

---

# Remember to be Curious: Episodic Context and Persistent Worlds for 3D Exploration

---

Lily Goli<sup>1,3,4</sup> Justin Kerr<sup>2</sup> Daniele Reda<sup>3</sup> Alec Jacobson<sup>1,4</sup>  
Andrea Tagliasacchi<sup>1,3,5</sup> Angjoo Kanazawa<sup>2</sup>

<sup>1</sup>University of Toronto   <sup>2</sup>UC Berkeley   <sup>3</sup>Wayve   <sup>4</sup>Vector Institute   <sup>5</sup>Simon Fraser University  
lily.goli@mail.utoronto.ca

[recuriosity.github.io](https://recuriosity.github.io)




Figure 1: **Remember to be Curious**. Test-time trajectories of our curiosity-driven policy trained with episodic context and a persistent world model. **Left**: Our agent explores an indoor scene purely from an image stream; waypoints show the agent’s views along the path. **Right**: Our end-to-end design enables fine-tuning to downstream sparse-reward tasks and zero-shot generalization to AI-generated worlds. Note: the model requires no explicit mapping at test time; trajectories are overlaid on bird’s-eye-view for visualization only.

## Abstract

Exploration is a prerequisite for learning useful behaviors in sparse-reward, long-horizon tasks, particularly within 3D environments. Curiosity-driven reinforcement learning addresses this via intrinsic rewards derived from the mismatch between the agent’s predictive model of the world and reality. However, translating this intrinsic motivation to complex, photorealistic environments remains difficult, as agents can become trapped in local loops and receive fresh rewards for revisiting forgotten states. In this work, we demonstrate that this failure stems from a lack of spatial persistence and episodic context. We show that effective curiosity requires a model of the world that is persistent and continuously updated, paired with an agent that maintains an episodic trajectory history to navigate toward novel regions. We achieve this using an online 3D reconstruction as a persistent model of the world, while the agent policy is parameterized as a sequence model over RGB observations to maintain episodic context. This design enables effective exploration during training while allowing the agent to navigate using solely RGB frames at deployment. Trained purely via curiosity on HM3D, our agent outperforms active-mapping baselines and generalizes zero-shot to Gibson and AI-generated worlds. Our end-to-end policy enables efficient adaptation to downstream tasks, such as apple picking and image-goal navigation, outperforming from-scratch baselines.


# 1 Introduction

In his seminal work on latent learning, Edward Tolman demonstrated that agents—such as rats navigating a maze—can acquire complex knowledge of their environment even in the absence of explicit rewards [1]. This suggests that intelligent agents possess an inherent drive to explore, with curiosity serving as a mechanism for engaging with novel and uncertain environments. This phenomenon is intuitively observable in any toddler at a playground: they are never paralyzed by the absence of an explicit goal. Rather than waiting for a sparse external reward, it is far more natural for them to engage in unstructured play, freely wandering and exploring their surroundings.

In the context of artificial agents, this exploratory drive has the practical benefit of densifying sparse rewards in long-horizon tasks, such as visuomotor navigation. Despite the progress in goal-conditioned navigation [2, 3], in real-world scenarios, these tasks fundamentally carry *sparse* success signals. Curiosity-driven exploration offers a compelling alternative, in which the agent derives intrinsic reward from surprise, a signal produced by the prediction error of a world model — trained alongside the agent — to anticipate the consequences of its actions [4, 5]. However, scaling curiosity with end-to-end policies to complex photorealistic environments remains challenging, as agents frequently collapse into repetitive behavior driven by cyclic curiosity rewards.

In this work, we show that enabling curiosity requires both a persistent model of the world and an agent equipped with episodic context. This dual requirement addresses a two-fold problem of “amnesiac” exploration: without historical context, agents repeatedly revisit the same locations; simultaneously, without a persistent and continuously updating world model, predictive errors spuriously arise in these revisited areas, yielding false novelty rewards for forgotten states.

To study this, we focus on agents exploring static 3D photorealistic scenes. Ideally, the world model driving curiosity should be dynamic and continuously refined through lifelong agentic experience—a paradigm where action-conditioned video models show promise. However, spatial persistence remains a critical bottleneck for these models; for instance, LingBot-World [6] fails to maintain scene consistency, generating an entirely different image after a simple in-place 360° turn (see inset figure). We therefore utilize a state-of-the-art online 3D reconstruction method (3DGS) as a proxy, as it enables spatial persistence while continuously updating itself online. Traditional methods like ICM [5] lack this property, as their learned world model acts as a statistical prior over lifelong experience rather than an episodic record of the environment. Indeed, our experiments confirm that when the persistent memory is artificially capped to a short-term window, exploration capability deteriorates significantly.



On the agent side, we build a transformer-based policy that operates purely on a sequence of RGB observations. By conditioning the policy on the agent’s recent RGB history, we enable it to form its own internal representation of the episodic context directly from pixels. This is in direct contrast to existing methods that bypass the exploration challenge by embedding *explicit geometric maps* to guide the agent [7–9], sometimes relying on ground-truth depth observations. While explicit mapping guarantees spatial awareness, such modularity often abstracts away rich semantic information, sacrifices the flexibility of end-to-end learning, and limits generalization to geometrically definable tasks. Our image-sequence agent, conversely, requires no depth sensor or localization at deployment and adapts flexibly to semantically diverse downstream tasks. Furthermore, we show that scheduled, temporary injections of random behavior into RL training suffice to overcome the reward-less stretches of long-horizon exploration [10]. This enables an emergent behavior where the agent learns to traverse already-seen regions to find novel “branches” of unseen places. Crucially, this is achieved without the imitation learning bootstrapping [11], or hierarchical goal selection of prior work [7–9], highlighting curiosity’s potential in sparse-reward settings.

Trained purely from curiosity on Habitat Matterport 3D (HM3D) [12], our agent outperforms RL-based active mapping baselines and generalizes zero-shot to Gibson [13], and to out-of-distribution scenes from an entirely AI-generated 3D world [14] never encountered during training. Crucially, because our policy does not require explicit mapping at test time, it remains fully end-to-end and highly adaptable. Fine-tuned for just a few episodes on external task rewards, it outperforms policies trained from scratch. Our work demonstrates that coupling a persistent model of the world with an episodic agent provides practical insights into how curiosity-driven exploration can be scaled to complex, photorealistic 3D environments.




Figure 2: **Method overview.** The **agent** (left) encodes each RGB frame into a per-frame token by fusing patch embeddings and DINOv2 features via a learnable query. Tokens are processed by causal temporal self-attention and a linear-attention module with a global hidden state, before an actor-critic head emits the next action and value estimate. The **environment** (right) executes the action and returns the next observation; in parallel, a persistent 3DGS *forward model* renders the same view from accumulated past experience, and the discrepancy with the observed view defines the *curiosity reward*. The forward model is used only at training time (blue path); at deployment the agent acts from RGB stream alone.

## 2 Method

Assuming a static 3D scene, we aim to train a camera-agent to efficiently explore its surroundings purely from visual inputs. We wish for novelty-seeking behavior to arise without task-specific rewards or external exploration labels, and so we formulate exploration as a self-supervised RL problem with a curiosity-driven learning signal derived from the agent’s own experience. We defer discussion of **related works** and **background** to Section 4 and Appendix C, and describe our method first.

