming data limitations for high-quality video diffusion models. In *Computer Vision and Pattern Recognition*, pp. 7310–7320. IEEE, 2024. doi: 10.1109/CVPR52733.2024.00698. URL https://arxiv.org/abs/2401.09047.

Yixiang Dai, Fan Jiang, Chiyu Wang, Mu Xu, and Yonggang Qi. Fantasyworld: Geometry-consistent world modeling via unified video and 3d prediction. *arXiv preprint arXiv:2509.21657*, 2025.

Yuan Gong, Youxin Pang, Xiaodong Cun, Menghan Xia, Yingqing He, Haoxin Chen, Longyue Wang, Yong Zhang, Xintao Wang, Ying Shan, et al. Talecrafter: Interactive story visualization with multiple characters. In *SIGGRAPH Asia 2023 Conference Papers*, pp. 1–10. ACM, 2023. doi: 10.1145/3610548.3618184. URL https://arxiv.org/abs/2305.18247.

Yuwei Guo, Ceyuan Yang, Anyi Rao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Sparsectrl: Adding sparse controls to text-to-video diffusion models. In *European Conference on Computer Vision*, pp. 330–348, 2023. doi: 10.48550/arXiv.2311.16933. URL https://arxiv.org/abs/2311.16933.

Hui Han, Siyuan Li, Jiaqi Chen, Yiwen Yuan, Yuling Wu, Chak Tou Leong, Hanwen Du, Junchen Fu, Youhua Li, Jie Zhang, et al. Video-bench: Human-aligned video generation benchmark, 2025. URL https://arxiv.org/abs/2504.04907.

Roberto Henschel, Levon Khachatryan, Hayk Poghosyan, Daniil Hayrapetyan, Vahram Tadevosyan, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Streamingt2v: Consistent, dynamic,

and extendable long video generation from text. In *Computer Vision and Pattern Recognition*, pp. 2568–2577, 2024. doi: 10.1109/CVPR52734.2025.00245. URL https://arxiv.org/abs/2403.14773.

Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P. Kingma, Ben Poole, Mohammad Norouzi, David J. Fleet, et al. Imagen video: High definition video generation with diffusion models, 2022a. URL https://arxiv.org/abs/2210.02303.

Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J. Fleet. Video diffusion models. In *Neural Information Processing Systems*, volume 35, pp. 8633–8646. Neural Information Processing Systems Foundation, Inc. (NeurIPS), 2022b. doi: 10.48550/arXiv.2204.03458. URL https://arxiv.org/abs/2204.03458.

Susung Hong, Junyoung Seo, Heeseong Shin, Sunghwan Hong, and Seungryong Kim. Direct2v: Large language models are frame-level directors for zero-shot text-to-video generation, 2023. URL https://arxiv.org/abs/2305.14330.

Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. In *International Conference on Learning Representations*, 2022. doi: 10.48550/arXiv.2205.15868. URL https://arxiv.org/abs/2205.15868.

Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qin Jin, Nattapol Chanpaisit, et al. VBench: Comprehensive benchmark suite for video generative models. In *Computer Vision and Pattern Recognition*, 2023. doi: 10.1109/CVPR52733.2024.02060.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):139–1, 2023.

Levon Khachatryan, Andranik Movsisyan, Vahram Tadevosyan, Roberto Henschel, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Text2video-zero: Text-to-image diffusion models are zero-shot video generators. In *IEEE International Conference on Computer Vision*, pp. 15908–15918. IEEE, 2023. doi: 10.1109/ICCV51070.2023.01462. URL https://arxiv.org/abs/2303.13439.

D. Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Rachel Hornung, Hartwig Adam, Hassan Akbari, Y. Alon, Vighnesh Birodkar, et al. Videopoet: A large language model for zero-shot video generation. In *International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 25105–25124. PMLR, 2023. doi: 10.48550/arXiv.2312.14125. URL https://arxiv.org/abs/2312.14125.

Yunxin Li, Haoyuan Shi, Baotian Hu, Longyue Wang, Jiashun Zhu, Jinyi Xu, Zhen Zhao, and Min Zhang. Anim-Director: A large multimodal model powered agent for controllable animation video generation. In *SIGGRAPH Asia 2024 Conference Papers*, pp. 1–11, 2024. URL https://arxiv.org/abs/2408.09787.

Han Lin, Abhay Zala, Jaemin Cho, and Mohit Bansal. VideoDirectorGPT: Consistent multi-scene video generation via llm-guided planning. *arXiv.org*, 2023. doi: 10.48550/arXiv.2309.15091.

Kuan Heng Lin, Zhizheng Liu, Pablo Salamanca, Yash Kant, Ryan Burgert, Yuancheng Xu, Koichi Namekata, Yiwei Zhao, Bolei Zhou, Micah Goldblum, Paul Debevec, and Ning Yu. Vista4D: Video reshooting with 4D point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 32671–32682, 2026. URL https://arxiv.org/abs/2604.21915.

Akide Liu, Jinbo Xing, Chaojie Mao, Ye Li, Zeyu Zhang, Yefei He, Weijie Wang, Zihan Wang, Yu Liu, Gholamreza Haffari, and Bohan Zhuang. ReCA: Multi-shot long video extrapolation via recursive context allocation. *arXiv preprint arXiv:2605.26525*, 2026. URL https://arxiv.org/abs/2605.26525.

Fuchen Long, Zhaofan Qiu, Ting Yao, and Tao Mei. VideoStudio: Generating consistent-content and multi-scene videos. In *European Conference on Computer Vision*, 2024. URL https://arxiv.org/abs/2401.01256.

Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, et al. Movie gen: A cast of media foundation models, 2024. URL https://arxiv.org/abs/2410.13720.

Xuanchi Ren, Tianchang Shen, Jiahui Huang, Huan Ling, Yifan Lu, Merlin Nimier-David, Thomas Müller, Alexander Keller, Sanja Fidler, and Jun Gao. Gen3c: 3d-informed world-consistent video generation with precise camera control. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6121–6132, 2025.

Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data, 2022. URL https://arxiv.org/abs/2209.14792.

Chenxi Song, Yanming Yang, Tong Zhao, Ruibo Li, and Chi Zhang. Worldforge: Unlocking emergent 3d/4d generation in video diffusion model via training-free guidance. *arXiv preprint arXiv:2509.15130*, 2025.

Zhenxiong Tan, Xingyi Yang, Songhua Liu, and Xinchao Wang. Video-infinity: Distributed long video generation, 2024. URL https://arxiv.org/abs/2406.16260.

Team HY-World, Chenjie Cao, Xuhui Zuo, Zhenwei Wang, Yisu Zhang, Junta Wu, Zhenyang Liu, Yuning Gong, Yang Liu, Bo Yuan, Chao Zhang, et al. Hy-world 2.0: A multi-modal world model for reconstructing, generating, and simulating 3d worlds. *arXiv preprint arXiv:2604.14268*, 2026.

Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(4):376–380, 1991. doi: 10.1109/34.88573.

