Novelty Assessment Report

Paper: MemAgent: Reshaping Long-Context LLM with Multi-Conv RL-based Memory Agent
PDF URL: https://openreview.net/pdf?id=k5nIOvYGCL
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2026-01-01

Abstract

Despite improvements by length extrapolation, efficient attention and memory modules, handling infinitely long documents without performance degradation during extrapolation remains the ultimate challenge in long-text processing. To solve this problem, We introduce a novel agent workflow, \method, which processes text in segments and updates memory through an overwrite strategy, addressing the challenge of long-context task through enhanced memory management. We further extend the DAPO algorithm to directly optimize memory ability in an end-to-end fashion, facilitating training via independent-context multi-conversation generation. Experimental results demonstrate that MemAgent has superb long-context capabilities, being able to extrapolate from an 8K context to a 3.5M QA task with a performance loss of less than 10\% and achieving over 95\% on the 512K NIAH test.

Core Task Landscape

This paper addresses: **Long-Context Language Model Processing with Memory Management**

A total of **50 papers** were analyzed and organized into a taxonomy with **37 categories**.

Taxonomy Overview

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

- **Memory Architecture and Representation**
- **Attention Mechanisms and Architectural Alternatives**
- **KV Cache Management and Optimization**
- **Context Compression and Encoding**
- **Retrieval-Augmentation and External Knowledge**
- **Training and Data Strategies for Long Context**
- **Dynamic Chunking and Segmentation**
- **Query-Aware and Selective Processing**
- **Agent-Based Long-Context Processing**
- **Evaluation and Benchmarking**
- ... and 3 more categories

Complete Taxonomy Tree

- Long-Context Language Model Processing with Memory Management Survey Taxonomy
- Memory Architecture and Representation
  - Neurobiologically-Inspired Memory Systems (3 papers)
  - [1] Hipporag: Neurobiologically inspired long-term memory for large language models (Yu Gu, 2024) View paper
  - [39] HEMA: A Hippocampus-Inspired Extended Memory Architecture for Long-Context AI Conversations (Ahn, 2025) View paper
  - [47] Human-inspired episodic memory for infinite context LLMs (Z. Fountas, 2025) View paper
  - Structured Memory Modules (4 papers)
  - [6] Memorybank: Enhancing large language models with long-term memory (Qiqi Gao, 2024) View paper
  - [11] Structured Memory Mechanisms for Stable Context Representation in Large Language Models (Xing Yue, 2025) View paper
  - [12] Cognitive memory in large language models (Shan LianLei, 2025) View paper
  - [23] Memory-Augmented Architecture for Long-Term Context Handling in Large Language Models (Usama, 2025) View paper
  - Parameter-Based Memory Encoding (2 papers)
  - [31] RMem: Bridging Memory Retention and Retrieval via Reversible Compression (X Wang, 2025) View paper
  - [50] LIFT: Improving Long Context Understanding of Large Language Models through Long Input Fine-Tuning (Mao, 2025) View paper
  - Hierarchical and Logarithmic Memory (1 papers)
  - [46] Logarithmic Memory Networks (LMNs): Efficient Long-Range Sequence Modeling for Resource-Constrained Environments (Taha, 2025) View paper
  - Memory Operating Systems (1 papers)
  - [25] Memos: A memory os for ai system (Li Zhiyu, 2025) View paper
- Attention Mechanisms and Architectural Alternatives
  - Linear and Subquadratic Attention (2 papers)
  - [10] Lizard: An Efficient Linearization Framework for Large Language Models (Zhang Rui-yi, 2025) View paper
  - [21] Efficient attention mechanisms for large language models: A survey (Sun Yutao, 2025) View paper
  - Sparse Attention Techniques (2 papers)
  - [9] Efficient streaming language models with attention sinks (Xiao, 2023) View paper

- [36] Memory-efficient Transformers via Top- Attention (A Gupta, 2021) View paper
- State Space Models and Recurrent Architectures (2 papers)
- [5] Retentive network: A successor to transformer for large language models (Sun Yutao, 2023) View paper
- [27] RecurrentGemma: Moving Past Transformers for Efficient Open Language Models (Botev, 2024) View paper
- Hybrid Architectures (3 papers)
- [24] Samba: Simple Hybrid State Space Models for Efficient Unlimited Context Language Modeling (Ren Liliang, 2024) View paper
- [34] InfiniteVL: Synergizing Linear and Sparse Attention for Highly-Efficient, Unlimited-Input Vision-Language Models (Hongyuan Tao, 2025) View paper
- [44] Taipan: Efficient and Expressive State Space Language Models with Selective Attention (C. Nguyen, 2024) View paper
- Attention Recalibration and Shaping (1 papers)
- [14] Neural attention shaping with contextual embedding recalibration in language models (Dorian Osatov, 2024) View paper

