Novelty Assessment Report

Paper: ReasoningBank: Scaling Agent Self-Evolving with Reasoning Memory

PDF URL: https://openreview.net/pdf?id=jL7fwchScm

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2026-01-01

Abstract

With the growing adoption of large language model (LLM) agents in persistent, real-world roles, they naturally encounter continuous streams of tasks and interactions. A key limitation, however, is their failure to learn from this accumulated experience, forcing them to discard valuable insights and repeat past errors. Unlike prior works that primarily store raw experience or successful routines, we propose ReasoningBank, a novel memory framework that allows an agent to self-curate generalizable reasoning strategies from both its successful and failed experiences for future leverage. This mechanism enables agents to generalize across tasks and become more capable over time. To accelerate and diversify this test-time learning process, we further propose memory-aware test-time scaling (MaTTS), which leverages a powerful synergy between memory and test-time scaling. On one hand, relevant memory from ReasoningBank guides the scaling process toward more effective exploration and improved reliability. On the other, scaling, in both parallel and sequential settings, generates abundant, diverse experiences that provide rich contrastive signals for synthesizing higher-quality memory. Experiments on web browsing and software engineering tasks show that ReasoningBank consistently outperforms existing memory mechanisms in both effectiveness and efficiency, with MaTTS further amplifying the gains. These findings position memory-driven experience as a new dimension of test-time scaling, where emergent behaviors naturally arise and agents acquire self-evolving capabilities.

Core Task Landscape

This paper addresses: **agent learning from experience through reasoning memory**

A total of **50 papers** were analyzed and organized into a taxonomy with **19 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Memory Architecture and Representation**
- **Memory-Based Learning Mechanisms**
- **Retrieval and Memory Operations**
- **Reasoning and Planning with Memory**
- **Cognitive and Neuroscience-Inspired Frameworks**
- **Domain-Specific Applications and Embodied Agents**
- **Surveys and Theoretical Foundations**

Complete Taxonomy Tree

- agent learning from experience through reasoning memory Survey Taxonomy
- Memory Architecture and Representation
  - Hierarchical and Structured Memory Systems (6 papers)
  - [8] Decentralizing AI Memory: SHIMI, a Semantic Hierarchical Memory Index for Scalable Agent Reasoning (Helmi, 2025) View paper
  - [9] AriGraph: Learning Knowledge Graph World Models with Episodic Memory for LLM Agents (Petr Anokhin, 2024) View paper
  - [10] ReaGAN: Node-as-Agent-Reasoning Graph Agentic Network (Guo Minghao, 2025) View paper
  - [12] Symagent: A neural-symbolic self-learning agent framework for complex reasoning over knowledge graphs (Ben Liu, 2025) View paper
  - [18] Agent kb: Leveraging cross-domain experience for agentic problem solving (Tang, 2025) View paper
  - [24] LiCoMemory: Lightweight and Cognitive Agentic Memory for Efficient Long-Term Reasoning (Huang Zheng-jun, 2025) View paper
  - Episodic and Sequential Memory Representations (5 papers)
  - [4] Memp: Exploring agent procedural memory (Liang Yuan, 2025) View paper
  - [15] Seeing, listening, remembering, and reasoning: A multimodal agent with long-term memory (Long Lin, 2025) View paper
  - [17] " my agent understands me better": Integrating dynamic human-like memory recall and consolidation in llm-based agents (Miyashita, 2024) View paper
  - [42] Efficient episodic memory utilization of cooperative multi-agent reinforcement learning (Hyungho Na, 2024) View paper
  - [45] MemVerse: Multimodal Memory for Lifelong Learning Agents (Junming Liu, 2025) View paper
  - Multimodal and Entity-Centric Memory (2 papers)
  - [23] ExpReS-VLA: Specializing Vision-Language-Action Models Through Experience Replay and Retrieval (Shahram Najam Syed, 2025) View paper
  - [32] EndoAgent: A Memory-Guided Reflective Agent for Intelligent Endoscopic Vision-to-Decision Reasoning (Tang Yi, 2025) View paper
- Memory-Based Learning Mechanisms
  - Self-Evolving and Lifelong Learning Agents (4 papers)