Fu-Yun Wang, Wenshuo Chen, Guanglu Song, Han-Jia Ye, Yu Liu, and Hongsheng Li. Gen-l-video: Multi-text to long video generation via temporal co-denoising. *arXiv preprint arXiv:2305.18264*, 2023.

Jiahao Wang, Luoxin Ye, TaiMing Lu, Junfei Xiao, Jiahan Zhang, Yuxiang Guo, Xijun Liu, Rama Chellappa, Cheng Peng, Alan Yuille, et al. Evoworld: Evolving panoramic world generation with explicit 3d memory. *arXiv preprint arXiv:2510.01183*, 2025a.

Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. *arXiv preprint arXiv:2503.11651*, 2025b.

Qianqian Wang, Yifei Zhang, Aleksander Holynski, Alexei A Efros, and Angjoo Kanazawa. Continuous 3d perception model with persistent state. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10510–10522, 2025c.

Weijie Wang, Haoyu Zhao, Yifan Yang, Feng Chen, Zeyu Zhang, Yefei He, Zicheng Duan, Donny Y. Chen, Yuqing Yang, and Bohan Zhuang. Latent spatial memory for video world models. *arXiv preprint arXiv:2606.09828*, 2026.

Haoyu Wu, Diankun Wu, Tianyu He, Junliang Guo, Yang Ye, Yueqi Duan, and Jiang Bian. Geometry forcing: Marrying video diffusion and 3d representation for consistent world modeling. *arXiv preprint arXiv:2507.07982*, 2025a.

Tong Wu, Shuai Yang, Ryan Po, Yinghao Xu, Ziwei Liu, Dahua Lin, and Gordon Wetzstein. Video world models with long-term spatial memory. *arXiv preprint arXiv:2506.05284*, 2025b.

Weijia Wu, Zeyu Zhu, and Mike Zheng Shou. Automated movie generation via multi-agent CoT planning. *arXiv preprint arXiv:2503.07314*, 2025c. URL https://arxiv.org/abs/2503.07314.

Zhifei Xie, Daniel Tang, Dingwei Tan, Jacques Klein, Tegawend F. Bissyand, and Saad Ezzini. DreamFactory: Pioneering multi-scene long video generation with a multi-agent framework. *arXiv.org*, 2024. doi: 10.48550/arXiv.2408.11788.

Jinbo Xing, Menghan Xia, Yong Zhang, Haoxin Chen, Gongye Liu, Xintao Wang, Tien-Tsin Wong, and Ying Shan. Dynamicrafter: Animating open-domain images with video diffusion priors. In *European Conference on Computer Vision*, 2023. doi: 10.48550/arXiv.2310.12190. URL https://arxiv.org/abs/2310.12190.

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

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. ReAct: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations*, 2022. arXiv:2210.03629.

Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4578–4587, 2021.

Hong-Xing Yu, Haoyi Duan, Charles Herrmann, William T Freeman, and Jiajun Wu. Wonderworld: Interactive 3d scene generation from a single image. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5916–5926, 2025.

Wangbo Yu, Jinbo Xing, Li Yuan, Wenbo Hu, Xiaoyu Li, Zhipeng Huang, Xiangjun Gao, Tien-Tsin Wong, Ying Shan, and Yonghong Tian. Viewcrafter: Taming video diffusion models for high-fidelity novel view synthesis. *arXiv preprint arXiv:2409.02048*, 2024.

Zhengqing Yuan, Ruoxi Chen, Zhaoxu Li, Weixiang Sun, Haolong Jia, Lifang He, Chi Wang, and Lichao Sun. Mora: Enabling generalist video generation via a multi-agent framework. *arXiv.org*, 2024. doi: 10.48550/arXiv.2403.13248.

Kaiwen Zhang, Liming Jiang, Angtian Wang, Jacob Zhiyuan Fang, Tiancheng Zhi, Qing Yan, Hao Kang, Xin Lu, and Xingang Pan. Storymem: Multi-shot long video storytelling with memory. *arXiv preprint arXiv:2512.19539*, 2025.

Canyu Zhao, Mingyu Liu, Wen Wang, Weihua Chen, Fan Wang, Hao Chen, Bo Zhang, and Chunhua Shen. MovieDreamer: Hierarchical generation for coherent long visual sequence. In *The Thirteenth International Conference on Learning Representations*, 2025a. URL https://openreview.net/forum?id=QHj2LL958o.

Jinjing Zhao, Fangyun Wei, Zhening Liu, Hongyang Zhang, Chang Xu, and Yan Lu. Spatia: Video generation with updatable spatial memory. *arXiv preprint arXiv:2512.15716*, 2025b.

Mingzhe Zheng, Yongqi Xu, Haojian Huang, Xuran Ma, Yexin Liu, Wenjie Shu, Yatian Pang, Feilong Tang, Qifeng Chen, Harry Yang, and Ser-Nam Lim. VideoGen-of-Thought: Step-by-step generating multi-shot video with minimal manual intervention. *arXiv preprint arXiv:2412.02259*, 2025. URL https://arxiv.org/abs/2412.02259.

Jinsong Zhou, Yihua Du, Xinli Xu, Luozhou Wang, Zijie Zhuang, Yehang Zhang, Shuaibo Li, Xiaojun Hu, Bolan Su, and Ying-Cong Chen. VideoMemory: Toward consistent video generation via memory integration. *arXiv.org*, 2026. doi: 10.48550/arXiv.2601.03655.

Yupeng Zhou, Daquan Zhou, Ming-Ming Cheng, Jiashi Feng, and Qibin Hou. StoryDiffusion: Consistent self-attention for long-range image and video generation. In *Advances in Neural Information Processing Systems*, volume 37, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/c7138635035501eb71b0adf6ddc319d6-Abstract-Conference.html.

Ziyue Zhu, Zhanqian Wu, Zhenxin Zhu, Lijun Zhou, Haiyang Sun, Bing Wan, Kun Ma, Guang Chen, Hangjun Ye, Jin Xie, et al. Worldsplat: Gaussian-centric feed-forward 4d scene generation for autonomous driving. *arXiv preprint arXiv:2509.23402*, 2025.

Cailin Zhuang, Ailin Huang, Wei Cheng, Jingwei Wu, Yaoqi Hu, Jiaqi Liao, Zhewei Huang, Hongyuan Wang, Xinyao Liao, Weiwei Cai, et al. ViStoryBench: Comprehensive benchmark suite for story visualization. In *arXiv.org*, 2025. doi: 10.48550/arXiv.2505.24862. arXiv:2505.24862.

APPENDIX

This appendix provides implementation details and supporting analyses for reproducing and interpreting ATAS. It covers screenplay compilation, scene and event memory, benchmark evaluation, controlled analyses, human judgments, and extended story rollouts.

A EVALUATION PROTOCOLS AND BENCHMARK DETAILS

For each standard benchmark, the aggregate score is accompanied by metric-level details rather than story- or sequence-level listings. All benchmark-specific tables use the same backbone and method grouping as Tab. 1. The entries are interpreted with the evaluator definitions given in the corresponding subsection.

