

Article

Not peer-reviewed version

From Storage to Experience: A Survey on the Evolution of LLM Agent Memory Mechanisms

Jinghao Luo[‡], Yuchen Tian[‡], Chuxue Cao, Ziyang Luo, [Hongzhan Lin](#), Kaixin Li, Chuyi Kong, Ruichao Yang, Jing Ma^{*}

Posted Date: 15 January 2026

doi: 10.20944/preprints202601.0618.v2

Keywords: LLM agents; agent architectures; memory mechanisms; experience learning



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

From Storage to Experience: A Survey on the Evolution of LLM Agent Memory Mechanisms[†]

Jinghao Luo^{1,‡}, Yuchen Tian^{2,‡}, Chuxue Cao³, Ziyang Luo², Hongzhan Lin², Kaixin Li⁴, Chuyi Kong², Ruichao Yang⁵ and Jing Ma^{2,*}

¹ South China Normal University, China

² Hong Kong Baptist University, Hong Kong

³ Hong Kong University of Science and Technology, Hong Kong

⁴ National University of Singapore, Singapore

⁵ University of Science and Technology Beijing, China

* Correspondence: majing@comp.hkbu.edu.hk

[†] Our continuously updated list of papers and resources is available at <https://github.com/FeishuLuo/From-Memory-to-Experience>.

[‡] Equal contribution.

Abstract

Large Language Model (LLM)-based agents have fundamentally reshaped artificial intelligence by integrating external tools and planning capabilities. While memory mechanisms have emerged as the architectural cornerstone of these systems, current research remains fragmented, oscillating between operating system engineering and cognitive science. This theoretical divide prevents a unified view of technological synthesis and a coherent evolutionary perspective. To bridge this gap, this survey proposes a novel evolutionary framework for LLM agent memory mechanisms, formalizing the development process into three stages: **Storage** (trajectory preservation), **Reflection** (trajectory refinement), and **Experience** (trajectory abstraction). We first formally define these three stages before analyzing the three core drivers of this evolution: the necessity for long-range consistency, the challenges in dynamic environments, and the ultimate goal of continual learning. Furthermore, we specifically explore two transformative mechanisms in the frontier Experience stage: proactive exploration and cross-trajectory abstraction. By synthesizing these disparate views, this work offers robust design principles and a clear roadmap for the development of next-generation LLM agents.

Keywords: LLM agents; agent architectures; memory mechanisms; experience learning

1. Introduction

In recent years, the rapid advancement of Large Language Models (LLMs) has fundamentally reshaped the landscape of artificial intelligence (Hurst et al. 2024; Touvron et al. 2023; Yang et al. 2025). To augment the capabilities of LLMs, researchers have developed LLM-based agents that integrate LLMs with external tools and modular components, thereby enabling planning, tool use, and environmental interaction (Luo et al. 2025; Qin et al. 2024; Yao et al. 2022). However, the inherent statelessness of LLMs poses a critical challenge: it hinders agents from maintaining logical consistency across complex, multi-step tasks and precludes learning from prior interactions, often resulting in recurring reasoning errors (Huang et al. 2023; Xiong et al. 2025). Consequently, the development of effective memory mechanisms has emerged as an architectural cornerstone. By mitigating this deficiency, memory mechanisms underpin the robust operation of LLM-based agents and pave the way for self-evolution Wang et al. (2023); Wu et al. (2025).

We identify two primary obstacles to advancing memory mechanisms for LLM agents: (i) *Paradigmatic Fragmentation*: Existing methodologies oscillate between two weakly integrated paradigms. One focuses on engineering, adopting design principles from operating systems for the management of

memory data (Hu et al. 2024; Kang et al. 2025; Packer et al. 2023), while the other draws inspiration from cognitive science and psychology to simulate mechanisms for the formation, consolidation, and retrieval of human memory (Hou et al. 2024; Xu et al. 2025; Zhong et al. 2023). This lack of synergistic progress results in a fragmented body of research, preventing the formation of a coherent and continuous trajectory of evolution. (ii) *The Absence of Technological Synthesis*: Although numerous methods address isolated stages of memory processing, the field lacks a cohesive summary of the critical technologies that have historically propelled memory mechanism advancement (Xu et al. 2025; Yang et al. 2025; Zhang et al. 2025). Existing surveys have not sufficiently isolated these key technical drivers from general methodologies Cao et al. (2025); Du et al. (2025); Wu and Shu (2025); Wu et al. (2025). Consequently, the core technologies remain obscure, leaving future researchers without a clear roadmap of which innovations are robust enough to build upon.

While recent surveys have examined memory mechanisms for LLM agent systems, they lack a unified evolutionary perspective. This limitation obscures the internal drivers of memory development and impedes the in-depth exploration of architectures for next-generation agents. Specifically, Zhang et al. (2024) focuses on the classification of engineering modules, but fails to systematically expound on the logic behind critical technological transformations throughout their development. Furthermore, while Hu et al. (2025) addresses the dynamic processes of memory, its perspective remains confined to static functional categorizations, failing to reveal the underlying principles of dynamic evolution inherent to memory mechanisms.

To address these limitations, we propose a framework for memory mechanisms in LLM-based agents centered on dynamic evolution. We formalize this evolutionary process into three distinct stages: (i) **Storage**, which constructs diverse storage modes focused on the faithful recording of historical interaction trajectories; (ii) **Reflection**, which introduces a loop for dynamic evaluation to actively manage and refine these records; and (iii) **Experience**, which implements prospective guidance by abstracting high-level behavior patterns and strategies from clustered interactions (§ 2).

Building upon the proposed three stages of memory mechanisms, this survey follows a "Why-How-What" logic to address three interconnected research questions: **RQ1: Why do memory mechanisms evolve?** reveals how the requirements for long-range consistency, dynamic environment interaction, and continual learning serve as core catalysts driving mechanistic evolution (§ 3); **RQ2: How do memory mechanisms evolve?** delineates the evolutionary path from Storage to Reflection and then to Experience, analyzing the fundamental structural shifts involved (§ 4); and **RQ3: What changes does Experience bring?** provides an in-depth analysis of how frontier paradigms in the Experience stage, such as proactive exploration and cross-trajectory abstraction, address the bottlenecks in agent adaptability and autonomy (§ 5).

Finally, we outline future directions for LLM agent memory mechanisms. First, we emphasize that memory mechanisms should adopt more dynamic triggering modes based on task types (§ 6). Second, we highlight that the construction of working memory is a vital core of memory mechanisms. Next, we advocate for the development of more comprehensive datasets for memory mechanisms, especially for the Experience stage. Finally, we establish the coordination of distributed shared memory and the fusion of multimodal memory as critical breakthroughs for future research.

The overview of this survey and related datasets is documented in Appendix A and C, respectively.

2. Background

2.1. The LLM Agent Framework

We formalize an LLM-based agent as a decision-making entity parameterized by θ , interacting with a dynamic environment \mathcal{E} . The agent's operation is governed by a policy π_θ , which maps the current context to a probability distribution over the action space \mathcal{A} .

At time step t , the agent receives an observation $o_t \in \mathcal{O}$ and retrieves relevant information m_t from its memory module \mathcal{M} . The generated action a_t is sampled as follows:

$$a_t \sim \pi_\theta(a_t \mid \mathcal{I}, o_t, m_t), \quad (1)$$

where \mathcal{I} denotes the static system instruction, and $m_t = \text{Retrieve}(\mathcal{M}, o_t)$ represents the context-specific memory. Crucially, we distinguish between the global memory repository \mathcal{M} and its retrieved instantiation m_t at time t . In this survey, we define ‘‘LLM agent memory’’ \mathcal{M} as an externalized repository that bridges the frozen parametric knowledge in θ and the evolving environmental dynamics.

2.2. Taxonomy

We classify the evolution of memory mechanisms into three tiers based on the level of information abstraction and cognitive processing.

Storage. Storage serves as the foundational layer. Unlike higher-level mechanisms, storage preserves trajectories with minimal transformation, maintaining a one-to-one correspondence between memory entries and execution traces. We define a trajectory τ as a chronological sequence of observation-action pairs within a task session:

$$\tau = \langle (o_1, a_1), \dots, (o_T, a_T) \rangle. \quad (2)$$

The raw storage \mathcal{M}_{raw} is formally defined as a cumulative set of historical trajectories:

$$\mathcal{M}_{raw} = \{\tau_i\}_{i=1}^N, \quad \tau_i \in \mathcal{T}, \quad (3)$$

where \mathcal{T} represents the space of all possible interaction trajectories.

Reflection. Reflection is modeled as a semantic transformation mapping $\mathcal{F}_{ref} : \mathcal{T} \rightarrow \mathcal{S}$, where \mathcal{S} denotes the space of evaluated or corrected reasoning paths. Similar to the storage phase, Reflection functions as a mechanism to populate the global repository \mathcal{M} , but with a focus on quality density rather than raw fidelity.

It operates by analyzing a completed trajectory τ_i to generate a refined memory unit m'_i , which encapsulates critiques or corrective insights:

$$m'_i = \mathcal{F}_{ref}(\tau_i | \phi), \quad (4)$$

where ϕ represents the evaluation criteria. The key distinction lies in the storage protocol: while standard Storage preserves raw interaction logs, Reflection acts as a semantic filter, injecting processed insights back into the repository ($\mathcal{M} \leftarrow \mathcal{M} \cup \{m'_i\}$). Once stored, m'_i becomes an independent memory entry, decoupling the valuable logic from the specific noise of the original trajectory τ_i and serving as a refined reference for future retrieval.

Experience. Experience represents the highest cognitive layer, characterized by cross-trajectory abstraction. This stage aims to satisfy the Minimum Description Length (MDL) principle by compressing redundant trajectories into generalized schemas. Let $\mathcal{T}_{batch} \subset \mathcal{M}_{raw}$ be a subset of topologically similar trajectories. We define the Experience function \mathcal{F}_{exp} as an inductive operator that extracts a set of universally applicable rules \mathcal{K} :

$$\mathcal{T}_{batch} = \{\tau_i | \text{Sim}(\tau_i, \tau_j) > \epsilon\}. \quad (5)$$

$$\mathcal{K} = \mathcal{F}_{exp}(\mathcal{T}_{batch}) \quad \text{s.t.} \quad |\mathcal{K}| \ll \sum_{\tau \in \mathcal{T}_{batch}} |\tau|. \quad (6)$$

Formally, \mathcal{K} serves as a policy prior that elevates π_θ beyond rule consistent actions, enabling decision-making at a higher level of abstraction than the extracted rules themselves.

Summary & Ideas - Defining the Stages of memory mechanisms



Defining the Stages of memory mechanisms We provide a formal definition for the workflow of actions in LLM agents and characterize the memory mechanisms within this process as three stages of evolution: Storage, Reflection, and Experience. (cf. Figure 1)



We define these mechanisms based on the degree of utilization of memory, without introducing other related technical definitions.



The three stages of evolution do not represent complete substitution but are framed within an evolutionary perspective; a mechanism for memory may still retain certain characteristics of the preceding stage, even as its essence has transitioned into the subsequent stage.

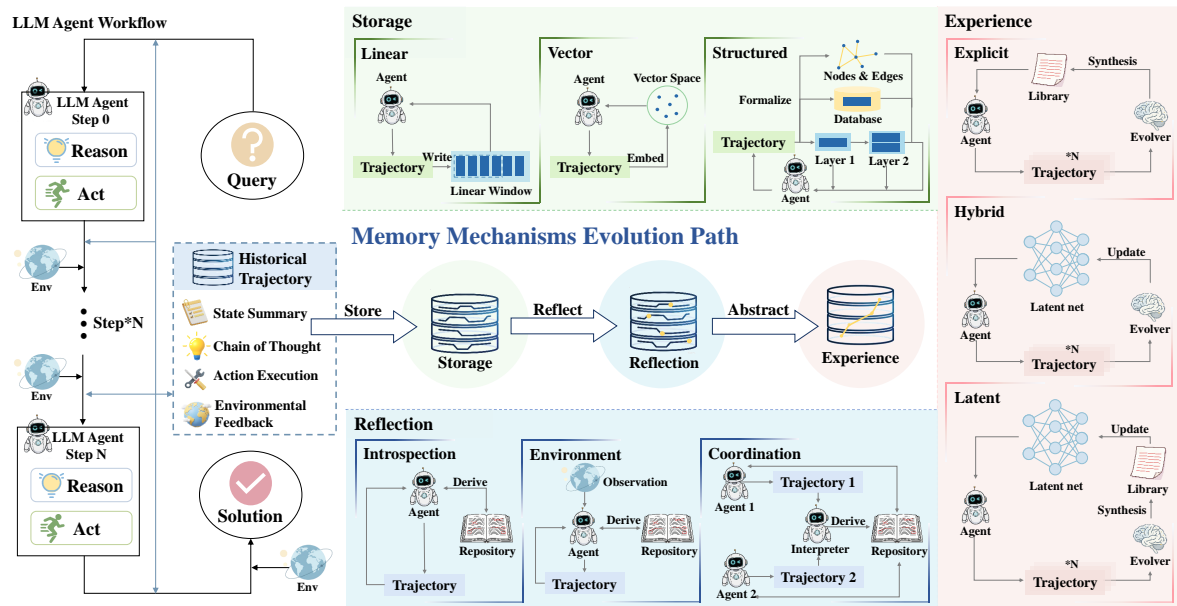


Figure 1. Overview of the LLM agent Memory Mechanisms System.

3. Evolutionary Drivers

To facilitate a comprehensive understanding regarding the evolution of memory mechanisms for LLM agents, we first address the fundamental question **RQ1: Why do memory mechanisms evolve?** In this section, we examine three core requirements for LLM agents to investigate how they drive the progression of memory mechanisms, thereby bridging the gap between models from pretraining and the real world.

3.1. Long-Term Consistency

Consistency across long horizons constitutes a prerequisite for the deployment of LLM agents within the real world and serves as the primary impetus for the early evolution of memory mechanisms. Although large language models exhibit robust local coherence within the context window, they frequently encounter issues such as redundant exploration, accumulation of errors, and discontinuities in reasoning during interactions involving multiple steps. We analyze the necessity of consistency over long durations through two dimensions: consistency of state and consistency of goals.

Coherence of State. The inherent statelessness of LLM agents results in a deficiency of internal mechanisms for explicit anchoring, which has catalyzed the emergence of modules for memory (Huang et al. 2023; Packer et al. 2023; Sumers et al. 2023). First, these modules maintain internal states for reasoning to ensure the coherence of thought (Sun et al. 2025; Yao et al. 2023); second, they synchronize the cognition of the agent with the external world to prevent erroneous decisions arising from inaccurate internal perceptions (Majumder et al. 2023; Yang et al. 2025); finally, they internalize interactions into persistent traits of the persona to ensure uniformity in behavior (Liang et al. 2025; Park et al. 2023; Westh auser et al. 2025).

Consistency of Goals. Due to the inherent nature of planning by the agent, LLM agents frequently optimize for actions with local consistency, which results in a departure from objectives at the global level (Everitt et al. 2025; Huang et al. 2024). Memory mechanisms mitigate this drift by providing persistent and explicit goals at a high level (Hu et al. 2024; Li et al. 2025). Furthermore, in systems with multiple agents, shared memory regarding goals can transform isolated behaviors into coordinated execution by the collective, thereby maintaining the unity of the final objective (Gao et al. 2024; Liu et al. 2025).

3.2. Dynamic Environments

The dynamic characteristics of the environment constitute a more enduring impetus for the evolution of memory mechanisms. In contrast to benchmarks of a static nature, the interplay between temporal validity and causality within the environment in settings from the real world renders fixed patterns for reasoning and static forms of storage rapidly fragile.

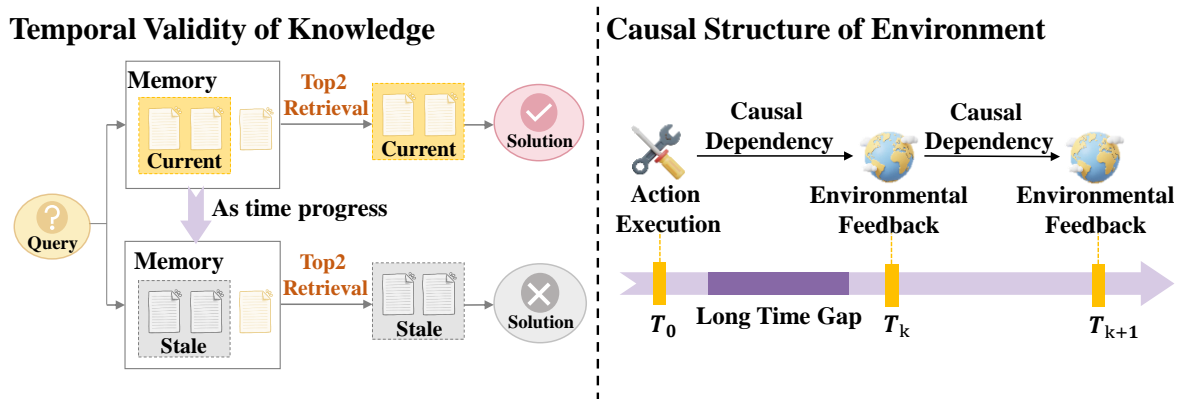


Figure 2. The Drivers in Dynamic Environments.