  - [3] Building self-evolving agents via experience-driven lifelong learning: A framework and benchmark (Cai Yu-Xuan, 2025) View paper
  - [13] SAGE: Self-evolving Agents with Reflective and Memory-augmented Abilities (Xuechen Liang, 2025) View paper
  - [16] Agent s: An open agentic framework that uses computers like a human (Agashe, 2024) View paper
  - [30] Smarter Together: Creating Agentic Communities of Practice through Shared Experiential Learning (Valentin Tablan, 2025) View paper
  - Reflection and Verbal Reinforcement Learning ★ (4 papers)
  - [0] ReasoningBank: Scaling Agent Self-Evolving with Reasoning Memory (Anon et al., 2026) View paper
  - [1] Reflexion: language agents with verbal reinforcement learning (Shinn, 2023) View paper
  - [33] R2D2: Remembering, Replaying and Dynamic Decision Making with a Reflective Agentic Memory (Tenghao Huang, 2025) View paper
  - [34] MetaReflection: Learning Instructions for Language Agents using Past Reflections (Gupta, 2024) View paper
  - Experience Replay and Trajectory Synthesis (4 papers)
  - [28] ERCI: An Explainable Experience Replay Approach with Causal Inference for Deep Reinforcement Learning (Jingwen Wang, 2025) View paper
  - [31] Agenttrek: Agent trajectory synthesis via guiding replay with web tutorials (Xu Yiheng, 2024) View paper
  - [47] Sample-Efficient Online Learning in LM Agents via Hindsight Trajectory Rewriting (Hu, 2025) View paper
  - [50] Memory-Augmented Agent Training for Business Document Understanding (Liu Jiale, 2024) View paper
  - Cross-Domain and Collaborative Experience Sharing (3 papers)
  - [2] Xolver: Multi-Agent Reasoning with Holistic Experience Learning Just Like an Olympiad Team (Rahman Salman, 2025) View paper
  - [6] Enhancing reasoning with collaboration and memory (Michelman, 2025) View paper
  - [44] AGENT KB: A Hierarchical Memory Framework for Cross-Domain Agentic Problem Solving (X Tang, 2025) View paper
- Retrieval and Memory Operations
  - Dynamic Memory Retrieval and Consolidation (2 papers)
  - [20] Review of Case-Based Reasoning for LLM Agents: Theoretical Foundations, Architectural Components, and Cognitive Integration (Hatalis, 2025) View paper
  - [22] Goal-Directed Search Outperforms Goal-Agnostic Memory Compression in Long-Context Memory Tasks (Yicong Zheng, 2025) View paper
  - Constant-Memory and Efficient Retrieval (1 papers)
  - [25] MEM1: Learning to Synergize Memory and Reasoning for Efficient Long-Horizon Agents (Zhou, 2025) View paper
  - Memory Management and Experience-Following Behavior (2 papers)
  - [11] How memory management impacts llm agents: An empirical study of experience-following behavior (Xiong, 2025) View paper
  - [37] From Knowledge to Noise: CTIM-Rover and the Pitfalls of Episodic Memory in Software Engineering Agents (Groh, 2025) View paper
- Reasoning and Planning with Memory
  - Long-Horizon and Multi-Step Reasoning (3 papers)
  - [26] Long-horizon Reasoning Agent for Olympiad-Level Mathematical Problem Solving (Songyang Gao, 2025) View paper
  - [39] DeepAgent: A General Reasoning Agent with Scalable Toolsets (Li Xiaoxi, 2025) View paper
  - [46] LLMs Are Not Good Strategists, Yet Memory-Enhanced Agency Boosts Reasoning (Y Wu, 2025) View paper
  - Knowledge-Augmented Reasoning (2 papers)
  - [14] GuardAgent: Safeguard LLM Agents via Knowledge-Enabled Reasoning (Zhen Xiang, 2024) View paper
  - [19] KG-Agent: An Efficient Autonomous Agent Framework for Complex Reasoning over Knowledge Graph (Jinhao Jiang, 2024) View paper
  - Hierarchical and Experience-Augmented Planning (1 papers)
  - [21] MC-GPT: Empowering Vision-and-Language Navigation with Memory Map and Reasoning Chains (Yu Lisha, 2024) View paper
- Cognitive and Neuroscience-Inspired Frameworks
  - Hippocampal and Compositional Memory Models (2 papers)
  - [7] Constructing future behavior in the hippocampal formation through composition and replay (Jacob J. W. Bakermans, 2025) View paper
  - [29] Constructing future behaviour in the hippocampal formation through composition and replay (Jacob J. W. Bakermans, 2024) View paper
  - Cognitive Architectures for Language Agents (3 papers)
  - [5] Cognitive architectures for language agents (Sumers, 2023) View paper
  - [38] Towards Neurocognitive-Inspired Intelligence: From AI's Structural Mimicry to Human-Like Functional Cognition (Noorbakhsh Amiri Golilarz, 2025) View paper
  - [40] Know rob 2.0â□□a 2nd generation knowledge processing framework for cognition-enabled robotic agents (Michael Beetz, 2018) View paper
  - Episodic Memory in Reinforcement Learning (2 papers)
  - [27] Reinforcement learning and episodic memory in humans and animals: an integrative framework (Samuel J. Gershman, 2017) View paper
  - [41] Transdreamer: Reinforcement learning with transformer world models (Chen Chang, 2022) View paper
- Domain-Specific Applications and Embodied Agents
  - Embodied Navigation and Visuomotor Control (1 papers)
  - [43] Visuomotor navigation for embodied robots with spatial memory and semantic reasoning cognition (Qiming Liu, 2024) View paper
  - Scientific and Engineering Problem Solving (2 papers)
  - [35] AgenticSciML: Collaborative Multi-Agent Systems for Emergent Discovery in Scientific Machine Learning (Qile Jiang, 2025) View paper
  - [49] Towards combining commonsense reasoning and knowledge acquisition to guide deep learning (M. Sridharan, 2023) View paper
- Surveys and Theoretical Foundations (2 papers)
  - [36] From language to action: A review of large language models as autonomous agents and tool users (Sadia Sultana Chowa, 2025) View paper