These breakdowns serve two purposes. First, they show which aspects of generation change when the rendering backbone changes, so that a single aggregate cannot be mistaken for a universal improvement. Second, they separate local rendering quality from the cross-shot behavior that the persistent state is intended to support. The dimensions are therefore read within a benchmark and within a backbone block. They are not combined into a new training objective or used to tune the reported agent.

A.1 SOURCE-GROUNDED DIAGNOSTIC EVALUATOR

The source-grounded diagnostic evaluator uses one fixed VLM judge and one versioned rubric for every method. The evaluator receives sampled source-event frames, the event description and interval, the target-camera request, an optional target-camera memory render, and sampled frames from the generated clip. It does not receive a synchronized target video. The memory render is treated as a requested-view condition rather than ground truth.

For each axis, the judge returns either a binary decision or an ordinal rating from 1 to 5. We map a binary decision to 0 or 1 and map an ordinal rating r to $(r-1)/4$. The normalized score for axis a is the mean over the rubric items for that axis. The Event Recapture Score is the arithmetic mean of the five normalized axes in Tab. 3: event identity, action and state preservation, temporal phase and order, character and object consistency, and viewpoint realization. We first average within each event request and then average across requests so that longer clips do not receive more weight.

The same rubric is used for all methods and memory configurations. A request without source-event evidence is excluded before aggregation. A request without a generated clip receives zero on the applicable axes. Human judgments on matched clips are used as a separate agreement analysis rather than as hidden calibration data.

The five axes intentionally distinguish what is recovered from the source event. Event identity asks whether the requested occurrence is the one being reproduced. Action and state evaluate what the participants and objects are doing. Temporal phase and order evaluate where the recapture lies within the source interval. Character and object consistency check persistent entities, and viewpoint realization checks the requested camera relation. A high score on one axis therefore does not imply a high score on the others. A clip can show the correct characters while depicting the wrong phase of the event, or preserve the action while failing to realize the requested view. This decomposition is why the aggregate is reported together with the per-axis values in Tab. 3.

The evaluator is source-grounded rather than pixel-grounded. It can judge whether the generated clip agrees with the supplied event evidence and camera request, but it cannot establish pixel-level agreement with an unobserved target video. The resulting score is appropriate for historical event access and view-conditioned recapture. It should not be read as a substitute for synchronized reconstruction when such a target is available.

Tab. A: **VBench dimension breakdown.** Each row is one reported VBench dimension, and each method/backbone column uses the official VBench evaluator and prompt suite. Higher is better unless the official dimension defines the opposite direction.

Dimension	Wan 2.7 ATAS	Wan 3.0 ATAS	MiniMax-H3 ATAS	HappyHorse ATAS
Aesthetic quality	0.6387	0.6731	0.6945	0.6792
Subject consistency	0.9481	0.9643	0.9576	0.9418
Background consistency	0.9554	0.9717	0.9680	0.9522
Temporal flickering	0.9932	0.9863	0.9847	0.9884
Motion smoothness	0.9888	0.9766	0.9714	0.9795
Dynamic degree	0.5488	0.6882	0.5842	0.6314

Tab. B: **ST-Bench metric breakdown.** The table reports official scores for aesthetic quality, global and single-shot prompt following, all-pair cross-shot consistency, and top-ten-pair cross-shot consistency.

Metric	Wan 2.7 ATAS	Wan 3.0 ATAS	MiniMax-H3 ATAS	HappyHorse ATAS	ReCA reference
Aesthetic	0.6387	0.6731	0.6945	0.6792	0.6517
PF-G	0.2174	0.2415	0.2351	0.2289	0.2216
PF-S	0.2154	0.2362	0.2308	0.2243	0.2192
C-All	0.5977	0.6844	0.6529	0.6615	0.5751
C-Top10	0.6253	0.7187	0.6908	0.6684	0.6015

A.2 VBench Dimensions

The VBench dimensions show that the rendering backbones trade different intrinsic properties. Wan 3.0 is strongest on subject and background consistency and dynamic degree, while Wan 2.7 is stronger on flicker and motion smoothness and MiniMax-H3 is strongest on aesthetic quality. The spread is largest for dynamic degree (0.5488–0.6882), whereas the temporal-flickering and motion-smoothness values remain high for every backbone. Thus the advantage in the standard multi-shot results cannot be explained by one uniformly better renderer: the backbones begin with different local strengths. This breakdown also gives a useful control for the later analyses. A change in cross-shot consistency should be interpreted alongside these local properties, rather than as evidence that ATAS improves every frame-level dimension equally. The table therefore separates backbone-dependent rendering behavior from the cross-shot gains measured by ST-Bench and ViStory.

A.3 ST-Bench Metrics

The ST-Bench breakdown is consistent with the aggregate comparison: the strongest ATAS columns are the cross-shot consistency measures, especially C-Top10. Prompt-following changes are smaller and more uniform than the consistency changes: for example, Wan 3.0 reaches 0.2415 on PF-G and 0.2362 on PF-S, while its C-All and C-Top10 values are 0.6844 and 0.7187. The gap between C-Top10 and C-All is positive for every ATAS backbone (0.0069–0.0379), so the two consistency summaries provide complementary checks rather than duplicate columns. Wan 3.0 gives the strongest overall cross-shot behavior, with MiniMax-H3 and HappyHorse retaining substantial gains over the reference column. The pattern is therefore more specific than a general prompt-following gain: it is consistent with state being queried at shot boundaries and helping preserve information across the sequence.

A.4 ViStory Metrics

The ViStory-Self dimensions show that the gains are not confined to one identity score. Wan 3.0 has the strongest appearance and temporal self-consistency and the highest aggregate, while the other ATAS backbones remain above the reference aggregate. Appearance consistency is the largest component for all four ATAS columns. Identity and temporal consistency move in the same direction, with Wan 3.0 reaching 0.6617, 0.8139, and 0.7908, respectively. A method that merely improves texture or color would not be expected to improve all three within-story dimensions together. The result indicates that the queried state supports both within-story identity and appearance continuity, while leaving room for future improvements in identity preservation.

Tab. C: **ViStory-Self metric breakdown.** Each row is an official ViStory-Self submetric. The final row reports its official aggregate when available.

Metric	Wan 2.7 ATAS	Wan 3.0 ATAS	MiniMax-H3 ATAS	HappyHorse ATAS	ReCA reference
Identity consistency	0.5494	0.6617	0.6421	0.5974	0.5608
Appearance consistency	0.6836	0.8139	0.7476	0.7772	0.7034
Temporal self-consistency	0.7648	0.7908	0.7392	0.7425	0.7186
Aggregate	0.6659	0.7555	0.7096	0.7057	0.6609

Tab. D: **ViStory-Cross metric breakdown.** Each row is an official ViStory-Cross submetric. The final row reports its official aggregate when available.