The Temporal Validity of Knowledge. Knowledge within environments of a dynamic nature is typically conditional rather than eternally valid (Jang et al. 2022; Ko et al. 2024; Lazaridou et al. 2021). As the environment progresses, strategies for action that were once correct may experience a gradual loss of utility. Crucially, knowledge that is outdated often fails without overt indication (Kalai and Vempala 2023; Kasai et al. 2024; Luu et al. 2022); although factually incorrect, such information may still exhibit significant relevance in its semantic representation. This necessity propels the evolution of memory mechanisms from the paradigm of static storage toward that of active management, integrating awareness of temporal factors, policies for decay, and methods for retrieval with enhanced flexibility (Du et al. 2025a; Houichime et al. 2025; Salama et al. 2025; Siyue et al. 2024; Zhong et al. 2023).

The Causal Structure of the Environment. Causal relationships within the complex real world involve delayed outcomes and cascading effects (Cui et al. 2025; Joshi et al. 2024; Liu et al. 2025). This necessitates that memory mechanisms transcend the mere documentation of interactions to construct dependencies for causality of a complex nature across steps in time (Du et al. 2025; Majumder et al. 2023; Raman et al. 2025). Consequently, planning with robustness is achieved through the realization of internal worlds characterized by consistency in causality (Bohnet et al. 2025; Kim and won Hwang 2025; Tang et al. 2024).

3.3. Continual Learning

Continual learning represents the ultimate requirement for LLM agents. Deployment within an open world inevitably involves encountering patterns that reside outside of the distribution of training. Without the effective internalization of these memories into actionable knowledge for reuse, the LLM agent will remain confined to repetitive cycles of trial and error. Therefore, memory mechanisms must not only enable the reproduction of historical trajectories but also address the bottlenecks of scaling and the requirements for abstraction inherent in dense memory.

Constraints on The Storage of Memory. Interaction with the real world over extended durations results in the linear expansion of memory in storage (Hu et al. 2023; Packer et al. 2023). Early memory mechanisms utilized techniques such as vectorization to scale storage capacity. However, recent research indicates that the unrestricted expansion of memory is detrimental to the performance of LLM agents, as errors propagate within the system for memory and contaminate the efficacy of learning (Srivastava and He 2025; Xiong et al. 2025). This necessitates the exploration of more strategic

policies for the addition and deletion of information within memory mechanisms [Du et al. \(2025\)](#); [Liu et al. \(2025\)](#).

The Requirement for Experience. The memory of most LLM agents is of an episodic nature and remains restricted to specific tasks ([Shinn et al. 2023](#); [Wang et al. 2023](#)). This limitation necessitates the transformation of raw clusters of memory into experience to provide guidance for behavior across future scenarios. Consequently, research on memory mechanisms has begun to explore various methodologies for the abstraction of experience ([Alakuijala et al. 2025](#); [Cai et al. 2025](#); [Guo et al. 2025](#); [Tang et al. 2025](#); [Xia et al. 2025](#)).

Summary & Ideas - Evolutionary Drivers



Consistency over long durations, dynamic environments, and the capacity for continual learning drive the evolution of LLM agents from static models into autonomous systems capable of operation within the real world over extended horizons.



The driving force for the evolution of memory mechanisms stems from a leap in research vision: from maintaining logical consistency in complex reasoning, to dynamic adaptation in environmental perception, and finally to the autonomous evolution of LLM agents.

4. Evolutionary Path

Building upon these evolutionary drivers, we conduct an investigation in-depth into **RQ2: How do memory mechanisms evolve?** We categorize the trajectory of evolution into three primary stages: storage, reflection, and experience.

4.1. Storage

The stage of storage serves as the starting point for memory mechanisms, where the primary objective is to resolve the contradiction between the limited window of context within Large Language Models and the continuously expanding history of interaction. Memory mechanisms during this phase are dedicated to the faithful preservation of interaction trajectories τ_i to the greatest extent possible to maintain consistency in the actions of the agent.

Linear. Linear storage represents the most direct method of recording, in which interaction trajectories are treated as a stream of tokens ordered by time and managed typically through a strategy of first-in, first-out (FIFO). Research focuses on the extension of the window of context via modifications to the mechanism of attention or the encoding of position ([Jin et al. 2024](#); [Ratner et al. 2022](#); [Xiao et al. 2023](#)), as well as the achievement of information sparsification through the mechanical reduction of noise ([Jiang et al. 2023](#); [Xiao et al. 2024](#); [Zhang et al. 2023](#)).

Vector. Vector storage encodes interaction trajectories into a high-dimensional space, which greatly expands the capacity for the storage of memory. Such methods shift the focus of research from the design of storage toward the optimization of retrieval, including retrieval based on semantic proximity ([Das et al. 2024](#); [Liu et al. 2024](#); [Melz 2023](#)) as well as weighted retrieval that incorporates temporal decay and scores for importance ([Park et al. 2023](#); [Zhong et al. 2023](#)).

Structured. Structured storage employs explicit data architectures to transcend the limitations on capacity inherent in linear storage and the ambiguity associated with vector retrieval. For instance, these methods utilize the tabular formats of relational databases for the storage of memory ([Hu et al. 2023](#); [Lee and Ko 2025](#); [Xue et al. 2023](#)), partition memory into distinct hierarchies to address the trade-off between storage capacity and speed of retrieval ([Lu et al. 2023](#); [Packer et al. 2023](#)), and directly model the history of interaction as a topological network of entities and relations ([Li et al. 2024](#); [Modarressi et al. 2024](#)).

4.2. Reflection

Mechanisms for storage fail to address the quality of memory, as raw trajectories are inevitably contaminated by hallucinations, errors in logic, and ineffective attempts. This limitation necessitates a transition of memory mechanisms toward reflection. In this phase, memory is transformed from a passive recorder into an active critic, utilizing various signals of feedback to perform correction and denoising of past trajectories to enhance the quality of the repository of memory.

Introspection. Introspective reflection conceptualizes the LLM agent as an autonomous critic that leverages the internal knowledge of the model to refine memory without the requirement for external feedback. Research in this area focuses on the correction of errors within trajectories (Bohnet et al. 2025; Liu et al. 2023; Zhang et al. 2025), the maintenance of the lifecycle of memory (Chhikara et al. 2025; Kang et al. 2025; Li et al. 2025), and the compression and distillation of long trajectories Han et al. (2025); Huang et al. (2025); Yang et al. (2025); Ye et al. (2025).

Environment. Environmental reflection treats signals from the external environment as the primary anchors for the reflection of memory to mitigate the issue of hallucinations. This approach focuses on the utilization of outcomes from the real world to proactively optimize policies for behavior (Sun et al. 2024; Yan et al. 2025,?) and calibrate internal models of the world (Sun et al. 2025, 2024; Xiao et al. 2025).

Coordination. Collaborative reflection extends this process to the collective, leveraging the division of roles and consensus to overcome bottlenecks in the cognition of individuals. This mechanism facilitates the reflection of memory through the construction of societies of heterogeneous agents Balestri and Pescatore (2025); Bo et al. (2024); Ozer et al. (2025); Wang et al. (2025).

4.3. Experience

Although reflection effectively mitigates noise and hallucinations, reflected memories are frequently fragmented and exhibit a high degree of dependence on context. This results in significant costs for retrieval and a heavy burden of inference for memory mechanisms when addressing new tasks. Moreover, recent research indicates that LLM agents often demonstrate a pronounced tendency to follow successful trajectories; corrected trajectories devoid of abstraction may still induce errors resulting from minor shifts in context. Consequently, memory in the stage of experience extracts universal heuristic wisdom by isolating similar trajectories from their specific contexts. This approach simultaneously compresses the originally vast repository of memory and enables a capacity for generalization to unknown environments through a form of intuition similar to that of humans.

Explicit. Explicit experience represents the integration of symbols, extracting human-readable and editable experiences with a high level of generalizability from clusters of interaction trajectories. This allows the LLM agent to achieve a highly interpretable process of self-evolution by either concretizing experiences into natural language policies (Cai et al. 2025; Hassell et al. 2025; Wan et al. 2025; Zhang et al. 2025) or directly abstracting them into executable entities (Shi et al. 2025; Wang et al. 2025; Zhang et al. 2025).

Implicit. Implicit experience internalizes interaction histories into model parameters, aiming to resolve the inference overhead and context limitations inherent in explicit memory. Implicit experience can be realized by directly converting experiences into the model's intrinsic capabilities through fine-tuning (Alakuijala et al. 2025; Tandon et al. 2025; Yu et al. 2025; Zhai et al. 2025; Zhang et al. 2025). Furthermore, the research community is exploring the transformation of experience into latent variables within the model's hidden layers, which are then dynamically invoked during the inference process (Zhang et al. 2025,?).

Hybrid. Hybrid experience establishes a dynamic cycle of accumulation and internalization. Through a mechanism for experience transfer, explicit experience is treated as a high-capacity cache, which is subsequently compressed and internalized into the implicit weights of the model through periodic updates of parameters (Anonymous 2025; Ouyang et al. 2025; Wu et al. 2025).

Summary & Ideas - Evolutionary Path

- This chapter delineates the evolutionary trajectory of memory mechanisms for LLM agents, which originates from the storage phase dedicated to the pursuit of faithful recording, proceeds through the reflection phase involving the correction of errors and the reduction of noise, and ultimately culminates in the experience phase defined by abstraction of a high order.
- 💡 Although a portion of the literature adopts the terminology of experience, these works may not fall within the conceptual scope of experience established for this survey.
- 💡 Due to space constraints, we do not provide detailed descriptions of these works within the evolutionary path; instead, some representative studies are presented in the Appendix B.

5. Transformative Experience

Following the exposition of the evolutionary trajectory for memory mechanisms, we address **RQ3: What changes does Experience bring?** In this section, we elucidate the distinct technical characteristics of experience as a novel stage in the development of memory mechanisms.

5.1. Active Exploration

Active exploration leverages memory mechanisms to transform LLM agents from passive recorders of information into collectors of experience driven by goals. In the stage of Experience, the core capability of memory mechanisms is no longer confined to the storage of history, but extends to the acquisition of valuable experience through the active exploration of the environment.

Exploration Mechanisms. The driving mechanisms for active exploration have transitioned from traditional strategies of random exploration toward more profound drivers of intrinsic motivation and feedback. Drivers based on signals of reward guide LLM agents to explore state spaces of greater value through the design and optimization of immediate reward functions (Pan et al. 2025; Sun et al. 2025; Zheng et al. 2024); drivers based on curricula facilitate exploration tasks of increasing difficulty through the dynamic generation and adjustment of sequences for tasks (Ahn et al. 2022; Wei et al. 2025); and drivers based on reuse enable exploration of high efficiency through the abstraction and reuse of trajectories from history (Cai et al. 2025; Wang et al. 2025).

Exploration Dimensions. The core of active exploration resides in the utilization of memory mechanisms to facilitate the expansion of the boundaries of capability for LLM agents. This process can be categorized into three critical dimensions: exploration of breadth aims to alleviate cognitive deficiencies of LLM agents in unfamiliar environments, transforming memory into experience that is structured through mechanisms of curiosity (Cheng et al. 2025; Qi et al. 2025; Zhai et al. 2025); exploration of depth focuses on the extraction of skills of a high order within vertical tasks, driving the evolution of memory from the basic following of instructions to complex experiential strategies (Liu et al. 2025; Xia et al. 2025); and exploration of strategy centers on the dynamic optimization of paths for decision making, leveraging the accumulation of experience to enhance the precision of decisions for LLM agents during planning over long horizons Bidochko and Vykylyuk (2026); Shi et al. (2025).

5.2. Cross-Trajectory Abstraction

Cross-trajectory abstraction compresses isolated trajectories into universal patterns, transforming scattered and episodic experiences into stable priors for policy. This enables LLM agents to transcend specific sequences of actions and engage in decision making at higher dimensions of abstraction, which provides prospective guidance for tasks that are unknown and facilitates an understanding of underlying regularities.

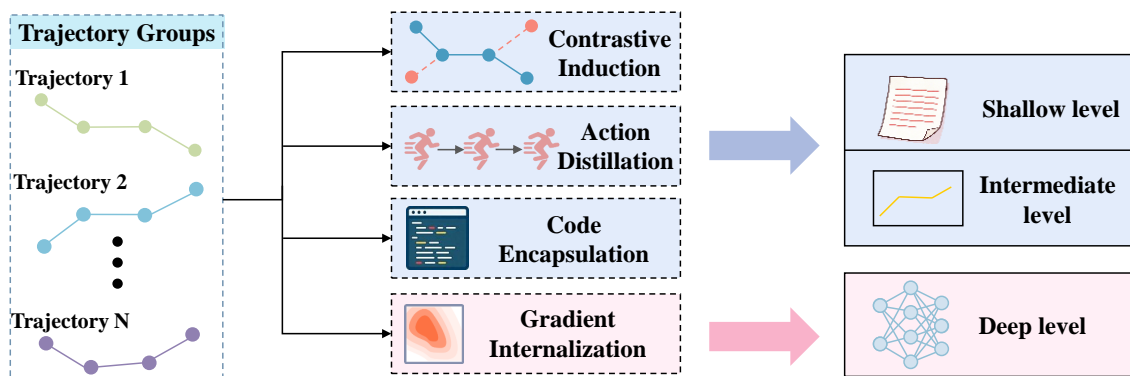


Figure 3. Overview of Cross-Trajectory Abstraction.

Abstraction Mechanisms. The mechanism for abstraction serves as the core operator for the transformation of groups of raw interaction trajectories into experience that is universal. In contrast to

mechanisms for reflection that focus on the correction of errors within single trajectories, abstraction during the stage of experience emphasizes the execution of inductive operations across trajectories. According to its operational logic, this process includes contrastive induction, which utilizes the opposition between successful and failed trajectories to delineate the boundaries of policy with precision (Forouzandeh et al. 2025; He et al. 2024); the distillation of actions of fine granularity into patterns of thought of a high order through the chunking and aggregation of behavioral sequences across multiple levels of granularity (Fang et al. 2025; Latimer et al. 2025); the encapsulation of recurring patterns of behavior into program functions that are reusable by leveraging the compositionality of code (Wang et al. 2025; Yang et al. 2025; Zhang et al. 2025); and the internalization of groups of trajectories into the parameters of the model through techniques for fine-tuning (Chen et al. 2025; Ding et al. 2025).

Abstraction Granularity. The hierarchy of abstraction for memory mechanisms determines the boundaries for generalization and the degree of interpretability for experience. Based on the degree to which the results of abstraction deviate from original trajectories, these can be categorized into three progressive levels: abstraction at the shallow level retains a portion of semantic logic, utilizing “rules” described in natural language as experience (Cao et al. 2025; Chen et al. 2025; Hayashi et al. 2025; Wei et al. 2025); abstraction at the intermediate level completely removes redundancies of natural language, extracting only modular skeletons for execution as experience (Liu et al. 2025; Wang et al. 2024; Yu et al. 2025); and abstraction at the deep level compresses the distribution of trajectories into the weights of the model, enabling the complete transformation of experience into intuition for decision making (Cheng et al. 2025; Luo et al. 2025; Wang et al. 2025).

Summary & Ideas - Transformative Experience

- Within the novel stage of memory mechanisms characterized by Experience, active exploration and abstraction across trajectories represent the two primary features of this phase.
- 💡 The feedback loop between exploration and abstraction serves as the central engine that drives the autonomous and continuous evolution of LLM agents during the phase of experience.
- 💡 Memories characterized by a high degree of abstraction during the phase of experience also enable agents to circumvent challenges frequently associated with conventional memory mechanisms (Gharat et al. 2025; Srivastava and He 2025; Xiong et al. 2025).

6. Future Directions

In this section, we discuss emerging prospects and promising directions for memory mechanisms for LLM agents.

