

RoboMemory: A Brain-inspired Multi-memory Agentic Framework for Lifelong Learning in Physical Embodied Systems

Mingcong Lei^{1,3*}, Honghao Cai^{1*}, Zezhou Cui^{1,3}, Liangchen Tan⁴,
Junkun Hong¹, Gehan Hu^{1,3}, Shuangyu Zhu^{1,5}, Yimou Wu³, Shaohan Jiang^{1,3}, Ge Wang^{1,3},
Zhen Li^{1,2}, Shuguang Cui^{1,2}, Yiming Zhao^{1,6,7}, Yatong Han^{1,7}

¹ FNii-Shenzhen ² SSE ³ The Chinese University of Hong Kong, Shenzhen

⁴ The University of Hong Kong ⁵ Harbin Institute of Technology, Shenzhen

⁶ Harbin Engineering University ⁷ Infused Synapse AI

yiming_zhao@hrbeu.edu.cn hanyatong@cuhk.edu.cn

Project website: <https://sp4595.github.io/robomemory/>

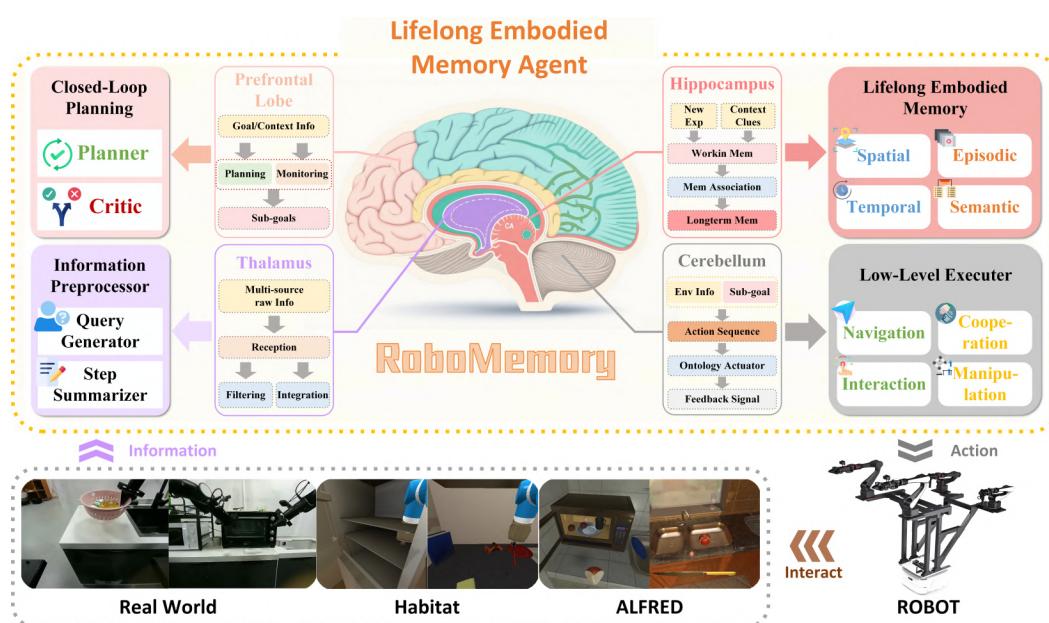


Figure 1: The brain-inspired architecture of the RoboMemory resembles the biological nervous system. It maps biological neural components, enabling the agent to interact with diverse environments (real-world, Habitat, ALFRED) and robotic hardware for long-term planning and lifelong learning.

Abstract

We present RoboMemory, a brain-inspired multi-memory framework for lifelong learning in physical embodied systems, addressing critical challenges in real-world environments: continuous learning, multi-module memory latency, task correlation capture, and infinite-loop mitigation in closed-loop planning. Grounded in cognitive neuroscience, it integrates four core modules: the Information Preprocessor (thalamus-like), the Lifelong Embodied Memory System

*Equal contribution

(hippocampus-like), the Closed-Loop Planning Module (prefrontal lobe-like), and the Low-Level Executer (cerebellum-like) to enable long-term planning and cumulative learning. The Lifelong Embodied Memory System, central to the framework, alleviates inference speed issues in complex memory frameworks via parallelized updates/retrieval across Spatial, Temporal, Episodic, and Semantic submodules. It incorporates a dynamic Knowledge Graph (KG) and consistent architectural design to enhance memory consistency and scalability. Evaluations on EmbodiedBench show RoboMemory outperforms the open-source baseline (Qwen2.5-VL-72B-Ins) by 25% in average success rate and surpasses the closed-source State-of-the-Art (SOTA) (Gemini-1.5-Pro) by 3%, establishing new SOTA. Ablation studies validate key components (critic, spatial memory, long-term memory), while real-world deployment confirms its lifelong learning capability with significantly improved success rates across repeated tasks. RoboMemory alleviates high latency challenges with scalability, serving as a foundational reference for integrating multi-modal memory systems in physical robots.

1 Introduction

Driven by the rapid advancements in Vision-Language Models (VLMs) [19, 4], VLM-based agents have been increasingly deployed in embodied tasks, leveraging pre-trained knowledge and multi-modal understanding capabilities to interact with physical environments [28, 18].

Current research primarily focuses on optimizing performance in single tasks within virtual simulators or controlled real-world setups [5, 2, 40], or relies on oversimplified memory frameworks. These frameworks lack mechanisms to model interdependencies between distinct tasks. In practice, real-world embodied agents must operate over a lifetime, handling sequential tasks where prior experiences (e.g., learning to “open a fridge” to retrieve an apple) directly shape subsequent performance (e.g., efficiently accessing other objects in similar containers). This demands long-term memory systems that not only retain experiences but also capture cross-task influences, enabling cumulative improvement rather than isolated task execution.

Prior efforts to integrate memory systems into embodied frameworks [33, 15, 36, 1, 13] have advanced long-term planning and lifelong learning. However, most of them concentrate on virtual environments and fail to generalize to the real world. While the embodied frameworks that are designed for the real world [42, 29] lack critical memory modules. But on the other hand, a complex memory architecture will suffer from prohibitive latency, which makes it hard to use.

To address these gaps, we propose **RoboMemory**, a brain-inspired multi-memory framework explicitly designed for lifelong learning in the real world. This highly parallelized, hierarchical architecture enables long-term planning and lifelong learning in real-world environments. Drawing on cognitive neuroscience [25], it comprises four core components (Fig. 1): an Information Preprocessor (thalamus-like) for multimodal integration; a Lifelong Embodied Memory System (hippocampus-like), the core part that adopts a three-tier structure (long-term, short-term, working memory) grounded in cognitive neuroscience to organize experiential data and object spatial relationships, with a parallelizable memory paradigm unifying information updating and retrieval across modules to mitigate latency; a Closed-Loop Planning Module (prefrontal lobe-like) for high-level action sequencing; and a Low-level Executer (cerebellum-like). For robust real-world deployment, RoboMemory employs a dual-system architecture [42, 29] where the upper layer (embodied agent) outputs abstract high-level actions, and a Vision-Language-Action (VLA) model with Simultaneous Localization and Mapping (SLAM) system translates these into low-level commands executable by robots.

We validated RoboMemory (excluding the executor) in EmbodiedBench, a long-horizon planning benchmark environment [38]. Results demonstrate substantial improvements: with Qwen2.5-VL-72b as the backbone, RoboMemory increases average success rates by 25% over the base model and 5% over the closed-source state-of-the-art (SOTA) Gemini-1.5-Pro. Moreover, we test RoboMemory in a real-world environment, where RoboMemory executed 15 diverse tasks twice (one for learning and one for testing). By validating the RoboMemory’s performance between two executions, we verify that the RoboMemory has basic lifelong learning ability in the real world. In addition, we

do ablation studies to quantify component-level contributions, and a comprehensive error analysis contextualizes limitations. In conclusion, our contributions fall into three aspects:

- Inspired by the brain’s unified memory mechanisms, we design a lifelong embodied memory system with four parallel modules (Spatial, Temporal, Episodic, and Semantic) under a unified framework. This framework supports parallelized update and retrieval across modules, mitigating latency accumulation in complex systems while facilitating coherent knowledge integration for lifelong learning.
- We propose a retrieval-based incremental Knowledge Graph (KG) update algorithm for dynamic spatial memory evolution. It enables efficient, consistent updates by retrieving relevant subgraphs, detecting local conflicts, and merging new information, alleviating scalability bottlenecks of traditional incremental strategies in dynamic environments.
- RoboMemory enables lifelong learning in real-world physical robots: it executes sequential diverse tasks without memory reset, with experience accumulation driving steady performance improvements. This demonstrates practical long-term autonomous learning in physical scenarios, with reduced reliance on simulated pre-training.