Metric	Wan 2.7 ATAS	Wan 3.0 ATAS	MiniMax-H3 ATAS	HappyHorse ATAS	ReCA reference
Cross-shot identity	0.5073	0.5992	0.5736	0.5428	0.4934
Cross-shot appearance	0.4871	0.5806	0.5392	0.5618	0.5013
Cross-shot scene consistency	0.4621	0.5729	0.5285	0.5138	0.4484
Aggregate	0.4855	0.5842	0.5471	0.5395	0.4810

The ViStory-Cross dimensions follow the same ordering as the main cross-shot results. Wan 3.0 leads identity, appearance, and scene consistency, and every ATAS backbone exceeds the reference aggregate in this table. Scene consistency is the most demanding of the three columns for each ATAS backbone, which is expected because it depends on preserving the shared environment rather than only the subject. The relative ordering is stable across identity, appearance, and scene consistency. No backbone obtains its aggregate advantage from a single isolated submetric. These dimensions provide a direct check that the behavior persists across different shots rather than only within individual clips.

B World-state compiler and bounded shot requests

The world-state compiler receives a high-level screenplay request and the current scene and event evidence. It recursively resolves scene transitions, temporal relations, historical-event references, and camera requirements until each bounded shot request fits the short-video generation budget. Each request provides a local shot description, an explicit scene coordinate and target camera, and a requested time interval with optional event information. The compiler does not render a video while an internal request is being expanded. It emits a bounded shot request and waits for its validated observation before deciding which state a subsequent request may use.

Fig. A gives a concrete compiler trace from screenplay request to bounded shot requests and validated state updates.

The trace makes the unit of generation explicit. The screenplay is resolved into requests that carry a local description together with scene, camera, time, and optional event fields. These fields determine which evidence can be queried before generation and which evidence can be attached to the resulting observation afterward. A request that changes location can therefore establish a new scene relation, while a request that revisits an earlier location can reuse its scene coordinates. An event request additionally carries the time interval needed to select dynamic evidence. The figure checks that decomposition precedes rendering and that a later request consumes the validated state produced by an earlier one.

C Memory Construction and Query

Each generated clip is retained as a shot-local observation with its request, source provenance, timestamps, scene and event identifiers, target camera, reconstruction artifacts, and independent film, registration, and state-update statuses. The scene memory stores registered static evidence after dynamic regions are excluded. The event memory stores time-local dynamic evidence bound to a source shot, an event identifier, and a source interval. The two memories are queried separately and combined when a request refers to a historical event.

Fig. B connects each shot-local observation to the static and event query paths.

The separation is useful because the two evidence types have different persistence requirements. Static geometry is registered once to a scene coordinate frame and can be rendered at a new camera,

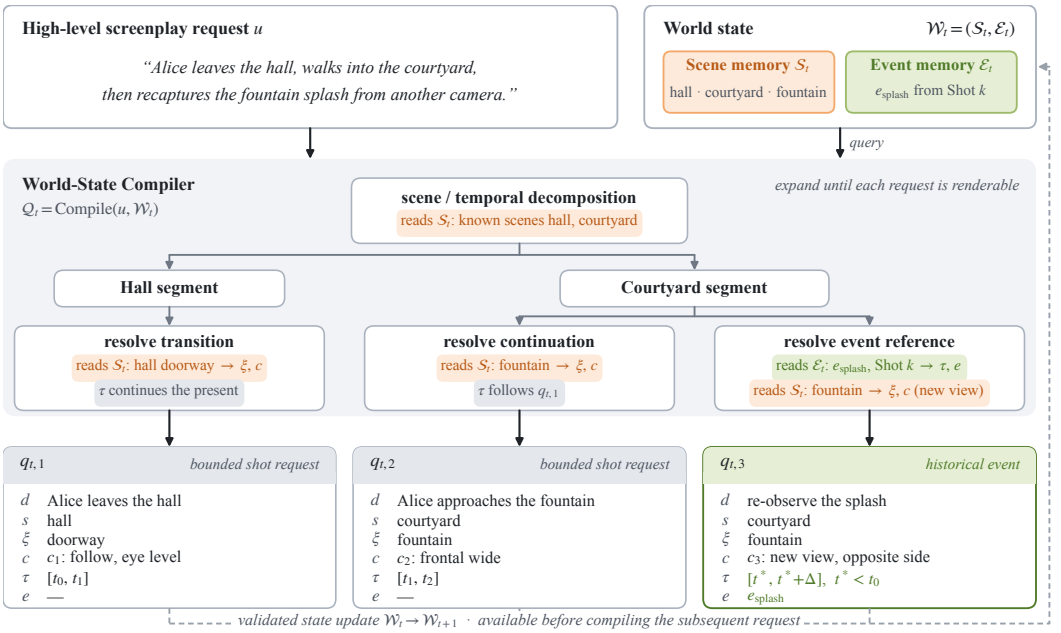

Fig. A: **Recursive screenplay compilation.** The compiler turns a high-level request into bounded shot requests and uses validated state updates to condition subsequent requests.

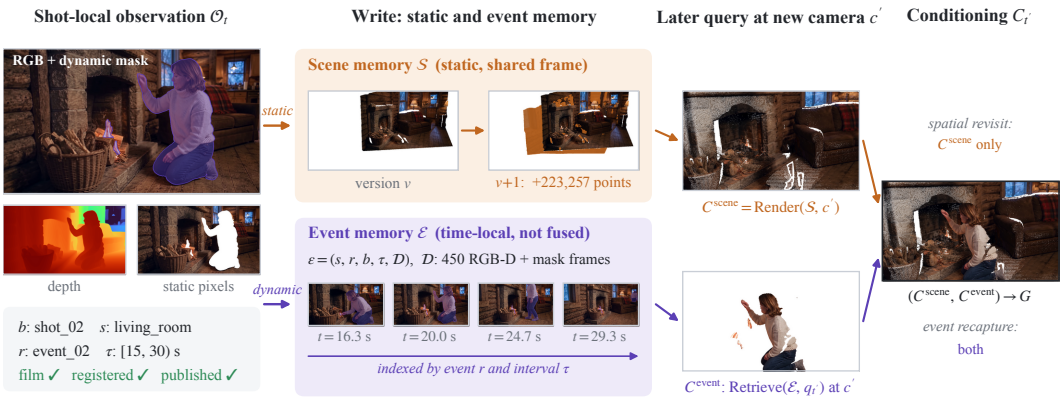

Fig. B: **Scene and event memory.** The figure follows one accepted shot-local observation as static evidence is written to scene memory and dynamic evidence is indexed in event memory, then shows how both references condition a later camera request.

whereas dynamic evidence must remain bound to an event and a time interval. Merging them into a single undifferentiated history would make a visually similar frame eligible even when it belongs to the wrong occurrence. The figure also clarifies the direction of information flow: reconstruction and registration create a candidate observation after generation, while rendering and event retrieval create conditioning references before the next generation. This distinction is central to interpreting the controlled memory rows in Tab. 5.