- KV Cache Management and Optimization
  - Dynamic KV Cache Eviction and Selection (2 papers)
  - [7] {InfiniGen}: Efficient generative inference of large language models with dynamic {KV} cache management (Lee, 2024) View paper
  - [38] D2O: Dynamic Discriminative Operations for Efficient Long-Context Inference of Large Language Models (Wan, 2024) View paper
  - Memory Paging and Allocation (1 papers)
  - [13] Efficient memory management for large language model serving with pagedattention (Woosuk Kwon, 2023) View paper
  - KV Cache Compression (1 papers)
  - [49] ZeroMerge: Parameter-Free KV Cache Compression for Memory-Efficient Long-Context LLMs (Liu Xin, 2025) View paper
  - Distributed KV Cache Systems (1 papers)
  - [41] Infinite-LLM: Efficient LLM Service for Long Context with DistAttention and Distributed KVCache (Lin Bin, 2024) View paper
  - Incremental and Adaptive Memory Management (1 papers)
  - [45] Memorize Step by Step: Efficient Long-Context Prefilling with Incremental Memory and Decremental Chunk (Zhi-yuan Zeng, 2024) View paper
  - KV Cache Management Surveys (1 papers)
  - [8] A Survey on Large Language Model Acceleration based on KV Cache Management (Li, 2024) View paper
- Context Compression and Encoding
  - Autoencoder-Based Context Compression (1 papers)
  - [16] In-context autoencoder for context compression in a large language model (Ge, 2023) View paper
  - Modality-Based Context Encoding (1 papers)
  - [42] Dolphin: Long Context as a New Modality for Energy-Efficient On-Device Language Models (Chen Wei, 2024) View paper
- Retrieval-Augmentation and External Knowledge
  - Retrieval vs. Long-Context Comparison (1 papers)
  - [2] Retrieval meets long context large language models (Xu Peng, 2023) View paper
  - Training-Free Context Memory (2 papers)
  - [18] Infllm: Training-free long-context extrapolation for llms with an efficient context memory (Xu Han, 2024) View paper
  - [33] InfLLM: Unveiling the Intrinsic Capacity of LLMs for Understanding Extremely Long Sequences with Training-Free Memory (Xiao, 2024) View paper
  - Long-Term Memory Augmentation (1 papers)
  - [43] Augmenting language models with long-term memory (Wang, 2023) View paper
- Training and Data Strategies for Long Context
  - Long-Context Training Recipes (1 papers)
  - [4] Longrecipe: Recipe for efficient long context generalization in large language models (Zhiyuan Hu, 2025) View paper
  - Long-Dependency Data Mining (1 papers)
  - [22] Long context is not long at all: A prospector of long-dependency data for large language models (He, 2024) View paper
- Dynamic Chunking and Segmentation (1 papers)
  - [17] Dynamic Chunking and Selection for Reading Comprehension of Ultra-Long Context in Large Language Models (Yao, 2025) View paper
- Query-Aware and Selective Processing (1 papers)
  - [15] Quickllama: Query-aware inference acceleration for large language models (Li, 2025) View paper
- Agent-Based Long-Context Processing
  - Multi-Agent Collaboration for Long Context (1 papers)
  - [32] Long context scaling: Divide and conquer via multi-agent question-driven collaboration (Lin Zi-xin, 2025) View paper
  - Workflow Memory for Agents (1 papers)
  - [40] Agent workflow memory (Mao, 2024) View paper
  - Hierarchical Working Memory for Agents (1 papers)
  - [28] Hiagent: Hierarchical working memory management for solving long-horizon agent tasks with large language model (Mengkang Hu, 2025) View paper
  - Memory-Augmented Multimodal Agents (1 papers)
  - [35] Jarvis-1: Open-world multi-task agents with memory-augmented multimodal language models (Zihao Wang, 2024) View paper
  - RL-Based Memory Agents ★ (1 papers)
  - [0] MemAgent: Reshaping Long-Context LLM with Multi-Conv RL-based Memory Agent (Anon et al., 2026) View paper
  - Context Compression for Agents (1 papers)
  - [37] Acon: Optimizing context compression for long-horizon llm agents (Kang, 2025) View paper
  - Theory of Mind in Multi-Agent Systems (1 papers)
  - [29] Theory of mind for multi-agent collaboration via large language models (Huao Li, 2023) View paper
  - Long-Horizon Task Planning for Agents (1 papers)
  - [48] FLTRNN: Faithful Long-Horizon Task Planning for Robotics with Large Language Models (Jiatao Zhang, 2024) View paper
- Evaluation and Benchmarking (1 papers)
  - [26] Loogle: Can long-context language models understand long contexts? (Jiaqi Li, 2024) View paper