2 Related Work

2.1 VLM/LLM-based Agentic Frameworks in Embodied Tasks

The rapid advancement of VLMs/LLMs has led to diverse agent frameworks in embodied environments [41, 31, 22]. Embodied tasks involve partial observability and long-horizon planning, requiring memory systems to retain context. Some use time-ordered context buffers for short-term memory (due to VLMs/LLMs’ limited long-context processing) [41, 27]; others adopt experience buffers as long-term semantic memory [13, 30]. For long-duration tasks, skill libraries serve as procedural memory, with agents accumulating skills via interaction [36, 33]. Recent efforts integrate diverse memories [43, 33, 1] but focus on virtual/GUI environments, leaving real-world multi-modal memory support for long-term planning under-explored.

2.2 Vision Language Action Model

Current work on VLA models uses imitation learning to output low-level controls from language and visuals [7, 45, 6, 20] but is limited to tabletop tasks and single actions, restricting long-horizon planning. VLAs lack long-term execution abilities, while high-level agents excel at planning. Recent works combine high-level frameworks with VLA executors, some augmented with simple memory [29, 32, 42, 39] for longer tasks. However, real-world robots need more sophisticated memory to handle continuous multi-task operations over extended periods.

2.3 Memory Frameworks

Many previous works improve long-term planning via memory systems: Voyager [36] uses a skill library in Minecraft but lacks diverse memory types; CoELA [43] includes procedural, semantic, and episodic memory with a task-specific 2D map; MSI-Agent [13] utilizes insight as long-term memory for in-task learning. Hippo Retrieval Augmented Generation (RAG) [16] mimics the hippocampus and introduces KGs as long-term memory indices [8, 9], enhancing retrieval. However, the previous approach is mainly focused on constructing a KG with static long context, such as a book, but it is hard to update the graph. We need to update the information in KG for the embodied task. Our approach builds a more general LLM-based memory system using a dynamic KG like Hippo RAG, which is designed for embodied tasks.

3 RoboMemory

As illustrated in Fig. 2, the key design of RoboMemory is a unified memory paradigm. This paradigm aims to streamline memory operations: during updates, only memory items relevant to new information are targeted, while retrieval leverages rule-based information gathering to enhance efficiency. In the following sections, we detail the specific design of each component and how they integrate within this framework to enable real-world lifelong learning.

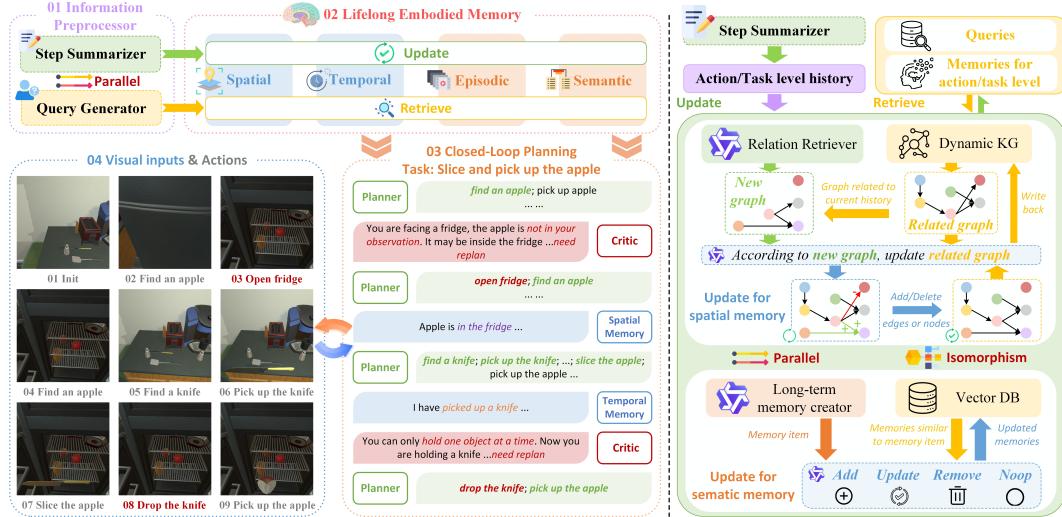


Figure 2: RoboMemory architecture with working pipeline and memory mechanisms. (a) Left: The agent’s pipeline. Parallel Step Summarizer and Query Generator in Information Processor (1) generate updates/queries for Lifelong Embodied Memory (2). These memories enable Closed-Loop Planning (3) for tasks like “slice and pick up the apple”—the Planner generates plans, while the Critic and memories adjust decisions via feedback from visual inputs/results (4). (b) Right: Spatial and Semantic memories operate in parallel with isomorphic updates. Internally, Spatial memory maintains a relevance/similarity-updated KG, and Semantic memory manages a Vector DB with analogous logic.

3.1 Information Preprocessor

For each time-step i , RoboMemory receives a visual observation \mathcal{O}_i —a single RGB frame in simulation or a short video snippet from the onboard camera on the real robot—capturing what the agent sees while acting. An Information Preprocessor constitutes the system’s perceptual front-end, converting this multimodal input into text that can be indexed and searched. To keep latency low, the preprocessor runs two lightweight modules in parallel:

1. *Step summarizer \mathcal{S}* : Transforms \mathcal{O}_i into a concise textual description s_i of the just-executed action. The string s_i is stored as the system’s working memory.
2. *Query generator \mathcal{Q}* : Derives a query q_i from the same observation \mathcal{O}_i , which is used to probe long-term memory for relevant episodes.

Together, \mathcal{S} and \mathcal{Q} provide a swift, text-based interface between raw sensory data and RoboMemory’s retrieval machinery.

3.2 Lifelong Embodied Memory System

RoboMemory incorporates a lifelong memory system with four modules: Spatial, Temporal, Episodic, and Semantic. The four-module structure is designed to enable continuous learning in dynamic real-world environments. To mitigate latency from repeated VLM invocations, we unify the four modules’ update and retrieval processes into a single paradigm with parallel implementation. As shown in Fig. 2, all four modules perform updates and retrievals in parallel. Thus, the framework avoids latency escalation despite multiple memory components.

Memory modules in Lifelong Embodied Memory System have three types of update frequency:

1. Action-level, which is updated once per action;
2. Task-level, which is updated only after the completion of each task;
3. Mixed-level, which is updated at both the action level and the task level.

3.2.1 Spatial-Temporal Memory System

To adapt to dynamic real-world environments, we design the Spatial-Temporal Memory. The spatial and temporal memory update in action-level frequency. The Temporal memory is a first-in-first-out (FIFO) buffer whose size is N . The buffer can store N summary steps $s_{[i:i+N]}$. When the buffer is full, it will be fully cleared. Then, we use an LLM to summarize the short-term memories in the discarded N step into a single entity, inserted as the first entry of the buffer.

The Spatial Memory is a dynamically updated KG-based module that addresses LLMs' limitations in implicitly extracting spatial information from temporal memory. This system dynamically records spatial relationships in the scene while efficiently maintaining a dynamic KG. Unlike previous KG construction algorithms [16, 12], which are designed for static information gathering. Our spatial memory needs to be updated frequently. To speed up the KG update process, we design a two-phase progressive approach: (1) *Rapid Response Phase*: New information is quickly buffered to avoid losing observations. (2) *Local Integration Phase*: On buffer saturation or conflict detection, affected local subgraphs are integrated (entity de-duplication, relationship merging, conflict detection).

This algorithm limits KG updates to the environment. For each update, we only consider relevant segments, controlling efficiency and update context length for better performance. The algorithm is shown in Algorithm 1. Meanwhile, to illustrate the growth of dynamic spatial memory, we provide a sample in Appendix E.

Our retrieval-based incremental KG update algorithm has provable efficiency guarantees. Specifically, for a KG with n nodes and maximum degree D , the number of nodes processed in each update is bounded by $O(D^K)$ where K is the retrieval hop distance (see Appendix D for formal proof). This ensures scalability even as the spatial memory grows over time.

3.2.2 Lifelong Learning System

In real-world scenarios, agents process sequential tasks over their lifetime, requiring continuous improvement through prior experiences. Drawing on cognitive psychology's classification of human long-term memory, we divide our system into episodic and semantic memory: the former records agent-environment interaction histories, while the latter extracts experiential insights to support long-term task reasoning. This update process mirrors human daily experience consolidation during sleep [23].

Episodic Memory: It captures task-level interactions, accounting for temporal interdependencies between sequential tasks in the same environment. The agent must memorize what it has done before to complete future tasks. Moreover, task-level interactions can be a reference that may help the agent improve its plan in the future.