- ◦ [48] A survey on agentic multimodal large language models (Zhang Ruifei, 2025) View paper

Narrative

Core task: agent learning from experience through reasoning memory. This field explores how agents accumulate, organize, and leverage past experiences to improve decision-making and reasoning over time. The taxonomy reveals a multifaceted landscape organized into seven main branches. Memory Architecture and Representation addresses how agents structure and encode experiential knowledge, ranging from episodic traces to knowledge graphs. Memory-Based Learning Mechanisms focuses on how agents extract lessons from stored experiences, including reflection-driven approaches and verbal reinforcement paradigms. Retrieval and Memory Operations examines the computational processes for accessing relevant past information, while Reasoning and Planning with Memory investigates how agents integrate retrieved experiences into forward-looking decision processes. Cognitive and Neuroscience-Inspired Frameworks draw on biological memory systems to inform agent design, and Domain-Specific Applications and Embodied Agents apply these principles to robotics, navigation, and interactive environments. Finally, Surveys and Theoretical Foundations provide overarching perspectives on the field's conceptual underpinnings.

Within Memory-Based Learning Mechanisms, a particularly active line of work centers on reflection and verbal reinforcement learning, where agents iteratively refine their behavior by generating natural language critiques of past actions. Reflexion[1] pioneered this direction by enabling agents to self-reflect on task failures and adjust strategies accordingly, while more recent efforts like R2D2[33] and MetaReflection[34] extend these ideas to multi-step reasoning and meta-level introspection. ReasoningBank[0] situates itself within this cluster by emphasizing the construction of a structured memory bank that captures reasoning traces and supports iterative learning from experience. Compared to Reflexion[1], which focuses on immediate self-correction, ReasoningBank[0] appears to prioritize the accumulation and reuse of reasoning patterns across episodes, aligning closely with works like Self-Evolving Agents[3] that explore long-term knowledge consolidation. This branch highlights an ongoing tension between lightweight, episode-specific reflection and more persistent, architecturally integrated memory systems that scale across diverse tasks.

Related Works in Same Category

The following **3 sibling papers** share the same taxonomy leaf node with the original paper:

1. Reflexion: language agents with verbal reinforcement learning

Authors: Shinn, Noah, Cassano, Federico, Noah Shinn, et al. (17 authors total) | **Year/Venue**: 2023 • Neural Information Processing Systems | **URL**: View paper

Abstract

Large language models (LLMs) have been increasingly used to interact with external environments (e.g., games, compilers, APIs) as goal-driven agents. However, it remains challenging for these language agents to quickly and efficiently learn from trial-and-error as traditional reinforcement learning methods require extensive training samples and expensive model fine-tuning. We propose Reflexion, a novel framework to reinforce language agents not by updating weights, but instead through linguistic...

Relationship Analysis

Both papers belong to the Reflection and Verbal Reinforcement Learning category, focusing on agents that learn through linguistic feedback without weight updates. While ReasoningBank stores distilled reasoning strategies from both successful and failed experiences in a structured memory bank for future task guidance, Reflexion maintains reflective text in an episodic memory buffer where agents verbally reflect on task feedback to improve subsequent trial decisions. The key difference is that ReasoningBank emphasizes extracting generalizable high-level reasoning patterns and integrates test-time scaling, whereas Reflexion focuses on self-reflection through verbal feedback incorporation across diverse task types.

---

2. R2D2: Remembering, Replaying and Dynamic Decision Making with a Reflective Agentic Memory

Authors: Tenghao Huang, Kinjal Basu, Ibrahim Abdelaziz, Pavan Kapanipathi, Jonathan May, et al. (7 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

The proliferation of web agents necessitates advanced navigation and interaction strategies within complex web environments. Current models often struggle with efficient navigation and action execution due to limited visibility and understanding of web structures. Our proposed R2D2 framework addresses these challenges by integrating two paradigms: Remember and Reflect. The Remember paradigm uses a replay buffer that aids agents in reconstructing the web environment dynamically, thus enabling the...