Our method has two key parts: a long-context transformer-based agent architecture which enables long-horizon exploration behavior, and a curiosity module based on 3D reconstruction which provides a stable reward for the agent to optimize with on-policy RL. The agent (Fig. 2, left) takes in *only* a stream of RGB observations from the 3D environment at hand, and outputs local camera motion per-timestep. It maintains internal memory long enough to enable it to discover exploration strategies like backtracking. The curiosity module (Fig. 2, right) builds a photorealistic 3D scene reconstruction from the agent’s stream of observations, then uses the disagreement between novel view renders and ground truth observation as reward. To optimize this curiosity-based reward the agent necessarily must learn to visually explore 3D scenes and navigate complex environments to seek novel viewpoints.

**Problem setup.** We frame exploration as interaction with an environment. At each time  $1 < t < N$ , the agent takes action  $a_t$  to move through environment  $\mathcal{E}$  and receives new observation  $o_{t+1}$ , transitioning from its current state to the next. While observations can include privileged sensory information at training time (denoted  $\tilde{o}_{t+1}$  when so augmented), the agent relies solely on a visual RGB stream at test time, enabling easy and general deployment. Specifically, the privileged inputs required at training time are the camera pose and depth image.

### 2.1 Persistent 3D forward model

For the agent to learn exploratory behavior, it should be rewarded for seeking novel regions of the environment. We achieve this via intrinsic curiosity by introducing an auxiliary model of the world during training. This *forward model* is tasked with predicting the next observation conditioned on an action, given the stream of observations so far:

$$\hat{o}_{t+1} = \mathcal{F}(a_t, \tilde{o}_{1:t}). \quad (1)$$

The forward model  $\mathcal{F}$  defines a curiosity reward that activates upon encountering novel regions not yet explained by this model, by comparing its prediction to the observed reality:

$$r_t^{\text{cur}} = d(\hat{o}_{t+1}, o_{t+1}), \quad (2)$$

where  $d$  is a visual prediction-error metric. A larger discrepancy indicates the new observation is poorly explained by the current forward model, corresponding to higher novelty.

We instantiate the forward model as an online 3D Gaussian Splatting (3DGS) model of the world, capable of incorporating streaming observations while maintaining a persistent model of what has been seen. We later show in Section 3 that this persistence is crucial for a reliable reward signal. We refer the reader to Appendix C.2 for background on 3DGS. Given the training-time privileged stream

of observations including RGB image, depth image and camera poses (i.e.,  $\tilde{o}_{1:t} = \{(I_i, D_i, p_i)\}_{i=0}^t$ ), we maintain a 3DGS representation  $\mathcal{G}_t$ . For each observed RGB-D frame, a new gaussian per pixel is added to the scene, initialized from its color, depth, and camera pose. At fixed intervals, the representation is optimized on randomly selected past frames to improve reconstruction quality, and pruned and densified following [15].

At each step, after the agent acts and the camera moves following environment dynamics, the forward model is queried at the new pose  $p_{t+1}$  to render a prediction of the incoming observation:

$$\hat{I}_{t+1} = \mathcal{F}(p_{t+1}|I_{1:t}, D_{1:t}, p_{1:t}) = \mathcal{R}(\mathcal{G}_t, p_{t+1}), \quad (3)$$

where  $\mathcal{R}$  denotes the differentiable 3DGS renderer. Since  $\mathcal{G}_t$  incorporates only observations up to time  $t$ , the prediction error primarily reflects whether the incoming view contains information not yet captured by the forward model. To avoid rewarding reconstruction error near high-frequency details in the scene that are not meaningful for exploration, we compute prediction error after low-pass filtering and down-sampling both the prediction and observation:

$$e_t = \frac{1}{|\Omega_s|} \sum_{(j,k) \in \Omega_s} \left\| \mathcal{D}_s(\mathcal{B}(I_{t+1}))^{j,k} - \mathcal{D}_s(\mathcal{B}(\hat{I}_{t+1}))^{j,k} \right\|_2^2, \quad (4)$$

where  $\mathcal{B}$  is a low-pass filter,  $\mathcal{D}_s$  is down-sampling operator by a factor of  $s$ , and  $\Omega_s$  is the down-sampled pixel set. In our implementation, the distance  $d$  in (2) is defined by thresholding the above filtered error to provide a binary reward:

$$r_t^{\text{cur}} = \begin{cases} r_{\text{new}}, & e_t > \tau, \\ r_{\text{old}}, & e_t \leq \tau, \end{cases} \quad (5)$$

where  $r_{\text{new}} > 0$  rewards sufficiently unexplained views, and  $r_{\text{old}} < 0$  gives a small penalty for views already explained by the forward model.

## 2.2 Agent Architecture

Our agent, decoupled from the forward model, is tasked with seeking novelty and therefore must maintain its own record of past observations to take actions that lead to unseen areas. While a geometric map is an effective abstraction for this, it discards all information beyond scene geometry. Instead we let the agent learn its own internal representation of the world, where both geometry and semantics are preserved. To this end, we use a transformer backbone whose actions are conditioned on the full sequence of past visual observations and actions throughout the exploration episode, hence giving the agent the *episodic memory* needed to learn a novelty-seeking policy  $\pi$ :

$$a_t \sim \pi(\cdot | o_{1:t}, a_{1:t-1}). \quad (6)$$

This memory has the further benefit of freeing the agent from being conditioned on any explicit geometric representation that abstracts the observations, making it straightforwardly adaptable to downstream tasks of varied nature that can benefit from exploratory behavior, through fine-tuning.

Concretely, the agent’s episodic memory is represented as the sequence of past RGB observations and actions. This sequence is fed into a transformer backbone connected to the actor and critic heads that output an action distribution and a value estimate for the state at each timestep. The architecture, shown in Fig. 2, consists of an image encoder, causal attention layers over a temporal sliding window, and a global linear-attention memory module. At time  $t$ , the input is  $\{(I_i, a_{i-1})\}_{i=1}^t$ , where each observation  $I_i$  is paired with the preceding action  $a_{i-1}$ . Each action is encoded geometrically as a Plücker-ray image [16] representing the intended camera transformation, independent of actual environment dynamics, and concatenated channel-wise with its corresponding RGB observation to form a unified RGB-action input.

This RGB-action image in each timestep is first compressed into a frame token through an image encoder. We also take the RGB image processed by DINOv2 [17] to provide richer visual features. A learnable query token cross-attends to the patch tokens and DINOv2 features:

$$z_i = \text{CrossAttn}(q, [u_i, f_i]), \quad (7)$$

where  $u_i$  denotes RGB-action patch tokens,  $f_i$  denotes DINOv2 features, and  $q$  is the learnable frame query. The resulting frame tokens are processed by sliding-window causal temporal attention:

$$\bar{z}_{1:t} = \text{SelfAttn}_W(z_{1:t}). \quad (8)$$


Causality ensures that each token uses only past and current observations, while the window size  $W$  keeps computation scalable over long episodes. Sliding-window attention provides efficient direct local context, but long-range information can only propagate indirectly through deeper layers. We therefore interleave selected temporal layers with a global linear-attention memory module inspired by TTT and LoGeR-style long-context architectures [18, 19]. The module reads from a running memory state before updating it with the current token:

$$\tilde{z}_i = \text{LinAttn}(\bar{z}_i, h_{i-1}), \quad h_i = \text{Update}(h_{i-1}, \tilde{z}_i). \quad (9)$$

The memory readout is added back into the temporal stream, giving the policy access to information beyond the sliding attention window without full attention over the episode. Finally, the current token is passed through the actor and critic heads to output action probabilities  $\pi_\theta(a_t | o_{1:t}, a_{1:t-1})$  and value estimate  $V_\theta(o_{1:t}, a_{1:t-1})$ . Action is then sampled from the probability provided by the  $\pi_\theta$  at both train and test times.