- Surveys and Taxonomies (2 papers)
  - [3] Beyond the limits: A survey of techniques to extend the context length in large language models (Wang Xindi, 2024) View paper
  - [19] Rethinking memory in ai: Taxonomy, operations, topics, and future directions (Du Yiming, 2025) View paper
- Large Memory Models (1 papers)
  - [30] LM2: Large Memory Models for Long Context Reasoning (J Kang, 2025) View paper
- Dynamic Contextual Memory Embedding (1 papers)
  - [20] Enhancing large language models through dynamic contextual memory embedding: A technical evaluation (Igor Dakat, 2024) View paper

Narrative

Core task: Long-context language model processing with memory management. The field addresses how language models can efficiently handle extended sequences that exceed typical context windows, organizing solutions into several major branches. Memory Architecture and Representation explores structured storage mechanisms such as episodic buffers and hierarchical memory systems (e.g., MemoryBank[6], Cognitive Memory[12]). Attention Mechanisms and Architectural Alternatives investigates alternatives to standard attention, including recurrent designs (Retentive Network[5], RecurrentGemma[27]) and hybrid approaches (Samba[24]). KV Cache Management and Optimization focuses on efficient key-value storage strategies (PagedAttention[13], Attention Sinks[9]), while Context Compression and Encoding develops methods to condense long inputs (Context Autoencoder[16]). Retrieval-Augmentation and External Knowledge integrates retrieval systems (HippoRAG[1], Retrieval Long Context[2]), and Training and Data Strategies for Long Context examines how models learn to process extended sequences (LongRecipe[4]). Dynamic Chunking and Segmentation, Query-Aware and Selective Processing, and Agent-Based Long-Context Processing address adaptive strategies for managing context, with the latter branch emphasizing reinforcement learning and agentic workflows.

A particularly active line of work contrasts architectural redesigns—such as recurrent or state-space models that avoid quadratic attention costs—with retrieval-augmented approaches that selectively fetch relevant context. Another tension lies between static compression techniques and dynamic, query-aware selection methods (InfLLM[18], Dynamic Chunking Selection[17]). Within the Agent-Based Long-Context Processing branch, MemAgent[0] employs reinforcement learning to train memory-management policies, positioning itself among works that treat memory as a learnable decision-making process rather than a fixed architectural component. This approach contrasts with retrieval-focused methods like HippoRAG[1], which rely on predefined indexing schemes, and with hierarchical agent frameworks such as HiAgent[28], which decompose tasks across multiple agents. By framing memory operations as RL-optimized actions, MemAgent[0] explores how adaptive policies can balance retention and eviction trade-offs in extended interactions.

Related Works in Same Category

No sibling papers were found in the same taxonomy leaf. A taxonomy-subtopic-level comparison will be produced instead.