Relationship Analysis

Both papers belong to the Reflection and Verbal Reinforcement Learning category, focusing on agents that learn through linguistic feedback and self-reflection without weight updates. They overlap in using memory mechanisms to store and retrieve past experiences for improving agent decision-making in interactive environments (web browsing tasks). However, ReasoningBank extracts high-level reasoning strategies from both successful and failed experiences into structured memory items, while R2D2 constructs a replay buffer as a directed graph of visited webpages and uses A* search for navigation, combining this with reflection on execution errors.

---

3. MetaReflection: Learning Instructions for Language Agents using Past Reflections

Authors: Gupta, Priyanshu, Kirtania, Shashank, Shi, et al. (12 authors total) | **Year/Venue**: 2024 • Conference on Empirical Methods in Natural Language Processing | **URL**: View paper

Abstract

The popularity of Large Language Models (LLMs) have unleashed a new age of Language Agents for solving a diverse range of tasks. While contemporary frontier LLMs are capable enough to power reasonably good Language agents, the closed-API model makes it hard to improve in cases they perform sub-optimally. To address this, recent works have explored techniques to improve their performance using self reflection and prompt optimization techniques. While techniques like self reflection work well in a...

Relationship Analysis

Both papers belong to the Reflection and Verbal Reinforcement Learning category, focusing on agents that learn through linguistic feedback and self-reflection without weight updates. They overlap in using memory mechanisms to store experiential learnings from past agent interactions to improve future performance. The key difference is that ReasoningBank extracts high-level reasoning strategies from both successful and failed experiences and integrates memory-aware test-time scaling, while MetaReflection uses offline reinforcement learning to augment semantic memory based on past reflections and focuses on prompt optimization across simpler task domains.

Contributions Analysis

Overall novelty summary. The paper proposes ReasoningBank, a memory framework enabling agents to self-curate generalizable reasoning strategies from both successful and failed experiences, and introduces memory-aware test-time scaling (MaTTS) to accelerate learning. Within the taxonomy, it resides in the 'Reflection and Verbal Reinforcement Learning' leaf under 'Memory-Based Learning

Mechanisms', alongside three sibling papers. This leaf represents a moderately populated research direction focused on linguistic feedback and self-reflection without parametric updates, situated within a broader branch containing four distinct learning paradigms across the field's 50 papers.

The taxonomy reveals neighboring research directions that contextualize this work's positioning. Adjacent leaves include 'Self-Evolving and Lifelong Learning Agents' emphasizing continuous capability improvement, 'Experience Replay and Trajectory Synthesis' leveraging stored trajectories for sample efficiency, and 'Cross-Domain Experience Sharing' enabling knowledge transfer across tasks. The scope note for the paper's leaf explicitly excludes parametric weight updates, distinguishing reflection-based approaches from gradient-driven methods. This boundary clarifies that ReasoningBank operates through memory curation rather than model fine-tuning, connecting it to verbal reinforcement paradigms while diverging from replay-based learning mechanisms in neighboring leaves.

Among 27 candidates examined across three contributions, the ReasoningBank framework shows one refutable candidate out of 10 examined, suggesting some overlap with prior memory architectures. The MaTTS contribution examined 7 candidates with none refutable, indicating relatively sparser prior work on memory-guided test-time scaling. The third contribution on memory-driven experience as a scaling dimension examined 10 candidates with none refutable, suggesting this framing may be less explored. The limited search scope means these statistics reflect top semantic matches rather than exhaustive coverage, and the single refutable case for ReasoningBank indicates at least one prior work addresses similar memory curation concepts within the examined set.

Based on the top-27 semantic matches examined, the work appears to introduce novel combinations of memory curation with test-time scaling, though the ReasoningBank framework itself shows some overlap with existing memory architectures. The analysis covers semantically proximate papers but does not guarantee exhaustive field coverage, particularly for works using different terminology or published in specialized venues. The taxonomy positioning suggests the paper bridges reflection-based learning with scaling paradigms, occupying a moderately explored niche within the broader agent learning landscape.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: ReasoningBank memory framework

Description: A memory framework that distills high-level reasoning strategies from both successful and failed agent experiences into structured, reusable memory items (with title, description, and content), enabling agents to generalize across tasks and evolve over time rather than storing only raw trajectories or successful routines.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Curating Demonstrations using Online Experience

URL: View paper

Brief Assessment

Online Experience Curation[58] focuses on filtering robot demonstration datasets using online experience to identify successful vs. unsuccessful policy rollouts, not on distilling high-level reasoning strategies from agent experiences into structured memory items for generalization across tasks.