The memory statistics show increasing scene and event coverage over the rollout while query and update time changes only modestly. From the initial state to the end of the rollout, scene size grows from 33.54 to 49.60 million units and event records from 1.0 to 4.8, while query and update time rises from 93.0 to 103.5. Storage increases from 0.95 to 2.44 in the reported units, so the state is not free. The cost grows more slowly than the amount of processed evidence in this short rollout. The processed-shot count also increases from 1.2 to 6.9, which makes the checkpoint comparison a measure of accumulation rather than a comparison of independent runs. These values should be read as an operating profile for the reported configuration, not as a general scaling law. They support

Tab. E: **Memory statistics.** This analysis table reports memory growth and query cost at rollout checkpoints. It is not a sample manifest.

Checkpoint	Scene size ($\times 10^6$)	Event records	Storage (GB)	Query and update time (s)	Processed shots
Initial state	33.54	1.0	0.95	93.0	1.2
After accepted update	40.92	2.1	1.38	97.8	2.8
End of rollout	49.60	4.8	2.44	103.5	6.9

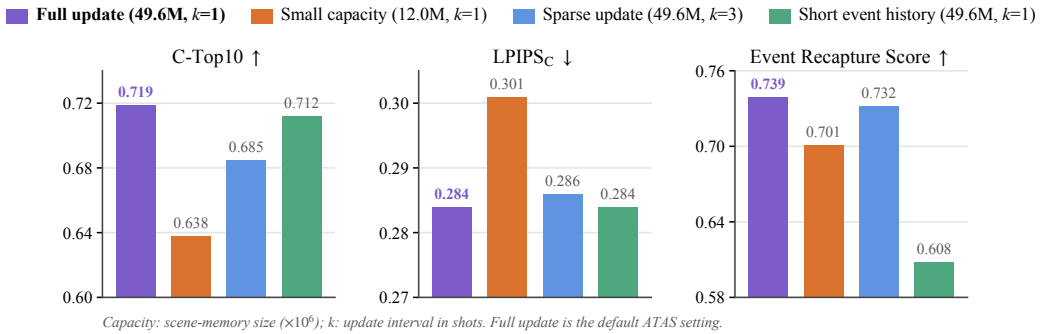

Fig. C: **Memory capacity and update sensitivity.** The figure shows whether quality gains require additional state size or frequent updates.

treating the state as a growing evidence store whose cost is measured separately from the video generator.

D Sensitivity and Mechanism Analysis

The sensitivity analysis keeps the prompt package, camera requests, generator, candidate budget, and retry policy fixed while varying memory capacity, update frequency, memory freshness, and state update check strictness. The figures diagnose how state quality changes. They do not introduce a second benchmark.

The sensitivity table shows that sparse updates and reduced scene capacity lower both cross-shot consistency and event quality, while shortening the event history mainly affects the event score. Relative to full update, sparse update reduces C-Top10 from 0.719 to 0.685 and the event score from 0.739 to 0.732. Reducing capacity has a larger spatial effect (C-Top10 0.638 and LPIPS$_C$ 0.301), while shortening event history reduces the event score to 0.608 with almost unchanged LPIPS$_C$. Spatial quality therefore changes less than event quality when the event memory is shortened, which is consistent with the separation between static scene evidence and time-local event evidence. The comparison also suggests that update frequency and memory capacity are distinct controls: more frequent updates cannot fully compensate for a substantially smaller scene store, and a large scene store cannot replace missing event history.

Quality declines as more shots are processed after the last accepted update. C-Top10 decreases from 0.736 at zero intervening shots to 0.651 after five or more, while LPIPS$_C$ increases from 0.276 to 0.306. The event score follows the same direction, from 0.762 to 0.678. The gradual decline across all three measures is more informative than any single endpoint. It suggests that stale state gradually loses usefulness rather than at one fixed horizon. These bins do not identify whether the loss comes from scene changes, event density, or accumulated registration error. They motivate the separate state-update analysis below.

Fig. E visualizes the quality and state-coverage trade-off of the state update check.

The state update check exposes a quality-coverage trade-off. A permissive setting admits 90.0% of candidate updates but produces lower downstream quality (C-Top10 0.696 and event score 0.708), whereas a strict setting admits 36.0% and also loses quality because fewer observations remain available (0.668 and 0.682). The default setting accepts 67.4% and obtains the strongest downstream scores, 0.719 and 0.739, together with the lowest LPIPS$_C$ of the three settings. This result does not imply that the default threshold is universally optimal. It shows that both indiscriminate admission

Tab. F: **Memory capacity and update sensitivity.** The table supplies the values used by Fig. C.

Configuration	Capacity ($\times 10^6$)	Update interval	C-Top10	LPIPS_C	Event score
Full update	49.60	1	0.719	0.284	0.739
Sparse update	49.60	3	0.685	0.286	0.732
Small capacity	12.00	1	0.638	0.301	0.701
Short event history	49.60	1	0.712	0.284	0.608

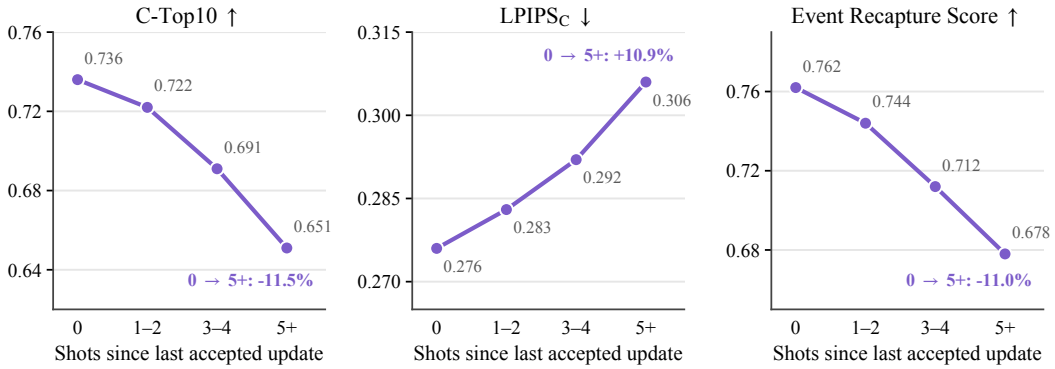

Fig. D: **Memory freshness and staleness.** The figure tests whether quality degrades as future requests operate farther from the most recent accepted observation.

and excessive rejection can reduce the evidence available to future requests. The check is therefore useful as a controlled admission mechanism, not as a replacement for improving reconstruction or retrieval.

Tab. G: **Memory freshness data.** The table supplies the values used by Fig. D.