Active Memory Perception. Currently, certain memory mechanisms still utilize modes of passive triggering, which necessitates that LLM agents perform the indiscriminate retrieval of a significant portion of the repository of memory (Rasmussen et al. 2025; Wang et al. 2025). More importantly, the persistent retrieval of irrelevant or obsolete memories can disrupt the coherence in reasoning of the LLM agent (Tan et al. (2025); Xu et al. (2025)). Recent work has begun to address this challenge through autonomous retrieval controllers (Du et al. 2025b). Future research demands that memory mechanisms autonomously evaluate whether a task requires the introduction of additional memory and determine the specific type of memory to be integrated, ensuring that memory mechanisms function as resources that are invoked on demand.

Organization of Working Memory. As the complexity of tasks and the horizons encountered by LLM agents continue to expand, the construction of working memory within tasks has emerged as a primary bottleneck. LLM agents must reconstruct trajectories from the past into memory intervals that are dynamic and plastic to facilitate the effective allocation of attention (Hu et al. 2024; Luo et al. 2025). Future research may focus on the isolation of interval memory, the retrospective integration of critical nodes for decision making, and the adaptive pruning of working memory (Nan et al. 2025; Sun et al. 2025; Zhang et al. 2025).

Benchmark for Experience. Existing datasets primarily evaluate the capacity for retrieval and de-noising of memory within the stages of storage and reflection, whereas the evaluation of the capacity for abstraction and generalization during the stage of experience remains significantly insufficient

(Appendix C). The assessment of the lifecycle of experience is closely linked to the capacity for meta-learning in LLM agents, which is essential for the realization of systems for self evolution based on active generalization (Behrouz et al. 2024; Wei et al. 2025). Consequently, the path of evolution we propose provides a valuable foundation for guiding the development of these benchmarks.

Distributed Shared Memory. Collaboration among multiple agents represents the essential pathway toward the realization of “Organizations.” This establishes distributed shared memory as central to current research Wu et al. (2025). At present, mechanisms for shared memory rely primarily on communication through explicit dialogue, which is not only constrained by bottlenecks in bandwidth but is also prone to the introduction of noise during the process of exchange (Liao et al. 2025; Tran et al. 2025; Zou et al. 2025). To overcome current communication constraints, future efforts should prioritize the development of consensus memory systems. These systems aim to achieve efficient synchronization between individual perspectives and collective knowledge, thereby fostering a more agile process of socialized experience evolution (Rezazadeh et al. 2025; Yuen et al. 2025).

Multimodal Memory. Multimodal memory represents a significant direction for the future development of memory mechanisms for LLM agents. This direction requires the integration of states of visual perception, processes of linguistic reasoning, and other perceptual modalities into memory units characterized by unified temporality and semantics (He et al. 2025; Liu et al. 2025; Zhou et al. 2025). For embodied intelligence in particular, the integrity of internal models of the world directly influences the planning and execution of tasks (Feng et al. 2025; Long et al. 2025). To achieve this objective, existing research explores novel memory mechanisms by investigating multimodal abstraction, temporal alignment across modalities, and the efficient consolidation of memory.

7. Conclusion

This survey provides a systematic review of memory mechanisms for LLM agents, establishing an evolutionary framework that encompasses three progressive stages: storage, reflection, and experience. Our analysis demonstrates that memory evolution is not merely the expansion of storage capacity but fundamentally involves the enhancement of information density and a transformation across cognitive abstraction dimensions. By introducing mechanisms such as active exploration and cross-trajectory abstraction, memory mechanisms within the experience stage enable agents to transcend situational constraints and acquire transferable behavioral experience. We hope this survey assists the community in designing more advanced memory mechanisms, guiding LLM agents toward the realization of true general artificial intelligence.

Limitations

This survey provides a comprehensive qualitative analysis of memory mechanisms for LLM agents; however, we acknowledge several limitations that warrant discussion.

Lack of Direct Quantitative Comparison. This survey adopts a qualitative analytical framework and lacks a comprehensive performance comparison of memory mechanisms. This is because the design objectives differ across the three stages of storage, reflection, and experience, and no unified benchmark currently exists for comprehensive evaluation across all stages. Moreover, variations in foundation models, environments, and prompts across original studies render direct numerical comparison potentially misleading.

Relation to Established Learning Paradigms. The experience stage, particularly implicit experience, intersects with fine-tuning, reinforcement learning, and meta-learning at a technical level. This taxonomy does not position experience as an entirely novel learning paradigm; rather, it emphasizes how these established techniques are deployed within memory-centric LLM agent architectures, serving as a critical intermediary between interaction trajectories and parameter updates.

Temporal Coverage and Recency Bias. Research on memory mechanisms for LLM agents has experienced rapid growth from 2024 to 2025, with the experience stage emerging as a coherent research direction only in the latter half of 2025. This temporal distribution is reflected in the coverage of this

survey and brings two methodological implications: (i) early influential works may not have received attention commensurate with their historical contributions, and (ii) some recent preprints included in this survey have not yet undergone formal peer review. To balance academic rigor with timeliness, this survey prioritizes works that propose novel architectures or demonstrate reproducible results.

A. Overview

We first provide a formal definition for the operational framework of LLM agents and the evolutionary paradigms of memory mechanisms. Furthermore, we categorize the primary drivers for this evolution into three dimensions, emphasizing that these driving forces constitute the fundamental factors supporting transformations in memory mechanisms and the inherent capabilities of LLM agents. Based on the depth to which historical trajectories are utilized, this survey proposes an evolutionary framework consisting of three distinct stages:

- **Storage:** As the foundational layer of evolution, this stage focuses on the faithful preservation of trajectories from interactions over a long duration to address constraints regarding the memory capacity of LLM agents.
- **Reflection:** Through the introduction of loops for dynamic evaluation, the memory mechanism transitions from a recorder of information to an evaluator, thereby mitigating issues related to hallucinations and logical errors within the memory of LLM agents.
- **Experience:** Representing the highest level of cognition, this stage employs abstraction across multiple trajectories to extract behavioral patterns of a higher order. This process compresses redundant memory into heuristic strategies that are transferable and reusable.

Furthermore, we provide an in-depth discussion on two pivotal technological shifts required for memory mechanisms to advance toward the stage of experience: active exploration and abstraction across trajectories. These advancements enable LLM agents to transition from passive recipients of information to collectors of experience driven by specific goals, thereby enhancing the capacity for proactive generalization in tasks that are unknown. Finally, this survey discusses several valuable directions for the future of memory mechanisms.

Summary of Contribution. As a survey, our primary objective is the synthesis and analysis of existing research, while providing novel insights and perspectives for researchers who aim to understand and design memory mechanisms. We believe that our work offers significant novelty in the following aspects:

- **Scope & Coverage:** To address the absence of a perspective on evolution and the significant fragmentation in contemporary research on memory mechanisms, this survey provides a comprehensive overview that is forward-looking. This work encompasses research that has been overlooked, the most recent advancements, and theoretical perspectives of a broader nature.
- **Organization & Structure:** This survey constructs an evolutionary framework in three stages to organize the manuscript. On this basis, we systematically delineate the drivers and pathways for the development of memory mechanisms, as well as characteristics at the frontier. This perspective provides novel insights for research within this domain.
- **Insights & Critical Analysis:** This survey provides original interpretations and an in-depth analysis of the existing literature. For instance, we propose a taxonomy from an evolutionary perspective, using the degree of utilization for trajectories of past interaction as a benchmark. Furthermore, we summarize two pivotal characteristics of memory mechanisms in the stage of experience and identify several issues that remain underexplored or unresolved.
- **Timeliness & Relevance:** In the inaugural year of LLM agents, this work represents the first survey to systematically examine memory mechanisms from a perspective of evolution, capturing research at the frontier through 2025. It addresses the urgent necessity for adaptation and learning as agents encounter the real world for the first time. Through the synthesis of

existing literature, we provide a new foundation for further exploration and innovation in this critical field.

Overall, this survey offers a novel perspective by providing comprehensive coverage and an innovative classification of memory evolution. We anticipate that these contributions will bridge the knowledge gap regarding LLM agent memory mechanisms and offer a foundational resource for future research.

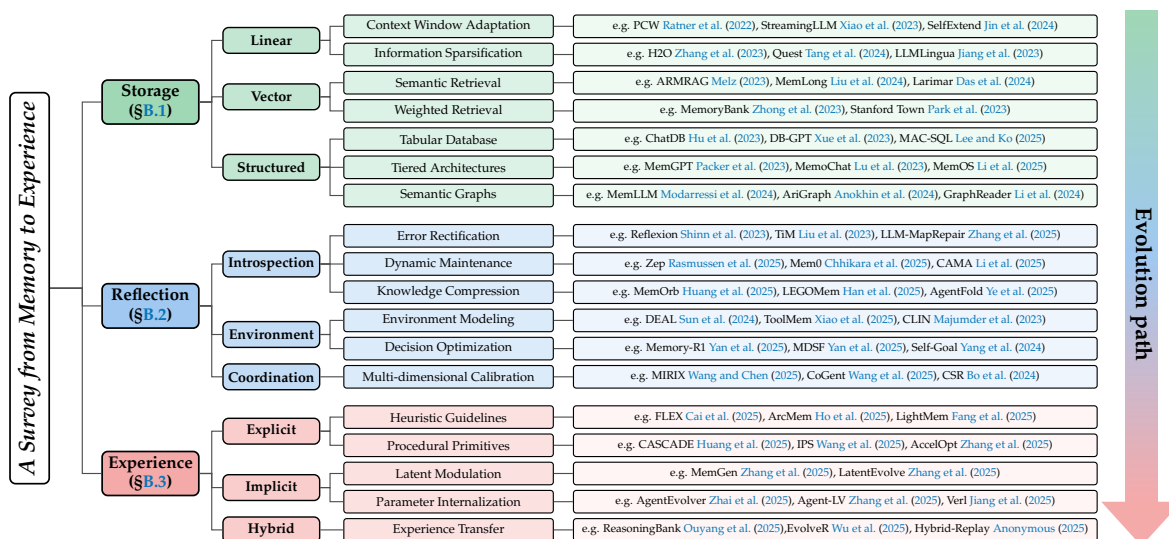


Figure 4. Taxonomy of LLM agent Memory Mechanisms

B. Detail Within the Evolutionary Path

Due to space constraints in the main text, we provide detailed exposition of representative works within each stage of the memory mechanism evolution in this section. For each work, we describe its core contribution, technical mechanism, and position within the evolutionary trajectory.

B.1. Storage

The primary objective of the storage phase is the precise preservation of trajectories to the maximum extent possible, which enables LLM agents to maintain an accurate perception of both internal and external states [Xi et al. \(2023\)](#). Although the memory mechanisms of the storage phase provide the context necessary for continuity and reasoning, they remain inherently susceptible to contamination from the stochasticity and hallucinations of the underlying model. Prior research has addressed the requirements for the writing, management, and retrieval of memory within various environments by constructing memory architectures across three technical categories: linear, vector, and structured.

Linear. Linear storage represents the most primitive and intuitive form of memory mechanisms. It treats the records of interaction for LLM agents as a continuous stream of tokens arranged in chronological order, managing memory within the context window through strict adherence to a strategy of first-in, first-out (FIFO). The work in this phase can be categorized into two components: adjustment of the context window and sparsification of information.

- Context Window Adaptation:** Context window adaptation techniques seek to extend the usable input length of LLMs by modifying attention mechanisms, positional encoding schemes, or input structures. Representative approaches include optimizing intrinsic attention computation [Xiao et al. \(2023\)](#), remapping positional encodings to enable longer sequences [Jin et al. \(2024\)](#), and restructuring inputs to mitigate length constraints [Ratner et al. \(2022\)](#). These methods expand raw storage capacity but do not alter the semantics of stored trajectories.

- **Information Sparsification:** Information sparsification treats memory compression as a mechanical denoising process independent of agent reflection. Methods typically rely on statistical or attention-based heuristics to remove low-utility tokens. For example, [Zhang et al. \(2023\)](#) evicts tokens based on cumulative attention scores, while [Tang et al. \(2024\)](#) and [Xiao et al. \(2024\)](#) retrieve salient memory blocks via query–key similarity. [Jiang et al. \(2023\)](#) further identifies redundant segments through perplexity estimation. While effective for efficiency, these methods operate without semantic abstraction.

Vector. Vector storage mitigates the constraints of capacity for memory storage by encoding interaction trajectories into spaces of high dimensionality. However, it also introduces a novel challenge: the efficient retrieval of memories relevant to the task from massive repositories. Consequently, the focus of research has transitioned toward the optimization of retrieval. We categorize these methodologies into two classes: semantic retrieval and weighted retrieval.

- **Semantic Retrieval:** Semantic retrieval constitutes the foundational approach to vector memory, where relevance is determined by geometric proximity in embedding space. [Melz \(2023\)](#) retrieves historical reasoning chains via semantic alignment, while [Liu et al. \(2024\)](#) integrates fine-grained retrieval-attention during decoding to sustain long-context reasoning. [Das et al. \(2024\)](#) further internalizes episodic memory into a latent matrix, enabling one-shot read–write operations. Despite improved recall, these methods treat retrieved content as flat historical evidence.
- **Weighted Retrieval:** Weighted retrieval extends semantic similarity by assigning differentiated importance to memories using multi-dimensional scoring signals. [Zhong et al. \(2023\)](#) models temporal decay via the Ebbinghaus Forgetting Curve, while [Park et al. \(2023\)](#) retrieves memories based on a weighted combination of relevance, recency, and importance. Such mechanisms improve prioritization but remain retrieval-centric rather than abstraction-driven.

Structured. Structured storage preserves memory through predefined structures of relationships. This paradigm emphasizes the integrity and enforcement of knowledge within memory, which facilitates precise operations, complex reasoning based on logic, and efficient retrieval across multiple hops. Based on the method of organization, we categorize these systems into three classes: tabular databases, tiered architectures and semantic graphs.

- **Tabular Database:** Database-backed memory systems leverage mature relational databases to store agent knowledge in structured tabular form. Early work frames databases as symbolic memory [Hu et al. \(2023\)](#), while subsequent approaches translate natural language queries into SQL via specialized controllers for secure and efficient access [Xue et al. \(2023\)](#). Multi-agent extensions further distribute database construction and maintenance across specialized roles [Lee and Ko \(2025\)](#).
- **Tiered Architectures:** Tiered memory architectures draw inspiration from computer storage hierarchies and human cognition to balance capacity and access latency. MemGPT [Packer et al. \(2023\)](#) introduces a dual-layer design separating main and external context, enabling virtual context expansion. Cognitive-inspired systems such as SWIFT–SAGE [Lin et al. \(2023\)](#) dynamically adjust retrieval intensity, while streaming-update architectures maintain long-term stability without exhaustive retrieval [Lu et al. \(2023\)](#); [Zhou et al. \(2023\)](#).
- **Semantic Graphs:** Graph memory represents interaction histories as networks of entities and relations, enabling structured reasoning beyond flat storage. Triplet-based extraction supports precise updates and retrieval [Modarressi et al. \(2024\)](#), while neuro-symbolic approaches integrate logical constraints into graph representations [Wang et al. \(2024\)](#). Graph-based world models further support environment-centric reasoning [Anokhin et al. \(2024\)](#), and coarse-to-fine traversal over text graphs enables efficient long-context retrieval [Lu et al. \(2023\)](#); [Zhou et al. \(2023\)](#).

B.2. Reflection

Although the storage stage explores diverse methods for the preservation of memory to ensure the consistency of LLM agents over the long term, these approaches do not fundamentally address the quality of memory. Raw trajectories of interaction inevitably conflate successful sequences with hallucinations, errors in logic, and attempts that are invalid [Ghasemabadi and Niu \(2025\)](#); [Zhang et al. \(2025, 2023\)](#). Without the application of evaluation, the passive storage of all trajectories leads to an accumulation of errors and the repetition of failures. The reflection stage incorporates introspection, the environment, and coordination as signals for feedback to rectify and denoise historical trajectories, thereby producing memory of higher quality.

Introspection. Introspective reflection represents an internal cognitive process that utilizes the LLM agent's own knowledge to evaluate, refine, and restructure memory without the need for external feedback. Current research achieves introspection through three distinct functional pathways: error rectification, dynamic maintenance, and knowledge compression.