---

2. MAPLE: Multi-Agent Adaptive Planning with Long-Term Memory for Table Reasoning

URL: View paper

Brief Assessment

MAPLE[59] focuses on table-based question answering with specialized cognitive agents and experience archiving, while the original paper addresses general agent self-evolution across web browsing and software engineering tasks with distilled reasoning strategies.

---

3. Reasoning, Memorization, and Fine-Tuning Language Models for Non-Cooperative Games

URL: View paper

Brief Assessment

Non-Cooperative Games[62] focuses on game-solving through tree-of-thoughts decomposition and multi-agent collaboration, not on distilling generalizable reasoning strategies from agent experiences into reusable memory items for continuous self-evolution across diverse tasks.

---

4. Xolver: Multi-Agent Reasoning with Holistic Experience Learning Just Like an Olympiad Team

URL: View paper

Brief Assessment

Xolver[2] focuses on multi-agent reasoning with holistic experience learning across diverse modalities (external retrieval, tool use, collaborative interactions), while the original paper's ReasoningBank specifically distills high-level reasoning strategies into structured memory items with title/description/content from agent trajectories. The technical approaches and memory schemas differ substantially.

---

5. Table-critic: A multi-agent framework for collaborative criticism and refinement in table reasoning

URL: View paper

Brief Assessment

Table-Critic[64] focuses on multi-agent collaborative criticism and refinement for table reasoning tasks, not on distilling generalizable reasoning strategies from agent experiences into structured memory items for cross-task learning.

---

6. From Experience to Strategy: Empowering LLM Agents with Trainable Graph Memory

URL: View paper

Prior Art Analysis

Graph Memory Strategy[61] demonstrates prior work that distills high-level reasoning strategies from agent experiences into structured, reusable memory items. Both papers propose frameworks that abstract raw trajectories into structured, human-interpretable strategic knowledge that can be reused across tasks. Graph Memory Strategy[61] explicitly describes distilling experiences 'into high-level, human-interpretable strategic meta-cognition' and making this memory 'adaptable' through optimization, which directly parallels the ORIGINAL paper's claim of distilling 'high-level reasoning strategies from both successful and failed agent experiences into structured, reusable memory items.'

Evidence

Evidence 1 - **Rationale**: Both papers propose novel memory frameworks that distill agent trajectories into high-level, structured strategic knowledge. Graph Memory Strategy[61] explicitly describes distilling trajectories 'into high-level, human-interpretable strategic meta-cognition,' which directly parallels the ORIGINAL paper's claim of curating 'generalizable reasoning strategies' from experiences. - **Original**: we propose reasoning bank , a novel memory framework that allows an agent to self-curate generalizable reasoning strategies from both its successful and failed experiences for future use. this mechanism enables agents to generalize across tasks and become more capable over time. - **Candidate**: we introduce a novel agent-centric, trainable, multi-layered graph memory framework and evaluate how context memory enhances the ability of llms to utilize parametric information. the graph abstracts raw agent trajectories into structured decision paths in a state machine and further distills them i...

Evidence 2 - **Rationale**: Graph Memory Strategy[61] identifies the same problem space as the ORIGINAL paper: the need to move beyond storing raw experiences to creating structured, reusable strategic knowledge. This establishes that the problem and solution direction were already explored in prior work. - **Original**: unlike prior works that primarily store raw experience or successful routines, we propose reasoning bank , a novel memory framework that allows an agent to self-curate generalizable reasoning strategies from both its successful and failed experiences for future use. - **Candidate**: a promising approach for improving the reasoning capabilities of llm agents is to better utilize prior experiences in guiding current decisions. however, llms acquire experience either through implicit memory via training, which suffers from catastrophic forgetting and limited interpretability, or e...

---

7. AgentEvolver: Towards Efficient Self-Evolving Agent System

URL: View paper

Brief Assessment

AgentEvolver[63] focuses on self-evolving mechanisms (self-questioning, self-navigating, self-attributing) for RL-based agent training, not on distilling high-level reasoning strategies from experiences into structured memory items for reuse across tasks.

---

8. Agent kb: Leveraging cross-domain experience for agentic problem solving

URL: View paper

Brief Assessment

Agent KB[18] focuses on cross-framework knowledge transfer through trajectory aggregation and hybrid retrieval APIs, rather than distilling high-level reasoning strategies from individual agent experiences into structured memory items with title/description/content schema.

---

9. SiriuS: Self-improving Multi-agent Systems via Bootstrapped Reasoning

URL: View paper

Brief Assessment

SiriuS[60] focuses on building an experience library of reasoning trajectories for multi-agent system optimization through training, while the original paper's ReasoningBank distills high-level reasoning strategies into structured memory items (with title, description, content) for test-time retrieval and reuse without training. These represent different approaches to learning from experience.