Semantic Memory: It accumulates step-by-step action usage experiences (based on invoked actions and outcomes) to inform action arrangement. Post-task, it summarizes temporal memory, distilling successes from completed tasks and identifying failure causes/improvement strategies from unsuccessful ones, thus enabling both action-level and task-level learning across both task and action levels.

In implementation, both episodic and semantic memory share the same RAG framework consisting of an extractor, updater, and RAG storage (each entry is a memory entity). Post-task, the Extractor summarizes the task's Spatial-Temporal Memory into a new memory entity. The RAG then retrieves similar existing entities (old information) from the RAG. Then, the Updater deletes, adds, or updates old memory entities according to new information. After that, we write the updated memory entities back to the RAG. Because we only update memory entities similar to new information, efficiency is ensured by restricting updates to old memory entities relevant to the new entity instead of all memory entities stored in the RAG.

3.3 Closed-Loop Planning Module for Dynamic Environment

The Closed-Loop Planning Module integrates information about the current task provided by the Spatial-Temporal Memory, Semantic and Episodic information recorded in long-term memory, and current observations to perform action planning. Each action is planned and passed on to the low-level executor for execution.

To enable closed-loop control in embodied environments, the Closed-Loop Planning Module adopts the Planner-Critic mechanism [21], which consists of the planner and the critic module. For each planning step, the planner generates a long-term plan consisting of multiple steps. However, due to the dynamics of embodied environments, the action sequence in the long-term plan may become outdated during the execution of the plan. Thus, before executing each step, we use the Critic model to evaluate whether the proposed action in this step remains appropriate under the latest environment. If not, the planner will replan based on the latest information. The demonstration of this process is shown in Fig. 2.

However, our experiments reveal that the original Planner-Critic mechanism may suffer from infinite loops. In the original mechanism, the first step of the action sequence output by the Planner is evaluated by the Critic before execution, which can lead to an infinite loop: if the Critic always demands replanning, no action will ever be executed. To address this, we modified the Planner-Critic mechanism so that the first step is not evaluated by the Critic. This ensures that even if the Critic persistently demands replanning, the RoboMemory will still execute actions.

3.4 Low-level Executer

The RoboMemory framework is a two-layer hierarchical agent framework. This design enables RoboMemory to accomplish longer-term tasks in the real world. The upper layer is responsible only for high-level planning, while the Low-level Executer carries out the actions planned by the upper layer in the real environment.

We use π_0 [7] as the action executor in the real world: actions planned by RoboMemory are converted into arm and chassis movements via the Low-level Executer, with LoRA fine-tuning applied to optimize performance in real-world robotic tasks. More details are provided in Appendix B.

Table 1: Comparison of Success Rates (SR) and Goal Condition Success Rates (GC) across difficulty levels between RoboMemory and baseline methods on EB-ALFRED.

Method	Type	Avg.		Base		Long	
		SR	GC	SR	GC	SR	GC
Single VLM-Agents							
GPT-4o	Closed-source	59.0 %	68.3 %	64.0 %	74.0 %	54.0 %	62.5 %
GPT-4o-mini		17.0 %	32.4 %	34.0 %	47.8 %	0.0 %	17.0 %
Claude-3.5-Sonnet		62.0 %	63.3 %	72.0 %	72.0 %	52.0 %	54.5 %
Gemini-1.5-Pro		64.0 %	69.7 %	70.0 %	74.3 %	58.0 %	65.0 %
Gemini-2.0-flash		60.0 %	63.9 %	62.0 %	65.7 %	58.0 %	62.0 %
Llama-3.2-90B-Vision-Ins	Open-source	27.0 %	33.9 %	38.0 %	43.7 %	16.0 %	24.0 %
InternVL2.5-78B		40.0 %	45.7 %	38.0 %	42.3 %	42.0 %	49.0 %
InternVL3-78B		37.0%	-	38.0%	-	36.0%	-
InternVL2.5-38B		31.0%	36.9%	36.0%	37.3%	26.0%	36.5%
Qwen2.5-VL-72B-Ins		42.0 %	-	50.0 %	-	34.0 %	-
VLM-Agent Frameworks							
Voyager (Qwen2.5-VL-72B-Ins)	Baselines	44.0%	63.7%	56.0%	73.2%	32.0%	54.2%
Reflexion (Qwen2.5-VL-72B-Ins)		29.0%	43.2%	48.0%	54.0%	10.0%	33.0%
Cradle (Qwen2.5-VL-72B-Ins)		43.0%	54.6%	54.0%	67.9%	32.0%	41.0%
RoboMemory (Qwen2.5-VL-72B-Ins)	Ours	67.0 %	78.4 %	68.0 %	75.5 %	66.0 %	81.3 %

4 Experiments

4.1 Experimental Setup

4.1.1 Benchmarks

To evaluate the task planning ability of RoboMemory, we select a subset of the EB-ALFRED benchmark from EmbodiedBench[38]. We selected the Base and Long subsets because they aim to test the agent’s planning ability. The Base and Long subset comprises 100 tasks for complex long-term embodied tasks. The EB-ALFRED environment provides a visually-grounded operational setting that closely mimics real-world conditions (see appendix B for environment details), enabling direct comparison with established baselines.

Additionally, we have also evaluated RoboMemory’s capabilities on the EB-Habitat benchmark [38]. For detailed results, please refer to Appendix A.

Moreover, we set up a real-world environment to test the lifelong learning ability of RoboMemory in the real world.

4.1.2 Settings & Baselines

To facilitate comparisons, we consider two types of baselines. First, we choose the advanced closed-source and open-source VLMs as a single agent. We compare their performance with RoboMemory. For closed source VLMs, we choose GPT-4o and GPT-4o-mini [26, 19], Claude3.5-Sonnet [3], Gemini-1.5-Pro and Gemini-2.0-flash [34, 11]. For open source VLMs, we choose LLama-3.2-90B-Vision-Ins [24], InternVL-2.5-78B/28B [10], InternVL-3-72B [46], and [4]. Secondly, we choose three agent frameworks: (1) Reflexion [30], which introduces a simple long-term memory and a self-reflection module. Reflexion uses the self-reflection module to summarize experiences as long-term memory, thereby enhancing the model’s capabilities. (2) Voyager [36], which utilizes a skill library as its procedural memory, is a widely used baseline for embodied agent planning. (3) Cradle [33], which proposes a general agent framework with episodic and procedural memory and gains good performances at various multi-model agent tasks.

In our experiments, each agent framework is tested using Qwen2.5-VL-72b-Ins [35]. The Qwen2.5-VL-72b-Ins represents a high-performing open-source alternative. Notably, the Qwen2.5-VL-72b-Ins demonstrates performance comparable to advanced closed-source VLMs in several benchmark tasks [37]. We use the Qwen3-Embedding model [44] to create embedding vectors for RAGs in RoboMemory. For the Low-level Executor, since EB-ALFRED provides high-level action APIs, we use the low-level executor provided by EmbodiedBench instead of the VLA-based method.

4.1.3 Evaluation metrics

We define two evaluation metrics to assess the performance: (1) Success Rate (SR), which is the ratio of completed tasks to the total number of tasks in each difficulty level. This metric reflects the agent’s ability to complete tasks across randomly generated scenarios. (2) Goal Condition Success Rate (GC), which is the ratio of intermediate conditions achieved to the maximum possible score in each scenario. An GC of 100% indicates that the task is completed in the given scenario. These two metrics can be computed as:

$$SR = \mathbb{E}_{x \in \mathcal{X}} [\mathbb{1}_{SCN_x = GCN_x}] \quad (1)$$

$$GC = \mathbb{E}_{x \in \mathcal{X}} \left[\frac{SCN_x}{GCN_x} \right] \quad (2)$$

Where \mathcal{X} denotes the test subset, and x represents a test task. The success condition number (SCN_x) refers to the number of conditions the agent has accomplished, while the global condition number (GCN_x) indicates the total number of conditions required for task completion. The task is considered successful if $SCN_x = GCN_x$.

4.2 Main results

As shown in Table 1, our model achieves significant improvements over both single VLM agents and Agent frameworks on the EB-ALFRED. Compared to the SOTA Single VLM-Agent model, Gemini1.5-Pro, RoboMemory with Qwen2.5-VL-72B-Ins backbone improves the average SR by 3% and GC by 8.7%. This demonstrates RoboMemory’s superiority over single VLM-Agents, proving that an Agent framework with open-source models can outperform closed-source SOTA models. Furthermore, when tested against other VLM-Agent frameworks, RoboMemory also shows substantial gains. This is because, unlike other frameworks, RoboMemory’s brain-like memory system provides embodied models with more accurate and persistent contextual information. Additionally, the Planner-Critic mechanism provides a closed-loop planning ability, which helps the RoboMemory gain better performance in long-term tasks. Because the RoboMemory can detect and try to overcome possible failures. And it is more robust when encountering unexpected situations.