### 2.3 Training and regularization

We optimize our actor-critic policy using PPO [20]. During training, rollouts are collected from the current policy, and the intrinsic reward  $r_t^{\text{cur}}$  defined in Sec. 2.1 is used as the only reward signal. A practical difficulty in curiosity-driven exploration is maintaining action diversity after rewards become sparse. Once nearby regions have been explored, reaching new parts of the scene may require temporarily moving through already-seen areas without receiving any intermediate reward, such as backtracking to an unseen branch. Regularizing action entropy with an additional loss term alongside the PPO losses helps, but it is sensitive to its coefficient and the optimization landscape: too small a coefficient can be overwhelmed by other loss terms and allow exploration to collapse (see inset), while too large a coefficient can destabilize optimization. Inspired by SAPG [21], which collects episode rollouts from a mixture of agents, we occasionally sample actions from a uniform random policy during the rollout collection in addition to a small entropy regularizer.



Specifically, at each training step, the executed action is sampled from a mixture of the learned policy and a uniform distribution over the discrete action set  $\mathcal{A}$ :

$$\mu_\theta(a_t | o_{1:t}, a_{1:t-1}) = (1 - \beta)\pi_\theta(a_t | o_{1:t}, a_{1:t-1}) + \beta\mathcal{U}(a_t), \quad (10)$$

where  $\beta$  is the probability of sampling from the uniform policy and  $\mathcal{U}(a_t) = 1/|\mathcal{A}|$ . This guarantees persistent exploratory behavior during rollout collection without requiring the learned policy distribution itself to remain highly stochastic. For the PPO update, we account for the mixed behavior distribution in the likelihood ratio by using the probability of the sampled action under the rollout behavior distribution in the denominator i.e.  $\rho_t(\theta) = \pi_\theta/\mu_{\theta_{\text{old}}}$ . The resulting ratio is then used in the standard PPO clipped surrogate objective as described in Appendix C.1, with the mixing coefficient annealed to zero over training (see Appendix D).

## 3 Experiments

In this section, we perform an experimental analysis of our curiosity-driven exploration. We first show how our agent can effectively explore indoor scenes achieving improved performance over previous RL-based active mapping methods (Section 3.1). To analyze the necessity of the agent’s episodic memory and world persistence, we ablate memory capacity variations across both models in (Section 3.2). We then demonstrate how our agent, equipped with exploratory behavior, can readily be fine-tuned for downstream navigation tasks, such as apple-picking and image-goal navigation (Section 3.3). Finally, we show our agent can generalize to AI-generated out-of-distribution scenes *without* further training (Section 3.4).

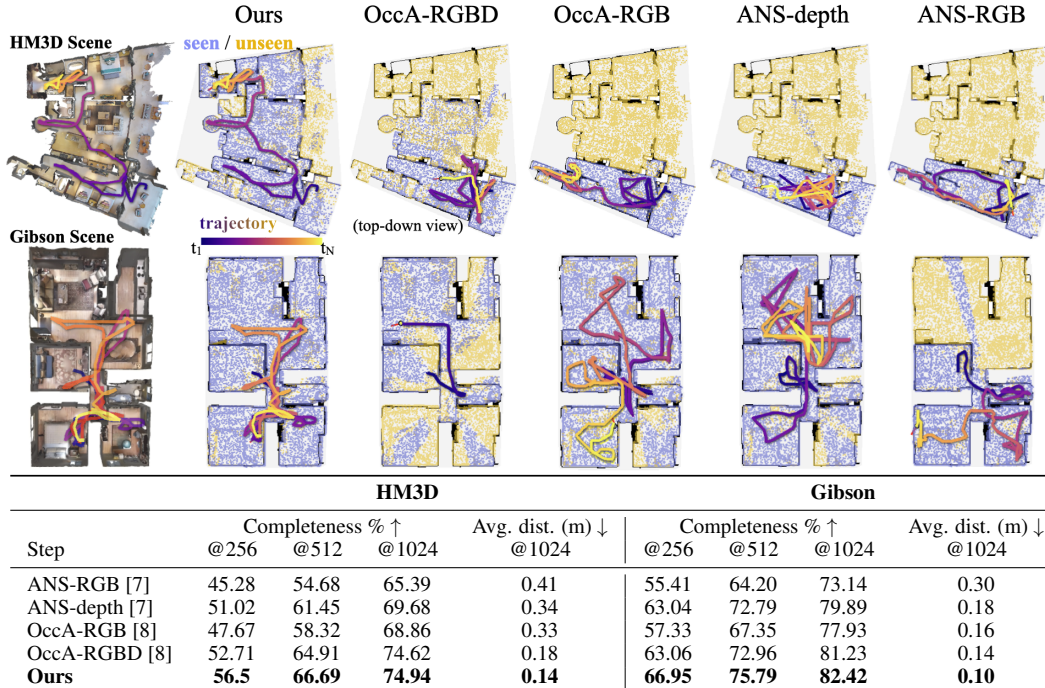


Figure 3: Our agent explores the scenes more thoroughly leading to a higher completeness in 3D scene coverage in HM3D [12] and Gibson [13] datasets. Figure shows agents trajectories and top-down projection of 3D points as **seen** or **unseen** by the agents.

### 3.1 Indoor scene exploration — Figure 3

To assess the capability of our agent, we show test-time exploration of *unseen* indoor scenes.

**Implementation details.** We train our transformer backbone with Adam [22] at a learning rate of  $10^{-5}$  for 110 million steps, with the random policy regularizer scheduled from 20% to zero over 5 million steps beginning at step 25 million, and the action entropy coefficient decayed at a rate of 0.99 from an initial value of 0.1. Reward values are set to  $r_{new} = 0.5$ ,  $r_{old} = 2 \times 10^{-4}$ . The 3DGS primitives are initialized from depth and color at every step, optimized on 10 random past views every 16 steps with Adam [22], and densified via 3DGS-MCMC [15]. Downsampling factor is set to 4. We simulate a drone agent with spherical embodiment spawned 1.25 meters above the floor, and a  $90^\circ$  FOV forward camera in Habitat [23]. It can perform four actions: *move forward* (0.25m), *look right* ( $15^\circ$ ), *look left* ( $15^\circ$ ), and *pause*. Our drone agent is not constrained to the scene navmesh. Instead, it moves freely in 3D, and collisions are computed directly by ray-tracing against the scene mesh and checking contact with the agent’s spherical embodiment. This avoids the navmesh-dependent sliding shortcut learning identified by Monaci et al. [24]. Actions are executed deterministically unless the agent collides, in which case it halts in place. All train- and test-time episodes are run for 1024 steps. Training runs for 5.5 days on a  $8 \times 80$ GB H100 GPUs.