---

10. GTR: Guided Thought Reinforcement Prevents Thought Collapse in RL-based VLM Agent Training

URL: View paper

Brief Assessment

GTR[65] focuses on preventing thought collapse during RL training of VLM agents through process guidance and automated correction, not on building a memory framework that curates generalizable reasoning strategies from past experiences for reuse across tasks.

Contribution 2: Memory-aware test-time scaling (MaTTS)

Description: A test-time scaling approach that creates bidirectional synergy between memory and scaling: memory guides scaling toward more promising explorations, while diverse rollouts from scaling provide rich contrastive signals for higher-quality memory curation in both parallel and sequential settings.

This contribution was assessed against **7 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Style-Adaptive Detection Transformer for Single-Source Domain Generalized Object Detection

URL: View paper

Brief Assessment

Style-Adaptive Detection[54] focuses on object detection with domain generalization using a memory bank for style adaptation, not on test-time scaling for agent systems with memory-guided exploration and contrastive learning from diverse rollouts.

---

2. Incremental model enhancement via memory-based contrastive learning

URL: View paper

Brief Assessment

Memory-Based Contrastive[51] focuses on domain adaptation during inference time for image classification tasks, not on agent systems with memory-guided test-time scaling for interactive decision-making tasks.

---

3. BECLR: Batch Enhanced Contrastive Few-Shot Learning

URL: View paper

Brief Assessment

BECLR[53] focuses on unsupervised few-shot learning with contrastive methods and optimal transport for sample bias, not on test-time scaling with memory guidance for interactive agent tasks.

---

4. Contrastive Test-Time Adaptation

URL: View paper

Brief Assessment

Contrastive Test-Time[57] focuses on test-time adaptation for domain shift without test-time scaling or memory-guided exploration. The original paper's MaTTS creates synergy between memory and scaling with diverse rollouts, which is not addressed in this candidate.

---

5. Demystifying Diffusion Policies: Action Memorization and Simple Lookup Table Alternatives

URL: View paper

Brief Assessment

Diffusion Policies[52] focuses on robot manipulation policies that memorize action sequences from demonstrations, not on test-time scaling approaches that use memory to guide exploration with contrastive signals from diverse rollouts in agent reasoning tasks.

---

6. SoftCoT++: Test-Time Scaling with Soft Chain-of-Thought Reasoning

URL: View paper

Brief Assessment

SoftCoT++[56] focuses on test-time scaling through continuous latent space perturbation and contrastive learning for diverse reasoning paths, without any memory mechanism. The original paper's MaTTS creates bidirectional synergy between memory and scaling, which is fundamentally different from SoftCoT++'s approach of perturbing latent thoughts via specialized tokens.

---

7. Multi-camera spatiotemporal deep learning framework for real-time abnormal behavior detection in dense urban environments

URL: View paper

Brief Assessment

Spatiotemporal Abnormal Detection[55] focuses on multi-camera surveillance for abnormal behavior detection in urban environments using graph attention networks and reinforcement learning for camera optimization. It does not address test-time scaling with memory guidance or contrastive signals from diverse rollouts in LLM agent systems.

Contribution 3: Memory-driven experience as a new scaling dimension

Description: The work establishes memory-driven experience as a novel dimension for test-time scaling in agent systems, demonstrating that agents can develop increasingly complex emergent reasoning strategies and self-evolving capabilities through the interaction of memory and scaling.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. VLM Can Be a Good Assistant: Enhancing Embodied Visual Tracking with Self-Improving Vision-Language Models

URL: View paper

Brief Assessment

Self-Improving VLM[74] focuses on embodied visual tracking with memory-augmented self-reflection for failure recovery in robotic tracking tasks, not on establishing memory-driven experience as a general test-time scaling dimension for agent systems across diverse reasoning tasks.

---

2. Memory Sharing for Large Language Model based Agents

URL: View paper

Brief Assessment

Memory Sharing[72] focuses on memory sharing among multiple agents for in-context learning enhancement, not on test-time scaling or emergent reasoning strategies through memory-experience interaction as a scaling dimension.

---

3. SAGE: Self-evolving Agents with Reflective and Memory-augmented Abilities

URL: View paper

Brief Assessment

SAGE[13] focuses on memory optimization using the Ebbinghaus forgetting curve with three collaborative agents. The candidate does not demonstrate that memory-driven experience as a scaling dimension for emergent reasoning was previously established.

---

4. Agentnet: Decentralized evolutionary coordination for llm-based multi-agent systems

URL: View paper

Brief Assessment

AgentNet[70] focuses on decentralized multi-agent coordination with retrieval-based memory for expertise refinement in distributed systems, not on test-time scaling through memory-driven experience accumulation in single-agent sequential learning scenarios.