- **Error Rectification:** targets hallucinations and multi-step reasoning errors by verifying and repairing stored trajectories through self-critique. [Shinn et al. \(2023\)](#) introduces Reflexion, which prompts agents to reflect on failed trajectories and distill corrective feedback into textual memory. This mechanism enables systematic error correction and sustained performance improvement across episodes, establishing introspective reflection as a central mechanism rather than a peripheral heuristic. Building on this paradigm, [Liu et al. \(2023\)](#) introduces a post-reasoning verification stage to retain only validated memories, while [Zhang et al. \(2025\)](#) detects contradictory or erroneous segments through introspective consistency checks, thereby limiting error accumulation and propagation.
- **Dynamic Maintenance:** Dynamic maintenance focuses on lifecycle management of memory content. [Li et al. \(2025\)](#) incrementally updates internal knowledge schemas via clustering, while [Rasmussen et al. \(2025\)](#) and [Chhikara et al. \(2025\)](#) maintain continuity by parsing and updating structured entity relations. At the system level, rule-based controllers inspired by operating systems strategically update and persist core memories [Kang et al. \(2025\)](#); [Packer et al. \(2023\)](#); [Zhou et al. \(2025\)](#).
- **Knowledge Compression:** Knowledge compression distills high-dimensional trajectories into compact and reusable representations. [Huang et al. \(2025\)](#) generates structured reflections to extract coherent character profiles, while [Han et al. \(2025\)](#) decomposes interaction sequences into modular procedural memories. Multi-granularity abstraction further aligns distilled memories with task demands [Tan et al. \(2025\)](#); [Yang et al. \(2025\)](#), and context-folding techniques preserve working-context efficiency during reasoning [Li et al. \(2025\)](#); [Sun et al. \(2025\)](#); [Ye et al. \(2025\)](#).

Environment. While introspective reflection leverages knowledge within the model to refine memory, it inherently carries a risk of inconsistency with factual reality. To mitigate this risk, reflection from the environment utilizes outcomes in the real world to actively optimize behavior and calibrate the internal knowledge of the model. Current research primarily proceeds along two trajectories: environment modeling and decision optimization.

- **Environment Modeling:** Environmental modeling aligns internal memory with dynamic external conditions such as environments, tools, and user preferences. [Sun et al. \(2024\)](#) enables agents to infer and validate world rules from demonstrations, while [Xiao et al. \(2025\)](#) summarizes tool behavior from execution outcomes. Preference-aware updates integrate short-term variation with long-term trends [Sun et al. \(2025\)](#), and EM-based formulations ensure memory consistency under distribution shifts [Yin et al. \(2024\)](#).
- **Decision Optimization:** Decision optimization treats memory management as a learnable policy guided by environmental rewards or execution feedback. [Yan et al. \(2025\)](#) learns discrete actions from outcome-based rewards, while [Yan et al. \(2025\)](#) refines memory quality

using value-annotated decision trajectories. For complex planning, interaction feedback is used to validate and prune goal hierarchies [Yang et al. \(2024\)](#).

Coordination. Collaborative reflection leverages the specialization of roles and mechanisms for consensus within systems of multiple agents to extend the process of reflection from the level of the individual to that of the collective. Through deliberation across multiple agents, this paradigm alleviates cognitive bottlenecks and hallucinations that are common in architectures with a single model during the processing of trajectories of complex memory.

- **Multi-dimensional Calibration:** Multi-dimensional calibration realizes distributed memory management through heterogeneous agent societies. [Wang and Chen \(2025\)](#) coordinates core, episodic, and semantic memory modules to process multimodal long contexts. [Wang et al. \(2025\)](#) decomposes graph reasoning into perception, caching, and execution roles to reduce context loss. Narrative-level coherence is achieved by integrating episodic and semantic memories across agents [Balestri and Pescatore \(2025\)](#). Furthermore, [Ozer et al. \(2025\)](#) and [Bo et al. \(2024\)](#) frameworks further enhance reasoning consistency and collaboration efficiency in agent societies by implementing collaborative reflection across diverse roles and personalized feedback mechanisms

B.3. Experience

Although reflection mechanisms mitigate hallucinations and noise through evaluation, their corrective efficacy remains at the level of trajectories and has not yet yielded knowledge at the level of strategy that is transferable. ([Renze and Guven 2024](#); [Shinn et al. 2023](#)) In addition, reflection focused on trajectories may lead to a linear expansion of the memory bank, which imposes a burden for inference and may potentially result in a characteristic of following trajectories. ([Fu et al. 2025](#); [Hong and He 2025](#); [Zhu et al. 2025](#)) Consequently, memory mechanisms must transcend the limitations of reflecting on the past and move toward a stage of experience for the guidance of the future. At this stage, memory mechanisms abstract universal wisdom that is independent of context from clusters of trajectories; through this process, LLM agents truly liberate themselves from memory banks that are verbose and complex, achieving zero-shot transfer to scenarios that are unknown by means of skills or rules that are intuitive. Research on the stage of experience achieves prospective wisdom through the abstraction of experience in forms that are explicit, implicit, and hybrid.

Explicit. Explicit experience abstracts patterns of knowledge that are readable by humans, editable, and generalizable from clusters of trajectories, framing experiential memory as insights of wisdom that allow for direct retrieval and reuse, analogous to the consultation of a reference manual or a library of functions. This methodology not only alleviates the pressure of inference but also provides LLM agents with capabilities for interpretability and self-evolution. Research on explicit experience can generally be categorized into Heuristic Guidelines and Procedural Primitives.

- **Heuristic Guidelines:** Heuristic guidelines serve to crystallize implicit intuition into explicit natural language strategies. In this domain, researchers focus on distilling experience into textual rules: [Ouyang et al. \(2025\)](#) abstracts key decision principles through contrastive analysis of successful and failed trajectories, while [Suzgun et al. \(2025\)](#) proposes dynamically generated "prompt lists" for real-time heuristic guidance. [Xu et al. \(2025\)](#) and [Hassell et al. \(2025\)](#) investigate rule induction from supervisory signals, achieving textual experience transfer via "cross-domain knowledge diffusion" and "semantic task guidance," respectively. To transcend linear text limitations in modeling complex dependencies, recent work shifts toward structured schemas. [Ho et al. \(2025\)](#) and [Zhang et al. \(2025\)](#) abstract multi-turn reasoning traces into experience graphs, leveraging topological structures to capture logical dependencies and enable effective storage and reuse of collaboration patterns and high-level cognitive principles. Moreover, [Cai et al. \(2025\)](#) organizes heuristic knowledge into modular and compositional units, enabling systematic reuse across tasks.

- **Procedural Primitives:** Procedural primitives represent the abstraction of complex reasoning chains into executable entities, designed to significantly reduce planning overhead. Wang et al. (2025) proposes a skill induction mechanism that encapsulates high-frequency action sequences into functions, enabling agents to invoke complex skills as atomic actions. Zhang et al. (2025) extends this executable paradigm to hardware optimization, enabling agents to accumulate kernel optimization skills that iteratively enhance accelerator performance. In this line of work, Huang et al. (2025) enables the composition and cascading execution of such procedural primitives, allowing agents to construct complex behaviors through structured skill invocation.

Implicit. Implicit experience eschews the retrieval of discrete text and abstracts the history of interactions into implicit priors, thereby addressing the overhead of inference and the constraints of context. Experiential memory is transformed into latent variables within spaces of high dimensionality or into parameters of the neural network. Based on the form of implementation for the transformation, implicit experience is categorized into two trajectories: Latent Modulation and Parameter Internalization.

- **Latent Modulation:** Latent modulation operates on the cognitive stream within continuous high-dimensional latent space. By encoding experience into latent variables or activation states, this paradigm "weaves" historical insights into current reasoning in a parameter-free manner, circumventing expensive parameter updates. Zhang et al. (2025) introduces the MemGen framework, employing a "Memory Weaver" to dynamically generate and inject latent token sequences conditioned on current reasoning state. Zhang et al. (2025) achieves smooth transfer from historical experience to current decision-making without altering static parameters, using alternating Fast Retrieval and Slow Integration within latent space.
- **Parameter Internalization:** Parameter Internalization transforms explicit trajectories into intrinsic capabilities within model weights. Through gradient updates, this mechanism instills adaptive priors into LLM agents, enabling effective navigation of complex environments. For context distillation, Alakuijala et al. (2025) proposes iterative distillation to internalize corrective hints into model weights. Liu et al. (2025) converts business rules into model priors, alleviating retrieval overload in RAG systems, while Zhai et al. (2025) introduces "Experience Stripping," eliminating retrieval segments during training to force internalization of explicit experience into autonomous reasoning capabilities independent of external auxiliaries. For Reinforcement Learning, Zhang et al. (2025) proposes a pioneering early experience paradigm, leveraging implicit world models and sub-reflective prediction to internalize trial-and-error experience into policy priors without extrinsic rewards. Lyu et al. (2025) achieves strategic transformation from Reflection to Experience by applying RL to student-generated reflections. Feng et al. (2025) proposes group-based policy optimization for fine-grained experience internalization across multi-turn interactions. Jiang et al. (2025) establishes standardized alignment between RL and tool invocation, enhancing agents' capacity to transmute tool-use experience into intrinsic strategies.

Hybrid. Hybrid experience aims to transcend the dichotomy between explicit and implicit paradigms by establishing a dynamic "Accumulate-Internalize" cycle. This paradigm directly addresses the challenges of "Storage Explosion" and "Retrieval Latency" encountered by explicit experience repositories during long-term interactions, while simultaneously mitigating the tension caused by parameter updates lagging behind environmental dynamics.

- **Experience Transfer:** Experience Transfer facilitates capability internalization by progressively decoupling agents from external retrieval reliance. Wu et al. (2025) employs offline distillation to abstract complex trajectories into structured experience for inference guidance, then uses these experiences to generate high-quality trajectories for policy optimization. By transferring knowledge from explicit experience pools into model parameters via gradient updates, this approach eliminates dependence on external retrieval systems. (Anonymous

2025; Ouyang et al. 2025) maintains an explicit experience replay buffer preserving high-value exploration trajectories. Through a hybrid On-Policy and Off-Policy update strategy, this framework leverages explicit memory for immediate exploration while encoding successful experiences into network parameters via offline updates, ensuring agents sustain optimal performance through internalized "intuition" without external support.

C. Datasets

Recently, the research community has developed various datasets to evaluate the consistency over the long term and the capacity for self evolution of LLM agents within dynamic environments. However, existing benchmarks still primarily assess the storage and retrieval of static data, which results in a lack of evaluation for other critical capabilities of memory within scenarios of dynamic interaction. Following our proposed path of evolution (§4), we categorize these benchmarks into stages of storage, reflection, and experience according to their primary areas of focus. Detailed information for these benchmarks is provided in Table 1.

Storage Stage. The storage stage serves as the cornerstone for mechanisms of memory, primarily evaluating the capacity of LLM agents for the accurate storage and retrieval of information over long distances across various scenarios, tasks, and modalities.

- **Extreme Context:** Extreme context types focus on probing the physical limits of memory in LLM agents, specifically the capacity for extracting and processing minute facts within massive volumes of distracting information. For instance, these benchmarks define the actual effective window of the model through the retrieval of multiple needles Hsieh et al. (2024), assess the capabilities of agents by embedding reasoning tasks within backgrounds of a million words Kuratov et al. (2024) and assessing the reliability of memory within a long context Yen et al. (2024), or extend these challenges to the domain of vision Wang et al. (2024). The core of this area is the evaluation of the authentic capacity for memory in the model.
- **Interactive Consistency:** Research regarding the category of Interactive Consistency is based on interaction across sessions, which requires LLM agents to maintain memory with consistency throughout such interactions. Examples include the provision of frameworks for coherent dialogue at the scale of ten million words Tavakoli et al. (2025), the direct evaluation of the update of knowledge and the capacity for rejection during continuous interaction Maharana et al. (2024), and the detection of how consistency and accuracy for personas are maintained over histories of long duration (Jia et al. 2025; Zhong et al. 2023). The core of this stage is the assessment of the capacity for memory with consistency over long distances.
- **Relational Fact:** Benchmarks of the relational fact category primarily evaluate the capacity of LLM agents for semantic association and reasoning across multiple hops. This involves testing the ability of the model for the integration of facts across documents and reasoning in multiple steps within the context of personal trivia (Yang et al. 2018; Zhang et al. 2025). Furthermore, certain frameworks focus on emotional support and interactive scenarios to evaluate the model's capacity for memory recall across proactive and passive paradigms He et al. (2024).

Reflection Stage. The core of the Reflection stage is the evaluation of how agents transform raw trajectories into memory of high quality, with an emphasis on the denoising and fidelity of memory, the deep alignment with characteristics of users, and the support for perception within complex environments.

- **Error Correction:** Error correction primarily evaluates whether errors or hallucinations emerge during the lifecycle of the memory system. For instance, it involves testing operations for the search, editing, and matching of memory Xia et al. (2025), examining the presence of hallucinations during the stages of extraction or update Chen et al. (2025).

- **Personalization:** Personalization focuses on the capacity for the extraction of deep personalization from the history of the agent, which includes the mining of latent information through reflection to identify implicit preferences (Huang et al. 2025; Jiang et al. 2025), traits of users (Du et al. 2024; Zhao et al. 2025), key information (Li et al. 2025; Yuan et al. 2023), and shared components (Kim et al. 2024; Tsaknakis et al. 2025).
- **Dynamic Reasoning:** Dynamic reasoning emphasizes the critical role of memory in reasoning across multiple steps and the perception of environments with high complexity. This involves the selective forgetting of memory (Hu et al. (2025)), backtracking on decisions (Wan and Ma (2025)), scenarios in the real world (Deng et al. 2024; Miyai et al. 2025), and the mechanisms for summarization and transition (Maharana et al. (2024)).

Experience Stage. The Experience Stage represents the pinnacle of the evolutionary path of memory mechanisms; at this phase, the focus shifts toward how LLM agents abstract general experience from fragmented trajectories within dynamic environments to facilitate continuous evolution through practical application. While benchmarks for this stage are relatively scarce, they possess a strong empirical character: Wu et al. (2024) and Ai et al. (2025) simulate environments for authentic deployment to evaluate the capacity of LLM agents for the extraction and internalization of experience within cycles of input and feedback; conversely, Wei et al. (2025) and Zheng et al. (2025) emphasize the capacity for the transfer of experience, measuring levels of abstraction and generalization by assessing the transfer of acquired experience to a diverse range of other tasks.

Table 1. Representative datasets for LLM agent memory mechanism research.

Stage	Dataset	Reference	Size	Description
Storage Stage Benchmark	LongBench	Bai et al. (2023)	4.7k	Evaluate faithful memory preservation and retrieval by performing information extraction and reasoning across multiple tasks with sequences up to 32k tokens.
	LongBenchv2	Bai et al. (2024)	503	Answer complex multiple-choice questions through the processing of extremely long sequences with lengths between 8k and 2M words for the purpose of evaluating the capacity of memory and the precision of reasoning.
	RULER	Hsieh et al. (2024)	Scalable	Evaluate the effectiveness of retrieval and synthesis within long contexts through tasks such as multi-needle extraction or multi-hop reasoning to identify true memory capacity.
	MMNeedle	Wang et al. (2024)	280k	Identify a target sub image within a massive collection of images through the analysis of textual descriptions and visual contents for the purpose of measuring the limits of multimodal retrieval.
	HotpotQA	Yang et al. (2018)	113k	Answer questions by performing reasoning across multiple hops over information scattered in diverse documents based on Wikipedia to provide accurate results and supporting facts.
	MemoryBank	Zhong et al. (2023)	194	Answer questions by recalling pertinent information and summarizing user traits across a history of interactions spanning ten days to evaluate the precision of retrieval and the maintenance of user portraits for long-term dialogues.
	BABILong	Kurатов et al. (2024)	Scalable	Answer questions by performing reasoning on facts scattered across extremely long documents of natural language to test the limits of memory and retrieval for contexts with length up to one million tokens.
	DialSim	Kim et al. (2024)	1.3k	Evaluate the precision of retrieval for memory by answering spontaneous questions across sessions of dialogue involving multiple parties with long durations.
	LongMemEval	Maharana et al. (2024)	500	Answer questions through the extraction of information from histories of interactive chat with multiple sessions for the purpose of evaluating the performance of retrieval and reasoning across dependencies of long range.
	BEAM	Tavakoli et al. (2025)	100	Evaluate the capacity of memory and the precision of retrieval by answering questions based on coherent and topically diverse conversations with length up to ten million tokens.
	MPR	Zhang et al. (2025)	108k	Answer complex questions by conducting reasoning across multiple hops of factual information specific to a user within a framework of explicit or implicit memory for the purpose of evaluating the precision of retrieval.
	LOCCO	Jia et al. (2025)	1.6k	Evaluate the persistence of memory and the retrieval of historical facts by analyzing chronological conversations across extended periods of time for the purpose of measuring information decay.
	MADial-Bench	He et al. (2024)	160	Evaluate the effectiveness of retrieval and recognition for historical events across multiple turns of interaction by simulating paradigms of passive and proactive recall for the purpose of providing emotional support.
	HELMET	Yen et al. (2024)	Scalable	Evaluate the effectiveness of models for long-context by performing information extraction and reasoning across seven diverse categories for sequences with lengths up to 128k tokens to provide a thorough assessment of memory capacity.