Figure 3: The reason why RoboMemory failed to complete the task.

4.3 Ablation Studies

We used the full Base and Long Subset from EB-ALFRED to validate RoboMemory’s effectiveness. We removed each component systematically and observed performance changes across task categories. We use the success rate as our metric. Results are shown in Table 2.

4.3.1 Long-term Memory

Adding long-term memory significantly improved RoboMemory’s success rate. The experiment shows that it enables continuous learning while it attempts to complete tasks. The semantic memory learns low-level skills’ properties, such as in what circumstances an action may fail. The temporal memory records all task attempts (successful/failed), providing valuable experience at the task level and giving insight into how to complete a task successfully. This helps the RoboMemory predict action outcomes and avoid ineffective attempts. This ability indicates that the RoboMemory has a lifelong learning capability.

4.3.2 Spatial Memory

Spatial memory is crucial for embodied agents, especially given that current pretrained VLMs have limited spatial understanding ability. Our novel dynamic KG update algorithm enables KG-based spatial memory in dynamic environments. This spatial reasoning helps RoboMemory handle partially observable embodied settings.

4.3.3 Critic Module

Table 2 shows performance without the critic module (55% vs 67% with full system). This drop highlights how the critic’s closed-loop planning adapts to dynamic environments. It helps RoboMemory recover from failures faster and handle unexpected situations better.

4.4 Real-world Robot Deployment

To evaluate RoboMemory’s lifelong learning capability in the real world, we designed a kitchen environment inspired by EB-ALFRED and EB-Habitat. The scene contains 5 navigable points, 8 interactive objects, and over 10 non-interactive (but potentially distracting) items. The environment is shown in Figure 4. In the real world, we use continuous video recordings during action execution (instead of snapshots after the action execution) as RoboMemory’s input. This provides a more temporally coherent perception. We created three task categories (5 tasks each, totaling 15 tasks) with difficulty matching EB-ALFRED’s Base subset (average oracle trajectory length from 5 to 10). Additional hardware experiment details are in Appendix B.

Table 2: Ablation Study on RoboMemory’s Success Rate (SR)

Method	Avg.	Base	Long
RoboMemory	67%	68%	66%
- w/o critic	55 %	60 %	50%
- w/o spatial memory	47 %	52 %	42 %
- w/o long-term memory	57%	66%	48%



Figure 4: Visualization of the experimental environment.

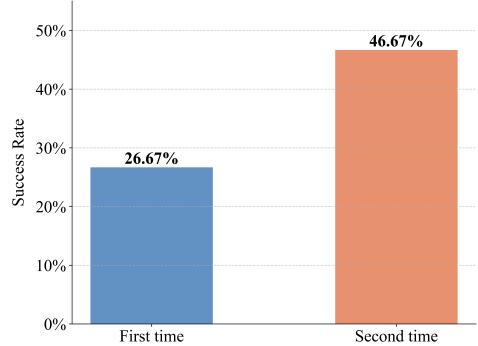


Figure 5: The improvement of RoboMemory after learning in the real world.

To test the lifelong learning ability of RoboMemory, we run each task twice without clearing long-term memory between attempts. Success rates for first and second attempts are shown in Figure 5.

The second attempt showed significantly higher success rates. This proves RoboMemory’s long-term memory effectively guides subsequent tasks in real embodied environments. Key observations include: (1) Closed-loop error recovery: RoboMemory retries failed actions when possible, even if the low-level executor (VLA model) fails. (2) Spatial reasoning: RoboMemory remembers object locations and spatial relationships using its memory. (3) Lifelong learning: RoboMemory analyzes failure causes reasonably. These analyses guide future decisions.

Detailed examples demonstrating these capabilities and further discussions are provided in Appendix E.

Moreover, we observed a significant drop in task success rates when deploying the agent with the Low-level Executor in real-world environments. This performance degradation primarily stems from the executor’s inherent limitations: (1) The VLA model exhibits unreliable instruction-following capabilities, frequently failing during grasping actions or selecting incorrect objects; (2) Pre-trained VLM models demonstrate inadequate video understanding - while capable of recognizing static objects, they struggle to interpret dynamic visual information such as action failures or state changes. These limitations collectively contribute to the reduced performance compared to simulated environments.

5 Conclusion and Future Work

In conclusion, RoboMemory, a brain-inspired multi-memory framework, enables lifelong learning and long-term planning in real-world embodied systems by addressing key challenges: continuous learning, memory latency, task correlation capture, and planning infinite loops. Experimental results on EmbodiedBench show it outperforms SOTA closed-source VLMs and agent frameworks, with ablation studies validating critical components like the Critic module and spatial/long-term memory. Real-world deployment confirms its lifelong learning capability via improved success rates in repeated tasks. While limited by reasoning errors and executor reliance, it paves the way for generalizable memory-augmented agents, with future work focusing on refining reasoning and enhancing execution robustness.

A key unsolved problem in current hierarchical agent research for embodied tasks, including ours, is about the interaction between high-level agents and low-level executors (e.g., VLAs). Most existing frameworks use language instructions solely as action instructions from high-level agents. However, some action details are hard to describe with language. Other modalities (e.g., vision) can better represent these details (e.g., grasp points). While our work focuses on long-term planning and lifelong learning for agents, future work can consider enhancing generalization within existing frameworks by improving the interaction method between the VLA and the Agent.

References

- [1] Saaket Agashe, Jiuzhou Han, Shuyu Gan, Jiachen Yang, Ang Li, and Xin Eric Wang. Agent s: An open agentic framework that uses computers like a human. *arXiv preprint arXiv:2410.08164*, 2024.
- [2] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- [3] Anthropic. Claude 3.5 sonnet, 2024.
- [4] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.
- [5] Suneel Belkhale, Tianli Ding, Ted Xiao, Pierre Sermanet, Quon Vuong, Jonathan Tompson, Yevgen Chebotar, Debidatta Dwibedi, and Dorsa Sadigh. Rt-h: Action hierarchies using language. *arXiv preprint arXiv:2403.01823*, 2024.
- [6] Johan Bjorck, Fernando Castañeda, Nikita Cherniadev, Xingye Da, Runyu Ding, Linxi Fan, Yu Fang, Dieter Fox, Fengyuan Hu, Spencer Huang, et al. Gr00t n1: An open foundation model for generalist humanoid robots. *arXiv preprint arXiv:2503.14734*, 2025.
- [7] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. π_0 : A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.
- [8] Neil Burgess, Eleanor A Maguire, and John O’Keefe. The human hippocampus and spatial and episodic memory. *Neuron*, 35(4):625–641, 2002.
- [9] Xiaojun Chen, Shengbin Jia, and Yang Xiang. A review: Knowledge reasoning over knowledge graph. *Expert systems with applications*, 141:112948, 2020.
- [10] Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*, 2024.
- [11] Google DeepMind. Introducing gemini 2.0: our new ai model for the agentic era, 2024.
- [12] Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*, 2024.
- [13] Dayuan Fu, Biqing Qi, Yihuai Gao, Che Jiang, Guanting Dong, and Bowen Zhou. Msi-agent: Incorporating multi-scale insight into embodied agents for superior planning and decision-making. *arXiv preprint arXiv:2409.16686*, 2024.
- [14] Zipeng Fu, Tony Z Zhao, and Chelsea Finn. Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation. *arXiv preprint arXiv:2401.02117*, 2024.
- [15] Marc Glocker, Peter Höning, Matthias Hirschmanner, and Markus Vincze. Llm-empowered embodied agent for memory-augmented task planning in household robotics. *arXiv preprint arXiv:2504.21716*, 2025.
- [16] Bernal Jiménez Gutiérrez, Yiheng Shu, Yu Gu, Michihiro Yasunaga, and Yu Su. Hipporag: Neurobiologically inspired long-term memory for large language models. *arXiv preprint arXiv:2405.14831*, 2024.
- [17] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.

[18] Yafei Hu, Quanting Xie, Vidhi Jain, Jonathan Francis, Jay Patrikar, Nikhil Keetha, Seungchan Kim, Yaqi Xie, Tianyi Zhang, Hao-Shu Fang, et al. Toward general-purpose robots via foundation models: A survey and meta-analysis. *arXiv preprint arXiv:2312.08782*, 2023.

[19] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.

[20] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.