Taxonomy-Level Summary

Sibling Subtopics

- **Context Compression for Agents** (leaves: 1, papers: 1)
- Scope: Frameworks optimizing compression of environment observations and action histories to reduce context length in agentic tasks.
- Exclude: Excludes general context compression and memory modules; those belong to Context Compression and Memory Architecture respectively.
- **Hierarchical Working Memory for Agents** (leaves: 1, papers: 1)
- Scope: Hierarchical memory management systems for agents to reduce context length and improve efficiency in long-horizon tasks.
- Exclude: Excludes flat memory and workflow-based memory; those belong to Structured Memory Modules and Workflow Memory respectively.
- **Long-Horizon Task Planning for Agents** (leaves: 1, papers: 1)
- Scope: Frameworks for faithful long-horizon task planning in robotics or other domains using LLMs with context management.
- Exclude: Excludes short-horizon planning and memory-only systems; those belong to other categories or Memory Architecture.
- **Memory-Augmented Multimodal Agents** (leaves: 1, papers: 1)
- Scope: Agents integrating memory modules with multimodal observations for open-world tasks requiring long-term context and embodied control.
- Exclude: Excludes text-only agents and non-embodied systems; those belong to other agent subcategories or Memory Architecture.
- **Multi-Agent Collaboration for Long Context** (leaves: 1, papers: 1)
- Scope: Frameworks using multiple agents with divide-and-conquer strategies or question-driven workflows to process long contexts collaboratively.
- Exclude: Excludes single-agent workflows and memory-augmented agents; those belong to Workflow Memory and Memory-Augmented Agents respectively.
- **Theory of Mind in Multi-Agent Systems** (leaves: 1, papers: 1)
- Scope: Evaluation of LLM-based agents' collaborative behaviors and theory of mind capabilities in multi-agent cooperative tasks.
- Exclude: Excludes single-agent systems and memory-focused studies; those belong to other agent subcategories or Memory Architecture.
- **Workflow Memory for Agents** (leaves: 1, papers: 1)
- Scope: Methods inducing and reusing task workflows or routines from past experiences to guide agent actions in long-horizon tasks.
- Exclude: Excludes multi-agent systems and hierarchical memory; those belong to Multi-Agent Collaboration and Hierarchical Working Memory respectively.

Contributions Analysis

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

Contribution 1: MEMAGENT agent workflow with overwrite-based memory management

Description: The authors propose MEMAGENT, a new agent workflow that handles long-context tasks by dividing documents into segments and iteratively updating a fixed-length memory using an overwrite strategy. This approach enables processing of arbitrarily long texts with linear time complexity while maintaining performance.

This contribution was assessed against **9 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Evaluating memory in llm agents via incremental multi-turn interactions

URL: View paper

Brief Assessment

Incremental Multi-Turn Evaluation[51] focuses on benchmarking and evaluating memory competencies in agents through multi-turn interactions, not on proposing a specific agent workflow with overwrite-based memory management for processing long texts.

---

2. Hiagent: Hierarchical working memory management for solving long-horizon agent tasks with large language model

URL: View paper

Brief Assessment

HiAgent[28] focuses on hierarchical working memory management using subgoals as memory chunks for long-horizon tasks, whereas MEMAGENT processes documents in segments with fixed-length memory updates. The architectural approaches and application domains differ fundamentally.

---

3. Step up your game: A research on two/multi-step summarisation of long, regulatory documents

URL: View paper

Brief Assessment

Multi-Step Summarisation[57] focuses on two/multi-step extractive-abstractive summarization of regulatory documents, not on agent workflows with overwrite-based memory management for long-context processing. The candidate addresses document summarization pipelines rather than memory-based agent architectures.

---

4. Dialog Generation Using Multi-Turn Reasoning Neural Networks

URL: View paper

Brief Assessment

Multi-Turn Reasoning Networks[55] focuses on dialog generation using multi-turn reasoning over conversation memories for response generation, not on processing arbitrarily long documents with overwrite-based memory management for long-context tasks.

---

5. Memory-Augmented Agent Training for Business Document Understanding

URL: View paper

Brief Assessment

Memory Augmented Training[54] focuses on business document understanding with domain-specific memory refinement for invoice processing, not general long-context handling with overwrite-based memory strategies for arbitrary text lengths.

---

6. Agentic Troubleshooting Guide Automation for Incident Management

URL: View paper

Brief Assessment

Agentic Troubleshooting Guide[53] focuses on automating troubleshooting guides for incident management using DAG-based execution and query preparation plugins, not on long-context processing with overwrite-based memory management for arbitrary-length text.

---

7. State and Memory is All You Need for Robust and Reliable AI Agents

URL: View paper

Brief Assessment

State and Memory[52] uses finite-state automata (FSA) for persistent state tracking in scientific workflows, not segment-based overwrite memory for long-context text processing. The technical approaches and application domains differ fundamentally.