Table 1. Representative datasets for LLM agent memory mechanism research (continued).

Stage	Dataset	Reference	Size	Description
Reflection Stage Benchmark	Minerva	Xia et al. (2025)	Scalable	Analyze the proficiency of LLMs in utilizing and manipulating context memory through a programmable framework of atomic and composite tasks for the purpose of pinpointing specific functional deficiencies and providing actionable insights.
	HaluMem	Chen et al. (2025)	3.5k	Evaluate the fidelity of memory by quantifying the occurrence of fabrication and omission during the stages of storage and retrieval across dialogues of multiple turns.
	MABench	Hu et al. (2025)	2.1k	Evaluate the competencies of accurate retrieval and learning at test time across sequences of incremental interactions with multiple turns.
	PRM	Yuan et al. (2023)	700	Evaluate the capability of personalized assistants to maintain a dynamic memory bank by preserving evolving user knowledge and experiences across long term dialogues for the purpose of generating tailored responses.
	PersonMemv2	Jiang et al. (2025)	1k	Generate personalized responses through the extraction of implicit personas from interactions with long context and thousands of preferences of users to evaluate the adaptation of agents.
	LoCoMo	Maharana et al. (2024)	50	Evaluate the reliability of memory by executing question answering and event summarization across sequences of conversation with lengths of up to nine thousand tokens spanning thirty-five sessions.
	WebChoreArena	Miyai et al. (2025)	451	Analyze the performance of memory for information retrieval and complex aggregation by performing multiple steps of navigation and reasoning across hundreds of web pages.
	MT-Mind2Web	Deng et al. (2024)	720	Evaluate the performance of conversational web navigation via multiturn instruction following across sequential interactions with both users and environment for the purpose of managing context dependency and limited memory space.
	StoryBench	Wan and Ma (2025)	80	Evaluate the capacity for reflection and sequential reasoning by navigating hierarchical decision trees within interactive fiction games to trace back and revise earlier choices independently across multiple turns.
	PerLTQA	Du et al. (2024)	8.6k	Answer personalized questions by retrieving and synthesizing semantic and episodic information from a memory bank of thirty characters to evaluate the accuracy of retrieval for memory.
	ImplexConv	Li et al. (2025)	2.5k	Evaluate implicit reasoning in personalized dialogues by retrieving and synthesizing subtle or semantically distant information from history encompassing 100 sessions to test the efficiency of hierarchical memory structures.
	Share	Kim et al. (2024)	3.2k	Improve the quality of interactions across long durations by extracting persona data and memories of shared history from scripts of movies to sustain a consistent relationship between two individuals.
	Mem-PAL	Huang et al. (2025)	100	Evaluate the capability of personalization for assistants oriented toward service by identifying subjective traits and preferences of users from histories of dialogue and behavioral logs across multiple sessions for the purpose of generating tailored responses.
	PrefEval	Zhao et al. (2025)	3k	Quantify the robustness of proactive preference following by evaluating the ability of models to infer and satisfy explicit or implicit user traits amidst long context distractions for sequences with length up to 100k tokens.
LIDB	Tsaknakis et al. (2025)	210	Discover the latent preferences of users and generate personalized responses through interactions across multiple turns within a framework of three agents for the purpose of evaluating the efficiency of elicitation and adaptation.	
Experience Stage Benchmark	StreamBench	Wu et al. (2024)	9.7k	Evaluate the capacity for continuous improvement and online learning via iterative feedback processing across diverse task streams to measure the adaptation of agents after deployment.
	MemoryBench	Ai et al. (2025)	20k	Evaluate the capacity for continual learning of Large Language Model systems by simulating the accumulation of feedback from users across multiple domains to measure the effectiveness of procedural memory.
	Evo-Memory	Wei et al. (2025)	3.7k	Evaluate the capacity for learning at test time and the evolution of memory by processing continuous streams of tasks for the purpose of reuse of experience across diverse scenarios.
	LABench	Zheng et al. (2025)	1.4k	Evaluate the lifelong learning ability and transfer of knowledge across sequences of interdependent tasks in dynamic environments for the purpose of measuring the acquisition and retention of skills.

References

- Ahn, Michael, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alexander Herzog, Daniel Ho, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Eric Jang, Rosario M Jauregui Ruano, Kyle Jeffrey, Sally Jesmonth, Nikhil Jayant Joshi, Ryan C. Julian, Dmitry Kalashnikov, Yuheng Kuang, Kuang-Huei Lee, Sergey Levine, Yao Lu, Linda Luu, Carolina Parada, Peter Pastor, Jornell Quiambao, Kanishka Rao, Jarek Rettinghouse, Diego M Reyes, Pierre Sermanet, Nicolas Sievers, Clayton Tan, Alexander Toshev, Vincent Vanhoucke, F. Xia, Ted Xiao, Peng Xu, Sichun Xu, and Mengyuan Yan. 2022. Do as i can, not as i say: Grounding language in robotic affordances. In *Conference on Robot Learning*.
- Ai, Qingyao, Yichen Tang, Changyue Wang, Jianming Long, Weihang Su, and Yiqun Liu. 2025. Memorybench: A benchmark for memory and continual learning in llm systems. *ArXiv abs/2510.17281*.
- Alakuijala, Minttu, Ya Gao, Georgy Ananov, Samuel Kaski, Pekka Marttinen, Alexander Ilin, and Harri Valpola. 2025. Memento no more: Coaching ai agents to master multiple tasks via hints internalization. *ArXiv abs/2502.01562*.

- Anokhin, Petr, Nikita Semenov, Artyom Y. Sorokin, Dmitry Evseev, Mikhail Burtsev, and Evgeny Burnaev. 2024. Arigraph: Learning knowledge graph world models with episodic memory for llm agents. In *International Joint Conference on Artificial Intelligence*.
- Anonymous. 2025. Exploratory memory-augmented LLM agent via hybrid on- and off-policy optimization. In *Submitted to The Fourteenth International Conference on Learning Representations*. under review.
- Bai, Yushi, Xin Lv, Jiajie Zhang, Hong Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. 2023. Longbench: A bilingual, multitask benchmark for long context understanding. *ArXiv abs/2308.14508*.
- Bai, Yushi, Shangqing Tu, Jiajie Zhang, Hao Peng, Xiaozhi Wang, Xin Lv, Shulin Cao, Jiazheng Xu, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. 2024. Longbench v2: Towards deeper understanding and reasoning on realistic long-context multitasks. *ArXiv abs/2412.15204*.
- Balestri, Roberto and Guglielmo Pescatore. 2025. Narrative memory in machines: Multi-agent arc extraction in serialized tv. *ArXiv abs/2508.07010*.
- Behrouz, Ali, Peilin Zhong, and Vahab S. Mirrokni. 2024. Titans: Learning to memorize at test time. *ArXiv abs/2501.00663*.
- Bidochko, Andrii and Yaroslav Vykylyuk. 2026. Thought management system for long-horizon, goal-driven llm agents. *Journal of Computational Science* 93, 102740. <https://doi.org/https://doi.org/10.1016/j.jocs.2025.102740>.
- Bo, Xiaohe, Zeyu Zhang, Quanyu Dai, Xueyang Feng, Lei Wang, Rui Li, Xu Chen, and Ji-Rong Wen. 2024. Reflective multi-agent collaboration based on large language models. *Advances in Neural Information Processing Systems* 37.
- Bohnet, Bernd, P. Kamienny, Hanie Sedghi, Dilan Gorur, Pranjal Awasthi, Aaron T Parisi, Kevin Swersky, Rosanne Liu, Azade Nova, and Noah Fiedel. 2025. Enhancing llm planning capabilities through intrinsic self-critique.
- Cai, Zhicheng, Xinyuan Guo, Yu Pei, Jiangtao Feng, Jiangjie Chen, Ya-Qin Zhang, Wei-Ying Ma, Mingxuan Wang, and Hao Zhou. 2025. Flex: Continuous agent evolution via forward learning from experience.
- Cao, Zouying, Jiaji Deng, Li Yu, Weikang Zhou, Zhaoyang Liu, Bolin Ding, and Hai Zhao. 2025. Remember me, refine me: A dynamic procedural memory framework for experience-driven agent evolution.
- Chen, Ding, Simin Niu, Kehang Li, Peng Liu, Xiangping Zheng, Bo Tang, Xinchu Li, Feiyu Xiong, and Zhiyu Li. 2025. Halumem: Evaluating hallucinations in memory systems of agents. *ArXiv abs/2511.03506*.
- Chen, Silin, Shaoxin Lin, Xiaodong Gu, Yuling Shi, Heng Lian, Longfei Yun, Dong Chen, Weiguo Sun, Lin Cao, and Qianxiang Wang. 2025. Swe-exp: Experience-driven software issue resolution. *ArXiv abs/2507.23361*.
- Chen, Shiqi, Tongyao Zhu, Zian Wang, Jinghan Zhang, Kangrui Wang, Siyang Gao, Teng Xiao, Yee Whye Teh, Junxian He, and Manling Li. 2025. Internalizing world models via self-play finetuning for agentic rl. *ArXiv abs/2510.15047*.
- Cheng, Jiali, Anjishnu Kumar, Roshan Lal, Rishi Rajasekaran, Hani Ramezani, Omar Zia Khan, Oleg Rokhlenko, Sunny Chiu-Webster, Gang Hua, and Hadi Amiri. 2025. Webatlas: An llm agent with experience-driven memory and action simulation.
- Cheng, Mingyue, Ouyang Jie, Shuo Yu, Ruiran Yan, Yucong Luo, Zirui Liu, Daoyu Wang, Qi Liu, and Enhong Chen. 2025. Agent-r1: Training powerful llm agents with end-to-end reinforcement learning.
- Chhikara, Prateek, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. 2025. Mem0: Building production-ready ai agents with scalable long-term memory. *ArXiv abs/2504.19413*.
- Cui, Shaobo, Luca Mouchel, and Boi Faltings. 2025. Uncertainty in causality: A new frontier. In *Annual Meeting of the Association for Computational Linguistics*.
- Das, Payel, Subhajit Chaudhury, Elliot Nelson, Igor Melnyk, Sarath Swaminathan, Sihui Dai, Aurélie C. Lozano, Georgios Kollias, Vijil Chenthamarakshan, Jirí Navrátil, Soham Dan, and Pin-Yu Chen. 2024. Larimar: Large language models with episodic memory control. *ArXiv abs/2403.11901*.
- Deng, Yang, Xuan Zhang, Wenxuan Zhang, Yifei Yuan, See-Kiong Ng, and Tat-Seng Chua. 2024. On the multi-turn instruction following for conversational web agents. In *Annual Meeting of the Association for Computational Linguistics*.
- Ding, Bowen, Yuhan Chen, Jiayang Lv, Jiyao Yuan, Qi Zhu, Shuangshuang Tian, Dantong Zhu, Futing Wang, Heyuan Deng, Fei Mi, Lifeng Shang, and Tao Lin. 2025. Rethinking expert trajectory utilization in llm post-training.
- Du, Xingbo, Loka Li, Duzhen Zhang, and Le Song. 2025a. Mem^{r3}: Memory retrieval via reflective reasoning for llm agents.

- Du, Xingbo, Loka Li, Duzhen Zhang, and Le Song. 2025b. Memr³: Memory retrieval via reflective reasoning for llm agents.
- Du, Yiming, Wenyu Huang, Danna Zheng, Zhaowei Wang, Sébastien Montella, Mirella Lapata, Kam-Fai Wong, and Jeff Z. Pan. 2025. Rethinking memory in ai: Taxonomy, operations, topics, and future directions. *ArXiv abs/2505.00675*.
- Du, Yiming, Baojun Wang, Yifan Xiang, Zhaowei Wang, Wenyu Huang, Boyang Xue, Bin Liang, Xingshan Zeng, Fei Mi, Haoli Bai, Lifeng Shang, Jeff Z. Pan, Yuxin Jiang, and Kam-Fai Wong. 2025. Memory-t1: Reinforcement learning for temporal reasoning in multi-session agents.
- Du, Yiming, Hongru Wang, Zhengyi Zhao, Bin Liang, Baojun Wang, Wanjun Zhong, Zezhong Wang, and Kam-Fai Wong. 2024. PerlTqa: A personal long-term memory dataset for memory classification, retrieval, and synthesis in question answering. *ArXiv abs/2402.16288*.
- Everitt, Tom, Cristina Garbacea, Alexis Bellot, Jonathan Richens, Henry Papadatos, Siméon Campos, and Rohin Shah. 2025. Evaluating the goal-directedness of large language models. *ArXiv abs/2504.11844*.
- Fang, Jizhan, Xinle Deng, Haoming Xu, Ziyang Jiang, Yuqi Tang, Ziwen Xu, Shumin Deng, Yunzhi Yao, Mengru Wang, Shuofei Qiao, Huajun Chen, and Ningyu Zhang. 2025. Lightmem: Lightweight and efficient memory-augmented generation. *ArXiv abs/2510.18866*.
- Fang, Runnan, Yuan Liang, Xiaobin Wang, Jialong Wu, Shuofei Qiao, Pengjun Xie, Fei Huang, Huajun Chen, and Ningyu Zhang. 2025. Memp: Exploring agent procedural memory. *ArXiv abs/2508.06433*.
- Feng, Lang, Zhenghai Xue, Tingcong Liu, and Bo An. 2025. Group-in-group policy optimization for llm agent training. *ArXiv abs/2505.10978*.
- Feng, Tongtong, Xin Wang, Yu-Gang Jiang, and Wenwu Zhu. 2025. Embodied ai: From llms to world models. *ArXiv abs/2509.20021*.
- Forouzandeh, Saman, Wei Peng, Parham Moradi, Xinghuo Yu, and Mahdi Jalili. 2025. Learning hierarchical procedural memory for llm agents through bayesian selection and contrastive refinement.
- Fu, Dayuan, Keqing He, Yejie Wang, Wentao Hong, Zhuoma Gongque, Weihao Zeng, Wei Wang, Jingang Wang, Xunliang Cai, and Weiran Xu. 2025. Agentrefine: Enhancing agent generalization through refinement tuning. *ArXiv abs/2501.01702*.
- Gao, Dawei, Zitao Li, Weirui Kuang, Xuchen Pan, Daoyuan Chen, Zhijian Ma, Bingchen Qian, Liuyi Yao, Lin Zhu, Chen Cheng, Hongzhu Shi, Yaliang Li, Bolin Ding, and Jingren Zhou. 2024. Agentscope: A flexible yet robust multi-agent platform. *ArXiv abs/2402.14034*.
- Gharat, Himanshu, Himanshi Agrawal, and Gourab K. Patro. 2025. From personalization to prejudice: Bias and discrimination in memory-enhanced ai agents for recruitment.
- Ghasemabadi, Amirhosein and Di Niu. 2025. Can llms predict their own failures? self-awareness via internal circuits.
- Guo, Jiacheng, Ling Yang, Peter Chen, Qixin Xiao, Yinjie Wang, Xinzhe Juan, Jiahao Qiu, Ke Shen, and Mengdi Wang. 2025. Genenv: Difficulty-aligned co-evolution between llm agents and environment simulators.
- Han, Dongge, Camille Couturier, Daniel Madrigal Diaz, Xuchao Zhang, Victor Rühle, and Saravan Rajmohan. 2025. Legomem: Modular procedural memory for multi-agent llm systems for workflow automation. *ArXiv abs/2510.04851*.
- Hassell, Jackson, Dan Zhang, Han Jun Kim, Tom Mitchell, and Estevam Hruschka. 2025. Learning from supervision with semantic and episodic memory: A reflective approach to agent adaptation. *ArXiv abs/2510.19897*.
- Hayashi, Hiroaki, Bo Pang, Wenting Zhao, Ye Liu, Akash Gokul, Srijan Bansal, Caiming Xiong, Semih Yavuz, and Yingbo Zhou. 2025. Self-abstraction from grounded experience for plan-guided policy refinement.
- He, Junqing, Liang Zhu, Rui Wang, Xi Wang, Gholamreza Haffari, and Jiaying Zhang. 2024. Madial-bench: Towards real-world evaluation of memory-augmented dialogue generation. In *North American Chapter of the Association for Computational Linguistics*.
- He, Kaiyu, Mian Zhang, Shuo Yan, Peilin Wu, and Zhiyu Chen. 2024. Idea: Enhancing the rule learning ability of large language model agent through induction, deduction, and abduction. In *Annual Meeting of the Association for Computational Linguistics*.
- He, Xingqi, Yujie Zhang, Shuyong Gao, Wenjie Li, Lingyi Hong, Mingxi Chen, Kaixun Jiang, Jiyan Fu, and Wenqiang Zhang. 2025. Rsgent: Learning to reason and act for text-guided segmentation via multi-turn tool invocations.
- Ho, Matthew, Chen Si, Zhaoxiang Feng, Fangxu Yu, Yichi Yang, Zhijian Liu, Zhiting Hu, and Lianhui Qin. 2025. Arcmemo: Abstract reasoning composition with lifelong llm memory. *ArXiv abs/2509.04439*.