---

5. Alita-g: Self-evolving generative agent for agent generation

URL: View paper

Brief Assessment

Alita-G[71] focuses on generating and curating domain-specific MCP tools from successful trajectories, not on memory-driven experience as a scaling dimension for emergent reasoning capabilities.

---

6. Chemagent: Self-updating memories in large language models improves chemical reasoning

URL: View paper

Brief Assessment

ChemAgent[73] focuses on a static library system for chemical reasoning tasks, where memory is constructed from a development set and used for retrieval during inference. This differs fundamentally from the original paper's concept of memory-driven experience as a test-time scaling dimension where agents continuously evolve through interaction and scaling generates diverse experiences for memory curation.

---

7. Evo-Memory: Benchmarking LLM Agent Test-time Learning with Self-Evolving Memory

URL: View paper

Brief Assessment

Evo-Memory[66] focuses on evaluating self-evolving memory in streaming task environments, not on establishing memory-driven experience as a test-time scaling dimension with emergent reasoning strategies as the original paper does.

---

8. A survey of self-evolving agents: On path to artificial super intelligence

URL: View paper

Brief Assessment

Self-Evolving Survey[68] is a survey paper that reviews self-evolving agents broadly, including memory mechanisms as one component. It does not present original empirical work demonstrating memory-driven experience as a scaling dimension with emergent reasoning strategies, which is the specific novel claim of the original paper.

---

9. MetaAgent: Toward Self-Evolving Agent via Tool Meta-Learning

URL: View paper

Brief Assessment

MetaAgent[67] focuses on meta tool learning through self-reflection and experience distillation for knowledge discovery tasks, not on establishing memory-driven experience as a test-time scaling dimension with emergent reasoning strategies as described in the original paper.