[21] Mingcong Lei, Ge Wang, Yiming Zhao, Zhixin Mai, Qing Zhao, Yao Guo, Zhen Li, Shuguang Cui, Yatong Han, and Jinke Ren. Clea: Closed-loop embodied agent for enhancing task execution in dynamic environments. *arXiv preprint arXiv:2503.00729*, 2025.

[22] Bill Yuchen Lin, Yicheng Fu, Karina Yang, Faeze Brahman, Shiyu Huang, Chandra Bhagavatula, Prithviraj Ammanabrolu, Yejin Choi, and Xiang Ren. Swiftsage: A generative agent with fast and slow thinking for complex interactive tasks. *Advances in Neural Information Processing Systems*, 36, 2024.

[23] Kourosh Maboudi, Bapun Giri, Hiroyuki Miyawaki, Caleb Kemere, and Kamran Diba. Retuning of hippocampal representations during sleep. *Nature*, 629(8012):630–638, 2024.

[24] Meta. Llama 3.2: Revolutionizing edge ai and vision with open, customizable models, 2024.

[25] David Milner. Cognitive neuroscience: the biology of the mind and findings and current opinion in cognitive neuroscience. *Trends in cognitive sciences*, 2(11):463, 1998.

[26] OpenAI. Gpt-4o mini: advancing cost-efficient intelligence, 2024.

[27] Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G. Patil, Ion Stoica, and Joseph E. Gonzalez. MemGPT: Towards llms as operating systems. *arXiv preprint arXiv:2310.08560*, 2023.

[28] Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22, 2023.

[29] Lucy Xiaoyang Shi, Brian Ichter, Michael Equi, Liyiming Ke, Karl Pertsch, Quan Vuong, James Tanner, Anna Walling, Haohuan Wang, Niccolo Fusai, et al. Hi robot: Open-ended instruction following with hierarchical vision-language-action models. *arXiv preprint arXiv:2502.19417*, 2025.

[30] Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36, 2024.

[31] Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M Sadler, Wei-Lun Chao, and Yu Su. Llm-planner: Few-shot grounded planning for embodied agents with large language models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2998–3009, 2023.

[32] Huajie Tan, Xiaoshuai Hao, Cheng Chi, Minglan Lin, Yaoxu Lyu, Mingyu Cao, Dong Liang, Zhuo Chen, Mengsi Lyu, Cheng Peng, et al. Roboos: A hierarchical embodied framework for cross-embodiment and multi-agent collaboration. *arXiv preprint arXiv:2505.03673*, 2025.

[33] Weihao Tan, Wentao Zhang, Xinrun Xu, Haochong Xia, Ziluo Ding, Boyu Li, Bohan Zhou, Junpeng Yue, Jiechuan Jiang, Yewen Li, et al. Cradle: Empowering foundation agents towards general computer control. *arXiv preprint arXiv:2403.03186*, 2024.

[34] Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.

- [35] Qwen Team. Qwen2.5: A party of foundation models, September 2024.
- [36] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*, 2023.
- [37] Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Ben Feuer, Siddhartha Jain, Ravid Shwartz-Ziv, Neel Jain, Khalid Saifullah, Siddartha Naidu, et al. Livebench: A challenging, contamination-free llm benchmark. *arXiv preprint arXiv:2406.19314*, 2024.
- [38] Rui Yang, Hanyang Chen, Junyu Zhang, Mark Zhao, Cheng Qian, Kangrui Wang, Qineng Wang, Teja Venkat Koripella, Marziyeh Movahedi, Manling Li, et al. Embodiedbench: Comprehensive benchmarking multi-modal large language models for vision-driven embodied agents. *arXiv preprint arXiv:2502.09560*, 2025.
- [39] Zhejian Yang, Yongchao Chen, Xueyang Zhou, Jiangyue Yan, Dingjie Song, Yinuo Liu, Yuting Li, Yu Zhang, Pan Zhou, Hechang Chen, et al. Agentic robot: A brain-inspired framework for vision-language-action models in embodied agents. *arXiv preprint arXiv:2505.23450*, 2025.
- [40] Zhutian Yang, Caelan Garrett, Dieter Fox, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. Guiding long-horizon task and motion planning with vision language models. *arXiv preprint arXiv:2410.02193*, 2024.
- [41] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*, 2022.
- [42] Haoqi Yuan, Yu Bai, Yuhui Fu, Bohan Zhou, Yicheng Feng, Xinrun Xu, Yi Zhan, Börje F Karlsson, and Zongqing Lu. Being-0: A humanoid robotic agent with vision-language models and modular skills. *arXiv preprint arXiv:2503.12533*, 2025.
- [43] Hongxin Zhang, Weihua Du, Jiaming Shan, Qinhong Zhou, Yilun Du, Joshua B Tenenbaum, Tianmin Shu, and Chuang Gan. Building cooperative embodied agents modularly with large language models. *arXiv preprint arXiv:2307.02485*, 2023.
- [44] Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, et al. Qwen3 embedding: Advancing text embedding and reranking through foundation models. *arXiv preprint arXiv:2506.05176*, 2025.
- [45] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.
- [46] Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, et al. Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025.

A Additional Experiments

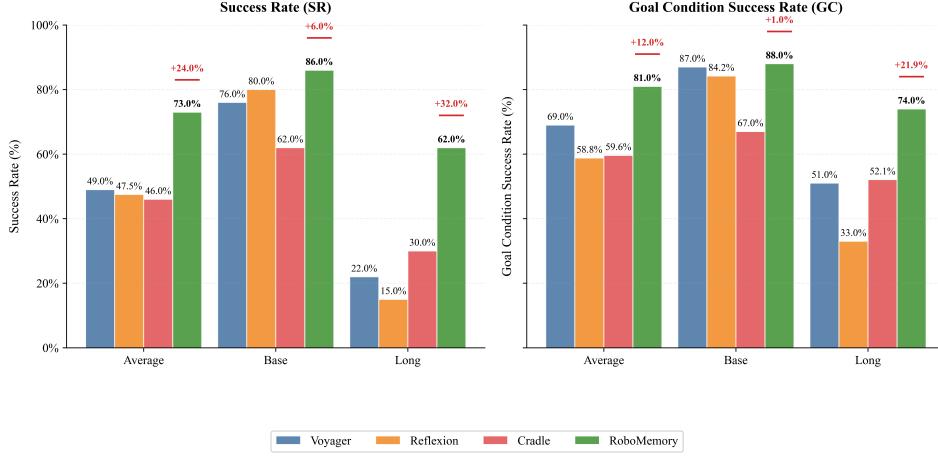


Figure 6: Comparison of Success Rates (SR) and Goal Condition Success Rates (GC) across difficulty levels between RoboBrain and baseline methods on EB-Habitat.

A.1 Evaluate on EB-Habitat

Similar to EB-ALFRED, we also deploy RoboMemory to EB-Habitat. Both EB-Habitat and EB-HALFRED are subsets of EmbodiedBench. We evaluate our model on the Base and Long subsets. Each subset contains 50 different trails. The results are shown in Fig 6.

The results demonstrate that our RoboMemory can adapt well to different environments. It achieves significant improvements over the baseline across various settings. On average, the success rate increases by 24% compared to the SOTA Multi-Agent Method. The goal-conditioned success rate improves by 12%.

These improvements indicate that RoboMemory enhances the agent’s embodied intelligence in various environments. The key factor is its complete memory system.

A.2 Error Analysis

We analyze RoboMemory trajectories for failed tasks. We identify error types based on the above definitions. A single task may contain multiple errors. We calculate the occurrence probability of each error type to show RoboMemory’s strengths and weaknesses. The results are shown in Figure 3.

We can observe that among all error types, the planning errors are the most common. This means that even though the memory modules can provide comprehensive information about the RoboMemory agent’s previous experience and spatial and temporal memory for the current task, the planner module may still not provide good action plans. This may be due to the capability of the pretrained base model.

The most common perception error is the hallucination error. We can observe that although some hallucinations can be handled by the critic module or memory information, there are still some cases in which the planner ignores all insights from memory and critic and fails to complete the task.

The detailed examples and discussions are provided in Appendix E.