- Hong, Chuanyang and Qingyun He. 2025. Enhancing memory retrieval in generative agents through llm-trained cross attention networks. *Frontiers in Psychology* 16.
- Hou, Yuki, Haruki Tamoto, and Homei Miyashita. 2024. "my agent understands me better": Integrating dynamic human-like memory recall and consolidation in llm-based agents. *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*.
- Houichime, Tarik, Abdelghani Souhar, and Younès El Amrani. 2025. Memory as resonance: A biomimetic architecture for infinite context memory on ergodic phonetic manifolds.
- Hsieh, Cheng-Ping, Simeng Sun, Samuel Krizan, Shantanu Acharya, Dima Rekish, Fei Jia, and Boris Ginsburg. 2024. Ruler: What's the real context size of your long-context language models? *ArXiv abs/2404.06654*.
- Hu, Chenxu, Jie Fu, Chenzhuang Du, Simian Luo, Junbo Jake Zhao, and Hang Zhao. 2023. Chatdb: Augmenting llms with databases as their symbolic memory. *ArXiv abs/2306.03901*.
- Hu, Mengkang, Tianxing Chen, Qiguang Chen, Yao Mu, Wenqi Shao, and Ping Luo. 2024. Hiagent: Hierarchical working memory management for solving long-horizon agent tasks with large language model. *ArXiv abs/2408.09559*.
- Hu, Yuyang, Shichun Liu, Yanwei Yue, Guibin Zhang, Boyang Liu, Fangyi Zhu, Jiahang Lin, Honglin Guo, Shihan Dou, Zhiheng Xi, Senjie Jin, Jiejun Tan, Yanbin Yin, Jiongnan Liu, Zeyu Zhang, Zhongxiang Sun, Yutao Zhu, Hao Sun, Boci Peng, Zhenrong Cheng, Xuanbo Fan, Jiabin Guo, Xinlei Yu, Zhenhong Zhou, Zewen Hu, Jiahao Huo, Junhao Wang, Yuwei Niu, Yu Wang, Zhenfei Yin, Xiaobin Hu, Yue Liao, Qiankun Li, Kun Wang, Wangchunshu Zhou, Yixin Liu, Dawei Cheng, Qi Zhang, Tao Gui, Shirui Pan, Yan Zhang, Philip Torr, Zhicheng Dou, Ji-Rong Wen, Xuanjing Huang, Yu-Gang Jiang, and Shuicheng Yan. 2025. Memory in the age of ai agents.
- Hu, Yuanzhe, Yu Wang, and Julian McAuley. 2025. Evaluating memory in llm agents via incremental multi-turn interactions. *ArXiv abs/2507.05257*.
- Huang, Jie, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. 2023. Large language models cannot self-correct reasoning yet. *ArXiv abs/2310.01798*.
- Huang, Xu, Junwu Chen, Yuxing Fei, Zhuohan Li, Philippe Schwaller, and Gerbrand Ceder. 2025. Cascade: Cumulative agentic skill creation through autonomous development and evolution.
- Huang, Xu, Weiwen Liu, Xiaolong Chen, Xingmei Wang, Hao Wang, Defu Lian, Yasheng Wang, Ruiming Tang, and Enhong Chen. 2024. Understanding the planning of llm agents: A survey. *ArXiv abs/2402.02716*.
- Huang, Yizhe, Yang Liu, Ruiyu Zhao, Xiaolong Zhong, Xingming Yue, and Ling Jiang. 2025. Memorb: A plug-and-play verbal-reinforcement memory layer for e-commerce customer service. *ArXiv abs/2509.18713*.
- Huang, Yunpeng, Jingwei Xu, Zixu Jiang, Junyu Lai, Zenan Li, Yuan Yao, Taolue Chen, Lijuan Yang, Zhou Xin, and Xiaoxing Ma. 2023. Advancing transformer architecture in long-context large language models: A comprehensive survey. *ArXiv abs/2311.12351*.
- Huang, Zhaopei, Qifeng Dai, Guozheng Wu, Xiaopeng Wu, Kehan Chen, Chuan Yu, Xubin Li, Tiezheng Ge, Wenxuan Wang, and Qin Jin. 2025. Mem-pal: Towards memory-based personalized dialogue assistants for long-term user-agent interaction.
- Hurst, Aaron, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Madry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, Alex Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Conneau, Ali Kamali, Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoonchian, Ananya Kumar, Andrea Vallone, Andrej Karpathy, Andrew Braunstein, Andrew Cann, Andrew Codisoti, Andrew Galu, Andrew Kondrich, Andrew Tulloch, Andrey Mishchenko, Angela Baek, Angela Jiang, Antoine Pelisse, Antonia Woodford, Anuj Gosalia, Arka Dhar, Ashley Pantuliano, Avi Nayak, Avital Oliver, Barret Zoph, Behrooz Ghorbani, Ben Leimberger, Ben Rossen, Ben Sokolowsky, Ben Wang, Benjamin Zweig, Beth Hoover, Blake Samic, Bob McGrew, Bobby Spero, Bogo Giertler, Bowen Cheng, Brad Lightcap, Brandon Walkin, Brendan Quinn, Brian Guarraci, Brian Hsu, Bright Kellogg, Brydon Eastman, Camillo Lugaresi, Carroll L. Wainwright, Cary Bassin, Cary Hudson, Casey Chu, Chad Nelson, Chak Li, Chan Jun Shern, Channing Conger, Charlotte Barette, Chelsea Voss, Chen Ding, Cheng Lu, Chong Zhang, Chris Beaumont, Chris Hallacy, Chris Koch, Christian Gibson, Christina Kim, Christine Choi, Christine McLeavey, Christopher Hesse, Claudia Fischer, Clemens Winter, Coley Czarnecki, Colin Jarvis, Colin Wei, Constantin Koumouzelis, and Dane Sherburn. 2024. Gpt-4o system card. *CoRR abs/2410.21276*. <https://doi.org/10.48550/ARXIV.2410.21276>.

- Jang, Joel, Seonghyeon Ye, Changho Lee, Sohee Yang, Joongbo Shin, Janghoon Han, Gyeonghun Kim, and Minjoon Seo. 2022. Temporalwiki: A lifelong benchmark for training and evaluating ever-evolving language models. In *Conference on Empirical Methods in Natural Language Processing*.
- Jia, Zixi, Qinghua Liu, Hexiao Li, Yuyan Chen, and Jiqiang Liu. 2025, July. Evaluating the long-term memory of large language models. In W. Che, J. Nabende, E. Shutova, and M. T. Pilehvar (Eds.), *Findings of the Association for Computational Linguistics: ACL 2025*, Vienna, Austria, pp. 19759–19777. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2025.findings-acl.1014>.
- Jiang, Bowen, Yuan Yuan, Maohao Shen, Zhuoqun Hao, Zhangchen Xu, Zichen Chen, Ziyi Liu, Anvesh Rao Vijjini, Jiashu He, Hanchao Yu, Radha Poovendran, Greg Wornell, Lyle Ungar, Dan Roth, Sihao Chen, and Camillo Jose Taylor. 2025. Personamem-v2: Towards personalized intelligence via learning implicit user personas and agentic memory.
- Jiang, Dongfu, Yi Lu, Zhuofeng Li, Zhiheng Lyu, Ping Nie, Haozhe Wang, Alex Su, Hui Chen, Kai Zou, Chao Du, Tianyu Pang, and Wenhui Chen. 2025. Verltool: Towards holistic agentic reinforcement learning with tool use. *ArXiv abs/2509.01055*.
- Jiang, Huiqiang, Qianhui Wu, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2023. Llmlingua: Compressing prompts for accelerated inference of large language models. In *Conference on Empirical Methods in Natural Language Processing*.
- Jin, Hongye, Xiaotian Han, Jingfeng Yang, Zhimeng Jiang, Zirui Liu, Chia yuan Chang, Huiyuan Chen, and Xia Hu. 2024. Llm maybe longlm: Self-extend llm context window without tuning. *ArXiv abs/2401.01325*.
- Joshi, Abhinav, Areeb Ahmad, and Ashutosh Modi. 2024. Cold: Causal reasoning in closed daily activities. *ArXiv abs/2411.19500*.
- Kalai, Adam Tauman and Santosh S. Vempala. 2023. Calibrated language models must hallucinate. *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*.
- Kang, Jiazheng, Mingming Ji, Zhe Zhao, and Ting Bai. 2025. Memory os of ai agent. *ArXiv abs/2506.06326*.
- Kasai, Jungo, Keisuke Sakaguchi, Yoichi Takahashi, Ronan Le Bras, Akari Asai, Xinyan Yu, Dragomir Radev, Noah A. Smith, Yejin Choi, and Kentaro Inui. 2024. Realtime qa: What's the answer right now?
- Kim, Eunwon, Chanho Park, and Buru Chang. 2024. Share: Shared memory-aware open-domain long-term dialogue dataset constructed from movie script. *ArXiv abs/2410.20682*.
- Kim, Jiho, Woosog Chay, Hyeonji Hwang, Daeun Kyung, Hyunseung Chung, Eunbyeol Cho, Yohan Jo, and Edward Choi. 2024. Dialsim: A real-time simulator for evaluating long-term dialogue understanding of conversational agents. *ArXiv abs/2406.13144*.
- Kim, Minsoo and Seung won Hwang. 2025. Coex - co-evolving world-model and exploration. *ArXiv abs/2507.22281*.
- Ko, Dayoon, Jinyoung Kim, Hahyeon Choi, and Gunhee Kim. 2024. Growover: How can llms adapt to growing real-world knowledge? In *Annual Meeting of the Association for Computational Linguistics*.
- Kurатов, Yuri, Aydar Bulatov, Petr Anokhin, Ivan Rodkin, Dmitry Sorokin, Artyom Y. Sorokin, and Mikhail Burtsev. 2024. Babilong: Testing the limits of llms with long context reasoning-in-a-haystack. *ArXiv abs/2406.10149*.
- Latimer, Chris, Nicol'o Boschi, Andrew Neeser, Chris Bartholomew, Gaurav Srivastava, Xuan Wang, and Naren Ramakrishnan. 2025. Hindsight is 20/20: Building agent memory that retains, recalls, and reflects.
- Lazaridou, Angeliki, Adhiguna Kuncoro, Elena Gribovskaya, Devang Agrawal, Adam Liska, Tayfun Terzi, Mai Giménez, Cyprien de Masson d'Autume, Tomás Kociský, Sebastian Ruder, Dani Yogatama, Kris Cao, Susannah Young, and Phil Blunsom. 2021. Mind the gap: Assessing temporal generalization in neural language models. In *Neural Information Processing Systems*.
- Lee, Seokhan and Hanseok Ko. 2025. Training a team of language models as options to build an sql-based memory. *Applied Sciences*.
- Li, Rui, Zeyu Zhang, Xiaohe Bo, Zihang Tian, Xu Chen, Quanyu Dai, Zhenhua Dong, and Ruiming Tang. 2025. Cam: A constructivist view of agentic memory for llm-based reading comprehension. *ArXiv abs/2510.05520*.
- Li, Shilong, Yancheng He, Hangyu Guo, Xingyuan Bu, Ge Bai, Jie Liu, Jiaheng Liu, Xingwei Qu, Yangguang Li, Wanli Ouyang, Wenbo Su, and Bo Zheng. 2024. Graphreader: Building graph-based agent to enhance long-context abilities of large language models. In *Conference on Empirical Methods in Natural Language Processing*.
- Li, Xintong, Jalend Bantupalli, Ria Dharmani, Yuwei Zhang, and Jingbo Shang. 2025. Toward multi-session personalized conversation: A large-scale dataset and hierarchical tree framework for implicit reasoning. *ArXiv abs/2503.07018*.

- Li, Xiaoxi, Wenxiang Jiao, Jiarui Jin, Guanting Dong, Jiajie Jin, Yinuo Wang, Hao Wang, Yutao Zhu, Ji-Rong Wen, Yuan Lu, and Zhicheng Dou. 2025. Deepagent: A general reasoning agent with scalable toolsets. *ArXiv abs/2510.21618*.
- Li, Ziyue, Yuan Chang, Gaihong Yu, and Xiaoqiu Le. 2025. Hiplan: Hierarchical planning for llm-based agents with adaptive global-local guidance. *ArXiv abs/2508.19076*.
- Li, Zhiyu, Shichao Song, Chenyang Xi, Hanyu Wang, Chen Tang, Simin Niu, Ding Chen, Jiawei Yang, Chunyu Li, Qingchen Yu, Jihao Zhao, Yezhaohui Wang, Peng Liu, Zehao Lin, Pengyuan Wang, Jiahao Huo, Tianyi Chen, Kai Chen, Ke-Rong Li, Zhenzhen Tao, Junpeng Ren, Huayi Lai, Hao Wu, Bo Tang, Zhenren Wang, Zhaoxin Fan, Ningyu Zhang, Linfeng Zhang, Junchi Yan, Ming-Zhou Yang, Tong Xu, Wei Xu, Huajun Chen, Haofeng Wang, Hongkang Yang, Wentao Zhang, Zhikun Xu, Siheng Chen, and Feiyu Xiong. 2025. Memos: A memory os for ai system. *ArXiv abs/2507.03724*.
- Liang, Jiafeng, Hao Li, Chang Li, Jiaqi Zhou, Shixin Jiang, Zekun Wang, Changkai Ji, Zhihao Zhu, Runxuan Liu, Taolin Ren, Jinlan Fu, See-Kiong Ng, Xia Liang, Ming Liu, and Bing Qin. 2025. Ai meets brain: Memory systems from cognitive neuroscience to autonomous agents.
- Liao, Callie C., Duoduo Liao, and Sai Surya Gadiraju. 2025. Agentmaster: A multi-agent conversational framework using a2a and mcp protocols for multimodal information retrieval and analysis. *ArXiv abs/2507.21105*.
- Lin, Bill Yuchen, Yicheng Fu, Karina Yang, Prithviraj Ammanabrolu, Faeze Brahma, Shiyu Huang, Chandra Bhagavatula, Yejin Choi, and Xiang Ren. 2023. Swiftsage: A generative agent with fast and slow thinking for complex interactive tasks. *ArXiv abs/2305.17390*.
- Liu, Jun, Zhenglun Kong, Changdi Yang, Fan Yang, Tianqi Li, Peiyan Dong, Joannah Nanjey, Hao Tang, Geng Yuan, Wei Niu, Wenbin Zhang, Pu Zhao, Xue Lin, Dong-Xu Huang, and Yanzhi Wang. 2025. Rcr-router: Efficient role-aware context routing for multi-agent llm systems with structured memory. *ArXiv abs/2508.04903*.
- Liu, Junming, Yifei Sun, Weihua Cheng, Haodong Lei, Yirong Chen, Licheng Wen, Xueming Yang, Daocheng Fu, Pinlong Cai, Nianchen Deng, Yi Yu, Shuyue Hu, Botian Shi, and Ding Wang. 2025. Memverse: Multimodal memory for lifelong learning agents.
- Liu, Jiateng, Zhenhailong Wang, Xiaojiang Huang, Yingjie Li, Xing Fan, Xiang Li, Chenlei Guo, Ruhi Sarikaya, and Heng Ji. 2025. Analyzing and internalizing complex policy documents for llm agents. *ArXiv abs/2510.11588*.
- Liu, Jiaqi, Kaiwen Xiong, Peng Xia, Yiyang Zhou, Haonian Ji, Lu Feng, Siwei Han, Mingyu Ding, and Huaxiu Yao. 2025. Agent0-v1: Exploring self-evolving agent for tool-integrated vision-language reasoning.
- Liu, Lei, Xiaoyan Yang, Yue Shen, Binbin Hu, Zhiqiang Zhang, Jinjie Gu, and Guannan Zhang. 2023. Think-in-memory: Recalling and post-thinking enable llms with long-term memory. *ArXiv abs/2311.08719*.
- Liu, Shukai, Jian Yang, Bo Jiang, Yizhi Li, Jinyang Guo, Xianglong Liu, and Bryan Dai. 2025. Context as a tool: Context management for long-horizon swe-agents.
- Liu, Weijie, Zecheng Tang, Juntao Li, Kehai Chen, and Min Zhang. 2024. Memlong: Memory-augmented retrieval for long text modeling. *ArXiv abs/2408.16967*.
- Liu, Xiaoyu, Paiheng Xu, Junda Wu, Jiabin Yuan, Yifan Yang, Yuhang Zhou, Fuxiao Liu, Tianrui Guan, Haoliang Wang, Tong Yu, Julian J. McAuley, Wei Ai, and Furong Huang. 2025. Large language models and causal inference in collaboration: A comprehensive survey. *ArXiv abs/2403.09606*.
- Liu, Yitao, Chenglei Si, Karthik R. Narasimhan, and Shunyu Yao. 2025. Contextual experience replay for self-improvement of language agents. In *Annual Meeting of the Association for Computational Linguistics*.
- Long, Xiaoxiao, Qingrui Zhao, Kaiwen Zhang, Zihao Zhang, Dingrui Wang, Yumeng Liu, Zhengjie Shu, Yi Lu, Shouzheng Wang, Xinzhe Wei, Wei Li, Wei Yin, Yao Yao, Jiangtian Pan, Qiu Shen, Ruigang Yang, Xun Cao, and Qionghai Dai. 2025. A survey: Learning embodied intelligence from physical simulators and world models. *ArXiv abs/2507.00917*.
- Lu, Junru, Siyu An, Mingbao Lin, Gabriele Pergola, Yulan He, Di Yin, Xing Sun, and Yunsheng Wu. 2023. Memochat: Tuning llms to use memos for consistent long-range open-domain conversation. *ArXiv abs/2308.08239*.
- Luo, Hongyin, Nathaniel Morgan, Tina Li, Derek Zhao, Ai Vy Ngo, Philip Schroeder, Lijie Yang, Assaf Ben-Kish, Jack O'Brien, and James R. Glass. 2025. Beyond context limits: Subconscious threads for long-horizon reasoning. *ArXiv abs/2507.16784*.
- Luo, Xufang, Yuge Zhang, Zhiyuan He, Zilong Wang, Siyun Zhao, Dongsheng Li, Luna K. Qiu, and Yuqing Yang. 2025. Agent lightning: Train any ai agents with reinforcement learning. *ArXiv abs/2508.03680*.
- Luo, Ziyang, Zhiqi Shen, Wenzhuo Yang, Zirui Zhao, Prathyusha Jwalapuram, Amrita Saha, Doyen Sahoo, Silvio Savarese, Caiming Xiong, and Junnan Li. 2025. Mcp-universe: Benchmarking large language models with real-world model context protocol servers. *CoRR abs/2508.14704*. <https://doi.org/10.48550/ARXIV.2508.14704>.