Shots since update	C-Top10	$LPIPS_C$	Event Recapture Score
0	0.736	0.276	0.762
1 to 2	0.722	0.283	0.744
3 to 4	0.691	0.292	0.712
5+	0.651	0.306	0.678

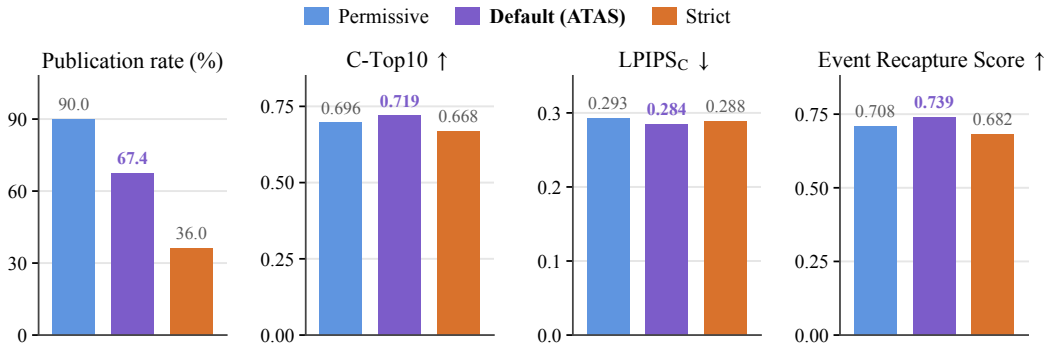

Fig. E: **State update check trade-off.** The figure shows how stricter admission changes state coverage and downstream quality.

Tab. H: **State update check sensitivity.** The table supplies the values used by Fig. E. It reports no independent benchmark.

State update setting	State update rate	C-Top10	$LPIPS_C$	Event score
Permissive	90.0	0.696	0.293	0.708
Default	67.4	0.719	0.284	0.739
Strict	36.0	0.668	0.288	0.682

E HUMAN EVALUATION

The human study follows the grouped-bar comparison in Fig. 5. Raters view anonymized clips in randomized order and score visual appeal, script faithfulness, character consistency, background consistency, physical law, and narrative coherence on a fixed five-point scale. The grouped bars report the six criteria, while this section reports score distributions and inter-rater agreement.

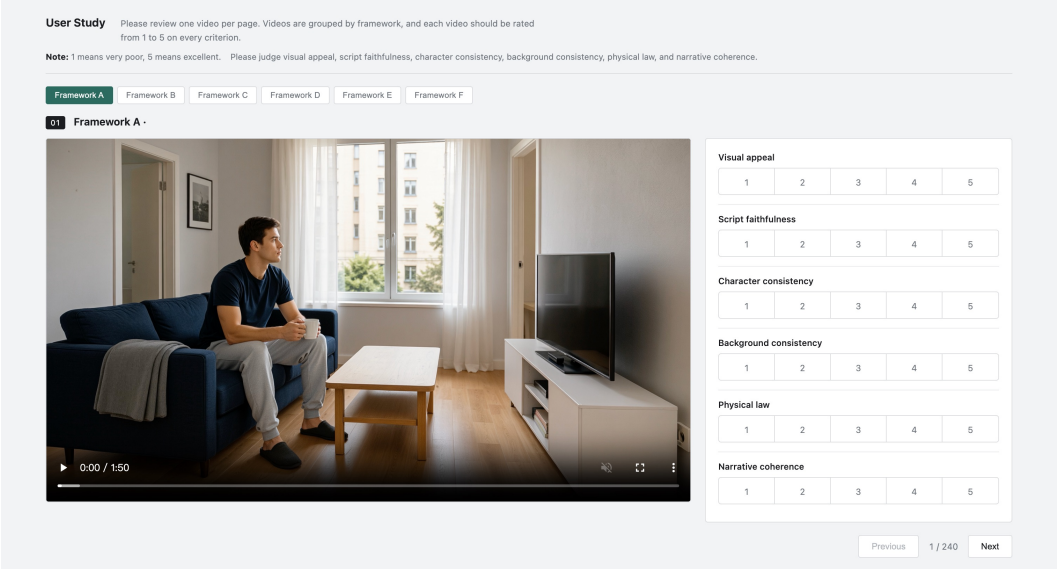

Fig. F: **Human evaluation interface.** Each page shows one anonymized long video with its story title and synopsis. Raters watch the whole video and score it on the six criteria from 1 (very poor) to 5 (excellent). The method identity is hidden, and all six criteria must be rated before moving to the next video.

Fig. F shows the rating interface. Each video is rated as a whole rather than shot by shot, so the consistency criteria reflect the complete film.

Fig. G shows score spread, and Fig. H compares the automatic diagnostics with human judgments.

The human-rating heatmaps show a consistent method ordering across criteria. ATAS has the highest mean rating in all six rows, with means from 3.62 to 4.40. ReCA is the strongest baseline, with means from 2.62 to 3.67. The fraction of ratings at least four points makes the same distinction: ATAS ranges from 57% to 95%, whereas ReCA ranges from 14% to 64%. The agreement coefficients shown beside the heatmap range from 0.58 to 0.72, so the comparison is not based on a single criterion with unusually unstable ratings. These values support the qualitative conclusion that the benefit is visible to raters. The spread across criteria shows that narrative coherence and physical-law judgments remain harder than script faithfulness and character consistency.

The agreement plot provides a second check on the automatic diagnostics. Correlation with human ratings is 0.61, 0.54, 0.49, 0.41, and 0.57 for event identity, action and state, temporal phase, character and object, and viewpoint, respectively. The spatial LPIPS diagnostic correlates at 0.44. The corresponding human–human ceilings are higher (0.67–0.80), leaving room for evaluator noise and for aspects of perception that are not captured by the automatic scores. The result supports using the automatic metrics as structured diagnostics rather than treating them as complete substitutes for human judgment.

F QUALITATIVE RESULTS

Figures I–K provide the qualitative audit, rollout timeline, and multi-story grid.

The baseline comparison highlights differences in shot-to-shot continuity. In the record-shop and bakery sequences, methods vary in whether room layout, object placement, and character state remain

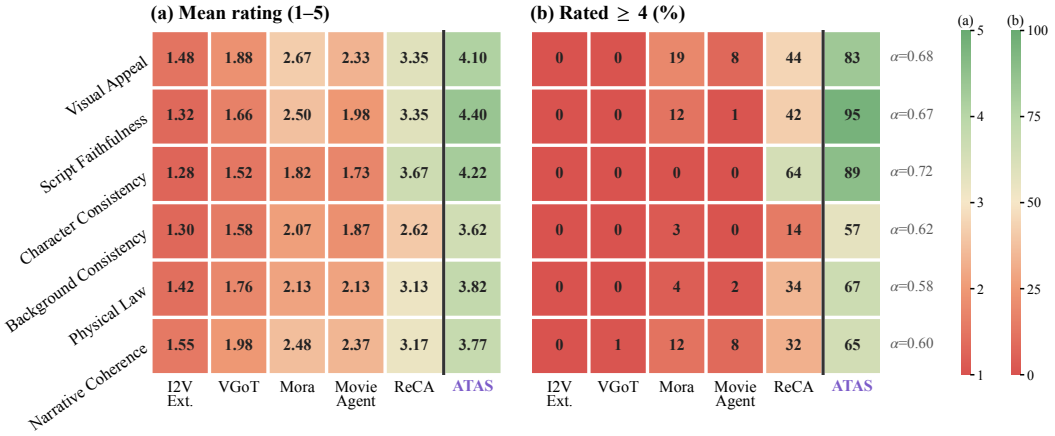

Fig. G: **Distribution of human judgments.** The heatmaps report mean ratings and the fraction of ratings above four for each criterion and method. The side panel summarizes inter-rater agreement.