**Datasets.** We train on the HM3D [12] training set (consisting of 800 scenes), with each episode initialized at a random location. We evaluate on the HM3D validation set (100 scenes, 2 random starting points each, 200 episodes total) and, to assess generalization beyond HM3D, on Gibson [13] (86 mostly office-space scenes, 1 starting point each).

**Baselines.** We compare against RL-based exploration methods, as our goal is to evaluate learned policies rather than online planners optimizing per-scene *geometric* objectives. Specifically, we compare against **Active Neural SLAM (ANS)** [7] and **Occupancy Anticipation (OccAnt)** [8], two representative map-driven RL approaches. We exclude GLEAM [9]: despite being more recent, it follows the same map-based hierarchical paradigm and assumes panoramic observations, a different action space, and a test-time collision-unaware A\* local planner, making it incompatible with our forward-facing RGB, collision-aware Habitat setting. The -RGB, -depth, and -RGBD variants denote

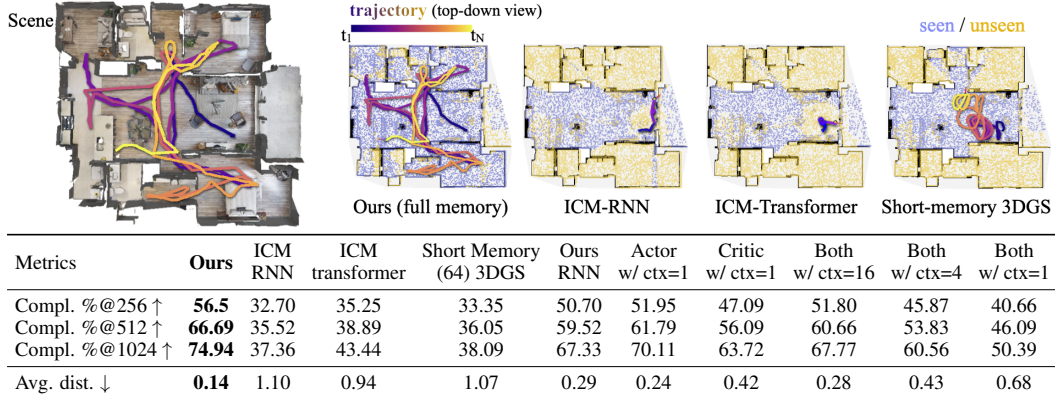


Figure 4: **Ablations** on the memory capacity of the forward model and the policy network on HM3D [12], show that a persistent forward model is necessary for effective exploration. Further, the agents equipped with memory capacity can achieve higher 3D scene coverage with less concentrated local loops in its trajectory.

different input modalities to their respective RL policy module: ANS predicts a top-down map from RGB; ANS-depth projects depth directly; OccAnt-rgb additionally predicts map completions beyond the projection; and OccAnt-rgb does so without depth. All baselines use default hyperparameters from the official OccAnt codebase and are trained for 110M steps on the same dataset and simulator.

**Metrics.** Following Yan et al. [25], we measure exploration via 3D scene completeness at 3 time horizons (256, 512, and 1024 steps): the percentage of points on the reachable ground truth mesh surface (sampled uniformly, 200k points) whose closest observed point, back-projected using ground-truth depth, lies within 5cm. We also report the average distance from each ground truth point to its nearest observed point.

**Analysis.** Our agent achieves greater 3D completeness faster than all baselines, while requiring only RGB input at test time, unlike OccAnt-RGBD and ANS-depth, which require ground-truth depth and are thus more limited in deployment. All baselines additionally rely on raw camera pose from agent sensors, subsequently corrected by a localization module, though this distinction is less critical in our deterministic motion setting. As shown in Figure 3, map-based methods exhibit characteristic failure modes: OccAnt and ANS-RGB can become locally trapped or stuck on the scene geometry due to erroneous map predictions, while ANS-depth over-invests in local coverage at the expense of long-horizon exploration efficiency. Qualitatively, our agent exhibits emergent behaviors such as seeking doorways, traversing corridors, and returning to junctions, while almost never resorting to the pause action during exploration.

### 3.2 Memory Ablations — Figure 4

We study the effect of memory in the forward model and in the agent. We start from the basic building blocks introduced by ICM [5] and gradually add to the memory capacities of each module. We use the same HM3D evaluation dataset and metric setup as in Section 3.1.

**Baselines.** Starting from ICM [5], we ablate the world model and agent memory independently. For the world model, we test ICM with its original inverse and forward dynamics modules, paired with either a capacity-matched RNN (**ICM-RNN**) as in [5], or our transformer (**ICM-Transformer**), isolating the effect of the world model alone. We then replace the forward model with a 3DGS limited to the last 64 frames (**Short Memory 3DGS**) at each time step, ablating short-horizon world memory. For agent memory, we compare a capacity-matched RNN, and transformers with limited context windows (ctx = 1, 4, 16). Finally, we ablate memory asymmetry between actor and critic (Actor w/ctx = 1; Critic w/ ctx = 1).

**Analysis.** Persistent memory consistently benefits exploration across all modules. ICM without persistent world memory catastrophically *collapses* spinning in place or getting stuck on geometry. Adding a persistent 3DGS forward model to RNN makes significant gains, and non-persistent Short Memory 3DGS causes agents to roam locally without venturing far. Agent memory ablations confirm that both RNN and short-context transformers degrade exploration, and notably asymmetric memory, even with only the critic retaining context, outperforms having no memory at all.

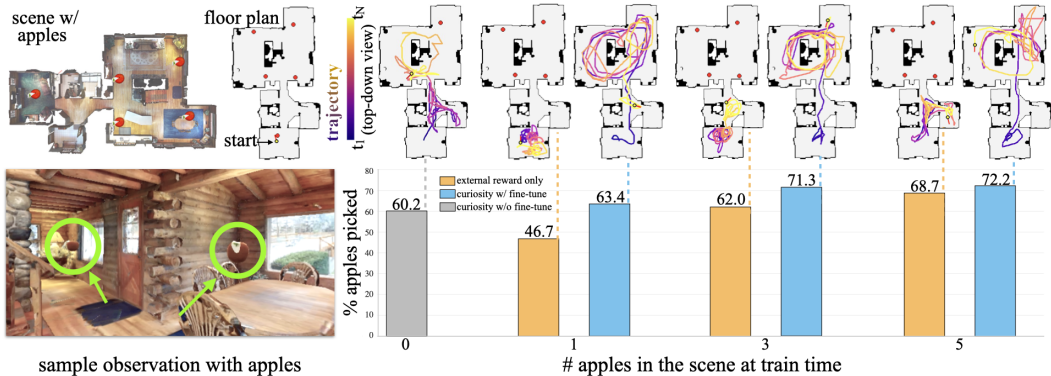


Figure 5: After a few fine-tuning episodes on the apple-picking reward, our exploration agent achieves higher task success than an agent trained from scratch using only the external task reward with the same total number of episodes. This advantage is larger when fewer apples are available per training scene, where rewards are sparser, and more reward-less exploration is needed.

### 3.3 Task-based fine-tuning — Figure 5 and Figure 6 (left)

We show that our map-free agent architecture transfers readily to downstream navigation tasks via fine-tuning, without any architectural modifications, leveraging its exploratory pretraining to achieve stronger performance in sparse-reward settings.