- Luu, Kelvin, Daniel Khashabi, Suchin Gururangan, Karishma Mandyam, and Noah A. Smith. 2022. Time waits for no one! analysis and challenges of temporal misalignment.
- Lyu, Yuanjie, Chengyu Wang, Jun Huang, and Tong Xu. 2025. From correction to mastery: Reinforced distillation of large language model agents. *ArXiv abs/2509.14257*.
- Maharana, Adyasha, Dong-Ho Lee, S. Tulyakov, Mohit Bansal, Francesco Barbieri, and Yuwei Fang. 2024. Evaluating very long-term conversational memory of llm agents. *ArXiv abs/2402.17753*.
- Majumder, Bodhisattwa Prasad, Bhavana Dalvi, Peter Alexander Jansen, Oyvind Tafjord, Niket Tandon, Li Zhang, Chris Callison-Burch, and Peter Clark. 2023. Clin: A continually learning language agent for rapid task adaptation and generalization. *ArXiv abs/2310.10134*.
- Melz, Eric. 2023. Enhancing llm intelligence with arm-rag: Auxiliary rationale memory for retrieval augmented generation. *ArXiv abs/2311.04177*.
- Miyai, Atsuyuki, Zaiying Zhao, Kazuki Egashira, Atsuki Sato, Tatsumi Sunada, Shota Onohara, Hiromasa Yamanishi, Mashiro Toyooka, Kunato Nishina, Ryoma Maeda, Kiyoharu Aizawa, and T. Yamasaki. 2025. Webchorearena: Evaluating web browsing agents on realistic tedious web tasks. *ArXiv abs/2506.01952*.
- Modarressi, Ali, Abdullatif Köksal, Ayyoob Imani, Mohsen Fayyaz, and Hinrich Schutze. 2024. Memllm: Finetuning llms to use an explicit read-write memory. *ArXiv abs/2404.11672*.
- Nan, Jiayan, Wenquan Ma, Wenlong Wu, and Yize Chen. 2025. Nemori: Self-organizing agent memory inspired by cognitive science. *ArXiv abs/2508.03341*.
- Ouyang, Siru, Jun Yan, I-Hung Hsu, Yanfei Chen, Ke Jiang, Zifeng Wang, Rujun Han, Long T. Le, Samira Daruki, Xiangru Tang, Vishy Tirumalashetty, George Lee, Mahsan Rofouei, Hangfei Lin, Jiawei Han, Chen-Yu Lee, and Tomas Pfister. 2025. Reasoningbank: Scaling agent self-evolving with reasoning memory. *ArXiv abs/2509.25140*.
- Ozer, Onat, Grace Wu, Yuchen Wang, Daniel Dosti, Honghao Zhang, and Vivi De La Rue. 2025. Mar:multi-agent reflexion improves reasoning abilities in llms.
- Packer, Charles, Vivian Fang, Shishir G. Patil, Kevin Lin, Sarah Wooders, and Joseph Gonzalez. 2023. Memgpt: Towards llms as operating systems. *ArXiv abs/2310.08560*.
- Pan, Yiyuan, Zhe Liu, and Hesheng Wang. 2025. Wonder wins ways: Curiosity-driven exploration through multi-agent contextual calibration. *ArXiv abs/2509.20648*.
- Park, Joon Sung, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*.
- Qi, Zehan, Xiao Liu, Iat Long Iong, Hanyu Lai, Xueqiao Sun, Wenyi Zhao, Yu Yang, Xinyue Yang, Jiadai Sun, Shuntian Yao, Tianjie Zhang, Wei Xu, Jie Tang, and Yuxiao Dong. 2025. Webrl: Training llm web agents via self-evolving online curriculum reinforcement learning.
- Qin, Yujia, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2024. Toollm: Facilitating large language models to master 16000+ real-world apis. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Raman, Vishal, R VijaiAravindh, and Abhijith Ragav. 2025. Remi: A novel causal schema memory architecture for personalized lifestyle recommendation agents. *ArXiv abs/2509.06269*.
- Rasmussen, Preston, Pavlo Paliychuk, Travis Beauvais, Jack Ryan, and Daniel Chalef. 2025. Zep: A temporal knowledge graph architecture for agent memory. *ArXiv abs/2501.13956*.
- Ratner, Nir, Yoav Levine, Yonatan Belinkov, Ori Ram, Inbal Magar, Omri Abend, Ehud Karpas, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2022. Parallel context windows for large language models. In *Annual Meeting of the Association for Computational Linguistics*.
- Renze, Matthew and Erhan Guven. 2024. Self-reflection in large language model agents: Effects on problem-solving performance. *2024 2nd International Conference on Foundation and Large Language Models (FLLM)*, 516–525.
- Rezazadeh, Alireza, Zichao Li, Ange Lou, Yuying Zhao, Wei Wei, and Yujia Bao. 2025. Collaborative memory: Multi-user memory sharing in llm agents with dynamic access control. *ArXiv abs/2505.18279*.
- Salama, Rana, Jason Cai, Michelle Yuan, Anna Currey, Monica Sunkara, Yi Zhang, and Yassine Benajiba. 2025. Meminsight: Autonomous memory augmentation for llm agents. *ArXiv abs/2503.21760*.

- Shi, Yuchen, Yuzheng Cai, Siqi Cai, Zihan Xu, Lichao Chen, Yulei Qin, Zhijian Zhou, Xiang Fei, Chaofan Qiu, Xiaoyu Tan, Gang Li, Zongyi Li, Haojia Lin, Guocan Cai, Yong Mao, Yunsheng Wu, Ke Li, and Xing Sun. 2025. Youtu-agent: Scaling agent productivity with automated generation and hybrid policy optimization.
- Shi, Zijing, Meng Fang, and Ling Chen. 2025. Monte carlo planning with large language model for text-based game agents. *ArXiv abs/2504.16855*.
- Shinn, Noah, Federico Cassano, Beck Labash, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: language agents with verbal reinforcement learning. In *Neural Information Processing Systems*.
- Siyue, Zhang, Yuxiang Xue, Yiming Zhang, Xiaobao Wu, Anh Tuan Luu, and Zhao Chen. 2024. Mrag: A modular retrieval framework for time-sensitive question answering. *ArXiv abs/2412.15540*.
- Srivastava, Saksham Sahai and Haoyu He. 2025. Memorygraft: Persistent compromise of llm agents via poisoned experience retrieval.
- Sumers, Theodore R., Shunyu Yao, Karthik Narasimhan, and Thomas L. Griffiths. 2023. Cognitive architectures for language agents. *Trans. Mach. Learn. Res.* 2024.
- Sun, Haoran, Yekun Chai, Shuohuan Wang, Yu Sun, Hua Wu, and Haifeng Wang. 2025. Curiosity-driven reinforcement learning from human feedback. In *Annual Meeting of the Association for Computational Linguistics*.
- Sun, Haoran, Zekun Zhang, and Shaoning Zeng. 2025. Preference-aware memory update for long-term llm agents. *ArXiv abs/2510.09720*.
- Sun, Weiwei, Miao Lu, Zhan Ling, Kang Liu, Xuesong Yao, Yiming Yang, and Jiecao Chen. 2025. Scaling long-horizon llm agent via context-folding. *ArXiv abs/2510.11967*.
- Sun, Zhiyuan, Haochen Shi, Marc-Alexandre Côté, Glen Berseth, Xingdi Yuan, and Bang Liu. 2024. Enhancing agent learning through world dynamics modeling. *ArXiv abs/2407.17695*.
- Suzgun, Mirac, Mert Yükekönül, Federico Bianchi, Daniel Jurafsky, and James Zou. 2025. Dynamic cheatsheet: Test-time learning with adaptive memory. *ArXiv abs/2504.07952*.
- Tan, Zhen, Jun Yan, I-Hung Hsu, Rujun Han, Zifeng Wang, Long T. Le, Yiwen Song, Yanfei Chen, Hamid Palangi, George Lee, Anand Iyer, Tianlong Chen, Huan Liu, Chen-Yu Lee, and Tomas Pfister. 2025. In prospect and retrospect: Reflective memory management for long-term personalized dialogue agents. *ArXiv abs/2503.08026*.
- Tandon, Arnub, Karan Dalal, Xinhao Li, Daniel Kocaja, Marcel Rod, Sam Buchanan, Xiaolong Wang, Jure Leskovec, Sanmi Koyejo, Tatsunori Hashimoto, Carlos Guestrin, Jed McCaleb, Yejin Choi, and Yu Sun. 2025. End-to-end test-time training for long context.
- Tang, Hao, Darren Key, and Kevin Ellis. 2024. Worldcoder, a model-based llm agent: Building world models by writing code and interacting with the environment. *ArXiv abs/2402.12275*.
- Tang, Jiaming, Yilong Zhao, Kan Zhu, Guangxuan Xiao, Baris Kasikci, and Song Han. 2024. Quest: Query-aware sparsity for efficient long-context llm inference. *ArXiv abs/2406.10774*.
- Tang, Xiangru, Tianrui Qin, Tianhao Peng, Ziyang Zhou, Yanjun Shao, Tingting Du, Xinming Wei, Peng Xia, Fang Wu, He Zhu, Ge Zhang, Jiaheng Liu, Xingyao Wang, Sirui Hong, Chenglin Wu, Hao Cheng, Chi Wang, and Wangchunshu Zhou. 2025. Agent kb: Leveraging cross-domain experience for agentic problem solving. *ArXiv abs/2507.06229*.
- Tavakoli, Mohammad, Alireza Salemi, Carrie Ye, Mohamed Abdalla, Hamed Zamani, and J. Ross Mitchell. 2025. Beyond a million tokens: Benchmarking and enhancing long-term memory in llms. *ArXiv abs/2510.27246*.
- Touvron, Hugo, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. Llama: Open and efficient foundation language models. *CoRR abs/2302.13971*. <https://doi.org/10.48550/ARXIV.2302.13971>.
- Tran, Khanh-Tung, Dung Dao, Minh-Duong Nguyen, Quoc-Viet Pham, Barry O'Sullivan, and Hoang D. Nguyen. 2025. Multi-agent collaboration mechanisms: A survey of llms. *ArXiv abs/2501.06322*.
- Tsaknakis, Ioannis C., Bingqing Song, Shuyu Gan, Dongyeop Kang, Alfredo Garcia, Gaowen Liu, Charles Fleming, and Mingyi Hong. 2025. Do llms recognize your latent preferences? a benchmark for latent information discovery in personalized interaction. *ArXiv abs/2510.17132*.
- Wan, Chunhui, Xunan Dai, Zhuo Wang, Minglei Li, Yanpeng Wang, Yinan Mao, Yu Lan, and Zhiwen Xiao. 2025. Loongflow: Directed evolutionary search via a cognitive plan-execute-summarize paradigm.
- Wan, Luanbo and Weizhi Ma. 2025. Storybench: A dynamic benchmark for evaluating long-term memory with multi turns. *ArXiv abs/2506.13356*.