---

8. DelTA: An Online Document-Level Translation Agent Based on Multi-Level Memory

URL: View paper

Brief Assessment

DelTA[56] focuses on document-level translation with multi-level memory (proper noun records, bilingual summary, long-term/short-term memory) for sentence-by-sentence translation, while MEMAGENT addresses general long-context tasks using a single fixed-length memory with overwrite strategy for arbitrary text processing.

---

9. Long context scaling: Divide and conquer via multi-agent question-driven collaboration

URL: View paper

Brief Assessment

Divide and Conquer[32] focuses on dynamic partitioning with question-driven workflows and selective partition replay for handling inverted-order structures. It does not use an overwrite-based memory strategy where a fixed-length memory is iteratively updated by overwriting previous content, which is the core mechanism of MEMAGENT.

Contribution 2: Multi-Conv DAPO algorithm for end-to-end memory optimization

Description: The authors extend the DAPO reinforcement learning algorithm to create Multi-Conv DAPO, which optimizes memory capabilities end-to-end by treating each context-independent conversation as an optimization objective. This enables training of agent workflows with multiple rounds of memory updates across independent contexts.

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. In prospect and retrospect: Reflective memory management for long-term personalized dialogue agents

URL: View paper

Brief Assessment

Reflective Memory Management[62] focuses on long-term personalized dialogue agents with prospective and retrospective reflection mechanisms for memory management, not on multi-conversation RL optimization for memory in independent contexts as described in the original paper's Multi-Conv DAPO contribution.

---

2. Doctoragent-rl: A multi-agent collaborative reinforcement learning system for multi-turn clinical dialogue

URL: View paper

Brief Assessment

DoctorAgent-RL[59] focuses on multi-turn clinical dialogue optimization through doctor-patient interactions, not on memory management across independent contexts. The RL framework optimizes questioning strategies for medical consultations rather than memory capabilities for long-context processing.

---

3. Planning without Search: Refining Frontier LLMs with Offline Goal-Conditioned RL

URL: View paper

Brief Assessment

Planning without Search[63] focuses on goal-conditioned value functions for multi-turn planning in interactive tasks, not on memory optimization across independent contexts. The candidate trains value functions over reasoning steps for decision-making, while the original extends DAPO for memory updates across context-independent conversations.

---

4. Secom: On memory construction and retrieval for personalized conversational agents

URL: View paper

Brief Assessment

Secom[58] focuses on memory construction and retrieval for conversational agents through segmentation and compression-based denoising, not on multi-conversation reinforcement learning optimization. The candidate does not address end-to-end RL training across independent contexts.

---

5. RAIDEN-R1: Improving Role-awareness of LLMs via GRPO with Verifiable Reward

URL: View paper

Brief Assessment

RAIDEN-R1[65] focuses on role-playing conversational agents with role-awareness rewards, not memory optimization across multi-conversation contexts. The technical domains are distinct.

---

6. History-Aware Cross-Attention Reinforcement: Self-Supervised Multi Turn and Chain-of-Thought Fine-Tuning with vLLM

URL: View paper

Brief Assessment

History-Aware Cross-Attention[66] focuses on cross-attention mechanisms in vLLM for multi-turn dialogue and chain-of-thought reasoning, not on memory optimization through multi-conversation DAPO reinforcement learning as described in the original paper.

---

7. Context-lite multi-turn reinforcement learning for LLM agents

URL: View paper

Brief Assessment

Context-Lite RL[64] focuses on multi-turn RL with customizable memory mechanisms and dual-discounting GAE for agent tasks, not on extending DAPO for end-to-end memory optimization across independent contexts as described in the original paper's Multi-Conv DAPO contribution.

---

8. Memory-T1: Reinforcement Learning for Temporal Reasoning in Multi-session Agents

URL: View paper

Brief Assessment

Memory-T1[67] focuses on temporal reasoning in multi-session dialogues with time-aware memory selection, while the original paper addresses general long-context processing through segment-based memory updates. The candidate's RL approach optimizes temporal consistency and evidence grounding for dialogue-specific tasks, not the multi-conversation independent context optimization described in the original contribution.

---

9. High-Resolution Visual Reasoning via Multi-Turn Grounding-Based Reinforcement Learning

URL: View paper

Brief Assessment

Visual Reasoning RL[61] focuses on visual grounding in multi-modal models through iterative image cropping, not memory optimization across independent text contexts. The multi-turn framework serves a different purpose (visual region selection) than MemAgent's memory management.