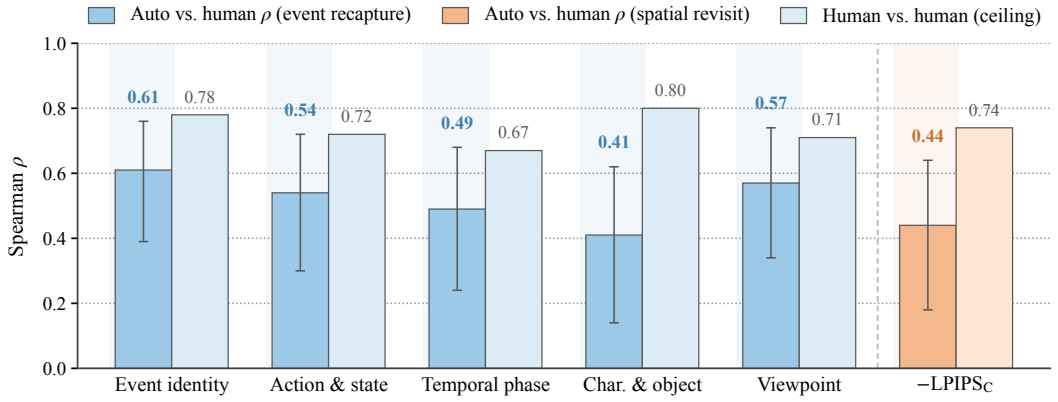

Fig. H: **Automatic-to-human agreement.** Spearman correlations compare the event-recapture and spatial-revisit diagnostics with human ratings on matched clips, alongside the human–human agreement ceiling.

consistent after intervening shots. The ATAS row provides a reference for this behavior. A clip can remain plausible locally while drifting in layout or entity state, which is not captured by a single global similarity score. The rollout timeline shows when a scene state is created, when a later request queries it, and when a returned observation is written back. The six-story grid shows that these behaviors span different settings, people, objects, and camera patterns.

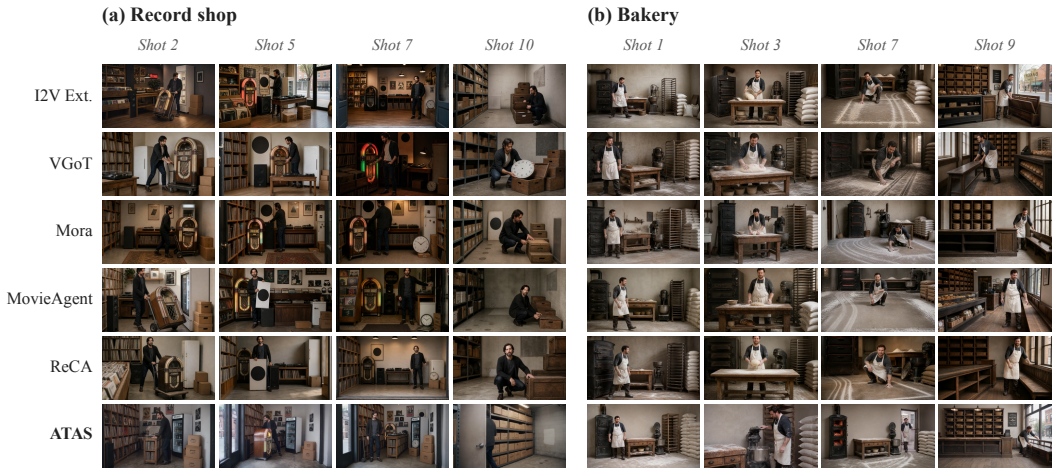

Fig. I: **Qualitative baseline audit.** Matched frames from the record-shop and bakery stories compare methods across shot continuation, scene layout, object state, and character consistency.



Fig. J: **Representative rollout timeline.** The timeline follows the generated film and memory condition through scene changes, a spatial revisit, and an event recapture request.



Fig. K: **Representative multi-story results.** The grid shows timestamped frames from six stories, illustrating the range of scenes, actions, and camera requests used in the qualitative evaluation.

G LONG-HORIZON RESULT GALLERY

This section presents three full-length ATAS films, one per subsection. Each story visits three locations and then returns to two of them after other locations have been shown in between. A return is where a shot-local generator has no access to what happened earlier in the film, so the returns are the direct test of whether the world state persists: the layout and objects must match the first visit, and changes made by the character must still be there, or must have progressed while the character was away.

Every story is presented in the same order. An overview figure first shows eight frames of the film in story order. The five visits are then described one by one, each with a large frame taken from that visit. A final figure pairs each return with the first visit to the same location. All frames are taken directly from the generated films, and timestamps give the film time in minutes and seconds.

G.1 SATURDAY LAUNDRY IN THE FLAT

A young man does his weekend laundry. The 1:50 film follows him from the living room to the laundry room and the bedroom, and then back to the living room and the laundry room. The story leaves two pieces of state behind: a mug set down on the coffee table, and a washer that is loaded and started before he leaves the laundry room.

Visit 1: living room (0:00–0:15). The film opens on a quiet Saturday morning. He sits on the dark blue sofa with a mug of coffee, facing the TV, with a gray throw over the arm of the sofa and a low wooden coffee table in front of him. Before leaving, he sets the mug down on the table. This is the first observation of the living room, and the room is reconstructed from it together with the position of the mug.

Visit 2: laundry room (0:15–0:45). He carries the basket into the small tiled laundry room, loads the clothes into the front-loading washer, adds detergent from the blue bottle on the shelf, and starts the machine. The shelf with the detergent, the plant, and the jar, the towel on the wall, and the folding drying rack by the door are established here. When he leaves, the washer is running and the basket is empty.

Visit 3: bedroom (0:45–1:15). While the wash runs, he goes to the bedroom, lies down on the unmade bed, gets up to open the curtains, and sits in the green armchair by the window. The bedroom is seen only once in the film, so this visit only adds a new room to the world state. It also separates the two returns from the rooms in which the state was left.

Visit 4: living room, return (1:15–1:40). He returns to the living room about a minute after leaving it. The stored living room conditions this visit, so the sofa, the throw, the coffee table, and the TV unit reappear in the same arrangement, and the mug he left behind is still standing on the table when he sits down again.



Fig. G.1: **Saturday Laundry in the Flat: overview.** He leaves the living room for the laundry room, adds detergent and closes the washer, sits and then lies down on the bed, opens the curtains, and later gets up from the sofa to walk back to the laundry room.