10. G-Memory: Tracing Hierarchical Memory for Multi-Agent Systems

URL: View paper

Brief Assessment

G-Memory[69] focuses on hierarchical memory architectures for multi-agent collaboration trajectories, not on test-time scaling or emergent reasoning through memory-scaling synergy as the original paper proposes.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] ReasoningBank: Scaling Agent Self-Evolving with Reasoning Memory View paper
- [1] Reflexion: language agents with verbal reinforcement learning View paper
- [2] Xolver: Multi-Agent Reasoning with Holistic Experience Learning Just Like an Olympiad Team View paper
- [3] Building self-evolving agents via experience-driven lifelong learning: A framework and benchmark View paper
- [4] Memp: Exploring agent procedural memory View paper
- [5] Cognitive architectures for language agents View paper
- [6] Enhancing reasoning with collaboration and memory View paper
- [7] Constructing future behavior in the hippocampal formation through composition and replay View paper
- [8] Decentralizing AI Memory: SHIMI, a Semantic Hierarchical Memory Index for Scalable Agent Reasoning View paper
- [9] AriGraph: Learning Knowledge Graph World Models with Episodic Memory for LLM Agents View paper
- [10] ReaGAN: Node-as-Agent-Reasoning Graph Agentic Network View paper
- [11] How memory management impacts llm agents: An empirical study of experience-following behavior View paper
- [12] Symagent: A neural-symbolic self-learning agent framework for complex reasoning over knowledge graphs View paper
- [13] SAGE: Self-evolving Agents with Reflective and Memory-augmented Abilities View paper
- [14] GuardAgent: Safeguard LLM Agents via Knowledge-Enabled Reasoning View paper
- [15] Seeing, listening, remembering, and reasoning: A multimodal agent with long-term memory View paper
- [16] Agent s: An open agentic framework that uses computers like a human View paper
- [17] " my agent understands me better": Integrating dynamic human-like memory recall and consolidation in llm-based agents View paper
- [18] Agent kb: Leveraging cross-domain experience for agentic problem solving View paper
- [19] KG-Agent: An Efficient Autonomous Agent Framework for Complex Reasoning over Knowledge Graph View paper
- [20] Review of Case-Based Reasoning for LLM Agents: Theoretical Foundations, Architectural Components, and Cognitive Integration View paper
- [21] MC-GPT: Empowering Vision-and-Language Navigation with Memory Map and Reasoning Chains View paper
- [22] Goal-Directed Search Outperforms Goal-Agnostic Memory Compression in Long-Context Memory Tasks View paper
- [23] ExpReS-VLA: Specializing Vision-Language-Action Models Through Experience Replay and Retrieval View paper
- [24] LiCoMemory: Lightweight and Cognitive Agentic Memory for Efficient Long-Term Reasoning View paper
- [25] MEM1: Learning to Synergize Memory and Reasoning for Efficient Long-Horizon Agents View paper
- [26] Long-horizon Reasoning Agent for Olympiad-Level Mathematical Problem Solving View paper
- [27] Reinforcement learning and episodic memory in humans and animals: an integrative framework View paper
- [28] ERCI: An Explainable Experience Replay Approach with Causal Inference for Deep Reinforcement Learning View paper
- [29] Constructing future behaviour in the hippocampal formation through composition and replay View paper
- [30] Smarter Together: Creating Agentic Communities of Practice through Shared Experiential Learning View paper
- [31] Agenttrek: Agent trajectory synthesis via guiding replay with web tutorials View paper
- [32] EndoAgent: A Memory-Guided Reflective Agent for Intelligent Endoscopic Vision-to-Decision Reasoning View paper
- [33] R2D2: Remembering, Replaying and Dynamic Decision Making with a Reflective Agentic Memory View paper
- [34] MetaReflection: Learning Instructions for Language Agents using Past Reflections View paper
- [35] AgenticSciML: Collaborative Multi-Agent Systems for Emergent Discovery in Scientific Machine Learning View paper
- [36] From language to action: A review of large language models as autonomous agents and tool users View paper
- [37] From Knowledge to Noise: CTIM-Rover and the Pitfalls of Episodic Memory in Software Engineering Agents View paper
- [38] Towards Neurocognitive-Inspired Intelligence: From AI's Structural Mimicry to Human-Like Functional Cognition View paper
- [39] DeepAgent: A General Reasoning Agent with Scalable Toolsets View paper
- [40] Know rob 2.0â□□a 2nd generation knowledge processing framework for cognition-enabled robotic agents View paper
- [41] Transdreamer: Reinforcement learning with transformer world models View paper
- [42] Efficient episodic memory utilization of cooperative multi-agent reinforcement learning View paper
- [43] Visuomotor navigation for embodied robots with spatial memory and semantic reasoning cognition View paper
- [44] AGENT KB: A Hierarchical Memory Framework for Cross-Domain Agentic Problem Solving View paper
- [45] MemVerse: Multimodal Memory for Lifelong Learning Agents View paper
- [46] LLMs Are Not Good Strategists, Yet Memory-Enhanced Agency Boosts Reasoning View paper
- [47] Sample-Efficient Online Learning in LM Agents via Hindsight Trajectory Rewriting View paper
- [48] A survey on agentic multimodal large language models View paper
- [49] Towards combining commonsense reasoning and knowledge acquisition to guide deep learning View paper
- [50] Memory-Augmented Agent Training for Business Document Understanding View paper
- [51] Incremental model enhancement via memory-based contrastive learning View paper
- [52] Demystifying Diffusion Policies: Action Memorization and Simple Lookup Table Alternatives View paper
- [53] BECLR: Batch Enhanced Contrastive Few-Shot Learning View paper
- [54] Style-Adaptive Detection Transformer for Single-Source Domain Generalized Object Detection View paper
- [55] Multi-camera spatiotemporal deep learning framework for real-time abnormal behavior detection in dense urban environments View paper

- [56] SoftCoT++: Test-Time Scaling with Soft Chain-of-Thought Reasoning View paper
- [57] Contrastive Test-Time Adaptation View paper
- [58] Curating Demonstrations using Online Experience View paper
- [59] MAPLE: Multi-Agent Adaptive Planning with Long-Term Memory for Table Reasoning View paper
- [60] SiriuS: Self-improving Multi-agent Systems via Bootstrapped Reasoning View paper
- [61] From Experience to Strategy: Empowering LLM Agents with Trainable Graph Memory View paper
- [62] Reasoning, Memorization, and Fine-Tuning Language Models for Non-Cooperative Games View paper
- [63] AgentEvolver: Towards Efficient Self-Evolving Agent System View paper
- [64] Table-critic: A multi-agent framework for collaborative criticism and refinement in table reasoning View paper
- [65] GTR: Guided Thought Reinforcement Prevents Thought Collapse in RL-based VLM Agent Training View paper
- [66] Evo-Memory: Benchmarking LLM Agent Test-time Learning with Self-Evolving Memory View paper
- [67] MetaAgent: Toward Self-Evolving Agent via Tool Meta-Learning View paper
- [68] A survey of self-evolving agents: On path to artificial super intelligence View paper
- [69] G-Memory: Tracing Hierarchical Memory for Multi-Agent Systems View paper
- [70] Agentnet: Decentralized evolutionary coordination for llm-based multi-agent systems View paper
- [71] Alita-g: Self-evolving generative agent for agent generation View paper
- [72] Memory Sharing for Large Language Model based Agents View paper
- [73] Chemagent: Self-updating memories in large language models improves chemical reasoning View paper
- [74] VLM Can Be a Good Assistant: Enhancing Embodied Visual Tracking with Self-Improving Vision-Language Models View paper