---

10. Experience replay-based deep reinforcement learning for dialogue management optimisation

URL: View paper

Brief Assessment

Experience Replay Dialogue[60] focuses on dialogue policy optimization in task-oriented spoken dialogue systems using A2CER (Advantage Actor-Critic Experience Replay), not on memory optimization across multiple independent conversations for long-context processing.

Contribution 3: RL-based approach for dynamically updated fixed-length memory in LLMs

Description: The authors introduce a reinforcement learning method that enables LLMs to maintain and update a fixed-length memory dynamically as they process text segment-by-segment. This allows the model to handle arbitrary text lengths while maintaining linear time complexity during processing.

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. Cognitive Architectures for Tomorrow: A Comprehensive Survey of Memory Management Paradigms in Agentic AI Systems

URL: View paper

Brief Assessment

Cognitive Architectures Survey[76] discusses context windows as short-term memory buffers and mentions deep reinforcement learning from human preferences, but does not present a specific RL method for dynamically updating fixed-length memory in LLMs for processing arbitrary text lengths with linear complexity.

---

2. Enhancing large language models through dynamic contextual memory embedding: A technical evaluation

URL: View paper

Brief Assessment

Dynamic Contextual Memory[20] focuses on contextual memory embedding techniques for language models, but the provided context is too limited to assess whether it addresses RL-based dynamic memory updates or fixed-length memory mechanisms similar to the original paper's approach.

---

3. Monte Carlo Planning with Large Language Model for Text-Based Game Agents

URL: View paper

Brief Assessment

Monte Carlo Planning[69] focuses on text-based game agents using MCTS with LLMs and memory mechanisms for game planning, not on general RL frameworks for processing arbitrary-length documents with fixed-length memory as in the original paper.

---

4. Large language models are semi-parametric reinforcement learning agents

URL: View paper

Brief Assessment

Semi-Parametric RL[71] focuses on using RL to update an external experience memory for decision-making tasks (e.g., WebShop, WikiHow), not on dynamically managing fixed-length memory for processing arbitrary-length text segments with linear complexity as in the original paper.

---

5. Amago-2: Breaking the multi-task barrier in meta-reinforcement learning with transformers

URL: View paper

Brief Assessment

Amago-2[75] focuses on meta-reinforcement learning with transformers for multi-task adaptation in control environments (e.g., robotics, video games), not on enabling LLMs to process text segments with dynamically updated memory for language understanding tasks.

---

6. Semantic coherence dynamics in large language models through layered syntax-aware memory retention mechanism

URL: View paper

Brief Assessment

Layered Syntax Memory[73] focuses on syntax-aware memory retention mechanisms rather than RL-based dynamic memory updates for handling arbitrary text lengths. The candidate's approach differs fundamentally in both methodology and application domain.

---

7. CtrlDiff: Boosting Large Diffusion Language Models with Dynamic Block Prediction and Controllable Generation

URL: View paper

Brief Assessment

CtrlDiff[72] uses RL to determine dynamic block sizes for diffusion-based generation, not for managing fixed-length memory in sequential text processing. The technical approaches and objectives differ fundamentally.

---

8. On-Device Large Language Models: A Survey of Model Compression and System Optimization

URL: View paper

Brief Assessment

On-Device LLM Survey[74] focuses on model compression techniques (quantization, pruning, distillation, low-rank factorization) and system optimizations for edge deployment. It does not address reinforcement learning methods for dynamic memory management in long-context processing.

---

9. Look back to reason forward: Revisitable memory for long-context llm agents

URL: View paper

Prior Art Analysis

Revisitable Memory[68] demonstrates that RL-based approaches for dynamically updated memory in LLMs were already explored before the original paper. Both papers employ reinforcement learning to train memory-augmented agents that process long documents by updating a memory corpus during document scanning. The candidate paper explicitly describes using 'reinforcement learning with multi-level rewards (rlmlr)' to train a 'memory-augmented agent' that maintains a 'memory corpus that is dynamically updated during a single-pass document scan.' This directly overlaps with the original paper's core contribution of using RL to enable 'dynamically updated fixed-length memory' for handling long contexts.