Fig. G.2: **Saturday Laundry, visit 1: living room (0:02).** The first view of the living room, with the sofa, the gray throw, the coffee table, and the TV unit.



Fig. G.3: **Saturday Laundry, visit 2: laundry room (0:22).** He loads the clothes from the basket into the washer.



Fig. G.4: **Saturday Laundry, visit 3: bedroom (1:07).** He sits in the armchair after opening the curtains.



Fig. G.5: **Saturday Laundry, visit 4: living room, return (1:25).** Back on the sofa, with the mug still on the coffee table where he set it down.



Fig. G.6: **Saturday Laundry, visit 5: laundry room, return (1:45).** He opens the washer he started more than a minute earlier.

Visit 5: laundry room, return (1:40–1:50). He goes back to collect the wash. The laundry room is rendered from its stored state after the wash was started: the shelf, the towel, the basket, and the drying rack are where they were, and the washer is still shut with the laundry inside. He opens the door of the finished machine.

G.2 Winter Fire in the Cabin

A woman lights the fire in a cold log cabin. The 2:10 film moves from the living room to the kitchen and the bedroom, and then back through the kitchen to the living room. Both returns test state that she created herself: a heart drawn in the frost on the kitchen window, and a fire lit in the hearth before she leaves the living room.

Visit 1: living room (0:00–0:30). She enters a cold living room with a stone fireplace, a log basket on the hearth, and a brown sofa with a red plaid throw. She kneels at the hearth, builds the fire from the logs in the basket, and lights it. When she leaves, a small fire is burning in a hearth that was cold at the start of the film.

Visit 2: kitchen (0:30–0:55). In the kitchen, frost covers the window above the counter. She draws a heart in the frost with her finger, then turns to the jar and the oil lamp on the counter beside the



Fig. G.7: **Saturday Laundry in the Flat: returns.** Top: the mug set down on the coffee table at the start is still there more than a minute later, next to the same sofa and throw. Bottom: the laundry room keeps its shelf, towel, basket, and drying rack, and the washer he started is still shut with the laundry inside.



Fig. G.8: **Winter Fire in the Cabin: overview.** She builds the fire from the log basket and watches it catch, stands by the frosted kitchen window with the heart she drew, rests in the bedroom, and returns to tend the fire, which has grown while she was away.



Fig. G.9: **Winter Fire, visit 1: living room (0:12).** She lights the fire in the stone hearth.



Fig. G.10: **Winter Fire, visit 2: kitchen (0:35).** She draws in the frost on the kitchen window.



Fig. G.11: **Winter Fire, visit 3: bedroom (1:10).** She lies on the bed under the lamp.

dresser of blue-and-white plates. The heart is a small, specific mark that the kitchen must still carry when she comes back.

Visit 3: bedroom (0:55–1:25). She rests in the bedroom, sitting on and then lying across the bed with the patchwork quilt, lit by the bedside lamp. The bedroom is seen only once. It adds a third room to the cabin and places half a minute of other content between the kitchen and its return.

Visit 4: kitchen, return (1:25–1:40). On her way back she passes through the kitchen. The stored kitchen conditions this shot: the stove, the counter with the oil lamp, the dresser, and the doorways are in the same places, and the heart she drew is still visible in the lower right pane of the frosted window.

Visit 5: living room, return (1:40–2:10). She returns to the living room about seventy seconds after lighting the fire. The fireplace, the log basket, the lamp, and the sofa with the red plaid throw match the first visit, and the fire she lit has kept burning and now fills the hearth. She settles into the armchair under a plaid blanket.

G.3 Running a Bath on a Winter Evening

A young woman, tired from working on her laptop, runs herself a bath. The 2:10 film moves from the bedroom to the bathroom and the living room, and then back to the bedroom and the bathroom. The



Fig. G.12: **Winter Fire, visit 4: kitchen, return (1:30).** The kitchen as she left it, with the heart still in the frost on the window.



Fig. G.13: **Winter Fire, visit 5: living room, return (2:05).** She sits by the fire she lit before leaving the room.

two returns test a static object left behind, the closed laptop on the bed, and a process that continues while she is away, the bath filling from the running tap.

Visit 1: bedroom (0:00–0:15). She sits cross-legged on the bed in a warmly lit bedroom with pink curtains, a white wardrobe, and a lamp on the nightstand, working on her laptop. She closes the laptop, leaves it on the bed, and walks out through the door to the bathroom.

Visit 2: bathroom (0:15–0:45). In the tiled bathroom she turns on the tap over the empty tub, pours in bath salts so that the water turns pink, and lights a candle on the glass shelf below the round mirror. She leaves with the tap running and the tub only partly filled.

Visit 3: living room (0:45–1:15). While the bath runs, she waits in the living room, looking out over the city lights, browsing the bookshelf, and resting on the sofa beside a small candle. The living room is seen only once, and it separates both returns from the rooms in which the state was left.

Visit 4: bedroom, return (1:15–1:40). She comes back to the bedroom a minute after leaving it. From the same view as the opening shot, the bed, the nightstand lamp, the wardrobe, and the pink curtains match the first visit, and the laptop she closed is still lying on the bed. She sits beside it before heading to the bath.



Fig. G.14: **Winter Fire in the Cabin: returns.** Top: the heart she drew in the frost is still in the lower right pane of the same kitchen window. Bottom: the living room keeps its stone fireplace, log basket, lamp, and sofa with the red plaid throw, while the hearth that was cold at the start now holds the fire she lit.



Fig. G.15: **Running a Bath on a Winter Evening: overview.** She leaves the laptop on the bed, turns on the tap, lights a candle as the tub fills with pink water, browses the bookshelf in the living room, opens the wardrobe back in the bedroom, and tests the water in the filled tub.



Fig. G.16: **Running a Bath, visit 1: bedroom (0:07).** She closes the laptop on the bed.



Fig. G.17: **Running a Bath, visit 2: bathroom (0:30).** The tub begins to fill with pink water.



Fig. G.18: **Running a Bath, visit 3: living room (1:07).** She rests on the sofa while the bath runs.



Fig. G.19: **Running a Bath, visit 4: bedroom, return (1:30).** She sits next to the laptop she left on the bed.



Fig. G.20: **Running a Bath, visit 5: bathroom, return (2:05).** She kneels by the tub, which has filled while she was away.



Fig. G.21: **Running a Bath on a Winter Evening: returns.** Top: from the same view, the bedroom matches the first visit, with the bed, nightstand lamp, wardrobe, and pink curtains, and the laptop she closed is still lying on the bed. Bottom: the tiles, glass shelf, mirror, and towel rail match her first entrance, while the empty tub has filled with pink water and the candle she lit is still burning.

Visit 5: bathroom, return (1:40–2:10). Back in the bathroom, the tiles, the glass shelf, the mirror, and the towel rail match her first entrance. The tub that was only partly filled when she left is now full of pink water, and the candle she lit is still burning. She kneels by the tub with a towel.