B Additional Environment settings

B.1 EB-ALFRED and EB-Habitat

We adopt the same environment parameters as in EmbodiedBench. The maximum steps per task are set to 30, with image inputs of size 500×500 . The temporal memory buffer length is set to

Table 3: Robot Action Command For different environments

Action Type	EB-ALFRED	EB-Habitat	Real World
Navigate to object/navigation point	find(obj)	navigate(point)	navigate_to(point)
Pick Up Object	pick_up(obj)	pick(obj)	pick_up(obj)
Drop to Ground	drop()	-	-
Place to Receptacle	put_down()	place(rec)	put_down_to(rec)
Open Object	open(obj)	open(obj)	open(obj)
Close Object	close(obj)	close(obj)	close(obj)
Turn On	turn_on(obj)	-	turn_on(obj)
Turn Off	turn_off(obj)	-	turn_off(obj)
Slice Object	slice(obj)	-	-
Task Complete	-	-	task_complete()

3. However, we modify the action formats of EB-ALFRED and EB-Habitat to better simulate real-world scenarios. Specifically, we define different action APIs (Python functions), where each action takes an object parameter indicating its target. We extract all possible objects from the environment as inputs to the Agent. The Agent must select appropriate actions and object parameters based on task requirements. Compared to the original interaction method in EmbodiedBench (which enumerates all possible actions, including both action names and target objects, and requires the Agent to choose), our approach offers greater flexibility. The detailed action APIs are presented in Table 3.

Since EB-ALFRED and EB-Habitat provide comprehensive high-level action APIs, we do not employ the VLA-Based Low-Level Executor in these environments. Instead, we utilize the built-in low-level controllers from EmbodiedBench.

B.2 Real-world experiments

We construct a common kitchen scenario to evaluate the RoboMemory framework’s lifelong learning capabilities in real-world settings. Using Mobile ALOHA [14] as our physical robotic platform, we design three categories of tasks: (1) Pick up & put down: The agent must locate a specified object among all possible positions and place it at a designated location. This task tests the model’s basic object-searching and planning abilities. (2) Pick up, operate & put down: Building upon the first task, the agent must additionally perform operations such as heating or cleaning the object. This task requires longer-term planning, which is crucial in embodied environments. (3) Pick up, gather & put down: The agent must place specified objects into a movable container and then move the container to a target location. This task evaluates the agent’s understanding of object relationships, requiring it to remember the positions of at least two objects (the container and the target item) and their spatial relationship.

To adapt to the real-world setup, we define high-level action APIs similar to those in EB-ALFRED and EB-Habitat. Additionally, we train a VLA-based model to execute tasks according to our action APIs. The detailed action APIs are presented in Table 3.

For the low-level executor, we use one main camera and two arm-mounted cameras as input, each with a resolution of 640×480 . The temporal memory buffer length is set to 3.

In our experiments, we set the maximum steps per task to 15. We also provide an API for actively terminating tasks. Since real-world environments lack direct success/failure feedback, RoboMemory must autonomously determine task completion. To prevent excessively long task execution, we enforce termination after 15 steps if no success is achieved. A single main camera (640×480 resolution) records video during action execution as input for RoboMemory’s higher-level processing.

B.3 Training Details of Low-Level Executer

We use the π_0 model as our foundation model. We collected 1,040 data samples over 10 types of tasks for fine-tuning. We use LoRA [17] fine-tuning to save resources during fine-tuning. The specific fine-tuning parameters and action types are given in Table 4. For tasks involving both pick-up and place actions, we split these tasks into separate pick-up and place actions. These are then treated as two distinct data samples during training. The separation of pick-up and place action allows the VLA to carry an object in its hand. For training, we used a server with six A100-80GB GPUs. The total training time was 12 hours.

Table 4: Dataset statistics and training hyperparameters for robotic manipulation tasks.

Dataset Statistics		Training Configuration	
Action Type	#Episodes	Parameter	Value
Turning on/off faucet	142	Optimizer	AdamW
Picking up & Placing basket on counter	63	Batch size	32×6
Picking up & Placing basket in sink	72	Training steps	10,000
Picking up & Placing banana into basket	114	Learning rate	6.12×10^{-5}
Throwing bottle into trash bin	132	warm up step	500
Placing gum box on dish	120	LoRA Configuration	
Picking up & Placing cup on plate	51	rank	16
Picking up & Placing dish into sink	69	α	16
Throwing paper ball into trash bin	135	Resource Usage	
Open/close oven	142	GPU	A100-80GB $\times 6$
Total episodes	1040	Training time	12 hours

Besides, we use the built-in LiDAR SLAM system of the Mobile ALOHA robot base as the navigation action actuator. We define five typical navigation points, similar to EB-Habitat. We used SLAM to navigate between these navigation points.

C Algorithm for Dynamic KG update algorithm

Algorithm 1 Retrieval-based Incremental Knowledge Graph Update Algorithm

Require: New triplet set T_{new} , main knowledge graph $G = (V, E)$, latest retrieval entities $Entity_t$, entity embeddings $\mathcal{E} : V \rightarrow \mathbb{R}^d$, similarity threshold θ , conflict detection VLM_{dt} , retrieval function $\text{Retrieve}(\cdot)$

Ensure: Updated consistent knowledge graph G'

- 1: Step 1: Retrieve relevant subgraph based on new triplets
- 2: $G_{retrieved} \leftarrow \text{Retrieve}(G, Entity_t, \mathcal{E}, \theta)$
- 3: Step 2: Create local working subgraph
- 4: $G_{local} \leftarrow G_{retrieved} \cup T_{new}$
- 5: Step 3: Conflict detection within local subgraph
- 6: $G_{NewSubgraph} \leftarrow VLM_{dt}(G_{local})$
- 7: Step 4: Generate updated subgraph
- 8: $G_{updated} \leftarrow \text{resolve all conflicts and integrate } T_{new} \text{ into } G_{NewSubgraph}$
- 9: Step 5: Merge updated subgraph back to main graph
- 10: $G' \leftarrow \text{merge } G_{updated} \text{ back into } G$
- 11: Update vector database V with new embeddings
- 12: **return** G'

D Proof of Dynamic Spatial Memory Update Algorithm

Theorem 1 (Upper Bound on K-hop Node Extraction in Directed Graphs). *Let $G = (V, E)$ be a finite directed graph with maximum out-degree $D \geq 1$, and let $\mathcal{S} \subseteq V$ be a set of M source nodes. Define the K-hop neighborhood $\mathcal{N}_K(s)$ of a node $s \in \mathcal{S}$ as the set of nodes reachable from s via directed paths of length at most K . Then the total number of distinct nodes in the union of all K-hop neighborhoods,*

$$\mathcal{N}_K(\mathcal{S}) = \bigcup_{s \in \mathcal{S}} \mathcal{N}_K(s),$$

Satisfies the following upper bound:

$$|\mathcal{N}_K(\mathcal{S})| \leq \begin{cases} M \cdot \frac{D^{K+1} - 1}{D - 1}, & \text{if } D > 1, \\ M \cdot (K + 1), & \text{if } D = 1. \end{cases}$$

Proof. For any node $s \in \mathcal{S}$, the number of distinct nodes reachable from s within i hops is at most D^i , assuming the worst-case scenario where each node encountered has the maximum out-degree D , and all neighbors are distinct and non-overlapping.

Thus, the size of the K-hop neighborhood of a single node satisfies:

$$|\mathcal{N}_K(s)| \leq \sum_{i=0}^K D^i = \begin{cases} \frac{D^{K+1} - 1}{D - 1}, & \text{if } D > 1, \\ K + 1, & \text{if } D = 1. \end{cases}$$

Since there are M such source nodes and assuming no overlaps between their K-hop neighborhoods (worst case), the union size satisfies:

$$|\mathcal{N}_K(\mathcal{S})| \leq M \cdot |\mathcal{N}_K(s)|.$$

Substituting the bound on $|\mathcal{N}_K(s)|$ gives the result. \square

Theorem 2 (Upper Bound for K-hop Node Extraction in Normalized Directed Graphs). *Let $G = (V, E)$ be a finite directed graph with $|V| = n$ nodes. Assume the maximum out-degree is at most $D_{\max} = Dn$, and the maximum in-degree is at most $N_{\max} = Nn$, where $D, N \in (0, 1]$ are constants. Let $\mathcal{S} \subseteq V$ be a set of M source nodes. Define $\mathcal{N}_K(\mathcal{S})$ as the union of all nodes reachable from \mathcal{S} via paths of length at most K , using only outgoing edges. Then the number of extracted nodes satisfies:*

$$|\mathcal{N}_K(\mathcal{S})| \leq \min \left\{ n, M \cdot \frac{(Dn)^{K+1} - 1}{Dn - 1} \right\}.$$

In particular, when $Dn \gg 1$, we have the approximation:

$$|\mathcal{N}_K(\mathcal{S})| \lesssim M \cdot (Dn)^K.$$

Proof. For each node $s \in \mathcal{S}$, the maximum number of reachable nodes within i -hops is at most $(Dn)^i$ under the assumption of maximum out-degree and no overlap.