We evaluate on two tasks. **Apple-picking** [26]: the agent must locate and approach apples scattered across the scene; picked apples are removed, and the agent receives a positive reward per pick and a small step penalty otherwise. Apples are placed at reachable positions sampled with probability proportional to their minimum distance to the already-selected set (i.e., avoiding clusters). **Image-goal navigation**: the agent is given a target RGB image and must reach the corresponding viewpoint; success requires  $\geq 50\%$  of the target’s 3D points to be visible and the agent to be within 1.5m of the target pose. The goal image is passed through image encoder and then provided as an additional token to the temporal module with dense attention to all other tokens.

**Experimental setting.** Both tasks use HM3D [12]. Apple-picking uses five apples per scene; image-goal navigation samples target poses reachable by the agent and sufficiently far from scene geometry. We compare against a matched agent trained from scratch on the task reward only for 118M steps; equal to our 110M pretraining plus 8M fine-tuning with external reward. We also report our non-fine-tuned exploration agent as a zero-shot baseline. We report success rate: average percentage of apples picked and percentage of image goals reached across the test set.

**Analysis.** Operating purely on RGB input without explicit geometric representations, our agent is flexible to adapt to both of the downstream semantic tasks through fine-tuning alone. In **apple-picking**, as shown in Fig. 5, the baseline degrades under sparser reward (fewer apples at train-time), its search confined to one room, while our fine-tuned agent scales consistently with apple count. Without fine-tuning, pure exploration hits some apples zero-shot. **Image-goal navigation** is a more extreme sparse-reward regime: the baseline agent collapses to a roaming-around behavior, while our fine-tuned model leverages its exploratory prior to find the target view more reliably; Fig. 6 (left). This is due to the fact that broader exploration increases reward exposure during training, enabling the re-targeting of general exploration to goal-based navigation.

### 3.4 Out-of-distribution scenes — Figure 6 (right)

To assess generalization beyond indoor Habitat scans, we apply our agent zero-shot to two AI-generated worlds from World Labs [14] (Hobbit World and Spaceship). These scenes differ not only in aesthetic but also in rendering pipeline, as they are rendered from a 3DGS representation of these worlds rather than the realistic mesh scans of HM3D. Nevertheless, the agent exhibits coherent exploratory behavior: it navigates corridors, discovers doors to new spaces, and avoids collisions. Specifically, the agent encounters only 2 and 3 collisions over 256 steps in the two worlds, suggesting robust internal representations of navigable space.

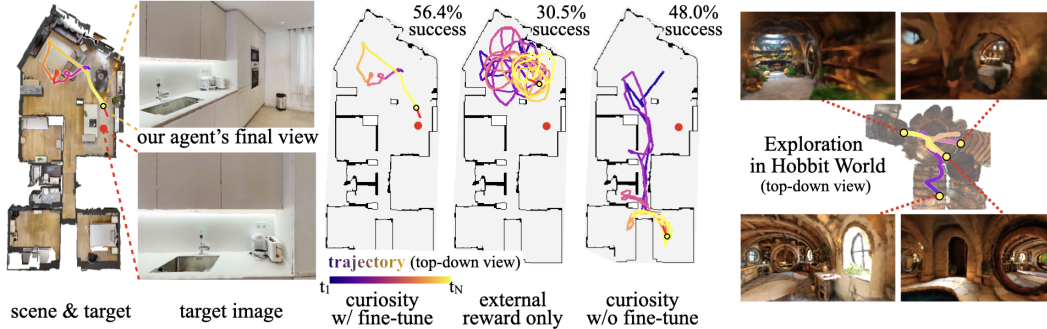


Figure 6: **Left:** Our exploration agent, when fine-tuned for a few episodes on image-goal navigation reward, outperforms an agent trained from scratch on this reward alone, redirecting its general-purpose exploration toward a targeted objective. **Right:** Our agent generalizes zero-shot to AI-generated worlds of different aesthetic and representation.

## 4 Related Work

We survey the most related works here; a more comprehensive discussion is provided in Appendix B.

**Intrinsic Motivation and Curiosity.** Curiosity-driven exploration often rewards prediction errors in a learned forward dynamics model [5, 27, 28] or maximizes episodic visitation proxies [29–32]. However, these approaches struggle with non-stationary reward signals and fundamentally lack a persistent, improving model of the world [33]. Furthermore, standard exploration policies are typically reactive or limited to short recurrent contexts, preventing agents from planning based on accumulated history [10]. We address both limitations by grounding curiosity in a persistent 3D reconstruction paired with a full-history sequence model for long-horizon planning toward novelty.

**Active Mapping and Next-Best-View.** Traditional next-best-view (NBV) methods greedily select viewpoints to maximize geometric information gain [25, 34–36]. Similarly, RL-based mapping methods often condition policies on explicit geometric maps [7–9] or rely on human demonstrations [11, 37]. By treating map-building as the end goal and explicitly conditioning on geometry, these approaches restrict their adaptability to semantically diverse tasks. Our policy, conversely, requires no explicit geometric map or depth sensor at deployment, operating purely from RGB streams to enable flexible downstream task adaptation.

**3D Representations as World Models.** Recent video-generative models offer a promising data-driven approach to modeling world dynamics [38, 6, 39], but currently suffer from severe spatial forgetting and open-loop generation [40]. Alternatively, explicit 3D reconstructions like 3D Gaussian Splatting [41, 42] serve as rudimentary, dynamics-free world models that guarantee exact geometric consistency. Rather than claiming 3DGS uniquely solves this problem, we use it as a controlled proxy demonstrating that spatial persistence and closed-loop online updating are strictly required for reliable curiosity signals, with dynamic scenes as a natural next frontier as generative models mature.

## 5 Conclusions

We introduced *Remember to be Curious*, a framework for end-to-end curiosity-driven exploration of static realistic 3D scenes, where the agent operates solely from a visual input stream. Through using a persistent online 3D reconstruction model as a proxy for the forward model, we demonstrate the necessity of two components for long-horizon exploration: a persistent forward model of the world for a reliable curiosity signal, and an episodic memory architecture for planning toward novelty.

While our use of 3DGS confines the agent to static scenes, we see promise in action-conditioned video models for addressing this limitation. We view our study as a guide for what these future world models require to enable intrinsic motivation. In particular, spatial persistence is key, alongside the ability to continuously update internal representations based on closed-loop online observations. We hope this work opens avenues for curiosity-driven exploration across diverse real-world environments and embodiments, extending to dynamic settings as world modeling matures.

## 6 Acknowledgments

We thank Tyler Bonnen, Antoine Guédon, Vincent Lepetit, Vassia Simaiaki, Jamie Shotton, and Alexei Efros for constructive feedback and helpful discussions. We also thank members of the DGP Lab at the University of Toronto, KAIR Lab at UC Berkeley, and Theia Lab at Simon Fraser University for their feedback, support, and discussions.

Our research is funded in part by NSERC Discovery (RGPIN–2022–04680), the Ontario Early Research Award program, the Canada Research Chairs Program, a Sloan Research Fellowship, the DSI Catalyst Grant program and gifts by Adobe Inc.