Evidence

Evidence 1 - **Rationale**: Both papers describe using dynamically updated memory during document processing. The candidate explicitly identifies this as existing work in the field, suggesting prior art exists for this approach. - **Original**: we propose a novel use of reinforcement learning (rl) to equip llms with a dynamically updated fixed-length 'memory', as illustrated in figure 2. during inference, the llm processes the input text segment-by-segment. as it reads each segment, the model proactively and selectively updates the memory - **Candidate**: existing works equip large language models with a memory corpus that is dynamically updated during a single-pass document scan, also known as the"memorize while reading"methods

Evidence 2 - **Rationale**: Both papers use reinforcement learning to optimize memory updates for question-answering tasks. The candidate's RL approach with rewards for memory management parallels the original's RL-based memory optimization. - **Original**: by viewing memory update in context processing for answer-generation tasks as part of the policy to be optimized by rl, we adopt the rlvr recipe (openai, 2024; guo et al., 2025; seed et al., 2025) to train memagent - **Candidate**: to further strengthen training, we propose reinforcement learning with multi-level rewards (rlmlr), which combines final-answer rewards with dense, step-level signals that guide effective memory use

---

10. A Category-Theoretic Framework for Wake-Sleep Consolidation in Dual-Transformer Architectures

URL: View paper

Brief Assessment

Wake-Sleep Consolidation[70] focuses on category-theoretic frameworks for dual-transformer architectures with wake-sleep consolidation mechanisms. The candidate's approach involves short-term and long-term transformers with different plasticity mechanisms, which is architecturally distinct from MemAgent's segment-based memory overwrite strategy using standard RL on single LLMs.

Appendix: Text Similarity Detection

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

References

- [0] MemAgent: Reshaping Long-Context LLM with Multi-Conv RL-based Memory Agent View paper
- [1] Hipporag: Neurobiologically inspired long-term memory for large language models View paper