Summing over hops from 0 to K , we get for each root:

$$|\mathcal{N}_K(s)| \leq \sum_{i=0}^K (Dn)^i = \frac{(Dn)^{K+1} - 1}{Dn - 1}.$$

Assuming no overlap among the M source node expansions (worst case), we have:

$$|\mathcal{N}_K(\mathcal{S})| \leq M \cdot \frac{(Dn)^{K+1} - 1}{Dn - 1}.$$

Since the total number of nodes in the graph is n , this quantity is also trivially bounded above by n , yielding the result. \square

E Supplementary Examples for Qualitative Analysis

E.1 Example of Dynamic Spatial Memory Update process

In RoboMemory’s Spatial Memory, the KG is dynamically constructed during environment exploration. As illustrated in Figure 7, we demonstrate the progressive expansion of the KG in Spatial Memory as the agent navigates through the environment. The figure indicates a continuous growth in the number of both nodes and edges of the KG as exploration progresses.

Notably, the KG undergoes dynamic updates through RoboMemory’s environmental interactions. For example, the initial KG state displays the relation “I am near the apple. But as the agent picks up the apple in the third step, in the fourth KG, the relationship becomes “I hold the apple”. This demonstrates RoboMemory’s capability for dynamic KG maintenance and expansion.

By querying this KG, the Planner-Critic module gains access to rich spatial information, empowering RoboMemory with robust spatial memory capabilities that significantly enhance its performance in both TextWorld and EmbodiedBench environments.

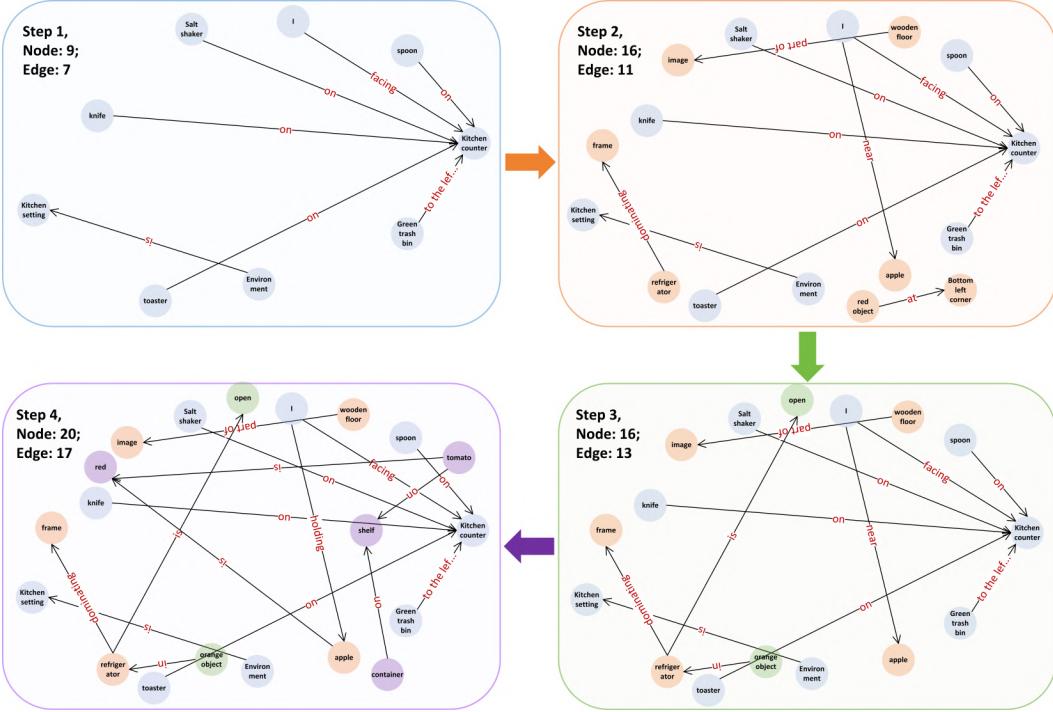


Figure 7: Visualization of Spatial Memory’s dynamic update process.

E.2 Real World

In Figure 8, we demonstrate an example of RoboMemory learning through trial and error in a real-world environment. Our task is “place a banana into the oven.” This task required RoboMemory to complete the objectives of finding the banana, picking it up, and transporting it to the oven. We observed that RoboMemory became stuck in an infinite loop during the first attempt. The banana was randomly placed on the “kitchen counter,” but RoboMemory overlooked this navigation target and remained trapped exploring other navigation targets instead.

However, based on this bad attempt, the semantic memory summarized that the robot should not repeatedly search in locations where the “banana” could not be found. Meanwhile, the episodic memory recorded what RoboMemory had done and the outcomes during the first attempt. Based on the information provided by semantic and episodic memory, in the second attempt, RoboMemory recognized that it had not previously tried navigating to the “kitchen counter.” After attempting this, it successfully completed the task. This example illustrates the role of RoboMemory’s long-term memory.

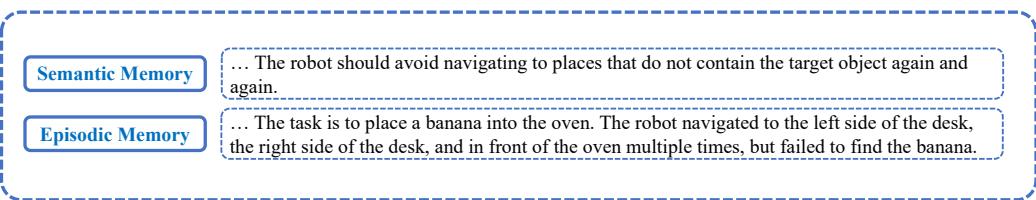
We also provide an example that completes the task in the first attempt. The example is shown in Figure 9. This example demonstrates that the RoboMemory has the ability to handle some relatively complex tasks in the real world. The task in this example is “Place a box of gum into the basket and put the basket on the kitchen counter”. Because two objects in different positions are involved in this task, RoboMemory has to memorize the position of at least one object to achieve the goal. With the help of the spatial memory, RoboMemory completes the task successfully.

E.3 EB-ALFRED

We select three examples in EB-ALFRED to show the errors that RoboMemory may encounter and the reasons why or why not RoboMemory can achieve the goal.

Task: Place banana into the oven

The first time



The second time

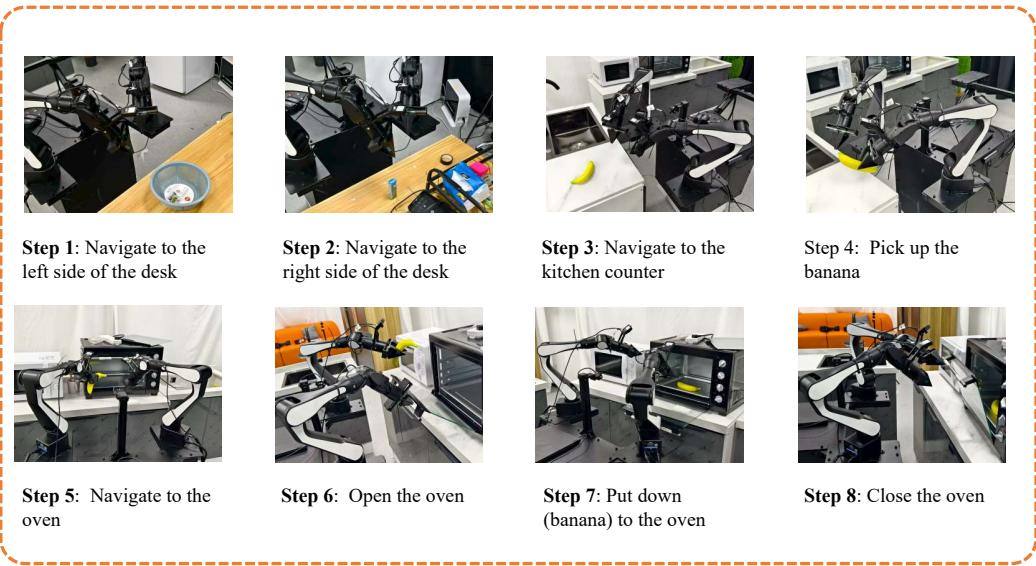


Figure 8: Case that a task is failed but the experience can help RoboMemory to succeed in the next try.

E.3.1 Successful example

We select a successful example to show how RoboMemory performed in the EB-ALFRED environment. The example trajectory is shown in Figure 10.

The task of this example is “set a plate with a spoon on it on the kitchen table”. However, in step 10, the Planner seems to ignore the temporal information from memory modules. RoboMemory thinks that it still needs to pick up the spoon (even though it has already placed a spoon in the plate). However, with the help of the critic, it finally becomes aware that picking up another spoon

Task: Place a box of gum into the basket and put the basket to the kitchen counter



Figure 9: Case that a task is successful.

is redundant, so RoboMemory goes back to the current trajectory and successfully completes the task at the end.