## References

- [1] Edward C Tolman. Cognitive maps in rats and men. *Psychological review*, 55(4):189, 1948. 2
- [2] Yuke Zhu, Roozbeh Mottaghi, Eric Kolve, Joseph J. Lim, Abhinav Gupta, Li Fei-Fei, and Ali Farhadi. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In *ICRA*, 2017. 2
- [3] Erik Wijmans, Abhishek Kadian, Ari Morcos, Stefan Lee, Irfan Essa, Devi Parikh, Manolis Savva, and Dhruv Batra. DD-PPO: Learning near-perfect pointgoal navigators from 2.5 billion frames. In *ICLR*, 2020. 2
- [4] Jürgen Schmidhuber. A possibility for implementing curiosity and boredom in model-building neural controllers. In *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*. MIT Press/Bradford Books, 1991. 2, 14
- [5] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *ICML*, 2017. 2, 7, 9, 14
- [6] Robbyant Team. Advancing open-source world models. *arXiv preprint arXiv:2601.20540*, 2026. 2, 9, 15
- [7] Devendra Singh Chaplot, Dhiraj Gandhi, Saurabh Gupta, Abhinav Gupta, and Ruslan Salakhutdinov. Learning to explore using active neural slam. In *ICLR*, 2020. 2, 6, 9, 14
- [8] Santhosh K. Ramakrishnan, Ziad Al-Halah, and Kristen Grauman. Occupancy anticipation for efficient exploration and navigation. In *ECCV*, 2020. 6, 14
- [9] Xiao Chen, Tai Wang, Quanyi Li, Tao Huang, Jiangmiao Pang, and Tianfan Xue. Gleam: Learning generalizable exploration policy for active mapping in complex 3d indoor scenes. In *ICCV*, 2025. 2, 6, 9, 14
- [10] Adrien Ecoffet, Joost Huizinga, Joel Lehman, Kenneth O. Stanley, and Jeff Clune. First return, then explore. In *Nature*, 2021. 2, 9
- [11] Tao Chen, Saurabh Gupta, and Abhinav Gupta. Learning exploration policies for navigation. In *ICLR*, 2019. 2, 9, 14
- [12] Santhosh Kumar Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alexander Clegg, John M Turner, Eric Undersander, Wojciech Galuba, Andrew Westbury, Angel X Chang, Manolis Savva, Yili Zhao, and Dhruv Batra. Habitat-matterport 3d dataset (HM3d): 1000 large-scale 3d environments for embodied AI. In *NeurIPS*, 2021. 2, 6, 7, 8
- [13] Fei Xia, Amir R. Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson Env: real-world perception for embodied agents. In *CVPR*, 2018. 2, 6
- [14] World Labs. Generating worlds. <https://www.worldlabs.ai/blog/generating-worlds> and <https://www.worldlabs.ai/blog/spark-2.0>, 2024. Accessed: 2025. 2, 8
- [15] Shakiba Kheradmand, Daniel Rebain, Gopal Sharma, Weiwei Sun, Yang-Che Tseng, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. 3d gaussian splatting as markov chain monte carlo. In *NeurIPS*, 2024. 4, 6, 15

- [16] Vincent Sitzmann, Semon Rezchikov, William T. Freeman, Joshua B. Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. In *NeurIPS*, 2021. 4
- [17] Maxime Oquab and Timothée Darcet et al. Dinov2: Learning robust visual features without supervision. In *TMLR*, 2023. 4
- [18] Yu Sun and Xinhao Li et al. Learning to (learn at test time): Rnns with expressive hidden states. In *ICML*, 2025. 5
- [19] Junyi Zhang, Charles Herrmann, Junhwa Hur, Chen Sun, Ming-Hsuan Yang, Forrester Cole, Trevor Darrell, and Deqing Sun. Loger: Long-context geometric reconstruction with hybrid memory. *arXiv preprint arXiv:2603.03269*, 2026. 5, 15
- [20] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 5, 15
- [21] Jayesh Singla, Ananye Agarwal, and Deepak Pathak. Sapg: Split and aggregate policy gradients. In *ICML*, 2024. 5
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [23] Xavi Puig and Eric Undersander et al. Habitat 3.0: A co-habitat for humans, avatars and robots, 2024. 6
- [24] Gianluca Monaci, Philippe Weinzaepfel, and Christian Wolf. What does really matter in image goal navigation? In *3DV*, 2026. 6
- [25] Zike Yan, Haoxiang Yang, and Hongbin Zha. Active neural mapping. In *ICCV*, 2023. 7, 9, 14
- [26] Piotr Mirowski, Razvan Pascanu, Fabio Viola, Hubert Soyer, Andrew J Ballard, Andrea Banino, Misha Denil, Ross Goroshin, Laurent Sifre, Koray Kavukcuoglu, et al. Learning to navigate in complex environments. *arXiv preprint arXiv:1611.03673*, 2016. 8
- [27] Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random network distillation. In *ICLR*, 2019. 9, 14
- [28] Deepak Pathak, Dhiraj Gandhi, and Abhinav Gupta. Self-supervised exploration via disagreement. In *ICML*, 2019. 9, 14
- [29] Nikolay Savinov, Anton Raichuk, Raphaël Marinier, Damien Vincent, Marc Pollefeys, Timothy Lillicrap, and Sylvain Gelly. Episodic curiosity through reachability. In *ICLR*, 2019. 9, 14
- [30] Adrià Puigdomènech Badia and Pablo Sprechmann et al. Never give up: Learning directed exploration strategies. In *ICLR*, 2020. 14
- [31] Mikael Henaff, Roberta Raileanu, Mingqi Jiang, and Tim Rocktäschel. Exploration via elliptical episodic bonuses. In *NeurIPS*, 2022. 14
- [32] Yao Fu, Run Peng, and Honglak Lee. Go beyond imagination: maximizing episodic reachability with world models. In *ICML*, 2023. 9, 14
- [33] Roger Creus Castanyer, Joshua Romoff, and Glen Berseth. Improving intrinsic exploration by creating stationary objectives. In *ICLR*, 2024. 9, 14
- [34] Wen Jiang, Boshu Lei, and Kostas Daniilidis. Fisherrf: Active view selection and uncertainty quantification for radiance fields using fisher information. In *ECCV*, 2024. 9, 14
- [35] Shiyao Li, Antoine Guédon, Shizhe Chen, and Vincent Lepetit. MAGICIAN: Efficient Long-Term Planning with Imagined Gaussians for Active Mapping. In *CVPR*, 2026. 14
- [36] Wen Jiang, Boshu Lei, Katrina Ashton, and Kostas Daniilidis. Multimodal llm guided exploration and active mapping using fisher information. In *ICCV*, 2025. 9, 14