- [2] Retrieval meets long context large language models View paper
- [3] Beyond the limits: A survey of techniques to extend the context length in large language models View paper
- [4] Longrecipe: Recipe for efficient long context generalization in large language models View paper
- [5] Retentive network: A successor to transformer for large language models View paper
- [6] Memorybank: Enhancing large language models with long-term memory View paper
- [7] {InfiniGen}: Efficient generative inference of large language models with dynamic {KV} cache management View paper
- [8] A Survey on Large Language Model Acceleration based on KV Cache Management View paper
- [9] Efficient streaming language models with attention sinks View paper
- [10] Lizard: An Efficient Linearization Framework for Large Language Models View paper
- [11] Structured Memory Mechanisms for Stable Context Representation in Large Language Models View paper
- [12] Cognitive memory in large language models View paper
- [13] Efficient memory management for large language model serving with pagedattention View paper
- [14] Neural attention shaping with contextual embedding recalibration in language models View paper
- [15] Quickllama: Query-aware inference acceleration for large language models View paper
- [16] In-context autoencoder for context compression in a large language model View paper
- [17] Dynamic Chunking and Selection for Reading Comprehension of Ultra-Long Context in Large Language Models View paper
- [18] Infllm: Training-free long-context extrapolation for llms with an efficient context memory View paper
- [19] Rethinking memory in ai: Taxonomy, operations, topics, and future directions View paper
- [20] Enhancing large language models through dynamic contextual memory embedding: A technical evaluation View paper
- [21] Efficient attention mechanisms for large language models: A survey View paper
- [22] Long context is not long at all: A prospector of long-dependency data for large language models View paper
- [23] Memory-Augmented Architecture for Long-Term Context Handling in Large Language Models View paper
- [24] Samba: Simple Hybrid State Space Models for Efficient Unlimited Context Language Modeling View paper
- [25] Memos: A memory os for ai system View paper
- [26] Loogle: Can long-context language models understand long contexts? View paper
- [27] RecurrentGemma: Moving Past Transformers for Efficient Open Language Models View paper
- [28] Hiagent: Hierarchical working memory management for solving long-horizon agent tasks with large language model View paper
- [29] Theory of mind for multi-agent collaboration via large language models View paper
- [30] LM2: Large Memory Models for Long Context Reasoning View paper
- [31] RMem: Bridging Memory Retention and Retrieval via Reversible Compression View paper
- [32] Long context scaling: Divide and conquer via multi-agent question-driven collaboration View paper
- [33] InfLLM: Unveiling the Intrinsic Capacity of LLMs for Understanding Extremely Long Sequences with Training-Free Memory View paper
- [34] InfiniteVL: Synergizing Linear and Sparse Attention for Highly-Efficient, Unlimited-Input Vision-Language Models View paper
- [35] Jarvis-1: Open-world multi-task agents with memory-augmented multimodal language models View paper
- [36] Memory-efficient Transformers via Top- Attention View paper
- [37] Acon: Optimizing context compression for long-horizon llm agents View paper
- [38] D2O: Dynamic Discriminative Operations for Efficient Long-Context Inference of Large Language Models View paper
- [39] HEMA: A Hippocampus-Inspired Extended Memory Architecture for Long-Context AI Conversations View paper
- [40] Agent workflow memory View paper
- [41] Infinite-LLM: Efficient LLM Service for Long Context with DistAttention and Distributed KVCache View paper
- [42] Dolphin: Long Context as a New Modality for Energy-Efficient On-Device Language Models View paper
- [43] Augmenting language models with long-term memory View paper
- [44] Taipan: Efficient and Expressive State Space Language Models with Selective Attention View paper
- [45] Memorize Step by Step: Efficient Long-Context Prefilling with Incremental Memory and Decremental Chunk View paper
- [46] Logarithmic Memory Networks (LMNs): Efficient Long-Range Sequence Modeling for Resource-Constrained Environments View paper
- [47] Human-inspired episodic memory for infinite context LLMs View paper
- [48] FLTRNN: Faithful Long-Horizon Task Planning for Robotics with Large Language Models View paper
- [49] ZeroMerge: Parameter-Free KV Cache Compression for Memory-Efficient Long-Context LLMs View paper
- [50] LIFT: Improving Long Context Understanding of Large Language Models through Long Input Fine-Tuning View paper
- [51] Evaluating memory in llm agents via incremental multi-turn interactions View paper
- [52] State and Memory is All You Need for Robust and Reliable AI Agents View paper
- [53] Agentic Troubleshooting Guide Automation for Incident Management View paper
- [54] Memory-Augmented Agent Training for Business Document Understanding View paper
- [55] Dialog Generation Using Multi-Turn Reasoning Neural Networks View paper
- [56] DelTA: An Online Document-Level Translation Agent Based on Multi-Level Memory View paper
- [57] Step up your game: A research on two/multi-step summarisation of long, regulatory documents View paper
- [58] Secom: On memory construction and retrieval for personalized conversational agents View paper
- [59] Doctoragent-rl: A multi-agent collaborative reinforcement learning system for multi-turn clinical dialogue View paper
- [60] Experience replay-based deep reinforcement learning for dialogue management optimisation View paper
- [61] High-Resolution Visual Reasoning via Multi-Turn Grounding-Based Reinforcement Learning View paper
- [62] In prospect and retrospect: Reflective memory management for long-term personalized dialogue agents View paper
- [63] Planning without Search: Refining Frontier LLMs with Offline Goal-Conditioned RL View paper
- [64] Context-lite multi-turn reinforcement learning for LLM agents View paper
- [65] RAIDEN-R1: Improving Role-awareness of LLMs via GRPO with Verifiable Reward View paper
- [66] History-Aware Cross-Attention Reinforcement: Self-Supervised Multi Turn and Chain-of-Thought Fine-Tuning with vLLM View paper
- [67] Memory-T1: Reinforcement Learning for Temporal Reasoning in Multi-session Agents View paper
- [68] Look back to reason forward: Revisitable memory for long-context llm agents View paper
- [69] Monte Carlo Planning with Large Language Model for Text-Based Game Agents View paper
- [70] A Category-Theoretic Framework for Wake-Sleep Consolidation in Dual-Transformer Architectures View paper

- [71] Large language models are semi-parametric reinforcement learning agents View paper
- [72] CtrlDiff: Boosting Large Diffusion Language Models with Dynamic Block Prediction and Controllable Generation View paper
- [73] Semantic coherence dynamics in large language models through layered syntax-aware memory retention mechanism View paper
- [74] On-Device Large Language Models: A Survey of Model Compression and System Optimization View paper
- [75] Amago-2: Breaking the multi-task barrier in meta-reinforcement learning with transformers View paper
- [76] Cognitive Architectures for Tomorrow: A Comprehensive Survey of Memory Management Paradigms in Agentic AI Systems View paper