In this example, RoboMemory successfully overcame the hallucination and eventually achieved the goal. This example demonstrates that the critic module can help RoboMemory to overcome error cases.

E.3.2 Failed example

We demonstrate a representative example of the Critical Error. The example trajectory is shown in Figure 11. In this example, the task involves slicing and heating a tomato and moving the heated tomato slice to the trash can. Initially, RoboMemory successfully sliced the tomato with a knife. But when the planner plans the whole sequence, it forgets to drop the knife before picking up the tomato (this is necessary because in EB-ALFRED, the robot can only hold one object at a time). The critic and the planner should notice this situation and ask the critic to replan, as RoboMemory failed to pick up a tomato slice. However, the critic module ignores this issue, and thus, after it heats the knife instead of a tomato slice, it stacks in an infinite loop.

Besides, we provide another example demonstrating a representative failure caused by inaccurate action planning. The example trajectory is shown in Figure 12. In the trajectory, RoboMemory is asked to place two CDs into the drawer. However, at step 6, the robot failed to select correct CD object. In this experiment, RoboMemory has already put CD_2 into the drawer, but it keeps picking up CD_2 even though the memory has clearly indicated that CD_2 has already been put down. So we classify this as inaccurate action error. This indicates that the planner failed to comprehensively integrate information from both the memory and information-gathering modules, resulting in inaccurate action planning.

Task: Set plate with a spoon in it on the kitchen table

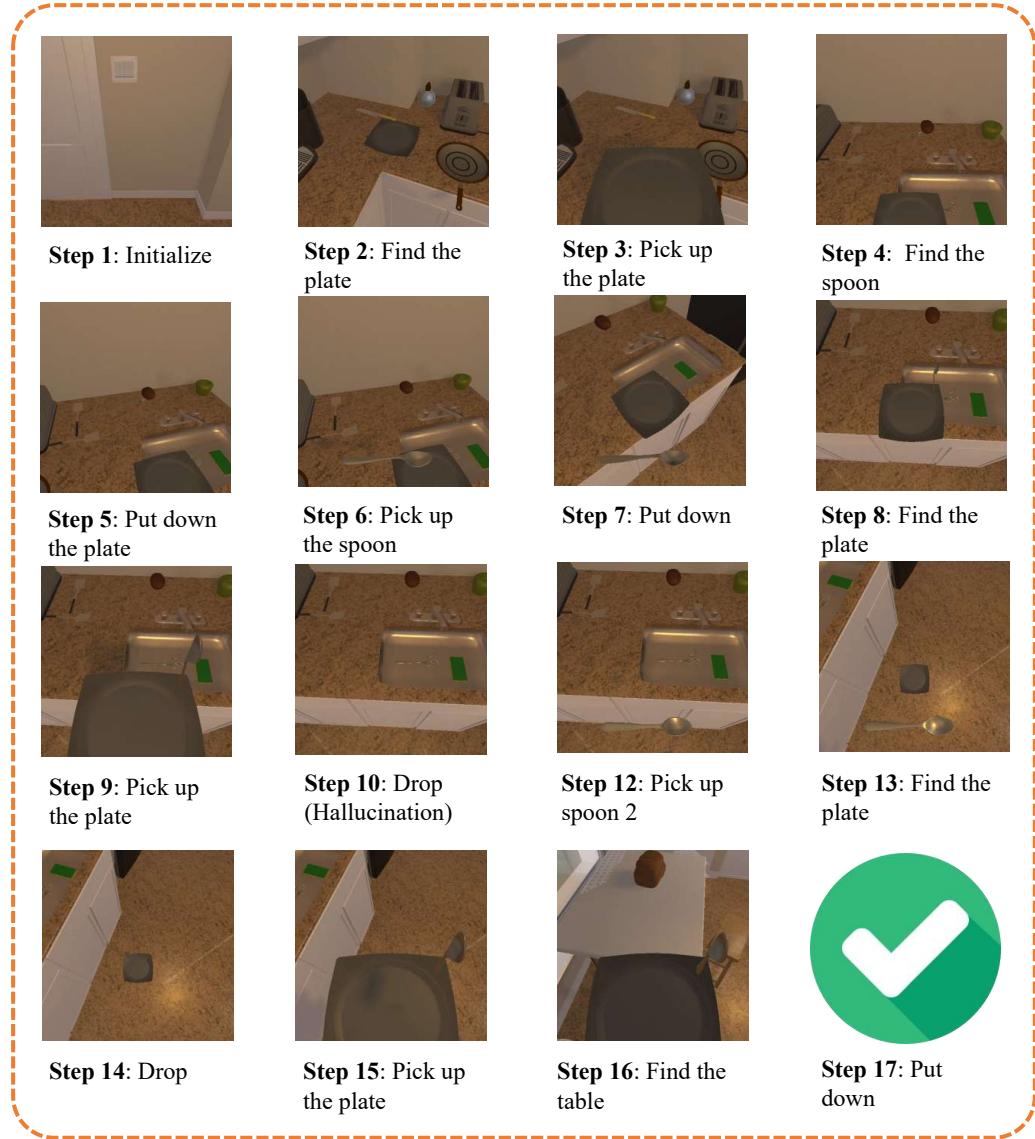


Figure 10: Case that a task is successful with the help of the critic and spatial memory modules.

Task: Cook a sliced tomato and throw it in the trash

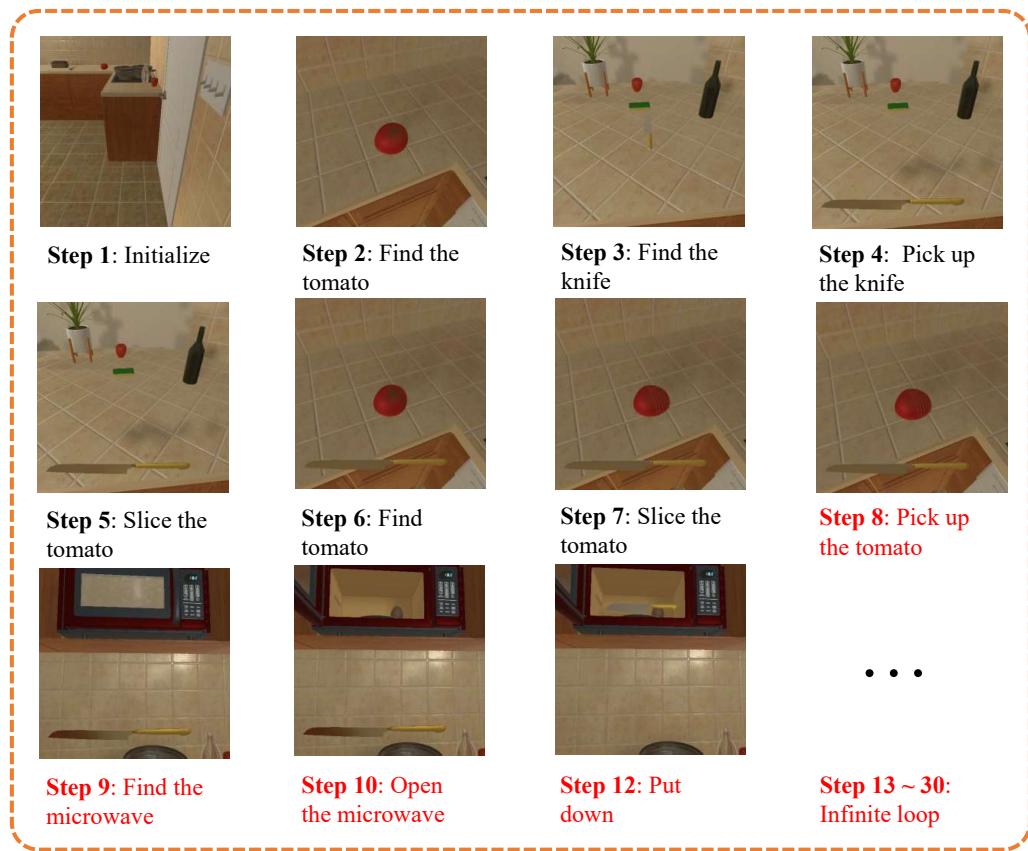


Figure 11: Case that a task fails in an infinite loop because the critic module failed to stop the agent when its planned action is no longer suitable.

Task: Move two CDs to the bottom drawer of the desk



Figure 12: Case that a task fails in an infinite loop because of inaccurate action planning.