- Wang, Hengyi, Haizhou Shi, Shiwei Tan, Weiwei Qin, Wenyuan Wang, Tunyu Zhang, Akshay Uttama Nambi, Tanuja Ganu, and Hao Wang. 2024. Multimodal needle in a haystack: Benchmarking long-context capability of multimodal large language models. In *North American Chapter of the Association for Computational Linguistics*.
- Wang, Jiong Xiao, Qiaojing Yan, Yawei Wang, Yijun Tian, Soumya Smruti Mishra, Zhichao Xu, Megha Gandhi, Panpan Xu, and Lin Lee Cheong. 2025. Reinforcement learning for self-improving agent with skill library.
- Wang, Lei, Chengbang Ma, Xueyang Feng, Zeyu Zhang, Hao ran Yang, Jingsen Zhang, Zhi-Yang Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Ji rong Wen. 2023. A survey on large language model based autonomous agents. *Frontiers of Computer Science* 18.
- Wang, Piaohong, Motong Tian, Jiaxian Li, Yuan Liang, Yuqing Wang, Qianben Chen, Tiannan Wang, Zhicong Lu, Jiawei Ma, Yuchen Eleanor Jiang, and Wangchunshu Zhou. 2025. O-mem: Omni memory system for personalized, long horizon, self-evolving agents.
- Wang, Rongzheng, Shuang Liang, Qizhi Chen, Yihong Huang, Muquan Li, Yizhuo Ma, Dongyang Zhang, Ke Qin, and Man-Fai Leung. 2025. Graphcogent: Mitigating llms' working memory constraints via multi-agent collaboration in complex graph understanding.
- Wang, Siyuan, Zhongyu Wei, Yejin Choi, and Xiang Ren. 2024. Symbolic working memory enhances language models for complex rule application. *ArXiv abs/2408.13654*.
- Wang, Yu and Xi Chen. 2025. Mirix: Multi-agent memory system for llm-based agents. *ArXiv abs/2507.07957*.
- Wang, Zihan, Kangrui Wang, Qineng Wang, Pingyue Zhang, Linjie Li, Zhengyuan Yang, Kefan Yu, Minh Nhat Nguyen, Licheng Liu, Eli Gottlieb, Monica Lam, Yiping Lu, Kyunghyun Cho, Jiajun Wu, Fei-Fei Li, Lijuan Wang, Yejin Choi, and Manling Li. 2025. Ragen: Understanding self-evolution in llm agents via multi-turn reinforcement learning. *ArXiv abs/2504.20073*.
- Wang, Zora Zhiruo, Apurva Gandhi, Graham Neubig, and Daniel Fried. 2025. Inducing programmatic skills for agentic tasks. *ArXiv abs/2504.06821*.
- Wang, Zora Zhiruo, Jiayuan Mao, Daniel Fried, and Graham Neubig. 2024. Agent workflow memory. *ArXiv abs/2409.07429*.
- Wei, Tianxin, Noveen Sachdeva, Benjamin Coleman, Zhankui He, Yuanchen Bei, Xuying Ning, Mengting Ai, Yunzhe Li, Jingrui He, Ed Huai hsin Chi, Chi Wang, Shuo Chen, Fernando Pereira, Wang-Cheng Kang, and Derek Zhiyuan Cheng. 2025. Evo-memory: Benchmarking llm agent test-time learning with self-evolving memory.
- Wei, Zhepei, Wenlin Yao, Yao Liu, Weizhi Zhang, Qin Lu, Liang Qiu, Changlong Yu, Puyang Xu, Chao Zhang, Bing Yin, Hyokun Yun, and Lihong Li. 2025. Webagent-r1: Training web agents via end-to-end multi-turn reinforcement learning. *ArXiv abs/2505.16421*.
- Westhäußer, Rebecca, Wolfgang Minker, and Sebastian Zepf. 2025. Enabling personalized long-term interactions in llm-based agents through persistent memory and user profiles. *ArXiv abs/2510.07925*.
- Wu, Cheng-Kuang, Zhi Rui Tam, Chieh-Yen Lin, Yun-Nung Chen, and Hung yi Lee. 2024. Streambench: Towards benchmarking continuous improvement of language agents. *ArXiv abs/2406.08747*.
- Wu, Rong, Xiaoman Wang, Jianbiao Mei, Pinlong Cai, Daocheng Fu, Cheng Yang, Licheng Wen, Xuemeng Yang, Yufan Shen, Yuxin Wang, and Botian Shi. 2025. Evolver: Self-evolving llm agents through an experience-driven lifecycle. *ArXiv abs/2510.16079*.
- Wu, Shanglin et al. 2025. Memory in llm-based multi-agent systems: Mechanisms, challenges, and collective intelligence. *TechRxiv preprint*.
- Wu, Shanglin and Kai Shu. 2025, December. Memory in llm-based multi-agent systems: Mechanisms, challenges, and collective intelligence. <https://doi.org/10.36227/techrxiv.176539617.79044553/v1>.
- Wu, Yaxiong, Sheng Liang, Chen Zhang, Yichao Wang, Yongyue Zhang, Huifeng Guo, Ruiming Tang, and Yong Liu. 2025. From human memory to ai memory: A survey on memory mechanisms in the era of llms. *ArXiv abs/2504.15965*.
- Xi, Zhiheng, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, Rui Zheng, Xiaoran Fan, Xiao Wang, Limao Xiong, Qin Liu, Yuhao Zhou, Weiran Wang, Changhao Jiang, Yicheng Zou, Xiangyang Liu, Zhangyue Yin, Shihan Dou, Rongxiang Weng, Wensen Cheng, Qi Zhang, Wenjuan Qin, Yongyan Zheng, Xipeng Qiu, Xuanjing Huan, and Tao Gui. 2023. The rise and potential of large language model based agents: A survey. *ArXiv abs/2309.07864*.
- Xia, Menglin, Victor Ruehle, Saravan Rajmohan, and Reza Shokri. 2025. Minerva: A programmable memory test benchmark for language models. *ArXiv abs/2502.03358*.
- Xia, Peng, Peng Xia, Kaide Zeng, Jiaqi Liu, Can Qin, Fang Wu, Yiyang Zhou, Caiming Xiong, and Huaxiu Yao. 2025. Agent0: Unleashing self-evolving agents from zero data via tool-integrated reasoning.

- Xiao, Chaojun, Pengle Zhang, Xu Han, Guangxuan Xiao, Yankai Lin, Zhengyan Zhang, Zhiyuan Liu, Song Han, and Maosong Sun. 2024. Infillm: Training-free long-context extrapolation for llms with an efficient context memory. *Advances in Neural Information Processing Systems* 37.
- Xiao, Guangxuan, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2023. Efficient streaming language models with attention sinks. *ArXiv abs/2309.17453*.
- Xiao, Yunzhong, Yangmin Li, Hewei Wang, Yunlong Tang, and Zora Zhiruo Wang. 2025. Toolmem: Enhancing multimodal agents with learnable tool capability memory. *ArXiv abs/2510.06664*.
- Xiong, Zidi, Yuping Lin, Wenya Xie, Pengfei He, Jiliang Tang, Himabindu Lakkaraju, and Zhen Xiang. 2025. How memory management impacts llm agents: An empirical study of experience-following behavior. *ArXiv abs/2505.16067*.
- Xu, Haoran, Jiacong Hu, Ke Zhang, Lei Yu, Yuxin Tang, Xinyuan Song, Yiqun Duan, Lynn Ai, and Bill Shi. 2025. Sedm: Scalable self-evolving distributed memory for agents. *ArXiv abs/2509.09498*.
- Xu, Wujiang, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. 2025. A-mem: Agentic memory for llm agents. *ArXiv abs/2502.12110*.
- Xue, Siqiao, Caigao Jiang, Wenhui Shi, Fangyin Cheng, Keting Chen, Hongjun Yang, Zhiping Zhang, Jianshan He, Hongyang Zhang, Ganglin Wei, Wang Zhao, Fan Zhou, Danrui Qi, Hong Yi, Shaodong Liu, and Faqiang Chen. 2023. Db-gpt: Empowering database interactions with private large language models. *ArXiv abs/2312.17449*.
- Yan, Sikuan, Xiufeng Yang, Zuchao Huang, Ercong Nie, Zifeng Ding, Zonggen Li, Xiaowen Ma, Hinrich Schutze, Volker Tresp, and Yunpu Ma. 2025. Memory-r1: Enhancing large language model agents to manage and utilize memories via reinforcement learning. *ArXiv abs/2508.19828*.
- Yan, Xue, Zijing Ou, Mengyue Yang, Yan Song, Haifeng Zhang, Yingzhen Li, and Jun Wang. 2025. Memory-driven self-improvement for decision making with large language models. *ArXiv abs/2509.26340*.
- Yang, An, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jian Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025. Qwen3 technical report. *CoRR abs/2505.09388*. <https://doi.org/10.48550/ARXIV.2505.09388>.
- Yang, Ruihan, Jiangjie Chen, Yikai Zhang, Siyu Yuan, Aili Chen, Kyle Richardson, Yanghua Xiao, and Deqing Yang. 2024. Selfgoal: Your language agents already know how to achieve high-level goals. *ArXiv abs/2406.04784*.
- Yang, Wei, Jinwei Xiao, Hongming Zhang, Qingyang Zhang, Yanna Wang, and Bo Xu. 2025. Coarse-to-fine grounded memory for llm agent planning. *ArXiv abs/2508.15305*.
- Yang, Yongjin, Sinjae Kang, Juyong Lee, Dongjun Lee, Se young Yun, and Kimin Lee. 2025. Automated skill discovery for language agents through exploration and iterative feedback. *ArXiv abs/2506.04287*.
- Yang, Zhilin, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Conference on Empirical Methods in Natural Language Processing*.
- Yao, Shunyu, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. *ArXiv abs/2305.10601*.
- Yao, Shunyu, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. *ArXiv abs/2210.03629*.
- Ye, Rui, Zhongwang Zhang, Kuan Li, Huifeng Yin, Zhengwei Tao, Yida Zhao, Liangcai Su, Liwen Zhang, Zile Qiao, Xinyu Wang, Pengjun Xie, Fei Huang, Siheng Chen, Jingren Zhou, and Yong Jiang. 2025. Agentfold: Long-horizon web agents with proactive context management. *ArXiv abs/2510.24699*.
- Yen, Howard, Tianyu Gao, Minmin Hou, Ke Ding, Daniel Fleischer, Peter Izsak, Moshe Wasserblat, and Danqi Chen. 2024. Helmet: How to evaluate long-context language models effectively and thoroughly. *ArXiv abs/2410.02694*.
- Yin, Zhangyue, Qiushi Sun, Qipeng Guo, Zhiyuan Zeng, Qinyuan Cheng, Xipeng Qiu, and Xuanjing Huang. 2024. Explicit memory learning with expectation maximization. In *Conference on Empirical Methods in Natural Language Processing*.
- Yu, Simon, Gang Li, Weiyan Shi, and Pengyuan Qi. 2025. Polyskill: Learning generalizable skills through polymorphic abstraction. *ArXiv abs/2510.15863*.

- Yu, Wenhao, Zhenwen Liang, Chengsong Huang, Kishan Panaganti, Tianqing Fang, Haitao Mi, and Dong Yu. 2025. Guided self-evolving llms with minimal human supervision.
- Yuan, Ruifeng, Shichao Sun, Zili Wang, Ziqiang Cao, and Wenjie Li. 2023. Personalized large language model assistant with evolving conditional memory. In *International Conference on Computational Linguistics*.
- Yuen, Sizhe, Francisco Gomez Medina, Ting Su, Yali Du, and Adam J. Sobey. 2025. Intrinsic memory agents: Heterogeneous multi-agent llm systems through structured contextual memory. *ArXiv abs/2508.08997*.
- Zhai, Yunpeng, Shuchang Tao, Cheng Chen, Anni Zou, Ziqian Chen, Qingxu Fu, Shinji Mai, Li Yu, Jiaji Deng, Zouying Cao, Zhaoyang Liu, Bolin Ding, and Jingren Zhou. 2025. Agentevolver: Towards efficient self-evolving agent system.
- Zhang, Guibin, Haotian Ren, Chong Zhan, Zhenhong Zhou, Junhao Wang, He Zhu, Wangchunshu Zhou, and Shuicheng Yan. 2025. Memevolve: Meta-evolution of agent memory systems.
- Zhang, Genghan, Shaowei Zhu, Anjiang Wei, Zhenyu Song, Allen Nie, Zhen Jia, Nandita Vijaykumar, Yida Wang, and Kunle Olukotun. 2025. Accelopt: A self-improving llm agentic system for ai accelerator kernel optimization.
- Zhang, Gui-Min, Muxin Fu, Guancheng Wan, Miao Yu, Kun Wang, and Shuicheng Yan. 2025. G-memory: Tracing hierarchical memory for multi-agent systems. *ArXiv abs/2506.07398*.
- Zhang, Gui-Min, Muxin Fu, and Shuicheng Yan. 2025. Memgen: Weaving generative latent memory for self-evolving agents. *ArXiv abs/2509.24704*.
- Zhang, Gui-Min, Fanci Meng, Guancheng Wan, Zherui Li, Kun Wang, Zhenfei Yin, Lei Bai, and Shuicheng Yan. 2025. Latentevolve: Self-evolving test-time scaling in latent space. *ArXiv abs/2509.24771*.
- Zhang, Kai, Xiangchao Chen, Bo Liu, Tianci Xue, Zeyi Liao, Zhihan Liu, Xiyao Wang, Yuting Ning, Zhaorun Chen, Xiaohan Fu, Jian Xie, Yuxuan Sun, Boyu Gou, Qi Qi, Zihang Meng, Jianwei Yang, Ning Zhang, Xian Li, Ashish Shah, Dat Huynh, Hengduo Li, Zi Xian Yang, Sara Cao, Lawrence Jang, Shuyan Zhou, Jiacheng Zhu, Huan Sun, Jason Weston, Yu Su, and Yifan Wu. 2025. Agent learning via early experience. *ArXiv abs/2510.08558*.
- Zhang, Kongcheng, Qi Yao, Shunyu Liu, Wenjian Zhang, Mingcan Cen, Yang Zhou, Wenkai Fang, Yiru Zhao, Baisheng Lai, and Mingli Song. 2025. Replay failures as successes: Sample-efficient reinforcement learning for instruction following.
- Zhang, Muru, Ofir Press, William Merrill, Alisa Liu, and Noah A. Smith. 2023. How language model hallucinations can snowball. *ArXiv abs/2305.13534*.
- Zhang, Puzhen, Xuyang Chen, Yu Feng, Yuhan Jiang, and Liqiu Meng. 2025. Constructing coherent spatial memory in llm agents through graph rectification. *ArXiv abs/2510.04195*.
- Zhang, Yuxiang, Jiangming Shu, Ye Ma, Xueyuan Lin, Shangxi Wu, and Jitao Sang. 2025. Memory as action: Autonomous context curation for long-horizon agentic tasks.
- Zhang, Zeyu, Quanyu Dai, Xiaohe Bo, Chen Ma, Rui Li, Xu Chen, Jieming Zhu, Zhenhua Dong, and Ji-Rong Wen. 2024. A survey on the memory mechanism of large language model-based agents. *ACM Transactions on Information Systems* 43, 1 – 47.
- Zhang, Zeyu, Quanyu Dai, Xu Chen, Rui Li, Zhongyang Li, and Zhenhua Dong. 2025. Memengine: A unified and modular library for developing advanced memory of llm-based agents. *Companion Proceedings of the ACM on Web Conference 2025*.
- Zhang, Zeyu, Yang Zhang, Haoran Tan, Rui Li, and Xu Chen. 2025. Explicit v.s. implicit memory: Exploring multi-hop complex reasoning over personalized information. *ArXiv abs/2508.13250*.
- Zhang, Zhenyu (Allen), Ying Sheng, Tianyi Zhou, Tianlong Chen, Lianmin Zheng, Ruisi Cai, Zhao Song, Yuandong Tian, Christopher Ré, Clark W. Barrett, Zhangyang Wang, and Beidi Chen. 2023. H2o: Heavy-hitter oracle for efficient generative inference of large language models. *ArXiv abs/2306.14048*.
- Zhao, Siyan, Mingyi Hong, Yang Liu, Devamanyu Hazarika, and Kaixiang Lin. 2025. Do llms recognize your preferences? evaluating personalized preference following in llms. *ArXiv abs/2502.09597*.
- Zheng, Junhao, Xidi Cai, Qiuke Li, Duzhen Zhang, Zhongzhi Li, Yingying Zhang, Le Song, and Qianli Ma. 2025. Lifelongagentbench: Evaluating llm agents as lifelong learners. *ArXiv abs/2505.11942*.
- Zheng, Qinqing, Mikael Henaff, Amy Zhang, Aditya Grover, and Brandon Amos. 2024. Online intrinsic rewards for decision making agents from large language model feedback. *ArXiv abs/2410.23022*.
- Zhong, Wanjun, Lianghong Guo, Qi-Fei Gao, He Ye, and Yanlin Wang. 2023. Memorybank: Enhancing large language models with long-term memory. *ArXiv abs/2305.10250*.

- Zhou, Wangchunshu, Yuchen Eleanor Jiang, Peng Cui, Tiannan Wang, Zhenxin Xiao, Yifan Hou, Ryan Cotterell, and Mrinmaya Sachan. 2023. Recurrentgpt: Interactive generation of (arbitrarily) long text. *ArXiv abs/2305.13304*.
- Zhou, Yanfang, Xiaodong Li, Yuntao Liu, Yongqiang Zhao, Xintong Wang, Zhenyu Li, Jinlong Tian, and Xinhai Xu. 2025. M2pa: A multi-memory planning agent for open worlds inspired by cognitive theory. In *Annual Meeting of the Association for Computational Linguistics*.
- Zhou, Zijian, Ao Qu, Zhaoxuan Wu, Sunghwan Kim, Sunghwan Kim, Alok Prakash, Daniela Rus, Jinhua Zhao, Bryan Kian Hsiang Low, and Paul Liang. 2025. Mem1: Learning to synergize memory and reasoning for efficient long-horizon agents. *ArXiv abs/2506.15841*.
- Zhu, Kunlun, Zijia Liu, Bingxuan Li, Muxin Tian, Yingxuan Yang, Jiaxun Zhang, Pengrui Han, Qipeng Xie, Fuyang Cui, Weijia Zhang, Xiaoteng Ma, Xiaodong Yu, Gowtham Ramesh, Jialian Wu, Zicheng Liu, Pan Lu, James Zou, and Jiaxuan You. 2025. Where llm agents fail and how they can learn from failures. *ArXiv abs/2509.25370*.
- Zou, Jiaru, Xiyuan Yang, Ruizhong Qiu, Gaotang Li, Katherine Tieu, Pan Lu, Ke Shen, Hanghang Tong, Yejin Choi, Jingrui He, James Zou, Mengdi Wang, and Ling Yang. 2025. Latent collaboration in multi-agent systems.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.