- [37] Yuhang He, Irving Fang, Yiming Li, Rushi Bhavesh Shah, and Chen Feng. Metric-Free Exploration for Topological Mapping by Task and Motion Imitation in Feature Space. In *RSS*, 2023. 9, 14
- [38] Philip J. Ball and Jakob Bauer et al. Genie 3: A new frontier for world models, 2025. URL <https://deepmind.google/blog/genie-3-a-new-frontier-for-world-models/>. 9, 15
- [39] Seonghyeon Ye and Yunhao Ge et al. Dreamzero: World action models are zero-shot policies. *arXiv preprint arXiv:2602.15922*, 2026. 9, 15
- [40] Ziqi Ma, Mengzhan Liufu, and Georgia Gkioxari. Out of sight, out of mind? evaluating state evolution in video world models. *arXiv preprint arXiv:2603.13215*, 2026. 9, 15
- [41] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. In *SIGGRAPH*, 2023. 9, 15
- [42] Nikhil Keetha, Jay Karhade, Krishna Murthy Jatavallabhula, Gengshan Yang, Sebastian Scherer, Deva Ramanan, and Jonathon Luiten. Splatam: Splat, track & map 3d gaussians for dense rgb-d slam. In *CVPR*, 2024. 9, 15
- [43] Younggyo Seo, Lili Chen, Jinwoo Shin, Honglak Lee, Pieter Abbeel, and Kimin Lee. State entropy maximization with random encoders for efficient exploration. In *ICML*, 2021. 14
- [44] Bhavya Sukhija, Stelian Coros, Andreas Krause, Pieter Abbeel, and Carmelo Sferrazza. Maxinfor!: Boosting exploration in reinforcement learning through information gain maximization. In *ICLR*, 2025.
- [45] Andres Campero, Roberta Raileanu, Heinrich Küttler, Joshua B. Tenenbaum, Tim Rocktäschel, and Edward Grefenstette. Learning with amigo: Adversarially motivated intrinsic goals. In *ICLR*, 2021. 14
- [46] Yuri Burda, Harri Edwards, Deepak Pathak, Amos Storkey, Trevor Darrell, and Alexei A. Efros. Large-scale study of curiosity-driven learning. In *ICLR*, 2019. 14
- [47] Roberta Raileanu and Tim Rocktäschel. Ride: Rewarding impact-driven exploration for procedurally-generated environments. In *ICLR*, 2020. 14
- [48] Roberto Bigazzi, Federico Landi, Silvia Cascianelli, Lorenzo Baraldi, Marcella Cornia, and Rita Cucchiara. Focus on impact: Indoor exploration with intrinsic motivation. In *RA-L*, 2022. 14
- [49] Lily Goli, Cody Reading, Silvia Sellán, Alec Jacobson, and Andrea Tagliasacchi. Bayes’ Rays: Uncertainty quantification in neural radiance fields. In *CVPR*, 2024. 14
- [50] Georgios Georgakis, Bernadette Bucher, Anton Arapin, Karl Schmeckpeper, Nikolai Matni, and Kostas Daniilidis. Uncertainty-driven planner for exploration and navigation. In *ICRA*, 2022.
- [51] Antoine Guédon, Tom Monnier, Pascal Monasse, and Vincent Lepetit. Macarons: Mapping and coverage anticipation with rgb online self-supervision. In *CVPR*, 2023.
- [52] Shiyao Li, Antoine Guedon, Clémentin Boittiaux, Shizhe Chen, and Vincent Lepetit. Nextbest-path: Efficient 3d mapping of unseen environments. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu, editors, *ICLR*, 2025.
- [53] Yuetao Li, Zijia Kuang, Ting Li, Qun Hao, Zike Yan, Guyue Zhou, and Shaohui Zhang. Activesplat: High-fidelity scene reconstruction through active gaussian splatting. In *RA-L*, 2025.
- [54] Liyan Chen, Huangying Zhan, Kevin Chen, Xiangyu Xu, Qingan Yan, Changjiang Cai, and Yi Xu. Activegamer: Active gaussian mapping through efficient rendering. In *CVPR*, 2025.
- [55] Yuezhan Tao, Dexter Ong, Varun Murali, Igor Spasojevic, Pratik Chaudhari, and Vijay Kumar. Rt-guide: Real-time gaussian splatting for information-driven exploration. In *RA-L*, 2025.

- [56] Cherie Ho, Seungchan Kim, Brady Moon, Aditya Parandekar, Narek Harutyunyan, Chen Wang, Katia Sycara, Graeme Best, and Sebastian Scherer. Mapex: Indoor structure exploration with probabilistic information gain from global map predictions. In *ICRA*, 2025.
- [57] Seungjae Baek, Brady Moon, Seungchan Kim, Muqing Cao, Cherie Ho, Sebastian Scherer, and Jeong Hwan Jeon. Pipe planner: Pathwise information gain with map predictions for indoor robot exploration. In *IROS*, 2025. 14
- [58] Sen Wang, Bangwei Liu, Zhenkun Gao, Lizhuang Ma, Xuhong Wang, Yuan Xie, and Xin Tan. Explore with long-term memory: A benchmark and multimodal llm-based reinforcement learning framework for embodied exploration. In *CVPR*, 2026. 14
- [59] Max Lodel, Bruno Brito, Alvaro Serra-Gómez, Laura Ferranti, Robert Babuška, and Javier Alonso-Mora. Where to look next: Learning viewpoint recommendations for informative trajectory planning. In *ICRA*, 2022. 14
- [60] Xiao Chen, Quanyi Li, Tai Wang, Tianfan Xue, and Jiangmiao Pang. Gennbv: Generalizable next-best-view policy for active 3d reconstruction. In *CVPR*, 2024. 14
- [61] Andreas Meuleman, Ishaan Shah, Alexandre Lanvin, Bernhard Kerbl, and George Drettakis. On-the-fly reconstruction for large-scale novel view synthesis from unposed images. In *TOG*, 2025. 15
- [62] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2024. 15
- [63] Qianqian Wang, Yifei Zhang, Aleksander Holynski, Alexei A. Efros, and Angjoo Kanazawa. Continuous 3d perception model with persistent state. In *CVPR*, 2025. 15
- [64] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *CVPR*, 2025. 15
- [65] Tianchang Shen and Sherwin Bahmani et al. Lyra 2.0: Explorable generative 3d worlds. *arXiv preprint arXiv:2604.13036*, 2026. 15
- [66] Siyuan Zhou, Yilun Du, Yuncong Yang, Lei Han, Peihao Chen, Dit-Yan Yeung, and Chuang Gan. Learning 3d persistent embodied world models. In *NeurIPS*, 2025.
- [67] 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*, 2025. 15
- [68] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. In *ICLR*, 2016. 15

## A Video Results

Please refer to our website, [recuriosity.github.io](https://recuriosity.github.io), for video results of our agent exploring diverse 3D environments and performing downstream tasks.

## B Related Work Discussion

**Reinforcement learning with intrinsic motivation.** In reinforcement learning (RL), agents learn by interacting with the environment and improving based on the rewards they receive. However, many real-world environments provide sparse or no extrinsic reward, motivating intrinsic-reward methods that encourage exploration through auxiliary objectives such as maximizing state coverage, information gain, or self-generated goals [43–45].

Curiosity is a form of intrinsic reward, first framed by Schmidhuber [4] as reward for prediction error in a learned world model. Pathak et al. [5] proposed a deep RL curiosity framework in which a latent forward dynamics model is trained to predict consequences of actions and the agent is rewarded to seek states with large prediction error. Follow-up works stabilized this curiosity reward further with forward dynamics ensemble disagreement [28] or fixed random network representations [27]. However, these rewards lack episodic persistence: the forward dynamics model only updates through gradient steps across training, so the agent can repeatedly revisit surprising states within an episode to accumulate rewards without having them updated [30], fixating on a narrow set of states while the forward model receives little diverse data to improve on. Further, the reward target becomes non-stationary as the learned latent representation of states evolve, making the signal noisy and optimization unstable [46, 33].

To address this, follow-up works augment curiosity with episodic visitation counts [47, 30, 32], but count-based methods do not scale to continuous, high-dimensional observations [48, 29], and efforts such as E3B [31] that use compact learned embeddings inevitably abstract away geometric and semantic structure. Whether grounded in forward dynamics or visitation counts, no method provides a persistent, improving model of the world for a stable curiosity signal in an exploration episode.

Moreover, the exploration policy which is architecturally decoupled from the forward dynamics model [5] is either solely reactive to the current observation or limited to a short recurrent context, preventing the agent from planning exploration based on its accumulated history. While our proposed 3D modeling of the world sacrifices a lifelong learned world model, we use it as a controlled proxy to study the role of persistence in curiosity-driven exploration.

**Active mapping and next-best-view selection.** While curiosity-driven methods learn a policy that seeks novelty as a behavioral objective, active mapping and next-best-view (NBV) methods treat exploration as a means to an end: building a map of the scene. Classical and learned NBV approaches [25, 34, 49–57] greedily select viewpoints by estimated information gain or coverage; MAGICIAN [35] extends this to a more comprehensive tree search over information gain maximizing trajectories, and recent works combine LLMs with geometric planning for even longer-horizon goal selection and planning [36, 58]. These methods do not learn from experience, are scene-specific and replan from scratch per scene, making them difficult to extend to downstream tasks beyond mapping.

On the other hand, a line of RL-based methods learns a mapping policy conditioned on an explicit geometric map, giving the agent spatial but not semantic memory. Chen et al. [11] and [59] learn exploration policies conditioned on egocentric depth-derived maps via imitation learning and coverage fine-tuning; ANS [7] and OccAnt [8] build hierarchical policies on SLAM-generated maps where a global RL module selects coverage-rewarded goals and a local planner executes them; and GLEAM [9] follows the same paradigm at larger scale with panoramic observations. By conditioning on a geometric map and optimizing for its coverage, these policies acquire no semantic understanding of the world and therefore can only be extended to tasks that can be defined geometrically, such as point-goal navigation. DeepExplorer [37] forgoes the map via a short recurrent context but still relies on expert demonstrations, and GenNBV [60] learns an object-centric NBV policy from occupancy and image features, targeting active object-level reconstruction rather than general exploration.

Our curiosity policy is conditioned on RGB observations alone, requiring no geometric map, depth sensor, or localization at test time, and is flexible to finetuning for downstream tasks such as image-goal navigation where exploratory behavior is a prerequisite.

**3D scene representations and world models.** 3D Gaussian Splatting [41] reconstructs static scenes as explicit Gaussian primitives with fast differentiable rasterization, with follow-up work improving optimization speed and quality [15, 61]. Online SLAM methods extend this to real-time incremental reconstruction from an image stream [42]. Feed-forward methods [62, 63, 19, 64] enable fast dense reconstruction from sequential frames without per-scene optimization, though most output point clouds rather than renderable representations. These reconstruction methods can be viewed as a rudimentary form of a world model: given a camera action, they synthesize the resulting observation, approximating next-state prediction in a geometrically grounded but dynamics-free manner.

Video generative world models [38, 6, 39] offer broader capacity by learning action-conditioned dynamics from data, showing promise for modeling non-stationary worlds for interactive tasks. However, current approaches suffer from spatial forgetting and temporal drift over long horizons [40], and approaches that address this with persistent 3D memory [65–67] mostly revert to a static world assumption to maintain consistency. Further, most generative methods condition on previously generated video chunks rather than real observations, leaving planning open-loop.

We show that curiosity-based exploration requires a forward model grounded in real, closed-loop observations with persistent memory. Rather than claiming this is uniquely solved by 3DGS, we use it as a controlled instantiation for static indoor scenes to study what properties a forward dynamics model must have to support reliable curiosity, opening up avenues for curiosity-based exploration in dynamic scenes as generative world models mature.

## C Background

### C.1 Proximal Policy Optimization

Proximal Policy Optimization (PPO) [20] is an on-policy actor-critic algorithm that optimizes a policy using trajectories sampled from the current policy. Given a policy  $\pi_\theta(a_t | s_t)$  over actions  $a_t$  conditioned on state  $s_t$ , a value function  $V_\theta(s_t)$ , and advantage estimates  $A_t$ , PPO updates the policy by maximizing the clipped surrogate objective

$$\mathcal{L}_{\text{PPO}}(\theta) = \mathbb{E}_t [\min(\rho_t(\theta)A_t, \text{clip}(\rho_t(\theta), 1 - \epsilon, 1 + \epsilon)A_t)], \quad \rho_t(\theta) = \frac{\pi_\theta(a_t | s_t)}{\pi_{\theta_{\text{old}}}(a_t | s_t)}. \quad (11)$$

Here,  $\rho_t(\theta)$  is the policy probability ratio between the updated policy and the rollout policy, and  $\epsilon$  is the clipping threshold. The clipping term limits the size of policy updates, improving training stability. The critic is trained to predict a discounted return target  $R_t$ , formed from accumulated future rewards and value bootstrapping [68]. Entropy regularization further encourages action diversity, where  $\mathcal{H}(\pi_\theta(\cdot | s_t))$  denotes the entropy of the policy distribution at state  $s_t$ .

$$\mathcal{L}_{\text{critic}}(\theta) = \mathbb{E}_t [(V_\theta(s_t) - R_t)^2], \quad \mathcal{L}_{\text{entropy}}(\theta) = -\mathbb{E}_t [\mathcal{H}(\pi_\theta(\cdot | s_t))]. \quad (12)$$

### C.2 3D Gaussian Splatting

3D Gaussian Splatting (3DGS) [41] represents a scene as a set of anisotropic 3D Gaussian primitives. Each Gaussian has a 3D mean  $\mathbf{x}_l$ , covariance  $\Sigma_l$ , opacity  $\alpha_l$ , and view-dependent color  $\mathbf{c}_l$ . During rendering, the Gaussians are projected into the image plane and alpha-composited in depth order to produce the pixel color:

$$I(\mathbf{u}) = \sum_{l=1}^L T_l(\mathbf{u}) \alpha_l g_l^{2D}(\mathbf{u}) \mathbf{c}_l, \quad T_l(\mathbf{u}) = \prod_{j<l} (1 - \alpha_j g_j^{2D}(\mathbf{u})), \quad (13)$$

where  $\mathbf{u}$  is an image-plane pixel location,  $g_l^{2D}$  is the projected 2D Gaussian footprint, and  $T_l$  is the accumulated transmittance from closer Gaussians. The Gaussian parameters are differentiable with respect to an image reconstruction loss, enabling efficient optimization of an explicit 3D scene representation. In RGB-D settings, depth can be used to initialize Gaussian means directly by unprojecting observed pixels into 3D, reducing geometric ambiguity and making online optimization faster and more stable, as used in RGB-D Gaussian SLAM systems such as SplaTAM [42].

## D Scheduled regularization annealing

Since mixing the learned policy with a uniform random policy shifts the action distribution seen by the critic, it can bias its value estimates. Therefore, we gradually anneal  $\beta$  from 0.2 to zero over a scheduled window during training. This provides additional exploration when the policy is weak and curiosity rewards become sparse locally, while reducing the bias introduced by random actions as the learned policy improves. At test time, we set  $\beta = 0$  and act directly from the learned policy.